



# UNIVERSITY OF TWENTE.

Faculty of Electrical Engineering,  
Mathematics & Computer Science



## Automated control prioritization based on generating sector-based threat profiles

Abe Winters  
Master's Thesis  
August 2023

---

**Supervisors:**  
dr. J. J. Santanna  
Mr. drs. ir. Paul Pols (Secura)

Design and Analysis of Communication Systems  
Faculty of Electrical Engineering,  
Mathematics and Computer Science  
University of Twente  
P.O. Box 217  
7500 AE Enschede  
The Netherlands

---



# Preface

This thesis concludes not only my master's degree, but also my time as a student, a period in my life that I really enjoyed. During this time I learned a lot from my studies, but gained maybe even more outside of that. I have met some wonderful people whom I am lucky to call my friends and I am grateful for the experiences I have had along the way.

I would like to thank my supervisor from the university, Jair Santanna, for guiding me through the process and helping me greatly with his insights and enthusiasm. He always managed to motivate me during our meetings and point me in the right direction when I was stuck.

I would also like to thank my supervisor from Secura, Paul Pols, for the guidance, the trust and the inspiration, and for the opportunity to conduct my research at Secura. This was a time I really enjoyed and I am happy that I will stay at Secura to work after my studies. After a short break, I will start my first full-time job there.

Furthermore, I would like to thank Florian Hahn and Michel Mollema for taking the time and effort to join my defense committee and for the help in arranging the colloquium.

I would like to thank the Secura colleagues who have been kind to have shared their knowledge with me and creating a fun work environment. Of course, I am obliged to give a shout out to my fellow interns, otherwise known as the *Council of Interns*. Special thanks goes to Raomi van Roozendaal, who coined the idea for sector-based threat profiles to support control prioritization in the future works section of her thesis.

My gratitude goes out to my friends, roommates, family and girlfriend for their support during the past months. Furthermore, I would like to thank people of the university, the staff, teachers, tutors, and, of course, my fellow students and the thesis support group for the study sessions.

Finally, I would like to thank Boiler Room, HÖR and all the DJs I have listened to along the way for providing me with the music that helped me through the hours of work that went into this thesis and fueling my passion for music.



# Summary

A threat can be defined as any circumstance that can adversely impact an organization. Threats may pose a risk, which is the probability that a threat causes harm. To protect against cyber threats, an organization can follow best practices from international standards, such as ISO 27001 or the NIST Cyber Security Framework. These standards provide guidelines to manage and improve the security for an organization, and contain sets of measures that may be implemented to manage the risks from cyber threats. Such measures are called controls, which contain descriptions for the implementations needed to satisfy that control. Due to resource and budget constraints, controls must be prioritized.

Prioritization of controls is often done following a risk assessment. These assessments often have subjective elements and can be an expensive practice. Therefore, they can benefit from a quantitative and automatic analysis as a support for scoring risks and recommendations. A way to do this is by including real-world threat information to create a risk profile from the active threat landscape. This threat landscape may not be the same for everyone. Threat actors have their preferred tactics, techniques, and procedures (TTPs) that they use and are known to target specific sectors. Therefore, the threat landscape can be different per sector. Knowing what threats might be more likely to target the organization can be used for proactive control implementation. Various studies have been conducted on (automatic) security control prioritization, mainly from a vulnerability perspective. These studies do not consider the active threat landscape that is relevant for the organization.

This study aims to define a methodology for automatic control prioritization based on active threat profiles for a sector. The research has been carried out as an internship at the Dutch cyber security company Secura. The main novelty of the study lies in this approach of automatically prioritizing security controls based on the active threat landscape for a sector, in the form of the TTPs used by active threat actors targeting that sector. This means that while most studies have an inside-out approach and focus mainly on all possible risks, this study takes a more outside-in approach by identifying the techniques of actors that are active in the threat landscape.

This work proposes a three-phased model. In the first phase, the active actors

are determined and filtered based on the sector. The second phase deals with ranking the actors and their TTPs, and the third phase handles the control prioritization. The supported control sets are from ISO 27001, NIST SP 800-53, NIST CSF, and the CIS Controls. The model is published on GitHub<sup>1</sup>.

Security controls are automatically prioritized against threat actors targeting a sector by first identifying and weighting the active actors targeting a sector. This is based on their operations within a time frame. Using MITRE ATT&CK, the TTPs that these actors use are collected. Threat landscape reports are considered as a source for determining active threat actors as well, but a method using the operation history can be fully automated, gives a more extensive threat landscape, and provides more transparency in the process, since they are directly linked to activity.

Following the results, the actors are best weighted from operations by using the product of the weighted operations within the time frame and their newness, based on the date the actors were first observed. Operations are weighted based on their date using an inverse function that prioritizes recent operations over older operations. The newness is a multiplier that compensates newer actors that do not have a lot of operations, since they can be active despite not having many operations yet. TTPs are weighted using the sum of the weights of the actors who use that particular TTP. This weighted threat profile forms the basis for the control prioritization. Mapping TTPs to controls is crucial in this step and can be done either directly or indirectly via other control frameworks. Using these mappings, controls can be weighted with the sum of the weighted techniques that they mitigate.

The results show that there are similarities in the top-weighted TTPs across sectors, even when the top-ranked actors are different. This means that there are TTPs that are used by many different actors. Furthermore, there are generic controls that arise to the top, regardless of the sector. This is explained by the shared TTPs among the actors and the large number of mappings that some controls have. Implementing these controls can serve as a wide base of mitigation against the shared techniques of different threat actors.

---

<sup>1</sup><https://github.com/AbeWinters/Threat-Profiles-for-Control-Prioritization>

# Contents

<b>Preface</b>	<b>iii</b>
<b>Summary</b>	<b>v</b>
<b>List of acronyms</b>	<b>xi</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Background</b>	<b>5</b>
2.1 Cyber threats . . . . .	5
2.2 Threat Actors . . . . .	6
2.2.1 Advanced Persistent Threats . . . . .	7
2.3 Cyber Threat Intelligence . . . . .	8
2.3.1 Tactics, Techniques and Procedures (TTPs) . . . . .	9
2.3.2 Threat Landscape Reports . . . . .	10
2.4 Security Controls . . . . .	11
2.4.1 ISO/IEC 27001 . . . . .	11
2.4.2 NIST . . . . .	12
2.4.3 CIS Controls . . . . .	12
2.5 Data Sources . . . . .	13
2.5.1 ETDA . . . . .	13
2.5.2 MITRE ATT&CK . . . . .	14
2.6 Related Work . . . . .	16
2.6.1 Security control prioritization . . . . .	16
2.6.2 Using NLP for cyber threat intelligence . . . . .	17
<b>3 Identifying threats for a sector</b>	<b>21</b>
3.1 Gather threat actor data . . . . .	22
3.2 Determining active actors . . . . .	22
3.2.1 Using threat landscape reports . . . . .	22
3.2.2 Using operations . . . . .	24
3.2.3 Methodology of comparison . . . . .	25

3.3	Filtering on a sector . . . . .	25
3.4	Results . . . . .	26
3.4.1	Gathered threat reports and operations . . . . .	26
3.4.2	Identified actors per sector . . . . .	27
3.5	Discussion . . . . .	30
3.5.1	Collecting reports . . . . .	30
3.5.2	Comparing the threat report analysis methods . . . . .	31
3.5.3	Threat reports versus operations . . . . .	31
3.6	Conclusion . . . . .	32
<b>4</b>	<b>Ranking threat actors and techniques</b>	<b>35</b>
4.1	Actors . . . . .	35
4.1.1	Threat report based . . . . .	36
4.1.2	Operation-based . . . . .	36
4.2	TTPs . . . . .	38
4.3	Results . . . . .	38
4.3.1	Inverse functions . . . . .	39
4.3.2	Ranking actors . . . . .	39
4.3.3	Ranking TTPs . . . . .	41
4.4	Discussion . . . . .	44
4.4.1	Inverse function . . . . .	44
4.4.2	Ranking actors . . . . .	45
4.4.3	Ranking TTPs . . . . .	45
4.5	Conclusion . . . . .	46
<b>5</b>	<b>Control prioritization from a threat profile</b>	<b>49</b>
5.1	Mapping controls to TTPs . . . . .	50
5.2	Prioritizing controls . . . . .	51
5.3	Results . . . . .	52
5.3.1	Targeted vs generic . . . . .	52
5.3.2	Prioritized controls . . . . .	53
5.4	Discussion . . . . .	55
5.4.1	Targeted vs generic . . . . .	55
5.4.2	Prioritizing controls . . . . .	56
5.5	Conclusion . . . . .	56
<b>6</b>	<b>Limitations</b>	<b>59</b>
<b>7</b>	<b>Conclusions</b>	<b>61</b>
7.1	Future Work . . . . .	63



---

<b>References</b>	<b>65</b>
<b>Appendices</b>	
<b>A Gathering and skimming threat reports</b>	<b>71</b>
A.1 Collecting threat reports . . . . .	71
A.2 Custom actor extraction method . . . . .	72
A.2.1 Create threat report overview . . . . .	72
A.2.2 Automated threat report scan . . . . .	72
A.3 Result list of companies . . . . .	74
A.4 Results from the threat report spreadsheet . . . . .	75
<b>B ETDA Threat Group Cards</b>	<b>77</b>
B.1 Victim Sector Data . . . . .	77



# List of acronyms

<b>APT</b>	advanced persistent threat
<b>CIS</b>	Center for Internet Security
<b>CSF</b>	Cybersecurity Framework
<b>CTI</b>	cyber threat intelligence
<b>NER</b>	named entity recognition
<b>NIST</b>	National Institute of Standards and Technology
<b>NLP</b>	natural language processing
<b>IEC</b>	International Electrotechnical Commission
<b>IG</b>	implementation group
<b>IOC</b>	indicator of compromise
<b>ISMS</b>	information security management system
<b>ISO</b>	International Organization for Standardization
<b>SP</b>	Special Publication
<b>TTP</b>	tactics, techniques and procedures



## Introduction

A *threat* can be defined as any circumstance that can adversely impact an organization [1]. Threats may pose a *risk*, which is the probability that a threat causes harm. These threats can come in the form of (targeted) cyber attacks by hackers, who are doing their best to exploit vulnerable systems. In 2021, the average number of cyber attacks and data breaches increased with 15.1% compared to the previous year [2]. In the same research by ThoughtLab, 29 % of CEOs and CISOs and 40 % of chief security officers believed that their organisations were not well prepared for the rapidly changing threat landscape.

A security maturity assessment can be used to assess how well an organization is prepared for potential security threats. It can help identify where the organization stands in terms of security and what measures should be prioritized to improve security levels. These reviews are regularly performed following international standards such as ISO/IEC 27001 [3] or the NIST Cybersecurity Framework (CSF) [4]. These standards provide best practices and guidelines for managing information security and dealing with cyber threats. They have their own requirements for compliance, and, if possible, how an accredited certification can be achieved. This is often required via an external audit. Cyber security standards like the aforementioned provide a list of measures that can be taken to help strengthen the cyber security posture. These measures are called controls.

*Controls* are measures to manage risks, which can be of administrative, technical, management or legal nature [1]. However, sometimes not all controls provided by a standard are applicable since they do not fit the risk profile of the organization. Furthermore, controls provide room for implementation. This can allow an organization to focus the implementation more on certain set of controls over the others. This can be needed since some standards come with a large number of controls. For example, ISO/IEC 27001:2013 contains 114 controls in 14 domains, and NIST SP 800-53 revision 5 has 322 base controls over 20 families [5]. These can be expanded to 1189 control enhancements. Therefore, since budget and resources are

limited, a prioritization of controls is needed.

Prioritizing these controls is usually done following a risk assessment. According to NIST SP 800-30, a realistic risk assessment identifies threats to an organization, vulnerabilities, the possible harm that may occur and the likelihood that harm will occur [6]. These assessments often have subjective components [7], where an expert evaluates the risks and assigns the corresponding risk scores. Based on these risks, security controls that mitigate these risks are prioritized. Subjective assessments can benefit from a quantitative analysis in scoring risks and recommendations for its objectivity. Furthermore, while a manual assessment by an expert can be accurate, it requires a fair amount of time, and therefore money. Smaller businesses may have smaller security budgets, so saving on the prioritization of controls reserves more budget for implementing those.

A way of incorporating a quantitative element is by using real-world threat information to create a threat-driven risk profile. Threat information can be used to define the threat landscape in which the organization operates. This threat landscape may not be the same for everyone. Some threat actors are known to focus on targets in specific sectors [8] and have their preferred tactics, techniques and procedures (TTPs) that they use [9]. Therefore, the threat landscape can be different per sector. Being aware of existing threats that may face an organization can help in proactive defense when this information is used in the prioritization of controls. Other studies have proposed ways of (automatic) control prioritization, but this is either vulnerability-based or based on all possible attack scenarios. These works do not take the active threat landscape and, therefore, realistic threats that an organization may face, into account. Section 2.6 discusses related work in more detail.

This research aims to define a methodology for prioritizing security controls for a sector via automatically generated threat profiles for this sector. To pursue this goal, the following research questions have been defined:

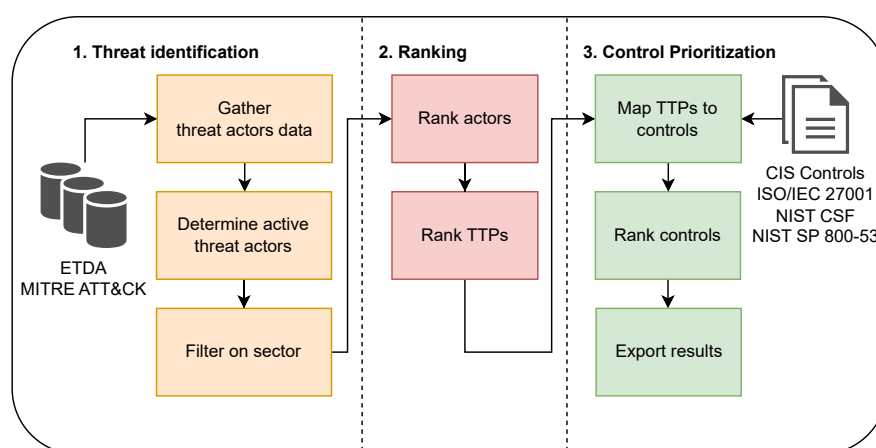
- **RQ1:** How to determine active threat actors targeting a sector?
- **RQ2:** How to determine the TTPs used by the active threat actors?
- **RQ3:** How can actors and their TTPs be prioritized for a sector?
- **RQ4:** How can controls be prioritized against threat actors targeting a sector?

This work proposes a three-phase model to achieve the goal of control prioritization via automatic generated sector-based threat profiles. The novelty lies within the fact that security controls are prioritized directly from real-life threats, in the form of the TTPs of active threat actors that target a sector. This way the implementing party has a direct link to the threats against which they are implementing these controls, thereby improving the understanding by deobfuscating the cyber threat landscape.

The model uses the ETDA Threat Group Cards and MITRE ATT&CK as primary sources of information on actors and their techniques. The techniques are depicted as TTPs from the MITRE ATT&CK database. A graphical overview of the model can be seen in Figure 1.1. Background information and related works are presented in Chapter 2. The data sources are discussed further in section 2.5.

The first phase aims to identify active threat actors that have been observed in a sector and retrieve their used techniques. Both threat landscape reports as well as publicly known operations from threat actors are evaluated as a potential source of threat actor activity. This phase aims to answer the first two research questions and is further discussed in Chapter 3. The second phase applies weighting functions to both the actors and their techniques. These weights are used to rank actors and their TTPs and are later used in the control prioritization. This phase is presented in Chapter 4 and answers the third research question. The third and final phase of the model is the control prioritization phase, answering the fourth research question. The weighted TTPs are used as input and are mapped to controls. These mappings are between MITRE ATT&CK TTPs and controls from cyber security standards. The controls are weighted on the basis of the sum of the TTP weights for the TTPs they mitigate.

These phase chapters contain their own introductions, methodology, results, discussions and conclusions. This model takes a sector and a time frame as an input, and will generate a threat profile based on the actors that have been identified as active in that period and are known to target that sector. Supported security control frameworks are: ISO 27001, NIST SP 800-53, NIST CSF and the CIS Controls.



**Figure 1.1:** High level overview of the model design

The limitations of the model are discussed in Chapter 6. The conclusions of the chapters are gathered in Chapter 7, where the general conclusions are drawn. This chapter also presents recommendations for future work.

This research has been carried out at the Dutch cyber security company Secura, as part of an internship, with the aim of including the proposed model into their service offerings. The final model is available online and published on GitHub.<sup>1</sup>

---

<sup>1</sup><https://github.com/AbeWinters/Threat-Profiles-for-Control-Prioritization>



# Background

This chapter provides background information for the topics covered in the study. Section 2.1 describes cyber threats and gives examples thereof, Section 2.2 provides background on threat actors and advanced persistent threats, which are the main focus of the thesis. Section 2.3 describes concepts from cyber threat intelligence and goes into more detail on TTPs and threat reports. Section 2.4 explains security standards and security controls, and Section 2.5 discusses the two main data sources used for the model. Section 2.6 discusses related work on prioritization of security controls and the use of natural language processing (NLP) for cyber threat intelligence.

## 2.1 Cyber threats

There are multiple definitions of a cyber threat. Within ISO 27000 it is defined as “potential cause of an unwanted incident, which may result in harm to a system or organization” [10]. NIST defines it as any circumstance that can adversely impact an organization [1]. Cyber threats are widely known in the form of a cyber attack, but come in many types and can have various origins [11]. Origins of a threat can be:

- **Deliberate:** e.g. a cyber attack
- **Accidental:** e.g. a human error or a machine failure
- **Environmental:** e.g. natural disaster
- **Negligence:** neglected factors compromising the safety

This study will focus on deliberate threats, or attacks, intentionally set out to do harm. Threats can be subdivided into various categories. Examples of categories of cyber threats are:

**Malware**, or malicious software, is software designed to do harm. The most prominent example is ransomware, which denies access to data and systems through encryption until a ransom is paid. Other examples of malware are spyware, viruses, worms, and trojans.

**Social Engineering** can be seen as a psychological attack. It tries to trick users into providing an entry point into the system, or unknowingly giving valuable information. A commonly used type of social engineering is phishing, where fraudulent emails or text messages are pretending to come from a trusted source and tricking the user to click a link, prompt credentials, or carry out an action to benefit the adversary, like transferring money to a bank account. These phishing scams often present themselves with a sense of urgency and can be personalized to the targeted user. The latter is called spear phishing.

A **Supply Chain Attack** targets an organization through its supply chain instead of directly targeting the organization. This exploits the trust that organizations may have in third party suppliers. Such an attack works by delivering a virus or other malware via a supplier or vendor. This can be either a software- or a hardware-based attack.

In an **Injection Attack** malicious code is inserted into a (web) application via untrusted user inputs or commands in order to expose sensitive information or compromise a system. A common vector is SQL injection, where SQL commands are executed to a back-end database by entering SQL queries in vulnerable user inputs. When successful, this could be used to steal data or modify a database. Another type of injection is cross-site scripting (XSS), where JavaScript code is injected into a web application.

**Man-In-The-Middle (MITM) Attack** is when communication between two endpoints is intercepted. The attacker can eavesdrop the communication, steal, or modify data, and impersonate the parties involved in the communication.

A **Denial of Service Attack** hinders the functioning of a system or renders a service inaccessible. This is typically accomplished by overloading the systems with traffic. In a distributed denial of service attack (DDoS), this traffic originates from many different sources.

As these examples show, cyber attacks come in many forms, and as new technologies emerge, attackers will find new attack vectors to use. A general term to describe these attackers is “threat actors”.

## 2.2 Threat Actors

A threat actor is an individual or a group posing a threat [1]. Various types of threat actors and their defining attributes are presented within the Threat Agent Library

(TAL) [12]. The goal of this library is to help in risk management to identify threat agent types relevant to assets, but on its own it gives a good overview of the various types of actors that can exist. The defining attributes within TAL are: intent, access, outcome, limits, resources, skills, objective and visibility.

The intent of a threat actor can either be non-hostile, like an untrained employee making a mistake, or hostile, e.g., a thief. Furthermore, the origin of access of threat actors can be either external or internal. An example of an internal threat actor is an unsatisfied employee holding a grudge. External actors are more common than internal actors: Verizon reports 80% of their examined breaches from 2008 up to 2022 have been caused by external actors [13]. The focus of this study lies on external (organized) actors.

Threat actors can have various motivations to carry out their attacks. A common motivation for threat actors is financial gain: retrieved data is sold to a third party or a ransom is demanded in a ransomware attack [13]. Nation-state actors are politically motivated and may perform cyber-espionage or seek to disrupt via attacks on critical infrastructure [14]. Hacktivists are often attacking for non-monetary reasons and motivated by ideology [15]. A special type of threat actor is an advanced persistent threat (APT), who can come in the form of a state actor performing espionage, but also as advanced cyber crime groups focusing on financial gain, e.g., using ransomware.

### **2.2.1 Advanced Persistent Threats**

An APT is a sophisticated and stealthy adversary who can remain undetected for a significant period of time. These attacks are often complex and performed by a well-resourced group, which may be state-sponsored [16]. Within most of these attacks, the adversary tries to remain a foothold in the systems of the victim for the duration of the attack.

In order to pursue this foothold, an APT attack typically takes a multi-step approach to gain and maintain access. An intrusion or attack can be broken down into distinct phases. These phases can be modeled in a “kill chain”, which describes the structure of an intrusion [17]. The Cyber Kill Chain® (CKC) from Lockheed Martin [17] is widely regarded as the industry standard and describes 7 consecutive phases of an APT attack. Another model is the Unified Kill Chain [18], which overcomes common critiques on the CKC by uniting and extending the Cyber Kill Chain and MITRE’s ATT&CK framework. The UKC, in contrast to the CKC, states that phases may be bypassed.

APTs are often known by various names due to cyber security companies using different naming conventions for attributed threat actors. For example: CrowdStrike

uses animals in their two-part cryptonym, like Cozy Bear, according to their nation state or motivation [15], Microsoft uses chemical elements, like Nobelium, and Mandiant uses a numerical convention, like APT 29. The three aforementioned names are all names for the same threat actor, which has been attributed to Russia's Foreign Intelligence Service [19]. A common naming convention does not yet exist, which complicates intelligence sharing. This problem has been addressed as one of the challenges of CTI in the 2020 position paper by Oosthoek and Doerr on the state of cyber threat intelligence [20]. When combining observations from different cybersecurity companies, it is important to be aware of this difference in naming conventions. Threat actor databases like the ETDA Threat Group Cards [8] or MITRE ATT&CK [9] have categorized actors and their associated names.

In order to be able to anticipate on threat actors and their attacks, defenders need an understanding of the threats that are facing them. Within the cyber security community, information on threat actor activity and trends in attack patterns are shared. This information can help in strengthening the defenses against threats by consuming the information and turning it into actionable intelligence. This is also known as cyber threat intelligence.

## 2.3 Cyber Threat Intelligence

cyber threat intelligence (CTI) is the actionable knowledge and insight into adversaries and their activities, enabling defenders to reduce harm [21]. The primary objective of CTI is to realize a knowledge advantage over adversaries [20]. Three main elements of CTI are relevant, timeliness, and actionable. The relevant threat data needs to be collected, analyzed and processed within a timely manner and the result should be actionable [22].

There are three broad levels of cyber threat intelligence: tactical, operational, and strategic, progressing from micro- to macro-level in terms of detail. Tactical threat intelligence focuses on the "what". The low-level and technical details of individual attacks and attackers are shared, like an indicator of compromise (IOC). This is often used within machine tools for the detection of threats and by incident responders and analysts searching for specific artifacts.

Operational threat intelligence focuses on the "how" and "where" and contains mid-level details of attacks and attackers. Shared intelligence on an operational level often consists of TTPs and provides information on the behavior of threat actors. This allows organizations to anticipate and prevent future attacks, but also assist in examining breaches.

Strategic threat intelligence is high-level intelligence on the threat landscape and the position of an organization therein [23]. It deals with the "who" and the "why"

and is the least technical level, being particularly useful for decision makers, such as CISOs and executives, to make informed decisions on mitigating risks posed by cyber threats. This level of threat intelligence includes expert opinions and insights based on combining both tactical and operational threat intelligence. Information includes mapping cyber attacks to geopolitical situations, targeting trends for industry sectors, statistics on breaches, and actor group trends [24].


Cyber threat intelligence has the ability to play a crucial role in defending against adversaries, since the defensive measures can be adapted to the current threat landscape. By using knowledge on the attackers and the types of attacks you can expect, you can act upon current trends in attacker behavior. The behavior of an attacker and the techniques used in an attack can be described using TTPs, or tactics, techniques, and procedures.

### **2.3.1 Tactics, Techniques and Procedures (TTPs)**

Tactics, techniques and procedures (TTPs) are a key concept in threat intelligence and can be seen as the behavior of an actor. A tactic represents the “why” and is the general end strategy of a threat action. A technique represents the “how” and describes the methods used to achieve the goal of the tactic. Procedures are the detailed descriptions of how the techniques are carried out. These TTPs are collected within MITRE ATT&CK (Adversarial Tactics, Techniques and Common Knowledge) [9]. This is a knowledge base of adversarial techniques based on real-world observations. An example record from ATT&CK is the Brute Force technique within the tactic Credential Access [25]. Furthermore, ATT&CK documents APT behavior by showing the techniques that are used by certain attackers.

David Bianco introduced a graphical representation of how acting on various types of threat indicators affects adversaries in his Pyramid of Pain [26], see Figure 2.1. This pyramid shows the relation between the indicators used to detect an adversary’s activities and how much “pain” it would bring an adversary when those indicators are denied. The closer an indicator is to the top, the harder it is for an attacker to change this for their next attacks. According to the pyramid, detecting and responding on the TTP level is the most painful for adversaries. Being able to deny actions on a TTP level means that you are operating on their behaviors instead of their tools. This means that adversaries need to change their behavior to try to pass the defense, and changing behavior can be difficult and takes time. Therefore being able to act on a TTP level is the most effective due to the hindering it gives adversaries.

Understanding TTPs is important in cyber threat intelligence and being able to act on them can prove very effective. TTPs can be retrieved from online databases,



**Figure 2.1:** The Pyramid of Pain [26]

like MITRE ATT&CK, but also from threat reports. These are reports from within the cyber security industry which describe observations and trends within the threat landscape.

### 2.3.2 Threat Landscape Reports

Within the cyber security community, some cyber security companies share their observations of the threat landscape in periodic threat landscape reports. These reports, often reporting on a full year, a half year, or a quarter of a year, contain information on the observed trends in active threats and threat actors within the reported period. The reported findings are often based on their own observations and may include observations of partners. The underlying data is often not public, only the trends are shared. These reports are a valuable source of information, but not necessarily a public source. While some companies offer their report online for anyone to download, others require an explicit request of access. Example reports are the annual ENISA Threat Landscape report [14] and the Verizon Data Breach Investigations Report [13].

Threat reports are written to be consumed by humans, so the essential information contains context for understanding. This allows humans to understand the full picture; the details are, however, wrapped in text. While this is very good for human understanding, it can be difficult to automatically spot the key information when trying to automatically analyze these reports. Threat reports are used in this study as a source for identifying active threat actors, and ways to extract this information are explored.

Cyber threat intelligence can help an organization to anticipate on cyber threats. Based on the intelligence, defensive measures can be put in place to act on these threats. Guidelines and best practices for such measures are included in cyber security standards and frameworks, where such measures are described in security controls. These controls have their implementation guidelines and requirements and are developed to help organizations follow best practices in cyber security.

## 2.4 Security Controls

A control is a measure to modify a risk, *including policies, procedures, guidelines, practices, or organizational structures, which can be of an administrative, technical, management, or legal nature* [27]. Lists of security controls may come with cyber security standards, which are a set of guidelines that organizations can use to improve their cyber security posture. When an organization is compliant with such a standard, it can achieve an accredited certification. To be certified, the organization must pass an external audit and meet all compliance criteria posed by the standard. Such a certification shows that the organization takes cyber security seriously and may be required by business partners. The requirements for certification differ per standard. For example, ISO 27001 requires a risk-based approach to select controls while compliance with a NIST standard requires a set minimum controls to be implemented. Reviewing and complying to these standards is not a one-time event, reviewing the compliance should be done regularly. Cyber security firms help organizations with their control implementation by performing risk assessments and proposing a prioritized list of controls to be implemented based on the results of that assessment.

There exist multiple cyber security standards, with ISO/IEC 27001 being a well-known example. For this work, ISO/IEC 27001, NIST CSF, NIST SP 800-53 and the CIS Controls are supported due to their popularity.

### 2.4.1 ISO/IEC 27001

ISO/IEC 27001 is an internationally recognized standard for information security management systems (ISMSs) that is developed by the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) [3]. It is part of the ISO/IEC 27000 family, which introduces guidelines for information security management. This family consists of multiple documents. The first document is ISO/IEC 27000, which is an overview and glossary of the standards in the family [10]. ISO/IEC 27001 follows and outlines the requirements for an information security management system. When performing a risk analysis for ISO/IEC 27001, the controls that correspond to the risks need to be implemented. Therefore, not all controls may be applicable to the organization. Annex A of ISO 27001 contains the controls corresponding to this standard. ISO/IEC 27002 [28] is an addition to ISO/IEC 27001 and goes more in-depth on the controls from Annex A, explaining how each control works, the objective and recommendations for implementation.

At the time of writing, ISO/IEC 27001:2013 is widely supported, but in late 2022, ISO/IEC 27001:2022 was published. The 2013 edition contains 114 controls across

14 domains, which have been reduced to 93 controls in the 4 domains of people, organizational, technological, and physical. From the moment the new version has been published, organizations have three years to update their processes and documentation to the 2022 version of the standard. Although the official abbreviation for this family of standards is ISO/IEC 27000, it is often written as ISO 27000 or ISO 27k.

## 2.4.2 NIST

The National Institute of Standards and Technology (NIST) is an institution from the United States, developing many standards and guidelines. The NIST Special Publication (SP) 800-53 is a publication that contains security and privacy controls [5] and is part of the SP 800 series on information security. At the time of writing, the most recent version of 800-53 is revision five. The controls in SP 800-53 are organized into 20 families, which contain base controls and control enhancements. These enhancements either add functionality to a control or increase the strength of a base control. In total there are 322 base controls, expanding to 1189 control enhancements. An example base control is SI-4: System Monitoring, part of the *System and Information Integrity* family. It has 25 control enhancements. An example of which is enhancement SI-4(1): System-wide intrusion detection system.

Another publication from NIST is the NIST Cybersecurity Framework (CSF) [4], also known under its official title: *Framework for Improving Critical Infrastructure Cybersecurity*. It provides an organizing structure for multiple approaches to cyber security by including guidelines, standards, and best practices. Although originally developed for critical infrastructure, it can be used by organizations in any sector. The framework core consists of the five areas identify, protect, detect, respond and recover. These areas, also called functions, are divided into 23 categories, which are then subdivided into 108 subcategories. While the NIST CSF does not prescribe controls, these subcategories contain requirements and link in their *informative references* to controls from other standards that illustrate a method to achieve the outcomes associated with that subcategory. An example of a subcategory is “ID.AM-4: External information systems are catalogued”. This is part of the Asset Management (ID.AM) category. While these are not officially called controls, in the model the prioritization of these subcategories is supported.

## 2.4.3 CIS Controls

The CIS Controls, previously known as Critical Security controls, are a set of recommended cyber security best practices originally developed by the SANS Institute



and currently owned by the Center for Internet Security (CIS) [29]. There are 18 CIS Controls, which all have their own safeguards. These safeguards can be seen as subcontrols. These describe the specific recommendations for defensive actions to be taken to implement that control. In total, there are 154 safeguards. An example is Control 06: Access Control Management, which has eight safeguards. One example of a safeguard for this control is safeguard 6.3: Establish an Access Revoking Process.

CIS controls are divided into three implementation groups implementation group (IG) to help organizations prioritize the controls. Each IG identifies a subset of controls that are assessed to be applicable to an organization with a similar risk profile. IG1 is defined as essential cyber hygiene, which serves as a foundation against general attacks, and the companies corresponding to IG1 are small to medium-sized with limited IT and cyber security expertise. IG2 builds upon IG1 and focuses on more complex environments, with organizations employing people responsible for managing and protecting IT infrastructure. Some of these safeguards require specialized expertise to install. An IG3 enterprise has security experts specializing in various aspects of cyber security. These safeguards must protect against targeted attacks from a sophisticated adversary.

CIS aims at cross-compatibility and offers mappings from CIS Controls to various other frameworks on their website [30], like controls from standards such as ISO 27001 and NIST SP 800-53, but also to MITRE ATT&CK. These available mappings are used in this study.

## 2.5 Data Sources

This section discusses the two main data sources used within the model presented in this study: the ETDA Threat Group Cards and MITRE ATT&CK. These data sources are merged to create a comprehensive base of information on threat actors and the TTPs they use.

### 2.5.1 ETDA

This study uses ETDA Threat Group Cards [8] for information on threat actors. This is an online “Threat Actor Encyclopedia” from ThaiCERT where information is cataloged on all known important threat actor groups. It was first published in 2019 as a free PDF, but is now available as an online portal. This portal aims to create full profiles of all threat groups worldwide that have been identified in research shared by anti-virus and security organizations over the years.

This information presented in the threat group cards is all the names of this actor, the country of origin, the sponsor type, motivation, the year they were first seen, a description, the observed victim sectors, the observed victim countries, their tools used, operations performed, counter operations, and links to general information. They also provide a link to the corresponding entry in MITRE ATT&CK. The dates of the listed operations are the dates when the stated activities started, not when they were reported. The information used to achieve the goal in this study is all the names of the actors, the year they were first seen, observed victim sectors, and the operations. Extra information presented in the output is the country of origin, victim countries, and motivations. All the information on the portal comes from public sources. Their main sources are:

- MISP Threat Actors galaxy [31]
- MITRE ATT&CK [9]
- Malpedia [32]
- AlienVault Open Threat Exchange (OTX) [33]
- Their own CTI archive and extensive searches on the internet.

There are a total of 443 threat groups in the database, 360 of which are APTs, 34 unknown and 49 fall in another category. There are 42 sectors included in the ETDA database. A complete list of sectors including the number of actors linked to that sector can be found in Table B.1 in Appendix B.

## 2.5.2 MITRE ATT&CK

MITRE ATT&CK (Adversarial Tactics, Techniques, and Common Knowledge) is a public knowledge base of adversary tactics and techniques based on observations from the real world [9], developed by the MITRE Corporation. It is a large collection of TTPs, which are indexed and described in detail, including the adversaries that have been observed using the TTPs. It evolves with the changing threat landscape and is a widely recognized knowledge base for the cyber security industry to understand attackers, their techniques, and mitigations. At a high level, ATT&CK is a behavioral model consisting of the following components [34]:

- Tactics, denoting the short-term end goal during an attack
- Techniques, describing the way the attacker achieves the tactical goal
- Sub-techniques, describing more specific means with which the attacker achieves the tactical goal at a lower level than in techniques

- Adversary usage of techniques

ATT&CK consists of three iterations, Enterprise, Mobile and ICS. This study uses ATT&CK for Enterprise, which focuses on traditional enterprise networks and cloud technologies. The documentation of adversary group behavior is essential for this study. This allows for linked techniques to the identified actors. ATT&CK collects information on threat groups, just like ETDA. Since ETDA is more extensive, mainly the link between adversaries and TTPs is used from ATT&CK, aside from the TTPs themselves. Further adversary information is gathered from ETDA. Since ATT&CK is such a widely recognized model, there exist many mappings to ATT&CK, and it is included in many studies. An example record in ATT&CK is

It should be noted that the information presented in these databases is based on public information on these threat actors and their operations. Since not all groups are as well documented as others, not every threat actor has the same extensive profile as others. This is because actors like to remain obscure and not all campaigns are documented in public. Therefore, it is very likely, and probably the case, that the listed threat actors are targeting more sectors than those listed, and probably have more operations than listed. However, it is very difficult to know everything. Being aware that the data is not the full picture is therefore important.

## 2.6 Related Work

This section discusses the related works to this study. Various studies on security control prioritization are discussed in Section 2.6.1. Since this thesis examines the possibility of using threat reports for determining relevant threat actors, related works on the use of natural language processing for cyber threat intelligence are included and discussed in Section 2.6.2.

### 2.6.1 Security control prioritization

Methods for the prioritization of controls in cyber security have been explored by different studies in the past and several models have been proposed. Gourisetti et al. (2020) have developed a prioritized mitigation framework (EPGA) that uses quantitative ranking techniques to perform prioritized vulnerability mitigation [35]. Cybersecurity controls from CSF are combined with multicriteria decision analysis and mathematical filters for a security analysis. Its novelty lies in the combination of quantitative ranking and security control dependency structures for prioritization of vulnerability mitigation. This model can be employed alongside frameworks, like CSF, to not only perform vulnerability analysis but also a prioritized vulnerability mitigation analysis to reach a desired cyber security maturity.

Al-Safwani et al. (2018) propose an information security control prioritization (ISCP) model that analyzes and prioritizes critical vulnerable controls based on assessment criteria [36]. Controls are assessed on threats, impact, and vulnerabilities, and their strength and weaknesses are tested. Based on these results and using multiple attribute decision making, the most effective and most vulnerable controls can be selected.

A prioritization method using digital twins has been proposed by Hadar et al. (2020) [37]. Based on attack graph analytics, security controls are gathered and prioritized automatically over active networks. The twin collects information on the network, connects it to possible attack tactics, measures the efficiency of implemented security controls, and detects the missing controls.

Lliansó (2012) has proposed a data-driven model to quantitatively justify investments in security control selection [38]. Priority is calculated using attack information, vulnerability impact, control cost, and the scoring of an expert. One step in the calculation consists of determining weights for controls based on their contribution to three attack-related areas: prevention, detection and response. These weights are based on the observation that controls that help to prevent an attack are more valuable than controls that later detect or respond to an attack.

In a study from Kwon et al (2020), the MITRE ATT&CK Matrix has been mapped

to the NIST Cybersecurity Framework to use threat information in defense implementations [39]. Versions of these frameworks made for the industrial sector are used. The result is a Cyber Threat Dictionary (CTD) that can be used both in reactive and proactive ways. Reactive in a way that when an attack is detected, the CTD can provide actions to mitigate the attack. It can be used in a proactive way to identify how controls will defend the organization against possible attacks. The authors suggest that more attack-defense mapping tools must be developed.

The methods mentioned before propose prioritization methods in various ways, where some do incorporate threat information, they do not consider the active threat landscape and thereby the possible active threats that may face the organization. The proposed methodology of this study prioritizes controls from a threat-driven perspective, whereas the above studies consider either vulnerabilities or general possible threats.

## 2.6.2 Using NLP for cyber threat intelligence

Multiple studies have been conducted on the extraction of information from text for a use in cyber threat intelligence. Depending on the goal of the study, they are designed for different types of text and can vary in the type of information they extract. A tabular comparison can be found in Table 2.1. Some of these studies and published tools focus on the extraction of indicators of compromise (IOCs) from technical reports [40], [41].

X. Liao et al. (2016) [40] have presented *iACE*, an approach for extracting indicators of compromise from technical reports using named entity recognition (NER) and relation extraction (RE). The approach is based on their observation that these IOCs are often described in a simple and predictable way: they are connected to a set of context terms, like “download” or “attachment”, through grammatical relations. The sentences containing IOC tokens are located using a set of regular expressions and common context terms. The relations between the tokens and context terms within a sentence are established by converting the sentence to a relation graph and applying a graph mining technique.

Other studies have focused specifically on extraction of threat actions from text [42]–[45]. Husari et al. presented TTPDrill [43], a publicly available model to extract threat actions from unstructured text using a novel and custom created threat-action ontology. They map the threat actions to MITRE tactics, techniques and procedures and their kill chain phase, representing the results in STIX format. In a later study, the same authors developed ActionMiner [42], a model that extracts chains of low-level treat actions from CTI reports.

The team at MITRE have developed the Threat Report Att&ck Mapper (TRAM)

with the goal to map threat intelligence reports to ATT&CK [46]. In the early development they explored fuzzy string matching. This worked really well for short technique names, for example *Mshita*. However, for longer names there was a very low fidelity [47]. This could be due to the fact that longer names of an attack contain verbs, and can be described in different words. A single-name attack has less possibilities in variation. Following this study, researchers from Microsoft have published the MitreMap Notebook, where text from submitted threat reports can be mapped to MITRE ATT&CK Enterprise techniques using a GPT2 model [48]. The ATT&CK ID is extracted, and indicators of compromise as well.

A study where both NLP algorithms and regular expressions are combined is CyNER [49]. The authors of the study provide three pre-trained models for prediction: a NER model trained on a cybersecurity corpus, which extracts five classes: malware, indicator, system organization, and vulnerability. The second model uses regular expressions for the extraction of IOCs, since these entities do not require an understanding of the context. Their third model is a generic NER model, to extract entities that do not fall under the cybersecurity concepts but may be of interest, like the targeted country.

The aforementioned studies focus on extracting intelligence from technical descriptions of an attack. An example of an origin of such a description is a technical threat report, where the details of an attack are shared. When reviewing related literature, no research has been found that is focused on extracting information from threat landscape reports. This is the type of report used in this study. Furthermore, most studies do not extract the names of threat actors. They are mainly focused on extracting either threat actions or IOCs. This study evaluates threat landscape reports as a source to determine threat actor activity and focuses on extracting threat actor names from these reports.

This study distinguishes itself from the existing literature in two ways. First, and the main contribution is that it prioritizes security controls based on the active threat landscape. This means that, while most studies have an inside-out approach and focus mainly on all possible risks, this study takes a more outside-in approach via threat profiles with the techniques from actors that have actually been observed in a sector.

The second contribution is that it analyses threat landscape reports. These reports are analyzed on threat actors, whereas most studies try to identify indicators of compromise or threat actions from technical reports.

**Table 2.1:** Published NLP tools for CTI

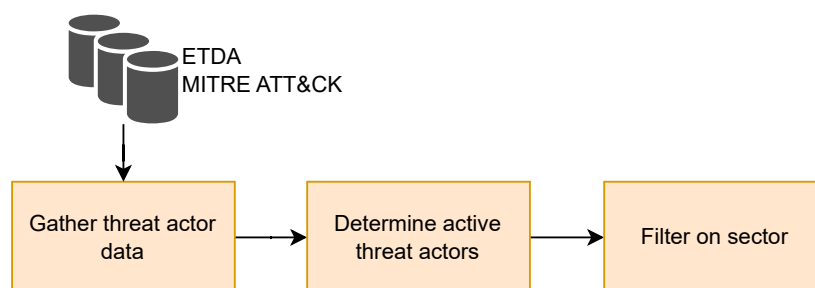
<b>model</b>	<b>properties extracted</b>	<b>public code</b>	<b>NLP models used</b>
TTPDrill [43]	Threat actions	x	TF-IDF
rcATT [50]	Threat actions	x	TF-IDF BoW, Linear SVC
CyNER [49]	Malware, Indicator, System, Organization, Vulnerability, IOCs, related non-cybersecurity concepts	x	
EXTRACTOR [44]	Threat actions	x	
g4ti NLP processor [41]	IOCs	x	
iACE [40]	IOCs		
ActionMiner [42]	Threat actions		
MitreMap [48]	ATT&CK TTPs, IOCs	x	Distill-GPT2
TRAM [46]	ATT&CK TTPs	x	





# Identifying threats for a sector

In order to prioritize controls against active threats, these threats must first be identified. This chapter aims to answer the first two research questions of the thesis: *How to determine active threat actors targeting a sector?* and *How to determine the TTPs used by the active threat actors?* The results of these two research questions are combined in the initial phase of the control prioritization model. The goal of this initial phase is to systematically identify the active threat actors that have targeted this sector in the past and collect the techniques that these actors are known to use in a way that can be automated. The phase consists of three main steps, as can be seen in Figure 3.1: Gathering data on threat actors, which includes information on used techniques, determining the active threat actors, and filtering the result set on the given sector. The result is a set of threat actors known to have targeted the given sector and the TTPs they use.



**Figure 3.1:** Overview of threat identification process

The three steps are discussed in Sections 3.1, 3.2 and 3.3 respectively. Two approaches are studied for determining active threat actors. Both using a collection of threat landscape reports to determine which actors have been active and using past operations to identify the degree of actor activity. These methods are described in Section 3.2.1 and Section 3.2.2 respectively.

## 3.1 Gather threat actor data

In the first step, data is collected from public sources to serve as a basis for further analysis. This data provides information on the sectors that actors are known to have targeted and the techniques they have used. This is the minimum required information for the further steps, such that a filter on a sector can be made and the shared techniques among the actors can be identified. The ETDA Threat Group Cards and MITRE ATT&CK are used to provide the necessary data. The data retrieved per source is:

- **ETDA Threat Group Cards** [8]: To provide information on threat actors. Data retrieved per threat actor is: name, alternate names, victim sector, victim country, performed operations, motivations, and the date first seen.
- **MITRE ATT&CK** [9]: To provide information on TTPs used by threat actors. Data retrieved per actor is: name, alternate names and TTPs used. Data retrieved per TTP is: the identifier (TID), tactic, technique and technique detection.

These two data sources are merged to create extensive threat actor profiles that contain actor metadata and their used techniques. The code used for this merge is forked from a repository by the GitHub user TropChaud called *Categorized Adversary TTPs* [51]. The existing merge did not contain the complete ETDA dataset, but includes only those actors who are present in the MITRE ATT&CK data. This repository is altered to include the complete ETDA actor list and gather the performed operations present in that same database. This retrieved information on actors and their techniques is used in the following step, where the active threat actors are determined.

## 3.2 Determining active actors

In this work, two sources are evaluated to determine active threat actors: threat landscape reports and operations. These approaches are described in the following sections and are compared in the results. The common variable of these approaches is the selected time frame. Threat landscape reports present the threat landscape of a certain time frame, and operations can be selected within a time frame.

### 3.2.1 Using threat landscape reports

Some cyber security companies publish threat landscape reports, where they report their observations in observed threats and threat actor activity. Vendors use


their own findings and may combine this with data from third parties, like Verizon does [13]. The idea behind this approach is that if a threat actor is mentioned within multiple threat landscapes, it is probably active. The assumption is that threat actors are mentioned either because of the number of operations within that time period or the level of impact they had. Combining the results of multiple threat landscape reports would provide results backed by multiple vendors in the cyber security industry.

The reports are collected by first identifying cyber security companies and then searching for threat landscape reports from those companies. Both are achieved by using a set of Google search queries. At first, the names of 50 well-known cyber security companies are collected. The names of these companies are then used to search for threat landscape reports within the specified time frame. The detailed methodology for collecting threat landscape reports can be found in Appendix A in Section A.1.


The reports are analyzed by counting the number of mentions per threat actor. An important step is to combine the results of synonyms for threat actors. Since threat actors are known by several names, the results for synonyms of the same actors must be combined. The scan itself is a simple string search that counts the number of occurrences of a string, and works with a supplied list of threat actor names to scan for. The steps taken in the scan are described in Section A.2.2 from Appendix A. Two methods of generating this list of names are compared: a method based on manually skimming the threat reports and a method that uses the actors' names from ETDA.

### **Analysis method 1: Manual skimming**

The first method to generate the initial list of threat actor names works by manually skimming the threat reports on the actors they mention. Skimming is a technique to rapidly go through a text, by not reading full sentences but scanning the pages on keywords and titles and leaving out details to extract the main essence of the author. The reasoning for skimming the reports is that reports may contain threat actor names that are not yet present in a database like that of ETDA or MITRE. Such a name must therefore be identified from the report, e.g. via skimming or reading. As a result of the skimming, a spreadsheet is created containing the report names and the actors mentioned per report. This approach considers threat reports as the base source for on actor identification and activity. An overview of the process of determining active actors by including skimming can be seen in Figure 3.2. A stepwise process of the skimming and building the spreadsheet, and how this is used in the scan, is described in Appendix A in Section A.2.



**Figure 3.2:** Actor determination method 1: Skimming threat reports



**Figure 3.3:** Actor determination method 2: Scanning threat reports

### Analysis method 2: Use actors within ETDA

The second method uses the threat actors included in the retrieved data from Section 3.1 as a basis for the names to scan for and thus considers threat reports as an enhancement for determining activity. ETDA is used as the basis for the actor identification. See Figure 3.3 for a graphical representation of the method. All alternative names, or synonyms, of threat actors are considered and the number of occurrences are counted per report.


The previous two analysis methods use threat landscape reports as the source of threat actor activity. In addition to using threat landscape reports, threat actor operations are also considered as a source of activity.

### 3.2.2 Using operations

Where many threat actors like to operate covertly, their actions can attract attention, especially when their operations have left a noticeable impact. The victim of an attack or a cyber security organization that has observed malicious behavior may decide to report on this publicly. If the operation is attributed to a certain threat actor group, this information can be included. Such known operations are collected in the ETDA Threat Group Cards per actor. This third method determines the active actors by selecting those actors with known operations within a set time frame. A graphical representation can be seen in Figure 3.4.

At first, the starting point is again the retrieved data from the MITRE ATT&CK and ETDA merge. This data contains the operations that are part of the ETDA dataset. Using a selected date interval, actors are chosen on the basis of operations within

that time frame. If they have a registered operation in the data within that time frame, the actor is selected.



**Figure 3.4:** Actor determination method 3: Select using operations within a time frame

### 3.2.3 Methodology of comparison

The methods are compared by generating results using the common variable between the different approaches: the time frame. From these results, the number of identified actors is compared, as well as their reported activity across five sectors. These five sectors are government, financial, education, IT, and industrial. These have been selected because they have a varying number of actors. Furthermore, the methods are compared by discussing the limitations of the used data sources, potential for automation, usability and the resulting threat landscapes. When the active actors have been determined and their metadata has been gathered, the selection can be filtered on the victim sector.

## 3.3 Filtering on a sector

Some threat actors are known to target organizations in specific sectors [8]. This means that the threat actors that one could expect differs per sector. Since threat actors can have differences in their way of working, commonly used techniques within a sector may differ as well. Therefore, the set of threat actors is filtered on a target sector to make the threat profile more personalized by leaving out threat actors that may not target this sector.

Sectors of the actors' victims are included within the ETDA Threat Group Cards and will be used as the filter property. A list of actors per sector can be found in Appendix B.

## 3.4 Results

This section presents the results from the aforementioned steps: gathering threat actor data, determining active actors, and filtering on a sector. The results of the different actor determination methods are compared on the shared time frame of the year 2021 across various sectors.

### 3.4.1 Gathered threat reports and operations

Using the proposed method, 12 reports have been collected from 11 cyber security companies that report on the year 2021. An overview of the reports, companies, and time periods is shown in Table 3.1. All companies that have been retrieved from the first step are listed in Table A.1. Obtaining reports from national organizations proved to be more challenging, since not all national organizations publish such reports and the naming schemes may differ. Although the national cyber security index is an interesting metric, it did not necessarily help identify national cyber security organizations that publish threat landscape reports. Therefore, following a manual search instead of the systematic method, the ENISA threat landscape report has been used. ENISA is the European Union Agency for Cybersecurity. Some reports that have been collected were not threat landscape reports as expected; for example, CTM360 publishes its "Cyber forecast" report [52]. This does not report on observations for the given year, but contains predictions of what the threat landscape may look like. Reports like this have been omitted in the subsequent steps.


Not all reports were directly available online; retrieving a report sometimes required a request form to be filled in. This would then grant access or send the report via email.

Looking at the time coverage of the reports, it can be seen that not all reports report on a full calendar year. Some reports do not start in the first quarter but start in the third quarter and report on the following four quarters, such as the ENISA threat landscape. In other cases, the report does not report on 4 quarters. For example, Bugcrowd reports up to and until the third quarter, since the report is released before the end of the year, and Fortinet publishes a report each half year. The year mentioned in the title of the report does not necessarily represent the year that has been reported on. Most often the year in the title is the publishing year, meaning that the report itself reports on the previous year.

Figure 3.5 shows the number of operations per year in the ETDA database. It can be seen that this number is the highest for the period 2020-2022, with the peak in 2020. There are 306 registered operations in 2021 by 80 threat actors.

**Table 3.1:** Used threat reports in the scan, with threat report names and reported time period

Organization	Name	Reported quarters
Blackberry	2022 Threat Report [53]	2021-1, 2021-2, 2021-3, 2021-4
Bugcrowd	Priority One Report 2022 [54]	2021-1, 2021-2, 2021-3
Crowdstrike	2022 Global Threat Report [15]	2021-1, 2021-2, 2021-3, 2021-4
deepwatch	Deepwatch Threat Intelligence 2022 [55]	2021-1, 2021-2, 2021-3, 2021-4
ENISA	ENISA Threat Landscape 2022 [14]	2021-3, 2021-4, 2022-1, 2022-2
Fortinet	Global Threat Landscape Report [56]	2021-1, 2021-2
Fortinet	Global Threat Landscape Report [57]	2021-3, 2021-4
IBM Security	X-Force Threat Intelligence Index [58]	2021-1, 2021-2, 2021-3, 2021-4
Microsoft	Microsoft Digital Defense Report [59]	2021-1, 2021-2, 2021-3
Palo Alto Networks	2022 Incident Response Report [60]	2021-2, 2021-3, 2021-4, 2022-1
Rapid7	Annual Vulnerability Intelligence Report [61]	2021-1, 2021-2, 2021-3, 2021-4
Verizon	Verizon Data Breach Report [13]	2021-1, 2021-2, 2021-3, 2021-4

**Figure 3.5:** Number of operations per year within ETDA database (on July 4th 2023)

### 3.4.2 Identified actors per sector

Table 3.2 shows the number of retrieved actors across five sectors, compared to the total number of actors in the data that target that sector. We can see that the operation-based method, Method 3, retrieves the most actors, and the threat report skimming method, Method 1, the least. The Industrial sector contains the least amount of actors, and here Method 2 retrieves one more actor than Method 3. En-

energetic Bear is mentioned by Microsoft and ENISA, but has no registered operations in 2021 within ETDA. Energetic Bear, or BROMIUM, is mentioned by Microsoft once within a paragraph, with an explanation that this actor is focusing more on the US state. It is not mentioned further, and no details are given on the operations, where this is the case for other actors in the following paragraphs of that report. Therefore, it is unclear whether the mention of Energetic Bear is related to an increase in activity or whether it is an example.

**Table 3.2:** Number of identified actors per determination method covering 2021. *Please note that the total number of actors present per sector in the “Total in ETDA” column have not been filtered on the time frame.*

	Total in ETDA	Method 1	Method 2	Method 3
<b>All sectors</b>	455	29	58	80
<b>Government</b>	192	17	35	46
<b>Financial</b>	102	11	21	31
<b>Education</b>	73	8	17	20
<b>IT</b>	36	5	9	11
<b>Industrial</b>	16	3	5	4

## Threat Reports

Figure 3.7 shows a heat map with the frequency of occurrence of threat actors per threat report resulting from Method 1: manual skimming. It can be seen that not all reports report on threat actors; 8 out of 13 reports mention threat actors. ENISA and CrowdStrike are mentioning the most threat actors. APT29, a group attributed to Russia’s Foreign Intelligence Service, is mentioned the most. For Method 2, the heat map is shown in Figure 3.6. The number of actors extracted from the threat reports is higher than that in Method 1. However, groups with names that are general words, such as Safe and Lead, are found to be present in most reports. This is because this word can not only be used to describe a threat actor group, but also as a general word in a sentence. Apart from these single words, we can see that it identifies more actors than the first method.






Figure 3.6: Actor coverage across threat reports using Method 2

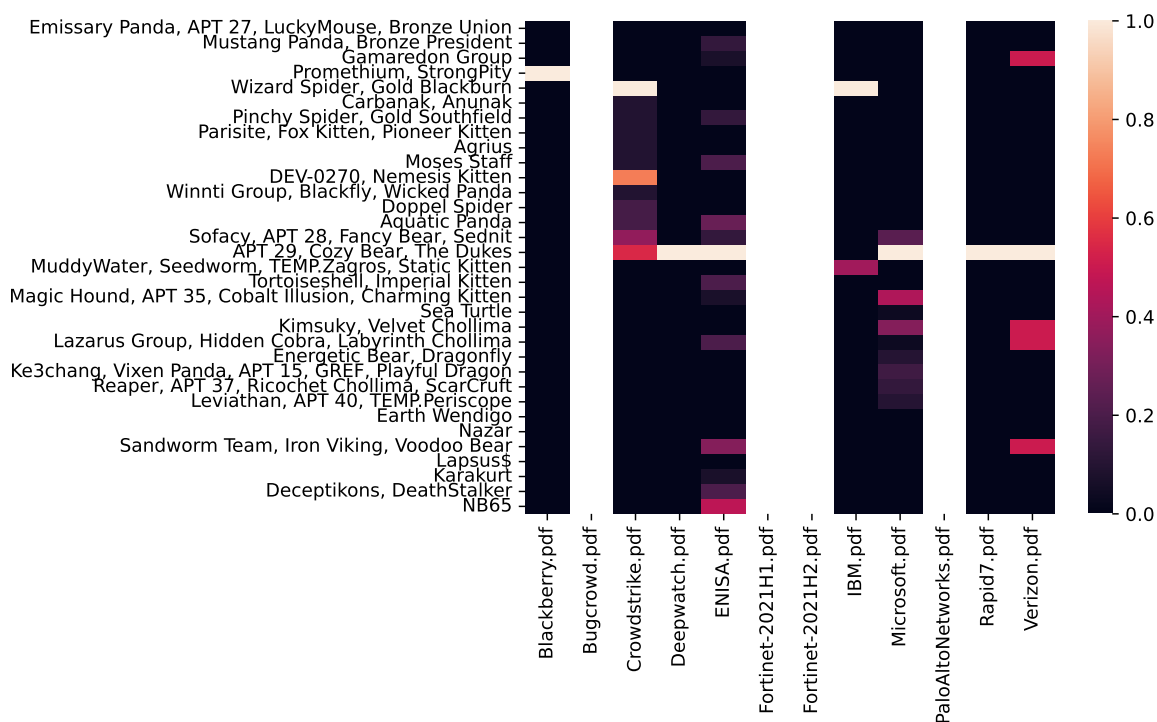


Figure 3.7: Actor coverage across threat reports using Method 1

## 3.5 Discussion

This section discusses the results presented in the previous section. First, the collection of the reports is discussed and the two ways for analyzing threat reports are compared. This is followed by a comparison between threat reports and operations for identifying active threat actors.

### 3.5.1 Collecting reports

Collecting reports is done via structured Google searches and can in theory be automated. However, in some cases a form needs to be filled out to request the report, rather than being readily available for download. Furthermore, since retrieving national cyber security organizations in a systematic way was difficult, the landscape report from the European Union Agency for Cyber Security was used. These challenges make automating the collection of threat reports somewhat harder. The first steps could be automated, where a tool retrieves organizations and website links, but would then be followed by a human operator that manually downloads and reviews the reports. This means that collecting threat reports would, for now, include a manual element.

### 3.5.2 Comparing the threat report analysis methods

The threat report scan requires a list of threat actor names to scan for in the reports. Method 1 built this list using manual skimming, and Method 2 retrieves this list from ETDA. Threat reports could potentially report on actors that are not (yet) in the ETDA and MITRE ATT&CK databases. This means that since Method 2 starts with a list of actor names from ETDA, it would not identify these actors in the reports. However, if the threat actor is not known in these databases, it would also mean that there is no TTP data available for this actor. This means that this actor cannot be used further for the control prioritization.

In theory, using Method 1 could result in at least as many actors as using Method 2. However, since the report is being skimmed rather than thoroughly read, the method is prone to accidentally missing an actor. We can see this happening when Method 2 discovers more actors in threat reports than Method 1, even extracting actors from reports where Method 1 has found no actors. These are actors that are mentioned once, like Aggah in the Fortinet 2021H2 report. Such mentions are easily overlooked. So, while Method 1 has the potential to identify new actors from reports, Method 2 actually identifies more actors in the same reports. A way to automatically extract these novel actors from reports without requiring a manual element is by using natural language processing to extract the names based on the context of the text. Building such a model was not in the scope of this thesis.

### 3.5.3 Threat reports versus operations

Both operations and threat reports have been considered as sources to identify active threat actors. It can be seen that not all threat landscape reports necessarily report on threat actors by name. If threat actors are mentioned, it is likely only a few. For example, when discussing a breach of that year. This could mean that probably only the most important ones are mentioned. Only CrowdStrike, Microsoft, ENISA and Verizon are reporting on a larger number of threat actors. Therefore, using threat reports as a source of threat actor identification results in a narrow threat landscape.

ETDA contains many more threat actors than the threat reports combined. The results in Table 3.2 show that a threat profile built from the operations in ETDA is more extensive in the number of identified actors than a threat profile built using threat reports. Across all sectors, 80 actors are defined as active using Method 3, compared to 58 using Method 2 and 29 using Method 1. A larger threat profile is favorable, since it contains more information and detail.

Since the results from Method 3 are directly based on operation data, the results are more transparent compared to parsing results of a closed data set, as is the case

for analyzing threat reports. However, the benefit of using threat reports is that since the results may come from a closed dataset, the results can be based on operations that are not publicly disclosed. Nevertheless, the manual steps hinder the suitability of threat reports for automatic threat profile generation. A method using operations can be fully automated since the data can be retrieved directly from an open-source dataset.

## 3.6 Conclusion

The goal of this chapter was to determine a way to systematically find active threat actors targeting a sector and obtain their used techniques in a way that can be automated, answering the first two research questions: *How to determine active threat actors targeting a sector?* and *How to determine the TTPs used by the active threat actors?*

At first a merge is made between the ETDA Threat Group Cards and MITRE ATT&CK to create a data set of extensive threat actor information and their used techniques. This is used for further actor analysis and listing the techniques used in a sector.

Both threat landscape reports and operations have been considered as a source for determining actor activity. 12 threat landscape reports from 11 cyber security organizations have been collected for the year 2021. Threat landscape reports proved difficult to automatically gather from the internet, since a request form may have to be filled in. This means that a manual step would be included in the collection of threat reports, meaning that it cannot be fully automated. The operations can be immediately retrieved from the ETDA database.

From these reports the number of actor mentions are counted. The list of actor names that are provided to this counter can either originate from skimming the reports beforehand or retrieving a list of actor names from ETDA. Threat reports may report on actors not present in ETDA, skimming the reports could identify these actors. The results show, however, that skimming is prone to missing actors. Scanning the reports on actor names from ETDA identifies more threat actors, even in reports where skimming has not found any actors. Furthermore, if an actor is found in a report that is not present in ETDA or MITRE ATT&CK, it means there is no information available on the TTPs used by that actor, thus it cannot be used in the later control prioritization. Therefore, scanning reports using a predetermined list of names from a source like ETDA works best. Using these methods it can be seen that not all threat landscape reports report on threat actors. Some might mention an actor once in an example, but mainly CrowdStrike, ENISA, Microsoft and Verizon discuss a wide variety of threat actors.


Although threat landscape reports have the potential to report on threat actors yet unknown to a database like ETDA, the threat profile resulting from threat landscape reports is smaller than when operations are used. Across all sectors, 80 actors have been identified as active in 2021 by operations, whereas at most 58 have been identified by using the threat reports. The results of a method that uses operations provide more transparency compared to threat reports, since threat reports themselves contain the results of a closed data set. Using threat operations is more suited for a fully automated approach. The data can be retrieved automatically, and threat actors with operations within a time period are selected as active.

Given that using operations to determine active threat actors gives a more extensive threat landscape, is more transparent, and can be fully automated, it is the preferred method over threat landscape reports. Thus, relevant threat actors targeting a sector can be determined by first identifying active threat actors based on their operations in a time frame and filtering on those actors that have targeted that specified sector previously. The TTPs that these actors use can be retrieved from MITRE ATT&CK.



# Ranking threat actors and techniques

This chapter introduces a ranking phase and aims to answer the third research question: *How can actors and their TTPs be prioritized for a sector?* In order to give something a priority, it is necessary to assign a weight or measure to use as the basis for the ranking. Within this phase, two weighting functions are developed for both the actors and TTPs respectively. The ranking phase consists of two steps: First, the actors are ranked using weighting functions. Second, the TTPs are ranked based on summing the weights of actors who use that TTP. An overview can be seen in Figure 4.1.




**Figure 4.1:** Overview of ranking phase

## 4.1 Actors

Since the different methods for determining actors provide different data, there are corresponding ranking functions for each method. For threat reports, the ranking is based on how frequently an actor is mentioned across reports, whereas the ranking based on operations takes the relative date and number of operations into account.

### 4.1.1 Threat report based

Actors that have been identified from threat reports are ranked based on the frequency with which they are mentioned. A stepwise overview of the weight calculation can be seen in Figure 4.2.



**Figure 4.2:** Determining weights for actors that have been identified using threat reports

The input to the ranking is a matrix with a count of occurrences per actor per report. These are absolute numbers, so they cannot be compared between reports since there are differences in writing styles and report lengths. Therefore, these numbers need to be normalized to similar ranges. For this, min-max normalization is used. The formula is depicted in Equation 4.1. The min-max normalization normalizes the results to a range between 0 and 1 and is chosen for its simplicity, since the main requirement is that the values should fall in the same range and are positive. Another simple normalization function is the z-score. This deals with outliers better than the min-max, but in this case outliers are of interest. An actor that is mentioned very frequently is deemed important enough by the authors to mention it so frequently.

$$x_{normalized} = \frac{x - x_{min}}{x_{max} - x_{min}} \quad (4.1)$$

After normalizing, the results of the reports are summed for each actor to obtain a total count per actor, as shown in Equation 4.2.

$$w_{actor}(x) = \sum_{r \in R} f_{xr} \quad (4.2)$$

with  $R$  being the set of threat landscape reports and  $f_{xr}$  being the normalized frequency of the actor  $x$  in report  $r$ . This count is then normalized once again using Equation 4.1, because these weights are later used as the basis for the TTP ranking. Actors determined by operations also use this normalization function.

### 4.1.2 Operation-based

Threat actors that are determined via operations are weighted according to those operations. The goal is to rank them based on their potential for future activity. Therefore, both the number of recent operations and the first date that an actor has



been seen are taken into consideration. Recent activity can be an indicator that an actor is currently active. An actor that has just been observed for the first time may not have as many registered operations as a more established threat actor, but can nevertheless be active. Solely the number of operations would not be a good indicator of activity for these actors. Therefore, the date that an actor is first seen, or the relative newness, is weighted as well.

The weight function of the operations uses an inverse function  $inv(x)$ , which assigns a weight to the operation based on its relative occurrence within the time frame. The result of this inverse function is in the range  $[0, 1]$ . The exact inverse function that will be used within the operation weighting is chosen based on the results in Section 4.3.1. For the calculation of the weights, the year that the operations took place is considered instead of calculating on a more granular basis like months or days. This is both for simplicity reasons and for not assigning a different priority to a few weeks difference. The time frame itself is also selected on the basis of whole years. The operations weighting function is defined as follows:

$$w_{operations}(x) = \sum_{o \in O_x} inv(y_{end} - y_o) \quad (4.3)$$

with  $O_x$  being the set of operations of actor  $x$  within the selected time frame. The input to the inverse function is the difference between the end of the time frame and the year of the operation:  $y_{end} - y_o$ , with  $y_{end}$  being the end year of the time frame and  $y_o$  being the year of the operation. If the selected time frame spans one year, all selected operations fall within this year. For this scenario, all operations will be given weight 1, so the result of the weight function would simply be the number of operations. Various options for this inverse function are presented in Equation 4.6 and are compared in the result section.

Actors who recently started out do not have a lot of operations but can nonetheless be very active. To account for this newness in the total weight, a “newness” function is established. This function weighs an actor based on their first operation. The more recent this is, the higher the weight. The newness acts as a multiplier and should compensate new actors for their lack of operations due to their recent appearance in the threat landscape. The choice for the inverse function to be used is made based on the results presented in Section 4.3.1. The function for the newness multiplier is defined as:

$$w_{newness}(x) = inv(y_{end} - y_{x,firstseen}) \quad (4.4)$$

with  $y_{x,firstseen}$  being the year this actor has been first seen. The total weight for an actor is the product of the weight of operations and the newness weight. The complete TTP weight function is defined in Equation 4.5:

$$w_{actor}(x) = w_{operations}(x) \times w_{newness}(x) \quad (4.5)$$

### Inverse weighting functions

The operations and newness are weighted using an inverse weight function  $inv(x)$ , such that small numbers get a higher weight. This allows for recent operations to be weighted heavier than older operations. Since division by 0 is not possible, the inverse function would not work if an operation falls in the same year as the end year of the time frame, so the constant 1 is added. Various weighting functions, as presented in Equation 4.6, are compared, and the best fit is chosen.

$$inv(x) = \frac{1}{(1+x)} \quad (4.6a)$$

$$inv(x) = \frac{1}{\sqrt{(1+x)}} \quad (4.6b)$$

$$inv(x) = 2^{-(1+x)} \quad (4.6c)$$

## 4.2 TTPs

The TTP weighting is based on the actor weights. One TTP can be used by multiple threat actors, so the TTP weight is the sum of the weights from the actors using that TTP. The weight function is depicted in Equation 4.7.

$$w_{ttp}(x) = \sum_{a \in A_x} w_{actor}(a) \quad (4.7)$$

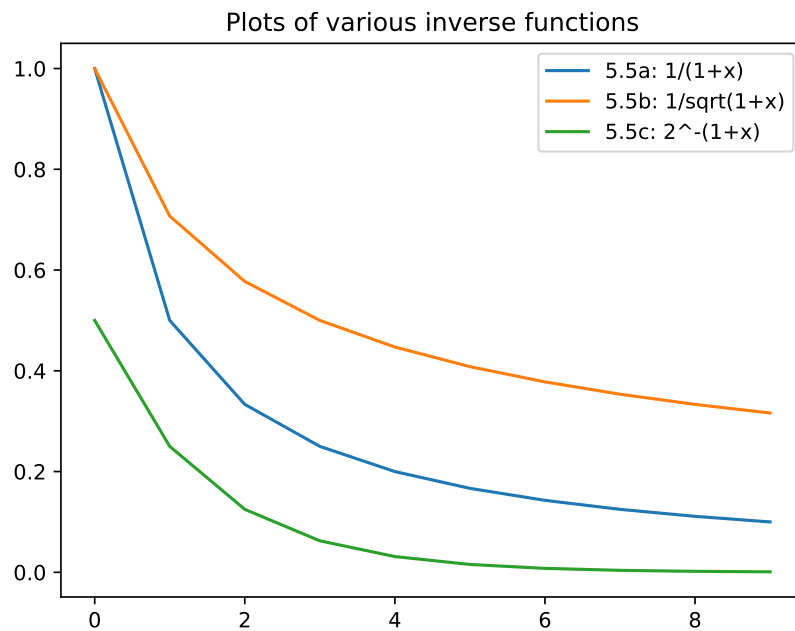
with  $A_x$  being the set of actors using that particular TTP  $x$ .

## 4.3 Results

This section presents the results from the steps discussed in the previous sections. Subsection 4.3.1 presents the results from using the various inverse functions, Subsection 4.3.2 presents the results from the two actor ranking methods and Subsection 4.3.3 presents the results from the TTP ranking.

### 4.3.1 Inverse functions

Figure 4.3 shows the graphs of the three inverse functions. We can see that both Equation 4.6a and 4.6b start at  $(0.0, 1.0)$ , whereas Equation 4.6c starts at  $(0.0, 0.5)$ . Equation 4.6a has a steeper slope between 0 and 2 than Equation 4.6b, after which the slopes are similar.



**Figure 4.3:** Determining weights for actors that have been identified using threat reports

### 4.3.2 Ranking actors

The top ten weighted threat actors from both threat reports and operations are compared on three sectors: Government, Financial and IT for the year 2021. These sectors are chosen to compare sectors that have a lot of actors, like the government, and sectors that have little actors targeting them, like IT. The financial sector is in terms of size in between the government and IT. Since the top rated actors have the most influence in the threat profile, the top ten are compared in this section instead of the full results, which can be up to 45 actors. The actors determined by operations are weighted according to Equation 4.8 for the operations as discussed in Section 4.4.1, following the results of the previous section. The results are presented in Table 4.1, 4.2 and 4.3. The actors that appear in both the top ten from operations and threat reports are presented in bold.

We can see that MuddyWater and Magic Hound are present in all three sectors for both operations and threat reports. This means that these actors are both very

active and targeting multiple sectors. Across the three sectors we can see that there is an overlap in the top tens of 50% for Government and IT, and 60% in the Financial sector. APT29 is clearly mentioned the most frequent in the threat reports, while it is lower rated by the operations. For the operations Wizard Spider is the most active threat actor for the Government and Financial sector. An interesting observation is that the highest rated actor from the threat reports is always present in the top ten from operations, but for example, Viking Spider, who is highest rated for the IT sector, is not present in the threat reports. Extending this observation somewhat broader, we can see that the majority of the top 5 from threat reports is present in the operations results, but not necessarily the other way around. The TTPs from these actors are presented in the next subsection, showing the top ranked TTPs across the same three sectors.

**Table 4.1:** Top ten ranked actors for Government sector in 2021

Operations		Threat Reports	
Threat Actor	Weight	Threat Actor	Weight
<b>Wizard Spider</b>	1.00000	<b>APT 29</b>	1.0000
Doppel Spider	0.35235	Safe	0.6490
<b>APT 29</b>	0.32429	<b>Wizard Spider</b>	0.5341
APT 41	0.21085	TA428	0.4612
<b>MuddyWater</b>	0.18335	<b>MuddyWater</b>	0.2800
Indrik Spider	0.16739	Aggah	0.2800
Lazarus Group	0.16739	Careto	0.2553
Reaper	0.12205	Sofacy	0.2506
<b>Magic Hound</b>	0.12205	<b>Magic Hound</b>	0.2494
<b>Kimsuky</b>	0.12205	<b>Kimsuky</b>	0.2059

**Table 4.2:** top ten ranked actors for Financial sector in 2021

Operations		Threat Reports	
Threat Actor	Weight	Threat Actor	Weight
<b>Wizard Spider</b>	1.0000	<b>APT 29</b>	1.0000
Carbanak	0.6293	<b>Wizard Spider</b>	0.5319
LockBit Gang	0.5685	<b>MuddyWater</b>	0.2766
FIN11	0.5473	Sofacy	0.2470
<b>APT 29</b>	0.3243	<b>Magic Hound</b>	0.2459
APT 41	0.2109	LightBasin	0.1277
<b>MuddyWater</b>	0.1834	<b>Indrik Spider</b>	0.1228
<b>Indrik Spider</b>	0.1674	FIN7	0.0957
<b>Lazarus Group</b>	0.1674	Parisite	0.0709
<b>Magic Hound</b>	0.1220	<b>Lazarus Group</b>	0.0662

**Table 4.3:** top ten ranked actors for the IT sector in 2021


Operations		Threat Reports	
Threat Actor	Weight	Threat Actor	Weight
Viking Spider	1.0000	<b>MuddyWater</b>	1.0000
<b>MuddyWater</b>	0.9639	<b>Sofacy</b>	0.8932
<b>Magic Hound</b>	0.6416	<b>Magic Hound</b>	0.8889
TA2101	0.4318	<b>Turla</b>	0.7133
<b>Sofacy</b>	0.2115	LightBasin	0.4615
<b>Turla</b>	0.1530	Tortoiseshell	0.4274
<b>Stone Panda</b>	0.1096	Parisite	0.2564
Bamboo Spider	0.0644	Energetic Bear	0.1197
Patchwork	0.0276	<b>Stone Panda</b>	0.0000
TaskMasters	0.0056		

### 4.3.3 Ranking TTPs

This section presents the results of the TTP ranking. Just like the previous section, the TTPs from the identified actors in the Government, Financial and IT sectors in 2021 are compared. At first, the TTP weight distributions are shown to compare the distribution in weighted TTPs. Second, the top ten TTPs are compared to see how much the slight differences in identified actors affect the TTP landscape. This also allows for the evaluation of frequently used TTPs across sectors.

Figures 4.4, 4.5 and 4.6 show the TTP weight distributions for the Government, Financial, and IT sector, respectively, for both operations and threat reports. It can be seen that all distributions are right-skewed. The distributions of both methods are very similar. For the IT sector, where there are a small number of actors identified, the peak of the distribution shifts more toward the right compared to the sectors with larger threat landscapes.

Tables 4.4, 4.5 and 4.6 show the top ten TTPs from the actors identified by operations and threat reports. The TTPs that appear on both sides of the table are presented in bold. We can see that there is a large overlap in shared top-rated TTPs. Across the three sectors, at least 7 of the 10 TTPs in the top ten are the same. Looking at the three tables, we can see PowerShell within the two highest rated TTPs across the three sectors. TTPs that appear in all three tables are PowerShell, Tool and Registry Run Keys / Startup Folder. The government sector and the financial sector show more overlap in TTPs with one another than with the IT sector. The TTPs for the IT sector show the most similarities between the two methods of the three tables, with 8 out of 10 TTPs being the same.




**(a):** TTP weight distribution resulting from threat reports

**(b):** TTP weight distribution resulting from threat reports

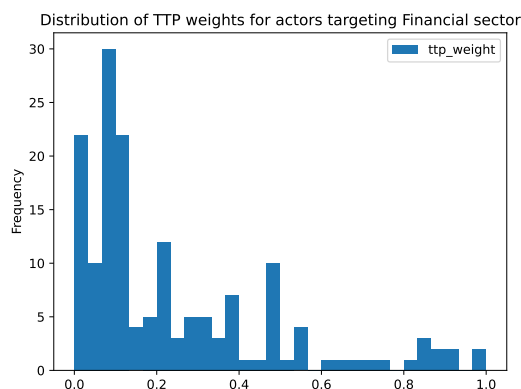
**Figure 4.4:** TTP weight distributions for actors targeting the Government sector in 2021

**Table 4.4:** Top ten TTPs in the Government sector by both operations and threat reports

Operations		Threat Reports	
TTP ID	Technique	TTP ID	Technique
<b>T1059.001</b>	<b>PowerShell</b>	<b>T1059.001</b>	<b>PowerShell</b>
<b>T1204.002</b>	<b>Malicious File</b>	<b>T1588.002</b>	<b>Tool</b>
<b>T1547.001</b>	<b>Registry Run Keys / Startup Folder</b>	<b>T1566.002</b>	<b>Spearphishing Link</b>
T1053.005	Scheduled Task	<b>T1204.002</b>	<b>Malicious File</b>
<b>T1566.001</b>	<b>Spearphishing Attachment</b>	T1105	Ingress Tool Transfer
T1059.003	Windows Command Shell	<b>T1547.001</b>	<b>Registry Run Keys / Startup Folder</b>
T1047	Windows Management Instrumentation	T1204.001	Malicious Link
<b>T1588.002</b>	<b>Tool</b>	<b>T1562.001</b>	<b>Disable or Modify Tools</b>
<b>T1566.002</b>	<b>Spearphishing Link</b>	<b>T1566.001</b>	<b>Spearphishing Attachment</b>
<b>T1562.001</b>	<b>Disable or Modify Tools</b>	<b>T1070.004</b>	<b>File Deletion</b>
<b>T1070.004</b>	<b>File Deletion</b>	T1036.005	Match Legitimate Name or Location

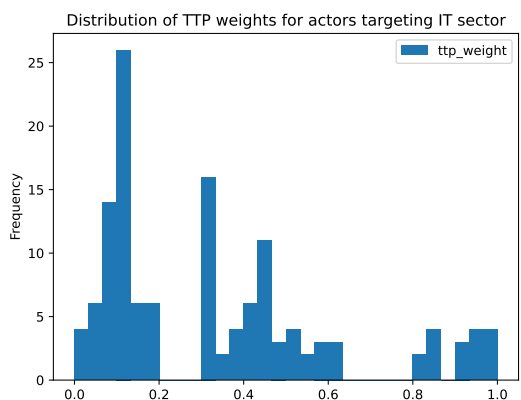


(a): TTP weight distribution resulting from operations




(b): TTP weight distribution resulting from threat reports

Figure 4.5: TTP weight distributions for actors targeting the Financial sector in 2021



(a): TTP weight distribution resulting from operations



(b): TTP weight distribution resulting from operations

Figure 4.6: TTP weight distributions for actors targeting the Financial sector in 2021

Table 4.5: Top ten TTPs in the Financial sector by both operations and threat reports

Operations		Threat Reports	
TTP ID	Technique	TTP ID	Technique
T1588.002	Tool	T1059.001	PowerShell
T1059.001	PowerShell	T1204.002	Malicious File
T1078	Valid Accounts	T1547.001	Registry Run Keys / Startup Folder
T1204.002	Malicious File	T1566.002	Spearphishing Link
T1053.005	Scheduled Task	T1588.002	Tool
T1566.001	Spearphishing Attachment	T1204.001	Malicious Link
T1047	Windows Management Instrumentation	T1047	Windows Management Instrumentation
T1547.001	Registry Run Keys / Startup Folder	T1053.005	Scheduled Task
T1059.003	Windows Command Shell	T1566.001	Spearphishing Attachment
T1036.004	Masquerade Task or Service	T1562.001	Disable or Modify Tools
T1566.002	Spearphishing Link	T1105	Ingress Tool Transfer

**Table 4.6:** Top ten TTPs in the IT sector by both operations and threat reports

Operations		Threat Reports	
TTP ID	Technique	TTP ID	Technique
T1059.001	PowerShell	T1083	File and Directory Discovery
T1083	File and Directory Discovery	T1059.001	PowerShell
T1588.002	Tool	T1059.003	Windows Command Shell
T1105	Ingress Tool Transfer	T1105	Ingress Tool Transfer
T1059.003	Windows Command Shell	T1560.001	Archive via Utility
T1560.001	Archive via Utility	T1547.001	Registry Run Keys / Startup Folder
T1204.001	Malicious Link	T1588.002	Tool
T1547.001	Registry Run Keys / Startup Folder	T1583.006	Web Services
T1566.002	Spearphishing Link	T1204.001	Malicious Link
T1583.006	Web Services	T1102.002	Bidirectional Communication
T1071.001	Web Protocols	T1057	Process Discovery

## 4.4 Discussion

This section discusses the results presented in the previous section. First, the results of the various inverse functions are discussed, and the inverse functions are chosen for the operation weighting and the newness multiplier. Then the results of the ranking of the actors and TTPs are discussed, including the similarities and differences between the results of the various determination methods.

### 4.4.1 Inverse function

In the graphs of the various inverse functions, we can see that Equation 4.6c does not start in 1.0, but in 0.5 for input 0. Adding a constant of 0.5 would transform the results to be too high. Therefore, this function is not suitable for the weighting of operations and newness. Equation 4.6a drops more steeply than Equation 4.6b. The rest of the graph lies below the graph of Equation 4.6b. This means that using Equation 4.6a would prioritize recent operations relatively heavier than Equation 4.6b as the resulting weights are lower for the same values. This suits the operation weighting, mainly prioritizing the very recent operations in a time frame and assigning less priority to the older operations. Therefore, the choice for the inverse function for the operation weighting is Equation 4.6a.

Equation 4.6b is less harsh than Equation 4.6a, which means that the curve is higher for the same values. This fits the description of the newness function. The characteristics of this equation allow it to compensate the very recent actors, but not affect older actors as much as Equation 4.6a. Therefore, the choice of inverse function for the newness will be Equation 4.6b.



This makes the total actor weight based on operations:

$$w_{actor}(x) = w_{operations}(x) \times w_{newness}(x) = \sum_{o \in O_x} \frac{1}{1 + y_{end} - y_o} \times \left( \frac{1}{\sqrt{1 + y_{end} - y_{x,firstseen}}} \right) \quad (4.8)$$

#### 4.4.2 Ranking actors

From the results we can see that both methods produce similar top-rated active actors, but also both have actors in their top ten that are scored much lower in the other method. The observation that the top 5 from threat reports is mostly present in the top ten from operations, but not the other way around, shows that via weighting operations those actors are ranked relatively lower than when threat reports are used. This is because from operations more actors emerge than from mentions in threat reports, and these additional identified actors can score higher.

Ranking actors based on their operations allows for more transparency in the ranking process and more assurance that the ranking is based on activity. A threat report may mention an actor for various reasons other than their high activity, e.g. a big decrease in activity compared to a previous year or the news that this threat actor has been dismantled or prosecuted. Therefore, weighting actors based on operations can be considered a more transparent method for ranking threat actor activity than threat reports because of the greater certainty that the actors are actually weighted based on activity.

#### 4.4.3 Ranking TTPs

The results of the TTP weight distributions show that all distributions are right-skewed. This makes sense since it shows that the majority of TTPs have a low weight and only a small amount have a high weight. This means that there is a small subset of techniques that are frequently shared among threat actors and a larger set that is more unique to the actors. PowerShell is present in every table in the top two weighted techniques. This is a broad technique, since PowerShell can be used for multiple actions, like discovery of information and execution of code [62]. The fact that it has such a high weight across multiple sectors means that it is a common technique among threat actors to use. The same holds for Tool and Registry Run Keys / Startup Folder, which are present across the three sectors as well.

We can also see a link between various top-rated TTPs, in how they might follow-up one another within an attack or how they are part of the same sequence of attack steps. An example can be seen in Table 4.4, with Spearphishing Attachment, Malicious File, PowerShell and File Deletion. The spearphishing attachment [63]

is a malicious attachment sent in a spearphishing mail. A follow-up technique is Malicious File [64], which tries to get a user to open a malicious file. This could trigger malicious code that could run within PowerShell [62]. At the end of the attack, an adversary may want to delete the files left behind by the actions. This is described in the File Deletion [65] technique. An attack consists of multiple steps and not of a single technique. Therefore, it makes sense that TTP chains can be recognized in the results.

Between the two methods, using operations or threat reports, there is a lot of overlap. This shows that while there was a bigger difference between the identified actors, the TTPs that result from these methods are nevertheless very similar. The fact that there is one TTP weighting function for both methods contributes to this because they are weighted by the same function. The overlap between sectors shows that there are a number of universal techniques that many actors use.

## 4.5 Conclusion

The goal of this chapter was to determine ways to rank threat actors and their TTPs and thus answer the third research question: *How can actors and their TTPs be prioritized for a sector?* To rank threat actors, two weighting functions have been developed for each of the two sources of activity identification: threat reports and operations. Ranking actors identified from threat reports is based on the frequency with which they are mentioned and ranking actors from operations is based on the weighted operations and the newness of the actors.

Operations are weighted based on their relative place in the time frame, the more recent, the higher the weight. Based on the results, Equation 4.6a is chosen as the inverse function for prioritizing recent operations heavier than older operations. The newness is a multiplier to compensate for threat actors that are new and do not have a lot of operations, even though they can be active. The year this actor has first been seen is weighted using the inverse square root (Equation 4.6b), as it prioritizes recent dates, but not as strongly as Equation 4.6a. The total actor weight function for operations is presented in Equation 4.8. TTPs are weighted using the sum of the weighted actors who use this TTP.

The two ranking approaches are compared across three sectors of varying actor sizes: Government, Financial and IT. For both the actors and the TTPs the resulting top ten results for a sector are compared. The results show that both determination methods produce similar top-rated actors, but not entirely equal. The five top results from the threat reports are always present in the 10 top results from using operations, but not the other way around. This is because more actors are identified via operations, which get a relative higher weight than some of those mentioned in

threat reports. The actors MuddyWater and Magic Hound are present in the top results of all three sectors.

The ranking of actors using operations provides more transparency than when threat reports are used and is directly linked to activity, whereas the mention of a threat actor within a threat report may not always be linked to activity. Therefore, the use of operations is better for determining and ranking the activity of threat actors than weighting using mentions in threat reports.

The prioritized TTPs show more overlap between the methods compared to the top results of the actors. This means that although the actors might differ, the techniques used are nonetheless similar. Among threat actors there seem to be general techniques that are commonly used. The PowerShell TTP ended up in the top two rated TTPs in the three sectors for both methods. Tool and Registry Run Keys / Startup Folder were present in the 10 top-rated techniques as well.

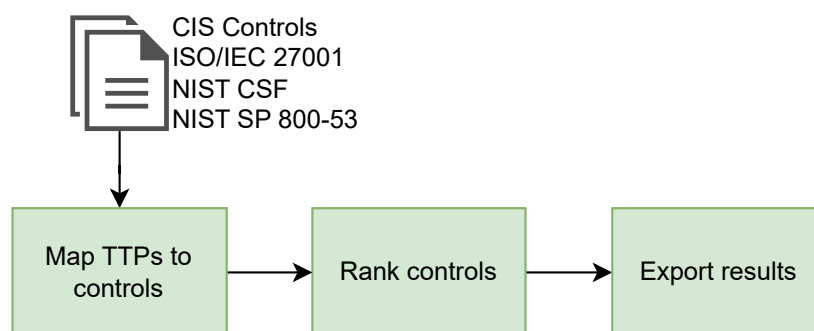
In conclusion, actors can best be prioritized based on operations via the product of a summation of the weighted operations and a multiplier for the newness of the threat actor. This multiplier compensates the new threat actors for their lack of operations. The operations are weighted on the basis of how recent they are. If actors are determined via threat reports, then simply the frequency of occurrence within these reports can be used for the ranking. TTPs can be prioritized by summing the actor weights of the actors who use that particular TTP.



# Control prioritization from a threat profile

This chapter describes the third phase of the model, the control prioritization phase, and aims to answer the fourth research question: *How can controls be prioritized against threat actors targeting a sector?* Regardless of the determination or ranking method, the input for the control prioritization phase is the same: a set of weighted TTPs. This phase consists of two steps. First, these TTPs are mapped to controls from one of the supported frameworks. Second, these controls are prioritized based on the weights of the related TTPs. The results are exported to an Excel file, with a sheet for the prioritized controls and a sheet containing the used TTP to control mappings. There is a sheet that contains the ranked TTPs including descriptions of technique detections from ATT&CK and a sheet containing the ranked actors with metadata from their threat group card in ETDA.

An overview can be seen in Figure 5.1. For universality, the goal is to support the prioritization of various common sets of security controls. Section 5.1 describes how the mappings between TTPs and controls are established, and Section 5.2 describes the prioritization of controls, which uses this mapping.

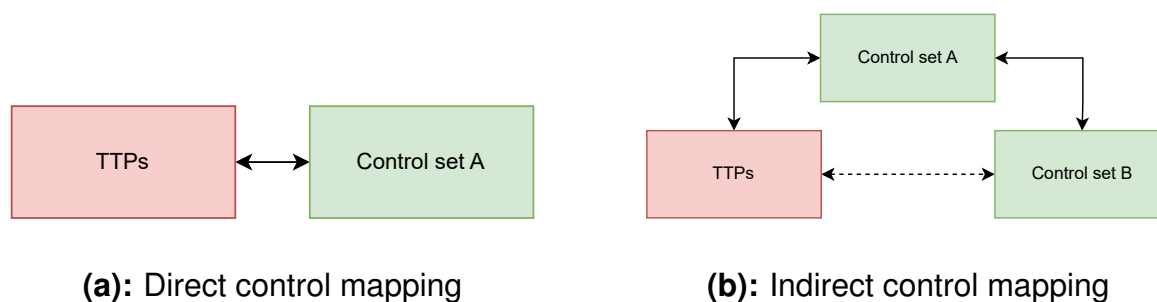


**Figure 5.1:** Overview of control prioritization phase

## 5.1 Mapping controls to TTPs

In order to prioritize controls based on a set of techniques, a mapping from these techniques to controls is needed. Such a mapping describes which controls mitigate that technique. For some control sets, there exist publicly available mappings from MITRE ATT&CK to those security controls. Since MITRE ATT&CK is a public source, it allows the public to create their own mappings. Two available mappings from MITRE ATT&CK to controls are those to CIS v8 [30] and NIST SP 800-53 revision 5 [66].

There exist mappings between different control sets as well, showing which controls from control set A satisfy the requirements of a control from set B and vice versa. The CIS Controls are an example of a control framework for which multiple mappings have been developed. These mappings are made and published by CIS themselves. For NIST SP 800-53 there exist various mappings to other control frameworks as well, including NIST CSF and ISO 27001. Using these mappings between controls, indirect mappings can be created between MITRE ATT&CK and a control set that does not have a direct mapping to their TTPs. Figure 5.2 shows the two types of mapping, direct or indirect.



**Figure 5.2:** Two approaches of mapping TTPs to controls via existing mappings

In this way, we can increase the number of control sets that can be prioritized via sector-based threat profiles. Table 5.1 lists possible mappings that map from MITRE ATT&CK TTPs to controls, either directly or indirectly. However, for ISO 27001 there exist multiple “routes” to ATT&CK. This can be done either via the CIS Controls or via NIST SP 800-53. CIS v8 maps to both MITRE ATT&CK and to ISO 27001:2022 [30]. NIST SP 800-53 rev. 5 maps to MITRE ATT&CK [66] and to ISO 27001:2013 [67]. When the mappings are made, controls can be prioritized.

**Table 5.1:** Mappings between various controls and MITRE ATT&CK

Control framework	Mapping to ATT&CK?
NIST SP 800-53 rev. 5	Directly
CIS v8	Directly
ISO 27001:2022	Indirectly, via CIS
ISO 27001:2013	Indirectly, via NIST SP 800-53
NIST CSF	Indirectly, via NIST SP 800-53

## 5.2 Prioritizing controls

Controls are prioritized based on the impact they have on the identified threat profile. Since there is no included metric of control effectiveness to a TTP within the mappings, the controls are prioritized by summing the weighted TTPs that are mitigated by a control. This prioritizes controls that either have a wide coverage within the set of TTPs from the threat profile or mitigate a set of high-ranked TTPs. This weighting function is depicted in Formula 5.1:

$$w_{control}(x) = \sum_{t \in T_x} w_{ttp}(t) \quad (5.1)$$

with  $T_x$  being the set of TTPs that are mitigated by control  $x$ , and  $w_{ttp}(t)$  being the weight of TTP  $t$  according to the TTP weighting function from Equation 4.7.

There are two ways to approach the control mapping, either a generic approach or a targeted approach. A generic approach takes the complete, or near-complete, set of TTPs used by the actors. For a targeted approach, the most high-rated TTPs are taken, which are then mapped to the controls. The reasoning for taking a targeted approach is the hypothesis that this will create a more tailored control prioritization, since some controls map to a very wide selection of TTPs. This could mean that the individual TTP ranking would be less important than the number of mappings between controls and TTPs, since the number of mappings could overshadow the TTP ranking. This is tested by taking only the top TTPs and evaluating the result of focusing only on this top selection. Selecting a subset of TTPs can be done by either filtering based on a threshold of the TTP weights or by selecting a top percentage of TTPs. The threshold method makes sure that all important TTPs are in there and that no TTPs with a similar ranking are cut out. A downside to this is that small sets of TTPs combined with a high threshold could result in only a handful of TTPs being selected, which then results in a scarce threat profile. This problem is not at play when selecting the subset by using the top percentage. This ensures that the subset is always populated. However, this could result in the cut-off point

being between TTPs of near-identical weights. The results are presented in the next section.

## 5.3 Results

This section presents the results of the control prioritization phase. Section 5.3.1 presents the results of taking a targeted approach versus a generic approach. Section 5.3.2 presents the results of the control prioritization of three sectors.

### 5.3.1 Targeted vs generic

This section presents the results of taking a targeted and a generic approach to the control prioritization. Subsets of the TTPs are made with both thresholds and percentages. The results are generated for the government sector with operations between 2020 and 2023 and the NIST SP 800-53 controls, as the government is the largest sector in terms of threat actors, and the NIST SP 800-53 provides the largest direct mapping between TTPs and controls. This large mapping can be interesting to see the difference between a large set and a significantly smaller set of TTPs and controls.

Table 5.2a shows the results of making a subset of the TTPs for the control mapping based on percentages and Table 5.2b shows the results from making a subset based on a threshold. We can see for both methods that the number of controls does not decrease at the same rate as the TTPs when the subset decreases in size. When comparing the two, it can be seen that taking the mean of the TTP weights is equal to taking 34.19% of the TTPs, but the number of controls are 80.70% of the full set.

**Table 5.2:** 10 top weighted controls from NIST SP 800-53 rev. 5 based using percentage-based subsets of the ranked TTPs in the Government sector between 2021 and 2023

<b>(a): Using percentages</b>			<b>(b): Using a threshold</b>		
<b>Percentage</b>	<b>Controls</b>	<b>TTPs</b>	<b>Threshold</b>	<b>Controls</b>	<b>TTPs</b>
100%	114	310	0.00	114	310
80%	113	248	0.25	82	79
50%	100	155	0.50	53	27
30%	90	93	0.75	34	10
10%	68	31	0.90	15	2
			0.19 (mean)	92	106



Table 5.3 shows the top ten controls when using all the TTPs and when taking a small subset of 10% to observe the difference a subset makes. We can see that the 10 top controls stay the same, although the order of ranking changes. System Monitoring and Configuration Settings stay the top two controls for both the complete set and the 10% subset.

**Table 5.3:** 10 top weighted controls from NIST SP 800-53 rev. 5 based using a generic and a targeted approach for the Government sector between 2021 and 2023

<b>(a): 100% of the TTPs (Generic)</b>		<b>(b): 10% of the TTPs (Targeted)</b>	
<b>Control ID</b>	<b>Control Description</b>	<b>Control ID</b>	<b>Control Description</b>
SI-4	System Monitoring	SI-4	System Monitoring
CM-6	Configuration Settings	CM-6	Configuration Settings
CM-2	Baseline Configuration	SI-3	Malicious Code Protection
SI-3	Malicious Code Protection	CM-2	Baseline Configuration
AC-6	Least Privilege	CA-7	Continuous Monitoring
CA-7	Continuous Monitoring	CM-7	Least Functionality
AC-3	Access Enforcement	AC-6	Least Privilege
CM-7	Least Functionality	AC-2	Account Management
AC-2	Account Management	AC-3	Access Enforcement
SC-7	Boundary Protection	SC-7	Boundary Protection

### 5.3.2 Prioritized controls

This section presents the results of control prioritization phase. Table 5.4 shows the number of actors, TTPs, controls, and the number of mappings between TTPs and controls for the government, financial, education, IT and industrial sector in 2021. The year 2021 is chosen to align with the previous two chapters. The control framework used is NIST SP 800-53 revision 5. We can see that, although the number of actors and TTPs decreases, the number of controls remains high. This is explained by the fact that some controls map to many TTPs.

Tables 5.5, 5.6 and 5.7 show the 10 top controls for these three sectors. We can see that the top three are ranked in the same order for all three sectors. The 10 top controls are nearly identical, with the IT sector containing AC-4: Information Flow Enforcement, which is not present in the other two sectors. The rest of the controls within the top 10 for the IT sector are shared by the other two sectors. The government and financial sector share the same 10 top ranked controls, with a slight ranking order difference.

**Table 5.4:** Number of items in the results of the threat profile and controls per sector

<b>Sector</b>	<b>Actors</b>	<b>TTPs</b>	<b>Controls</b>	<b>TTP - Control Mappings</b>
<b>Government</b>	46	300	113	2322
<b>Financial</b>	31	261	112	2102
<b>Education</b>	20	253	111	1988
<b>IT</b>	11	194	101	1632
<b>Industrial</b>	4	123	95	1147

**Table 5.5:** 10 top weighted controls from NIST SP 800-53 rev. 5 in the Government sector in 2021

<b>Control ID</b>	<b>Control Description</b>
SI-4	System Monitoring
CM-6	Configuration Settings
CM-2	Baseline Configuration
SI-3	Malicious Code Protection
AC-6	Least Privilege
AC-3	Access Enforcement
CA-7	Continuous Monitoring
CM-7	Least Functionality
AC-2	Account Management
SC-7	Boundary Protection

**Table 5.6:** 10 top weighted controls from NIST SP 800-53 rev. 5 in the Financial sector in 2021

<b>Control ID</b>	<b>Control Description</b>
SI-4	System Monitoring
CM-6	Configuration Settings
CM-2	Baseline Configuration
AC-6	Least Privilege
SI-3	Malicious Code Protection
AC-3	Access Enforcement
CA-7	Continuous Monitoring
CM-7	Least Functionality
AC-2	Account Management
SC-7	Boundary Protection

**Table 5.7:** 10 top weighted controls from NIST SP 800-53 rev. 5 in the IT sector in 2021

Control ID	Control Description
SI-4	System Monitoring
CM-6	Configuration Settings
CM-2	Baseline Configuration
SI-3	Malicious Code Protection
CA-7	Continuous Monitoring
CM-7	Least Functionality
AC-6	Least Privilege
AC-4	Information Flow Enforcement
AC-3	Access Enforcement
SC-7	Boundary Protection

## 5.4 Discussion

This section discussed the results presented in the previous section. The differences between taking a targeted versus a generic approach in selecting TTPs for control prioritization are discussed in Section 5.4.1. The results of the actual prioritization of the controls are discussed in Section 5.4.2.

### 5.4.1 Targeted vs generic

The results shows that taking a targeted approach, by selecting a subset of the TTPs to include in the control mapping, decreases the number of controls, but not as quickly as the TTPs are reduced in size.

Taking a targeted approach using a threshold allows control over a minimum degree of importance to be taken into account. Selecting a subset based on percentages allows for more control over the TTP subset size. Although the number of TTPs and controls decrease, the top ranked controls stay the same in a targeted approach. Since the number of controls is larger in a generic approach, and the top controls are the same to a targeted approach, it is better to take a generic approach in the prioritization. The larger control list allows for more flexibility in implementing, since controls that have already been implemented can be crossed off and the staff from organization can decide how many controls they wish to implement. Furthermore, a larger TTP set is also a wider representation of the threat actor activity.

## 5.4.2 Prioritizing controls

Table 5.4 shows that when the number of actors and TTPs decreases, the number of controls stays high due to the large number of mappings between TTPs and controls. This large number of mappings explains the similarities in top controls in the government, financial, and IT sectors. There are some controls that map to a large number of (commonly used) TTPs. For NIST SP 800-53 these are SI-4, CM-6, CM-2, SI-3, CA-7, CM-7, AC-6, AC-3 and SC-7. Since these are the top controls in the various sectors, they can be seen as universal controls to implement. This shows that although the set of actors that target a sector might be different, there are no big differences in the resulting prioritized controls per sector. This can be explained by the similarities in the top TTPs per sector, which are the TTPs commonly shared by the weighted actors. This means that a select set of controls offers mitigation capabilities to shared techniques of a wide variety of actors. Of course, the more controls are implemented, the better the overall mitigation capabilities.

## 5.5 Conclusion

The goal of this chapter was to define a way to prioritize security controls against threat actors targeting a sector. From the previous two chapters a ranked set of threat actors and their techniques was the result. For the prioritization, mappings from the TTPs to control frameworks are used. These can either be direct, or indirect by first mapping TTPs to control set A, and then mapping control set A to control set B. The controls are prioritized using the sum of the weights of the TTPs that a control mitigates.

The set of TTPs considered for the prioritization of controls can be either targeted, meaning that only a top set is taken, or generic, when all are considered. The results show that differentiating in the TTP subset size shows little difference in top prioritized controls. Therefore, a generic approach is considered best due to its support for a wide threat landscape.

The results show that there are a set of controls that can be seen as universal since they arise to the top across various sectors. This can be explained by the set of TTPs shared among the actors. Implementing these controls would therefore provide a basis of mitigation against commonly used techniques from various threat actors, regardless of the sector.

In conclusion, controls can be prioritized against threat actors targeting a sector by first identifying and weighting the active actors in a sector and weighting the techniques they use. By mapping these techniques to controls, either directly or indirectly, the controls can be prioritized based on the weighted techniques they

mitigate. Since some controls map to a wide set of TTPs, implementing these top controls can serve as a basis against the shared techniques of threat actors.



# Limitations

The proposed model is very dependent on its two main data sources, the ETDA Threat Group Cards and MITRE ATT&CK. Since the proposed model does not collect its own observations but makes use of the collected information within these data sets, it relies fully on the availability of these sources. Furthermore, the victim sectors presented in ETDA do not follow a standard that describes sectors and their methodology is not fully transparent. This makes the results not fully explainable.

Data is limited to the observations included in the data sets and is not representative of all attacks around the world. The attacks that are reported are the tip of the iceberg of the attacks that are actually carried out. The data used in the study contains publicly known breaches and incidents, and may therefore be prone to a reverse survivorship bias. A big part of this data is based on successful or observed attacks. Threat actors are performing more attacks than just their successful ones; these are, however, harder to document since they are not always observed. Therefore, defenses using the resulting control prioritization will not be based on all adversary activity, but on the *observed* techniques used by the *observed* actors in a sector.

This study considers all TTPs linked to an actor in ATT&CK, there is no time of usage linked to the TTPs. Therefore the TTP landscape can contain techniques that are not used anymore by these actors. The same holds for the observed sectors. The operations themselves are not linked to a sector. Filtering is based on all sectors that an actor has once been observed to target. Therefore, no ranking can be made on how prevalent an actor is in that specific sector. Only on how active an actor is in general. A counter-argument could be that if an actor was once using a technique or targeting a sector, they could still do it again. Furthermore, it could be that they are targeting their preferred list of sectors, but have the most success in a specific sector. The choices for inverse functions of the operation weighting is done on a subjective basis, and is not supported by literature.

The control prioritization makes use of mappings between TTPs and controls.

There is no information on the level of impact this control has on a TTP. Therefore, the prioritization is not based on the effectiveness of the controls against TTPs, solely on the number of weighted TTPs linked in the mappings.



# Conclusions

Organizations can follow cyber security standards to help improve their security posture. These standards provide sets of security controls that may be implemented to comply with this standard. Since there are many controls, and not all may be applicable, a prioritization must be made. This prioritization is usually done via a risk assessment, a lengthy and thorough process where the risks of an organization are assessed. Based on the identified risks, controls are prioritized. Such an assessment often takes a qualitative approach and can, therefore, benefit from a quantitative and automatic approach. Some studies propose methods for automatic control prioritization that are vulnerability-based, or try to generate all possible attack paths.

The goal of this thesis was to take a threat-driven approach to control prioritization and define a method to prioritize security controls for a sector using automatically generated threat profiles. To achieve this goal, we defined four research questions, which resulted in a model consisting of three phases.

The first phase is based on answering the first two research questions: *How to determine active threat actors targeting a sector?* and *How to determine the TTPs used by the active threat actors?* Relevant threat actors are determined by first gathering information on threat actors, determining who are active, and filtering this set on those actors having targeted that sector in the past. But first, threat actor data is collected from the ETDA Threat Group Cards and MITRE ATT&CK, including information on alternate names, victim sectors, operations and used TTPs. For determining active threat actors, two sources are evaluated. Both using threat landscape reports and selecting actors based on their operation history. While threat landscape reports have the potential to provide insights not present in published operations, since they report results from a (partly) closed data set, the number of reports mentioning actors is low and the gathering and analysis requires manual steps. On the contrary, using operations to identify active threat actors can be fully automated, gives a more extensive threat landscape, and provides transparency in the data. Therefore, threat operations are more suitable as a source for identifying

active threat actors than threat landscape reports. The set of active threat actors can be filtered on the sectors they are known to have targeted in the past. This information is included in the ETDA database. In conclusion, active threat actors targeting a sector can be determined by first identifying active threat actors based on their operations in a time frame and then filtering the set on the actors that have targeted the specified sector in the past. TTPs of these actors can be determined by retrieving them from MITRE ATT&CK based on the actors in the resulting threat profile.

The second phase is designed to answer the third research question: *How can actors and their TTPs be prioritized for a sector?* To prioritize actors and TTPs, weighting functions are developed for both activity determination methods. When actors are determined using threat reports, the frequency with which they are mentioned is the metric used for the ranking. The results, however, show that actors can best be prioritized via operations. This actor weighting is the product of a summation of the weighted operations and a multiplier for the newness of the threat actor. This multiplier compensates new threat actors for their lack of operations and is based on the year they are first seen. The relative occurrence of this year in the time frame is weighted using the inverse square root. Operations are weighted using the inverse function  $\frac{1}{1+x}$ , prioritizing recent operations. TTPs can be prioritized via a summation of the actor weights of the actors that use that particular TTP. Following these prioritization methods, results show that there is some overlap in top actors in the sectors, but there is more overlap in TTPs. This shows that there is a general set of TTPs that is universally shared, even when the top-ranked actors differ. Among these universal TTPs are PowerShell, Tool, and Registry Run Keys / Startup Folder.

The third phase, control prioritization from a threat profile, aims to answer the fourth and final research question: *How can controls be prioritized against threat actors targeting a sector?* Controls can be prioritized against threat actors targeting a sector by first identifying and weighting the active actors targeting a sector based on their operations within a time frame, and identifying the TTPs these actors use. This set of actors is then filtered on the sector. This sector-based threat profile forms the basis for the control prioritization. Mapping TTPs to controls is crucial in this step. There exist direct mappings from MITRE ATT&CK to CIS Controls and NIST SP 800-53. Through these mappings, indirect mappings to other control frameworks can be made, such as ISO 27001 and NIST CSF. Using these mappings, the controls can be weighted based on the sum of the weighted techniques they mitigate. The results show that there are generic controls that arise to the top, regardless of sector. This is explained by the shared TTPs among the actors and the large number of mappings that some controls have. Implementing these controls can serve as a wide base of mitigation against the shared techniques of threat actors.

## 7.1 Future Work

Future research can focus on building a more accurate threat profile, e.g. using natural language processing on reports and descriptions of actor operations. Performing a deeper analysis of actor operations to identify the used TTPs within that operation and/or metadata such as motive, targeted sector and country allows for a threat profile based on recent behavior. With these results, trends in the usage of techniques by actors can be analyzed to see how actors change their way of working over time. This can provide insights in whether there are certain trends in TTP usage or whether actors continue to use the same techniques. Furthermore, doing a deeper analysis on operations could introduce a metric on the severity of this operations. This can build a threat landscape where actors are ranked according to their sophistication and the impact of their operations.

Another place where impact can be included is the TTP ranking. This work used the number of actors that use a TTP and their respective weights, but aspects like their place within a kill chain are properties that can be used to give a more granular prioritization to TTPs. One TTP might be more critical than the other when executed, so being able to include this in the prioritization will improve the accuracy of the results.

This work provides a first step of demonstrating the benefits and consequences of implementing specific security controls. But within the control prioritization, there is currently no information on the degree to which a control mitigates a TTP. Further research could improve the control prioritization by looking into control effectiveness. Explainability can be improved by trying to merge descriptions of the controls and descriptions of mitigating a TTP from within ATT&CK. This could again be done with the help of natural language processing. This can both be used to try to get an indication how effective that control is against a TTP, and also form an implementation guideline.

Apart from basing control prioritization on just the threat landscape, information from within the company should be included as well, like their assets and potential vulnerabilities. Further research can look in to way on how to include a quantitative set of results like those in this study, with the results from a classic risk assessment.



# Bibliography

- [1] National Institute for Standards and Technology (NIST). Glossary. [Online]. Available: <https://csrc.nist.gov/glossary>
- [2] ThoughtLab, “Cybersecurity solutions for a riskier world.” ThoughtLab, May 2022. [Online]. Available: <https://thoughtlabgroup.com/cyber-solutions-riskier-world/>
- [3] “ISO/IEC 27001:2022 - Information security, cybersecurity and privacy protection — Information security management systems — Requirements,” 2022.
- [4] National Institute for Standards and Technology (NIST), “Framework for Improving Critical Infrastructure Cybersecurity Version 1.1,” National Institute for Standards and Technology (NIST), Standard, April 2018.
- [5] “Security and privacy controls for information systems and organizations,” National Institute of Standards and Technology, Gaithersburg, MD, Standard NIST Special Publication (SP) 800-53, Rev. 5, Includes updates as of December 10, 2020, September 2020.
- [6] “Guide for conducting risk assessments,” National Institute of Standards and Technology, Gaithersburg, MD, Standard NIST Special Publication (SP) 800-30, Rev. 1, September 2012.
- [7] F. Redmill, “Risk analysis-a subjective process,” *Engineering Management Journal*, vol. 12, no. 2, pp. 91–96, 2002.
- [8] Electronic Transactions Development Agency (ETDA), “Threat group cards: A threat actor encyclopedia.” [Online]. Available: <https://apt.etda.or.th/cgi-bin/aptgroups.cgi>
- [9] The MITRE Corporation, “MITRE ATT&CK Matrix.” [Online]. Available: <https://attack.mitre.org/matrices/enterprise/>
- [10] “ISO/IEC 27000:2018 Information technology — Security techniques — Information security management systems — Overview and vocabulary,” 2018.

- [11] “ISO/IEC 27005:2022 Information security, cybersecurity and privacy protection — Guidance on managing information security risks,” 2022.
- [12] T. Casey, “Threat agent library helps identify information security risks,” 2007.
- [13] Verizon, “2022 data breach investigations report,” 2022.
- [14] European Union Agency for Cybersecurity (ENISA), “ENISA Threat Landscape 2022,” 2022.
- [15] CrowdStrike, “2022 global threat report,” 2022.
- [16] A. Alshamrani, S. Myneni, A. Chowdhary, and D. Huang, “A survey on advanced persistent threats: Techniques, solutions, challenges, and research opportunities,” *IEEE Communications Surveys and Tutorials*, vol. 21, pp. 1851–1877, 4 2019.
- [17] E. M. Hutchins, M. J. Cloppert, and R. M. Amin, “Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains,” *Leading Issues in Information Warfare & Security Research*, vol. 1, p. 80, 2011.
- [18] P. Pols, “The unified kill chain,” 2017.
- [19] MITRE ATT&CK. Apt29. [Online]. Available: <https://attack.mitre.org/groups/G0016/>
- [20] K. Oosthoek and C. Doerr, “Cyber threat intelligence: A product without a process?” pp. 1–16, 2020.
- [21] S. Caltagirone, “Industrial control threat intelligence,” *von Dragos Inc. online*, 2018.
- [22] M. S. Abu, S. R. Selamat, A. Ariffin, and R. Yusof, “Cyber threat intelligence – issue and challenges,” *Indonesian Journal of Electrical Engineering and Computer Science*, vol. 10, pp. 371–379, 4 2018.
- [23] IBM. What is threat intelligence? [Online]. Available: <https://www.ibm.com/topics/threat-intelligence>
- [24] Anomali. What is threat intelligence? [Online]. Available: <https://www.anomali.com/resources/what-is-threat-intelligence>
- [25] MITRE ATT&CK. Brute force, technique t1110. [Online]. Available: <https://attack.mitre.org/techniques/T1110/>

- [26] D. J. Bianco, "The pyramid of pain," March 2013. [Online]. Available: <http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html>
- [27] ISACA, "ISACA Glossary of Terms." [Online]. Available: <https://www.isaca.org/resources/glossary>
- [28] "ISO/IEC 27002:2022 - Information security, cybersecurity and privacy protection — Information security controls," 2022.
- [29] CIS Controls. [Online]. Available: <https://www.cisecurity.org/controls>
- [30] Center for Internet Security (CIS), "CIS Critical Security Controls Navigator." [Online]. Available: <https://www.cisecurity.org/controls/cis-controls-navigator>
- [31] A. Dulaunoy, F. Roth, T. Schreck, and T. Steffens, "MISP Threat Actor Galaxy." [Online]. Available: <https://github.com/MISP/misp-galaxy/blob/master/clusters/threat-actor.json>
- [32] Fraunhofer FKIE, "Malpedia." [Online]. Available: <https://malpedia.caad.fkie.fraunhofer.de/>
- [33] AlienVault Inc., "AlienVault Open Threat Exchange (OTX)." [Online]. Available: <https://otx.alienvault.com/>
- [34] B. E. Strom, A. Applebaum, D. P. Miller, K. C. Nickels, A. G. Pennington, and C. B. Thomas, "Mitre att&ck@: Design and philosophy."
- [35] S. N. G. Gourisetti, M. Mylrea, and H. Patangia, "Cybersecurity vulnerability mitigation framework through empirical paradigm: Enhanced prioritized gap analysis," *Future Generation Computer Systems*, vol. 105, pp. 410–431, 4 2020.
- [36] N. Al-Safwani, Y. Fazea, and H. Ibrahim, "Iscp: In-depth model for selecting critical security controls," *Computers and Security*, vol. 77, pp. 565–577, 8 2018.
- [37] E. Hadar, D. Kravchenko, and A. Basovskiy, "Cyber digital twin simulator for automatic gathering and prioritization of security controls' requirements," vol. 2020-August. IEEE Computer Society, 8 2020, pp. 250–259.
- [38] T. Llansó, "Ciam: A data-driven approach for selecting and prioritizing security controls," 2012, pp. 91–98.
- [39] R. Kwon, T. Ashley, J. Castleberry, P. McKenzie, and S. N. G. Gourisetti, "Cyber threat dictionary using mitre att&ck matrix and nist cybersecurity framework mapping." Institute of Electrical and Electronics Engineers Inc., 10 2020, pp. 106–112.

- [40] X. Liao, K. Yuan, X. Wang, Z. Li, L. Xing, and R. Beyah, "Acing the ioc game: Toward automatic discovery and analysis of open-source cyber threat intelligence," vol. 24-28-October-2016. Association for Computing Machinery, 10 2016, pp. 755–766.
- [41] Y. Ghazi, Z. Anwar, R. Mumtaz, S. Saleem, and A. Tahir, "A supervised machine learning based approach for automatically extracting high-level threat intelligence from unstructured sources." Institute of Electrical and Electronics Engineers Inc., 1 2019, pp. 129–134.
- [42] G. Husari, X. Niu, B. Chu, and E. Al-Shaer, "Using entropy and mutual information to extract threat actions from cyber threat intelligence," 2018, pp. 1–6.
- [43] G. Husari, E. Al-Shaer, M. Ahmed, B. Chu, and X. Niu, "Ttpdrill: Automatic and accurate extraction of threat actions from unstructured text of cti sources," vol. Part F132521. Association for Computing Machinery, 12 2017, pp. 103–115.
- [44] K. Satvat, R. Gjomemo, and V. N. Venkatakrishnan, "Extractor: Extracting attack behavior from threat reports," 4 2021. [Online]. Available: <http://arxiv.org/abs/2104.08618>
- [45] G. Ayoade, S. Chandra, L. Khan, K. Hamlen, and B. Thuraisingham, "Automated threat report classification over multi-source data." Institute of Electrical and Electronics Engineers Inc., 11 2018, pp. 236–245.
- [46] MITRE Engenuity, "Threat Report ATT&CK Mapping (TRAM)." [Online]. Available: <https://mitre-engenuity.org/blog/2021/09/30/threat-report-attck-mapper-tram/>
- [47] S. Yoder, "Automating mapping to att&ck: The threat report att&ck mapper (tram) tool," 2019. [Online]. Available: <https://medium.com/mitre-attack/automating-mapping-to-attack-tram-1bb1b44bda76>
- [48] Microsoft, "MitreMap - Inferring MITRE Technique from Threat Intel Data ." [Online]. Available: <https://github.com/Azure/Azure-Sentinel-Notebooks/tree/master/mitremap-notebook>
- [49] M. T. Alam, D. Bhusal, Y. Park, and N. Rastogi, "Cyner: A python library for cybersecurity named entity recognition," 4 2022. [Online]. Available: <http://arxiv.org/abs/2204.05754>
- [50] V. Legoy, M. Caselli, C. Seifert, and A. Peter, "Automated retrieval of att&ck tactics and techniques for cyber threat reports," 4 2020. [Online]. Available: <http://arxiv.org/abs/2004.14322>



- [51] TropChaud, "Categorized adversary ttps." [Online]. Available: <https://github.com/tropChaud/Categorized-Adversary-TTPs>
- [52] CTM360, "The Cyber Forecast - Top 9 cybersecurity threats for 2021."
- [53] BlackBerry, "BlackBerry 2022 Threat Report," 2022.
- [54] Bugcrowd, "Priority One Report 2022," 2022.
- [55] Deepwatch, "Deepwatch Threat Intelligence 2022," 2022.
- [56] Fortinet, "Global Threat Landscape Report," August 2021.
- [57] —, "Global Threat Landscape Report," February 2022.
- [58] IBM Security, "X-Force Threat Intelligence Index 2022," 2022.
- [59] Microsoft, "Microsoft Digital Defense Report," 2021.
- [60] Unit 42 by Palo Alto Networks, "Incident Response Report," 2022.
- [61] C. Condon, J. Baines, S. McIntyre, and B. Watters, "Rapid7 2021 Vulnerability Intelligence Report."
- [62] MITRE, "Command and Scripting Interpreter: PowerShell," T1059.001. [Online]. Available: <https://attack.mitre.org/techniques/T1059/001/>
- [63] —, "Phishing: Spearphishing Attachment," T1566.001. [Online]. Available: <https://attack.mitre.org/techniques/T1566/001/>
- [64] —, "User Execution: Malicious File ," T1204.002. [Online]. Available: <https://attack.mitre.org/techniques/T1204/002/>
- [65] —, "Indicator Removal: File Deletion," T1070.004. [Online]. Available: <https://attack.mitre.org/techniques/T1070/004/>
- [66] MITRE Center for Threat Informed Defense, "NIST 800-53 CONTROLS TO ATT&CK MAPPINGS." [Online]. Available: <https://mitre-engenuity.org/cybersecurity/center-for-threat-informed-defense/our-work/nist-800-53-control-mappings/>
- [67] National Institute for Standards and Technology (NIST), "NIST SP 800-53, Revision 5 Control Mappings to ISO/IEC 27001."
- [68] e-Governance Academy, "National Cyber Security Index (NCSI)." [Online]. Available: <https://ncsi.ega.ee/ncsi-index/>



# Gathering and skimming threat reports

## A.1 Collecting threat reports

Before being able to analyse threat reports, they need to be collected. This is done using the following steps:

1. **Finding companies:** A list of 50 cyber security companies is made with the use of the search engine Google. The following Google search terms are used to find information on the top cyber security companies: “Top cyber security companies” and “Best cyber security companies”. Articles in the results are reviewed and used to build the list.
2. **Finding national organizations:** Next to the companies, national organizations are considered. The top 10 countries from the National Cybersecurity Index (NCSI) [68] are listed. The NCSI is a live global index that provides an assessment of the country’s cyber security. The top of this list is used with the assumption that countries that have a higher cyber security capacity have the resources to measure their national cyber security landscape.
3. **Finding threat reports:** For each organization in the list, a Google search query is used to find threat reports. The format used is:

*“organization name” AND (“threat report” OR “threat landscape” OR “year in review”) AND year*

An example query for a threat report from Microsoft on 2021 is:

*“Microsoft” AND (“Threat report” OR “Threat landscape” OR “year in review”) AND 2021*

If no threat reports are made by the company, this is marked in the list and the company will be filtered out.

## A.2 Custom actor extraction method

To serve the goal of this study, a custom threat report analyzer is developed that extracts threat actor names using a simple string search with a provided list of threat actor names. After collecting reports, a list of the current active threat actors is composed by manual skimming threat reports. This is to create a general understanding of the actors that these companies report on and how these reports are structured. After creating this initial list, the reports are scanned in an automated fashion using the generated lists from the first step. The number of occurrences per keyword are counted.

### A.2.1 Create threat report overview

This overview is made in the following way:

1. **Manual skimming threat reports:** The threat reports are skimmed on threat actor names. Skimming is a technique to rapidly go through a text, by not reading full sentences but scanning the pages on keywords, titles and leaving out details to extract the main essence of the author.
2. **Compile threat actor lists:** The unique threat actors of all the reports are combined in their respective lists.

### A.2.2 Automated threat report scan

In order to get a better understanding what threat actors are mentioned in threat reports and speed up the process, an automated scan is done to count the exact number of times the threat or actor occurs per report. After normalizing the results it can be estimated whether the report actually reports on the actor as active, with the number of occurrences as an indicator of the importance. This is following the assumption that when a threat or threat actor is discussed in and across reports more than other threats, it is likely more important than other threats. This also depends on the source organization, since one might mention an important threat multiple times in-text, while another may report in a more statistical sense and list the most active ones.

The scan makes use of the ETDA Threat Group Cards [8] of information on the retrieved actors. This is an online "Threat Actor Encyclopedia" from ThaiCERT

where information on all known important threat actor groups is cataloged. The information is based on public sources. This information on actors includes their synonyms, but also sectors and countries of the victims, motivations, country of origin and more.

This scan done as follows:

1. **Collect reports in a folder:** The reports resulting from the method described in Section A.1 are collected in a folder as PDF files.
2. **Preprocessing text:** Per report, the text from all pages is collected, converted to lowercase and symbols are removed. Lowercase is used for the case insensitive comparison later, such that different usage in capitals across reports does not inflict upon the results.
3. **Scan threat reports using the compiled lists:** The compiled lists from the spreadsheet contain an overview over the reported threat actors. Using the items in this list, a scan is done on the threat reports to count the number of occurrences per actor per report.
4. **Combine synonyms:** Using the threat actor encyclopedia from ETDA [8], the results from alternative names of threat actors are combined to one actor. This is due to the reason that the same threat actor can be known under various names, since organisations have different naming schemes. So in order to compare the actual actors with one another, the results of synonyms of the same actor need to be combined.
5. **Data normalization:** Reports vary in length and structure, some reports are a summary of the findings and straight to the point whereas other reports are longer and contain more textual explanation. The data needs to be normalized in order to compare the data between reports. The data is normalized using min-max normalization: this is one of the most common methods of normalization and uses the following formula:  $x_{normalized} = \frac{x - x_{min}}{x_{max} - x_{min}}$  The closer the number is to 1, the more this term is discussed within the report.
6. **Sum actor scores across reports:** Sum the scores across reports, such that the result is a summed score per actor.
7. **Sort results:** The results are sorted such that the most mentioned actors are on top.

### A.3 Result list of companies

Table A.1 contains the list of companies retrieved from the first step described in Section A.1 and lists whether this company publishes threat reports.

*Table A.1: Retrieved companies*

<b>Name</b>	<b>Published threat reports</b>
AlgoSec	FALSE
AppGuard	FALSE
Avast	TRUE
Avira	TRUE
Blackberry	TRUE
Bugcrowd	TRUE
CA Technologies	FALSE
Check Point Software	TRUE
Cisco	TRUE
Cobalt Iron	FALSE
CrowdStrike	TRUE
CTM360	TRUE
CyberArk	TRUE
DataDome	FALSE
deepwatch	TRUE
Fortinet	TRUE
Herjavec Group	FALSE
Hillstone Networks	FALSE
IBM Security	TRUE
iboss.com	FALSE
Identiv	FALSE
ImmuniWeb	FALSE
Imperva	FALSE
Infosec	FALSE
KnowBe4	FALSE
McAfee	FALSE
Microsoft	TRUE
OccamSec	FALSE
OPSWAT	FALSE
Palo Alto Networks	TRUE
Perimeter 81	FALSE

---

QAwerk	FALSE
Raytheon	FALSE
ReversingLabs	FALSE
Sapphire	FALSE
Sectigo	TRUE
Secure Code Warrior	FALSE
SecurityHQ	FALSE
SEKOIA	FALSE
SlashNext	FALSE
Symantec	FALSE
ThreatLocker	FALSE
ThreatQuotient	FALSE
Trend Micro Inc.	TRUE
Oracle Corporation	TRUE
Juniper Networks	FALSE
Rapid7	TRUE
Verizon	TRUE

---

## A.4 Results from the threat report spreadsheet

The resulting threat report spreadsheet is scanned and generates the following combined list of actors: LuckyMouse, Mustang Panda, Gamaredon, Promethium, HAFNIUM, Wizard Spider, Bitwise Spider, Carbon Spider, Pinchy Spider, Pioneer Kitten, Deus, BlackShadow, Moses Staff, Nemesis Kitten, Wicked Panda, Doppel Spider, Aquatic Panda, Fancy Bear, Cozy Bear, nei, barf, inthematrix1, UNC2452, Astro Locker Team, MuddyWater, ITG23, LemonDuck, Tortoiseshell, Charming Kitten, Fox Kitten Parasite, ControlX, APT40, APT5, APT15, APT31, Sea Turtle, Kimsuky, Konni, Lazarus, APT32, Energetic Bear, APT28, Nobelium, Nickel, Thallium, Phosphorus, Cerium, Gadolinium, Strontium, Bromine, TeamTNT, Water Pamola, Earth Wendigo, Earth Vetala, Iron Tiger, PlugX, Sandworm, APT41, Ghostwriter, Black Shadow, REvil, CI0p, NetWalker, LockerGoga, MegaCortex, BlackCat, Conti, DEV-0537, Karakurt, DeathStalker, Candiru, Anonymous, TeamOneFirst, GhostSec, Against The West, NB65, Belarusian Cyber Partisans, KILLNET, XakNet and The Red Bandits





# ETDA Threat Group Cards

## B.1 Victim Sector Data

This section shows the number of actors in the ETDA Threat Group Cards that are identified to have targeted a particular sector.

*Table B.1: Number of actors per sector, sorted on the number of actors*

<b>Sector</b>	<b>Number of actors</b>
Government	192
None Provided	151
Defense	113
Financial	102
Energy	86
Telecommunications	85
Education	73
Media	71
Healthcare	58
Manufacturing	52
High-Tech	37
IT	36
Transportation	34
Technology	34
Aerospace	29
Aviation	28
Hospitality	27
Oil and gas	26
Engineering	24
Retail	23

Continuation of Table B.1

<b>Sector</b>	<b>Number of actors</b>
Pharmaceutical	23
NGOs	21
Construction	20
Shipping and Logistics	18
Think Tanks	17
Industrial	16
Embassies	14
Chemical	13
Utilities	13
Automotive	12
Food and Agriculture	12
Research	11
Law enforcement	9
Non-profit organizations	8
Casinos and Gambling	6
Entertainment	6
Mining	5
Online video game companies	5
Critical infrastructure	5
Maritime and Shipbuilding	5
Satellites	4
Petrochemical	3
Gaming	3