UiO : University of Oslo

Siri Bromander

Understanding Cyber Threat Intelligence - Towards Automation

Thesis submitted for the degree of Philosophiae Doctor

Department of Informatics
Faculty of Mathematics and Natural Sciences

mnemonic AS

University of Oslo

2021

Series of dissertations submitted to the
Faculty of Mathematics and Natural Sciences, University of Oslo

No. 2376

ISSN 1501-7710

Cover: Hanne Baadsgaard Utigard.
Print production: Reprosentralen, University of Oslo.

Preface

This thesis is submitted in partial fulfillment of the requirements for the degree of *Philosophiae Doctor* at the University of Oslo. The research presented here has been performed at mnemonic and at the Informatics Department at the University of Oslo under the supervision of professor Audun Jøsang and Martin Eian.

The work has been funded by mnemonic and the Research Council of Norway under the Industry PhD program and the project Threat Ontologies for Cyber Security Analytics (TOCSA).

The thesis consists of an introduction and overview of the research, accompanied by 6 publications produced as part of the project.

Acknowledgements

I would like to thank my supervisors prof. Audun Jøsang and Martin Eian for their guidance. Audun, your knowledge of and guidance in the academic world has kept me afloat. Martin, your knowledge and curiosity is inspiring and will continue to be what I aspire to.

I would like to thank Morton Swimmer, Lilly Muller and Vasileios Mavroeidis for invaluable discussions and cooperation in writing papers. Thanks to everyone in the Digital Security Group in the Informatics Department at UiO for support and cooperation.

I could not have completed this project without the contribution, support and presence of my colleagues at mnemonic Research and Development. I especially want to thank Geir Skjøtskift and Fredrik Borg for their overwhelming skills and expertise.

Friends and family generates balance. I would like to thank my family and friends for support and encouragement. In particular: thank you mamma and pappa for always being proud of me. Thank you Ella and Sofie, my ambitions are high because of you. Thank you Erik, it all works because of you.

Siri Bromander
Oslo, February 2021

List of Papers

Paper I

Siri Bromander, Audun Jøsang and Martin Eian 'Semantic Cyberthreat modelling'. In *Proceedings of the Eleventh Conference on Semantic Technology for Intelligence, Defense, and Security (STIDS 2016)*, pp. 74–78.

Paper II

Siri Bromander 'Ethical considerations in sharing cyber threat intelligence'. In: *Proceedings of the Tenth Norwegian Information Security Conference (NISK 2017)*, pp. 54–62.

Paper III

Vasileios Mavroeidis and Siri Bromander 'Cyber threat intelligence model: An evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence'. In: *Proceedings of the 2017 European Intelligence and Security Informatics Conference (EISIC 2017)*, pp. 91–98.

Paper IV

Siri Bromander, Lilly Muller, Martin Eian and Audun Jøsang 'Examining the Known Truths of Cyber Threat Intelligence - the case of STIX'. In: *Proceedings of the 15th International Conference on Cyber Warfare and Security (ICCWS 2020)*, pp. 493–XII.

Paper V

Siri Bromander, Morton Swimmer, Martin Eian, Fredrik Borg and Geir Skjøtskift 'Modelling Cyber Threat Intelligence'. In: *Proceedings of the 6th International Conference on Information Systems Security and Privacy (ICISSP 2020)*, pp. 273-280. DOI: 10.5220/0008875302730280.

Paper VI

Siri Bromander, Morton Swimmer, Lilly Muller, Martin Eian, Fredrik Borg and Geir Skjøtskift 'Investigating sharing of Cyber Threat Intelligence and proposing a new data model for enabling automation in knowledge representation and

exchange'. *Submitted for publication in ACM Digital Threats: Research and Practice (DTRAP)*, October 2020.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The digital transformation of society entails opportunities and increased efficiency, and can be beneficial for democracy and the economy in general. The sustainability of this evolution depends on adequate security on all levels, which unfortunately does not exist today.

Cyberattacks now have the realistic potential of causing serious harm to humans, their assets and business processes. Cybersecurity is aimed at blocking or mitigating such threats, by preventing, detecting and recovering from harmful incidents in cyberspace. The task of implementing adequate cybersecurity is already daunting, and becomes increasingly challenging every day. The threat landscape is continuously changing, it is often difficult to distinguish between friend and foe, and attribution of attacks is often uncertain. This is a situation of moving targets where existing security approaches used by 'white hats' quickly become outdated and ineffective against the next generation of attack strategies by 'black hats'.

Early security analytics tools such as NIDES[2] introduced in 1986 typically had rule-based expert systems to detect known types of network intrusions, as well as statistical anomaly detection components based on profiles of users and systems. In 1987, John McAfee released the first version of the VirusScan tool to detect malicious software based on virus signatures. This type of security technology served its purpose for many years. However, around 2005 it became clear that traditional attack detection and malware filtering based on fixed rules and signatures could no longer protect against more advanced attack methods. Machine-learning methods were then introduced to automatically classify events and potential security attacks[33]. Similarly, behavioral malware filtering was introduced to block malware based on what it does, not just on how it looks. However, attackers are naturally reacting to this trend, they become smarter and start to use unexpected and deceptive attack methods.

In the overwhelming majority of identified security incidents there is currently no understanding of who the threat actor is, why they attack or how they operate. The result is a lack of ability to make informed decisions when it comes to protection and countermeasures. The threat actors most often are not identified and made responsible for their actions, resulting in continuous criminal behaviour. We simply do not understand our opponent and can identify - if even that- only the results of the opponent's actions.

Too often security professionals are only observing the evidence of cyber-attacks – trails of information that are the long left-behind remnants from an attacker's past actions. When defending against these attacks, priority is understandably placed on recovering from the current attack, with identifying the attackers as an afterthought. The repercussion is that attackers are rarely

identified, seldom prosecuted, and able to operate with an almost free-reign.

Treat intelligence has played a key role in keeping networks secure for as long as computers have communicated. With the aim to collaboratively defend against the increasing threats in and from cyberspace, Cyber Threat Intelligence (CTI) has risen in popularity, and in correlation a growth of concepts and terms within the field is observable.

To have success in digital defense, we must exchange knowledge and experiences from fighting an increasing number of threats with an increasing amount of sophistication. The “footprints” or signatures of threat actors are mainly found through detecting them in networks where attacks have taken place. The availability of data about the actors is thus often tied to the infrastructure within which they operate. Hence, the information needed to defend preemptively is dependent on collaboration. Our ability to effectively describe a threat actor’s modus operandi influences our ability to analyze, share and consume it.

The leak code-named "Vault 7" by WikiLeaks[1], is the largest ever publication of confidential documents on the U.S. Central Intelligence Agency. These types of publications have made advanced tools for targeted cyber attacks available for the rest of the world, and can be argued to have empowered less advanced threat actors with abilities to strengthen their offensive capabilities.

With these types of advanced tools publicly available to anyone, detecting the signature of a particular tool used in malicious activity is by itself insufficient to attribute the activity to a specific threat actor. More details describing the use of the tools are necessary to know who one is fighting - and what to expect from the same actors in future scenarios.

We often use the abbreviation CTI for Cyber Threat Intelligence when discussing digital threat intelligence. We use the term CTI to refer to both the process of creating CTI and the shareable results of such processes[9].

1.1 Motivation and objectives

Cyber threat intelligence has emerged as an essential part of every cyber defense team across the digital world. We routinely rely on quality CTI in order to defend against the ever increasing amount of cyber threats. How effective we are depends on the tools and processes we use. To increase our efficiency and effectiveness, we need to improve our tools and processes. Automating tasks will free more time for analysts and hence increase their efficiency. Creating possibilities for computers to process and aid in threat intelligence operations will open up for use of computer-based analysis with larger computational power, which will enable automated sharing and processing, which in turn will increase the analysis capabilities and overall effectiveness in defending against cyber attacks.

With practical experience as an incident handler and threat intelligence analyst, I have noticed that the term CTI is being discussed and used in different

[1] https://wikileaks.org/ciav7p1/

settings with different meaning. Hence, there seems to be a certain confusion in the meaning of the term *cyber threat intelligence* in the community.

The specific meaning and interpretation of the term CTI provides the foundation for how the tools we use to conduct CTI are developed and used.

Understanding the challenge one want to solve before solving it has always been a leading star within computer systems development. Hence, a need to fully understand the use and consequences of usage of the term cyber threat intelligence became a key part of my research efforts when wanting to improve the way we operate our cyber threat intelligence efforts.

The objective of the research conducted therefore emerged: *Increase the effectiveness and efficiency of our cyber threat intelligence operations by improving the foundation of how we collect, analyze, enrich and share cyber threat intelligence.*

1.2 Research questions

The following research questions have been specified and used for the duration of our research activities:

RQ1 What is a meaningful definition and interpretation of CTI?

RQ2 How is CTI conducted in practice?

RQ3 Which standards and best practices are relevant, and how are they used?

RQ4 Can we automate tasks related to CTI?

1.3 Research method

In order to answer the research questions, we used the following research approach:

1. Literature review. We conducted a literature review in order to determine the current status within the field of structuring cyber threat intelligence. Literature reviews were done using Google Scholar, reference lists of relevant industry and academic publications and also by normal Google searches where relevant.

2. Data collection and analysis. Through a combined approach drawing on ethnography[16] from the CTI community, interviews and questionnaires we examined current practices within CTI.

    - A questionnaire was created and used to study how practitioners understand and execute CTI work. Questionnaire design and execution were done in three stages: two rounds of testing with limited test groups, and improvement of the questionnaire design for each step. The design was initially designed with the guidelines in [21] and [37].

- Semi-structured interviews were conducted to verify the questionnaire findings. The questionnaire is found in Appendix A and both the questionnaire and the semi-structured interview guide are published on GitHub[2].

3. Proposing a data model. We applied theoretical formalism to the description of CTI when creating our data model. The creation of the data model followed an iterative method, implementing and testing the data model after each improvement, adding new data sources alongside new adjustments.

4. According to Oxford Languages[22], the definition of ontology within computer science is *a set of concepts and categories in a subject area or domain that shows their properties and the relations between them.* The data model resulting from step 3 can be classified as an ontology in this definition. The field of ontology gives a range of tools for use within ontology development and evaluation. The evaluation of the data model was hence done using known evaluation approaches from ontology[24].

1.4 Structure of thesis

This work is written in the form of a cumulative thesis, compiling the results of six research papers.

The thesis consists of two parts, where Part 1 contains a summary description of the research project and results through Chapter 1-4 and Appendices A and B. Part 2 contains the publications written and published during the research project.

[2]https://github.com/sbrom/sharingCTI

Chapter 2

Background

This section provides background information needed for understanding the rest of the thesis. The background information is describing related and relevant work and research results.

2.1 Cyber Threat Intelligence

Cyber threat intelligence, often referred to as CTI, is a concept which receives increasing interest and attention. However, CTI is not new. An early example of CTI is the Phage email list[39] created as a response to the Morris Worm[38] in 1988. As the outbreak started, technical personnel exchanged information on how the worm had been identified, how to handle infections, and how to protect against infection. This type of collaboration is exactly what we now refer to as CTI, meaning the people behind the mailing list were pioneers in the field of CTI.

As defined by Gartner *Threat intelligence is evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice, about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject's response to that menace or hazard* [14]. We emphasize *evidence-based knowledge* as CTI needs to be based upon evidence in order to be trusted. It is also of importance that this type of knowledge can be used for mitigating threats, and hence needs to be actionable. The Allied Joint Procedures published and used by NATO[1], concur with this definition of threat intelligence, elaborating on the differences between threat data, threat information and threat intelligence. Threat data can be processed to become threat information, which needs structure and adoption for a given audience in order to qualify as threat intelligence.

CTI operations largely consists of collecting data, information and knowledge from different sources, much of which are from outside your own organization. Some are closed sources, both paid and free and often relying on networks and acquaintances, and some are open sourced intelligence (OSINT).

Chismon and Ruks [13] proposed a model representing technical, operational, tactical and strategic CTI and its properties, and to what degree they were detailed and of long-term use. This is illustrated in Figure 2.1. We find that few published initiatives are covering the full range of relevant knowledge, but conducting research in a subsection of CTI.

The model of Chismon and Ruks illustrates that the more detailed the knowledge is, the more certainty about the threat actor's presence and identity can be obtained, and that the more robust (long term) it is, the longer the knowledge is useful for the defenders. Tactical threat intelligence, consisting of

tactics, techniques and procedures, is of most value in the attempt to detect and prevent future attacks. Yet, this is one of the least developed areas within CTI, which we assume has to do with the availability of structured and well defined data available in this area. Research within tactical CTI seems to be given the most attention with the field of research.

Working with CTI is in the border between IT and management. My impression is that standards used by non-technical personnel are less concerned with consistency as humans will be part of reading and interpreting the content. In contrast, IT staff find standards to be valuable when they make no room for adjustments or inconsistency. IT staff typically use standards to allow for automation. Within CTI we find vast amounts of data, the challenge is to collect, merge and analyze all of it, and also be able to share the results of these analyses. Automation is key to do this at large scale, and in this context, standards with little room for adjustments or inconsistencies in use is preferable.

Collaboration on digital defense is a global effort. English language is often used for communication. There are however communities and published CTI which are not reached due to the lack of skills in a wider range of human languages. The research presented in this thesis has solely worked with English language.

Trust and confidence are important aspects of CTI, especially when sharing. These aspects are given limited focus in the work presented in this thesis, even though they have not been neglected. The topic was visited in the [9] research paper, and the ACT platform and the data modeling presented in [10] have been created with this as a prerequisite. This topic is an open research question which

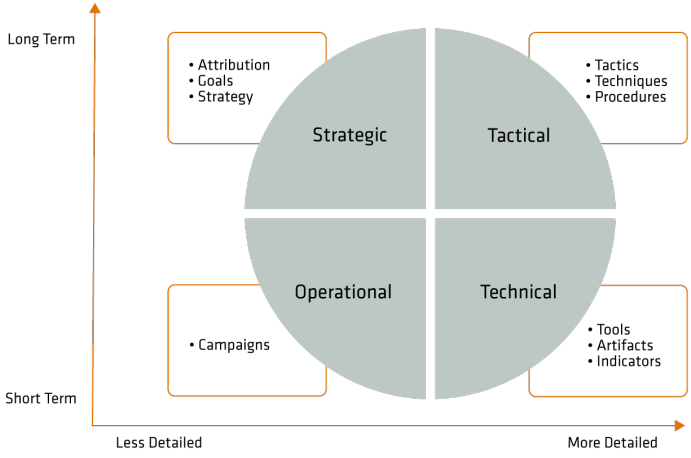

Figure 2.1: Chismon and Ruks model of CTI [13]

is outside the scope of this thesis. Further investigation into this topic is left to future research.

2.2 Available standards, models and tools

The following section will chronologically present the most important results in terms of structuring and explaining security related concepts relevant to CTI. They serve as a foundation of the current status within the CTI communities and also the foundation for the work presented in this thesis.

2.2.1 Models and standards

In 2011 Lockheed Martin proposed the *Cyber Kill Chain* model [19] which gives an overview over the phases an advanced persistent threat will traverse in order to mount a successful attack. The model consists of seven phases, all representing a stage where the attack can be "killed" in order to prevent the attacker to gain success. The further up the chain an attack is "killed", the less foothold the attacker will gain. To be able to defend in the early stages of the attack one cannot rely on detection as there is no traces of the attacker in the defender's environment.

The Diamond Model is explained in detail in [12]. The model was already mentioned as part of the foundation for ontology development by Obrst et al. in 2012 [32] having been briefed on the Diamond Model in 2010 [20] which appears to be the original source of the model and the foundation for the 2013 publication. The Diamond Model consists of four corners, Victim, Infrastructure, Capability, and Actor (the one threatening the victim), which account for all the major dimensions of a malicious cyber threat[32]. The model complements the Kill Chain-model with a broader perspective of intrusion activity.

In 2012, Sean Barnum et al. presented the Structured Threat Information Expression (STIX) [4]. It was an XML-based standard covering relevant concepts to CTI. STIX was maintained by MITRE until 2017, when OASIS took over and published the STIX standard in version 2 [30], now as a JSON-based format with improvements based upon usage and experience, including increased possibility to express relationships. To this date OASIS is still maintaining STIX, and the standard is now in version 2.1. The latest version was made public in 2020 and contains several valuable improvements adhering to the development of the field. Future research projects should evaluate the impact this has on practitioners' sharing of CTI.

To exchange STIX content, MITRE and subsequently OASIS, have developed and maintain Trusted Automated Exchange of Intelligence Information (TAXII), which is an application protocol for exchanging CTI over HTTPS [1]. This protocol is using STIX, and is therefore not part of the structuring of CTI, rather a tool for rapid exchange. The tool is closely related to STIX, and is therefore mentioned here for completeness.

[1] https://oasis-open.github.io/cti-documentation/taxii

The Pyramid of Pain has become a frequently cited model within the security community, published as a blog post by David Bianco in 2014[6]. The Pyramid of Pain reflects the pain a defender can inflict upon the attacker corresponding to different characteristics. The higher up the pyramid, the more painful it will be for the attacker to escape detection. The top of the pyramid is "TTPs", tactics, techniques and procedures, the tactical CTI referred to by Chismon and Ruks in Figure 2.1.

Another blog post which has received much attention and is often referred to is the Detection Maturity Level Model (DML) by Ryan Stillions, first presented in 2014[40]. The model was slightly extended and included in [8] as seen in Figure 2.2. The model gives a hierarchical presentation meant for evaluation of an incident response function's ability to consume and act upon received threat data, threat information and CTI. The version published in our paper is modified to aid in structuring of the same type of content.

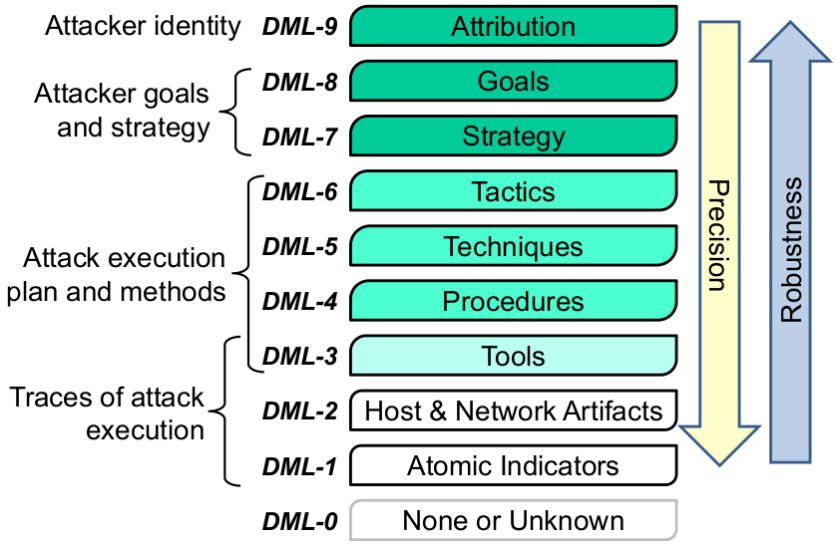

Figure 2.2: Modified DML model as found in [8]

There have been several studies on evaluation and comparing CTI relevant standards, taxonomies and languages. As part of this thesis, [23](Paper III) gives a good overview until 2017. At the time of writing the last publication doing this type of evaluation is Ramsdale et al.[34], which also describes the work published after our 2017 publication. This latest paper concurred with the findings in this thesis and the 2019 publication by Sauerwein et al.[35]: there is no agreed and standardized way of sharing CTI. The most referred and used standard is STIX, but the usage and implementations are varying and not consistent. Most CTI is shared in a variety of formats, often JSON, especially created for a given scenario. Also, which is an important finding, they emphasize the need for "origin" of received CTI, which corresponds well with the data model presented

later in this thesis.

2.2.2 Tools, knowledge-bases and platforms for automation

In 2011, the Malware Information Sharing Platform (MISP) project started. The project has evolved to become the MISP Threat Sharing Platform and is a community-driven project with efforts of development of the platform and sharing of threat information. MISP[45] is found on GitHub[26]. One of the benefits with MISP is their focus on rapid exchange of information, where the structure is tied to their platform, and therefore, the MISP platform is by some seen as an alternative to using STIX. MISP allows for exportation to STIX format.

MITRE ATT&CK[41] was published in 2016 as an online knowledge base of adversary tactics and techniques linked to tools and threat actors based upon real world observations. The content is published by the ATT&CK team, and the information sources are always public. MITRE also provides relevant and freely available tools for use within CTI, like Caldera, Caret, Car, ATT&CK Navigator and Tram, all of which helps a user to take advantage of the ATT&CK knowledge base[28]. ATT&CK is as far as we can see the first large attempt to publish a somewhat structured description of tactics and techniques, which by itself is a large step in the right direction of structuring tactical CTI. Their newer publication of sub-techniques, and non-structured descriptions of procedures is a large step into tying the lower level indicators to the higher abstraction levels of the DML model. This work is an important next step for future research.

In 2017, the Semi-Automated Cyber Threat Intelligence (ACT) project [2] was initiated and is now a platform allowing for consumption, analysis, enrichment and sharing of CTI. The research presented in this thesis has been part of the development of the ACT platform as described in [10].

During the Spring of 2019, OpenCTI[3] was published. The platform has similar ideas as the ACT project, building a graph based platform with strong query possibilities, where combination of sources is handled and where the data model enables linked data. The chosen open source license[3] deviates from the ACT platform[4].

There is a range of threat intelligence platforms which are in use, but that are closed source, without publications or that are subject to payment for receiving descriptions. The selection given in the above is a collection which I find to be sufficient to give an overview of the context of which the research presented in this thesis was conducted. There has been done some research on the different platforms for CTI, including commercial products. The results until 2017 can be found in Sauerwein et al.'s publication[36]. There are several interesting findings from this research, and their Key Finding 2 marked an important influence on our research: STIX is the de-facto standard for describing threat intelligence. This has been further studied and in [35] the same authors finds that only 22

[2] https://github.com/mnemonic-no/act-platform
[3] https://github.com/OpenCTI-Platform/opencti/blob/master/LICENSE
[4] https://github.com/mnemonic-no/act-platform/blob/master/LICENSE

percent of the identified security sources rely on standardized representations of security information.

2.3 Research directions

The field of CTI is influenced by the current available standards, tools and platforms. These serve as the foundation for structuring CTI, and have influenced many of the research directions which are relevant to the research presented in this thesis. The following section describes the research directions applied to the field of CTI, relevant to the work presented in this thesis.

2.3.1 Ontologies, taxonomies and vocabularies

An ontology, in the field of computer science, is a formal description of concepts and how they are related to each other, often referred to as classes and properties. In turn, ontologies provide computational meaning to data by building semantic and logic relations in the ontology which enables us to use reasoning methods (such as induction or deduction) on our data in our knowledge base. While there are many implementations of knowledge bases and ontologies, the World Wide Web Consortium (W3C) chose a triplet model for facts and calls this the Resource Description Framework (RDF)[5]. RDF also allows us to implement the RDFS schema language[6] and OWL[7], the web ontology language.

Description logics (DL) are formal languages designed for knowledge representation and reasoning of which most of the languages are decidable fragments of first order logic[17]. DL are the basis of the W3C definition of OWL[18], even though dialects of OWL which does not base upon DL exists[31].

Ontology can be argued to be a subfield of Artificial Intelligence (AI), due to a computer's capability to reason and infer new knowledge through applying an ontology on a given set of data.

When working not only with the concepts within a field, but also with the relationships between them, it is natural to apply ontology-based techniques for modeling and analysis. Ontology gives the opportunity to express knowledge, not only data, and when structuring Cyber Threat Intelligence which we argue to be knowledge, the use of ontology is a benefit.

Ontology-based modeling and analysis are very general and can be valuable in many fields of study. In the field of CTI, we have found there to be several directions that have received recent attention and that are relevant to our objective of automation within CTI, as summarized in the following list:

1. Definition of concepts - what do the concepts mean and how do they relate to each other? In the field of CTI common concepts like "malware", "file" and "campaign" are subject to different interpretations.

[5] https://www.w3.org/TR/rdf-concepts/
[6] https://www.w3.org/TR/rdf-schema/
[7] https://www.w3.org/TR/owl2-overview/

2. Taxonomy and vocabulary- creating a taxonomy with connected vocabularies to be used within the field. The STIX vocabularies and the ATT&CK knowledge base are examples of this development.

3. Reasoning - infer new knowledge based on available content. Little results have been seen in the fields of CTI based upon logical descriptions alone, the use of rule languages like SWRL[44] is more common. Examples are the ontology-based cybersecurity framework for the internet of things by Mozzaquatro et al.[29] and SIMON: Semantic Inference Model for Security in Cyber Physical Systems using Ontologies published by Ventaka et al. in 2019[43].

Since the publication of [23] (Paper III) there have been several publications contributing to this field. Examples of ontologies which are evolving and have been maintained since 2017 are Unified Cyber Ontology (UCO)[5] which set the foundation for CASEwork[8] and Unified Cybersecurity Ontology (UCO)[42] (last updated in 2019) which provides a mapping between different online sources of CTI to enable cross-searches. UTIM [9] is also created and maintained in order to map different concepts within CTI and allow for searches across distributed sources. UTIM is developed in parallel with the work presented in this thesis and it is hoped that one day data from our model will be freely interchangeable with data modeled with UTIM. Menges et al. [25] propose a unified CTI model which covers the concepts used within CTI. This work should be tested with data and in practical use. The lack of defined differences in relationships in this model may be a challenge when expressing knowledge and makes it closer a taxonomy than an ontology.

2.3.2 Machine Learning (ML)

Machine learning can also be argued to be a subfield of AI, due to the possibility of computers performing classification similar to human reasoning capabilities. In ML the use of advanced statistics and large computational power are used to make predictions and derive new knowledge about the available data. ML can be compared to ontology as the choice of letting machines and statistics find the patterns and from that describe the world, instead of describing the world through a language of descriptive logic, and understanding the available data through this description.

Machine learning is relevant to the present work in this thesis because ML and ontologies are possible to combine. An example is to classify whether a domain name is "benign" or "bad" which would be helpful for technical and possibly tactical CTI. It is possible to use descriptive logic to initialize a process of weak supervision and then a knowledge base drawn partially from the same description to serve as the supervisor for the ML engine. These types of techniques are not investigated in this thesis, but is suggested as future research.

[8] https://github.com/casework/CASE
[9] https://github.com/mswimmer/utim

Chapter 3

List of Research Papers

3.1 Included Papers

The following list describes the papers included in this thesis and gives a brief summary of their content.

Paper I is an initial publication describing the intended direction and focus of my research project. The paper contributes with models and examples explaining the value and place for semantics in the field of CTI. It creates an expanded version of the well known DML (Detection Maturity Level) model of Stillions[40] which serves as an important foundation for the concepts found and used within CTI.

Paper II focuses on the ethical aspects of practical work with sharing of CTI. The paper evaluates the different options threat intelligence personnel faces when met with the choice and decision of sharing CTI. The paper concludes that the question of sharing is typically not suited for predefined answers as the situations are often complex with small changes making large impacts on the consequences. The main contribution of the paper is the description of the complexity of the choices threat intelligence personnel are facing and why the easiest choice can be to not share, and may with this explain some of the reason why there is less sharing than the collective defense may desire.

Paper III gives a rich overview of the ontologies, standards and taxonomies found within the field, and a comprehensive overview of the available literature within the field of research. The main contribution of this article is that there is no consensus in the field, but an increasing amount of research efforts within ontology.

Paper IV investigates how threat intelligence personnel work with CTI in practice. The background for the research was the experience of meeting "known truths" when meeting practitioners from different communities, and aiming at finding the answer to these contradicting views of how practitioners understands and conduct their work. The main contribution of this article is the finding that STIX (Structured Threat Information Expression) is not well suited as a foundation for automation within CTI, and that even though many claim to be using STIX, they do not actually create or consume STIX in a standardized way and as part of their threat intelligence efforts. This paper was written as a collaboration of researchers in the border between technology and international relations.

The paper includes a context and a description of consequences of the findings exceeding the technical world.

Paper V presents a data model for CTI well suited for consuming, analyzing and sharing CTI without loss of information or knowledge. The main contribution of this work is an open-sourced implementation of the data model, and the description of how the model handles known issues within the field of CTI. The work leading to the publication was a collaboration between different communities, research departments and security vendors which is a strength for the results.

Paper VI has combined paper Paper IV and Paper V, and added an evaluation of the suggested data model based upon relevant literature on ontology evaluation. Both of the combined papers were updated and content that was excluded due to length limitations in the conference proceedings was included.

3.2 Other contributions

- Presentation at NSM sikkerhetskonferansen 2017. "Threat Intelligence: samarbeid for å beskytte mot fremtidige angrep".
- Presentation at the 29th Annual FIRST Conference 2017. Title "Threat Ontologies for Cybersecurity Analytics (TOCSA)". The presentation was done by my supervisor Martin Eian as I was held back from traveling for personal reasons.
- Key Note presentation at the 18th European Conference on Cyber Warfare and Security. "Cyber Threat Intelligence (CTI)".
- Presentation "Threat Ontologies" at the Oslo 2018 FIRST Technical Colloquium.
- Presentation of TOCSA and threat intelligence in general on breakfast seminar given by Tekna in February 2019. "Trusseletterretning i det digitale rom".
- Poster and lightning talk at the DFRWS EU conference 2019. "Structuring Cyber Threat Intelligence".
- Lightning talk at the 31st FIRST Annual Conference 2019. "How do we share CTI?".
- Presentation at NG-SOC workshop in Canterbury, August 2019. "ACT: Cyber Threat Intelligence Platform".
- Presentation at Security Divas 2020. Title "Trusseletterretning i det digitale rom – mer enn data".

- Program committee of the following conferences: NISK 2017, NordSec 2018, CyberHunt 2019, CyberHunt 2020 and IFIP SEC 2021.
- Supervisor for Master student Mari Grønberg : "An Ontology for Cyber Threat Intelligence" [15].

Chapter 4

Conclusion

This chapter returns to the research questions formulated in Section 1 and discusses them in connection with the contributions of this thesis. This chapter also suggests directions for future work.

4.1 Summary of contributions

Research questions	Papers contributing to answer
RQ1	Paper I, Paper II, Paper III, Paper IV
RQ2	Paper I, Paper II, Paper IV
RQ3	Paper I, Paper III, Paper IV, Paper V, Paper VI
RQ4	Paper IV, Paper V, Paper VI

Table 4.1: Papers answering research questions.

4.1.1 RQ1: What is a meaningful definition and interpretation of CTI?

CTI has grown to become a wide field of practice and the term is used for both the processes and the results of research revolving around threat actor activity[9] (Paper IV).

Definitions and usage of the term CTI are not always in agreement. The term is used for both data, information and knowledge in addition to the process of creating and using such content[9] (Paper IV).

There are several definitions of the term CTI and the published definitions are mostly agreeing that CTI is more than data and information. Both knowledge and intelligence are used as a description of the level of processing data and information needs to reach before it can be called actual CTI[23, 9] (Paper III, Paper IV).

We should strive for knowledge as opposed to data and information (which we need to create knowledge), and continuing to treat the field as a search for knowledge will keep the field evolving[10] (Paper V).

4.1.2 RQ2: How is CTI conducted in practice?

The process of CTI includes consumption, analysis, enrichment and sharing of content describing threat actor activity. All areas influence each other.

Sharing and receiving CTI relies on the access to quality sources. The access to closed sources is still influenced by which relationships the CTI function of an organization has. Sharing CTI is often conducted in a complex context and the decision to not share is easier to make than that of sharing[7] (Paper II). Confidentiality of CTI is often referred to as a reason for not sharing[9] (Paper IV).

The complexity of CTI lies partially with the range of different sources. This creates difficulties for enrichment and analysis as the different sources are not presenting the same CTI in the same way[9] (Paper IV). Also, trust in sources and confidence in given CTI influences how to value consumed CTI[11] (Paper VI).

CTI is still to a large degree unstructured, both in sharing and storage. There is knowledge of, but not extensive use of, standards. The understanding of how and what a standard means varies. The lack of standardization reduces the ability to automate[9] (Paper IV).

4.1.3 RQ3: What standards and best practices are relevant, and how are they used?

The development of the field of CTI has been influenced by both industry and academic publications. The usage and popularity of them vary considerably. An overview of the main contributions is given in [23] (Paper III) and usage of the most referenced and popular standard, STIX, has been investigated in [9] (Paper IV).

No standard or best practice publication is covering the whole field of CTI. Either it is limited by the type of CTI it is covering(strategic, tactical, operational or technical), or by the intended application(consuming, enriching, analyzing or sharing).

No single standard or best practice publication has gained practical application by all practitioners within the CTI community. The importance of using them correctly may not be well enough presented or understood by the range of users, but it is just as possible that the publications are not solving the challenges faced within the field and therefore gain less momentum than desired. Significant contributions to evolving the field of CTI can be exemplified by the Cyber Kill Chain Model[19], the DML model[40], STIX[4] and MITRE ATT&CK[27].

4.1.4 RQ4: Can we automate tasks related to CTI?

Effective automation of CTI relies on having a mature level of standardization which currently does not exist[9, 10, 11] (Paper IV, Paper V, Paper VI). We have shown a possible data model implementation which enables consumption of unstructured content to become structured[10] (Paper V). With this structure we show that automation of enrichment, analysis and sharing can be done.

The main strengths of the proposed data model are its coverage of concepts and strictness in terms of both enforcing data ingestion solely with facts (triplets) and the absence of open fields. The data model is not implemented using a

standardized ontology language, which limits the usage of available reasoners to verify the consistency of the data model and infer additional knowledge[11] (Paper VI).

4.2 Future work and open research questions

Further development of tactical and strategic CTI will lead to a better understanding of the threats we are facing. Looking to the DML model[40], the definition of "Procedures" is missing, which could connect a single threat actor to sequences of tasks and to use of specific tools in a given scenario. This information is missing in a structured form today and should be part of future research efforts to standardize and automate CTI.

In the short term, a further standardization and development of our understanding of the tactical CTI is relevant. Building on the foundation laid by efforts like STIX[30] may be useful, as the current understanding of CTI is impacted by available efforts and results. Our work to define the commonly used terms and concepts within tactical CTI must be continued. This work lays the foundation for achieving practical automation within tactical CTI.

Further, the use of ontologies to better enable reasoning on available CTI may lead to more effective and efficient use of currently available CTI professionals of which we know there is a limitation. This requires creation of ontologies with the aim of inferring new knowledge, which is a step forward in proceeding from only using ontologies for agreeing on terms and concepts. One feasible approach is to describe the field of CTI using descriptive logics in a language like Web Ontology Language (OWL), maybe with help from rule based languages like Rule Markup Language for the Semantic Web (RuleML).

In a longer term, creating bridges which enable utilization of low-level indicators in the work with higher level CTI, like tactical and strategic CTI, should be a goal. Making bridges to discover and describe causality across the different layers within the field of CTI should also be a goal. Finally, it should be a goal to utilize the knowledge we have on a strategic level to improve our technical abilities, and making sure technical observations influence the evaluation of probability and consequences of threat actor activity instead of human evaluation alone. This will allow for impact on strategic choices through for example risk management. The connections in these layers are not well investigated, but could lead to utilizing threat intelligence in strategic decision-making in an organization, effectiveness in a world where CTI professionals are limited in numbers, and deeper understanding of the currently available CTI.

CTI will continue to be performed in large by human analysts. Further research into how this work is performed and how the different tools and available standards are utilized is important in steering the research directions within the CTI community. In specific, an evaluation of actual use of our ACT data model is relevant input to the further development as usability of the data model is best evaluated through evaluation of actual usage.

Bibliography

[1] AJP, N. S. 2.0 allied joint doctrine for intelligence, counterintelligence and security doctrine.

[2] Anderson, D., Frivold, T., and Valdes, A. Next-generation intrusion detection expert system (nides): A summary.

[3] ANSSI, Luatix, and CERT-EU. The OpenCTI Platform. https://github.com/OpenCTI-Platform, 2019.

[4] Barnum, S. Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX™). *MITRE Corporation 11* (2012).

[5] Barnum, S. Unified cyber ontology. https://github.com/ucoProject/uco, 2016.

[6] Bianco, D. The pyramid of pain. http://detect-respond.blogspot.no/2013/03/the-pyramid-of-pain.html, 2014.

[7] Bromander, S. Ethical considerations in sharing cyber threat intelligence. *NISK Journal* (2017), 54–62.

[8] Bromander, S., Jøsang, A., and Eian, M. Semantic cyberthreat modelling. In *STIDS* (2016), pp. 74–78.

[9] Bromander, S., Muller, L. P., Eian, M., and Jøsang, A. Examining the" known truths" in cyber threat intelligence–the case of stix. In *Proceedings of the International Conference on Cyber Warfare and Security* (2020), Academic Conferences International Limited, pp. 493–XII.

[10] Bromander, S., Swimmer, M., Eian, M., Skjøtskift, G., and Borg, F. Modeling cyber threat intelligence. In *Proceedings of the 6th International Conference on Information Systems Security and Privacy - Volume 1: ICISSP* (2020), INSTICC, SciTePress, pp. 273–280.

[11] Bromander, S., Swimmer, M., Muller, L., Jøsang, A., Eian, M., Borg, F., and Skjøtskift, G. Investigating sharing of cyber threat intelligence and proposing a new data model for enabling automation in knowledge representation and exchange, 2020.

[12] Caltagirone, S., Pendergast, A., and Betz, C. The diamond model of intrusion analysis. Tech. rep., DTIC Document, 2013.

[13] Chismon, D., and Ruks, M. Threat intelligence: Collecting, analysing, evaluating. *MWR InfoSecurity Ltd* (2015).

[14] Gartner. Definition: Threat intelligence. https://www.gartner.com/en/documents/2487216/definition-threat-intelligence, 2013.

[15] Grønberg, M. An ontology for cyber threat intelligence. Master's thesis, 2019.

[16] Hammersley, M. Ethnography. *The Blackwell encyclopedia of sociology* (2007).

[17] Horrocks, I. Description logics in ontology applications. In *International Conference on Automated Reasoning with Analytic Tableaux and Related Methods* (2005), Springer, pp. 2–13.

[18] Horrocks, I. Description logic: A formal foundation for ontology languages and tools.

[19] Hutchins, E. M., Cloppert, M. J., and Amin, R. M. Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. In *6th International Conference on Information Warfare and Security (ICIW2011)* (2011).

[20] Ingle, J. Organizing intelligence to respond to network intrusions and attacks. In *Briefing for the DoD Information Assurance Symposium* (2010).

[21] Krosnick, J. A. Questionnaire design. In *The Palgrave Handbook of Survey Research*. Springer, 2018, pp. 439–455.

[22] Languages, O. Definition: Ontology.

[23] Mavroeidis, V., and Bromander, S. Cyber threat intelligence model: an evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence. In *2017 European Intelligence and Security Informatics Conference (EISIC)* (2017), IEEE, pp. 91–98.

[24] McDaniel, M., and Storey, V. C. Evaluating domain ontologies: Clarification, classification, and challenges. *ACM Computing Surveys (CSUR) 52*, 4 (2019), 1–44.

[25] Menges, F., Sperl, C., and Pernul, G. Unifying cyber threat intelligence. In *International Conference on Trust and Privacy in Digital Business* (2019), Springer, pp. 161–175.

[26] MISP. The MISP platform. https://github.com/MISP/MISP, 2019.

[27] MITRE. Adversarial Tactics, Techniques and Common Knowledge (ATT&CK). https://attack.mitre.org/.

[28] MITRE. MITRE ATT&CK GitHub. https://github.com/mitre-attack, 2021.

[29] Mozzaquatro, B. A., Agostinho, C., Goncalves, D., Martins, J., and Jardim-Goncalves, R. An ontology-based cybersecurity framework for the internet of things. *Sensors 18*, 9 (2018), 3053.

[30] OASIS CTI TC. Structured threat information expression (STIX™) 2.0. https://oasis-open.github.io/cti-documentation/, 2017.

[31] Obrst, L. Ontologies for semantically interoperable systems. In *Proceedings of the twelfth international conference on Information and knowledge management* (2003), pp. 366–369.

[32] Obrst, L., Chase, P., and Markeloff, R. Developing an Ontology of the Cyber Security Domain. In *STIDS* (2012), pp. 49–56.

[33] Pinto, A. Secure because math: A deep-dive on machine-learning-based monitoring. *Black Hat Briefings USA* (2014).

[34] Ramsdale, A., Shiaeles, S., and Kolokotronis, N. A comparative analysis of cyber-threat intelligence sources, formats and languages. *Electronics 9*, 5 (2020), 824.

[35] Sauerwein, C., Pekaric, I., Felderer, M., and Breu, R. An analysis and classification of public information security data sources used in research and practice. *Computers & Security 82* (2019), 140–155.

[36] Sauerwein, C., Sillaber, C., Mussmann, A., and Breu, R. Threat Intelligence Sharing Platforms: An Exploratory Study of Software Vendors and Research Perspectives.

[37] Smyth, J. D., Dillman, D. A., Christian, L. M., and Stern, M. J. Comparing check-all and forced-choice question formats in web surveys. *Public Opinion Quarterly 70*, 1 (2006), 66–77.

[38] Spafford, E. H. The internet worm program: An analysis. *ACM SIGCOMM Computer Communication Review 19*, 1 (1989), 17–57.

[39] Spafford, G. Phage list. http://securitydigest.org/phage/.

[40] Stillions, R. The DML Model. http://ryanstillions.blogspot.com/2014/04/the-dml-model_21.html, 2014.

[41] Strom, B. E., Applebaum, A., Miller, D. P., Nickels, K. C., Pennington, A. G., and Thomas, C. B. Mitre att&ck: Design and philosophy. *Technical report* (2018).

[42] Syed, Z., Padia, A., Mathews, M. L., Finin, T., and Joshi, A. UCO: A Unified Cybersecurity Ontology. In *Proceedings of the AAAI Workshop on Artificial Intelligence for Cyber Security* (2016), AAAI Press.

[43] Venkata, R. Y., Maheshwari, R., and Kavi, K. Simon: Semantic inference model for security in cyber physical systems using ontologies. *ICSEA 2019* (2019), 61.

[44] W3C. Semantic Web Rule Language. https://www.w3.org/Submission/SWRL/, 2004.

[45] Wagner, C., Dulaunoy, A., Wagener, G., and Iklody, A. Misp: The design and implementation of a collaborative threat intelligence sharing platform. In *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security* (2016), ACM, pp. 49–56.

Appendices

Appendix A

The Questionnaire

The questionnaire was published in a web solution, Nettskjema, provided by the University of Oslo. Based upon answers, the respondent only received relevant questions. The complete set of questions are given in the following page.

Questionnaire: Sharing Cyber Threat Intelligence

We are attempting to better understand how (cyber) threat intelligence is shared within the security community. This questionnaire is prepared to give data that may give valuable insight. We will publish all results of our data analysis.

Definition of flat file:
"A file having no internal hierarchy. Typically email content, .txt, .csv, flat json."
"A flat file contains records that have no structured interrelationship. A flat file typically consists of a text file, from which all word processing or other structure characters or markup have been removed."

If you want to share data with no hierarchy or interrelationships, you may use flat files. If you want to share information or knowledge this would arguably require the use of describing relationships between data points, and this would require some other way of communicating. STIX gives the opportunity to do this (but is not the only option).

Part 1: About the respondent

1. What is your role in your organization (examples can be incident responder, threat hunter, security analyst)?
2. Are you in a role where sharing cyber threat intelligence is part of your role/tasks?
3. What is the size of your organization (approx. number of employees)?
4. What sector do you represent?
5. Which country are you from?

Part 2: About sharing CTI

1. Are you or your organization a producer or a consumer of threat intelligence?
2. In what format was the last piece of threat intelligence you SHARED with others?
3. In what format was the last piece of threat intelligence you CONSUMED?
4. Please estimate what percentage(%) of your consumed threat intelligence over the last 6 months was WITHOUT structured interrelationships within the data:
5. Please estimate what percentage(%) of your consumed threat intelligence over the last 6 months was WITH structured interrelationships within the data:
6. Do you STORE the consumed threat intelligence in a structured and easily accessible way?
7. Many use threat intelligence directly for defence purposes, for example, directly in block lists. Further analysis of the threat intelligence may add value. Do you use all your consumed threat intelligence in your organization for analysis?
8. What is the single most common reason, related to personal or professional circumstances, for you NOT to share threat intelligence?

Part 3: About STIX

1. Have you used STIX (any version) in the past 6 months?
2. If you use STIX, which version are you mainly using?
3. Have you manually consumed a STIX file in the past 6 months?
4. Have you automatically consumed a STIX file in the past 6 months?
5. If automatic consumption: do you manage to consume all information enclosed in any STIX file?
6. Have you created a STIX file in the past 6 months?
7. List of all SDOs and Relationships in STIX 2.1. with “used”/”not used” options.

Published with nettskjema.no in june-august 2019.

Appendix B

Data model

The data model can be represented as a graph as seen in Figure B.1 auto-generated with the use of Graphviz [1].

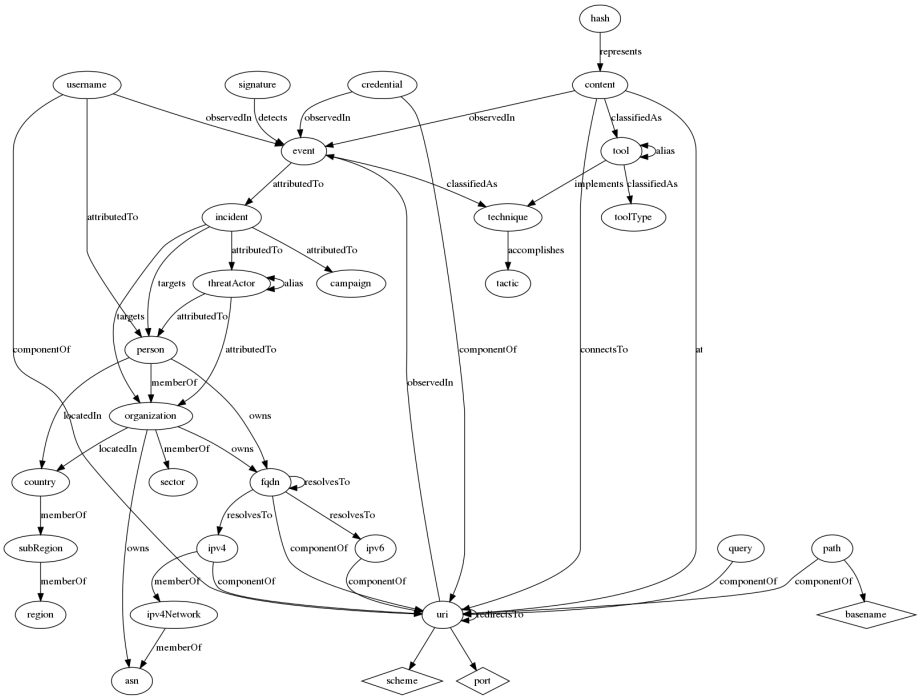

Figure B.1: The complete datamodel.

[1] https://graphviz.org/

Papers

Paper I

Semantic Cyberthreat modelling

Siri Bromander, Audun Jøsang, Martin Eian

Proceedings of the Eleventh Conference on Semantic Technology for Intelligence, Defense, and Security (STIDS 2016).
George Mason University, Fairfax VA, USA.
ISSN: 1613-0073, CEUR Workshop Proceedings, volume 1788.

Semantic Cyberthreat Modelling

Siri Bromander
mnemonic
Norway
siri@mnemonic.no

Audun Jøsang
University of Oslo
Norway
josang@ifi.uio.no

Martin Eian
mnemonic
Norway
meian@mnemonic.no

***Abstract*—Cybersecurity is a complex and dynamic area where multiple actors act against each other through computer networks largely without any commonly accepted rules of engagement. Well-managed cybersecurity operations need a clear terminology to describe threats, attacks and their origins. In addition, cybersecurity tools and technologies need semantic models to be able to automatically identify threats and to predict and detect attacks. This paper reviews terminology and models of cybersecurity operations, and proposes approaches for semantic modelling of cybersecurity threats and attacks.**

I. Introduction

When security incidents occur there is typically limited understanding of who the threat agent is, why they attack and how they operate, which makes it difficult to make well informed decisions about countermeasures. Threat agents who are not identified and made responsible for their actions will continue their criminal behaviour. When we do not understand the attacker we can only see - if even that- the results of the attacker's actions. Improved cybersecurity requires digital threat intelligence - structured and semi-automated analysis and sharing of information. In order to make sense out of increasingly large and complex datasets related to cybersecurity we see the potential in developing models and tools for automated or semi-automated classification and discovery of cyberthreats based on ontologies.

Semantic technologies and ontologies are a relatively new logic-based landscape of technologies and tools aimed at giving better meaning to large and unstructured corpuses of data. Interesting research challenges are for example to investigate semantic representations of relevant concepts in the domain of cybersecurity big data, in order to facilitate advanced machine learning, search and discovery.

The potential benefit of this approach is that the developed tools and related technologies will provide a flexible framework for representing and structuring the large variety of data with which security analysts are confronted. The framework can further be used for the implementation of cybersecurity analytics tools.

II. Cybersecurity Threat and Risk Models

Cybersecurity is the body of technologies, processes and practices designed to protect networks, computers, programs and data from attack, damage or unauthorized access. Cybersecurity thus assumes that some actors, typically called *threat agents*, have the intent and capacity to produce attacks, gain unauthorized access and cause damage. The magnitude of the perceived potential damage caused by cyber attacks is typically interpreted as security risk.

A. Specific Security Risk Model

Cybersecurity risks are caused by threats. However, the concept of a threat can be ambiguous in the sense that it can mean the threat agent itself, or it can mean the thing that a threat agent (potentially) produces, typically called a *threat scenario*. Figure 1 illustrates a specific risk model which integrates the concepts of threat agent and threat scenario.

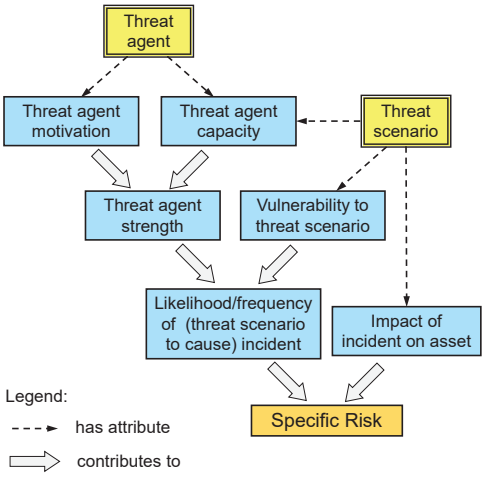

Figure 1. Specific risk model including threat agent and threat scenario

The specific risk model of Figure 1 emphasizes the risk dimension of threats, i.e. how threats lead to risk.

It can be seen that the threat agent and the threat scenario have very different attributes, but in combination they both contribute to risk. A threat agent can be modeled as a real agent with a motivation or goal as well as with a capacity to execute a specific threat scenario. Together, the motivation and capacity produce the strength of the threat agent. The threat agent strength can be modelled according to the weakest

This research was supported by the research projects TOCSA, ACT and Oslo Analytics funded by the Research Council of Norway.

link, i.e. the attacker is only as strong as the weakest of its motivation and capacity.

A threat scenario can be modelled as a sequence of attack steps which can be stopped by defence and security mechanisms. However, when the defence mechanisms fail to stop a specific threat scenario, we say that there are *vulnerabilities*.

The more severe the vulnerabilities and the greater the strength of the threat agent, the greater the likelihood that the threat scenario will cause a security incident and lead to damage, as illustrated in Figure 1. The actual risk of a specific threat scenario emerges by including the amplitude of the expected damage in case the security incident actually occurs. Risk assessment models such as in [1] are based on this interpretation of security risk.

There can of course be many different threat scenarios leading to the same goal when seen from the attacker's perspective. Each scenario represents the dynamic execution of a *tactic*. The attacker might consider multiple tactics, and then decide to use the one which is assumed to produce the greatest expected result with the least effort.

The threat scenario is an abstract set of steps executed in sequence, which from the victim/defender's perspective can cause damage to its assets. A threat scenario becomes a *cyber attack* when the scenario is actually executed. Behind every attack there is thus a specific threat scenario executed by an attacker or a group of attackers. However, a threat scenario by itself is abstract, and does not become an attack unless it is actually executed.

A threat scenario can therefore be interpreted as the blueprint for attacks. For cyber defenders there is thus a fundamental difference between detecting real attacks and identifying threat scenarios which only represent potential attacks.

B. Stillions' Detection Maturity Level Model

A model for the maturity of cyberthreat detection has been proposed by Ryan Stillions in several blogpostings [2]. A slightly extended version of Stillions' Detection Maturity Level (DML) model is illustrated in Figure 2. We have added the additional *DML-9 Attacker Identity* which can be important in certain contexts. We have also added *precision* and *robustness* to illustrate the qualitative aspects of features at each level. The DML model emphasizes the increasing level of abstraction in the detection of cyber attacks, where it is assumed that a security incident response team with low maturity and skills only will be able to detect attacks in terms of low level technical observations in a network, without necessarily understanding the significance of these observations. On the other hand, a security incident response team with high maturity and skills is assumed to be able to interpreted technical observations in networks in the sense that the type of attack, the attack methods used and possibly the identity of the attacker can be determined.

The levels of the DML model are briefly explained below. The focus is on what the IR team (incident response team) is

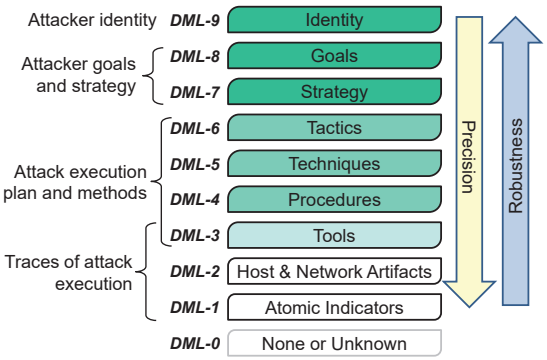

Figure 2. Detection Maturity Level Model [2]

capable of doing at each level. Our description is a summary and interpretation of Stillions' description [2].

- **DML-0 None or Unknown**. There is no IR team, or they are totally clueless.
- **DML-1 Atomic indicators of compromise (IOCs)**. These are elementary pieces of host & network artifacts, which might have been received from other parties. The value of atomic IOCs is limited due to the short 'shelf life' of this type of information.
- **DML-2 Host & Network Artifacts**. This is the type of information which can be collected by network and endpoint sensors. With high capacity links the amount of information collected can be overwhelming and requires good analytical tools to analyse and understand the attack at higher levels of abstraction.
- **DML-3 Tools**. Attackers install and use tools within the victim's network. The tools often change, so that a tool detected and analysed in a previous security incident might be similar but not exactly the same in new attacks. DML-3 means that the defender can reliably detect the attacker's tools, regardless of minor functionality changes to the tool, or differences in the artifacts and atomic indicators left behind by the tool.
- **DML-4 Procedures**. Detecting a procedure means detecting a sequence of two or more of the individual steps employed by the attacker. The goal here is to isolate activities that the attacker appears to perform methodically, two or more times during an incident. In the military jargon, procedures mean *"Standard, detailed steps that prescribe how to perform specific tasks"* [3].
- **DML-5 Techniques**. Techniques are specific ways of executing single steps of an attack. In the military jargon, techniques mean *"Non-prescriptive ways or methods used to perform missions, functions, or tasks"* [3].
- **DML-6 Tactics**. To detect a tactic means to understand how the attack has been designed and executed in terms the techniques, procedures and tools used. In the military jargon, tactics mean *"the employment and ordered*

arrangement of forces in relation to each other" [3].

- **DML-7 Strategy**. This is a non-technical high-level description of the planned attack. There are typically multiple different ways an attacker can achieve its goals, and the strategy defines which approach the threat agent should follow.
- **DML-8 Goals**. The motivation for the attack can be described as a goal. Depending on how the attacker is organised, the goal might not be known for the attack team executing the attack, the team might only receive a strategy to follow.
- **DML-9 Identity**. The identity of the attacker, or the threat agent, can be the name of a person, an organisation or a nation state. Sometimes, the identity can only be linked to other attacks without any other indication of who they are or from where they operate. The attacker identity might not be relevant to the defender if they only want to get the attacker out of the network. However, it is often important to be able to connect multiple attacks to the same actor in order to predict strategy, tactics, techniques and procedures expected to be used. This is an additional level defined by us, the original DML model [2] only consists of the levels 0–8.

The challenge is to leverage observed attack features detected at low levels to determine derivative causes at higher levels.

Assume that a given company *B* has as goal to beat company *A* in the open market. This goal might cause company *B* to use unethical means, with a strategy to steal secret information from company *A* in order to improve their own products and market position. Company *B*'s tactics may be to gain access to company *A*'s internal servers based on an attack plan with techniques, procedures and tools. Finally, the execution of the plan causes traces of the attack to be left in the network of victim *A*.

The cyber incident response team will first detect the traces, and from there must try to figure out what has happened and then decide the appropriate response. The traces are indicators, and the task of determining what really happened is a form of abductive reasoning which consists of using the indicators as classifiers to determine the nature and origin of the attack.

Most incident response teams of today are working on DML-1 and DML-2. Some are working on DML-3 and partly DML-6. However, the further up the stack you get the more seldom you find machine readable results from the analysis and work that is done. Defining semantic models for the type of information gathered in the higher levels of the DML model and the relations between them will enable more teams to increase their maturity level. Information sharing will also be facilitated by this development.

III. Elements of Semantic Threat Modelling

Discovering the real nature of a threat given a set of data or information requires a semantic model to represent all aspects of the threats with no room for ambiguous input. The further down the DML model you get, the more precise an identification can be done. The further up, the more costly a change is for the attacker and the more robust your conclusion of identity may become. Both aspects are useful for different roles and situations throughout a security incident. SIEM (Security Incident and Event Management) tools typically use semantic representation of host & network artifacts at the lower levels of the DML model, but rarely provide semantic representations of high level aspects. It is thus necessary to standardise the semantic representations of high level aspects in the DML model. This will allow automated reasoning to leverage the potential of machine learning and classifiers to do advanced cybersecurity analytical reasoning.

A. A Semantic Threat Classification Model

The primary focus of the DML model is to indicate levels of maturity in cyberthreat detection. However, the same model can be used as a basis for the design of cyberthreat classifiers, and we call this new model the *semantic threat classification model* (STCM).

Figure 3 shows the STCM which consists of a compact representation of the DML model combined with *classifiers* representing the analytical relationships from low level features to high level features.

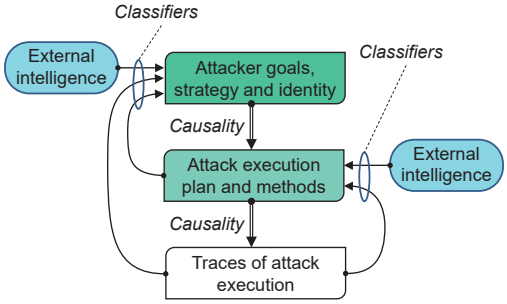

Figure 3. Semantic Threat Classification Model

Note that there are causal relationships from high level features to low level features. Hence, classifiers are used to reason in the opposite direction to that of causal relationships.

In machine learning and statistics, classifiers are used to determine categories to which some observation belongs, on the basis of a training set of data containing observations (or instances) whose category membership is known. For cybersecurity analytics, a classifier can e.g. be used to determine which type of attack a set of network artifacts belong to (i.e. are caused by), the goal of the attacker or even the identity of the attacker.

Note that contextual information can also be used as input indicators for classifiers. Contextual intelligence can e.g. be political events covered by the media. A political conflict between nation states can make it more likely that states launch specific types of cyberattacks against each other.

The challenge for developing reliable classifiers is to identify appropriate semantic features and their variables at each

level of abstraction, and to have available sufficient amount and type of data in order to give the classifiers sufficient training for reliable detection and classification.

The design of classifiers for machine learning is heavily dependent on statistical methods, and several authors have pointed out the importance of mathematics for cybersecurity [4].

B. Semantic Feature Extraction

Stillions' DML model [2] uses English prose to informally define each level of abstraction. The use of classifiers, however, requires formal definitions of the features at each level of abstraction. Our approach is to gather informal descriptions of goals, strategies, tactics, techniques and procedures from the literature. Through analysis of these informal descriptions, we derive tuples that describe each level of abstraction. In the following, we illustrate this process for the abstraction level "Goals".

Stillions mentions the following goal as an example:

> Replicate Acme Company's Super Awesome Product Foo in 2 years or less [2]

If we ignore the time dimension of this goal, then we can derive the 2-tuple ("Replicate", "Product") from the informal description.

From Mandiant's APT1 report [5], we can derive the following goals: ("Replicate", "Product"), ("Replicate", "Manufacturing process"), ("Obtain", "Business plan"), ("Obtain", "Policy position").

Another goal can be derived from Symantec's blog post on the "Cadelle" and "Chafer" APT groups [6]: ("Monitor", "Individuals").

By generalising the examples above, we get the following definition of a goal: (Action, Object). When we observe the 2-tuples from the examples, we identify two challenges. The first challenge is that we use strings to describe each element of the tuple. If we use 2-tuples of strings in a system where a multitude of analysts and classifiers identify and record new goals, then the result will be duplicated by synonyms resulting in an explosion of features. In order to avoid this, our goal is to define a formal taxonomy of goals, where each tuple contains references to the taxonomy.

The second challenge is that the second element of the 2-tuple is too general. To alleviate this, we must define sub-elements that are more specific, e.g. that the "Product" in the first example is manufactured by "Acme company", and that the specific product is "Super Awesome Product Foo". In the last example, "Individuals" could have a sub-element "Iranian Citizens". Note that in some cases we will not be able to determine these sub-elements due to insufficient data.

Applying this approach to all the layers of abstraction in the extended DML model requires a monumental amount of effort. We believe that in order to achieve this, a community effort is needed. Thus, one of our primary goals is to lay the foundations for such an effort. Furthermore, re-using existing standards and taxonomies where applicable can significantly reduce the amount of work needed. A good example of such re-use can be observed for the abstraction level "Techniques". The MITRE ATT&CK taxonomy [7] has already defined more than 100 techniques used by adversaries in the post-compromise phases of an attack.

C. Current Initiatives for Cyberthreat Representation

There are several initiatives currently being used for representation and sharing of data on the different levels of the DML model. The following initiatives are seen as useful and may be used when selecting features for representation on the different levels:

- **INTEL Threat Agent Library** (TAL) [8] was suggested in 2007 and provides a consistent reference describing the human agents that pose threats to IT systems and other information assets. This library may serve as a feature of "Identity" in our semantic threat modelling.
- **STIX** [9] is a language for having a standardized communication for the representation of cyberthreat information. It is well known in the incident response community, but not serving the purpose of describing all aspects of cyber threats. The main shortcoming in the current version is the lack of separation between tactics, techniques and procedures.
- **CAPEC** The objective of the Common Attack Pattern Enumeration and Classification (CAPEC) [10] effort is to provide a publicly available catalog of common attack patterns classified in an intuitive manner, along with a comprehensive schema for describing related attacks and sharing information about them. CAPEC is run by MITRE and is openly available for use and development for the public. For our semantic threat modelling it may be used when describing 'Tactics' and 'Techniques'.
- **ATT&CK** is a common reference for post-compromise tactics, techniques and tools [7] run by MITRE. ATT&CK and CAPEC are related and do not exclude use of each other.

IV. Example applications of semantic cyberthreat models

In this paper, we argue that semantic cyberthreat models can help cybersecurity professionals to be more effective and efficient. This section presents some concrete examples from our own experience that support this hypothesis.

A. Incident response

Breaches due to attacks from advanced persistent threats (APTs) are often detected post-compromise. APTs quickly initiate lateral movement after the initial compromise, so assessing the scope of the breach can be challenging. In order to assess the scope of the breach, we need to know how the threat agent operates and what kind of indicators, artifacts, tools, tactics, techniques and procedures (TTPs) we should search for. The incident response analysis process typically consists of the following steps:

1) Evidence collection
2) Analysis of evidence

3) Identification of new indicators, artifacts, tools and TTPs
4) Threat agent attribution

Steps 1-3 are performed in an iterative fashion. The analysis results may indicate that we need to collect more evidence, or that we should search the existing evidence for new indicators. If we are able to perform step 4 and attribute the breach to a known threat agent, then we can leverage our historical knowledge of this threat agent. We can use this knowledge to guide our evidence collection and analysis. We have used the MITRE ATT&CK taxonomy [7] to be able to quickly compare our evidence to known threat agents during incident response. By manual analysis, we found threat agents that used tools and techniques very similar to what we observed in our evidence. The ATT&CK taxonomy [7] has a loose semantic model connecting threat agents, tactics, techniques and tools. It does not model procedures, artifacts or indicators. In order to automate the analysis of threat agent similarities, we implemented a simple semantic model using a graph database. The model linked threat agents to observed indicators, artifacts, tools and TTPs. We then used the graph database to find all subgraphs that connected the findings from our incident to known threat agents. The result enabled us to attribute the evidence from our incident to a known threat agent, and the results helped guide our evidence collection and analysis. Another great advantage of using such a model is that the attribution hypothesis can be re-tested as more knowledge is added to the graph, in order to avoid confirmation bias. Our experience from this incident was that we were able to attribute the evidence to a known threat agent much more rapidly than by using manual analysis. We were also able to fully document all relations between our evidence and the threat agent by issuing a simple graph query.

B. Requests for information

A common task for threat intelligence analysts is to find all information related to a single data point, e.g. an IP address, a malware sample or a threat agent. Having a semantic model implemented as a graph makes it possible to complete such a task quickly and reliably by issuing a single graph query.

C. Intrusion detection

Current intrusion detection systems operate at DML-1, DML-2 and/or DML-3. One of the challenges with operating at DML-4 and above is that TTPs are commonly described using English prose, i.e. as unstructured data. This makes it challenging to translate the description to intrusion detection signatures, and signature development must be performed manually. Defining formal models for TTPs makes it possible to automatically generate signatures from structured data when a new TTP is defined. One concrete example is the procedure described in [11]:

> An example would be an adversary running **net time**, followed by the **AT.exe** command to schedule a job to kick off just one minute after the current local time of the victim system. [11]

Given an endpoint security solution that logs process execution with arguments and command inputs/outputs, a human analyst could write a signature to detect this procedure. The signature would have to detect the following:

1) Execution of **net.exe** with **time** as the first argument and **victim system** as the second argument
2) Timestamp returned by the command in step 1
3) Execution of **at.exe** with **victim system** as the first argument and ((timestamp from step 2) + 1 minute) as the second argument

Interpreting the description "to schedule a job to kick off just one minute after the current local time of the victim system" is easy for a human, but very difficult for a computer. A formal definition of this procedure would make it possible for a computer to automatically generate signatures for the procedure by applying transformation rules.

V. Conclusion

Semantic modelling of threats is a promising approach for automated threat and attack detection at multiple levels of abstraction. A semantic model of threats will enable security analysts to work faster and more efficiently in terms of identifying threat agents and take advantage of previous experience and gathered intelligence when handling incidents caused by known or unknown threat agents. The task of extracting semantic features for all levels of abstraction in our suggested extended DML model is an undertaking of daunting proportions. In order to make this task manageable the reuse of related standards and taxonomies is required.

References

[1] ISO, *ISO/IEC 27005:2011 - Information technology – Security Techniques – Information security risk management (second edition)*. ISO/IEC, 2011.
[2] R. Stillions, "The DML Model," http://ryanstillions.blogspot.com/2014/04/the-dml-model_21.html, 22 April 2014.
[3] U. DoD, *Department of Defense Dictionary of Military and Associated Terms*. Joint Chiefs of Staff, 2010.
[4] A. Pinto, "Secure because of Math: A deep-dive on Machine Learning-Based Monitoring," Black Hat Briefing. BlackHat Conference, 2014.
[5] Mandiant, "Mandiant APT1," https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf, 18 February 2013.
[6] S. S. Response, "Iran-based attackers use back door threats to spy on Middle Eastern targets," http://www.symantec.com/connect/blogs/iran-based-attackers-use-back-door-threats-spy-middle-eastern-targets, 7 December 2015.
[7] MITRE, "Adversarial Tactics, Techniques and Common Knowledge (ATT&CK)," https://attack.mitre.org/.
[8] T. Casey, "Threat agent library helps identify information security risks," *Intel White Paper, September*, 2007.
[9] S. Barnum, "Standardizing cyber threat intelligence information with the Structured Threat Information eXpression (STIX)," *MITRE Corporation*, vol. 11, 2012.
[10] MITRE, "Common Attack Pattern Enumeration and Classification (CAPEC)," https://capec.mitre.org/.
[11] R. Stillions, "On TTPs," http://ryanstillions.blogspot.com/2014/04/on-ttps.html, 22 April 2014.

Paper II

Ethical considerations in sharing cyber threat intelligence

II

Siri Bromander

Proceedings of the Tenth Norwegian Information Security Conference (NISK 2017).
Oslo, Norway.
ISSN: 1894-7735.

Ethical considerations in sharing cyber threat intelligence

Siri Bromander
mnemonic as
University of Oslo (UiO), Norway
siri@mnemonic.no

Abstract

Sharing information with others is always a choice. In the world of cyber defense, sharing information with others can help others defend themselves, and with this increase the joint defense our society needs to have in order to stay safe. Several factors influences the choice of sharing valuable cyber threat intelligence, and the ethical considerations are argued to be a prominent part of this.

When encountering a situation where a choice of sharing information is emerging, the choice will be twofold: 1. what information should be shared?, and 2. with whom should the information be shared? The ethical challenges of the choices is primarily tied to who you have obligations to. The consequences of the choices will potentially affect the society in variable degrees, your employer, your colleagues, your friends and obviously yourself.

This article discusses the ethical considerations cyber security personnel is facing making these types of decisions.

The first part of the article explains details of cyber threat intelligence and its community architecture. Following this, the article describes what influences the choice personnel is facing when having the possibility to share valuable information with others, tying the considerations to known research within knowledge management and ethics of knowledge sharing. An example is given to discuss the possible choices and the ethical considerations within all possible choices. Towards the end a short note is done on sharing information in the aftermath of incidents instead of during an incident. The articles concludes that not sharing valuable information at all is immoral, but how much and with whom needs to be a consideration made special in each case, leaving a deontological approach unsuitable.

1 Cyber Threat Intelligence

Cyber infrastructure encompasses many aspects of our daily lives. Our homes are an increasing part of the 'Internet of Things', our society is increasingly digitalized and

This paper was presented at the NIK-2017 conference; see http://www.nik.no/.

our workplaces are all, to some degree, using available cloud services as convenient and efficient solutions for us to perform at our best. Everything connected to the internet is made available to the rest of the world. The rest of the world are not always having only good intentions. Everything available on the internet makes possible targets for cyber threats[1], and the consequences are possibly lethal; physical damage to for example a dam or a nuclear power plant could kill numerous people and provide environmental changes beyond repair within our lifetime. Defending ourselves has never been more important, and will be increasingly important in the years to come.

Cyber-attacks are becoming more common, sophisticated and damaging. The stories of Stuxnet, the 2016 US Election and the more recent ransomware WannaCry reveals the concerning fact that highly skilled threat agents are capable of sabotage, espionage and subversion to the degree of nation state concern. Some argue there is no such thing as cyber warfare[1], but in July 2016 NATO recognized cyberspace as a domain of operations in which NATO must defend itself as effectively as it does in the air, on land and at sea[2]. The terminology used may not be as interesting as the discussion revealing capabilities and consequences of conflict in the field of cyber. Recent history shows us that critical infrastructure can be taken down, elections may be influenced and critical parts of society can be disrupted for days caused by attacks happening in cyber space alone.

The need for advanced and rapid response is increasing. Seeing the battlefield is far less visual than that of physical war, the need for sharing and communicating known intelligence between defending partners increase. Sharing information and knowledge is a field of its own. A field where technical and strategic obstacles are discussed and debated, but where I would argue that ethical considerations are just as important to address.

Cyber threat intelligence started out as something the larger and best computer incident response environments did to succeed in their day to day job. Good work made good results which could be useful to others trying to defend against the same adversary. To receive good information one needed to share good information, and so started the large communities of cyber threat intelligence. The commercial value of this type of work was quite fast seen by other environments, and the market for threat intelligence grew very fast[3]. Today we are facing a severe amount of businesses offering threat intelligence products.

Threat intelligence is evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice, about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject's response to that menace or hazard[4]. In the world of cyber security, this means sharing everything from smaller network artifacts, indicators of compromise and samples of malware or infection vectors, to descriptions of how attackers are operating, their capabilities and intents. In some cases also the identity of attackers are shared. The format of sharing this type of information is going from emails, files and chat-channels, to automatic feeds. The timing of such sharing is often very important as much of the information shared is only valid for a limited amount of time and can be crucial to receive in order to be able to handle whatever threat is targeting you. More often than not is this type of sharing based on personal

[1] Oxford dictionary: 'Cyber threat: The possibility of a malicious attempt to damage or disrupt a computer network or system.'

relationships. The choice of what to share and who to share it with is often a decision made or at least initiated in the moment by the same personnel doing incident response and threat hunting activities.

The information we gain when investigating an incident in a thorough manner is knowledge, and knowledge management is a relatively young but evolved field of research. In relation to knowledge and business ethics it is seen that an unwillingness to share knowledge that may hurt an organization's survival is seen as being seriously unethical[5]. In general, we could argue that the knowledge you have about an adversary that may seriously injure another organization is something we are morally obliged to share, but in our circumstances the decision to share may also give the consequences of severely injuring your own ability to defend yourself against serious consequences. The debate is therefore somewhat more complex in our circumstances.

We need to look at the general facts and circumstances that will influence the choice of action when facing situations where sharing is an option.

The first aspects that will influence the decision you make is how knowledgeable the receiver of your information or knowledge is. If organizational members believe in other members' expertise and skills, the intention to share individual knowledge increases[6]. Often the use of shared threat intelligence can ruin the information itself, if the information you provide them is used in a manner the adversary is capable of detecting. 'Blowing' the intelligence received is seen as likely if the recipient has little knowledge or experience in both handling the technical details as well as the stress related to a serious incident. The intent is seldom to ruin the information when using it to defend yourself, but the possible consequences are nevertheless there.

Secondly, also related to knowledge, is the ability for managers to understand the consequences of sharing. Sharing is ultimately seen as a good deed, but sometimes information shared is carrying metadata which can reveal more than initially thought. This is something the technical personnel may try to explain to their managers, but not knowing the business sides missing to understand the consequences of and therefore not being able to explain in a sufficient way. The uncertainty of not knowing how much is actually revealed with sharing a given set of information, often influences the choice of sharing towards not sharing.

Thirdly, who can be affected by your decision is relevant, and how you are obliged to them. Working for a company most always encounter a contract of work, making you obliged to follow company policy and the instructions of your managers. Being a citizen of a given country you are obliged by law to follow the rules set in that country. As part of a volunteer community you are expected by your peers to contribute, and you may even owe someone in the community or elsewhere a favor after significant help in the past. In some cases you may be in a situation where certain others seems to deserve help as they are either a special type of organization or critical of nature, and the overall obligation to prevent bad behavior and criminal acts as part of the society is always present.

From knowledge management we know that knowledge alliances motivate managers to enter into strategic alliances with other firms in order to balance knowledge deficiencies, obtain necessary competencies and create new knowledge[7]. This is exemplified by cyber threat intelligence and the vast amounts of sharing and collaboration networks that exists. Some of these are based upon contracts and legal obligations, like reporting to governmental parties when handling critical

infrastructure, and some are just based on spoken agreements and the desire to share and collaborate like companies working in the same industry collaborating when it is found useful. Either way the collaboration agreements you are faced with will influence your choice of sharing information.

Fourthly, the culture of both the country, organization and community in which you reside affects the willingness to share knowledge. Ethical decision-making is affected by culture through an individual's deontological and teleological evaluations. Although individuals may regard a particular activity as ethical, they may follow a different course of action because of the desirable outcome. Because people make different assumptions about personal knowledge, it can therefore not be assumed that workers in all cultural value systems will view their own decision not to share their personal knowledge, or a decision to act out of self-interest in the face of internal competition, as unethical or immoral[6]. Within cyber threat intelligence, these challenges can be exemplified with the differences between corporate organizations and military organizations. Even though personnel from both have obligations to their residing country, military personnel would arguably act from a stronger obligation to national interests, simply because of their training, experience and choice of work place. Consequently, their evaluation of consequences on national security plays stronger than those of for example financial loss or personal gain.

Finally, timing is of importance. The nature of threat intelligence is that is most often is only valid for a limited amount of time, and that the receiver needs it as soon as possible to increase their ability to defend. An example is information on the infrastructure used for attacking a given organization. An advanced attacker would change the used infrastructure on a regular basis, leaving information on IP addresses and domains useless as soon as they swap. In many cases this is a matter of hours. Sometimes this means that there is not enough time to go all the rounds internally to get approval before you share, and also that the time spent on considering all consequences of the action could be a waste of time you do not have. Whether you share classified information when sharing, or whether the attacker is pushed to change its infrastructure sooner and leaves yourself unable to know where the attack is coming from next, are considerations that requires time consuming analysis. If conducting all analysis before sharing, the information may no longer be worth sharing. In these terms it could be argued that following rules, adhering to duties (deontological approach) is far better for the time sensitive matter of sharing threat intelligence than that of considering consequences.

2 An example: sharing while enforcing your own defense

A normal situation for a security analyst to be in is given as an example to illustrate the challenges related to deciding who to share with. The described situation illustrates the influencing factors relevant to the ethical consideration the analyst must make.

Imagine the scenario: you are a young security analyst, skilled and with experience from several organizations, both work related and as part of the volunteer security community. You are popular both because you are knowledgeable, but also because you on several occasions have helped others in succeeding with handling difficult incidents in the past. You are active within several cyber security communities on your spare time.

At work, you take place in the handling of a severe security incident. An adversary has successfully compromised your computer network, but you have detected the attacker and are monitoring their every move together with your team. You do not know for sure what the adversary is after, but based on the business of which your company is working, you have a fair idea. If your suspicion is right, the adversary in question would be able to do severe physical damage through your computer systems if not stopped. You have several friends in organizations likely to be targeted by the same adversary, both private and public sector, and the adversary is likely to be after valuables of national interests and capable of sabotaging critical infrastructure. You also know enough about the adversary to conclude it is an advanced attacker with the ability to change its behavior to the extent that you are no longer able to either detect or monitor them anymore. Hence you are depending on your knowledge not to be leaked to the adversary in any way to be sure you can defend against them yourself. You prepare information on how to detect and monitor the adversary for sharing with others and approach your manager. The discussion that follows is difficult: Should we share this information? And with whom? Who are we obliged to help and to what extent can we morally defend putting our own defense before others? Is personal obligations something you can set aside or is that relevant as well?

There are technical and practical aspects that we set aside for this discussion, like the ability to share the information in a relevant manner. For our purposed we are looking into the ethical aspects of the decision of sharing/not sharing with different parties.

To debate what is morally right and wrong in our example we need to examine the possible actions and related rules (deontological approach), and the consequences of our possible actions, both the direct consequences and the long term consequences (teleological approach, in this case consequentialism).

So the possible choices to make regarding our piece of information in this situation are the following:

- **Doing nothing.** A general, positive rule is that 'we share information that can help others'. In these terms the act of not sharing information is unethical. However, if sharing that information encounters possibly sharing metadata covered by laws and regulations in the county in question, you are breaking a more prominent rule of 'do not break national laws'. However, if skilled at incident response you know what information that is ok to share, and the deontological approach would tell you that the act of not sharing is unethical.

  The consequences of not sharing information is directly that you do not spend time on it, which may help you do better at actual defense. In addition, you are certain that you keep all company information safe. On the negative side you find several consequences, but the worst would be that several others are not able to defend themselves and that it could lead to severe physical damage. With this as considerations it is not ethically possible to defend not sharing information with anyone. I consider the action of 'waiting and sharing later' to give the same discussion as above.

- **Sharing information with national capabilities only.** In Norway (and most countries with defined national cyber capabilities) there are laws and regulations stating that incidents which can affect national security shall be

reported to the authorities. This means that handling an incident in your environment if your environment is part of for example critical infrastructure, is something that should be reported as soon as the incident has been detected. The decision of not helping anyone else does however mean you do break the rule of 'we share information that can help others' as stated above.

Company policies are usually having statements in lines of 'we shall always adhere to laws and regulations, but any circumstances where we suspect possible prosecution as a consequence shall be run by legal'. In the time sensitive circumstances of incident response, this often means you need to break either company or national laws when deciding. Most companies will officially state that national laws are first in line, but in real world scenarios we see that this is not always as straight forward. If considering our society it is hard to argue that not sharing with national capabilities is morally right.

If evaluating consequences the most prominent positive are that they can protect our national interest. They can decide on further sharing, which for many means you have done your duty. But knowing the authorities does not have the same network as yourself, you know that not everyone you could have helped is being helped. This is still breaking your obligations to the society and the security community and still many organizations being defenseless must be seen as a negative consequence still. Further, on negative side of actually sharing with the authorities, is the uncertainty of how well they will treat the information you provide them with. Their interests and their skill level is probably unknown to you, and you risk them ruining your own defense. The information you provide can put you in a position where they will investigate you further, and also, you risk that they will classify your information which makes it harder for you to use in the further.

Sharing with the authorities is seldom argued immoral due to the laws and their ability to help national interests, but not sharing with others may be defended as unethical.

- **Sharing information with those you know.** The situation where people that are close to you get hurt is harder to accept. The rule of helping others is strengthened, but will only survive as long as laws or regulations are not forbidding you to share information. Seeing you know those close to you, your evaluation of skill level is related to less uncertainty, which decreases the negative consequences of sharing with them. Breaking laws to share would encounter prosecution and would not be considered ethical if lives are not at stake. Following both deontological and teleological approaches will therefore likely give the same conclusion: you should share as long as the recipients are allowed to by law.

- **Sharing information with the relevant sectors.** Within the security industry several sharing collaborations has been created in order to share with relevant partners in different incidents. The groups are often created in different industry sectors and based on voluntary participation. It can be seen as closed sharing with participation restrictions, but without personal knowledge of the group of recipients, their skill level and where their loyalty lies.

In our example we have a situation where the analyst is not obliged by rule or formal contract to share. The rule of sharing information to help others is still present and I can find no other rule that is strong enough to contradict this. However, looking at the possible consequences of sharing you have larger degree of uncertainty related to how the information is being treated and therefore you have potential consequences ruining the information not only for you, but also for national capabilities and others you have shared with. Given this evaluation I find it to be defendable to limit the amount of sharing done to those you know can help the most and maybe can assume can handle the information best.

- **Sharing information with everyone.** The act of sharing information with everyone is good alone. It follows the rule of 'you are not keeping to yourself information that can be valuable to someone else'. Seeing it is impossible for you to know everyone who can gain value out of the information you share, broadcasting the information in ways that makes everyone interested capable of finding it is therefore the right way to go. Technically, this means creating a public report or similar and publishing it somewhere online. However, sharing with everyone also means sharing with your adversaries, and knowing this breaks the unwritten rule of 'not telling your enemies how you work or what you know'. Which rule is the most prominent of these? The uncertainty of the latter and the size of the benefit it serves other victims will judge this.

  Following the consequences of these actions, one can argue that the good of sharing with everyone is both that more people may be able to protect themselves, but also that by sharing intelligence more people can learn and the general skill level is increased. The flip side is as indicated earlier, the adversaries may change their patterns, improve and be even harder to protect against in the future[8].

  In these terms the 'doctrine of double effect' comes into relevance. If sharing the information with everybody, then it is likely that everyone will benefit in the short run, but the advanced communities will lose eyes on the adversary as soon at the adversary knows their details are known. This is known, but the good of more organizations being able to defend themselves in the short run, outweighs the fact that the adversary is able to escape detection by changing the details now known in the broader communities. This is seen as ethically defendable as long as the intention is that of helping more organization defending themselves, and not to help the adversaries in improving their methods.

3 Sharing information about past cyber incidents

Another relevant question concerns sharing information about a cyber-threat in the aftermath of an incident. This can still contain valuable information about an adversary, but often not technical information that can be of direct help to a certain incident. IP addresses are no longer in use by the adversary, but they still use the same procedures when attacking a new victim. Even though the technical details may be ruined and useless, information about cyber threat methodology stays robust over time and may be shared in the aftermaths of an incident to contribute to the base of experience the rest of the security communities can benefit from. The

information can therefore not be ruined by recipients that lacks sufficient knowledge. Also, to share information is often a decision made by others than the personnel themselves, for example marketing or legal, but is still an interesting and related debate to address. The consequences of such an action is less influenced by that of 'ruining information' and hence not as related to the knowledge of the recipients, but more so of the long term consequences of the community we live in. Sharing knowledge between defending parties makes our community better prepared and more likely able to protect its citizens. The conclusions made within knowledge management theory[5] is therefore more valid here, and it is possible to state that not sharing such knowledge is considered unethical.

4 Conclusions

The position you are in/the context will always influence the decision you make and also the difficulties of acting morally right. The most prominent aspects influencing the decision of sharing threat intelligence is who you are obliged to, who can be affected by your decision and how much damage the information can have among the wrong recipients. In our cyber security world we see that both employer and friends in the security community have high influence. Trust in and skillset of the recipients is of high importance when evaluating the possible consequences. When sharing sensitive information one needs to trust the recipient to protect it from the adversaries and not to ruin it. Sharing information itself is a good deed, but if the negative consequences are easily understood, then the decision not to share is the easiest. The ethical challenges of acting against laws and regulations seems to be the strongest positive influence on the choice of sharing information, which may come as a result of the little analysis needed to understand the consequence. The extent of both negative and positive consequences of sharing cyber threat intelligence is otherwise requiring more extensive analysis and may not be possible to even estimate due to time constraints and lack of available knowledge. Your assumed adversary may be able to severely damage your organization and your peers may be able to both use the received information and treat it with care. When uncertain people often has a tendency not to act.

Creating a given rule to follow in any such case is impossible due to the large degree of uncertainty in the above stated aspects influencing the choice of action. A solely deontological approach is therefore not a suitable ethical framework to deal with such cases.

As part of a society the long term consequences of not sharing information at all is making the act immoral. With the knowledge of severe negative consequences of sharing a piece of information, the sharing can defendable be done within closed communities where the recipients are known to treat the information right, like sector specific sharing groups or sharing with groups consisting of members based on 'invite only'. However, in the aftermath of an incident, when the incident has been handled, I can see no good argumentation to defend the act of not sharing valuable threat intelligence. The consequences of not sharing information or revealing real incidents will in the long run mean that less people understand the severity of the cyber-attacks in our region, and consequently do not spend resources protecting against them. For us as a society that is a major security issue.

References

[1] Thomas Rid. Cyber war will not take place. *Journal of strategic studies*, 35(1):5–32, 2012.

[2] NATO. Warsaw summit communiquÃ©. `https://ccdcoe.org/sites/default/files/documents/NATO-160709-WarsawSummitCommunique.pdf`, 2016.

[3] Brian Bartholomew and Juan Andres Guerrero-Saade. Wave your False Flags! Deception Tactics Muddying Attribution in Targeted Attacks. In *Virus Bulletin Conference*, 2016.

[4] Gartner. Definition: Threat intelligence. `https://www.gartner.com/doc/2487216/definition-threat-intelligence`, 2013.

[5] Chieh-Peng Lin. To share or not to share: Modeling tacit knowledge sharing, its mediators and antecedents. *Journal of business ethics*, 70(4):411–428, 2007.

[6] Nina Evans and Mary McKinley. Ethical paradoxes in knowledge management. *Vie & sciences de l'entreprise*, (2):57–71, 2011.

[7] Richard Baskerville and Alina Dulipovici. The theoretical foundations of knowledge management. *Knowledge Management Research & Practice*, 4(2):83–105, 2006.

[8] Gadi Evron and Inbar Raz. Apt reports and opsec evolution, or: these are not the apt reports you are looking for. `https://www.virusbulletin.com/uploads/pdf/magazine/2017/VB2016-EvronRaz.pdf`, 2016.

Paper III

Cyber threat intelligence model: an evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence

Vasileios Mavroeidis, Siri Bromander

Proceedings of the 2017 European Intelligence and Security Informatics Conference (EISIC 2017).
Attica, Greece.
ISBN: 978-1-5386-2385-5.

III

Cyber Threat Intelligence Model: An Evaluation of Taxonomies, Sharing Standards, and Ontologies within Cyber Threat Intelligence

Vasileios Mavroeidis
University of Oslo
Norway
vasileim@ifi.uio.no

Siri Bromander
mnemonic
University of Oslo
Norway
siri@mnemonic.no

***Abstract*—Threat intelligence is the provision of evidence-based knowledge about existing or potential threats. Benefits of threat intelligence include improved efficiency and effectiveness in security operations in terms of detective and preventive capabilities. Successful threat intelligence within the cyber domain demands a knowledge base of threat information and an expressive way to represent this knowledge. This purpose is served by the use of taxonomies, sharing standards, and ontologies.**

This paper introduces the Cyber Threat Intelligence (CTI) model, which enables cyber defenders to explore their threat intelligence capabilities and understand their position against the ever-changing cyber threat landscape. In addition, we use our model to analyze and evaluate several existing taxonomies, sharing standards, and ontologies relevant to cyber threat intelligence. Our results show that the cyber security community lacks an ontology covering the complete spectrum of threat intelligence. To conclude, we argue the importance of developing a multi-layered cyber threat intelligence ontology based on the CTI model and the steps should be taken under consideration, which are the foundation of our future work.

***Index Terms*—cyber threat intelligence, threat information sharing, cyber security, threat intelligence ontologies, cyber attack attribution, cyber threat detection, cyber threat prevention, knowledge representation**

I. Introduction

The capabilities, persistence, and complexity of adversarial attacks in the present threat landscape result in a speed race between security analysts, incident responders, and threat actors. Coordinated cyber crime is at each peak. PwC's global economic crime survey of 2016 [1] reports that there are organizations that had suffered cybercrime losses over $5 million, and of these nearly a third reported losses in excess of $100 million. In addition, Juniper Research [2] reports that cybercrime will increase the cost of data breaches to $2.1 trillion globally by 2019; four times the estimated cost of breaches in 2015.

In the Proceedings of the European Intelligence and Security Informatics Conference (EISIC 2017), Attica, Greece, September 11-13, 2017. DOI: 10.1109/EISIC.2017.20

This research was supported by the research projects Oslo Analytics, TOCSA, and ACT funded by the Research Council of Norway.

Security analysts and incident responders need the right skills to recognize attacks before performing defense efforts. The development of adequate controls require a thorough threat analysis, but small and medium sized businesses most of the times have inadequate capabilities due to lack of skilled personnel and budget constraints.

Threat intelligence is referred to as the task of gathering evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice, about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject's response to that menace or hazard [1]. Threat information reported and shared between security teams is overwhelming making difficult its absorption and correlation to existed stored knowledge; as a result, threat intelligence vendors are increasingly shifting to ways of automating this process making threat analysis a viable task.

Analyzing and sharing threat data and threat information in an effective way requires common representation, standard formats and protocols for sharing, and a common understanding of the relevant concepts and terminology. A solution approach to this need is the use of artificial intelligence (AI) and particularly the use of ontologies. An ontology is a form of knowledge representation that can integrate information coming from different sources.

Working towards an ontology for cyber threat intelligence is not an easy task. Our research reports the following as the largest difficulties:

- Vaguely defined terminology leads to confusion among experts and additional work to extend or unify ontologies.
- Lack of formal standardized representation of relevant information results in strings of English prose, with no standard pattern. Standardizing well defined taxonomies can eliminate this barrier.
- Lack of coherent relationships between the different layers of abstraction in ontologies. Modular ontologies containing several sub-ontologies need sound relationships

[1] https://www.gartner.com/doc/2487216/definition-threat-intelligence

between the different data points to leverage the power of semantics and reasoning. For example, to understand the behavior and the capabilities of a threat actor the connections and relationships between pieces of information must be sound.

This article evaluates taxonomies, sharing standards, and ontologies relevant to the task of creating an ontology for use within cyber threat intelligence. Some of the ontologies potentially can aid threat intelligence but initially have been introduced to address a specific domain within cyber security. Additionally, we pinpoint the relationship between our own Cyber Threat Intelligence model (CTI), the taxonomies, the sharing standards, and the ontologies discussed, aiming to classify them in terms of expressivity. Finally, we critically discuss the shortcomings of the present cyber threat intelligence ontology approaches and we address the directions that should be followed for their advancement.

II. Methodology

This section introduces two models related to threat detection maturity and cyber threat intelligence, respectively. The two models overlap and both can meet different needs that are explained in the next two subsequent subsections. The Cyber Threat Intelligence model is the basis of the evaluation process conducted in this paper.

A. The Detection Maturity Level Model - DML

Ryan Stillions proposed the DML model in several blog postings in 2014 [3]. The model was originally used to describe the maturity of an organization in terms of their ability to consume and act upon given threat information. Threat information can include indicators of compromise, tactics techniques and procedures of an actor (TTPs), threat intelligence reports and many more. In 2016, we extended this model by adding an additional level (9) "Identity" and presented it for use in semantic representation of cyber threats [4].

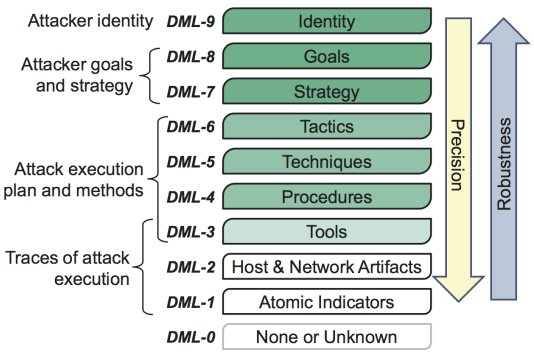

Fig. 1. Modified Detection Maturity Level Model [4] [3]

The DML model emphasizes the increasing level of abstraction in the detection of cyber attacks, where it is assumed that a security incident response team of low maturity and low skills would be able to detect attacks in terms of low level technical observations in a network, without necessarily understanding the significance of these observations. On the other hand, a security incident response team of high maturity and high skills is assumed to be able to interpret technical observations in networks in the sense that the type of attack, the attack methods used, and possibly the identity of the attacker can be determined.

Detection maturity, threat information, and threat intelligence overlap in a way that high or low detection maturity consequently can produce rich or poor threat information that can result in rich or poor threat intelligence. However, rich threat intelligence can aid the detecting and preventing capabilities of teams of low maturity by absorbing advanced threat intelligence shared from teams with higher detection capabilities.

B. The Cyber Threat Intelligence Model - CTI

For the purpose of evaluating and classifying taxonomies, sharing standards, and ontologies relevant to threat intelligence we identified the need to develop a new model that can suitably characterize threat intelligence. The Cyber Threat Intelligence model is not hierarchical like the DML model, but mainly a way to represent what types of information are needed for advanced threat intelligence and potential attack attribution. Acquisition of the Cyber Threat Intelligence model in the security operations of an organization strengthens the security posture of the organization itself by enabling advanced detective and preventive capabilities.

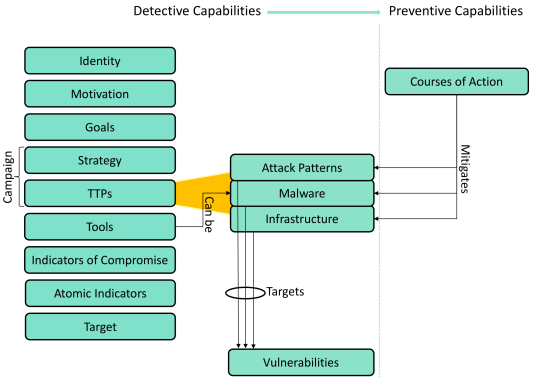

Fig. 2. Cyber Threat Intelligence Model

The remaining of the section is devoted in specifying the definitions of the elements comprising the CTI model.

Identity: The identity of a threat actor can be the name of a person, an organization, or a nation state. Sometimes, the identity can only be linked to other attacks without actual attribution or even location of their operations. However, it is important to be able to connect multiple attacks to the same

actor in order to determine any strategy, tactics, techniques, and procedures expected to be used.

Motivation: Motivation can be described as the driving force that enables actions in the pursuit of specific goals. Motivation may be derived from the benefits achieving a goal. The goals of an attacker may change, but the motivation most of the times stays the same. Knowing a threat agent's motivation narrows down which targets that agent may focus, helps defenders focus their limited defense resources on the most likely attack scenarios, as well as shapes the intensity and the persistence of an attack [5]. Examples of motivation can be ideological (human rights, ethnic etc.), military, financial and many more.

Goals: According to Fishbach and Ferguson [6] "a goal is a cognitive representation of a desired endpoint that impacts evaluations, emotions, and behaviors". A goal consists of an overall end state and the behavior objects and plans needed for attaining it. The activation of a goal guides behaviors (strategy). Depending on how the attack is organized the goal might not be known for the attack team executing the attack. The team might only receive a strategy to follow. In present cyber threat intelligence most of the times goals are described in prose. A goal can be defined as a tuple of two: (Action, Object), but work needs to be done to create a consistent taxonomy at an adequate level of detail [4]. Typical examples of goals are to "steal intellectual property", "damage infrastructure", and "embarrass competitor".

Strategy: This is a non-technical high-level description of the planned attack. There are typically multiple different ways an attacker can achieve its goals, and the strategy defines which approach the threat agent should follow. In present cyber threat intelligence, strategies are most of the times described in prose. It is our belief that the introduction of a formal taxonomy describing relationships between motives, goals, and strategies would be advantageous for the advancement of cyber threat intelligence, as well as the risk assessment processes. Part of our future research is the development of such a taxonomy.

TTPs: Tactics, techniques, and procedures characterize adversary behavior in terms of what they are doing and how they are doing it.

1) Attack Patterns: Attack Patterns are a type of TTP that describe ways the adversaries utilize to compromise targets.

2) Malware: Malware is a type of TTP and refers to a software that is inserted into a system with the intent of compromising the target in terms of confidentiality, integrity, or availability.

3) Infrastructure: Infrastructure is a type of TTP and refers to the resources of the attackers available to perform attacks. Examples of adversarial infrastructure include command and control servers, malware delivery sites, and phishing sites.

Tools: Attackers install and use tools within the victim's network. The tools often are modified so that a tool detected and analyzed in a previous security incident might be similar, but not exactly the same in new attacks. Malware is a subcategory of tools. In addition, tools might be non-malicious software (e.g., vulnerability scanners, network scanning tools) used for malicious reasons.

Indicators of Compromise: IOCs are detective in nature and describe how to recognize malicious or suspicious behavior that directly detects campaigns, TTPs, attack patterns, malware, tools, and threat actors. To create a good IOC it is desirable to combine different types of information, like atomic indicators, behavioral indicators, and computed indicators related to TTPs often referred to as "ABC" [7].

Atomic Indicators: The value of atomic indicators is limited due to the short shelf life of this type of information and can include file hashes, domain names, IPs and many more. This is the type of data and information that has the longest history in cyber threat intelligence and many threat intelligence efforts are based upon.

Target: Targets can represent organizations, companies, sectors, nations, and individuals.

Courses of Action: Courses of Action refer to measures that can be taken to prevent or respond to attacks.

C. Evaluation Criteria

In the next section of the article we cover taxonomies, sharing standards, and ontologies relevant to threat intelligence and analyze them based on the following criteria:

- Data and concepts covered based on the CTI model (Table 1).
- Connections (relationships) with other taxonomies and ontologies (Sections III, IV).
- Critical analysis of the ontologies based on the description provided in their publications or documentation, as well as their source files (Sections IV, V).

Some identified articles present ontologies which are not described in great detail and have no reference to the actual ontology (rdf/owl files), thus making their evaluation a hard task to achieve. Furthermore, some available ontologies do not offer an additional publication and most of the times not even proper documentation.

Table 1 shows concisely the results of our research conducted on taxonomies, sharing standards, and ontologies based on the CTI model.

III. Taxonomies and Sharing Standards

This section provides an overview of taxonomies and sharing standards that are used or can be used in cyber threat intelligence. We categorize them as enumerations, scoring systems, and sharing standards.

A. Enumerations

Threat Agent Library (TAL) [8] is a set of standardized definitions and descriptions to represent significant threat agents. The library does not represent individual threat actors, thus it is not intended to identify people, or investigating actual security events. The goal of TAL is to help in risk management and specifically to identify threat agents relevant to specific assets. In that way security professionals pro-actively can build defenses for specific threats. In our opinion, the defined

"hostile" threat actor types in TAL library can be used in combination to Mitre's ATT&CK taxonomy which provides a collection of known threat actors and their known tactics and techniques. The connection of the two aforementioned taxonomies would result in the introduction of a new taxonomy which classifies threat actors. An example is state actors that have government resources and their skill are considered adept.

Casey in 2015 [5] introduced a new taxonomy for cyberthreat motivations. The taxonomy identifies drivers that cause threat actors to commit illegal acts. Knowing these drivers could indicate the nature of the expected harmful actions.

Mitre's Common Vulnerabilities and Exposures (CVE) [9] dictionary provides common identifiers for publicly known information-security vulnerabilities in software packages.

NIST's National Vulnerability Database repository (NVD) [10] includes databases of security checklists, security related software flaws, mis-configurations, product names, and impact metrics (CVSS). NVD is built upon CVE and integrates CPE, as well as CWE into the scoring (impact metrics) of CVE entries.

Mitre's Common Platform Enumeration (CPE) [11] specification defines standardized machine readable methods for assigning and encoding names to IT product classes (software and hardware).

Mitre's Common Weakness Enumeration (CWE) [12] is a dictionary of software security weaknesses and vulnerabilities based in part on CVE aiming to better understand flaws in software and to propose adequate countermeasures. Their dictionary includes summaries of the attacks, the prerequisites of launching these attacks, and mitigation solutions.

Mitre's Common Attack Patterns Enumerations and Characteristics (CAPEC) [13] provides a collection of the most common techniques (methods) used in cyber attacks resulting from CWE. Like CWE, CAPEC includes summaries , attack prerequisites, and solutions (countermeasures) of the most common attack patterns.

Mitre's Adversarial Tactics, Techniques and Common Knowledge (ATT&CK) [14] provides a collection of known actors, their known tactics (10 tactic categories), and post-compromise techniques to achieve their objectives. The difference between CAPEC and ATT&CK is that the first one enumerates a range of attack patterns across the entire cyber attack life cycle whereas the latter provides comprehensive coverage across a range of post-compromise techniques. In addition to the techniques observed in ATT&CK includes tools that have been used by specific threat actors which are connected with specific techniques. Overall, these taxonomic connections help us to correlate identified indicators and TTPs to threat actor identities.

B. Scoring Systems

NIST's NVD Common Vulnerability Scoring System (CVSS) [15] is a measurement standard aiming to score vulnerabilities accurately based on their severity; as a result, CVSS enables prioritization vulnerability remediation activities.

Mitre's Common Weakness Scoring System (CWSS) [16] is part of CWE and it provides a mechanism for scoring weaknesses (CWEs) using 18 different factors. Worthy to mentioning is that Mitre's Common Weakness Risk Analysis Framework (CWRAF) can be used in conjunction with CWSS to identify the most important CWEs applying to a particular business and their deployed technologies. The difference between CVSS and CWSS is that the first one targets specific software vulnerabilities scoring, whereas the latter one targets CWE scoring.

C. Sharing Standards

A study of existing threat intelligence sharing initiatives [17] concludes that structured threat information eXpression (STIX) is currently the most used standard for sharing structured threat information. STIX [18] is an expressive, flexible, and extensible representation language used to communicate an overall piece of threat information. STIX architecture is comprised of several cyber threat informations such as cyber observables, indicators, incidents, adversaries tactics, techniques, procedures, exploit targets, courses of action, cyber attack campaigns, and threat actors. Furthermore, STIX was recently redesigned and as a result, omits some of the objects and properties defined in the first version. The objects chosen for inclusion in the second version represent a minimally viable product (MVP) that fulfills basic consumer and producer requirements for CTI sharing. Both standards can be used and adapted based on an organization's needs. It is worth pointing out that MITRE offers additionally Malware Attribute Enumeration and Characterization (MAEC) [19], which is a very expressive malware sharing language for encoding and communicating high-fidelity information about malware based upon attributes such as behaviors, artifacts, and attack patterns. MAEC can be integrated in STIX or used as a standalone.

OpenIOC, originally developed by Mandiant, is an extensible XML schema that enables you to describe the technical characteristics that identify a known threat, an attacker's methodology, or other evidence of compromise. The types of information covered directly by OpenIOC are derived mainly by low level atomic indicators, comprising indicators of compromise, thus covering the IOC category of the CTI model.

IV. Ontologies

Since the work of Blanco et al. [20] in 2008, we have not found any overviews of existing ontologies within the cyber security domain. The authors remark that the scientific community has not accomplished a general security ontology because most of the works are focused on specific domains or the semantic web. The same conclusion was drawn by Fenz and Ekelhart [21]. Additionally, Blanco et al. [20] emphasize the complication of combining their identified ontologies due to the non-common interpretation and different terms applied

for similar concepts in different ontologies. Our study confirms the same almost 10 years after the study of Blanco et al. [20].

Cyber threat information is a small subsection of the information relevant to cyber security and the full security domain. While several ontologies relevant to cyber security and security analytics exist, few ontologies related to threat information and threat intelligence can be identified. We have listed the ontologies discovered relevant to cyber threat intelligence and some more general security ontologies that look promising, at least conceptually, to be taken under consideration when working towards a full cyber threat intelligence ontology. In addition, for many ontologies, relation to specific CTI categories is a tough assignment due to their limitation of being described at a very high level. For most of the ontologies we were unable to find the relevant rdf/owl files even though many of them are called ”open-source” by the authors. The ontologies analyzed hereafter are listed chronologically based on the publication date.

Stefan Fenz and Andreas Ekelhat [21] described an information security ontology that can be used to support a broad range of information security risk management methodologies. The high level concepts of the ontology are based on the security relationship model described in the National Institute of Standards and Technology Special Publication 800-12 and is comprised of the threat, vulnerability, control, attribute, and rating concepts to represent the information security domain knowledge. In addition, concepts such as asset, organization, and person are necessary to formally describe organizations and their assets. Lastly, the concept of location is integrated combined with a probability rating concept to interrelate location and threat information in order to assign priory threat probabilities. Like most of the works the authors have difficulties to connect unambiguous concepts from different standards such as the distinction between threats and vulnerabilities.

Wang and Guo [22] proposed an ontology for vulnerability management and analysis (OVM) populated with all existing vulnerabilities in NVD. The basis of the ontology is built on the results of CVE and its related standards such as CWE, CPE, CVSS, and CAPEC. OVM captures the relationships between the following concepts which constitute the top level of the ontology; vulnerability, introduction phase (software development life cycle - time periods during which the vulnerability can be introduced), active location (location of the software where the flaw manifests), IT product, IT vendor, product category (such as web browsers, application servers, etc.), attack (integration of CAPEC), attack intent, attack method, attacker (human being or software agent), consequence, and countermeasure.

Obrst et al. [23] suggested a methodology for creating an ontology based on already well-defined ontologies that can be used as modular sub-ontologies. In addition, they remark the usefulness of existing schemas, dictionaries, glossaries, and standards as a form of knowledge acquisition of the domain by identifying and analyzing entities, relationships, properties, attributes, and range of values that can be used in defining an ontology. Their suggested ontology is based on the diamond model of malicious activity [24], which expresses the relationships between an adversary (actor), the capabilities of the adversary, the infrastructure or resources the adversary utilizes, and the target of the adversary (victim). The authors state that they developed first the aspects of infrastructure and capabilities, but they are still not in the level of detail they desire. In addition, their current ontology is focused on malware and some preliminary aspects of the diamond model.

A good argumentation for transitioning from taxonomies to ontologies for intrusion detection was made in 2003, by Undercoffer et al. [25]. They suggested an ontology that would enable distributed anomaly-based host IDS sensors to contribute to a common knowledge-base, which again would enable them to quicker identify a possible attack.

Based on this, More et al. [26] in 2012, suggested to build a knowledge-base with reasoning capabilities to take advantage of an extended variety of heterogeneous data sources, to be able to identify threats and vulnerabilities. Their data sources suggest that data retrieved and included in the ontology is within the atomic indicators category of the CTI model.

Oltramari et al. [27] proposed a three layer cyber security ontology named ”CRATELO” aiming to improve the situational awareness of security analysts, resulting to optimal operational decisions through semantic representation. Following the methodology of [23], the authors build upon existing ontologies and expand them. Specifically, CRATELO includes the top level ontology DOLCE-SPRAY extended with a security related - middle level ontology (SECCO) capable to capture details of domain specific scenarios such as threat, vulnerability, attack, countermeasure, and asset. The low level sub-ontology, cyber operations (OSCO), is the extension of the middle level ontology.

Gregio et al. [28] suggested an ontology to address the detection of modern complex malware families whose infections involve sets of multiple exploit methods. To achieve this, they created a hierarchy of main behaviors each one of them consisting of a set of suspicious activities. Then they proposed an ontology that models the knowledge on malware behavior. They state that a given program behaves suspiciously if it presents one or more of the six events (main behaviors) described below which consist of several characteristics. The events are attack launching, evasion, remote control, self-defense, stealing, and subversion. When new set of process actions with malicious behaviors appear (input from ”transformed” log files), the ontology can be inferred to see if an instance of suspicious execution is linked to a malware sample.

Salem and Wacek [29] designed a data extraction tool called TAPIO (Targeted Attack Premonition using Integrated Operational data) which is specialized in extracting data (natural language processing) and automatically map them into a fully linked semantic graph accessible in real time. Part of TAPIO is a cyber security ontology going by the name Integrated Cyber Analysis System (ICAS) that ingests extracted data (logs and events) from several sources to provide relationships across an enterprise network. The tool aims to help incident response teams in connecting and correlating events and actions into an

ontology for automatic interpretation. ICAS is a collection of 30 sub-ontologies specializing in specific conceptual areas as part of host based and network based conceptual models.

Iannacone et al. [30] described their STUCCO ontology, which is developed to work on top of a knowledge graph database. The STUCCO ontology design is based upon scenarios of use by both human and automated users and incorporates data from 13 different structured data sources with different format. The data included in the current STUCCO ontology fall into the categories identity, TTPs, tools, and atomic indicators of the CTI model. Their future work included extending the ontology to support STIX.

Gregio, Bonacin, de Marchi, Nabuco, and de Geus [31] expanded the work of Gregio et al. [28] and introduced the malicious behavior ontology (MBO). MBO is capable of detecting modern complex malware families whose infections involve sets of multiple exploit methods, by applying SWRL rules to the ontology for inferencing. In addition, these rules also apply metrics to specify whether a program is behaving maliciously or not and specifically, how suspicious the execution of a program is. The authors state that their model is able to detect unknown malicious programs even in cases where traditional security mechanisms like antivirus are not, by performing automatic inference of suspicious executions in monitored target systems. However, the current state of the ontology has some limitations such as performance issues, cannot detect malware in real time, and false positives and negatives.Based on its operation MBO can provide useful indicators of compromise for malware.

Fusun et al. suggested ontologies for quantifying attack surfaces [32]. Their Attack Surface Reasoning (ASR) gives a cyber defender the possibility to explore trade-offs between cost and security when deciding on composition of their cyber defense. Ontologies created include those of attacks, systems, defenses, missions and metrics. ASR is mainly modeled after the Microsoft STRIDE [33] threat classification framework, which categorizes attack steps into 6 categories and is to the extent of our knowledge not the preferred framework within threat intelligence community due to its lack of details. In comparison, CAPEC and CPE have around 500 and 1000 ”categories” respectively.

As part of their study on using security metrics for security modeling, Pendelton et al. suggested the Security Metric Ontology [34]. Their ontology includes four sub-ontologies; vulnerability, attack, situations and defense mechanisms, and describes the relationship between them. The terminology used is somewhat different than that of known taxonomies, and their aim at modeling metrics is more prominent than that of analysis and reasoning. Their ontology is published on GitHub[2].

Unified Cybersecurity Ontology was suggested by Syed et al. [35] in 2016. It serves as a backbone for linking cyber security and other relevant ontologies. There are mappings to aspects of STIX, and references to CVE, CCE, CVSS, CAPEC, STUCCO and KillChain. The mappings are loosely connected at a very high level. It is worthy to note that they do not make use of OWL constructs which reduces the reasoning capabilities of the ontology. In addition, their use of domain and range restrictions would result in faulty classification when used with a reasoner. Their ontology is published on GitHub[3].

Unified Cyber Ontology has been introduced on GitHub[4], without any academic publications to date and no actual rdf/owl files yet. Their model ontology is however interesting as it originates from the creators of STIX, which is currently the most used format for sharing threat intelligence [17]. The content of their work is driven primarily from the initial base requirements of expressing cyber investigation information and is the product of input from the Cyber-investigation Analysis Standard Expression (CASE) community (CASE)[5].

Without any publication we find the Cyber Intelligence Ontology (CIO), published only on GitHub[6] to be relevant. This GitHub repository includes most of the mentioned taxonomies and sharing standards in this article, extended and transformed in OWL. The limitation of those ontologies is that they are not connected or unified. For the aforementioned reason we do not classify CIO, since details can be found by checking the relevant taxonomies and sharing standards described in this paper.

V. Discussion

Threat intelligence demands great attention from any organization entailing advanced cyber-threat detective and preventive capabilities. The goal of threat intelligence is to gain rich evidence that can aid decision making, thus the maturity, the skills, and the information sources of a security team define their capability to produce accurate and actionable threat information [39] [40]. Security teams of any maturity and skills can benefit from information sharing activities allowing someone's detection to become another's prevention. By exchanging threat information, organizations can leverage the collective knowledge to get a better understanding of threats an organization might face and consequently, improve their security posture.

Relationships and reasoning: For leveraging the power of ontologies and description logic all the abstraction layers of the CTI model need to be introduced and taken under consideration by formalizing their relationships. In the ontologies evaluated we concluded that lack of OWL constraints is a common phenomenon. Constraints (restrictions) are what makes OWL powerful and enable its reasoning capabilities by inferring information from asserted information. In addition, most of the ontologies explicitly target specific sub-domains of threat intelligence, thus limiting the decision making process in the presence of an observed threat. Successful integration, operation, and advanced reasoning capabilities require existing taxonomies, standards, and vocabularies interconnected in

[2]https://github.com/marcusp46/security-metrics-ontology

[3]https://github.com/Ebiquity/Unified-Cybersecurity-Ontology
[4]https://github.com/ucoProject/uco
[5]https://github.com/casework/case
[6]https://github.com/daedafusion/cyber-ontology

TABLE I
CTI EVALUATION: TAXONOMIES, SHARING STANDARDS, AND ONTOLOGIES

		Identity	Motivation	Goal	Strategy	TTP	Tool	IOC	Atomic Indicator	Target	COA
Taxonomies	TAL [8]	*									
	Threat Agent Motivation [5]	*	*								
	CVE [9]								*		
	NVD [10]								*		
	CPE [11]								*		
	CWE [12]					*			*		*
	CAPEC [13]					*		*			*
	ATT&CK [14]	*				*	*				
	CVSS [15]								*		
	CWSS [16]								*		
Sharing Standards	STIX 1 [18]	*	*	* (Intended Effect:taxonomy)	*	*	*	*	*	*	*
	STIX 2 [36]	*	*	* (Objectives:string)	*	*	*	*	*	*	*
	MAEC [19]							*			
	OpenIOC [37]					*	*	*	*		
Ontologies	Fenz & Ekelhat (2009) [21]								*		
	Wang & Guo (2009) - OVM [22]					*			*		*
	Orbst et al. (2012) [23]	*					*		*	*	
	More et al. (2012) [26]					*			*		
	Oltramari et al. (2014) - CRATELO [27]	*				*			*	*	
	Gregio et al. (2014) [28]						* (malware)		*		
	Salem & Wacek (2015) - ICAS [29]					*			*		
	Iannacone et al. (2015) - STUCCO [30]	*				*	*		*		
	Gregio et al. (2016) - MBO [31]						* (malware)	*	* (it may provide)		
	Fusun et al. (2015) - ASR [32]					*			*	*	
	Pendelton et al. (2016) - Security Metrics Ontology [34]					*					*
	Syed et al. (2016) - UCO [35]	*	*	*	*	*	*	*	*		*
	Unified Cyber Ontology (2016) - UCO [38]	*	*	*	*	*	*	*	*	*	*

addition to human domain expertise to create an ontology that reasons between all the abstraction layers. Furthermore, we cannot ignore the lack of interconnection between taxonomies related to motivations, goals, and strategies of the attackers which can be used multi-purposely. The importance of these taxonomies can be seen in cases that we want to identify which threat actors target particular sectors, and ways of infiltration often used based on their motives, goals, strategies, and TTPs.

Knowledge collection: Much of the knowledge used by most of skilled analysts today is residing only in their heads. If we manage to model this knowledge and express it in an ontology, not only more analysts would be able to consume this type of intelligence, but the analytics would be executed in a consistent way, contradicting the ”confirmation bias” often referred to in an investigation. To be able to express the knowledge of skilled analysts through an ontology, we need to gather their knowledge without expecting them to have prior knowledge of ontologies. We suggest doing this in a iterative way, conducting interviews to pinpoint the process, data sources and actual reasoning points used by highly skilled personnel.

Attribution: Attribution of attacks is the most important element of threat intelligence both for direct recipients and general public. To be able to attribute an attack, evidence of operations is needed that can be linked to an attacker. This entails data and information from different categories of the CTI model. Relating the data points to each other is the task of data enrichment.

The current most common bases for attribution claims include [41] timestamps in executable files; strings, debug paths, and metadata in binary sources such as malware and infected documents; reuse of infrastructure and back-end connections; malware families; code reuse; reused passwords (email accounts, encrypted pieces of code); exploits (0-days); targets (states, secret agencies, etc.).

According to Bartholomew and Guerreri-Saade [41] for the aforementioned bases only sloppy actors or careless operators will provide more data than they should, like debug paths and language strings, or reuse infrastructure from previous attacks. Rid and Buchanan [39] agree with the most common bases for attribution claims but also state that language indicators remain a worthy part of the attribution process. The authors [39] additionally remark that the attackers often re-use software to accomplish basic tasks in their operations for efficiency reasons.

Trust and uncertainty: Attribution is related to uncertainty since intermediate to advanced threat actors are aware of attribution methods and adapt several masquerading techniques. In addition, we find different degrees of knowledge among those sharing threat intelligence enabling sharing of possibly faulty or inaccurate threat intelligence. Added to that, comes the absence of having standardized requirements related to the quality of evidence before shared, which in many cases creates just as large amplifications as the actual threat itself.

For the reasons described we need to take into consideration the level of certainty related to a single piece of information

and the level of trust we have in a given source, that being human or computer device. To address this issue we suggest the use of subjective logic [42] in modeling trust of sources, and confidence in pieces of intelligence that can result in expanding the situational awareness of a security analyst.

VI. Conclusion

Our study concludes that there is not any existing ontology readily available for use within cyber threat intelligence. The main shortcoming is the lack of expressiveness resulting from their poor development and the fact that none of them covers all the relevant data and information (abstraction layers) needed for effective cyber threat intelligence. We suggest several tasks that need addressing in order to create a multi-layered cyber threat intelligence ontology. First, formal terminology (definitions) and vocabularies should be described. Second, all the abstraction layers of the cyber threat intelligence model should be included and expressed properly in the ontology. Third, knowledge coming from domain expertise in a structured way should be gathered and formally represented in the ontology to facilitate advanced reasoning based on relationships between data. Fourth, constraints should be defined and constructs should be used in the ontology enabling the reasoning capabilities lying within the OWL language. Finally, the use of subjective logic to model trust in sources and confidence in information.

References

[1] K. McConkey, "Cybercrime: A Boundless Threat," https://www.pwc.com/gx/en/services/advisory/forensics/economic-crime-survey/cybercrime.html.

[2] S. Smith, "Cybercrime will Cost Businesses over $2 Trillion by 2019," https://www.juniperresearch.com/press/press-releases/cybercrime-cost-businesses-over-2trillion.

[3] R. Stillions, "The DML Model," http://ryanstillions.blogspot.com/2014/04/the-dml-model_21.html.

[4] S. Bromander, A. Jøsang, and M. Eian, "Semantic Cyberthreat Modelling," in *STIDS*, 2016, pp. 74–78.

[5] T. Casey, "Understanding cyber threat motivations to improve defense," *Intel White Paper*, 2015.

[6] A. Fishbach and M. J. Ferguson, "The goal construct in social psychology." 2007.

[7] SANS, "Security Intelligence: Attacking the Cyber Kill Chain," https://digital-forensics.sans.org/blog/2009/10/14/security-intelligence-attacking-the-kill-chain/.

[8] T. Casey , "Threat Agent Library Helps Identify Information Security Risks," *Intel White Paper, September*, 2007.

[9] MITRE, "Common Vulnerabilities and Exposures," https://cve.mitre.org.

[10] NIST, "National Vulnerability Database," https://nvd.nist.gov/.

[11] Mitre, "Common Platform Enumeration," https://cpe.mitre.org/specification/.

[12] MITRE, "Common Weakness Enumeration," https://cwe.mitre.org.

[13] MITRE, "Common Attack Pattern Enumeration and Classification," https://capec.mitre.org/.

[14] MITRE, "Adversarial Tactics, Techniques and Common Knowledge," https://attack.mitre.org/.

[15] NIST, "Common Vulnerability Scoring System," https://nvd.nist.gov/vuln-metrics/cvss.

[16] MITRE, "Common Weakness Scoring System," https://cwe.mitre.org/cwss/cwss_v1.0.1.html.

[17] C. Sauerwein, C. Sillaber, A. Mussmann, and R. Breu, "Threat Intelligence Sharing Platforms: An Exploratory Study of Software Vendors and Research Perspectives," 2017.

[18] S. Barnum, "Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX ™)," *MITRE Corporation*, vol. 11, 2012.

[19] Mitre, "Malware Attribute Enumeration and Characterization," https://maec.mitre.org.

[20] C. Blanco, J. Lasheras, R. Valencia-García, E. Fernández-Medina, A. Toval, and M. Piattini, "A Systematic Review and Comparison of Security Ontologies," in *Availability, Reliability and Security, 2008. ARES 08. Third International Conference on*. Ieee, 2008, pp. 813–820.

[21] S. Fenz and A. Ekelhart, "Formalizing Information Security Knowledge," in *Proceedings of the 4th international Symposium on information, Computer, and Communications Security*. ACM, 2009, pp. 183–194.

[22] J. A. Wang and M. Guo, "OVM: An Ontology for Vulnerability Management," in *Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research: Cyber Security and Information Intelligence Challenges and Strategies*. ACM, 2009, p. 34.

[23] L. Obrst, P. Chase, and R. Markeloff, "Developing an Ontology of the Cyber Security Domain," in *STIDS*, 2012, pp. 49–56.

[24] S. Caltagirone, A. Pendergast, and C. Betz, "The diamond model of intrusion analysis," DTIC Document, Tech. Rep., 2013.

[25] J. Undercoffer, A. Joshi, and J. Pinkston, "Modeling Computer Attacks: An Ontology for Intrusion Detection," in *International Workshop on Recent Advances in Intrusion Detection*. Springer, 2003, pp. 113–135.

[26] S. More, M. Matthews, A. Joshi, and T. Finin, "A Knowledge-Based Approach to Intrusion Detection Modeling," in *Security and Privacy Workshops (SPW), 2012 IEEE Symposium on*. IEEE, 2012, pp. 75–81.

[27] A. Oltramari, L. F. Cranor, R. J. Walls, and P. D. McDaniel, "Building an Ontology of Cyber Security," in *STIDS*. Citeseer, 2014, pp. 54–61.

[28] A. Grégio, R. Bonacin, O. Nabuco, V. M. Afonso, P. L. De Geus, and M. Jino, "Ontology for Malware Behavior: a Core Model Proposal," in *WETICE Conference (WETICE), 2014 IEEE 23rd International*. IEEE, 2014, pp. 453–458.

[29] M. B. Salem and C. Wacek, "Enabling New Technologies for Cyber Security Defense with the ICAS Cyber Security Ontology," in *STIDS*, 2015, pp. 42–49.

[30] M. Iannacone, S. Bohn, G. Nakamura, J. Gerth, K. Huffer, R. Bridges, E. Ferragut, and J. Goodall, "Developing an Ontology for Cyber Security Knowledge Graphs," in *Proceedings of the 10th Annual Cyber and Information Security Research Conference*. ACM, 2015, p. 12.

[31] A. Grégio, R. Bonacin, A. C. de Marchi, O. F. Nabuco, and P. L. de Geus, "An Ontology of Suspicious Software Behavior," *Applied Ontology*, vol. 11, no. 1, pp. 29–49, 2016.

[32] M. B. Fusun, A. S. Yaman, T. Marco, E. Carvalho, and C. N. Paltzer, "Using Ontologies to Quantify Attack Surfaces."

[33] A. Shostack, *Threat Modeling: Designing for Security*. John Wiley & Sons, 2014.

[34] M. Pendleton, R. Garcia-Lebron, J.-H. Cho, and S. Xu, "A Survey on Systems Security Metrics," *ACM Computing Surveys (CSUR)*, vol. 49, no. 4, p. 62, 2016.

[35] Z. Syed, A. Padia, M. L. Mathews, T. Finin, and A. Joshi, "UCO: A Unified Cybersecurity Ontology," in *Proceedings of the AAAI Workshop on Artificial Intelligence for Cyber Security*. AAAI Press, 2016.

[36] OASIS CTI TC, "Structured Threat Information Expression (STIX ™) 2.0," https://oasis-open.github.io/cti-documentation/, 2017.

[37] Mandiant Corporation, "Sophisticated Indicators for the Modern Threat Landscape: An Introduction to OpenIOC," http://www.openioc.org/resources/An_Introduction_to_OpenIOC.pdf, 2013.

[38] "Unified Cyber Ontology," https://github.com/ucoProject/uco, 2016.

[39] T. Rid and B. Buchanan, "Attributing Cyber Attacks," *Journal of Strategic Studies*, vol. 38, no. 1-2, pp. 4–37, 2015.

[40] C. Johnson, L. Badger, D. Waltermire, J. Snyder, and C. Skorupka, "Guide to cyber threat information sharing," *NIST Special Publication*, vol. 800, p. 150, 2016.

[41] B. Bartholomew and J. A. Guerrero-Saade, "Wave your False Flags! Deception Tactics Muddying Attribution in Targeted Attacks," in *Virus Bulletin Conference*, 2016.

[42] A. Jøsang, "Artificial Reasoning with Subjective Logic," in *Proceedings of the second Australian workshop on commonsense reasoning*, vol. 48. Perth:[sn], 1997, p. 34.

Paper IV

Examining the known truths of Cyber Threat Intelligence - the case of STIX

Siri Bromander, Lilly Muller, Martin Eian, Audun Jøsang

Proceedings of the 15th International Conference on Cyber Warfare and Security (ICCWS 2020).
Center for Cybersecurity Education and Research (CCSER), Old Dominion University (ODU), Norfolk, Virginia, USA.
ISSN:2049-9870

IV

Examining the "Known Truths" in Cyber Threat Intelligence – the Case of STIX

Siri Bromander(1)(3), Lilly Pijnenburg Muller(2), Martin Eian(3), Audun Jøsang(1)
(1)Department of Informatics, University of Oslo, Oslo, Norway
(2)War Studies Department, Faculty of Social and Political Science, Kings College London, United Kingdom
(3)Research and Development, mnemonic, Oslo, Norway
siri@mnemonic.no
lilly.muller@kcl.ac.uk
meian@mnemonic.no
josang@ifi.uio.no

Abstract
Treat intelligence has played a key role in keeping networks secure for as long as computers have communicated. With the aim to collaboratively defend against the increasing threats in and from cyberspace, Cyber Threat Intelligence (CTI) has risen in popularity, and in correlation a growth of concepts and terms within the field is observable. While largely benefiting the ability to secure networks, this growth has also led to assumptions and confusion surrounding how technical experts work in practice. This paper examines how the technical threat intelligence community share CTI and their methods and practices. Through a combined approach drawing on ethnography, interviews and questionnaires we examine aspects of knowledge sharing in the CTI community based on using Structured Threat Information Expression (STIX). It is found that while sharing threat intelligence is deemed to be crucial, the adoption of STIX is hindered by strict policies for classification and trust, unclear use of terminology and too much flexibility within STIX. While the flexibility in STIX has played a vital part in the ability to structure CTI and build bridges between CTI communities, this paper argues that with growing amounts of data, new possibilities and tools for data-based analysis, increased precision within terminology and definitions is needed to advance the field. The flexibility of STIX allows wide applicability, but also causes lack of precision in the expression of CTI. This reduces the possibility for data analytics and creates a potential for false expectations in the development of the field of CTI, thereby representing a global security concern for private and state actors alike.

Keywords: Cyber Threat Intelligence, CTI, Knowledge representation, Security, Sharing CTI, STIX

1. Introduction

The cyber domain is recognized as the new battlefield where modern day conflicts are fought. To defend ourselves we need to cooperate, through executing and sharing threat intelligence. In July 2016, NATO recognized cyberspace as a domain equal to air, land and sea (NATO, 2016). To achieve successful defense it is essential to establish collaboration and active sharing of intelligence between partner organizations. This is similar to traditional threat intelligence, which is the foundation for cyber threat intelligence. Although they use different tools, the content and processes are relatively similar. Yet, some crucial difference are important to note. Firstly, in contrast to physical-world threat intelligence, in cyberspace it is possible to change the terrain or battlefield to one's own benefit, as digital networks are privately owned. Secondly, the "footprints" or signatures of threat actors are mainly found through detecting them in networks where attacks have taken place. The availability of data about the actors is thus often tied to the infrastructure within which they operate. Hence, the information needed to defend preemptively is dependent on collaboration. Exchange of intelligence regarding the digital battlefield is arguably less relevant than intelligence about the actual threats. In this way, the scope of cyber threat intelligence is slightly narrower than physical/traditional intelligence. In this paper we use the abbreviation CTI (Cyber Threat Intelligence) when discussing digital threat intelligence. We use the term CTI to refer to both the process of creating CTI and the shareable results of such processes.

1.1 The Emergence of CTI

Exploring available search engine data shows that the term "cyber threat intelligence" has had growth since 2013, while "threat intelligence" has been searched for continuously, but somewhat increasingly since 2004 (Google, 2019). A strong peak is visible in February 2015, following the White House publication of the "Cyber Threat Intelligence Integration Center" (White House, 2015). Even though the term "Cyber Threat Intelligence" is relatively new, we have observed threat intelligence work in the digital domain for as long as computers have been connected to networks. One of the first public cases documenting the use thereof is during the Morris-worm in 1988: the first worm to spread itself across digital networks, with significant consequences for the global computer network of that time. An analysis of the malware after the incident revealed a simple program, which despite its simplicity had enormous consequences due to its innovative way of spreading (Spafford, 1989). As the outbreak started, a technical collaboration was created through an emailing list where technical personnel exchanged information on how the worm had been identified, how to handle infections, and how to protect against infection (Phage, 1988). This type of collaboration is exactly what we now refer to as CTI, meaning the people behind the mailing list were pioneers in the field of CTI.

1.2 Tactical CTI

"Threat intelligence is evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice, about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject's response to that menace or hazard" (Gartner, 2013). We emphasize "evidence-based knowledge", as CTI needs to be based upon evidence in order to be trusted. It is also of importance that this type of knowledge can be used for mitigating threats, and hence needs to be actionable. The Allied Joint Procedures published and used by NATO (NATO, 2016), concur with this definition of threat intelligence, elaborating on the differences between threat data, threat information and threat intelligence. Threat data can be processed to become threat information, which needs structure and adoption for a given audience in order to qualify as threat intelligence.

In "The Pyramid of Pain" from 2013 (Bianco, 2013) David Bianco presents a model for representation of CTI-relevant data. The model is presented in different formats in most contexts where CTI is discussed. The pyramid emphasizes what data and knowledge to prioritize for robustness against digital threats. At the bottom of the pyramid are hash values of malware samples, which is simple to detect for the defenders, but equally simple to change for the attackers. At the top of the pyramid, Bianco places TTPs (Tactics, Techniques and Procedures) as the hardest for the attackers to change (to avoid detection), but also the hardest for defenders to identify. "TTP" is a concept with long traditions in intelligence, but is challenging to use within CTI, as it is difficult to specify what exactly tactics, techniques and procedures are when sharing and (automatically) processing CTI. The development of capabilities for detecting threats based on TTPs are nevertheless of utmost importance for further development of our collective defense in the cyber domain. In 2014 Ryan Stillions published a similar model, the "Detection Maturity Level" (DML) model (Stillions, 2014) seen in Figure 1.

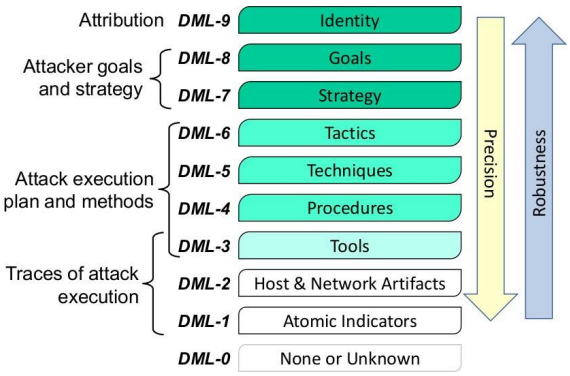

Figure 1 A slightly expanded version of the DML model (Bromander et al., 2016)

Stillions' DML model specifies hierarchical levels/types of data an organization should be able to consume to be at given level. In this case, to be "able to consume" means not only receive and understand the data or information, but also to be able to take useful action based on it. The model has proven to be useful in several contexts, for example for describing the relevant data exchanged within CTI (Bromander S. M., 2016) where tactics, techniques and procedures are behavioral indicators.

Chismon and Ruks (2015) proposed a model representing the current types of CTI, and to what degree they were detailed and of long-term use. This is illustrated in Figure 2.

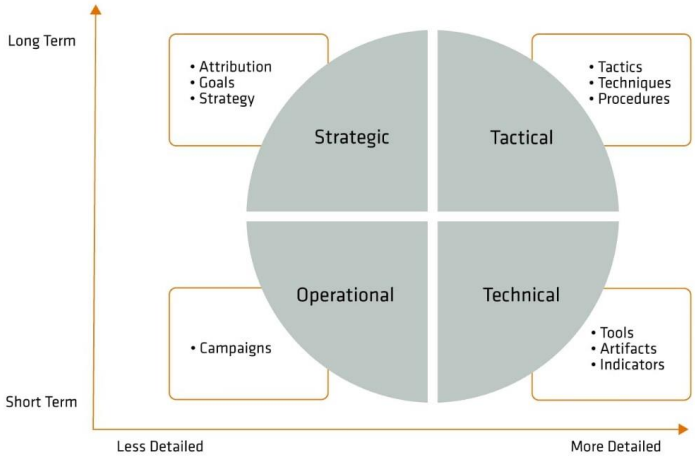

Figure 2 (David Chismon, 2015).

The model of Chismon and Ruks illustrates that the more detailed the knowledge is, the more certainty about the threat actor's presence and identity can be obtained, and that the more robust (long term) it is, the longer the knowledge is useful for the defenders. Tactical threat intelligence, consisting of TTPs, is of most value in the attempt to detect and prevent future attacks. Yet, this is one of the least developed areas within CTI, and this is the area with the strongest focus in research and development within CTI.

1.3 Structured Threat Intelligence and the Myth of "Everyone Uses STIX"

To stimulate development in the field of CTI we need a structured way of representing data, information and knowledge about cyber threats. Several initiatives exist in this regard, and the one gaining most attention has been Structured Threat Information Expression (STIX). Proposed by Barnum et al. in 2012 (Barnum, 2012), STIX is an Extensible Markup Language (XML)-based language created to make sharing of CTI more than just sharing data. Since its first release STIX has been updated, through 2.0 in 2017 (OASIS, 2019) with changes in the language structure based on usage experience and feedback from the community. Another major change was to use JavaScript Object Notation (JSON) instead of XML. As a result, STIX now consists of 12 different STIX Domain Objects (SDO) and two different relationship types which have a list of suggestions for relationship names and which SDOs they may connect. STIX is argued to be the de facto standard for representing cyber threat intelligence within the technical CTI community (Sauerwein C. S., 2017). Yet, while STIX is the most used standard for representing CTI, this article questions if the features STIX is praised for are actually used in practice.

Drawing inspiration from anthropology and ethnographical studies following the professionals working with CTI, there are indications that the amount of shared CTI without the use of standards is indeed high, with possible impact on the developments of new standards, models, software and processes of CTI. Sauerwein et al. (2019) confirm this. From this initial analysis, the following the research questions emerged: How do practitioners

understand CTI and how do they share it? Do their understanding and sharing of CTI align with the current development of the field of tools and formal representations of CTI?

2. Methodology

This research paper base on a triangulation of preliminary ethnographical observation of CTI environments, a questionnaire, and semi-structured interviews. In addition, Eclectic IQ performed and presented a statistical analysis on STIX usage presented in (Polzunov, 2019). These results are used in the discussion together with our own results. The most relevant results from their study are described in Section 3.3.
For our questionnaire we have used (Krosnick, 2018) and (Smyth, 2006) as guidance, with the most important choices listed underneath.

- In order to minimize the respondent's fatigue, we chose to keep the number of questions as low as possible. This limits the amount of information we can extract, but presumably increases the quality of the results.
- We chose questions that were easy to reply to, requiring a limited amount of interpretation of the questions, and also a limited amount of retrieval of relevant information in their memory or integration of this information in judgments and final answers (choice of option).
- Where we ask the respondents to estimate percentages we chose to make the questions open in order to escape biases. We chose the classes to be reported based upon the answers we got.
- We chose open questions where possible to mitigate possible biased responses in the case of a non-exhaustive list. The use of "Other, please specify" is not recommended as a solution to this (Krosnick, 2018) and hence not chosen.
- Experimental evidence suggests that checklists should be structured in "did – did not" format as opposed to "check-all-that-apply", partially because respondents take longer to answer forced choice items, and partially because forced-choice answers are easier to interpret (Smyth, 2006). This is the argument for choosing this option in our questionnaire.

The questionnaire was developed in three stages, with an initial version tested on a reference group of four people. Improvements and a new version was tested on a new reference group of 11 persons during the 2019 FIRST CTI Symposium. The final version was created with the input and evaluation of this 2nd version. In order to set the frame for the questionnaire, a limited text was included to introduce it, as seen in (Bromander S. , 2019). The questionnaire was published using a web portal available from the University of Oslo (UiO, 2019). The questionnaire was online from June 15th to August 8th 2019. There is a limited number of CTI professionals, and the method for attracting many to participate was to advertise in public and to time the data collection during the 31st Annual FIRST Conference where many CTI practitioners meet.

Interviews were conducted post-questionnaire, with three participants. The participants were volunteers since the respondents were anonymous, and hence could not be asked or selected for participation. The participants were asked to comment on the questionnaire questions, and the answers were used to better understand the results. The interview guide can be found in (Bromander S. , 2019). As part of the interviews, the participants were asked to represent a given piece of information in STIX, used as a case study in this article.

3. Results

The results from the questionnaire are presented in Section 3.1 through Section 3.3 and are divided into three parts corresponding to the questionnaire design.

The definition of sectors and organization sizes are found in (S&P, 2019) and (EC, 2019). There were 36 respondents to the complete questionnaire. The relatively small number of respondents means that all results must be seen as indicative, not as conclusive. All percentages are rounded off.

The semi-structured interviews contributed to the choice of included results in this paper, and to the discussion found in Section 4.1 and Section 4.2.

Section 3.4 presents the results of a case study to explain the problem with flexibility within STIX, performed as part of the semi-structured interviews.

3.1 About the respondents

Of the 36 respondents, 30 answered that they shared CTI as part of their role. The respondents were representing the sectors, organization size and countries as seen in Figure 3 to Figure 5.

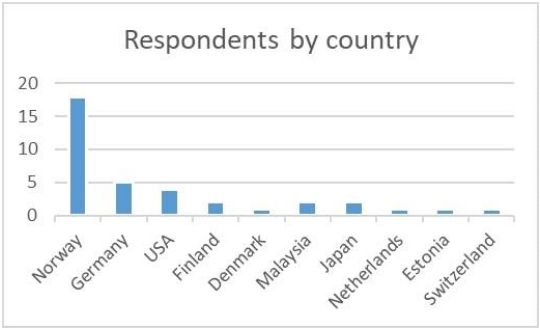

Figure 3 Respondents by country.

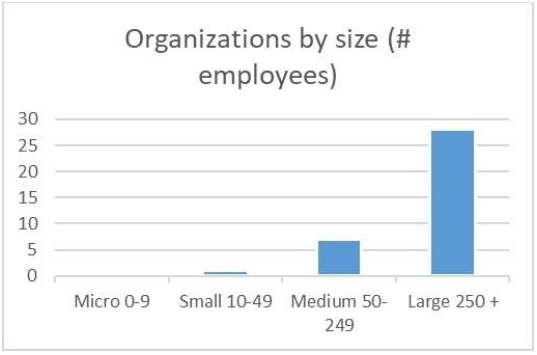

Figure 4 The size of the organizations represented.

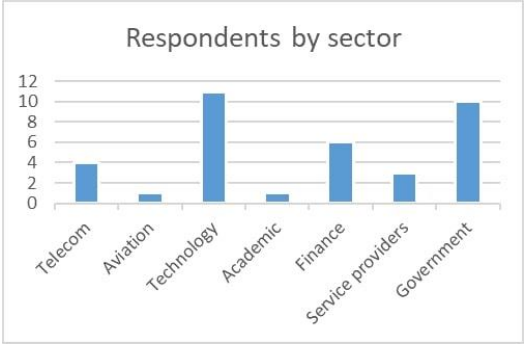

Figure 5 Respondents by sector.

3.2 Sharing CTI

The main conclusions possible to draw from the responses are related to the format of shared CTI, the use of consumed CTI and the reasons for not sharing. All the respondents considered themselves consumers of CTI, and ~70% considered themselves producers of CTI.

3.2.1 Format of shared CTI

The format of consumed CTI ranges from text files, PDF files and articles, to Malware Information Sharing Platform (MISP), STIX, JSON, XML, cvs and txt files. Of the 30 consumers of CTI, the median of the amount of consumed CTI without interrelationships within the data is 80%. The corresponding median for the amount of consumed CTI with interrelationships within the data is 20%. The sectors or the sizes of the organizations do not influence the results.

3.2.2 Storing and Using the Consumed CTI

83% of the respondents reply that they partially or fully store the received data in a structured way, and 37% reply that they use the consumed CTI for further analysis.

3.2.3 Obstacles for Sharing

The Traffic Light Protocol (TLP) (FIRST, 2019) is used for handling and sharing CTI. TLP describes four levels (white, green, amber, red) which dictate how labeled information must be protected and who may receive access. Knowledge of TLP is limited outside the technical community and may hence be problematic to use in communities where TLP is uncommon.

In our results, 77 % of respondents state that the most prominent reasons for not sharing CTI are related to privacy, confidentiality and classification of data. Out of these, TLP is mentioned by 26 % directly. 30 % refer to classification in general, without specifying a specific classification scheme. For 43% of the respondents, the lack of time, motivation and resources are the most prominent reason for not sharing CTI.

3.3 Using STIX

70% of the respondents replied that they have not used STIX in the past six months. Of the 30% that did use STIX, only 66% created a STIX file/bundle. Of those who responded that sharing of CTI was part of their role, 80% had not created a STIX file/bundle in the last six months. The questionnaire results gave little insight into which SDOs and relationship types were in use due to the low number of respondents.

There is a limited amount of STIX feeds available (Polzunov, 2019). The available feeds evaluated by (Polzunov, 2019) indicates that there exists “good” and “bad” use of STIX, and they provide a rich set of metrics for evaluating this. Their analysis implies that there is a large span in the number of object types in use, and that the more objects types are utilized, the more custom fields are included as well. The same is true for the use of relationships. We suggest that utilizing more object and relationship types implies advanced use, and that need of custom fields implies shortcomings in the standard. The results then indicates that the more advanced STIX usage becomes, the more shortcomings emerge.

3.4 Using STIX in CTI Processes

STIX represents one of the most thorough standards for describing CTI. Yet with the rapid developments in the field of CTI the standard as it currently stand suffers from limitations. Firstly, the absence of a top-level element to represent and structure specific company assets such as IT systems affected by an incident is a limitation (Böhm, 2018). STIX relationships are flexible, and without restrictions for usage. Hence, the same relationship may be used between different STIX Domain Objects (SDOs). Since relationships are described within the properties of one object, this gives a potential for confusion in using them. According to (Polzunov, 2019) relationships are only used to a limited extent.

A fast-growing knowledge base for CTI is the ATT&CK framework supplied by Mitre, (Mitre, 2019), which describes threat actors, tools, techniques and tactics. This knowledge base is a major step towards structuring knowledge of tactical CTI in terms of the TTPs previously described. However, the use of STIX to represent the content of ATT&CK is not straight forward, due to tactics not being represented in STIX as a term. MITRE publishes their ATT&CK knowledge base with the use of STIX, but they have had to add custom fields within the “Attack Pattern” SDO in order to include all the information. The lack of explicit options to express this relevant knowledge base within CTI is a shortcoming in STIX.

Secondly, the large flexibility of the STIX language is by itself a weakness. An example piece of information typically shared is “Sad Panda has used 123.456.789 for command and control”. When asked to represent this with STIX 2.0, three different threat intelligence analysts came up with three different representations, and we cannot exclude the possibility that additional different representations would be suggested if additional personnel were asked. Figure 1 presents the simplified code for the three representations.

```
object:
        "type" = "threat-actor"
        "name" = "sad panda"
        "description" = "uses 123.456.789 for c&c."
        "object_refs":
                "indicator-123"
                "attack-pattern-123"

        "type" = "indicator"
        "id" = "indicator123"
        "name" = "IPaddress=123.456.789"

        "type" = "attack-pattern"
        "id" = "attack-pattern-123"
        "name" = "c&c attack pattern used by sad panda"
```

```
object:
        "type" = "campaign"
        "name" = "sad panda"
        "description" = "sad panda uses 123.456.789 for c&c."
```

```
object:
        "type" = "threat-actor"
        "name" = "sad panda"
        "description" = "uses 123.456.789 for c&c."

        "type" = "indicator"
        "name" = "IPaddress 123.456.789"

        "type" = "attack-pattern"
        "name" = "c&c attack pattern used by sad panda"
```

Figure 1Three different representations of the same information using STIX.

The three analysts all place information in description fields using English prose, rather than structured information consumable by a computer. In addition, there are different object types in use, which means that someone consuming this CTI has to look several places in order to ensure consuming everything. Due to the different ways of representing the same information, the possibility of automatic consumption and computer-based analysis becomes limited. If a computer cannot identify information because the information type is not normalized, “Big Data”-style analysis is not possible. As a result, large amounts of manual work is needed to interpret, correct and analyze the data. Furthermore, different ways of representing the same information, whether it is consumed manually or automatically, will result in loss due to the normal behavior of both humans and computers to look for what is known, and then discard the rest.

These limitations of STIX have emerged as a result of new requirements for precise CTI representation, which seem to be in conflict with flexibility, that had to be included in the early versions of STIX in order for practitioners to use it. However, as the maturity and capabilities for precise CTI are increasing, we are in need of a more strict and precise model than what STIX currently offers.

4. Discussion

4.1 CTI in Theory and Practice

Shared data and information may be used to perform or extract CTI, but do not always classify as such. NATO and Gartner are aligned in defining CTI as produced knowledge, which cannot be transferred using standalone data points. If CTI is shared using only standalone data points, vital knowledge regarding the threat actors and their habitats becomes lost in the process, because the interrelationships and context represents the knowledge.

The results of the questionnaire indicate that many professionals share CTI, but when digging deeper into formats and actual shared data, it is found that what is shared is simply threat data, and in some cases threat information. The concept of CTI loses its meaning when used for sharing of threat data. As it currently stands, if an actor claims to send CTI through STIX one cannot trust that one will receive more than simple threat data.

4.2 What does it mean to “use STIX”?

The questionnaire results indicate that many claim to use STIX, but few have actually created STIX bundles. There is no need for creating STIX bundles in order to use it, but to create STIX bundles is how one can actually decide on content and how one can specify how STIX is used.

To “use STIX” in terms of consumption of STIX bundles should involve more than accepting a JSON file and manually consuming it, because the same content may then just as well be shared as English prose using a text file, PDF, emails or a csv file.

The potential benefits of using STIX are to structure information and create a standardized way of sharing CTI suitable for automation. To "use STIX" would then arguably entail fulfilling at least one of the two, and preferably both. The results reported in Section 3.4 show how the usage of different SDOs and relationships varies, which entails that true structure is generally not in place. There are typically different representations of the same CTI, and hence the value of structure is degraded. In addition, as shown in Section 3.4, not all information can be represented with the default SDOs as they are currently defined and hence custom fields and properties must be used. This entails a lack of standardization of all relevant information, which poses a problem for automation. We therefore argue that while many claim to "use STIX", in most cases it is not used as a standardized way of sharing CTI suitable for automation, even when a STIX bundle or file has been created.

4.3 What value does STIX give if not used as a strict format?

The value of using a standardized format is limited if it is not used consistently. This especially holds for IT applications where standards are essential for distributed applications.

The value of representing data or information in a standardized format includes knowing where in a file a certain type of data is found, and what it looks like. If a standard provides flexibility about where to place a certain piece of data, or its format, it reduces the ability of other parties to identify and use this piece of data. If a standard allows ad hoc extensions of representation, the extended parts of the standard requires additional work for other parties to consume. In both cases, the standard will significantly reduce its value for computers and human operators using it for sharing. We argue that STIX currently has this deficiency. This means that the current usage of STIX is not superior to any other standardized way of sharing data that two or more parties have agreed on, based partly on a common vocabulary. The STIX vocabularies are valuable contributions.

If CTI is shared in a format that the recipient does not understand, the recipient's ability to consume the knowledge is limited. The labor-intensive task of agreeing on how to use the standard between two parties can solve this. If no agreement between two parties has been made, or a parsing task has not been conducted on the receiving end, then there will be a significant loss of information in transfer. Most importantly, if the standards are not used consistently the threat intelligence community's ability to know if they are talking about the same threats/information is hindered. Standardized sharing is key for CTI, and for STIX to be useful it needs to be used in line with intent.

4.4 The issue with claimed usage of sharing standards – when ideas do not match reality

Through examination of one of the "known truths" in CTI, this study finds that assumptions made regarding the use of STIX are not valid.

The three main consequences these types of assumptions lead to are all results of using "known truths" as guidance for prioritizing the work and development within a field. Prioritizing a task or a fact means deprioritizing something else.

Firstly, we find that training personnel to use STIX as it stands today takes valuable time away from other types of training that can potentially hold more value. This influences the shortage of technical security personnel (Vogel, 2016) (Crumpler, 2019) in a negative direction as personnel may be less capable and less efficient to do the required work.

Secondly, choice and development of tools and procedures need to adhere to the reality. The field of CTI is highly dependent on technical tools and solutions, and the effectiveness and capabilities of the collective workforce rely on informed choices. Priorities based on imprecise information can lead to a decrease in effectiveness and capability.

Lastly, research and development needs to focus on real world problems and prioritize based upon a rational foundation. Assumptions and "known truths" need validation, and if an assumption is found to be wrong, it must be disseminated to the community. This validation is necessary to steer ongoing research in a direction that can benefit our collective cyber defenses.

Although the research presented here is a micro study this paper touches on critical foundations of the processes of CTI and methods within the field. The ramifications identified can extend to international cooperation. Agreements for International cooperation to improve cybersecurity beyond country borders for a secure international cyberspace can be made on an international level (NATO, 2016). However, if cooperation on the technical level is hindered, the potential for collaboration at a political level is limited. In the current landscape of cyber threats, it is essential to have a common language for sharing and acting on CTI through public and private efforts with the aim to secure cyberspace. Without this, the ability to develop automated tools that are able to utilize the CTI is restricted, and the global community's ability to defend against threats in and from cyberspace is unnecessarily hindered.

5. Conclusion

Through questioning the difference between data, information and intelligence in models and standards used in CTI, this paper finds that while sharing threat intelligence is deemed to be crucial, classification and trust, unclear use of terminology, as well as large flexibility within STIX hinders developments in the field of CTI. While the current flexibility of STIX has allowed for inclusion of a variety of users, the lack of precision reduces the possibility for knowledge transfer and data analytics. To improve this situation, stricter definitions and greater specificity are called for. Increased precision and clearer guidelines can enable a full use of STIX without loss of vital information, which in turn can create possibilities to share knowledge beyond flat files. Such a use of CTI can improve the ability to defend collectively in the cyber domain. While the international community calls for shared efforts to secure cyberspace, little will be successful if the technical ability is not in place to do so.

Acknowledgements

This work was supported by the Research Council of Norway and mnemonic under the ACT and TOCSA projects. The authors would like to thank the anonymous reviewer for valuable input that improved the paper and the respondents of the questionnaire who contributed to this study.

References

Barnum, S. (2012). Standardizing cyber threat intelligence information with the Structured Threat Information eXpression (STIX). *MITRE Corporation, 11*, 1-22.

Bianco, D. (2013). *The Pyramid of Pain*. Retrieved from Enterprise Detection and Response: http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html

Bromander, S. (2019). *Questionnaire: Sharing Cyber Threat Intelligence.* Retrieved from GitHub: https://github.com/sbrom/sharingCTI/

Bromander, S. M. (2016). Semantic Cyberthreat Modelling. (pp. 74-78). STIDS.

Böhm, F. M. (2018). Graph-based visual analytics for cyber threat intelligence. *Cybersecurity, 1*, 16.

Crumpler, W. a. (2019). The Cybersecurity Workforce Gap. Center for Strategic and International Studies. Retrieved from https://csis-prod.s3.amazonaws.com/s3fs-public/publication/190129_Crumpler_Cybersecurity_FINAL.pdf

David Chismon, M. R. (2015). *Threat Intelligence: Collecting, Analysing, Evaluating.* The National Cyber Security Centre UK.

EC. (2019). Retrieved from What is an SME?: https://ec.europa.eu/growth/smes/business-friendly-environment/sme-definition_en

FIRST. (2019). *Traffic Light Protocol.* Retrieved from first.org: https://www.first.org/tlp/

Gartner. (2013). *Definition: Threat Intelligence.* Retrieved from www.gartner.com: https://www.gartner.com/doc/2487216/definition-threat-intelligence

Google. (2019, October 6). *Google Trends*. Retrieved from Trend search for "Cyber threat intellgence": https://trends.google.com/trends/explore?date=all&q=Cyber%20threat%20Intelligence,threat%20intelligence

Krosnick, J. A. (2018). Questionnaire design. In *The Palgrave Handbook of Survey Research* (pp. 439-455). Springer.

Mitre. (2019). *Mitre ATT&CK*. Retrieved from Mitre ATT&CK: https://attack.mitre.org/

NATO. (2016). Allied Joint Doctrine for Intelligence, Counterintelligence and Security Doctrine. *AJP 2.0.*

NATO. (2016, June). *Warsaw Summit Communiqué, paragraph 70+71.* Retrieved from https://ccdcoe.org/sites/default/files/documents/NATO-160709-WarsawSummitCommunique.pdf

OASIS. (2019). *STIX*. Retrieved from https://oasis-open.github.io/cti-documentation/stix/intro.html

Phage. (1988, November). *Phage mailinglist*. Retrieved from listen: http://securitydigest.org/phage/

Polzunov, S. a. (2019, March 5). *EVALUATE OR DIE TRYING - A Methodology for Qualitative Evaluation of Cyber Threat Intelligence Feeds.* Retrieved from FIRST CTI Conference: https://www.first.org/resources/papers/london2019/EVALUATE-OR-DIE-TRYING-Abraham-Polzunov.pdf

S&P. (2019). *Global Industry Classification Standard.* Retrieved from https://www.spglobal.com/marketintelligence/en/documents/112727-gics-mapbook_2018_v3_letter_digitalspreads.pdf

Sauerwein, C. a. (2019). An analysis and classification of public information security data sources used in research and practice. *Computers & Security, 82*, 140-155.

Sauerwein, C. S. (2017). Threat Intelligence Sharing Platforms: An Exploratory Study of Software Vendors and Research Perspectives.

Smyth, J. D. (2006). Comparing check-all and forced-choice question formats in web surveys. *Public Opinion Quarterly, 70*, 66-77.

Spafford, E. H. (1989). The Internet worm program: An analysis. *ACM SIGCOMM Computer Communication Review, 19*(1), 17-57.

Stillions, R. (2014, 04). Retrieved from The DML Model: http://ryanstillions.blogspot.com/2014/04/the-dml-model_21.html

UiO. (2019). *Nettskjema*. Retrieved from https://www.uio.no/english/services/it/adm-services/nettskjema/

Vogel, R. (2016). Closing the cybersecurity skills gap. *Salus Journal, 4*(2), 32.

White House, O. (2015, February 25). *FACT SHEET: Cyber Threat Intelligence Integration Center*. Retrieved from Obama White House Archives: https://obamawhitehouse.archives.gov/the-press-office/2015/02/25/fact-sheet-cyber-threat-intelligence-integration-center

Paper V

Modeling Cyber Threat Intelligence

Siri Bromander, Morton Swimmer, Martin Eian, Geir Skjøtskift, Fredrik Borg

Proceedings of the 6th International Conference on Information Systems Security and Privacy (ICISSP 2020).
Valetta, Malta.
ISBN: 978-989-758-399-5

V

Modeling Cyber Threat Intelligence

Siri Bromander[1,2], Morton Swimmer[3], Martin Eian[1], Geir Skjotskift[1] and Fredrik Borg[1]
[1]*mnemonic AS, Oslo, Norway*
[2]*Department of Informatics, University of Oslo, Norway*
[3]*Trend Micro Research*
siri, meian, geir, fredrikb@mnemonic.no, morton_swimmer@trendmicro.com

Keywords: Cyber Threat Intelligence, Security, Knowledge graph, Ontology

Abstract: For a strong, collective defense in the digital domain we need to produce, consume, analyze and share cyber threat intelligence. With an increasing amount of available information, we need automation in order to be effective. We propose a strict data model for cyber threat intelligence which enables consumption of all relevant data, data validation and analysis of consumed content. The main contribution of this paper is the strictness of the data model which enforces input of information and enables automation and deduction of new knowledge.

1 INTRODUCTION

In recent years we have seen several initiatives to structure and streamline Cyber Threat Intelligence (CTI). Organizations share CTI, and most of the CTI in an average organization comes from external sources. Consuming, normalizing and analyzing CTI from heterogeneous sources are major challenges for CTI analysts. Successful defense against threats depends on automation and to make available CTI more useful. Big data analysis and advanced reasoning may be applied, but these rely on consistent and structured data. We propose the ACT data model to address these challenges.

1.1 Research Motivation

Threat intelligence is served in several formats and channels, and with a varying degree of structure. Having worked with threat intelligence and incident response we had a need for a data model that enabled automation and analysis of our available threat intelligence. Combining all available data in one place, allowing for different data sources to be combined and analyzed, will increase the analysis capability of an analyst and remove repetitive tasks.

We find import and export of CTI from a system to be trivial given the data is stored in a consistent and structured manner, covering all relevant data. How we model our data is hence the foundation for everything else.

A key requirement for automation and analysis is data quality. Data quality is both content and format. We cannot enforce quality of content, but we can enable an analyst to evaluate this. Format consistency can be enforced by a strict data model. This means that a computer knows where to find a certain data type in a data set, and that the data found in that place always is the same type of data. With flexibility within the schema of a data model, this requirement will not be met, removing the ability to automate consumption and analysis across different platforms.

Threat intelligence analysis traditionally requires a large amount of knowledge from the analyst. Adding knowledge into the data model will make the knowledge available to more analysts.

Threat intelligence depends on collaboration between a range of organizations and communities. Any tool or system used by collaborators should be openly available to the community without restrictions, which has been a key motivation for this project.

2 RELATED WORK

There are several attempts at structuring cyber threat intelligence (CTI). The motivations for the different approaches seems to differ and these influence the results.

Barnum et al suggested the Structured Threat Information Expression (STIX) (Barnum, 2012) in

2012. This was created with the motivation of *sharing* CTI, preferably as more than just data. STIX was intended as a data exchange format and not a suggestion for how to store the data. STIX was published in version 2.0 in 2017 (OASIS CTI TC, 2017) and argued to be the de facto standard for representing CTI (Sauerwein et al., 2017). At the same time, critics of STIX argue that the flexibility of STIX makes it less useful for automation. As there are different possibilities of expressing the same data and information in addition to a fair amount of data included in custom fields or as comments using English prose (Polzunov and Abraham, 2019), automating consumption and further analysis is difficult.

The Malware Information Sharing Platform (MISP) (Wagner et al., 2016) is a platform for rapid sharing of indicators of compromise and sightings of indicators. The MISP data model is under continuous development (MISP, 2019). The data contained within MISP platforms correspond well with the suggested data model in this paper, however the popularity of the platform and loose data governance has led, in our experience, to a decrease in the consistency of the data.

ATT&CK (MITRE, 2019) is a framework and knowledge base for describing adversary behavior through enumerating adversary groups, tactics, techniques and tools and the relationships between them. The knowledge base is maintained by MITRE, and it is published online. ATT&CK uses a data model with defined relationships for structuring their knowledge base.

The OpenCTI platform (ANSSI et al., 2019) was published in late spring 2019 and is a platform aiming at consuming, analyzing and sharing cyber threat intelligence. The OpenCTI platform is including STIX observables and STIX relationships in its data model. Grakn[1] is used to enable graph querying of the data and includes rule-based reasoning to infer new relationships. To the best of our understanding, OpenCTI is limited to the scope of STIX and thus limits the possibilities of consumption and analysis within the platform.

An ontology, in the field of computer science, is a formal description of concepts and how they are related to each other, often referred to as classes and properties. In turn, ontologies provide computational meaning to data by building relations to the logic in the ontology and thus enables us to use reasoning methods (such as induction or deduction) on our data in our knowledge base. While there are many implementations of knowledge bases and ontologies, the World Wide Web Consortium (W3C) chose a triple model for facts and calls this the Resource Description Framework (RDF). RDF also allows us to implement the RDFS schema language[2] and OWL[3], the ontology language which builds upon RDFS.

There are several ontologies built with the aim to structure security relevant data. They cover a range of data and motivations like data validation, transformation or logical reasoning. An overview of available ontologies may be found in (Mavroeidis and Bromander, 2017). To the extent of our knowledge, none of the available ontologies are suitable for solving our problem alone, however the UTIM[4] ontology is being developed in parallel with this model and it is hoped that one day data from the ACT model will be freely interchangeable with data modeled with UTIM.

The rest of the paper is structured as follows: First we describe the methodology of our work in Section 3. Then we explain the details of the data model, with argumentation for our choices in Section 4, which includes a graph representation of the data model. We discuss our findings in Section 5, and conclude in Section 6.

3 METHODOLOGY

We developed the model using an iterative process basing our design on the relevant threat intelligence data we had available and then testing and updating as needed.

The platform we have used to implement the data model for prototyping and testing has been developed using agile development principles. This is a good fit for our iterative process of data model development.

3.1 Limitations

While the data model is an ontology, it is not implemented in RDFS or OWL, but all content can be exported as triplets. Initial testing of implementing the data model using Protégé[5] has been done in order to find improvements, but the desired reasoning capabilities lead to the need for rule based reasoning, which can be performed on top of the proposed data model with other tools as well.

We need a strict data model to avoid bad data in the knowledge base. The proposed data model requires a certain amount of work to consume new sources of data because of this chosen strictness.

[1] https://grakn.ai/

[2] https://www.w3.org/TR/rdf-schema/

[3] https://www.w3.org/TR/owl2-overview/

[4] Unified Threat Intelligence Model. See: http://www.ti-semantics.com

[5] https://protege.stanford.edu

4 RESULTS

We have created a data model and implemented it using an Apache Cassandra[6] and Elasticsearch[7] backend. We have implemented an Apache TinkerPop[8] graph engine which enables graph querying with the use of the graph query language Gremlin[9].

Note that a graph view is not the same as a graph database. You can display any kind of data, even a flat text file, as a graph, but you cannot use graph queries unless you have a graph engine interfacing with your data.

An implementation of the data model can be found on GitHub[10] under the ISC license. An openly available instance of the same implementation can be found online[11].

We have divided the results section into three: The foundation of our data model and the discussion leading to it, the schema improvement due to additional data, and the choice of allowing placeholder objects.

4.1 Foundation: Objects and Facts

The foundation for our work has been a data model consisting of objects and facts. We can define different object types and different fact types. Thinking of graphs, objects are the vertices and facts are the edges. Objects can be described as nodes and facts may be described as relationships. In the following we use the terms objects and facts.

Figure 1: Objects and fact.

The specifications and restrictions to this model is given in the next sections.

4.1.1 Immutable Objects - Retraction of Facts

Objects are defined globally and are immutable. There are no properties linked to an object, everything you know about one object is stored as facts. A fact may connect to one or two objects. A fact is directed, and can be bidirectional.

Deleting a fact is also not possible, however a new fact can be added that retracts the old one. In this way we make sure nothing is deleted and we can prevent repudiation. This way, we also preserve history and check the history of the data set.

[6] http://cassandra.apache.org/

[7] https://www.elastic.co/

[8] http://tinkerpop.apache.org/

[9] https://tinkerpop.apache.org/gremlin.html

[10] https://github.com/mnemonic-no/act-platform

[11] https://act-eu1.mnemonic.no/

4.1.2 Time

Because facts cannot be deleted, we are able to traverse the available data back and forth in time. Using the available threat intelligence in an incident response setting, this is useful for two reasons:

Firstly, knowing exactly what we knew at a given point in time. In situations where a range of decisions are made within a time frame of months, it is useful to be able to turn back time in order to know what information were available at the time when the decision was made. When incorrect decisions have been made, the ability to go back in time and see what information was available at that time will provide the ability to learn from mistakes.

Secondly, knowing how a threat has evolved over time. To know what infrastructure, behavior and resources a given threat actor has used at different times is useful in order to separate threat actors from each other, to identify copycats or impersonation and in order to evaluate how advanced the threat actor is. A threat actor using novel techniques, but abandoning them when they become normal behavior may be considered more advanced than others.

4.2 Data Model

Based on our object/fact foundation, we have defined a set of object types and fact types that are relevant and necessary for our domain.

The initial selection of object types were done influenced by STIX (Barnum, 2012), the Detection Maturity Model (Stillions, 2014), the Diamond Model (Caltagirone et al., 2013), available Open Source Intelligence extracted with the use of Natural Language Processing (NLP) and our own experience.

Fact types were added as we found them useful, with an increasing attention to the semantics and the characteristics of each of them. As our use cases for querying the data expanded, we saw the usefulness of differentiating between fact types.

Figure 2 shows the complete data model schema as a graph. The diamond shapes represent the values of fact types connected to only one object type.

We have populated the data model with a range of sources. A list of openly available sources used so far may be found in Table 1. The data model has been developed and improved along with introduction of new data.

In the following we explain the background and reasoning for the choices we have made, and include results from importing different data sources.

Table 1: The sources influencing the data model.

Source	Relevant object types
ATT&CK	Tactics, Techniques, Tools, Threat Actor
VirusTotal	IPv4, IPv6, FQDN, URI, Content, Hash, Tool, ToolType, Query, Path, Scheme, Basename
Shadowserver ASN	IPv4, IPv4Network, ASN[12], Organization, Country
Passive DNS	IPv4, IPv6, FQDN
MISP Galaxies	Tool, Threat Actor, Sector
STIX vocabularies	Sector
Open Source Intelligence extracted with NLP	All

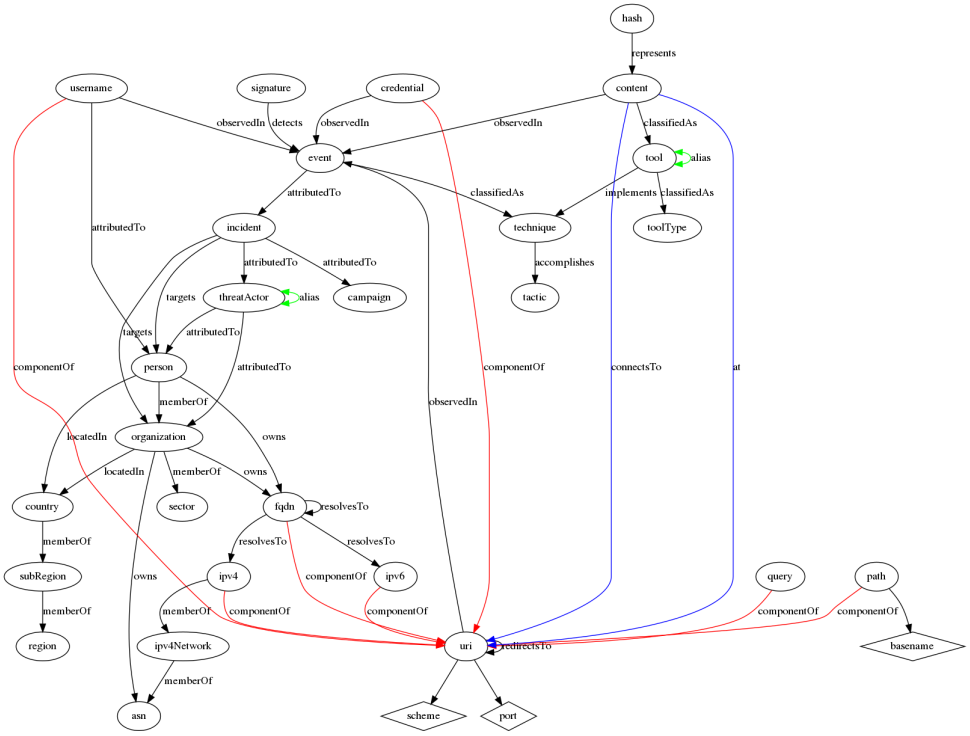

Figure 2: The complete data model represented as a graph.

4.2.1 Consistent Data

There are restrictions on which fact types can be used to link which object types as seen in Figure 2. These restrictions enforce data consistency by preventing different representations of the same data, which is a known problem with current attempts to model CTI. If we allow flexibility in how different data can be represented and introduce a range of users, then the data quality in terms of format is quickly reduced.

The granularity of the data model is intended to be aligned with pivot points normally used by analysts. This has been achieved by covering all sources in the currently available models and sources of CTI,

described in the introduction of this section.

We started with differentiation between *malware*, *tool* and *utility*, all instances of software. However, we saw that the definitions of the different groups varied in different sources and became difficult to maintain. This is consistent with the known problem of classifying malware, and we do not attempt to solve this in our data model. When the content of the CTI was not consistent we found that there was no value of using it at all. Therefore in our platform these concepts were all rolled up into *tool*, with the possibility of tagging them as *malware* or *utility* as appropriate.

4.2.2 Enrichment and Query/Analysis Across Sources

One of our first observations was that our graph ended up being a series of subgraphs, and we wanted to be able to connect them. The simple solution was enrichment. As we added more enrichment sources, the graph gradually became more and more interconnected, and we could find new connections between clusters of information that were originally separate.

Pivoting on an object is useful, as it lets you find related information and give you a more comprehensive context. One simple example is from DNS: start with a domain name, find all of the IP addresses that it has resolved to, and then find all other domain names that have resolved to those IP addresses.

Passive DNS (pDNS) data is a historic record of DNS lookup resolutions and is important for an investigation. From 2013 mnemonic has collected pDNS data. By 2017, when we had the initial version of the platform ready for data consumption, we had a TLP:White data set of approximately 100 GB of data. By analyzing super nodes in the data set, we have discovered new and unknown sinkholes. We tag known sinkholes with a fact connecting to the object in order to filter them out when traversing the graph further.

A more advanced solution was to use classifiers to bridge technical, tactical, operational and strategic threat intelligence. An example of this is using VirusTotal to bridge technical indicators to tactical information in MITRE ATT&CK. We extracted the malware family name from anti virus signatures and normalized it. We then normalized the Software entries from MITRE ATT&CK, e.g. “TrickBot” became “trickbot”. Automated enrichment with VirusTotal then connects file hashes and network infrastructure to the trickbot object, which is again linked to the tactical threat intelligence in ATT&CK.

We also observed that we could create uncommon pivot points, and our URI object type is an example of this. A URI object is just a UUID connecting different components to each other for a complete URI. Figure 2 shows the facts connecting to a URI in red and blue color. Given a URL, we split it into the host (domain/IP) part, the path and the query parameters. Pivoting on query parameters proved useful when tracking spam campaigns with specific phishing kits, as all of the other pivot points changed for each spam run, but the query parameter stayed the same.

4.2.3 Aliasing

Our data model allows for aliasing different names for the same object.

Instead of giving a threat actor a primary name, like in MISP Galaxy, we use *alias* as a fact type between threat actor names that are known or suggested to be the same. This may also be seen in Figure 2 with green color. Adding information on any threat actor’s name is then done by linking to the name given at the source. In this way, if an alias turns out to be wrong, you only need to retract that one alias, and the rest of your information is still correct.

The problem of different names for the same object is a common situation in CTI. Often, we find different providers of CTI gives a primary name for the object, and connect all information about this object to that name. For instance, if selecting “APT28” as the main name for a threat actor, and receive information about “Fancy Bear” (an alias for APT28), then such a solution will connect the information to “APT28”. This information can be wrong. If you at some point in the future decide that “Fancy Bear” is not an alias for “APT28”, then you would have a large manual task in correcting your data.

The *alias* fact type is used between threat actors and tools and might be applied to other object types in the future.

4.2.4 What is Content?

The concept of *content* is an example of where we need to be precise in order to enable automation. In the context of CTI, we handle not just files, but also stream segments, text strings and parts of content that has been found in memory. This is all “content”, but should not all be classified as files. Furthermore, even in the case of a file, we find that it is seen as unique based on more than one property. We argue that the file name, the actual content, and the location of the content together is what we refer to when we describe something as a unique file.

To illustrate the above we use the example of two files with the file system path `/etc/hosts` on two different Linux machines. In a given situation, the name and content may be the same, but they are still not

the same file due to the fact that they reside on different machines. In a different scenario you can find two files with the same name on the same machine, but with completely different content. In both cases, everyone agrees on the files being different from each other.

To be able to describe these things in a precise manner, and to identify similarities and identical objects, we saw the need for splitting them. The result was *content* linked to *uri* with the fact types as seen in Figure 2 with blue color. Basename (which includes the filename) is included within the URI.

The *at* fact found connecting a content object to a *uri* object is in the meaning of *seen at* and *downloaded from*. The general *at* was selected so as to not exclude any of the terms. The additional *connectsTo* fact represents a content which has been seen connecting to a *uri* and show the two very different scenarios where there is a link between the two object types. This is an example of the importance of semantics when handling CTI.

4.3 Placeholders to Preserve Information

In the ACT data model, you cannot link objects without a defined fact type between the object types. From adding new sources in various structures and formats, we found ourselves in need of adding more fact types based solely upon the information we wanted to consume. This resulted in a vast amount of fact types, and no consistency in representation of information. This is one of the most commonly mentioned weaknesses of the structure given in STIX, where there are several ways of representing the same CTI, resulting in problems digesting all information, especially without manual work and deduplication.

Looking for solutions we found the need for describing "things we know exist, but know little about". Blank nodes has been a solution for this problem in the field of ontologies (Hogan et al., 2014) and is part of the standardized W3C RDF Semantics (W3C, 2014). We introduced the same thought in our data model, by using what we called "placeholders". The idea is that the user may find information about the object in the future, and then replacing the placeholder with an actual object through a new fact. In this way, we were able to strictly define how the data are truly connected to each other, without worrying about having all data in a chain in order to consume it.

As an example, Figure 3 - 5 explains a typical scenario when working with CTI.

After implementing placeholders in our data model and restricting the fact types' possible connections, we found that adding and searching the data gave us an easy overview over what data is missing. This is a very interesting benefit for security analysts receiving or searching data on a relevant incident, both to know what data you do not have, but also to know what data others will need to be in possession of when sending you data. In evaluation of different CTI sources, this is a relevant analysis to perform.

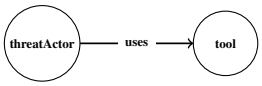

Figure 3: A typical piece of information received as CTI.

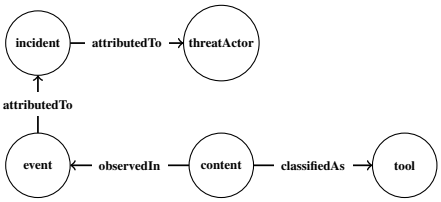

Figure 4: The information needed to give the statement in Figure 3.

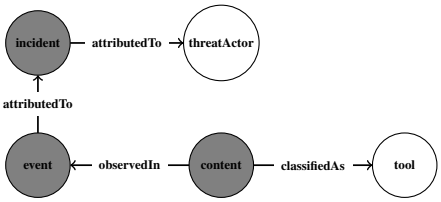

Figure 5: The need for placeholders: the information we know exist in gray, but is often not available for sharing.

5 DISCUSSION

The data model that we propose is strict: it restricts which relationships may be added to connect two objects, and it enforces that objects may not be added directly but through facts. The main benefit from this is a consistent data set which enables automation and improves data quality. It reduces the computational load of graph queries. It also provides for easier graph queries as there is no need to know the data you query so long as the user understands the data model. As an example there is a limited amount of traversals of the graph between *threat actor* and *technique*. Knowing this makes it trivial to find all connections between known *threat actors* and the *techniques* we know it has used without missing any available data. With

this, we argue that building the data model has transferred some of the advanced knowledge from CTI professionals into the model itself, which enables less skilled professionals to analyze the same data with consistent results.

In the course of developing this model we have discussed various solutions, implemented them and been surprised by some of the findings. The following discourse will look into the most relevant insights.

5.1 Data Validation

Large data sets often include some data that does not comply to the given specification. Adding data to our model, data outside specifications will be identified fast as they will fail upon consumption.

Bad CTI may lead to bad results when analyzing both as they may cause incorrect conclusions but also because they may ruin some of the other data. Most times errors exists due to mistakes entered at the source, because of the complexity of the subject matter or because multiple authors use different methods or terminologies.

We have found that the model allows for data validation. As an example, when querying the data from ATT&CK using our data model, we found that there actually was one technique called *Shared Webroot* without a link to any threat actors or any tools, which in threat intelligence is an interesting observation. Knowing that ATT&CK only includes data they have a reported observation of, means that this technique has been observed, but not described by openly available sources. This was obvious when we applied our model.

Adding MISP Galaxy for threat actors[13] where there is a range of users adding data with limited restrictions on data inclusion, we found that all threat actors were listed under a main name, with all information about them linking to this name. There are aliases listed underneath, but with no capability of reasoning on these aliases, the result is that a large portion of the threat actors actually are connected and seen as one. This meant that the value of the information was diluted as almost all information known about one threat actor was also stated to be valid for a large amount of other threat actors. This is an example of validation that may be used for evaluation of CTI sources, and it shows the importance of the chosen solution of aliasing as chosen in our model.

[13] https://github.com/MISP/misp-galaxy/blob/master/clusters/threat-actor.json

5.2 Evaluation of CTI Sources

When evaluating different sources of CTI, it is useful to evaluate the quality of the offered data. Our data model may be used for this purpose. Firstly, by adding context and knowledge to your data, which enables you to interpret the data you receive. Extensive aliasing, wrongful classifications or attributions may be easily found through such evaluation. Secondly, it helps finding data with errors, inconsistencies or bad formatting. The strictness of the data model excludes the possibility of importing data with errors, inconsistencies or bad formatting. When working to include new data sources these shortcomings will surface. Thirdly, to check what data is missing. When utilizing the data model with a given data set, if there is missing data it can be identified by identifying missing data in between data points. We can also find what object and fact types are used in that data to evaluate the range of CTI provided from the source.

5.3 Agreeing on Terms and Relationships

The terms and concepts within CTI are often referred to with different understanding. An example of this is *campaign* which often is used to describe standalone incidents and relevant threat actors in addition to the collection of incidents by the same threat actor targeting a given sector or geographical location. When connecting each concept to other concepts in a defined way, the data is given context, and with this additional meaning to a user. In this way we argue that ambiguity in terms and definitions will be reduced.

5.4 Differences in Object Types and Fact Types

There is a difference between objects that may be observed directly, and objects that are a result of human decision or analysis. Example of these types are *incident* and *tool* (not *content* or *hash*). The relationships going to and from these may also imply analysis, like *classifiedAs* and *attributedTo*. These facts are not a directly observable link. The trust we have in the source of these facts is thus more significant.

The differences in meaning of the different fact types shows the importance of semantics. There are object types which have multiple possible fact types connecting them, and where the semantics of the chosen fact type significantly differentiates.

An example of this is *content* $\xrightarrow{connectsTo}$ *URI* and *content* $\xrightarrow{at}$ *URI* as described in Section 4.2.4.

5.5 Sharing CTI

Newer publications suggest that still about 78% of shared CTI is unstructured (Sauerwein et al., 2019). Without any structure, we can only automate sharing of *data* as no relationships are present. With the choice of only adding information as facts (relationships) in ACT, we force all CTI to be stored with/as relationships. With this baseline we can automate sharing of triplets which is a significant improvement from sharing data and allows for sharing of graphs.

6 CONCLUSIONS AND FURTHER STUDY

We have proposed a strict data model based on objects and relationships, with the ability to represent available CTI. We have populated it with relevant data, and have identified new information through analysis enabled by the data model. The most prominent results from the data model is data validation, seamless enrichment, excellent analysis capabilities and flexibility of CTI ingest.

Future development of the data model will include hierarchical object types and fact types (using relationships borrowed from ontologies such as *subClassOf* and *subPropertyOf*) which will enable inheritance, more precision and reasoning.

In the implemention of our data model we allow external workers to access the content and add new facts. In this context we are exploring the use of an OWL-implemented version of our data model to infer new facts based on rule based reasoning using Semantic Web Rule Language (SWRL) (W3C, 2004).

ACKNOWLEDGEMENTS

This work was supported by the Research Council of Norway and mnemonic under the ACT and TOCSA projects.

The authors would like to thank the anonymous reviewers for valuable input that improved the paper.

REFERENCES

ANSSI, Luatix, and CERT-EU (2019). The OpenCTI Platform.

Barnum, S. (2012). Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX^{TM}). *MITRE Corporation*, 11.

Caltagirone, S., Pendergast, A., and Betz, C. (2013). The diamond model of intrusion analysis. Technical report, DTIC Document.

Hogan, A., Arenas, M., Mallea, A., and Polleres, A. (2014). Everything you always wanted to know about blank nodes. *Journal of Web Semantics*, 27:42–69.

Mavroeidis, V. and Bromander, S. (2017). Cyber threat intelligence model: an evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence. In *2017 European Intelligence and Security Informatics Conference (EISIC)*, pages 91–98. IEEE.

MISP (2019). The MISP platform.

MITRE (2019). Adversarial Tactics, Techniques and Common Knowledge (ATT&CK). https://attack.mitre.org/.

OASIS CTI TC (2017). Structured threat information expression (STIX^{TM}) 2.0. https://oasis-open.github.io/cti-documentation/.

Polzunov, S. and Abraham, J. (2019). EVALUATE OR DIE TRYING - A Methodology for Qualitative Evaluation of Cyber Threat Intelligence Feeds.

Sauerwein, C., Pekaric, I., Felderer, M., and Breu, R. (2019). An analysis and classification of public information security data sources used in research and practice. *Computers & Security*, 82:140–155.

Sauerwein, C., Sillaber, C., Mussmann, A., and Breu, R. (2017). Threat Intelligence Sharing Platforms: An Exploratory Study of Software Vendors and Research Perspectives.

Stillions, R. (2014). The DML Model. http://ryanstillions.blogspot.com/2014/04/the-dml-model_21.html.

W3C (2004). Semantic Web Rule Language.

W3C (2014). RDF 1.1 semantics.

Wagner, C., Dulaunoy, A., Wagener, G., and Iklody, A. (2016). Misp: The design and implementation of a collaborative threat intelligence sharing platform. In *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*, pages 49–56. ACM.