Identifying Adversaries' Signatures Using Knowledge Representations of Cyberattack Techniques on Cloud Infrastructure

Gilliam van der Merwe, Christian Muller, Wilhelm van der Merwe and Dewald Blaauw
Department of Information Sciences, Stellenbosch University, Cape Town, South Africa
gvdm.work@gmail.com
chrismullerpsa@gmail.com
dwilhelmvdm@gmail.com
dnblaauw@sun.ac.za

Abstract: Advanced Persistent Threats (APTs) have increased in parallel to growing cloud infrastructure and cloud Software-as-a-Service (SaaS) needs, exposing new vulnerabilities within the cloud environment. Moreover, APT groups are becoming more sophisticated and organised which needs to be addressed by the research community to enable faster response and more importantly, prevent threats within the domain. The MITRE ATT&CK Cloud framework offers one of the leading structured inventories within this context. Our research is to expose patterns and signatures of a select group of APT's on the MITRE Cloud Framework by using Formal Concept Analysis (FCA) to construct a "lattice graph" and an ontology. The goal is to develop a better conceptualisation of the MITRE ATT&CK Cloud Matrix framework for cyber security experts to be able to proactively act upon adversary techniques. The MITRE ATT&CK framework was retrieved, cleaned, and pre-processed to construct the lattice and ontology using data cleaning methods, FCA tools such as Concept Explorer, and the Web Ontology Language (OWL), with additional symbolic reasoning and inference generation. This resulted in knowledge representations/graphs, which are highly efficient representations of this knowledge field. The underlying linkages between techniques and targets specific to those APTs are further exposed and enriched and presented visually and integrated into the ontology. The ontology gives formalisation to associations and implications between techniques, tactics, and APTs – enabling cyber security practitioners to forecast potential targets and techniques based on their scenario, but also to attribute certain technique patterns and signatures to individual APTs. Cyber security practitioners can query from this knowledge graph and formulate strategic proactive measures. From these findings, the applications and constraints of the APTs' cyber-attack techniques and their associated patterns were determined. The findings provide a guideline for future additional research in the field of AI knowledge representation in cybersecurity, as well as highlighting certain limitations in this field of research.

Keywords: Cloud Infrastructure, Cybersecurity, Knowledge Representation, Advanced Persistent Threats

1. Introduction

Identifying APTs through signatures derived from utilizing a knowledge representation of cyberattack techniques on cloud infrastructure uncovers previously unknown information which could in turn prevent losses and transcend the knowledge of the cyber security community. The extended use of existing information given on the MITRE ATT&CK cloud matrix framework aims to expose underlying patterns and signatures of techniques that are of a similar fashion and nature. The MITRE ATT&CK Cloud matrix framework, an underlying branch of the MITRE knowledge base, lays out several adversary techniques formulated from actual world-wide attacks on Cloud infrastructure. An APT is a clandestine cyber-attack found on a network with the goal to exfiltrate and/or damage information. These attacks can be imposed by small to large groups, in some cases covert government -backed entities. The research team uses a Knowledge Representation in AI to highlight these underlying patterns which become apparent through this form of representation. The representation will pay particular attention to include correlations within the techniques used by six of these malicious groups exclusively, referred to as Advanced Persistent Threats (APT's), as to display the extensive potential value this research holds but without the burden of substantial complexity. Although current methods exist to detect APTs through their signatures, most of these methods are black box methods. This research uses KR to gain a deeper understanding of the relationships between attack techniques as well as threat actors (APTs) that make use of a certain set of techniques. By mapping the APTs in a KR the study hopes to predict what the next stage within a sequence of attacks will be. More specifically, the FCA implication and association sets from the KR will be utilised to generate a glass box model of these relationships in order to formulate a visual interpretation in the form of an ontology which enforces advanced comprehension, and in turn could allow for a faster response or proactive defence against these attacks.

In comment to the remainder of this paper, which structure reflects as follows: a presentation of background information which provides context surrounding Cloud Infrastructure and Cloud SaaS, MITRE ATT&CK Cloud matrix framework, and the formulation of an ontology enabled through FCA. Followed by an in -depth gradual

discussion of the data exploration and analysis methodology used to extract the sought -after and meaningful insights described within the outset of this research paper. In closing, research results and a description of possible future research is presented.

2. Background and Related works

2.1 Cloud Infrastructure and cloud SaaS

Cloud Infrastructure is a term encompassing the different components that are required for cloud computing. To be able to host varying cloud applications and services, specific infrastructure is required. Cloud infrastructure works by way of virtualisation. This is the process by which the processing capabilities of hardware are utilised for cloud computing, by means of software that allocates resources to environments that users can access anytime and scale to what the users require. (Red Hat, 2021)

2.2 MITRE ATT&CK framework

The MITRE ATT&CK framework is a publicly accessible trusted and standardised framework that provides a knowledge base of information on cyberattack tactics and techniques. These tactics and techniques are based on actual cyberattack examples. This knowledge base also contains information on the Advanced Persistent Threats (APTs) that utilise these tactics and techniques. This knowledge base serves as a basis for the development of cyber threat models and methodologies in various sectors both public and private (Strom, et al., 2020).

The MITRE ATT&CK Cloud Matrix is a more specialised alternative to the MITRE ATT&CK Enterprise framework as it focuses on adversary techniques in a Cloud infrastructure context.

2.3 Formal Concept Analysis

FCA is, according to Priss (2006), a method for data analysis, knowledge representation, and information management. In short, it allows us to hierarchically group objects according to their common attributes by using the mathematical theory of concepts and conceptual hierarchies, which forms the basis for FCA (Ganter & Obiedkov, 2016). The mathematical theory for FCA is beyond the scope of this paper but can be viewed in (Ganter & Obiedkov, 2016). Furthermore, FCA is based on the theory that a concept consists of two parts namely, its extension, which includes all the objects belonging to the concept, and its intension which consists of all attributes shared by these objects (Obitko, Snasel, & Smid, 2004). Additionally, formal concept lattices enable the visualisation of relationships between concepts, which in turn allows for a more intuitive manner to interact with the data (Kuznetsov & Watson, 2017).

In terms of FCA tools, various implementations exist (Priss, 2006), however, ConExp and Lattice Miner was best suited in relation to the objectives of this paper. Lattice Miner is an open-source Java tool that allows representation and manipulation of input data, lattices, implication sets, and association rules (Bertet, Borchmann, Cellier, & Ferre, 2007). ConExp works in a similar fashion, however, the tool works better for attribute exploration (Priss, 2006).

Association rules allow relationships within the concept lattices to be categorised depending on their confidence and support values. Support signifies the frequency of a relationship occurring, whereas confidence represents the probability of relationships to co-occur (Pasquier, 2000).

2.4 Ontology Engineering

Semantic technologies allow computers to interpret data in order to derive meaning thereof, this is made possible using formal semantics by defining concepts and how they relate. To allow computers to interpret a knowledge base an ontology is utilised to represent the said knowledge base (Gronberg, 2019). More specifically, ontologies model and define a domains entities, classes, properties and relationships (Gruber, 2009).

By using an ontology, knowledge exchange is facilitated by implementing a shared vocabulary. Moreover, ontologies can be used in conjunction with inference engines to enable reasoning, it is also used for user interface, database, and application components (Gronberg, 2019). The Protégé ontology editor is an open-source tool that enables ontologies to be built with all the necessary implementations and it conforms to the OWL 2 Web Ontology Language and RDF standards from the World Wide Web Consortium (Mussen, 2015). Protégé also allows the use of inference engines to enable reasoning.

2.5 Related Work

Recent research on the application of Knowledge Representation, and more specifically Formal Concept Analysis (FCA), on the MITRE ATT&CK Framework showed promising results for the Cyber Security domain. Watson & Watson (2020) explored the use of FCA as a means for exploratory data science focusing on the relationship between adversary attack techniques and procedures of the mobile/IOS device subset in the ATT&CK framework. This approach enabled them to not only detect which components are shared between threats but also the possible identification of attribution signatures, prediction of new threats and enabling security practitioners to take precautionary measures.

Similarly, Singels, et al. (2020), applied FCA to the Industrial Control Systems (ICS) ATT&CK framework to enable 'earlier and strategic detection of cyberthreats'. They went further by encoding the FCA results to a formal ontology, in this case using the Web Ontology Language (OWL), allowing machine interpretation and semantic web reasoning. Through this, security practitioners are able to query the FCA results facilitating the discovery of knowledge.

3. Problem statement and research objectives

Although current methods exist to detect APTs through their signatures, most of these methods are black box methods. This research aims to use KR to gain a deeper understanding of the relationships between attack techniques as well as threat actors (APTs) that make use of a certain set of techniques. By mapping the APTs in a KR the study hopes to predict what the next stage within a sequence of attacks will be. More specifically, the FCA implication and association sets from the KR will be utilized to generate a glass box model of these relationships to see if a visual interpretation could aid in a better understanding, and in turn allow for a faster response or proactive defence against these attacks.

Questions that this research aims to address:

1. Are Knowledge Representations an accurate and efficient method of illustrating cyber-attack techniques?
2. What implications can be drawn from the knowledge representations of cyberattack techniques?
3. What insight can be gained from the researched cyberattack toolkits?
4. Is it possible to derive (*useful and identifying)* patterns and signatures of APTs from mapped data of cyber-attack techniques?

4. Methodology

4.1 Data pre-processing

First and foremost, raw data was extracted from the MITRE ATT&CK framework, specific to the groups and techniques in the Cloud Matrix. The dataset itself consisted of 86 groups and 53 techniques (sub-techniques included).

4.2 Data Mapping

For our procured data to be compatible with our FCA tools, the raw datasets were cleaned and processed using the R programming language, which lead to a formal context being exported as a .csv file for use within ConExp and .cex file in Lattice Miner. Lattice Miner and ConExp was utilised as our conceptual analysis tools with APT's/groups transposed as attributes and techniques as objects. Following the creation of the hierarchical structure of APT's and their applicable techniques it became apparent that mapping all 86 groups with applicable techniques would not be feasible within the scope of this particular research endeavour. In light of this, the approach to use only six of the most prominent APT's and the techniques associated with these six individual groups was deemed more effective. The approach enabled a more detailed view and underlying relationships could be showcased under a greater lens of focus.

4.3 Data Exploration

Viewing the relationships within Lattice Miner and ConExp exposed association and implication sets categorised and ranked according to their support and confidence levels. Low confidence is categorised as a percentage between 0.10 - 0.45. Medium confidence is seen as a percentage between 0.46 – 0.74 and high confidence, between 0.75 - 1.00. Antecedents and consequents of techniques and groups were found and noted. The level of support indicates the frequency of instances, and confidence the number of antecedents –consequent relationships occurring in each particular group and technique.

The Protégé ontology editor was used to substantiate these findings and showcase inferences within the data, in the form of an ontology. Under the main domain of the ontology concepts were initialised as follows: “APT”, “Proposition”, “Tactic” and “Technique”.

“APT” refers to the malicious entities implementing adversarial tactics. “Tactic” is the ultimate strategy which APT’s use to implement techniques. The six APT’s used are “Fox Kitten”, “Operation Wocao”, “Sandworm Team”, “APT32”, “APT41” and “Turla”. Each APT employs a technique through “techniqueEmployedBy”. “Technique” refers to the specific attack methods categorised under each tactic. Some techniques have sub -techniques indicated through “isSubTechniqueOf”. Each “Tactic” contains a certain number of techniques through “containsTechnique”. “Collection” “CredentialAccess”, “DefenceEvasion”, “Discovery”, “Execution”, “Exfiltration”, “Impact”, “InitialAccess”, “LateralMovement”, “Persistance” and “PrivilegeEscalation”. “Proposition” represents the implementation which constitutes the relationship between the antecedent and consequent. “Proposition” contains six instances namely, “Adversary1”, “Adversary2”, “Adversary3”, “AR1”, “AR2” and “AR3”. Each proposition relationship includes the support and confidence level bound to it in the form of data property assertions, depicted as “hasSupport” and “hasConfidence”, followed by an assigned level ranging from “Low”, “Medium”, and “High”.

The given ontology, labels groups (APT’s) and techniques in a fashion which identically reflects its description on the MITRE ATT&CK framework. Further specification surrounding object properties include “asserts”, “hasProbability”, “implies”, “isAntecedentIn” and “isConsequentIn”. These object properties are associative entities defining linkages between the concepts/instances within classes.

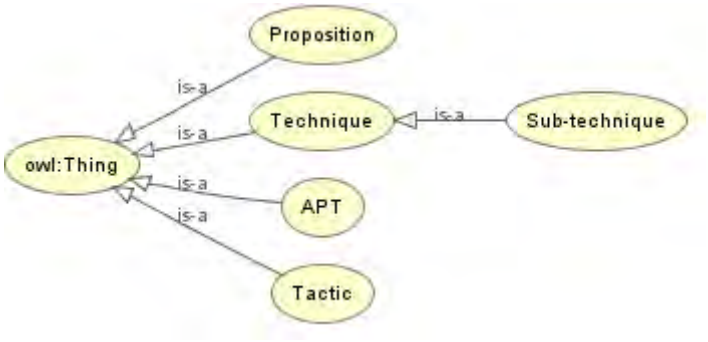

Figure 1: Diagram illustrating the Protégé OWL ontology

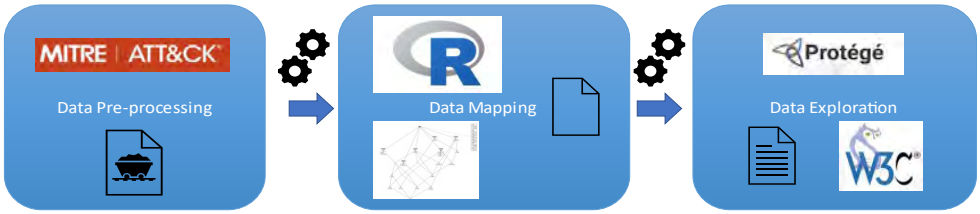

Figure 2: Diagram illustrating the research methodology

5. Data Analysis

5.1 Lattice Interpretation

Utilizing cyberattack techniques and APT groups from the MITRE ATT&CK Cloud matrix framework, a formal context, that focuses on the cyberthreats to cloud infrastructure and architecture was created. In this formal context the cyberattack techniques are objects, and the APT groups are attributes. In both ConExp and Lattice Miner the nodes of the lattice were labelled with names of the techniques used and the APT groups that used these techniques. In both figures (Figure 3 and Figure 4) a sample of six APT groups were chosen for an in-depth analysis of the formal context. In Figure 4 there are depicted concentrations of implications in the upper right region. The concentration of these implications indicates insightful connections on how APT groups and the techniques that they use relate to one another. From these connections, a more sophisticated comprehension

of the way in which APT groups utilize cyberattack techniques on cloud infrastructure and architecture is obtained. Considering the upward forming edges between the different nodes, it stands to reason that more information can be derived about which techniques are used by specific APT groups. This allows for more efficient prediction and mitigation of cyberattacks. Therefore, the concept lattices provide a visualization that allows for greater insight gained solely from the MITRE ATT&CK Cloud matrix.

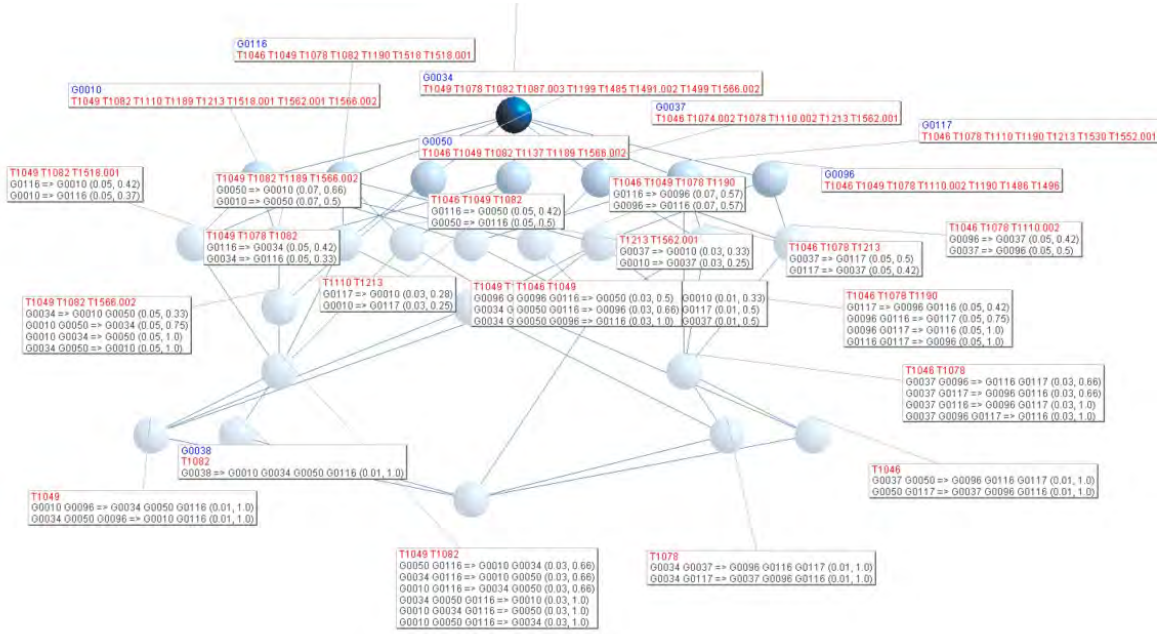

Figure 3: Lattice Miner lattice showing the connection between APT groups and cyberattack techniques as determined by the MITRE ATT&CK Cloud matrix. It illustrates which APT groups make use of which techniques and which APT groups have techniques in common with other APT groups

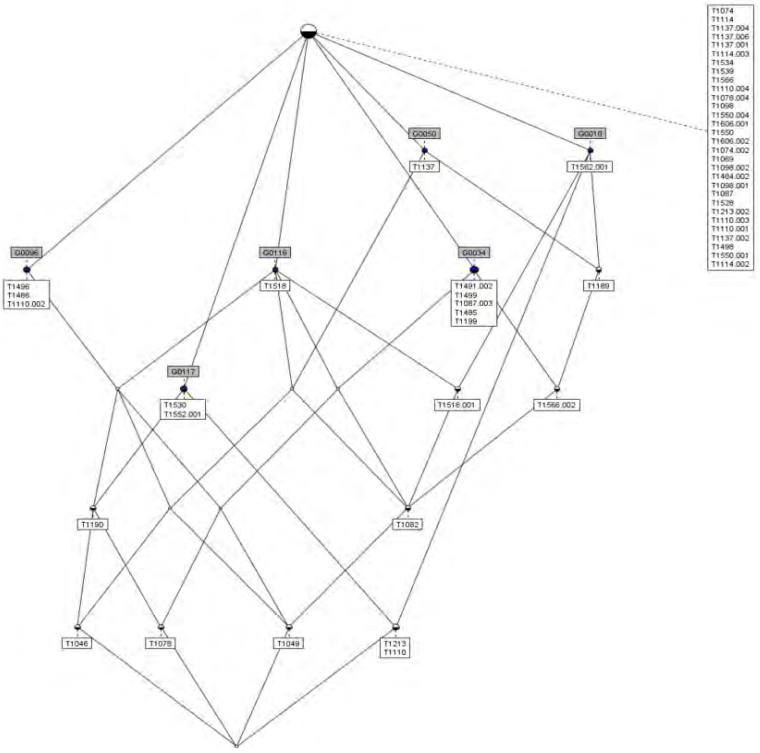

Figure 4: ConExp lattice showing the connection between APT groups and cyberattack techniques as determined by the MITRE ATT&CK Cloud matrix. It illustrates which APT groups make use of which techniques and which APT groups have techniques in common with other APT groups

5.2 Implications and associations

It is imperative to evaluate the implications and associations between the APT groups in the lattices presented above. A relationship can be identified between the APT groups G0096(APT41), G0087(APT39), G0116(Operation Wacao) and G0117(Fox Kitten). The implication *G0096, G0117 → G0087, G0116* with three objects supporting the implication. In reference to the MITRE ATT&CK framework, the implication can be read as *APT41, Fox Kitten → APT39, Operation Wacao*. A deeper analysis indicates that all three groups made use of the technique T1190 (Exploit Public-Facing Application). This technique is implemented by APT groups exploiting vulnerabilities in internet-exposed computers or programs through the utilization software, commands or data in order to cause the application to operate not as intended. G0087 and G0117 are suspected to be based in Iran and G096 and G0116 are suspected to be based in China.

In the utilization of implications provided by concept lattices, cyberattacks can be identified swiftly and efficiently as well as being able to determine the next steps or techniques that will be used in a cyberattack and aid in creating mitigations to prevent future cyberattacks.

The confidence level of association rules provides us with the probability that the implications will hold. The Lattice Miner solvers required a minimum support and minimum confidence value which were set as 0.0% and 100% respectively. Since there is a large number of associations, only a select number are highlighted. The different associations are grouped according to their confidence level percentage. The solvers show that the implication *G0096, G0117 → G0116* has a confidence level of 1.0 (100%), which indicates that there is a high probability that the techniques used by G0096 and G0117, will also be used by G0116. This illustrates that a greater focus would need to be placed on associations that have high probability values.

5.3 Ontology Queries and inferences

It is necessary to showcase the ontology utilization which provides the means for identifying APT's signatures or patterns in their cyberattacks on cloud infrastructure using knowledge representation. Example queries to illustrate the above mentioned will now be provided and discussed.

A security analyst suspects that a system has been attacked. The analyst determines the cyberattack technique used was T1078 (Valid Accounts). The analyst wants to determine which APT groups have utilized the same technique in order to identify the attacker. The analyst makes use of the ontology and writes the following query in order to achieve that:

employsTechnique some {T1078}

A security analyst wants to determine which tactic a technique falls under in order to better understand which phase an attack will undergo next. The query below indicates that the technique (T1078) is utilized in the tactics: initial access, persistence, privilege escalation and defence evasion.

containsTechnique some {T1078}

A security analyst can determine the APT group that utilizes techniques in implications with high confidence levels.

(employsTechnique some Technique) and isConsequentIn some (hasConfidence value "High")

A security analyst can determine which APT groups are likely to employ the same or similar techniques from querying implications and associations. AdversaryAR1 refers to the implication *G0096, G0117 → G0116.*

isConsequentIn some {AdversaryAR1}

A security analyst is able to derive a list of cyberattack techniques from APT groups and from these techniques, a mitigation strategy can be designed and implemented.

techniqueEmployedBy some {G0096} or {G0117} or {G0116} and Technique

6. Results and Discussion

The MITRE ATT&CK Cloud Matrix has information on various cloud platforms such as Azure AD, Office 365, Google Workspace, SaaS, IaaS (MITRE ATT&CK, 2021), however the matrix does not contain all the information on cloud platforms as there are new cyberattacks daily with some going unnoticed and not recorded into the framework. By making use of Description Logic (DL) queries one is able to analyse relationships that go beyond just implications and associations between APT groups and cyberattack techniques. Taking into account support and confidence levels of associations, it enables for the identification of high priority cyberattacks. This is done through the utilisation of the ontology and formal context analysis results derived from the Cloud Matrix data. The concept lattice provides analysts with the necessary information on cyberattack technique concentrations in respect APT groups, allowing for analysts to better deal with the cyberattacks. The ontology is used to complement the concept lattice for better understanding and insight into how cyberattacks occur and proceed. Considering the fact that the ontology utilises open world reasoning as opposed to the FCA's closed reasoning, there is risk that unforeseen factors affect the ontological queries. Therefore, it should be noted that APT groups, techniques, implications and their associated support and confidence levels change with time. However, the usefulness of the concept lattice and ontology is still prevalent enough to make use of them in order to identify APT groups' signatures or patterns in their attacks.

7. Conclusion and Future Research

By supplementing the MITRE ATT&CK Cloud Matrix with FCA and implementing an ontological approach, the patterns and relationships between APTs and their techniques implemented could be analysed with greater understanding. This provides cybersecurity stakeholders a much more intuitive interaction with the ATT&CK Cloud Matrix, enabling faster response times and reduced impact as a result of cyberattacks.

The information and modus operandi depicted within this body of work contains foundational value as it could, without a doubt, be built upon in the sense to include more relevant and up to date information but also in the context of its' sophistication as an aid for cyber security practitioners. Further research should be conducted into expanding the body of research concerned with utilising AI knowledge representations in order to ascertain patterns or signatures of cyberattack techniques in order to identify APTs. By further refining this research and coupling it with better programmatic testing, new knowledge and insight can be achieved.

References

Bertet, K., Borchmann, D., Cellier, P., & Ferre, S. (2007). Lattice Miner - A Formal Concept Analysis Tool. Rennes: Supplementary Proceedings of ICFC.

Ganter, B., & Obiedkov, S. (2016). *Conceptual Exploration.* Berlin: Springer.

Gronberg, M. (2019). *An Ontology for Cyber Threat Intelligence.* Ph.D. thesis, University of Oslo.

Gruber, T. R. (2009). Ontology. In L. Liu, & M. T. Ozsu (Eds.), *Encyclopedia of Database Systems.* (pp. 1963-1965). Springer US.

Kuznetsov, S., & Watson, B. (2017). Proceedings of the International Workshop on Formal Concept Analysis for Knowledge Discovery (FCA4KD). *1921.* CEUR Workshop Proceedings.

MITRE ATT&CK. (2021). *Cloud Matrix.* Retrieved June 8, 2021, from https://attack.mitre.org/matrices/enterprise/cloud/

Mussen, M. (2015, June). The Protégé project: A look back and a look forward. *AI Matters, 1*(4), 4-12.

Obitko, M., Snasel, V., & Smid, J. (2004). Ontology Design with Formal Concept Analysis. 111-119.

Pasquier, N. (2000). Mining association rules using formal concept analysis. *1867*(12), pp. 259-264.

Priss, U. (2006). Formal Concept Analysis in Information Science. *Annual Review of Information Science and Technology, 40*, 521-543.

Red Hat. (2021). *What is cloud infrastructure?* Retrieved June 7, 2021, from https://www.redhat.com/en/topics/cloud-computing/what-is-cloud-infrastructure

Singels, L., Biebuyck, C., & Malukele, L. (2020). A Formal Concept Analysis Driven Ontology for ICS Cyberthreats. (pp. 247-263). Hatfield: SACAIR 2020 Organising Committee.

Strom, B., Applebaum, A., Miller, D., Nickles, K., Pennington, A., & Thomas, C. (2020, March). *MITRE ATT&CK: Design and Philosophy.* Retrieved June 2, 2021, from MITRE ATT&CK: https://attack.mitre.org/docs/ATTACK_Design_and_Philosophy_March_2020.pdf

Watson, L., & Watson, B. (2020). Exploratory data science on ATT&CK. *KM Conference 2021.* Leipzig: International Institute for Applied Knowledge Management.