Defender-centric Conceptual Cyber Exposure Ontology for Adaptive Cyber Risk Assessment

Lamine Aouad[1] and Muhammad Rizwan Asghar[2]
[1]*Tenable Network Security, U.S.A.*
[2]*The University of Auckland, New Zealand*

Keywords: Cyber Exposure, Vulnerability Management, Ontology, Cyber Risk.

Abstract: A major gap in cybersecurity studies, especially as it relates to cyber risk, is the lack of comprehensive formal knowledge representation, and often a limited view, mainly based on abstract security concepts with limited context. Additionally, much of the focus is on the attack and the attacker, and a more complete view of risk assessment has been inhibited by the lack of knowledge from the defender landscape, especially in the matter of the impact and performance of compensating controls. In this study, we will start by defining a conceptual ontology that integrates concepts that model all of cybersecurity entities. We will then present an adaptive risk reasoning approach with a particular focus on defender activities. The main purpose is to provide a more complete view, from the defender perspective, that bridges the gap between risk assessment theories and practical cybersecurity operations in real-world deployments.

1 INTRODUCTION

In today's interconnected digital world, one of the major challenges facing organisations is securing their assets from potential cyber attacks. While organisations would like to mitigate cybersecurity risks, the growing number of vulnerabilities, the frenetic expansion of the attack surface, the rise in sophistication of attacks (Symantec, 2019), in addition to multiple other technological and human factors, make reducing the uncertainty around cyber risk a very challenging task.

In order to efficiently assess cyber risk and to improve the overall security posture, the cybersecurity ecosystem needs more visibility into the interactions of complex systems, and defensive controls and cyber operations driven by threat analysis and context of use. A necessary first step towards achieving this is identifying and organising security concepts and their relationships into a formal knowledge representation; hence, there is the need for an adequate ontological representation that models all relevant entities. In addition, appropriate selection of defences and measuring their performance is essential to enable adaptation and improvements. This will lead to ameliorate the level of rigour and automation associated with assessing the effectiveness of cyber defences, which has been essentially a manual task, prone to errors.

2 BACKGROUND AND LITERATURE REVIEW

In philosophy, ontology refers to the basic description of "things" in the world. From the Information Technology (IT) perspective – specifically Artificial Intelligence (AI) – Guarino (Guarino, 1998) defines an ontology as "an *engineering artefact*, constituted by a specific *vocabulary* used to describe a certain reality. In simple terms, an ontology is a set of concepts and their relationships and its importance is being recognised in a variety of research fields including information and knowledge engineering and representation. Ontologies promote the creation of a unique standard to represent entities and concepts within a specific knowledge domain. It is also a useful tool for reasoning about relationships between its entities, *e.g.,* vulnerabilities and compensating controls, and help answer questions related to those relationships, *e.g.,* which compensating controls would protect against exposure of data, or a given vulnerability.

In order to support cybersecurity ontologies, there is already a large body of work of definitions, taxonomies, and enumerations related to core terms and concepts. These include vulnerabilities and weaknesses, threats and exploitation, malware and its types, attack patterns, *etc.* One of the most active organisations in this field, MITRE, maintains dictionar-

Aouad, L. and Asghar, M.
Defender-centric Conceptual Cyber Exposure Ontology for Adaptive Cyber Risk Assessment.
DOI: 10.5220/0009826205800586
In *Proceedings of the 17th International Joint Conference on e-Business and Telecommunications (ICETE 2020) - SECRYPT*, pages 580-586
ISBN: 978-989-758-446-6

ies around vulnerabilities and weaknesses (MITRE, 2019a) (MITRE, 2019b), a knowledge bases of adversary tactics and techniques (ATT&CK) (MITRE, 2020a), and a similar classification of attack patterns, more focused on application security (CAPEC) (MITRE, 2020b). Other organisations, including NIST (NIST, 2020b) and NVD (NVD, 2020), also maintain repositories of standards, and threat and vulnerability management data. There are also initiatives around standardising the description and sharing of cyber threat information and sources, such as STIX (OASIS, 2020a) and TAXII (OASIS, 2020b).

These taxonomies focus on information about system observables, such as vulnerabilities, security events, and Indicators of Compromise (IoCs), for the purpose of presenting and sharing specific information that provide enhanced vulnerability and threat detection and protection. In contrast, a conceptual ontology is expressed at a higher level of abstraction and focuses on design-level assessments of a cybersecurity operation. This is similar to the use of abstractions to help thinking about risk when doing threat modelling (Shostack, 2014). It is necessary because IT systems and deployments are widely different to each other and abstracting away details will provide a better look at the big picture. Additionally, the taxonomies focus on independent system and threat representation, but provide no means for representing controls and defence capabilities.

In the literature, few cybersecurity and Vulnerability Management (VM) ontologies have been proposed. Blanco *et al.* were among the first to conduct a review study of security ontologies (Blanco et al., 2008). Following studies in (Fenz and Ekelhart, 2009) and (Mavroeidis and Bromander, 2017) agree with the observations made by Blanco *et al.* that a general cybersecurity ontology is yet to be defined by the community due to a number of reasons, including the lack of structure in the knowledge coming from domain experts for advanced reasoning. A study by Syed *et al.* (Syed and Zhong, 2018) has integrated an intelligence element to their VM ontology, but it has only considered Twitter data. Also, to the best of our knowledge, none of the existing ontologies includes context as a concept, which is the main driver of trade-offs when it comes to the interplay between security requirements, and vulnerabilities and threats, and their mitigation.

In (Syed et al., 2016), the authors introduced a Unified Cybersecurity Ontology (UCO) as an extension to a previously developed Intrusion Detection System ontology. They built UCO by semantically linking various aspects of STIX, CVE, CCE (Martin, 2008), CVSS, CAPEC, STUCCO (Iannacone et al., 2015), and the kill chain. The STUCCO ontology itself had initially incorporated data from 13 different sources. In (Wang and Guo, 2009), the Ontology for Vulnerability Management and Analysis (OVM) was also built on existing standards and taxonomies such as CVE, CWE, CPE, CVSS, and CAPEC.

For aspects more related to threat modelling, the Attack Surface Reasoning (ASR) ontologies, proposed in (Fusun et al., 2016), gives a cyber defender the possibility to explore trade-offs between cost and security when deciding on the composition of their cyber defence. Ontologies created include those of attacks, systems, defences, missions, and metrics.

The huge diversity of the theory and practice of cybersecurity also accounts for the large variety of underlying concepts and principles used in previous studies, and is the main reason (and cause) of the continued effort in this area. In (Syed et al., 2016), for instance, all concepts link back to the *attack*, and should be more referred to as a cyber attack ontology. We strongly believe that a common language, which abstracts out low-level system observables into a set of basic concepts is essential in developing a shared understanding of the cyber security ecosystem, and further expand it into an ontology. This was one of our primary motivations to propose a more comprehensive foundational conceptual representation.

On the other hand, in the area of measuring and managing information risk, various frameworks have been proposed, some of which are ontology-based. The FAIR institute (FAIR, 2020), for instance, proposes standards, best practices, and also a risk ontology (or rather a taxonomy). It starts from the top-level concept *risk* and steps down to key factors that derive risk in FAIR, including loss frequency, magnitude, and exposure, and vulnerability and threat. The modelling is centered around quantifying the threat factors for the risk associated with a given scenario, while other important concepts such as controls, policies, assets, or intelligence are not explicitly defined (Freund and Jones, 2014). Note that most risk management frameworks rely on flat terminologies and lack the richness and flexibility that can be provided by relationships in an ontological representation.

There are multiple other frameworks, proposed by standard bodies and researchers, to understand, measure, and assess risk. These include CIRA (Conflicting Incentives Risk Analysis) (Rajbhandari, 2013; Snekkenes, 2013), in which risks are modelled in terms of conflicting incentives between risk owner and other stakeholders. CIRA does not directly conduct vulnerability and control identification, but threats and stakeholders are at the core of the method. Another risk framework is CORAS (Den Braber et al.,

2006; Lund et al., 2011), which is based on modelling threat scenarios related to assets. CORAS does not provide any steps for identifying and assessing existing controls and also lacks advanced threat intelligence activities for risk estimation. A number of other risk analysis and management tools have been proposed, including CRAMM (CCTA Risk Analysis and Management Method) (Yazar, 2002), OCTAVE (Operationally Critical Threat and Vulnerability Evaluation) (Caralli et al., 2007), ISO/IEC 27005 (ISO, 2011), and NIST SP 800-30 (Blank, 2011). For a more detailed description and comparison of different frameworks, we refer the interested reader to (Wangen, 2017).

3 CYBER EXPOSURE ONTOLOGY

We propose a cyber exposure ontology that builds upon the classic components of risk assessment – vulnerability, compensating controls, threat, and asset – incorporating three additional core concepts including intelligence, context, and defence policy. Figure 1 shows these core concepts and their relationships. Attributes of the threat and context concepts are presented as examples in Figures 2 and 3. Cybersecurity operations (including cyber defence workflows and procedures) will operate over this set of concepts that will describe all aspects related to the exposure of assets under assessment, controls and defences, and the detailed capabilities and context around threats and adversaries.

For the threat modelling and classification aspect, the MITRE taxonomies can be used. Compensating controls can then be categorised using the ATT&CK threat matrix and CAPEC. If a control does not cover or contribute to prevent a threat, at least one of the CIA (Confidentiality, Integrity, and Availability) triad requirements will be violated. Alternative models, including other factors besides the three facets of the CIA triad, can also be considered depending on the type of events or scenarios.

The proposed additional concepts have a capital enrichment role that will allow adaptation and better defence ability. The *intelligence* concept will be informed by the intelligence source attributes, *e.g.,* technical blogs, reports, academic or analyst notes, blacklists, *etc.* It presents the high-level summary of intelligence. The main attributes are related to the source of the data, source type (machine- or human-driven), content type (*e.g.,* availability and maturity of exploits, attack vectors), and high-level technical details (*e.g.,* known campaigns or actors using the vulnerability, IoCs).

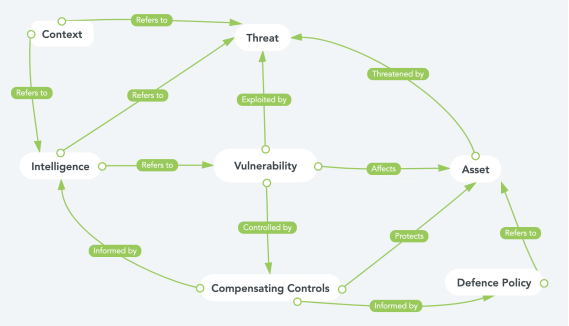

Figure 1: Core concepts of the proposed cybersecurity ontology.

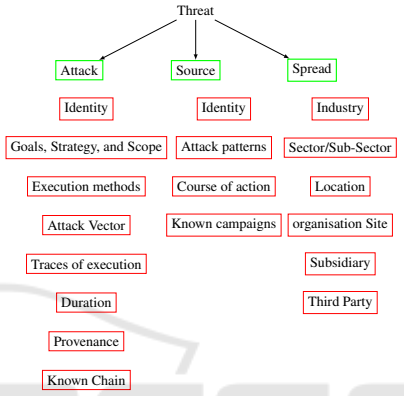

Figure 2: Threat.

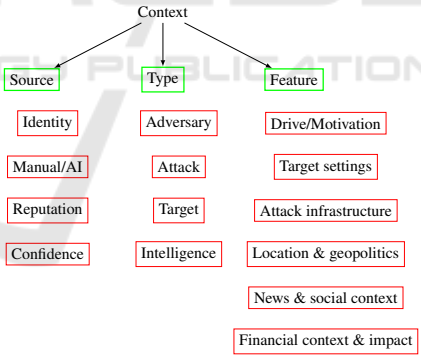

Figure 3: Context.

On the other hand, the *context* concept will be informed by contextualised information around attacks and targets (assets) of interest. This especially relates to threat and intelligence attributes, and influences their interpretation. The main attributes there are related to the source of the data. That is, what is known about the source, *e.g.,* in terms of reputation and known False Positives (FPs), and content type. It also includes details around the attacks or threat intelligence, *e.g.,* the root cause of an attack, target and attack settings, what data has triggered the intelligence or incident (explainability of an alert), or any other

details related to the incident including, for instance, news, course of action, and kill chain details and progression.

As to the *defence policy* concept, it represents security goals, and defence and risk management strategies and implementations, *e.g.,* the frequency of assessment and updates, the presence or not of offensive testing, or training programs. It has a level and type attributes and is informed by existing standards, guidelines, planning strategies, and best practices to manage cybersecurity risks, *e.g.,* the NIST cybersecurity framework (NIST, 2020a).

4 ADAPTIVE RISK ASSESSMENT

Implementing security solutions requires understanding and visibility into both the resources (*processes - people - technology*) and the defences in place supporting those resources to provide necessary mission functionality. While most practitioners recognise the threat landscape as highly dynamic, very few pay attention to the intricacies of interactions between threats and the system. There is a need for iterative and adaptive reasoning, which comes from the fact that the threats, and assets/ entities themselves, are connected to each other in non-trivial ways, *i.e.* any operation on the system, or new threat observation, might trigger a control direction or rule that was not previously considered. This adaptation will eventually require a more comprehensive knowledge representation, hence the supporting ontological representation we propose above.

4.1 Analysis Process

We represent cyber risk by taking hypothetical *risk events* and identify elements that influence the likelihood and impact of an event into three different stages. Figure 4 summarises the three stages of an adaptive cyber risk process. The proposed three stages model allows to have a holistic and dynamic view, from prevention through response and recovery. This essentially means that all possible outcomes of an event and variation in attributes should be considered in risk evaluation. The objective is to detail risk events and activity from the defender perspective, where more weight is given to the role of cyber defences and other tools supporting a certain set of mission against cyber attacks. The three event stages are described as follows.

- **Pre Event.** Mainly focused on situation awareness of the organisation.

Figure 4: Adaptive Risk Assessment Framework.

- **During Event.** Enables risk modellers to understand the organisation protective mechanism (proactive) and response mechanism (reactive).
- **Post Event.** Once the security incident is contained, organisations have to address the damage done due to the event.

For the analysis process, deriving relevant risk events will primarily depend on how people and organisations think about risk, including their risk profile and tolerance. Risk events are inevitably seen differently under different circumstances. Supported by the concepts proposed in the cyber exposure ontology in Section 3, we can define a risk event as a function based on the following factors:

- Asset A_t: vulnerability exposure and criticality of an asset.
- Threat T_t: threat landscape, vulnerability intelligence, and context of the threat under assessment.
- Control C_t: mission type and setup of the control.
- Impact I_t: nature of effects over the risk event, whether technical or business.
- Residual risk tolerance RT_t: risk appetite and tolerance over the risk event given controls performance and capital in place.

where t defines the stage of a risk event, *i.e. pre, during, and post*.

$$Risk(E_t) = \sum_{t=pre}^{post} f(A_t, T_t, C_t, I_t, RT_t) \qquad (1)$$

Next step in the process of adaptive risk assessment is risk quantification of a particular risk event via assigning values to the parameters in Equation 1. The purpose of quantification is to assess the maturity level of

Table 1: Evaluation metrics over risk event configurations.

Controls Metrics	Security Metrics
Coverage - Total number of assets covered/monitored	Asset exposure (Public, Restricted, Private)
Performance (Fail—Partial—Pass)	Exploitation (None, PoC, Functional, High)
Mean time to affect	Cumulative effect, Daisy-chaining
Mean time to detect	Total number of attack vectors for known attack
Mean time to respond	Total number of other possible entry points

the organisation associated with the risk event, especially as it relates to the threat landscape and a given defence composition and layout. The state of various assets, threat, control, impact, and tolerance will determine the feasibility of threat success with respect to the evaluation factors including compensating controls and security metrics. This is primarily a deterministic approach. The maturity can then be deduced as a weighted sum, or an aggregate index over the factors quantification. However, one aim of this framework is to move away from current, predominately, number-based systems (with aggregated risk scores) where the rational of those numbers is often lost, making adaptation hard to achieve. Table 1 shows a list of possible compensating controls performance and security metrics used to characterise event factors.

Exposure, including exploitation status and criticality (for both vulnerabilities and assets) are informed by a vulnerability management solution (Tenable Network Security, 2020a; Tenable Network Security, 2020b). Finding attack vectors or entry points, *i.e.* all applicable ways an actor may exploit a technical exposure, can be informed by threat or breach detection and simulations (whether automated or manual) (Picus Security, 2020; Shostack, 2014). This would also inform controls performance. Ideally, however, this would be deduced by a rich ontology via a set of specific questions/queries captured either in the threat concept (potentially linked to intelligence and context) or the compensating controls and defence policy concepts for instance; (i) starting positions and attack steps given a vulnerability, (ii) specific systems being targeted, (iii) threat actor objectives, and (iv) controls or defences known to mitigate the vulnerability and their performance.

4.2 Motivating Example

Risk event definition is the lowest level of granularity in the assessment process. The scoping of risk events can be based on a very specific question, *e.g.,* what is the impact of a given vulnerability on an environment? Or it can be a more generic scenario, *e.g.,* what is the risk associated with the vulnerability management style? Taking into account the process maturity and delays in the assessment and remediation cycles. The three stages (pre, during, and post) are assessed for a risk event, especially as it relates to controls, to inform the estimation of the impact for the analysis. Table 2 presents typical core control components related to event stages. As an example, we will present the common attack surface case mentioned above, *i.e.* risk associated with vulnerability exploitation in a given environment.

Based on real-world data from a vulnerability management firm, Table 3 shows some of the top exploitable vulnerability exposure (by the number of affected environments) in scans from the mid 2019. We chose the Microsoft ActiveX Data Objects RCE vulnerability (CVE-2019-0888) for this purpose. Prerequisites of risk assessment may include: (a) an efficient asset inventory programme to gather and understand all the associated inventory interacting with the asset; (b) state of the network map, dependency graph of devices and various configuration and policies in the environment. Table 4 shows an example evaluation that identifies all metrics associated with the risk event and the controls that are part of the analysis, at each stage, and their performance.

It is also possible to go further by either aggregating indices over the evaluation metrics or estimating the loss values associated with potential impacts. However, this will lead back to a unique score or a band (with minimum and maximum exposure). The factors view as presented in Table 4 is much more suitable and make it easier, especially as it relates to the technical impact, to identify events with the highest risk potential. In this example, the assessment points to a *high risk* event given both the partial failure of controls and the low residual risk tolerance. Note that it is sometimes hard to interpret controls performance, especially when it comes to behavioural-based protection. In this example, it might wrongly be interpreted as success given the detection hits, however, reports of activity point to the fact that suspicious activity has happened, which might have led to a successful compromise. Understanding the methods and technology powering controls is also a necessity.

5 CONCLUSION

Building a comprehensive knowledge model of the cybersecurity domain remains a major objective for

Table 2: Core control components of risk event stages.

Pre Event	During Event	Post Event
Vulnerability Management program Controls {Firewall, IDS/IPS[1], EDR[2]} Training programs Policies & Administrative controls Threat intelligence platforms BAS[3] Platforms, Pen testing, red teaming activity	Controls {SIEM[4], SOAR[5], EDR} Recovery plans Forensic capabilities	Recovery plans Communication strategies After action reports

[1] Intrusion Prevention and Detection Systems
[2] Endpoint Detection and Response
[3] Breach and Attack Simulation
[4] Security Information and Event management
[5] Security Orchestration Automation and Response

Table 3: Top CVEs (by the number of affected environments) in mid 2019.

Description	CVE
Microsoft ActiveX Data Objects (ADO) RCE vulnerability	CVE-2019-0888
Elevation of privilege vulnerabilities affecting Windows 10	CVE-2019-1064, CVE-2019-1069
Meltdown & Spectre	CVE-2017-5715, CVE-2017-5753, CVE-2017-5754

the cybersecurity community. Various taxonomies, dictionaries, and terminologies have been proposed, covering multiple aspects of the cyber landscape. However, there is still a discrepancy between what is needed from the defender perspective and current models of cybersecurity, whether in relation to ontological representations or risk modelling.

In this work, we addressed the lack of adaptive and balanced treatment of defender efforts to better understand the uncertainty around risk and mitigate attacks, which is generally neglected in present cyber risk and cyber operations modelling, largely focused on the attack and the attacker. In fact, more focus on compensating controls is a much needed addition to the defender landscape to efficiently express its associated security and attached risks. Our future work will include further details and implementation of the ontology, its usage in compensating controls evaluation, and comparisons of existing automation in threat intelligence and context, and breach modelling and simulation.

REFERENCES

Blanco, C., Lasheras, J., Valencia-García, R., Fernández-Medina, E., Toval, A., and Piattini, M. (2008). A systematic review and comparison of security ontologies. In *2008 Third International Conference on Availability, Reliability and Security*, pages 813–820. IEEE.

Blank, R. M. (2011). Guide for conducting risk assessments.

Caralli, R., Stevens, J. F., Young, L. R., and Wilson, W. R. (2007). Introducing octave allegro: Improving the information security risk assessment process.

Den Braber, F., Brændeland, G., Dahl, H. E., Engan, I., Hogganvik, I., Lund, M., Solhaug, B., Stølen, K., and Vraalsen, F. (2006). The CORAS model-based method for security risk analysis. *SINTEF, Oslo*, 12:15–32.

FAIR (2020). Factor analysis of information risk. https://www.fairinstitute.org. Online; accessed February 2020.

Fenz, S. and Ekelhart, A. (2009). Formalizing information security knowledge. In *Proceedings of the 4th International Symposium on Information, Computer, and Communications Security*, ASIACCS '09, pages 183–194, New York, NY, USA. ACM.

Freund, J. and Jones, J. (2014). *Measuring and Managing Information Risk: A FAIR Approach*. Butterworth-Heinemann, Newton, MA, USA, 1st edition.

Fusun, M. B., Yaman, A. S., Marco, T., Carvalho, E., and Paltzer, C. N. (2016). Using ontologies to quantify attack surfaces.

Guarino, N. (1998). Formal ontology and information systems. pages 3–15. IOS Press.

Iannacone, M. D., Bohn, S., Nakamura, G., Gerth, J., Huffer, K. M., Bridges, R. A., Ferragut, E. M., and Goodall, J. R. (2015). Developing an ontology for cyber security knowledge graphs. *CISR*, 15:12.

ISO, E. (2011). IEC 27005: 2011 (EN) information technology–security techniques–information security risk management switzerland. *ISO/IEC*.

Lund, M. S., Solhaug, B., and Stølen, K. (2011). Risk analysis of changing and evolving systems using CORAS. In *International School on Foundations of Security Analysis and Design*, pages 231–274. Springer.

Martin, R. A. (2008). Making security measurable and manageable. In *MILCOM 2008-2008 IEEE Military Communications Conference*, pages 1–9. IEEE.

Mavroeidis, V. and Bromander, S. (2017). Cyber threat intelligence model: An evaluation of taxonomies, shar-

Table 4: A vulnerability risk event characterisation example.

Risk Event	Asset	Threat	Control	Impact	Risk Tolerance
$E_{CVE-2019-0888}$	Client host (Corporate, Information) - Restricted	Compromise; malicious payload/ backdoor Entry points: IE (& apps using the IE kernel), Office, potentially other VBScript apps (IIS, WSH) Attack vectors: Web pages, E-mail attachments Exploitation: High (in the wild)	All assets covered Pre: EDR (Partial/ suspicious PowerShell activity), IPS/IDS (Partial/ Download reported), Phishing awareness training (Fail) During: Time to Respond - 14 days No forensic capabilities Post: Full recovery from backup External communication - linked to 'Time to Respond' and mitigation	Technical: (ATT&CK): Initial Access, Persistence, Exfiltration Business: Productivity, Response, Replacement	Low

ing standards, and ontologies within cyber threat intelligence. In *2017 European Intelligence and Security Informatics Conference (EISIC)*, pages 91–98. IEEE.

MITRE (2019a). Common vulnerabilities and exposures. http://cve.mitre.org. [Online; accessed December 2019].

MITRE (2019b). Common weakness enumeration. http://cwe.mitre.org. [Online; accessed December 2019].

MITRE (2020a). The attack matrix. http://attack.mitre.org. [Online; accessed January 2020].

MITRE (2020b). Common attack pattern enumeration and classification. http://capec.mitre.org. [Online; accessed January 2020].

NIST (2020a). Framework for improving critical infrastructure cybersecurity version 1.1. https://www.nist.gov. [Online; accessed February 2020].

NIST (2020b). National Institute of Standards and Technology. https://www.nist.gov. [Online; accessed February 2020].

NVD (2020). National vulnerability database. https://nvd.nist.gov. [Online; accessed January 2020].

OASIS (2020a). Structured Threat Information eXpression (STIX). https://stixproject.github.io. [Online; accessed January 2020].

OASIS (2020b). Trusted Automated eXchange of Indicator Information (TAXII). https://taxiiproject.github.io. [Online; accessed January 2020].

Picus Security (2020). Continuous & real world cyber-threat simulation. https://www.picussecurity.com/platform.html#how-does-picus-work. [Online; accessed February 2020].

Rajbhandari, L. (2013). Risk analysis using "conflicting incentives" as an alternative notion of risk.

Shostack, A. (2014). *Threat Modeling: Designing for Security*. Wiley Publishing, 1st edition.

Snekkenes, E. (2013). Position paper: Privacy risk analysis is about understanding conflicting incentives. In *IFIP Working Conference on Policies and Research in Identity Management*, pages 100–103. Springer.

Syed, R. and Zhong, H. (2018). Cybersecurity vulnerability management: An ontology-based conceptual model.

Syed, Z., Padia, A., Finin, T., Mathews, L., and Joshi, A. (2016). UCO: A unified cybersecurity ontology. In *Workshops at the Thirtieth AAAI Conference on Artificial Intelligence*.

Symantec (2019). Internet security threat report 2019.

Tenable Network Security (2020a). The modern attack surface. https://www.tenable.com/cyber-exposure. [Online; accessed February 2020].

Tenable Network Security (2020b). Tenable lumin cyber exposure analytics. https://www.tenable.com/products/tenable-lumin. [Online; accessed February 2020].

Wang, J. A. and Guo, M. (2009). 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*, page 34. ACM.

Wangen, G. B. (2017). Cyber security risk assessment practices: Core unified risk framework.

Yazar, Z. (2002). A qualitative risk analysis and management tool–CRAMM. *SANS InfoSec Reading Room White Paper*, 11:12–32.