Information & Security

V. Posea et al.
vol. 53, no. 1 (2022): 33-44
https://doi.org/10.11610/isij.5303

Published since 1998

ISSN 0861-5160 (print), ISSN 1314-2119 (online)

Research Article

Towards Unified European Cyber Incident and Crisis Management Ontology

Vlad Posea [1], *George Sharkov* [2] (✉), *Adrian Baumann* [3], *and Georgios Chatzichristos* [4]

[1] *Politehnica University of Bucharest, Bucharest, Romania https://upb.ro/en/*

[2] *European Software Institute – Center Eastern Europe, Sofia, Bulgaria https://esicenter.bg/*

[3] *Federal Office of Information Technology, Systems and Telecommunication, Switzerland, https://www.bit.admin.ch/*

[4] *ENISA (European Union Agency for Cybersecurity), https://www.enisa.europa.eu/*

ABSTRACT:

ENISA highlighted the need for a common reporting taxonomy for cybersecurity incidents to be used by cybersecurity analysts across Europe. The analysis of the domain revealed a large number of taxonomies for different areas of the cybersecurity domain (types of attacks, vulnerabilities, sectors, harm), but those needed to be linked together in a model that allows a cybersecurity officer to report and track an incident fast and accurately. The taxonomy should also treat the cybersecurity domain not only from the technical point of view but also from the socio-economical aspect. This document describes the taxonomy, how we propose to use it, and the methodology used to develop it.

ARTICLE INFO:

RECEIVED: 29 JULY 2022
REVISED: 08 SEP 2022
ONLINE: 21 SEP 2022

KEYWORDS:

cybersecurity taxonomy, ontology, incident response, threat, crisis, risk assessment, harm, interoperability, Cybersecurity Act

✉ Corresponding Author: E-mail: gesha@esicenter.bg

Introduction

An important aspect of a cybersecurity analyst's work is reporting on cybersecurity events worldwide. These events can be identified directly by observing a system under attack or as an effect of a natural phenomenon, also by incident coverage in public media. Depending on the level of implication, the analyst can be interested in several types of information, from the very low-level technical details to social or personal harm that occurs at the highest level. The analysts need help in describing the incident using this 360° approach. They need a hierarchically structured model and rich vocabularies that allow them to express all details of their report. Finally, such a report should be exportable in a format that allows interoperability with other systems and organizations. The hierarchical structure of the information in a tree (or taxonomy) allows the analyst to describe facts that can be further analyzed in an automatic or semi-automatic fashion. However, a taxonomy contains only relationships like class-subclass or class-instance. Arguments have been made that taxonomies are not an adequate way to represent knowledge as they cannot express different types of relationships and constraints. Gruber introduced the term ontology [1] as a "specification of a conceptualization," basically a system of classes and relationships that describe the data structure. Different domain experts might create different ontologies to model the same domain depending on the purpose and use of the ontology.

When starting this work on modeling the Cybersecurity domain to facilitate reporting of cybersecurity incidents across Europe, we soon discovered that while taxonomy was clearly not expressive enough, a fully formal ontology would have been extremely difficult to model and use as most information sources available were not well structured. Then we decided to model the Cybersecurity domain by a lightweight ontology. A "lightweight ontology" is a term introduced by Davies [2] to describe an expressive way to model information using classes and relationships without considering aspects such as depth and computational tractability. This approach is in line with how most of the domain's existing information is modeled. It allows us to generate a rich, expressible model in interoperable formats and covers a large area of the domain knowledge.

Methodology and State of the Art

Many methodologies have been developed across time for ontology development as the domain of ontology engineering has existed since the end of the twentieth century. But still, the most cited and used methodology for ontology development is "Ontology 101" by Noy and McGuiness from 2001.[3] It considers that the steps for building an ontology should be: determine the scope, consider reusing existing ontologies, enumerate important terms, define the classes and the hierarchy, the properties of classes (slots), define the facets of the slots, create instances. The advantages of this approach are that it is easy to explain and involve the domain experts, especially in the first steps. The scope of the

ontology is determined by asking competency questions. However, for the second part of the work, we applied another, more agile methodology – Upon Lite,[4] which relies more on domain experts, a collaborative approach, and tools to build comprehensive models faster. The main steps in this model are terminological level, glossary level, taxonomy level, predication level, parthood level (the PartOf relationships), and ontology level. We also considered it important to apply a user-centered approach and build a usable taxonomy that meets real needs. We combined the two methodologies using the steps of determining the scope and the existing ontologies from the Ontology 101, and the more agile and expert-intensive approach of Upon Lite for the rest. We followed these six steps:

1. *Define the scope and refine through the competency questions.*
2. *Identify other ontologies/taxonomies that can be used/reused.*
3. *Identify sources from where we can extract the knowledge to model.*
4. *Define the main concepts and the relationships between them.*
5. *Define the properties for the concepts, involving the domain experts.*
6. *Implement the ontology using OWL.*

Steps 1-5 were performed iteratively using regular sprints and workshops, and data was formalized using shared Google sheets. The Miro[5] collaborative board was used to model the relationships between the top-level concepts.

Other Relevant Ontologies/Taxonomies

JRC Taxonomy, Based on NIST

The JRC Taxonomy[6] represents extensive research on taxonomies, documents, and regulations in the field. In focus are the domains of cybersecurity, economic sectors that can be affected, and technologies and use cases. The JRC taxonomy was looked upon to be used to classify domains for incidents and for the economic sectors impacted or where a possible impact could be. Including the JRC taxonomy in our model would allow us to answer the following questions:

- What is the cybersecurity domain of an incident/risk?
- What is the sector in which an incident has an impact?
- What are technologies and/or use cases that could be involved?

UCO (Unified Cybersecurity Ontology)

UCO[7] is an ontology used for "standardized information representation across the cyber security domain/ecosystem." It is extremely complex and can be used as a reference for a formal implementation, as other ontologies also refer to it. As a downside, UCO does not seem to consider the impact, and even if the paper describing UCO says it contains concepts like Exploit and Attack, the last version published on Github does not seem to contain them. Interesting aspects of the UCO ontology could be the "vocabulary" namespace and the "observable" namespace. UCO is highly specific, containing even classes for windows and unix

threads. However, UCO does not link to other ontologies so that we could not reuse it for our work.

ENISA Cyber Incident Taxonomy (Blueprint, High-level)

This high-level taxonomy was proposed to classify cyber security incidents at the strategic political level to meet the Cybersecurity Act requirements. It was developed by Stream 7 of NIS Cooperation Group on "Large scale cybersecurity incidents."[8] This taxonomy is to be used under the Integrated Political Crisis Response (IPCR) framework. The taxonomy addresses only the "naming" of cybersecurity incidents and not the "processes" (e.g., for notifying or escalating incidents). It does not exclude the use of additional taxonomies when a more specific classification is needed. The taxonomy has two core parts: "The nature of the incident" with Root cause category (system failure, natural phenomena, human error, malicious action, third party failure) and Severity of threat (high, medium, low), and the second "Impact" part with Sector impacted, Scale of impact (red, yellow, green, white), and Outlook (improving, stable, worsening).

ENISA Threat Taxonomy

A comprehensive, structured threats taxonomy was first proposed in 2015 as a general classification of threats in "internet infrastructure" and subsequently used and updated for ENISA annual "Threat Landscape Reports." These reports identify prime threats, trends, threat actors and attack techniques, and appropriate measures. The threat taxonomy, along with internet infrastructure taxonomy are described in detail in "ENISA Threat Landscape and Good Practice Guide for Internet Infrastructure."[9] Threat types have been defined based on the cause in nine groups. The guide also refers to a risk assessment equation based on the classical risk components, as outlined in ISO 27005. The Risk is "calculated" based on the following three elements: ***Asset*** *(Vulnerabilities, Controls),* ***Threat*** *(Threat Agent Profile, Likelihood), and* ***Impact.***

We have listed the general groups for our purposes, but the types of threats are open for updates, especially for malicious activities. Sector-specific threat types are added or updated by ENISA in regular Thematic Landscapes (e.g., supply chains, smart grids, AI, 5G, IoT, etc.). A valuable "proximity" classification is given following the classification for the EU Common Security and Defence Policy (CSDP) with four levels: near, mid, far, and global. This classification of threat categories is the basis for structuring the information in OSINT (Open-Source Intelligence) work by ENISA in the area of Situational Awareness.

ENISA Reference Incident Classification Taxonomy

This taxonomy resulted from collaboration initiatives such as the annual ENISA/EC3 Workshop involving CSIRTs, LEAs, ENISA, and EC3. The proposed Reference Taxonomy is based on a tight correlation with the "Common Taxonomy for LE and CSIRTs," an adaptation of the CERT.PT taxonomy, and the eCSIRT.net mkVI, being also well mapped to other taxonomies. This taxonomy is limited to

incidents related to human-caused cyber offenses and attempted offenses. References to the legal framework are also continuously updated to provide a basis for prosecuting the incidents. The taxonomy is widely used for info sharing between CSIRTs, Law Enforcement Agencies, and Europol. The popular MISP Platform [10] provides a mapping of existing taxonomies which allows analysts to use the appropriate namespace and values.

Threat Intelligence Taxonomy: MITRE ATT&CK and Cyber Kill Chain

MITRE ATT&CK (Adversarial Tactics Techniques and Common Knowledge) is a knowledge base of adversary tactics and techniques based on accumulated real-world observations. It is a structured list of known attacker behaviors that have been compiled into tactics and techniques and expressed in many matrices as well as via STIX/TAXII. The list is a comprehensive representation of behaviors attackers employ when compromising networks. The model is technically more detailed than the generic Cyber Kill Chain model and, therefore, more applicable to our ontology. The basic ATT&CK categories of tactics are 14, like: Reconnaissance, Resource Development, Execution, Privilege Escalation, etc. It should be noted that this list, along with the associated known techniques, does represent the cyberattack "lifecycle model," and not every attack uses all tactics, but the observable evidences for applying tactics and techniques are the most detailed technical model in use. As for modeling the lifecycle of a cyber attack, Lockheed Martin's Cyber Kill Chain framework is widely used both for investigating and simulating (e.g., by "red teams").

MITRE CVE, CAPEC, CWE

The mission of the MITRE CVE (Common Vulnerabilities and Exposures) Program is to identify, define, and catalog publicly disclosed cybersecurity vulnerabilities. The CVE Records are used by the professionals to ensure they are discussing the same issue and to coordinate their effort. This repository is the global reference for vulnerabilities. MITRE supports other reference taxonomies and repositories, such as CAPEC (Common Attack Pattern Enumeration and Classification taxonomy) and CWE (Common Weakness Enumeration).

OASIS Cyber Threat Intelligence Taxonomies. STIX

STIX is a schema that defines a taxonomy of cyber threat intelligence. It is developed and hosted by the Technical Committee for Cyber Threat Intelligence (CTI) of OASIS (Organization for the Advancement of Structured Information Standards). The specifications published and maintained include Trusted Automated eXchange of Indicator Information (TAXII), Structured Threat Information eXpression (STIX), and Cyber Observable Expression (CybOX). The STIX model describes an adversary and adversary activities in appropriate data structures. STIX Domain Objects cover Threat Actor, Malware, Tools, Campaign, Intrusion Set, and Attack Pattern (referencing CAPEC). A Campaign is defined as a grouping of adversarial behaviors, attack patterns, malware and/or tools. The model of a Threat Actor has several relevant properties, such as goals, sophistication,

resource level, etc. The STIX domain model also includes observables, indicators, and courses of action. However, STIX does not specify a specific set of kill chain phases, but those could be indicated as optional properties of Attack Pattern, Indicator, Malware, and Tool objects. STIX is excellent for the technical part of cybersecurity but lacks impact evaluation on economic/political levels. Cyber threat actors identified by high sophistication, advanced capabilities, and tools are named Advanced Persistent Threats (APTs). Large-scale incidents and campaigns are associated with one or more APTs, thus providing better ground for risk assessment, prevention, and response.

The Cybersecurity Taxonomy

The Competency Questions

The competency questions for an ontology represent those questions that we'd like our ontology to be able to answer. These questions are also helpful for the knowledge engineer to understand what the domain experts and the stakeholders want to achieve from this ontology. Analyzing these starting questions, we proposed the following set. These questions were discussed and iterated during the development of the ontology. Some of the questions were not answered in the first draft of our ontology but remained for methodological purposes.

The formulated **Incident-related** 14 questions further detailed the cybersecurity incident, like What is the nature of an incident? How severe is the incident? What types of incidents? When did an incident occur?

Then we had 7 **Risk-related** questions: How do we define risk? What are the types of risk? What are risk mitigation measures for each type of risk? What is the connection between a risk and a threat? What is the possible cost of not mitigating a risk?

The four **Threat-related** questions addressed the definition of a threat, types, associated risks, and the possible technical/operational/political impact.

Since we identified a more common understanding of an impact, we have introduced the notion of **Harm** with three generic questions.

The nine final **Competency** questions were the most important for the reporting process and tightening it all together. They address the evidences and observations, the links between incidents and suspected campaigns, the related harm(s), assets and targets, etc.

The Model and the Top Level Elements

Based on the competency questions and the domain analysis, the model presented in figure 1 has been sketched. The story behind this model is that the analyst can observe an incident or can identify a threat. The incident produces harm(s), affects one or more assets, can be caused by one or more actors, can have one or more targets, can be caused by a vulnerability, and can be determined by a threat that was already observed. An incident can be similar to one or more other incidents and directly related to one or more other incidents. An incident can have one or more responses, which can also be observed by the analysts. The observations can be based on evidence, whether reports or log

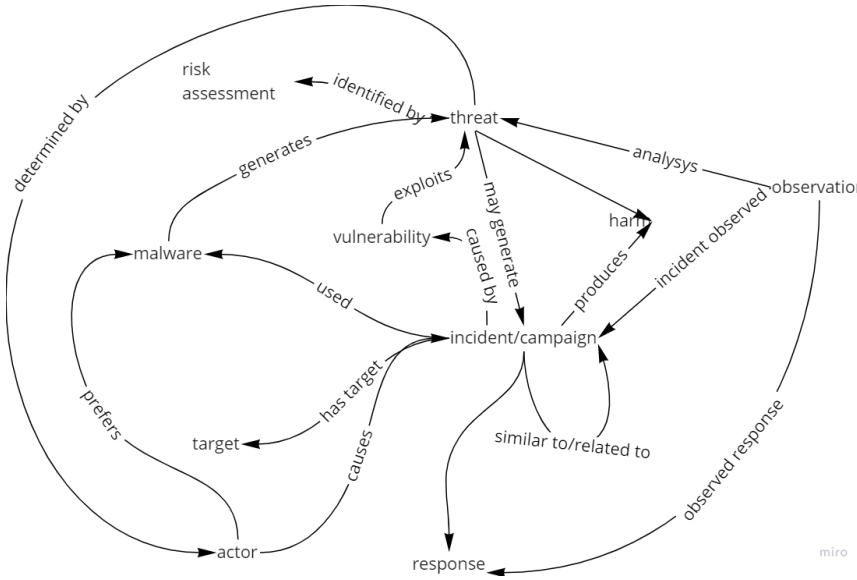

Figure 1: Top-level visualization of the ontology (miro whiteboard).

files, or even news items or tweets. The threats are identified based on risk assessments that are made by experts or through direct observation.

The model is defined at a higher level, and less focused on the low-level technical details of the attack. The top-level approach allows us to assess targets and harm produced and may, in the future, allow us to better estimate risks and prepare better defenses. The top-level classes with relevant taxonomies are defined explicitly or preferably by referencing existing maintainable taxonomies. Mapping between different taxonomies is also foreseen (e.g., for economic sectors, EU NIS2 and US CISA).

The primary classes identified and outlined with relations in Figure 1 are:

Observation. Observation is the key instrument of the cyber security analyst. It can be direct by the expert or indirect, reported by the analyst based on articles, tweets or other sources of information. An analyst can observe a Threat while performing a system analysis or can observe a security Incident or what he believes to be a Campaign. The analyst can correlate the observation with others on the same topic and can also observe a response. The analyst also indicates the status of the risk/incident, allowing further tracking of events over time. The detection can be direct (signature-based), indirect (anomaly-based), or based on other evidence. The evidence might have a link attached and a type (OSINT, closed source). Building the taxonomy around what is observable is essential as we can't report on what we can't observe somehow. And the purpose of the taxonomy is to serve a common reporting framework.

Incident and Campaign. While the Observation is the core entity for the analyst, the Incident is the core of the technical part of the system. The Incident is what the analysts mostly report upon, what technical teams need to address with their Responses, and what causes various types of Harm to the direct and indirect Targets. According to the ISO/IEC 27000:2018 an information security incident is a "single or a series of unwanted or unexpected information security events that have a significant probability of compromising business operations and threatening information security." The Incident might belong to different classes. The OSINT type property was kept for backward compatibility. The "nature" property is used to describe incidents not caused by nefarious activity. Two classifications of the Sectors (US and EU) are used. In case the nature is "malicious action" the MITRE attack type property is used. Incidents could also be linked in various ways - similar patterns, attackers, cascading effects, etc. An incident can be caused by a Threat Actor and could affect multiple Assets. It can be generated by a threat and caused by a vulnerability. The Campaign (or "scenario") can be formed by a series of Incidents. It can have one or more Targets and be performed by one or more Threat Actors. The Campaign can cause Harms to Assets different from those caused by its Incident parts. Campaigns can be similar to other Campaigns.

Harm. Another key concept of the taxonomy is Harm. Cybersecurity incidents often focus on the harm done to systems, networks, and possibly organizations. The damage, however, is often done at multiple levels. Users can be hurt, and employees can be hurt physically, economically, and socially. Whole groups of people or various organizations can suffer different types of damage. We need to allow our analysts to observe and report different types of harm caused by an Incident or a Campaign not only when the Incident was initially observed but also days, months, and years afterward. We modeled harm using the taxonomy in the "Cyber harm taxonomy."[11] Harm can be physical, emotional, economic, political, reputational, and cultural. For each type of harm there can be several types of damage. The type of harm depends on the subject type, which can be individual, organizational, property or infrastructure, and national. For this first version of our taxonomy, we didn't implement the constraints suggested.

Threat. According to the same ISO cited above (ISO/IEC 27000:2018) a "threat" is a potential cause of an unwanted incident, which may result in harm to a system or organization. The Threat concept is already defined in the ENISA Threat Taxonomy. From this taxonomy, we decided to keep the following categories: high-level threats, threats, and descriptions. A Threat will have a property called threat category which will have values in a Threat Category concept. The Threat can be observed by an analyst or from a risk assessment report. It can have severity levels (high, medium, and low) and can be attached to an identified Threat Actor.

Vulnerability. The concept of Vulnerability in cybersecurity is defined in ISO/IEC 27000 as the weakness of an asset or control that one or more threats can exploit. The Vulnerability is usually regarded from the point of view of a software or hardware issue that can be exploited, and we have a very good model of the Vulnerability concept in STIX. Vulnerability has a name, a description and in STIX is connected to concepts like infrastructure, course-of-action, malware and campaign. In our model, we have added for the vulnerability two properties - a type which could initially have only two values - cybersecurity and other. The instances of the Vulnerability class with the type equal to cybersecurity will come from the CVE repository of MITRE and will have as an identifier the CVE ID defined by MITRE. In our ontology the Vulnerability is connected to a Threat as the Threat can be based on a specific vulnerability of a system and it is connected to an Incident as the Incident could have occurred because of a specific vulnerability.

Risk Assessment. The Risk is defined (ISO/IEC 27000:2018) as an effect of uncertainty on objectives. In our current ontology, we only define a concept called Risk assessment which is an estimation of a risk. The risk assessment is attached to a Threat, and intuitively it is an estimation for a threat to actually generate harm. The risk assessment is still a work in progress and will be extended in future versions.

Threat Actor. In case the Incident is caused with malicious intent, a Threat Actor most likely caused it. The Threat Actor is modeled in quite a few taxonomies, but we felt the one that best fitted the needs of this one, and it is also important from the interoperability perspective, was the model from STIX v2 . Besides the name and description, we included properties defined in STIX open vocabularies like threat actor type, actor, sophistication, roles, resource level, and motivation. We also considered it relevant to be able to link a Threat Actor to an APTs as they are defined in multiple places with multiple names. APTs are also more formal than the aliases property, which is the way for the STIX to express something similar. The Threat Actor is directly connected to the Incident and with the Campaign. In the future, connections between Threat Actors can be detected and/or inferred based on similar behaviors of the system.

Target. The Target of an attack is not always known, and sometimes it is guessed, or alternatively, sometimes we can't be sure if the Target is the person/organization that was harmed, or it was just collateral harm. Therefore, the Target is an entity strongly connected to the Harm, as most of the time, we detect the target through the damage we observe. There is also a direct connection between the Incident and the Target for the purposes when the analyst is not interested in logging a Harm, but they know there is a connection between an Incident and a Target. The Target also has a name, a description, and can have a URL. A Target can belong to one or more Sectors (both EU and US versions) and can have one or more countries where it is localized and a distance (near, far, global).

Response. The Response concept has been considered connected to an Incident/Campaign and also to an Observation. The analyst can observe incidents and then log the reaction to them. In time lessons can be learned by centralizing information about responses to incidents. Responses can occur at different technical, political, and operational levels and could address specific Harm. However, the decision was that the Response part of the taxonomy will be sketched in the next development iteration.

Sector. The Sector represents the area of human activity that can be affected by a cybersecurity incident. There are multiple classifications for sectors, and we chose to use two of them, representing mainly the "critical infrastructures" - the NIS2 Directive and Cyber Resilience Act proposed classification and the US classification maintained by CISA. Both categories will be attached to the Incident concept, and the analyst will choose which of them to fill in (or even both) when describing an Incident. The subclassification is almost identical for the NIS2 and the US, splitting the sectors into subsectors. The NIS2 classification also has the concept of an entity, and the entities are attached to subsectors.

Malware. Malware is defined in STIX as a type of tactic, technique, or procedure that "represents malicious code" and "generally refers to a program that is inserted into a system, usually covertly." Malware appears extremely often in cybersecurity reporting because its existence makes it easy for many types of threat actors to create incidents. Although the scope of this taxonomy wasn't to describe the domain at a very low and technical level, we decided to include the concept of Malware as it is ubiquitous. Malware has a name, a description, a list of aliases, and a link to an external site where more info could be found. The Malware is connected to the Incident it causes, to the Threat it represents, and to a Threat Actor who might use it. In cases of extending the ontology to cover technical aspects of cybersecurity, the Malware concept should be made a subclass of TTP and further details regarding its functionality and distribution.

Implementation and Results. Further steps

The implementation of the ontology was performed in two steps, according to the methodology presented. The first step consisted of filling spreadsheets with the domain experts and defining the concepts and the values from existing vocabularies. The next step was converting the sheets into a different format that could be easily transformed into an OWL file afterward. The ontology is provided in two formats: as tables of classes together with their properties and the various vocabularies that are identified, and the second format consists of an OWL file, which will be published on ENISA server having an URL identical with the namespace of the ontology.[12]

The work presented was focused on gathering and structuring knowledge from multiple domains related to the field of cybersecurity to facilitate reporting across institutions inside and outside the EU. For better expressiveness, we

developed it beyond a taxonomy in a "lightweight ontology" containing concepts in the domain and introducing a large number of relevant relationships between them. This model is directly actionable towards a reporting tool.

For this first version of the model, there are areas which were just sketched and not sufficiently developed, like the Response to an Incident or the Countermeasures. Appropriate metrics for different categories are also under development to propose harmonized qualitative/quantitative "formulas" for assessments of Risk, Impact, Harm, alert levels, etc.

Acknowledgements

This research is co-funded by ENISA (EU Agency for Cybersecurity).

References

[1] Thomas R. Gruber, "A translation approach to portable ontology specifications," *Knowledge Acquisition* 5, no. 2 (1993): 199-220, https://doi.org/10.1006/knac.1993.1008.

[2] John Davies, "Lightweight Ontologies," in *Theory and Applications of Ontology: Computer Applications*, eds. Poli, R., Healy, M., Kameas, A. (Dordrecht: Springer, 2010), https://doi.org/10.1007/978-90-481-8847-5_9.

[3] Natalya F. Noy and Deborah L. McGuinness, "Ontology development 101: A guide to creating your first ontology," 2001.

[4] A. De Nicola and M. Missikoff, "A lightweight methodology for rapid ontology engineering," *Communications of the ACM* 59, no. 3 (2016): 79–86.

[5] MIRO: online collaborative whiteboard platform, 2022, https://miro.com

[6] Igor Nai Fovino, Ricardo Neisse, Ramos Hernandez, Luis Jose, Nineta Polemi, Gian Luigi Ruzzante, Malgorzata Figwer, and Alessandro Lazari, *A proposal for a European cybersecurity taxonomy* (European Commission, Joint Research Centre, Publications Office, 2019), https://doi.org/10.2760/106002.

[7] Zareen Syed, Ankur Padia, Tim Finin, Lisa Mathews, and Anupam Joshi "UCO: A unified cybersecurity ontology," *Workshops at the thirtieth AAAI conference on artificial intelligence,* Technical Report WS-16-03, 2016.

[8] "Cybersecurity Incident Taxonomy," CG Publication 04/2018, NIS Cooperation Group, July 2018, https://ec.europa.eu/information_society/newsroom/image/document/2018-30/cybersecurity_incident_taxonomy_00CD828C-F851-AFC4-0B1B416696B5F710_53646.pdf.

[9] ENISA, "Threat Landscape and Good Practice: Guide for Internet Infrastructure," January 2015, https://www.enisa.europa.eu/publications/iitl/@@download/fullReport.

[10] Malware Information Sharing Platform – Open Source Threat Intelligence Platform & Open Standards for Threat Information Sharing, 2022, https://www.misp-project.org/.

[11] Ioannis Agrafiotis, Maria Bada, Paul Cornish, Sadie Creese, Michael Goldsmith, Eva Ignatuschtschenko, Taylor Roberts, and David M. Upton, "Cyber harm: concepts, taxonomy and measurement," Saïd Business School, WP 23-2016, August 4, 2016.

[12] ENISA, Threat and Risk Management – Taxonomy, 2022, https://enisa.europa.eu/topics/threat-risk-management/extended-cybersecurity-taxonomy.

About the Authors

Vlad Posea is an assistant professor at the Politehnica University of Bucharest and an experienced consultant in knowledge and product management. His main topics of interest are open data, semantic web and knowledge modeling with a focus on how these tools can be used to actually improve users' experience. https://orcid.org/0000-0002-6864-0165

George Sharkov has been CEO of the European Software Institute CEE since 2003 and Head of the Cybersecurity Lab at Sofia Tech Park. He was an adviser to the Bulgarian Minister of Defense (2014-2021) and National Cybersecurity Coordinator, leading the development of the national Cyber Resilience Strategy. Member of the EU AI High-Level Expert Group, SMEs voice at ETSI Technical Committees CYBER and ISG "Securing AI," ENISA Group on Secure AI, ENISA Stakeholders Cybersecurity Certification Group. He holds Ph.D. in AI and is lecturing at four leading universities on software quality, cybersecurity and resilience, and active security. https://orcid.org/0000-0001-5086-311X

Adrian A. Baumann is a senior cyber security officer in the Swiss public administration. Originally a linguist with focus on the English Language, he sidelined into IT during his studies and eventually decided to switch careers. He has been with the Swiss civil service for 15 years, of which he has been working to improve security for the last eight. From 2017 to 2020 he was acting as deputy head of security and risk management before joining ENISA as the first Swiss seconded national expert from 2020 to 2022. He is currently back at the Federal Office of IT and Telecommunications. https://orcid.org/0000-0002-0731-6126

Georgios Chatzichristos is a Cybersecurity Officer at the Operational Cooperation Unit of the European Union Agency for Cybersecurity. He is working on Cyber Operations and Situational Awareness. His academic background includes Postgraduate Studies on Artificial Intelligence. Mr Chatzichristos is a retired Commander of the Hellenic Navy. Before joining ENISA, Georgios was the Planning & Exercises Cell Leader of the Cyber Defence Branch at Allied Command Operations (SHAPE,Belgium).