CSIHO: An Ontology for Computer Security Incident Handling

Guilherme Baesso Moreira[1], Vanusa Menditi Calegario[2], Julio Cesar Duarte[1], Anderson Fernandes Pereira dos Santos[1]

[1]Instituto Militar de Engenharia (IME) – Rio de Janeiro – RJ – Brazil

[2]Centro Dom Vital – Rio de Janeiro – RJ – Brazil

gbaesso@gmail.com, vmenditi@yahoo.com.br, {duarte, anderson}@ime.eb.br

Abstract. *The information technology advancements in the last decades led the society to a growing process of dependency on computer systems and Internet-based services. This complex and dynamic scenario implies more challenging cyberdefense initiatives, but, although the industry is applying countless efforts to ensure the Information Security, considerable growth in frequency and severity of incidents is still observed. The primary objective of this work is to present a new model for incident handling, described as an ontology, which is easily extensible and integrable with other models, besides allowing logical inferences and simplifying the knowledge transfer within a collaborative cyberdefense context. Among its contributions, the creation of the Computer Security Incident Handling Ontology (CSIHO), in OWL format, can be highlighted. In order to demonstrate the applicability of the ontology, SPARQL queries were created based on competency questions derived from CSIHO, which, as far as we know, is the first cyber security ontology that focuses on incident handling and defines and implements the fundamental concepts of security events while also supporting the recording of temporal aspects of an incident.*

1. Introduction

Recent cybersecurity related news and reports reveal that, besides the high investments in Information Security initiatives, cyber attacks are still causing significant losses to the companies [O Globo 2015], and the number of incidents keeps growing year after year [CERT.br 2018]. Most of the current cyber attacks have political or financial motivations [Kharraz et al. 2015] and many of them are funded by States, as part of a cyberwarfare context [Healey 2016]. As there is no perfect security solution, an attacker with more proficiency, resources or motivation will very likely succeed and, consequently, most organizations will eventually suffer a security incident.

Despite this, traditional security approaches are much more focused on detection and prevention than incident response [Baskerville et al. 2014] and the literature is inconsistent when describing incident management processes [Ab Rahman and Choo 2015], leading most companies to be unprepared to respond to security incidents [Computer World 2017, F-Secure 2017] and act in an improvised way, poorly documenting its incidents [Grispos 2016, Moreira et al. 2017]. Not surprisingly, exercises of "lessons learned" in this field are frequently ignored [Cichonski et al. 2012].

These problems show that there is a need for a robust security model that implements incident response process, benefiting from web semantics, enabling proper knowledge representation as well as allowing some level of inference.

Incident response and treatment process is a subset of the incident management macro-process that also includes vulnerabilities treatment, artifacts treatment, events management, announcements and alerts, as illustrated in Figure 1.

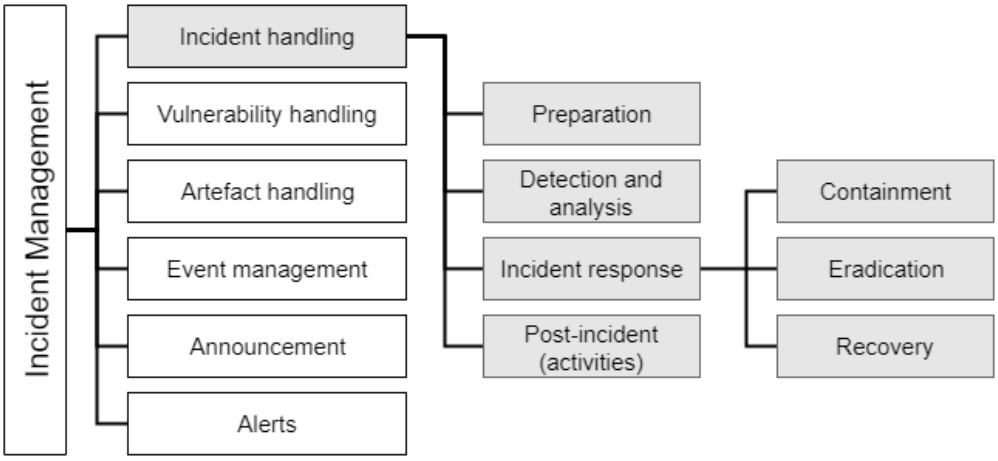

Figure 1. Incident Management macro-process. Adapted from [Ab Rahman and Choo 2015]

The incident treatment process consists, generally, in the following steps:

1. An event (or set of events) with a security violation suspicion is detected;
2. The situation is further investigated, and, if the violation is confirmed, the occurrence of an incident is characterized;
3. Immediate actions are applied to interrupt the violation;
4. Additional actions are applied to avoid violation recurrence;
5. Lastly, the environment is restored to the state it was before the violation.

In this paper, we adopted the following taxonomy to describe the incident handling process: IDENTIFICATION (steps 1 and 2 above), RESPONSE and FOLLOW-UP. The response phase comprises the phases CONTAINMENT, ERADICATION and RECOVERY, respectively consisting of steps 3, 4 e 5 of the generalization above. The FOLLOW-UP process consists of all the tasks related to the general documentation and follow-up of the incident, focused on building knowledge to support the RESPONSE process during its occurrence, to allow insights and metrics for process improvements on later "lessons learned" exercises and also to provide useful information for technical and executive *post-mortem* reports.

The higher the level of knowledge about the violation, the more effective the response actions will be. Furthermore, it is essential to maximize the knowledge about incident causes, consequences and sources since its confirmation.

With the exponential growth of data published on the Internet and in computer systems, the ability to extract useful information related to a specific topic has been decreasing rapidly. One of the limitations found in traditional systems search is in the inability

to recognize ambiguities in terms and to perform efficient queries in environments with uncontrolled format [Jakus et al. 2013].

[Shadbolt et al. 2006] defines Web Semantic as an evolving extension of World Wide Web in which web content can be expressed not only in natural language but also in a format that can be read and used by software agents, enabling them to find, share and integrate information more efficiently. The Semantic Web may be considered a Web of Metadata, that is, data about the details of the various websites, data sources and their relationships. It can be defined as a framework for integrating multiple data sources from different structures (structured, semi-structured and unstructured) providing them with the proper semantics.

The first formal definition of a standard for Web Semantic happened in 1997 when the World Wide Web Consortium (W3C) introduced the Resource Description Framework (RDF) specification, a data model with a simple but powerful representation, based on triples (subject, predicate, object). The subject is the entity that is being represented, the predicate defines an attribute or a relation of the subject with another object, that can be a literal or a resource itself, allowing the formation of a "web of data" (Figure 2). To avoid ambiguities, subjects and predicates must be URIs (Uniform Resource Identifiers), which means they belong to specific namespaces. URIs are utilized to identify any entity in Web Semantics in the same way URLs are used to identify the sites on the World Wide Web.

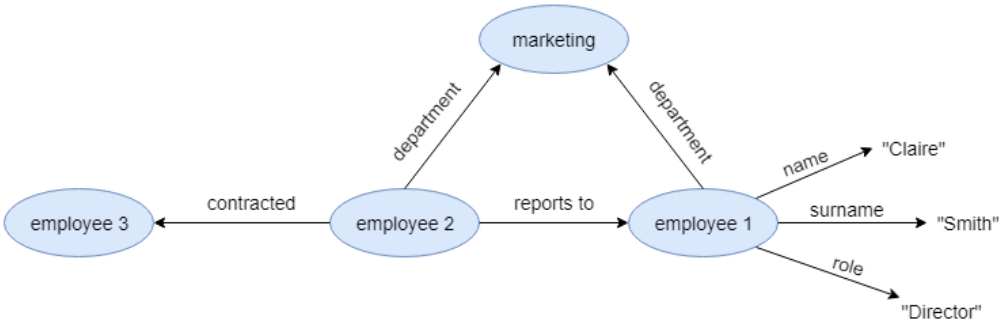

Figure 2. A graph illustrating connected triples

SPARQL (SPARQL Protocol And RDF Query Language) enables queries execution in RDF repositories, as presented in the sample use cases in section 4.

The Web Ontology Language (OWL) was later introduced by the W3C as a language to define ontologies. It is considered one of the most important standards defined to bring semantics for the web [DuCharme 2013]. Using fundamentals of Description Logics, the OWL extends the RDF model, allowing more expressiveness on the definition of concepts and object relations. It also offers support for some kinds of inferences, typically classification and subsumption, through the use of several available reasoning tools, like the HermiT Reasoner, which uses a state of the art algorithm.

This work is structured as follows. On the next section, related works are shown. Then, CSIHO is described and an evaluation is done with competency questions using SPARQL queries from the knowledge base. Lastly, the conclusions and suggested future works are presented.

2. Related Work

[Blackwell 2010] proposes "a security ontology for incident analysis", but it sounds more like a taxonomy than an ontology. The article defines some concepts of incident, cyber-defense and attack, but does not formalize them as an ontology. The use of OWL is listed as a future improvement.

[Mundie et al. 2014] proposes an ontology based on an earlier work that defines a meta-model with essential Incident Management processes. The goal is to document, compare and analyze Computer Security Incident Response Teams (CSIRTs). The processes modeled are at a high level, more focused on the relationships and roles among teams, without going into the technical details of the incident.

[Silva and Fagundes 2014] proposes an ontology for incident management of Information Security based on ISO/IEC 27035:2011 and aim to assist in the training of incident response teams. It compares other works using as criteria: classification (vocabulary, type of ontology, taxonomy, etc.), references, whether the methodology for the construction of the ontology is documented and its public availability for the community. It presents eight superclasses and the properties of the *Incident* class, but it does not present any use cases. It also does not define the concept of "Event" although it is a fundamental part of ISO 27035. The ontology is no longer available in the repository indicated by the article, but it can be partially reproduced based on the illustrations.

[O'Sullivan and Turnbull 2015] introduces an "open source" ontology, available in github, that represents cybernetic assets with the purpose of modeling scenarios for simulation of defense of computer networks, increasing their resilience. It divides the ontology into two: Cyber Simulation Terrain (CST) which defines concepts such as computers, network connectivity, users, software, vulnerabilities and exploits; and Cyber Effects Simulation Ontology (CESO) which models the impacts of a cyber attack on systems that are part of a complex network. It is the most complete work among the reviewed ones, but the major focus of the use cases is the identification of vulnerabilities and the impact related to them. It does not consider security events nor temporal aspects.

[Syed et al. 2016] proposes an open ontology, available in github, with the objective of "incorporating and integrating heterogeneous data and knowledge schemes from different cybersecurity systems and commonly used standards for information sharing". The work has a huge goal and is presented as "the first ontology of cybersecurity to support a wide range of use cases". The paper presents four use cases, all of them related to the identification of software vulnerabilities similar to those presented by the Australian work [O'Sullivan and Turnbull 2015]. The proposal is coherent and has well-defined concepts, some of them reusable in the context of this work. However, it does not answer the competency questions proposed here, it does not define "security event," and it does not consider temporal aspects, but proposes this last point as a future improvement.

[Mavroeidis and Bromander 2017] consolidates a model with common definitions of *Cyber Threat Intelligence* and evaluates taxonomies, patterns and ontologies related to the concept. The model also aims to support organizations in measuring the maturity level of their capabilities to perform *Threat Intelligence*.

Except for [Syed et al. 2016] (Unified Cyber Ontology), which has a comprehensive usage proposal, none of the other referenced papers appears in the comparison pu-

Table 1. Related work comparison

Ontologies	Defines Incident	Defines Event	Temporal Aspects	OWL Format	Publicly Available	Competency Questions
[Blackwell 2010]	✓					-
[Mundie et al. 2014]	✓	✓				-
[Silva and Fagundes 2014]	✓					-
[O'Sullivan and Turnbull 2015]	✓				✓	45
[Syed et al. 2016]	✓			✓	✓	4
This work	✓	✓	✓	✓	✓	9 (3)

blished by [Mavroeidis and Bromander 2017], since its purpose is focused explicitly on Cyber Threat Intelligence. As observed in the present paper, the authors also note that many of the surveyed ontologies do not provide the relevant RDF/OWL files, even those that are "open source".

3. Computer Security Incident Handling Ontology (CSIHO)

In order to design an initial meta-model, the following fundamental requirements have been defined: an incident can only be characterized if there is **at least one related event**. Although an incident can be automatically identified, it is usually **validated by an analyst**, so an incident will always be related to a responsible analyst. The handling of an incident is divided into phases, and an analyst must define its current phase. An incident needs response actions, and these actions are defined and applied by one or more analysts. The analyst should also be responsible for reporting the **incident status**, for example, if it is in progress or has already been completed.

It is desirable to maximize knowledge about causes, consequences and origins of the incident since the moment of its confirmation, in addition to minimize the time between occurrence and detection of an event and also between discovery and completion of the incident. Aiming to address these issues, the model must implement a timeline with a historical record of all occurrences related to the incident, i.e., not only associated events and applied actions, but also observations on the results of the actions, relevant information obtained from external and internal sources, communications intra-team and inter-teams, among others.

From these preliminary definitions the first classes of the model were described:

ACTION: an activity performed to respond to the incident. Other ontologies and standards often refer to this concept as *Course of Action*.

ANALYST: a person somehow involved in the incident handling process.

EVENT: the identification of a possible violation in the information security policy and its controls, according to the definition of the standard [ISO/IEC 27035 2011]. In practice, they are logs of security tools, such as firewalls, host/network intrusion prevention/detection systems (IPS/IDS), anti-virus, Web servers, directory services (e.g. Active Directory), DNS, DHCP and so on. The object properties of this class will depend on the tool that was used to generate the record.

Although it is an essential concept in the context of incident response, none of the reviewed ontologies presents a definition for SECURITY EVENT.

PHASE: the possible phases of an incident response will be CONTENTAINMENT, ERADICATION and RECOVERY. The phase is not necessarily a class of the model, but it is desired to record each phase transition, which means, the date/time it was defined, as well as the analyst who determined it.

INCIDENT: the identification of one or more information security events that have a significant probability of compromising the operations of the business and threaten the security of the information, as defined by the [ISO/IEC 27035 2011] standard. Once the incident is confirmed, it should be documented appropriately and associated with other related events.

OCCURRENCE: allows one to document the historical record of all occurrences related to the incident in a timeline, whatever they are applications of defined response actions or any other relevant information. This is one of the most important contributions of the model since no other related ontology deals with temporal aspects. Each occurrence recorded by an analyst and must contain at least the date/time, incident phase and description attributes.

The primary tool adopted for the development of CSIHO was Protégé, a free software developed at Stanford University, widely used by the academic community [Mundie et al. 2014]. Protégé is in version 5.2 and is a highly extensible, open source environment that enables rapid prototyping and development. It has been actively supported by a community of users and developers for more than 17 years and supports the most current specification of the W3C, the Web Ontology Language (OWL) 2.

CSIHO was created "from scratch" reusing some relevant definitions from other ontologies but introducing new concepts. The data used to create the knowledge base came from a case study of an incident with the Wannacry ransomware [Moreira et al. 2017].

The following premises were established for the development of the ontology: the incident documentation should contain detailed logs, not just textual descriptions; a collaborative work mode should be foreseen, encouraging research in external entities; temporality is important, i.e., at which moment events and actions occurred.

The ontology was created using the *Ontology Development 101* methodology [Noy and McGuinness 2000], according to the steps described ahead. Figure 3 presents the CSIHO meta-model designed based on this process.

Step 1 - Determine the domain and scope of the ontology

CSIHO covers the incident handling macro-process domain in order to improve the processes for responding to cybersecurity incidents in organizations. It should provide a foundation for proper incident documentation and support the definition of response actions in different scenarios. It is hoped that the ontology will be used and maintained by researchers in the area of incident handling, incident response solution developers and incident response analysts. The following competency questions have been initially defined (non-exhaustive):

1. What types of events are related to a given incident?
2. What affected assets are related to a given incident?

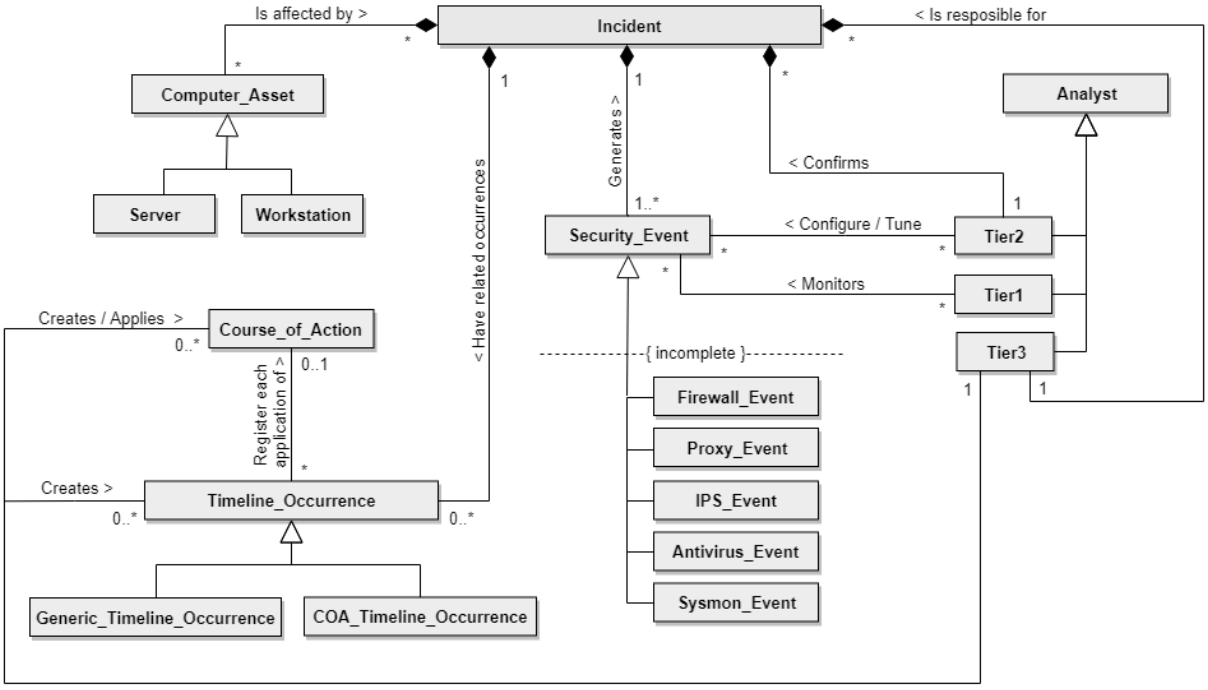

Figure 3. CSIHO meta-model

3. What are the incidents related to a particular asset?
4. What events are related to a given incident?
5. What are the response actions related to a given incident?
6. Given the characteristics of an event, which incidents were related?
7. Given a new incident, what previous actions could be reapplied?
8. How many actions were taken in each phase of a given incident?
9. What is the total time between the detection and conclusion of a given incident?

Step 2 - Reuse of existing ontologies

The following concepts were incorporated from the related works:

- **Event**: any occurrence that can have negative consequences for security [Mundie et al. 2014];
- **Incident**: event that has had confirmed negative security consequences [Mundie et al. 2014];
- **Course of Action**: Responsive or preventative actions to mitigate an attack [O'Sullivan and Turnbull 2015, Mavroeidis and Bromander 2017].

Step 3 - List important terms in the ontology

The following terms were listed based on the preliminary modeling: incident, event, action, phase, status, threat, vulnerability, preparation, identification, response, follow-up, eradication, containment and recovery.

Step 4 - Define the classes and their hierarchy

The following classes were initially defined:

- **Computer_Asset**: is the node in the defended network. In CSIHO, it is divided into two categories: *Workstation* and *Server*.
- **Course_of_Action**: is the incident response action.
- **Incident**: is the main class that represents the incident and is associated with all other elements of the ontology.
- **Person**: class that defines the analysts related to the incident. It is classified into three different levels: Tier1, Tier2 and Tier3.
- **Security_Event**: class that defines the security events associated with the incident. It is divided into subclasses according to the nature of the event: Antivirus_Event, Firewall_Event, IPS_Event, Proxy_Event and Sysmon_Event.
- **Timeline_Occurrence**: is the class that defines all the historical occurrences related to the incident (who, when, why). It is classified into two subclasses:
  - COA_Timeline_Occurrence: application record of response actions;
  - Generic_Timeline_Occurrence: Any other type of record related to the incident. For example: defining a plan of action, observations on the outcome of an action and relevant information obtained from internal and external sources.

Step 5 - Define class properties

The properties of the classes are presented in Table 2. Since the "Security_Event" class has a subclass for each type of security system, additional properties are included in accordance with the nature of the event.

Step 6 - Define the facets of each property

It is how the methodology refers to the definitions of cardinalities, data types (or range) and domain (one or more classes to which the property is related). Based on the preliminary definitions, the following cardinalities were defined:

- A **Security_Event** can be associated with 0 or 1 **Incident**;
- An **Incident** has at least 1 associated **Security_Event**;
- An **Incident** is confirmed by some **Tier3_Analyst**;
- An **Incident** has at least 0 **Course_of_Action**;
- A **Course_of_Action** is created by some **Tier3_Analyst**;
- A **Timeline_Occurrence** must be associated with an **Incident**;
- A **Timeline_Occurrence** must be entered by a **Tier3_Analyst**;
- A **Timeline_Occurrence** can be related to 0 or 1 **Course_of_Action**.

Step 7 - Create the instances (individuals)

The data used to initially populate the ontology were created from the case study on an incident with the Wannacry ransomware [Moreira et al. 2017]. The case history was organized chronologically in a table with identifiers and then transported to the ontology.

Table 2. Class properties (step 5) and data types (step 6)

Class	Properties	Data type (range)
Computer_Asset	hostname	string
	IP address	string
Course_of_Action	creation date/time	dateTime
	objective	string
	description	string
Person	name	string
	surname	string
	reg_id	string
	role	string
	function	string
	phone number	string
Security_Event	detection date/time	dateTime
	source system	string
	event type	string
	response	string
	source IP	string
	source hostname	string
	source port	integer
	target IP	string
	target hostname	string
	target port	integer
Incident	confirmation date/time	dateTime
	conclusion date/time	dateTime
	description	string
	type	string
	status	string
Timeline_Occurrence	occurrence date/time	dateTime
	incident phase	'Containment', 'Eradication' or 'Recovery'
	description	string

4. Ontology Evaluation

To demonstrate the effectiveness of CSIHO, we define some competency questions for the application of the ontology. These questions were mainly derived from the use case that was used for the initial load of the ontology, the Wannacry Incident, and were mapped to SPARQL queries for easy data handling and comprehension. For each query, it is assumed that the user made a previous and more comprehensive search to identify the object of interest, for example, obtaining a list of all registered incidents or searched for a specific keyword to identify the name assigned to a particular incident.

It is important to note that the competency questions presented, as well as the SPARQL queries built to answer them, are generalizable to any incident. To illustrate its application, however, we have used example records inspired mainly in the case study of the WannaCry ransomware incident [Moreira et al. 2017].

A total of nine competency questions were defined on the original dissertation, but due to space limitations three questions were selected for this paper.

Question 2) What are the affected assets related to a given incident?

The query, presented by the Listing 1, obtains the incident identifier “WannaCry 2017-05”, pinpoints the related events and then which assets were affected by each of these events. The desired attributes are then extracted from each asset and displayed. The results are shown in Table 3.

This information can support the response team in the activities during the incident treatment and also in the stages of “lessons learned” after the incident, such as root cause study, vulnerability analysis, among others.

```
PREFIX csiho: <http://anon.com.br/csiho.owl#>

SELECT ?hostname ?ip
WHERE {
    ?incident_id csiho:incident_name "WannaCry 2017-05" ;
                 csiho:has_event ?event_id .
    ?event csiho:has_affected_asset ?affected_assets .
    ?affected_assets csiho:ip_address ?ip ;
                     csiho:hostname ?hostname .
}
ORDER BY ASC(?affected_assets)
```

Listing 1: SPARQL query for the competency question 2

Table 3. Results from SPARQL query 2

hostname	ip
newton	10.0.51.110
euler	10.0.19.247
gauss	10.0.70.112
fermat	10.0.90.156
einstein	10.0.87.91
turing	10.0.20.125
neumann	10.0.74.188
pascal	10.0.48.254
fibonacci	10.0.16.80
hardy	10.0.9.108
nash	10.0.42.67
lovelace	10.0.2.102
hawking	10.0.92.190

Question 4) What events are related to a given incident?

This is an extension of the first query, including the attributes detection date, description, source network and source address for each listed event, as seen in Listing 2.

The WHERE statements form a design pattern where all conditions must be met for some information to return, that is, data must exist in all fields. In cases where a field may or may not contain data, as in the case of this query, you must use the “OPTIONAL” clause. The results are displayed in Table 4.

This query provides greater technical detail on events related to an incident, providing a better understanding of the attack techniques and strategies, and thus subsidizing response actions.

Question 5) What are the response actions related to a given incident?

```
PREFIX csiho: <http://anon.com.br/csiho.owl#>

SELECT (STR(?det_dt) AS ?detection_dt) ?event_type ?event_description ?source_network
  ↪  ?source_ip
WHERE {
  ?incident_id csiho:incident_name "WannaCry 2017-05" ;
               csiho:has_event ?event_id .
  ?event_id csiho:event_detection_dateTime ?det_dt ;
            rdf:type ?event_type .
  FILTER (?event_type != owl:NamedIndividual) .

  OPTIONAL { ?event_id csiho:event_description ?event_description . }
  OPTIONAL { ?event_id csiho:event_src_net ?source_network . }
  OPTIONAL { ?event_id csiho:event_src_address ?source_ip . }
}
ORDER BY ?detection_dt
```

Listing 2: SPARQL query for the competency question 4

Table 4. Results from SPARQL query 4

detection_dt	event_type	event_description	source_network	source_ip
2017-05-12T09:00:00	csiho:Sysmon_Event	Service STOPPED “h6hy63y2uhs”		10.0.1.51
2017-05-12T10:30:00	csiho:Sysmon_Event			
2017-05-12T11:45:00	csiho:IPS_Event		10.7.0.0/16	
2017-05-12T11:46:00	csiho:IPS_Event		10.12.0.0/16	

This query, presented by Listing 3, locates the incident identifier “WannaCry 2017-05” and obtains, from the “TimeLine_Occurrence” object, the response actions applied, retrieving its occurrence date and descriptions of the action and its application. In order to limit the volume of data returned only the actions applied in the CONTAINMENT phase were considered. The output is shown in Table 5.

This information can be used during the incident handling process to track actions already taken, or at a later time, for reporting, “lessons learned” exercises and identifying actions that can be reapplied in a new scenario.

```
PREFIX csiho: <http://anon.com.br/csiho.owl#>

SELECT (STR(?oc_dt) AS ?occurrence_dt) ?occurrence_notes ?coa_description
WHERE {
    ?incident_id csiho:incident_name "WannaCry 2017-05" .
    ?occurrence csiho:tl_occurrence_has_incident ?incident_id ;
                csiho:tl_occurrence_has_coa ?coa ;
                csiho:tl_occurrence_datetime ?oc_dt ;
                csiho:tl_occurrence_notes ?occurrence_notes ;
                csiho:tl_occurrence_incident_phase "Contenção" .
    ?coa csiho:coa_description ?coa_description .
}
ORDER BY ?occurrence_dt
```

Listing 3: SPARQL query for the competency question 5

5. Conclusion and Future Work

Advances in information technology have profoundly changed human relations, professions, the media, organizations and the economy, at the same time raised the dependency

Table 5. Results from SPARQL query 5

occurrence_dt	occurrence_notes	coa_description
2017-05-12T10:35:00	Applied coa_001	Isolate the network segment 10.0.1.0/24
2017-05-12T10:38:00	Applied coa_002	Request action for antivirus provider (vaccine/signature)
2017-05-12T10:51:00	Samples sent to antivirus vendor	Send samples of the binary from infected machines to antivirus vendor.
2017-05-12T10:51:00	Applied coa_003	Send samples of the binary from infected machines to antivirus provider.
2017-05-12T12:11:00	Blocking port 445 in connection to Subsidiary RJ1	Block port 445 (SMB) on edge firewalls on connections to Subsidiaries networks in case of suspect traffic.
2017-05-12T12:40:00	Blocking port 445 in connection to Subsidiary RJ2	Block port 445 (SMB) on edge firewalls on connections to Subsidiaries networks in case of suspect traffic.
2017-05-12T14:00:00	Blocked execution of known ransomware binaries, based on samples (hashes) collected and others obtained by the antivirus vendor.	Block execution of known ransomware binaries, based on samples (hashes) collected and others obtained by the antivirus vendor;
2017-05-12T15:33:00	New distribution of the patch that fixes the vulnerability in Windows SMB (bulletin MS17-010), with mandatory boot requested.	Request new distribution of the patch that fixes the vulnerability in Windows SMB (bulletin MS17-010), with mandatory boot.
2017-05-12T19:40:00	Distributed new “vaccine”to WannaCry released by antivirus vendor.	Distribute new “vaccine”to WannaCry released by antivirus vendor.
2017-05-13T11:00:00	Performed hosts survey: (i) Without SMB patches; (ii) With outdated antivirus (without the new vaccine).	Perform hosts survey: (i) Without SMB patches; (ii) With outdated antivirus (without the new vaccine).
2017-05-13T18:00:00	Disconnected (or powered off) stations with potential for contamination, identified by coa_008.	Disconnect (or shutdown) stations with potential for contamination, identified by coa_008.
2017-05-14T19:00:00	Implemented internal web server responding to WannaCry Kill Switch URLs (honeypot).	Implement internal web server responding to WannaCry Kill Switch URLs (honeypot).

of humanity on information systems and the Internet. Combined with the expansion of computing resources, this has led to a scenario of growingly complex and interconnected systems, expanding its exposure to risks and making its protection an increasingly challenging task.

Although the global perception is that the occurrence of incidents is almost inevitable, as the number of incidents continues to rise in frequency and severity, the literature shows that Information Security initiatives are more focused on detection and prevention than incident response, with many organizations often poorly prepared and ignoring key incident handling processes.

As a way to provide means to overcome these problems, the purpose of this work is to propose a model for handling incidents, inspired by the ideals of the Semantic Web and described in the form of an ontology. It provides a foundation for the incident handling process, in addition to enabling logical inferences and simplifying the process of transferring information and knowledge within a collaborative work context.

The Computer Security Incident Handling Ontology (CSIHO) was created based on simple principles, but with solid references. An example dataset was included with logs derived from a case study of an incident with the WannaCry ransomware. In order to support the evaluation of the ontology, nine general competency questions were defined. Following its collaborative philosophy, the ontology was published on GitHub for use by the professional community or development of future work by the academy **(https://github.com/moreiragb/csiho)**.

Thus, to demonstrate the contributions of the work as well as the fulfillment of the defined requirements, a group of experiments with specific scenarios were established. The experiments showed how the CSIHO ontology would allow a complete and efficient follow-up of an incident. Using simple SPARQL queries, it was possible to answer all the defined competency questions, and this was done using fewer human resources and considerably less time than those reported by the referenced case study [Moreira et al. 2017]. CSIHO is also the only ontology that defines and implements the concept of a "security event" as well as allows the historical recording of occurrences related to the incident, ensuring the construction of a timeline.

As of suggestions for future work, in order to contribute to the expansion of the ontology, new use cases can be included in its repository, as well as new competency questions can be derived from these new use cases.

Also, it is proposed to expand the subclasses of "events" and the development of connectors (parsers) for ingestion of logs of real-time security tools to explore scenarios of logical inference (reasoning) and the application of Artificial Intelligence techniques for automatic incident detection, which could also motivate the creation of an ontology dedicated to the treatment of events, connected to CSIHO.

Referências

Ab Rahman, N. H. and Choo (2015). A survey of information security incident handling in the cloud. *Computers & Security*, 49:45–69.

Baskerville, R., Spagnoletti, P., and Kim, J. (2014). Incident-centered information security: Managing a strategic balance between prevention and response. *Information & management*, 51(1):138–151.

Blackwell, C. (2010). A security ontology for incident analysis. In *Proceedings of the Sixth Annual Workshop on Cyber Security and Information Intelligence Research*, page 46. ACM.

CERT.br (2018). Estatísticas dos incidentes reportados ao cert.br - valores acumulados de 1999 a 2017. 03 abr. de 2018.

Cichonski, P., Millar, T., Grance, T., and Scarfone, K. (2012). Computer Security Incident Handling Guide : Recommendations of the National Institute of Standards and Technology. Technical Report NIST SP 800-61r2, National Institute of Standards and Technology. DOI: 10.6028/NIST.SP.800-61r2.

Computer World (2017). Maioria das empresas brasileiras não tem plano de resposta a incidentes. 27 mar. de 2017.

DuCharme, B. (2013). *Learning SPARQL: querying and updating with SPARQL 1.1*. O'Reilly Media, 2 edition.

F-Secure (2017). F-secure state of cyber security 2017. 17 fev. de 2017.

Grispos, G. (2016). *On the enhancement of data quality in security incident response investigations*. PhD thesis, University of Glasgow, Glasgow.

Healey, J. (2016). Winning and losing in cyberspace. In *International Conference on Cyber Conflict (CyCon)*, pages 37–49. IEEE. 08 fev. de 2017.

ISO/IEC 27035 (2011). *Information technology – Security techniques – Information security incident management*. ISO/IEC, Genebra, Suíça.

Jakus, G., Milutinović, V., Omerović, S., and Tomažič, S. (2013). *Concepts, ontologies, and knowledge representation*. Springer.

Kharraz, A., Robertson, W., Balzarotti, D., Bilge, L., and Kirda, E. (2015). *Cutting the Gordian Knot: A Look Under the Hood of Ransomware Attacks*, volume 9148, pages 3–24. Springer International Publishing, Milan.

Mavroeidis, V. and Bromander, S. (2017). Cyber threat intelligence model: An evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence. In *European Intelligence and Security Informatics Conference (EISIC)*, Athens, Greece. IEEE. 03 mar. de 2018.

Moreira, G. B., Calegario, V. M., Duarte, J. C., and dos Santos, A. F. P. (2017). A era dos crypto ransomwares: um estudo de caso sobre o wannacry. In *Simpósio Brasileiro em Segurança da Informação e de Sistemas Computacionais*, pages 509–516, Brasília. Sociedade Brasileira de Computação. 19 nov. de 2017.

Mundie, D. A., Ruefle, R., Dorofee, A. J., Perl, S. J., McCloud, J., and Collins, M. (2014). An incident management ontology. In *STIDS*, pages 62–71.

Noy, N. F. and McGuinness, D. L. (2000). Ontology development 101: A guide to creating your first ontology. 05 nov. de 2017.

O Globo (2015). Investimento em segurança da informação cresce mais no país - 2015. 08 fev. de 2017.

O'Sullivan, K. and Turnbull, B. (2015). The cyber simulation terrain: Towards an open source cyber effects simulation ontology. In *Australian Information Warfare Conference*, pages 14–23. Security Research Institute, Edith Cowan University. 05 nov. de 2017.

Shadbolt, N., Berners-Lee, T., and Hall, W. (2006). The semantic web revisited. *IEEE intelligent systems*, 21(3):96–101.

Silva, P. C. d. and Fagundes, L. L. (2014). Simo: Security incident management ontology. In *Simpósio Brasileiro em Segurança da Informação e de Sistemas Computacionais*, pages 302–305, Brasília. Sociedade Brasileira de Computação. 05 nov. de 2017.

Syed, Z., Padia, A., Finin, T., Mathews, M. L., and Joshi, A. (2016). Uco: A unified cybersecurity ontology. In *AAAI Workshop: Artificial Intelligence for Cyber Security*. 05 nov. de 2017.