An Incident Management Ontology

D. Mundie, R. Ruefle, A. Dorofee, J. McCloud, S. Perl, M. Collins
CERT®
Software Engineering Institute | Carnegie Mellon University
4500 Fifth Ave., Pittsburgh, PA, United States of America

***Abstract*—This paper outlines the need for and the development of an Incident Management Ontology. The Incident Management Ontology is derived from an Incident Management Meta-Model. We describe the shortcomings of the Incident Management Meta-Model and how the Incident Management Ontology addresses these shortcomings. The development of the Incident Management Ontology is outlined and the need for such an ontology is discussed. Related work is described and the Incident Management Ontology's potential uses and applications are presented.**

Keywords—Ontology, Incident Management, Description Logic

I. Introduction

When the JASON[1] Program within MITRE looked at the scientific community for ways to make cybersecurity "more scientific", their very first conclusion was that the security community needed "a common language and a set of basic concepts about which the security community can develop a shared understanding" [1], or in other words, a Cybersecurity ontology. The work described in this report is part of an ongoing effort within CERT® to build such an ontology for incident management.

We believe that such formal models are the best way for the community to evolve towards a "science of cybersecurity", and that our incident management ontology can play a crucial role in improving incident management. The ontology's purpose is to create a common language for describing the processes and functions associated with CSIRTs. We intend to use the ontology to analyze existing CSIRTs, to define a standard set of processes and services that should be offered by CSIRT teams, to formalize roles and responsibilities, and to build an ontology based competency model for the knowledge, skills, and abilities required of team members.

This paper describes the evolution of our work on characterizing incident security teams from a natural-language text document to a formal ontology and analyzes the benefits that accrued in the process. When creating our ontology, we chose to use the W3C Ontology Web Language - OWL[2] due to its formalism and increasing use in the Semantic Web community. We feel this work may be a useful case study for others who are thinking about formalizing their own information security knowledge.

[1] "JASON is an independent scientific advisory group that provides consulting services to the U.S. government on matters of defense science and technology. [In 2010] JASON was asked by the Department of Defense to examine the theory and practice of cyber-security, and to evaluate whether there are underlying fundamental principles that would make it possible to adopt a more scientific approach." (http://fas.org/irp/agency/dod/jason/)

II. The Incident Management Meta-Model

In previous work [2], we aggregated a wide variety of incident management process models such as ISO 27002 [3] and NIST 800-61 [4]. From those sources we abstracted a generalized meta-model that captured the essential processes involved in incident management.

This meta-model was at the heart of what we previously called an Incident Management Body of Knowledge (IMBOK). It broke incident management activities into 18 high-level tasks organized by the incident management life cycle phases as Prepare, Protect, and Respond. It also included five non-procedural, crosscutting capabilities that constrain all the other tasks. The following outlines the phases and tasks and 5 crosscuts of the IMBOK:

A. The phases and tasks

1) Prepare

- Develop trusted relationships with external experts
- Provide staff with appropriate education and training
- Develop policies, processes, procedures
- Measure incident management performance
- Provide constituents with security education, training, and awareness
- Develop an incident response strategy and plan
- Improve defenses

2) Monitor and Detect

- Assist constituents with correcting problems identified by vulnerability assessment activities
- Detect and report events
- Monitor networks and information systems for security
- Perform risk assessments and vulnerability assessments on constituent systems

3) Respond

[2] OWL is based upon description logics. OWL supports those users who want the maximum expressiveness while retaining computational completeness (all conclusions are guaranteed to be computable) and decidability (all computations will finish in finite time). (http://www.w3.org/TR/owl-features/)

- Triage Incident
- Collect and preserve evidence
- Restore and validate the system
- Perform a postmortem review of incident management actions
- Integrate lessons learned with problem management process
- Analyze incident, including artifacts, causes, and correlations
- Determine and remove the cause of the incident

B. The 5 crosscuts

1) Manage information

2) Properly handle collected evidence following best practices

3) Manage the incident management team

4) Communicate incidents

5) Track and document incidents from initial detection through final resolution

C. Drawbacks to the Incident Management Meta-Model

Although the Incident Management Meta-Model provides a considerable simplification and consolidation of prior knowledge, it suffers from a number of drawbacks due to its knowledge representation formalism:

- The use of imperative verb forms expressing infinitive constructions means that each task is only partially represented, because the subject is implicit. This obfuscates, for example, the fact that some of the tasks (e.g. managing the team) are carried out by the team's managers, not by the incident responders.
- In general, the use of natural language makes machine processing of this knowledge representation difficult.
- In particular, there is no easy way to use this representation to perform modeling and simulation, nor to build applications on top of it.
- To keep the process model manageable, concepts have been abstracted to an unusable level, with no graceful way to expand them into a more detailed form. There is no way within this system, for example, to say what is meant by "defenses" in "improve defenses".
- Apart from including a glossary, this representation does not facilitate the use of a standardized vocabulary.
- Also to keep the process model manageable, related concepts have been combined, as in "restore and validate the system".
- Despite its relative compactness, this representation violates the "7 plus or minus 2" law [5] and is hard for users to take in at a glance and internalize.

III. From Meta-Model to Ontology

Recently we realized that many of the drawbacks of the IMBOK could be remediated by moving beyond the informal natural-language format of the body of knowledge, and instead building a formal ontology using OWL.

A. Ontologies

An ontology is simply a set of shared, precisely-defined concepts in a given domain, along with the relationships among those concepts. OWL (the Web Ontology Language) is a W3C recommendation that builds on earlier languages from DARPA and elsewhere [6], is a key component of the Semantic Web [7], and is currently the leading knowledge representation and reasoning language in computer science. OWL is descended from earlier attempts at usable knowledge representation systems such as expert systems, logical programming languages, frame-based reasoning systems, modal logic, KL-One [8], entity-relationship modeling, and the like [9]. Description Logics emerged as a flexible yet powerful knowledge representation tool as the relationships among these approaches were better understood and new ways to engineer logics and reasoning systems were discovered. Description Logics have been used projects ranging from the International Catalogue of Diseases [10] to Google's Knowledge Graph [11].

To build our IM ontology, we decomposed the 18 high-level tasks in the IMBOK meta-model into component concepts and their respective relationships.The concepts, also known as classes in the Description Logic community, are organized into a strict hierarchy of subclasses. The incident management tasks are composed of relationships among those classes. This separation of classes from relationships is the key to most modern knowledge formalisms, from KL-One [8] to OWL [12].

B. N-ary Relationships

The only relationships inherent in the Description Logic on which OWL is built are binary relationships consisting of two concepts (or objects) and a relationship between them. However, many of the relationships we want to model in incident management are "n-ary" relationships among more than just two objects. For example, training requires a relationship among at least three objects: the training itself, a trainer, and a trainee. There are a number of ways to handle this situation in OWL; for the IM ontology we used one of the techniques recommended by the W3C [13]. This technique consists of creating a new class that holds the relationships among the training concepts.

This requires a slight adjustment to our ways of thinking about relationships. To illustrate, the original meta-model tasks

(IM leaders) Develop trusted relationships with external experts.

(trainers) Provide staff with appropriate education and training.

become

developing external relationships:
involves external groups
produces trusted relationships
is performed by IM leaders

staff training:
is provided by either external or internal trainers
is provided to IM personnel

Once the reified relationships are in place, it becomes straightforward to enhance them with additional information. In full, these two classes actually are as follows in the ontology:

developing external relationships:
belongs to the prepare process
involves external groups
produces trusted relationships
is subject to the incident management crosscuts
is performed by IM leaders

staff training:
is provided by either external or internal trainers
is provided to IM personnel
is a training service
is part of the prepare process
is subject to the incident management crosscuts

The table In Appendix B gives a simple summary of the relationships in the ontology.

Figure 1 shows a screenshot of the IM ontology being edited in Protégé [14], the ontology development tool from Stanford that is widely used in the community. The display contains five panes giving five views of the ontology. The upper left pane shows the class hierarchy. The two most important classes are "activities" and "crosscuts". The activities are simply the tasks carried out by the incident management staff, while "crosscuts" or "principles" as Beebe and Clark call them [15] are pervasive constraints on the activities. In addition to those main classes, we needed eight auxiliary classes to describe the activities in full: incident components, IT components, knowledge assets, life cycle phases, organizational groups, quality standards, relationships, and team resources. These classes were identified using traditional ontology-mining techniques: we started with the terms in the meta-model, then clustered them and introduced class hierarchies based on our knowledge of the domain.

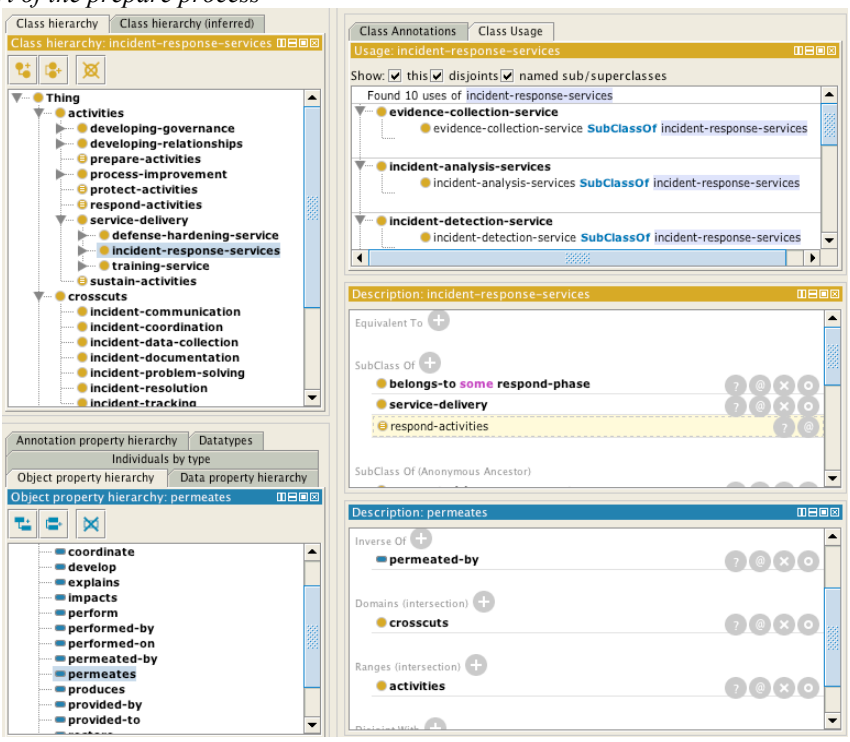

Figure 1 - The Incident Management Ontology Being Edited in Protégé. Note that the "crosscuts" class has grown and its members renamed during the development process.

The top two right-hand panes of the display show additional information about the selected class in the class hierarchy, in this case "incident response services". The top pane shows the usage of the selected class, while the second pane shows information about the class in terms of its subclasses, its superclasses, its members, any equivalent classes, and so forth.

The pane at the lower left of the screen shows the hierarchy of relationships, called "object properties" in

OWL. The "permeates" relationship has been selected. The lowest pane on the right describes that relationship in the ontology, showing that its domain is "CSIRT managers" and its range is "team resources", capturing the fact that CSIRT team managers acquire the team's resources.

C. *Overcoming the Drawbacks*

We believe that this formal IM ontology solves the problems noted in Section 2 for the IMBOK meta-model.

- The use of classes and relationships ensures that the knowledge is represented completely.
- This representation is machine-processable; Figure 2 shows a simple graphic automatically generated from the IM ontology using the OntoGraf tool [16] with a GraphViz post-processing script.
- The use of Description Logic (DL) ontologies for modeling and for constructing applications is well understood [9].
- The use of a strict class hierarchy gives us a user-friendly way to talk about concepts at any needed level of abstraction without complicating the IM ontology as a whole: we can talk about "security training", or "training", or "proactive services", or "incident management services", and the reasoning system will infer properties and type relationships as needed.
- The use of OWL annotations to capture definitions makes the IM ontology usable as a dictionary.
- Because of the class hierarchy and the formality of the system, there is no pressure to collapse concepts to keep the document small.
- Finally, the separation of entities from relationships reduces the complexity of the representation, and makes the structure of the IM ontology easier to absorb.

Figure 2 gives a high-level breakdown of the incident management activities. The "service delivery" activities are the most important, and Figure 2 expands that class to a further level of detail. Figure 3 shows a close-up of the root cause analysis environment, showing that it is performed by incident management personnel, that its goal is to explain root causes, that it is an incident analysis service, and so forth.

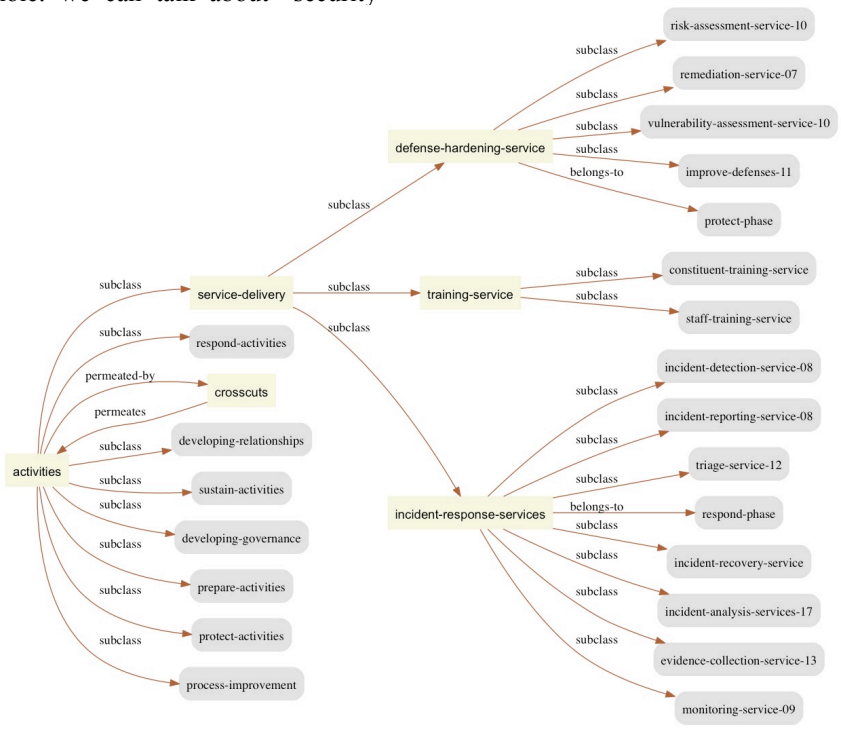

Figure 2 - The Activity Classes in the Ontology, with the Service-Delivery Activity Expanded

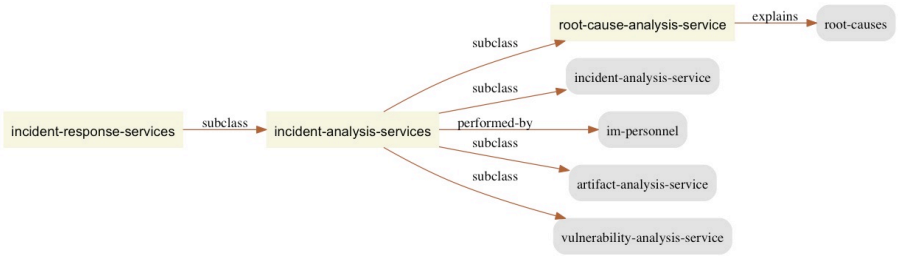

Figure 3 - A Detailed View of Root Cause Analysis

D. Other Benefits

In addition to solving the difficulties we had with the Meta-Model, moving to a formal ontology had several other advantages.

1) Very flexible typing

We quickly grew fond of the ability to create new types simply by specifying the necessary and sufficient conditions for membership in the type. Earlier we had used a multidimensional organization system called facet maps [17] to achieve multiple categorizations for the Meta-Model, but class expressions are much more lightweight and flexible. They are like a very disciplined tagging system. To cite just one example, we realized at some point that although we want to retain the classification of activities by the life-cycle phase in which they are used (prepare, protect, detect, respond, etc.), there is no need to build the life-cycle phases into the class hierarchy. Instead we simply assert a "belongs-to" relationship between an activity and a life-cycle phase. Then we can define a "protect-activities" class where the membership condition is "an activity that belongs-to the protect phase" and the reasoner will automatically compute the members of the class.

2) More powerful Modeling

The n-ary relations that use binary relations to "reify" relations among individuals turned out to be a very effective method for packaging up domain knowledge in a taxonomic hierarchy. When it seemed clear that the different types of incident analysis were characterized by the goal of their analysis, it was trivial to add "explains" and "explained-by" relationships.

3) Improved knowledge visualization

A shortcoming of our Incident Management Meta-Model was the absence of a satisfactory visualization. After converting the Meta-Model into a formal ontology, we used OntoGraf [16] to export files in the GraphViz DOT format [18]. DOT is a text-based format that allows for customizable graphics.

E. Individuals

The real power of Description Logic ontology comes when an ontology is populated by individuals and reasoning is enabled. "Reasoning" is a key-functionality of semantic technologies and allows automatic inferences to be made using the rules and classes described by the ontology. The ability of OWL to be used at internet scale comes from the highly optimized and logically precise handling of both terminological, or *taxonomic*, knowledge in what the Description Logic community calls the TBox, and the contingent *assertional* knowledge about individuals in what the community calls the ABox [9].

We have not yet formally extended the Incident Management Ontology to real world individuals, but Figure 4 shows an example using two fictitious individual CSIRTs in the ontology. The Acme team, focused only on incident response, provides monitoring, incident detection, incident reporting, and incident analysis services. The National Team from Borduria focuses on vulnerability assessment, vulnerability analysis, incident analysis, performance measurement, and relationship building. As the diagram makes clear, the only service these two CSIRTs have in common is incident analysis.

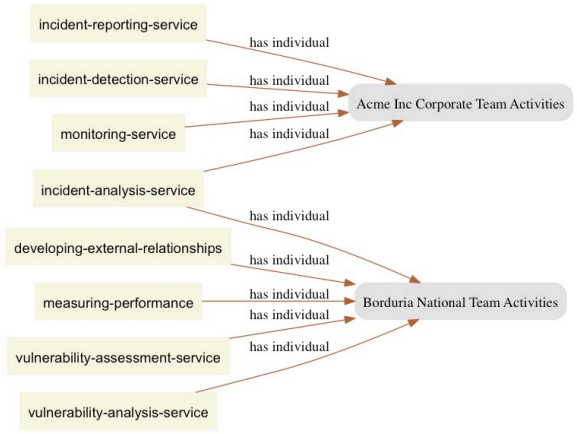

Figure 4 - A Comparison of Two Fictitious Incident Management Teams. Note that "has individual" denotes the membership of individuals in classes. Thus Acme is an individual member of the class of incident reporting services.

IV. Related Research

The seminal paper Formalizing Information Security Knowledge by Fenz et al. [19] presents the rationale for capturing information security best practices in an OWL ontology. Though it served as an influence for our ontology, the work addresses information security in general while our work focuses on incident management.

There have been many proposals for standardized incident handling process models; for a summary of the models that were used for our meta-model, see [2]. Although they incorporated much collective wisdom, none of them were based on a formal knowledge representation. Like our meta-model, the forensic framework of Beebe and Clark [15] aimed to assimilate existing practice into a comprehensive framework. The distributed, loosely-coupled incident response model of Millar, Osorno, and Reger [20] is a deeply-reasoned attempt to analyze and improve upon existing incident management practices based on scientific theory and simulation, but is not based on a formal ontology.

Furthermore, we found that many proposed ontologies that exist fail to capture all of the important relationships between members of organizations and the organizations themselves. These representations arise from an internal focus of an organization who has been victim to attack, and many ignore the roles and relationships between a CSIRT and incidents that occur.

Magklaras and Furnell [21] observe that incidents occur through misuse by individuals, but do not propose a formalized ontology of a taxonomy including this human-misuse concept. Classifications of individuals are made more

distinguished based on behavior (e.g. accidental or intentional), and possible consequences of misuse correlated to these actions.

Wang and Guo's [22] research in developing OVM (Ontology for Vulnerability Management) identifies individuals responsible for attacks, but the relationships amongst these individuals is not made clear. The formalizations within their work capture knowledge sufficient to answer questions about the assets targeted in an incident and mechanisms by which an incident takes place. While organization and individuals are clear in this work, further subdivisions of organizations and groups of individuals are not. No concept of trust appears in the ontology's class hierarchy, making the risk of agents difficult to reason about. Chiang [23] proposed mapping the IT Security EBK [24] and ISO/IEC 27001 [25] standard to an incident ontology. The construct is similar to OVM, but has the benefit of subdivision of roles amongst individuals and groups. Subdivisions, however, are limited and the ontology will require additional, higher-level concepts to subsume various sibling classes of the hierarchy.

The most complete formalization framework in security that gathers all necessary information to incident management might be Ekelhart's [26] move from simple security taxonomy to ontology. This work acknowledges the different threats and means for attacks, along with measurable reductions when safeguards are introduced. Even relationships amongst individuals in an organization and the roles they take are represented clearly. However, this research does not model subdivisions of an organization and the roles that multiple organizations can have (both within and in relation to one another). Different subdivisions of service types and measures of trust are not represented.

V. Next Steps

Future work on the Incident Management Ontology will focus on evaluting the ontology and using it to categorize incident response organizations. This work names CSIRT processes but does not yet describe them in full detail. Future work may include using existing standards, such as the Process Specific Language[27], to model the process flows for each service offered by a CSIRT in greater detail. We plan to evaluate the ontology by using it to analyze the processes performed by and services offered by incident response teams. A hypothesis we would like to test is whether there is a difference between the functions of CSIRTs and Coordination Centers. We are collecting data on both types of organizations and plan to analyze it using the ontology. We also plan to improve the ontology by adding axioms, more defined classes, and taking more advantage of reasoning capabilities.

VI. Acknowledgements

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. This material has been approved for public release and unlimited distribution. Carnegie Mellon® and CERT® are registered marks of Carnegie Mellon University. DM-0001433

References

1 McMorrow, D.: 'Science of Cyber-Security', in Editor (Ed.)^(Eds.): 'Book Science of Cyber-Security' (DTIC Document, 2010, edn.), pp.
2 Mundie, D.A., and Ruefle, R.: 'Building an Incident Management Body of Knowledge', in Editor (Ed.)^(Eds.): 'Book Building an Incident Management Body of Knowledge' (Citeseer, 2012, edn.), pp. 507-513
3 ISO, I., and Std, I.: 'ISO 27002: 2005', Information Technology-Security Techniques-Code of Practice for Information Security Management. ISO, 2005
4 NIST: 'Special Publication 800-61, Revision 2', Computer Security Incident Handling Guide, 2012, pp. 800-861
5 Miller, G.A.: 'The magical number seven, plus or minus two: some limits on our capacity for processing information', Psychological review, 1956, 63, (2), pp. 81
6 Motik, B., Patel-Schneider, P.F., Parsia, B., Bock, C., Fokoue, A., Haase, P., Hoekstra, R., Horrocks, I., Ruttenberg, A., and Sattler, U.: 'Owl 2 web ontology language: Structural specification and functional-style syntax', W3C recommendation, 2009, 27, (65), pp. 159
7 Hitzler, P., Krötzsch, M., Parsia, B., Patel-Schneider, P.F., and Rudolph, S.: 'OWL 2 web ontology language primer', W3C recommendation, 2009, 27, (1), pp. 123
8 Brachman, R.J., and Schmolze, J.G.: 'An Overview of the KL - ONE Knowledge Representation System*', Cognitive science, 1985, 9, (2), pp. 171-216
9 Baader, F.: 'The description logic handbook: theory, implementation, and applications' (Cambridge university press, 2003. 2003)
10 Organization, W.H.: 'International classification of diseases (ICD)', 2012
11 Singhal, A.: 'Introducing the knowledge graph: things, not strings', Official Google Blog, May, 2012
12 Antoniou, G., and Van Harmelen, F.: 'Web ontology language: Owl': 'Handbook on ontologies' (Springer, 2004), pp. 67-92
13 Noy, N., Rector, A., Hayes, P., and Welty, C.: 'Defining n-ary relations on the semantic web', W3C Working Group Note, 2006, 12, pp. 4
14 Ontology, P.: 'Knowledge Acquisition System', See http://protege. stanford. edu, 2007
15 Beebe, N.L., and Clark, J.G.: 'A hierarchical, objectives-based framework for the digital investigations process', Digital Investigation, 2005, 2, (2), pp. 147-167
16 http://protegewiki.stanford.edu/wiki/OntoGraf2014
17 facetmap.com/2014
18 Ellson, J., Gansner, E., Koutsofios, L., North, S.C., and Woodhull, G.: 'Graphviz—open source graph drawing tools', in Editor (Ed.)^(Eds.): 'Book Graphviz—open source graph drawing tools' (Springer, 2002, edn.), pp. 483-484
19 Fenz, S., and Ekelhart, A.: 'Formalizing information security knowledge', in Editor (Ed.)^(Eds.): 'Book Formalizing information security knowledge' (ACM, 2009, edn.), pp. 183-194
20 Osorno, M., Laurel, M., Millar, T., Team, E.R., and Rager, D.: 'Coordinated Cybersecurity Incident Handling', in Editor (Ed.)^(Eds.): 'Book Coordinated Cybersecurity Incident Handling' (2011, edn.), pp.
21 Magklaras, G., and Furnell, S.: 'Insider threat prediction tool: Evaluating the probability of IT misuse', Computers & Security, 2001, 21, (1), pp. 62-73
22 Wang, J.A., and Guo, M.: 'OVM: an ontology for vulnerability management', in Editor (Ed.)^(Eds.): 'Book OVM: an ontology for vulnerability management' (ACM, 2009, edn.), pp. 34
23 Chiang, T.J., Kouh, J.S., and Chang, R.-I.: 'Ontology-based Risk Control for the Incident Management', IJCSNS International Journal of Computer Science and Network Security, 2009, 9, (11), pp. 181-189
24 Division, O.o.C.a.C.N.C.S.: 'Information Technology (IT) Security Essential Body of Knowledge (EBK): A Competency and Functional Framework for IT Security Workforce Development', 2007
25 Lambo, T.: 'ISO/IEC 27001: The future of infosec certification', ISSA Journal, Information Systems Security Organization (http://www. issa. org), 2006
26 Ekelhart, A., Fenz, S., Klemen, M., and Weippl, E.: 'Security ontologies: Improving quantitative risk analysis', in Editor (Ed.)^(Eds.): 'Book Security ontologies: Improving quantitative risk analysis' (IEEE, 2007, edn.), pp. 156a-156a
27 Schlenoff, C., Gruninger, M., Tissot, F., Valois, J., Lubell, J., and Lee, J.: 'The process specification language (PSL) overview and version 1.0 specification' (Citeseer, 2000. 2000)

APPENDIX A: AN OVERVIEW OF THE ACTIVITIES CLASS IN THE ONTOLOGY

- Incident management (IM) leaders develop trusted relationships with external groups
- Both internal and external trainers provide training to IM personnel
- internal trainers provide awareness training to partners
- IM leaders develop governance artifacts
- IM leaders perform management functions on IM personnel
- IM leaders develop planning artifacts
- IM personnel provide vulnerability remediation to constituents
- IM personnel provide incident detection to constituents
- IM personnel provide incident communication to constituents
- IM personnel provide defense hardening to constituents
- IM personnel perform triage
- incident data collectors perform incident data collection
- IT personnel restore IT components
- IT personnel validate IT components
- IM personnel coordinate analyzing lessons learned
- IM incident handlers perform incident analysis
- IM personnel perform incident resolution
- IM leaders perform management functions
- IM personnel perform incident tracking

APPENDIX B: THE CLASS HIERARCHY OF THE IM ONTOLOGY

This appendix contains the class hierarchy in the Incident Management Ontology.

- **activities**: functions performed by a CSIRT
  - **developing-governance**: establishing the operational guidelines for an organization
    - **developing-plans**: establishing and maintaining the business and operational plans for an organization
    - **developing-policies**: establishing and maintaining the policies that guide the organizational activities
    - **developing-procedures**: establishing and maintaining implementations of organizational policies
    - **developing-processes**: establishing and maintaining organizational processes
    - **develop-data-collection-processes**: establishing logs and monitoring to provide insight into incidents
  - **developing-relationships**: identifying and communicating with essential business partners
    - **developing-external-relationships**: developing relationships with external parties
    - **developing-internal-relationships**: developing relationships with internal parties
  - **prepare-activities**: activities that are typically carried out during the prepare phase of the incident life cycle
  - **process-improvement**: activity whose goal is to improve the efficiency, reproducibility, reliability, or other quality attribute of business processes
    - **integrating-lessons-learned**: feeding the results of a postmortem review into the organization's problem-solving process
    - **postmortem-review**: an examination of an event to discover factors that affected the quality of the handling of the event
    - **measuring-performance**: collecting metrics that assess the quality of a process for process improvement purposes
  - **protect-activities**: activities that are typically carried out during the protect phase of the incident life cycle
  - **respond-activities**: activities that are typically carried out during the respond phase of the incident life cycle
  - **service-delivery**: the activity of providing a service to a constituent
    - **defense-hardening-service**: assisting with improving the security defenses of a constituent
      - **improve-defenses:** hardening defenses by improving the security controls in place
      - **remediation-service**: hardening defenses by removing known vulnerabilities and risks
      - **risk-assessment-service**: hardening defenses by identifying threats
      - **vulnerability-assessment-service:** hardening defenses by identifying vulnerabilities
    - **incident-response-service:** providing assistance in responding to and recovering from incidents
      - **evidence-collection-service:** gathering and maintaining information concerning an event
        - **diagnostic-data-collection-service:** incident-data-collection to support diagnosis and restoration activities
        - **forensics-data-collection-service:** incident-data-collection to support legal activities
      - **incident-analysis-services:** using collected data to uncover the causes and time-line of an event
        - **artifact-analysis-service:** incident analysis applied to artifacts
        - **incident-analysis-service:** general incident analysis
        - **root-cause-analysis-service:** incident analysis with the goal of determining the root cause of an event
        - **vulnerability-analysis-service:** incident analysis applied to the vulnerability that enabled an event
      - **incident-detection-service:** proactive steps to ensure events and incidents are discovered and reported as soon as possible
      - **incident-recovery-service:** reactive activities with the goal of restoring an affected system to the state before an event
        - **system-restoration-service:** restoring an affected system to the state before an event
        - **system-validation-service:** verifying that an affected system has been restored
      - **incident-reporting-service:** communicating information about an event or incident in accordance with an incident reporting policy
      - **monitoring-service:** maintaining an automated infrastructure to detect events and report incidents
    - **training-service:** a proactive service to ensure that stakeholders have the knowledge, skills, and abilities they need
      - **constituent-training-service:** training for constituents that helps them protect their infrastructure
      - **staff-training-service:** training for staff that helps them perform their jobs
        - **team-training-coordination:** ensuring adequate training for staff

sustain-activities: activities whose goal is to prevent the CSIRT's posture from declining over time

crosscuts: constraints or principles that apply to activities

incident-communication: communicating information about the effects of a incident to staff and constituents

incident-coordination: ensuring that all IM stakeholders are with a shared plan

incident-data-collection: collection of data relevant to an incident

incident-documentation: documenting the results of incident-analysis

incident-problem-solving: using generic or specialized methods in an orderly manner to find solutions to problems

incident-resolution: an action taken to repair the root cause of an incident or to implement a workaround

incident-tracking: managing and maintaining a database of information on incidents and constituents

incident-components: the various elements that constitute the conceptual model of an event

artifacts: any entities left behind after an incident takes place; for example, malicious code or logfiles

events: any occurrences that may have negative security consequences

incidents: events that have been confirmed to have negative security consequences

root-causes: the earliest occurrence in the causal chain leading to an incident

vulnerabilities: the weaknesses in the system that were exploited by an incident

IT-components: the various elements that constitute the conceptual model of an IT system

information-system: collection of technical and human resources that provide storage, computing, and distribution for enterprise information

network: collection of host computers together with the sub-network or inter-network through which they can exchange data

security-tools: hardware and software that improves the security of the information-system in which they are installed

incident-detection-tools: security-tools that perform incident-detection

av-systems: incident-detection-tools that work by analyzing virus signatures

ids-systems: incident-detection-tools that work by analyzing activity on the network

network-monitors: security-tools that work by observing network activity

knowledge-assets: the various types of documents that constitute the intellectual capital of the organization

governance-artifacts: documents that are used in the process of governing

policies: abstract documents that express decisions made by management about the running of the organization

procedures: concrete documents that implement policies

processes: workflows that implement policies and procedures

incident-reports: documents that inform the CSIRT about events and incidents

incident-tracking-documents: case management documents that trace the progress of an event through the incident-handling process

incident-assignments: tagging of incidents with the names of IM-personnel responsible for handling them

incident-categorization: tagging of incidents with the classification into which they fall

information: general documents that do not fall in any other category

lessons-learned: documents that capture the results of analyzing-lessons-learned

other-knowledge-assets: any information not included in other categories

planning-artifacts: abstract documents that prepare IM-personnel for incident response

incident-response-plans: planning-artifacts that reflect decisions made about incident-response within the organization

incident-response-strategies: technical documents that guide IM-personnel in responding to incidents

training-materials: documents that are used to provide training

life-cycle-phase: the temporal periods into which incident response is divided

prepare-phase: educating personnel and providing them with the tools needed to perform their jobs

protect-phase: applying controls and otherwise hardening the infrastructure to resist attack

respond-phase: detecting, analyzing, and recovering from incidents

sustain-phase: ensuring that the capability of the CSIRT does not degrade over time

organizational-groups: stakeholders in the incident management process

external-groups: stakeholders not within the administrative boundaries of the organization

external-csirts: incident management teams outside the boundaries of the organization

external-trainers: educational personnel outside the organization

law-enforcement-agencies: external groups performing law enforcement functions

other-external-groups: any other external group

partners: groups or sets of individuals with close relationships to the organization

constituents: the groups or sets of individuals for whom incident management is being performed

staff: stakeholders contained within administrative boundaries of the organization

IM-personnel: groups or sets of individuals tasked with performing incident management

IM-incident-handlers: individuals responsible for responding to and recovering from incidents

IM-forensics-analyst: an IM-incident-analyst specializing in analysis for legal purposes

IM-incident-analyst: an IM-incident-handler specializing in the analysis of incident-components

IM-malware-analyst: an IM-incident-analyst specializing in reverse engineering

IM-leaders: individuals responsible for leading the incident management personnel

incident-data-collectors: individuals responsible for collecting data about incidents

diagnostic-data-collectors: incident-data-collectors that collect data for diagnostic purposes

forensic-data-collectors: incident-data-collectors that collect data for forensic purposes

internal-trainers: educational personnel within the organization

IT-personnel: members of the it staff that carry out security functions such as infrastructure hardening

management: individuals responsible for governing

line-management: managers at the low end of the chain of command

mid-level-management: managers in the middle of the chain of command

senior-management: managers at the high end of the chain of command

risk-assessors: individuals responsible for assessing risks to the organization

vulnerability-assessors: individuals responsible for identifying vulnerabilities in the organization's infrastructure

quality-standards: normative requirements for ensuring the high quality of the CSIRT's activities

document-management-standards: standards that constrain the way information is handled within the organization

appropriately-dissemination-standards: standards that govern the provision of information to the appropriate audiences

confidentiality-preserving-standards: standards that govern how information is withheld from inappropriate audiences

forensic-standards: quality standards that ensure the admissibility of the analysis in a court of law

preserving-chain-of-custody: documenting that there has been no opportunity for forensic evidence to be tampered with

other-quality-standards: quality standards not included in other categories

relationships: connections between individuals or groups of individuals

trusted-relationship: relationships among entities that are willing to share confidential data

untrusted-relationship: relationships among entities that are willing to share confidential data

team-resources: anything needed for the CSIRT activities or the operations of IM-personnel

funding: financial resources necessary for the operations of IM-personnel

IT-infrastructure: information security assets necessary for the operations of IM-personnel

staffing: human resources necessary to ensure the operations of IM-personnel