

Technical Report: Creating a Preliminary Cyber Ontology for Insider Threats in the Financial Sector *

Gökhan Kul
Department of Computer Science and
Engineering
The State University of New York at Buffalo
Buffalo, New York 14260
gokhanku@buffalo.edu

Shambhu Upadhyaya
Department of Computer Science and
Engineering
The State University of New York at Buffalo
Buffalo, New York 14260
shambhu@buffalo.edu

ABSTRACT

Insider attack has become a major threat in financial sector and is a very serious and pervasive security problem. Currently, there is no insider threat ontology in this domain and such an ontology is critical to developing countermeasures against insider attacks. In this paper, we create an ontology focusing on insider attacks in the banking domain targeting database systems. We define the taxonomy used in this ontology and identify the relationships between the ontology classes. The resulting structure is a domain ontology mapped onto SUMO, FOAF and Finance ontologies to make the our work integrable to the systems that use these ontologies and to create a broad knowledge base. The attack types we formulate in the ontology are masquerade, privilege elevation, privilege abuse and collusion attacks. Our model could be used to systematically evaluate any insider threat detection schemes in a realistic way and discover attacks that share similarities with previously identified attacks.

Keywords

Cyber ontology; financial sector; relational database systems; taxonomy

1. INTRODUCTION

Most security systems are built to protect from external threats. Networked systems and information technology systems are changing with rapid innovations and they have been claiming more crucial roles in critical infrastructures. This rapid development has made more apparent the issue of information security due to the information contained in these systems.

Insider attacks are becoming an extremely serious security problem for financial institutions due to the threat they pose

*Date published: 07/29/2015

to the monetary assets and the sensitive customer data they handle. The threat that leads to insider attack is called an insider threat and the RAND report [1] addresses insider threats as “malevolent (or possibly inadvertent) actions by an already trusted person with access to sensitive information and information systems.”

Due to the nature of the banking domain, most of the employees like tellers and customer representatives can access very sensitive information. An attack that is coming from an employee can go unnoticed for a very long time [2]. There may be multiple insider attacks consequently within an organization with either the same or different intentions. According to 2014 U.S. State of Cybercrime Survey [3], 37% of organizations have experienced an insider incident, and 32% of the respondents to the research conducted say that damage caused by insider attacks is more damaging than outsider attacks. In 82% of these cases private or sensitive information was unintentionally exposed, 76% of incidents confidential records were compromised or stolen. In 71% and 63% of these incidents, respectively, customer and employee records were compromised or stolen. It is expressed in this report that 28% of electronic crime events are known or suspected to have been caused by insiders and in 46% of electronic crimes, insider attacks were more costly or damaging to their organization. The report also shows that 75% of cases were handled internally without legal action or law enforcement, mostly because of lack of evidence or not enough information to prosecute. Only 10% of cases were handled internally with legal action and 12% of the cases were handled externally with notification of law enforcement while only 3% of cases were handled externally by filing a civil action. This raises the questions on the reliability of security systems toward identifying insider threats.

We focus on insider attacks on relational database management systems for a variety of reasons. First, keeping the focus on a specific but important domain allows us to contain the scope of the model to a more manageable level. Second, even though there are other data preservation techniques and systems, relational databases are heavily used in back-end servers to store financial data, which consists of a lot of sensitive information. This makes relational databases a primary target for cyber crime and identity theft, namely, criminal activity. The aim of this effort is to develop an ontology of this area, expressed in the Web Ontology Lan-

guage (OWL) that ensures integration with other knowledge domains and enables data integration across different data sources. Semantic web applications are becoming more popular by the day and ontology is the most important enabling technology of these applications. Basically, it describes terms and different relationship types between terms. In this paper, we create a taxonomy of insider threats and identify the relationships between the entities we define in the taxonomy. These entities and relationships are used to create an insider threats ontology which is then mapped onto upper ontologies and domain ontologies that are commonly used in financial systems. The contribution of this work is both creating a framework of a cyber ontology for insider threats in the financial sector focusing on relational database management systems, and integrating this ontology with commonly used ontologies SUMO [4] [5], FOAF [6] [7] and Finance [8] to make it applicable and integrable to the systems that use these ontologies.

Section 2 reviews related work in the literature. Section 3 creates a taxonomy on insider threats and Section 4 explains insider threat scenarios. Section 5 gives the required information to create the knowledge base and for evaluation of the ontology. Section 6 discusses the advantages and contributions of our research, and finally Section 7 presents the future work.

2. RELATED WORK

The prior research in this area is along two different directions. One considers psychological aspects and the other considers physical aspects of insider threats. The phrases “insider” and “insider threats” are terms that have ambiguous definitions, but are known to many for what they mean. Most of the research which has been done on insider threats is mainly on the psychological structure and incentives of these attacks and how to prevent them on general cases. There are cases in which researchers have focused on physical threats from insiders. For example, the work performed on [9] focuses on how to protect datacenters from physical attacks and insider threats. Some instances may include a recently terminated employee, a user on a computer that is logged in, or even a janitor. No matter who the insider is, the potential threat to an organization is a problem that many organizations need to account for.

Hunker and Probst go into detailing what exactly an insider and insider threat are, while giving examples of solutions to the problem of insider threats [10]. They give definitions for “insider”, “insider threat”, as well as detailing issues that arise when managing insider threats and the lack of data on the topic while describing the multiple approaches to an insider attack. They describe the technical and socio-technical approaches to dealing with insider attacks and discuss the sociological, psychological, and organizational approaches to dealing with insider attacks. The authors explore the wide range of different people that can be insiders and they describe all of the aspects that go into making someone an insider. This level of detail is carried into the description of an insider attack, showing how there can be different types including accidental threats as well as malicious attacks.

Mathew et al. state that Insider attacks pose a serious threat due to the fact that current security systems are aimed

at prevention of unauthorized access [11]. They focus on the fact that not only can threats come from trusted entities within an organization, but a successful attack may be the result of multiple entities working together, termed insider collusion. Therefore, this is said to justify a call for monitoring and detection methods which take into account these potential interactions between entities. From here, they go on to detail the use of a new system, called ICMAP (Information-Centric Modeler and Auditor Program), which generates CAGs (Capability Acquisition Graphs) to represent information about physical locations of data, difficulty of access to components of the data system, etc. This graph allows for feasible analysis of which paths to insider abuse targets are the least difficult to traverse. The CAG holds information about the potential difficulty of accessing certain nodes in the system, and can therefore determine the path of least resistance. This can allow for security analysts to bolster the defenses of the systems along that path, or simply to monitor activity along these nodes for suspicious behavior. It is noted that the cost of creating, updating, and analyzing a CAG is considerably high and thus impractical to maintain in real-time. The proposed solution is to only update the CAG periodically (termed “CAG milestones”), as well as search for paths vulnerable to attack using a greedy algorithm that may not give the absolute most vulnerable path in the system, but is likely to after a number of runs. They provide an example of a situation in which a collusion attack could be carried out undetected, with malicious activity performed under the guise of being legitimate work tasks. Therefore, such a scenario would be difficult to catch in the act. However, a CAG generated by ICMAP can trace the means through which somebody with only public access could obtain information with top-secret security restrictions.

In their paper [12], Costa et al. detail their creation of an ontology for use in describing the indicators of insider threats. The primary reason cited for focusing on this area is that it had been uncommon for information about these insider threats to be circulated outside of the businesses that were typically subject to them; without a standardized method to abstract the data, doing so would have meant releasing confidential data related to the attack. Without public circulation of this information, progress in determining methods to prevent these insider threats has been severely hampered despite increasing focus in this area of research. The ontology was developed with the aid of over 800 cases of malicious insider activity compiled from various sources, all of which were natural language descriptions of the incident. These cases were analyzed using a semi-automated method which had output relationships between common concepts which were used as the basis of the classes for the ontology. The top-level classes used were “Actor, Action, Asset, Event, and Information.” The ontology can then describe scenarios by showing the relationships between subclasses of these top-level classes. The paper goes on to give a series of examples for how to use the ontology to further the field of insider threat detection. While it starts by restating the usefulness of this level of abstraction for publicizing information related to threats without also disclosing organization-sensitive information, the paper also goes on to note that the semi-automation of data collection that the ontology implementation paves the way for others to develop detectors for

indicators of insider threats. Also, the paper states that it would be possible for this work to be extended such that event logs (and other operational data) as well as information that organizations keep about insiders that is not as a result of direct interaction with information technology (human-resources data) could be translated and parsed in order to automatically create ontology individuals. If these processes are automated, then this would make it possible for a semantic reasoner to be constructed to classify insiders as instances of subclasses within the ontology, which would provide a clear view of specific indicators of threats. Ultimately, the development of this ontology appears to be a valuable stepping stone to further progress in the area of insider threat detection, but its greatest benefits will be lost if it is not widely used in a standard form.

3. TAXONOMY AND ONTOLOGY

Taxonomy and ontology are two common terminologies that are being used in information management and there are cases that people treat them as synonyms.

The term “taxonomy” could refer to a hierarchical classification or categorization system, or to an organization of concepts of knowledge, as well as a knowledge organization system designated to include term lists and classifications [13]. Except for some rare cases, defining the relationships between entities is not a concern when defining taxonomies, other than a hierarchical relationship between entities.

The term “ontology” other than its philosophical meaning, is a formal framework to represent knowledge in computer and information science. Ontologies define classes, properties of these classes and relationships between these classes within their domain. Using the relationships, we can extract other information from these information entities and use them to identify other previously unidentified relationships between them. The authors of [14] classify taxonomies as linguistic/terminological ontologies. However, taxonomies can also be used to define ontologies when the relationships between the classes are defined and a formal structure of an ontology can be constructed with them. How to develop an ontology is summarized in [15] as

- Defining classes in the ontology
- Arranging the classes in a taxonomic hierarchy
- Defining slots and describing allowed values for these slots, namely, creating properties of the classes
- Filling in the values for slots for instances.

This section identifies the methodology employed in the taxonomy and ontology development process and explains the details of the construction of ontology classes.

3.1 Methodology

The ontology development process we employed in this work is a top-down analysis which requires understanding the semantics of the end-users who will actually use the resulting ontology. It starts with creating a list of terms which will be used to construct the taxonomy of the structure. The

taxonomy needs to include the terms that define the classes in the domain. The taxonomy needs to be limited with what the resulting ontology will cover, what will it be used for, and what types of questions the ontology will answer to. Following the creation of the taxonomy, and the hierarchy within the taxonomy, the properties of the classes should be defined along with the relationships between classes.

The validation of the ontology structure is performed through competency questions. These questions assure the targeted value of the structure is achieved. They serve as procedures that indicate when the ontology development is sufficiently complete. The competency questions aim to ensure that the results are accurate, sufficient, and has the right level of granularity which is identified by the subject matter expert. They also ensure that the scope of the ontology is still within the limits.

It is essential to integrate the ontology created with other ontologies, as it will integrate the domain with the rest of the world. Considering that ontologies are a web of knowledge, integrating the ontology with other ontologies will create a bigger knowledge base and extend the opportunities of integrating this ontology with the existing systems. However, to increase data and information quality within a domain, we need to create an ontology that can represent that domain successfully, and creating an ontology requires expert knowledge within that specific domain as well as the skills required to create it. To create an ontology, ontology developers and domain experts need to work together. Ontologies that are created by people who lack either expert knowledge or ontology development skills may result in serious problems and wrong results. However, not all research projects have enough resources to hire people who have these skills. Also, even if the resources are sufficient, project teams may not think it is necessary.

3.2 Taxonomy

The efforts we have put into creating a taxonomy on finance domain has resulted with the taxonomy shown in Figure 1. As a result of the top-down analysis we performed with the domain experts of our collaborator banking institution, we have identified the taxonomy classes based on basic scenarios given in Section 4. The validation of these classes are tested through mapping between classes and instances gathered from the mentioned scenarios. The taxonomy we created for insider threats is given below:

Location: A position or site occupied or available for occupancy or marked by some distinguishing feature [16]

Institution: An established organization or corporation especially of a public character [16]

Bank: An establishment for the custody, loan, exchange, or issue of money, for the extension of credit, and for facilitating the transmission of funds [16]

Credit Union: A cooperative association that makes small loans to its members at low interest rates and offers other banking services [16]

Branch: A local office at a specific location of an institution

Person: One (as a human being, a partnership, or a corporation) that is recognized by law as the subject of rights and duties [16]

Customer: One that purchases a commodity or service [16]

Hacker: A person who illegally gains access to and sometimes tampers with information in a computer system [16]

Personnel: The people who work for a particular company or organization [16]

Role: A function or part performed especially in a particular operation or process [16]

Teller: A person who works in a bank and whose job is to receive money from customers and pay out money to customers [16]

Banker: A person that engages in the business of banking [16]

Branch Manager: A person that is responsible for managing a branch of an institution and the personnel working at that branch

Data: Facts or information used usually to calculate, analyze, or plan something [16]

Public Data: Data that can be accessed by anyone who is interested. The access and usage rights may vary and can be accessed with various ways

Sensitive Data: Data that calling for care and caution which can usually cause problems in case that someone else uses it

Account Data: Data that belongs to personal or business accounts which includes but not limited to name, address, account number etc. [16]

Login Credentials: Data that belongs to personal or business accounts which includes login usernames and passwords, security questions and answers

Transaction Data: Data of “a communicative action or activity involving two parties or things that reciprocally affect or influence each other” [16]

Database: A usually large collection of data organized especially for rapid search and retrieval (as by a computer) [16]

Attack: To set upon or work against forcefully [16]

Threat: Someone or something that could cause trouble, harm, etc. [16]

Security Issue: A matter or event of threat or attack

External Threat: A threat that is posed by someone or something that is not from the personnel of an institution

Insider Threat: Malevolent (or possibly inadvertent) ac-




Figure 1: Ontology classes created from initial terminology

tions by an already trusted person with access to sensitive information and information systems

3.3 Ontology

There are several types of ontologies that we can base the rules of our ontology framework.

Upper level ontology: The ontologies that belong to this level describe concepts that are the same across all knowledge domains which provides a high level of semantic interoperability.

Domain ontology: The ontologies that belong to this level describe concepts in a specific field or in a part of the world. This specific field or part of the world represents the domain that the ontology describes. Since the concepts belong to the domain, they may or may not be compatible with a concept that has the same name in a different domain ontology.

Hybrid ontology: The ontologies that belong to this level describe concepts that can be both mentioned in domain and upper level ontologies. Especially by working on integration of different systems together, the hybrid approach makes it easier to work with multiple ontologies. Some concepts can be defined universally but some concepts are described according to the domain related limitations.

Our goal is to provide a web of knowledge by integrating commonly used upper ontologies into our ontology. To achieve this task, we created a domain ontology on insider attacks focusing on financial sector, and then we identified

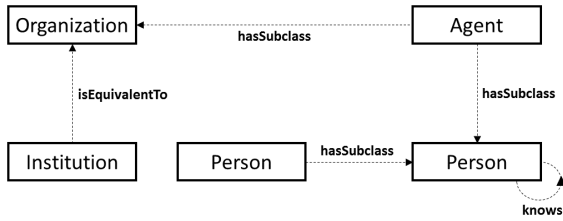


Figure 2: Integration of FOAF ontology classes

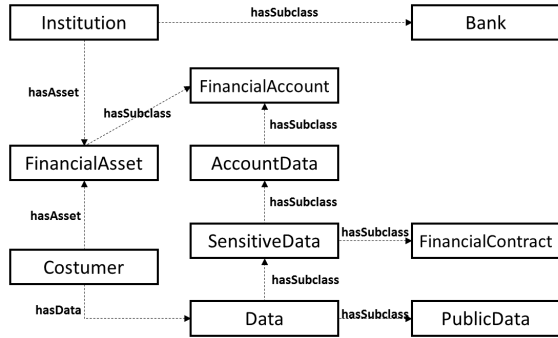


Figure 3: Integration of SUMO ontology classes

some ontologies that are commonly used by academia and industry that may possibly have similar classes that we identified in our ontology.

Friend of a Friend (FOAF) ontology [6] [7] describes people, their activities and relationships between each other and other objects. It allows groups of people to create social networks, which we are using to describe the relationships between customers, bank personnel and roles and hierarchy within the organizations. The common terms that we imported from this ontology are “Organization” and “Person” classes as can be seen in Figure 2. After importing these classes, we have expanded these terms with the domain specific subclasses, to define the banking environment.

The Suggested Upper Merged Ontology (SUMO) [5], has a broad range of domain areas included in it. However, it only provides a structure and a set of general concepts upon which domain ontologies could be constructed. Financial concepts are among these concepts, too. The common terms that we imported from this ontology are “FinancialAccount”, “FinancialContract”, “financial asset” and all of their subclasses. The relationships that these terms have with the other classes in our ontology can be seen in Figure 3.

Finance ontology [8] is an ontology on financial instruments, involved parties, processes and procedures in securities handling. We are using this ontology to define the financial instruments and involved parties within organizations, so that the main concern of our ontology stays as insider attacks instead of expanding into defining banking domain itself. The common terms that we imported from this ontology are “Address”, “Party” and all of the subclasses of “Party” as can be seen in Figure 4.




Figure 4: Integration of Finance ontology classes

Therefore, we integrated our ontology with FOAF to base our Person and Organization structure on universally defined terms and we expanded these terms. On the other hand, we imported classes from SUMO and Finance ontology to use the classes that are already defined in financial domain, so that we didn’t need to define new classes in the finance domain. The graph of the resulting ontology is shown in Figure 5.

4. SCENARIOS

The proposed insider threat ontology includes the following insider attack types: masquerade attacks, privilege elevation attacks, privilege abuse attacks and collusion attacks. Some attacks may appear in various cases and they may seem very different than each other even if they belong to the same classification, some can even happen unintentionally. The following part explains these attack types.

4.1 Masquerade attacks

In a masquerade attack, the attacker illegitimately assumes the identity of a legitimate user [17]. Before launching an attack like this, the attacker must gather the credentials to access the system. However, the gathering phase of the attack is outside of the scope of this work, and so we will make the assumption that the attacker has already gained the credentials necessary to access the system. Here, it is clear that the attackers try to hide their identity and make the victim responsible for any action they take.

Scenario 1: Amelia and Ben are tellers and work at the same branch of a bank. Amelia goes to use the restroom and leaves her computer logged in to the system thinking that she will be gone for a very short time. Ben takes advantage of her absence, and uses her computer to look up a few persons’ social security numbers and transactions to use them later. He returns back to his seat before Amelia comes back.

Scenario 2: Ben is helping a customer when his computer begins to freeze. He sees that Amelia has left her computer open, so he borrows her computer to help the customer. He accidentally forgets to switch accounts, and changes the details for another customer’s account.

4.2 Privilege elevation attacks

In a privilege elevation (also known as privilege escalation) attack, a user with insufficient permissions accesses the information that only a more privileged user can see. The attackers usually exploit a vulnerability of the system to escalate granted permissions [18], so that they can use these




Figure 5: Main components of the ontology

new permissions to access information.

Scenario 1: Amelia is a branch manager and Ben is a banker at the same branch of a bank. Amelia has an emergency and have to leave the bank for a few hours. Since there may be emergencies that require her attention or her privileges, Amelia leaves her computer logged in to the system and trusts Ben to manage the bank in her absence. Ben uses her privileges to look up a few persons' sensitive information to use them later along with filling in for her at the bank.

Scenario 2: Ben is a teller at a bank. He finds out that the software allows him to see all of the sensitive information of the bank's customers and thinking that he is supposed to see them, he doesn't notify his superiors. He uses this information when he needs to without permission of his superiors.

4.3 Privilege abuse attacks

In a privilege abuse attack, the user uses his/her permissions to retrieve information that they are not supposed to see.

Scenario 1: Amelia is a branch manager of a branch at a bank in Buffalo, NY. She also looks up sensitive information of some people from New York, NY. As not to look suspicious, she chooses the customers of a specific branch, keeping in mind that people from the same household tend to travel together and have bank accounts from the same branch.

Scenario 2: Carl is performing routine updates on the bank's database system. He is given full access to perform the updates, and decides to download all of the customer information from a specific branch to check out if the updates were performed correctly. Somehow, he forgets to delete the file from his computer.

4.4 Collusion attacks

In a typical collusion attack, there are usually more than one people with different privileges collaborating to access and harvest information. Since this data is usually supposed to

include more relations and be more extensive, the impact of these attacks is usually higher.

Scenario 1: Amelia and Ben are bankers and work at the same branch of a bank. Both have different responsibilities, hence different privileges in the system. They collaborate together to harvest sensitive information to use them later.

Scenario 2: Carl is a branch manager and Karen is a secretary at the same branch of a bank. Carl leaks information from the database systems and from internal documents of the bank to a rival company. However, after gathering them, Carl hides the information along with a lot of other documents and directs Karen to send it to specific addresses. Karen doesn't know that she is collaborating but she doesn't check out what she is sending.

5. DATA SOURCES

We have collaborated with a financial institution to create our ontology structure and consulted with banking experts on how we can start our initial efforts. After this phase, we are currently working on taking our collaboration to a next level. For regulatory reasons, the databases at high risk of insider attacks have not yet been revealed for us to observe, log and collect information for evaluation. Because of these reasons, we currently don't know what will be available to us, and what will not be, to be able to validate our initial proof of concept. Hence, in this section, we can only list what we can expect from the financial institution to make available for us to create instances in the ontology.

5.1 Databases and log files

We need to understand access patterns for one database (and eventually more) in a financial institution database system. This includes a snapshot of the data in the database, as well as query logs including:

- The look up or update query being issued

- (Anonymized) identification of the user or role that accessed the data
- (Anonymized) physical or IP address of the machine issuing the query

What we try to achieve here is to see a view of the daily life of that database. The size and traffic information of the database can be gathered from the log entries.

5.2 Permission structure

To have a better understanding of the internal structure of the environment, we created a taxonomy to identify roles and differentiate between different attack types. This taxonomy for the user roles includes the following information:

- User roles
- Access Permissions
- Manages
- Managed by

For example, what does a teller do? What privileges does a teller have on the database? Who manages the teller and who has more privileged access to the system in that branch? Who can change the privileges given to a user? We need to ask these questions for selected types of users. Some of the relationships that are represented in the ontology does not have to present in order to trust the resulting structure. This missing values can be created in time, or left blank.

5.3 Previously identified attacks

We will use attacks identified in the past to create new insider attack scenarios and simulate them on the example databases. Real examples provided by financial institutions will be a guideline preparing these scenarios. The data, log files, and scenario details about some insider attacks that are detected before should be very useful in this phase.

These scenarios can either be used to *create attack models*, or *simulate an attack* to see if the system accurately identifies an insider attack. Although this information would be very beneficial if it existed in order to see how a real attack represented in the ontology, the experts of the institution we are collaborating expressed that this information is one of the most classified data types.

6. DISCUSSION

We have presented a preliminary cyber ontology focusing on insider attacks in banking domain targeting database systems. As indicated before, the prior efforts in insider threats branches to two different directions. These branches are psychological aspects and physical aspects of insider threats. Our work takes the initiative to start efforts on building a cyber ontology for insider threats in the financial sector, as it is critical to developing countermeasures against insider attacks in this domain. The contribution of our work is,

- creating a cyber ontology framework for insider threats in the financial sector focusing on relational database management systems
- ensuring the integration with other knowledge domains to enable data integration.

The literature survey we have performed shows us that this ontology fills the gap in ontological structuring of insider threat research in financial sector. The ontologies developed on insider threat research generally focus on defining insider threat and incidents [12]. The work in [12] leads us to experiment on specific domains and use the domain specific knowledge to create a semantic structure. This structure defines the insider threat in financial sector more conclusively. Even if we have collaborated with financial sector experts, we know that there is still a lot to do to expand the capability of our ontology, since we still cannot gather real data from banking databases.

The preliminary cyber ontology we created has classes from FOAF and SUMO ontologies, which are universally defined, and the terms in them mean the same across all knowledge domains. In this sense, our ontology provides a high level of semantic interoperability. When fully developed, we believe that this integration with other domains and semantic structures approach can prove effective to addressing more factors about insider threats as it could be used by researchers to test and evaluate their detection and mitigation schemes, as well as identifying similar attacks by using previously identified attacks.

7. FUTURE WORK

The major threat of insider attacks drives both academia and industry to find better solutions. As we continue our research on insider threats, we will need to extend the ontology that we developed and create a knowledge base. As indicated in Section 5, we should create our knowledge base from real working systems to be able to validate the ontology that we constructed. We are working on building collaborations with financial institutions to gather the data required to validate the current structure. After the validation phase, we are looking forward to iteratively building on the ontology to improve both scope and capability. The validation phase will be performed with competency questions to test if the ontology contains enough information to answer the questions, if the answers it provides have a sufficient level of detail, or if they represent the domain well enough.

Additionally, the risk analysis of insider attacks can be very beneficial, especially when performed from both defender's and the attacker's perspective. Conceptually, the likelihood of an attack happening is correlated to the cost to the attacker [19]. Defenders consider the entire system and take security measures considering the system as a whole. However, the attackers focus on a part of the system and attack that specific part, usually preferring the appropriate type of attack specific to that part. These attacks need a preparation time and effort, which is considered a cost to the attacker. We will aim to exploit this information to create a more effective insider attack semantic structure.

8. ACKNOWLEDGMENTS

This material is based in part upon work supported by the National Science Foundation under award number CNS - 1409551. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. We would like to thank Thomas Mitchell, and Patrick Coonan for their review of related articles and their help in editing.

9. REFERENCES

- [1] Robert H. Anderson and Richard Brackney. Understanding the insider threat: Proceedings of a March 2004 workshop. Santa Monica, CA, USA, 2004. RAND Corporation. Also available in print form.
- [2] The insider threat: An introduction to detecting and deterring an insider spy.
<http://www.fbi.gov/about-us/investigate/counterintelligence/the-insider-threat>. Accessed: 2015-06-15.
- [3] CERT Insider Threat Center. 2014 U.S. state of cybercrime survey. July 2014.
- [4] Ian Niles and Adam Pease. Towards a standard upper ontology. In *Proceedings of the international conference on Formal Ontology in Information Systems-Volume 2001*, pages 2–9. ACM, 2001.
- [5] Adam Pease, Ian Niles, and John Li. The suggested upper merged ontology: A large ontology for the semantic web and its applications. In *In Working Notes of the AAAI-2002 Workshop on Ontologies and the Semantic Web*, 2002.
- [6] Jennifer Golbeck and Matthew Rothstein. Linking social networks on the web with FOAF: A semantic web case study. In *AAAI*, volume 8, 2008.
- [7] Li Ding, Lina Zhou, Tim Finin, and Anupam Joshi. How the semantic web is being used: An analysis of FOAF documents. In *System Sciences, 2005. HICSS'05. Proceedings of the 38th Annual Hawaii International Conference on*, page 113c. IEEE, 2005.
- [8] Eddy Vanderlinden. Finance ontology documentation.
<http://fadyart.com/ontologies/documentation/finance/index.html>. Accessed: 2015-06-15.
- [9] Jakub Szefer, Pramod Jankhedkar, Diego Perez-Botero, and Ruby B. Lee. Cyber defenses for physical attacks and insider threats in cloud computing. In *Proceedings of the 9th ACM Symposium on Information, Computer and Communications Security*, ASIA CCS '14, pages 519–524, New York, NY, USA, 2014. ACM.
- [10] Jeffrey Hunker and Christian W. Probst. Insiders and insider threats – an overview of definitions and mitigation techniques. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, 2(1):4–27, 2011.
- [11] S. Mathew, S. Upadhyaya, D. Ha, and H.Q. Ngo. Insider abuse comprehension through capability acquisition graphs. In *Information Fusion, 2008 11th International Conference on*, pages 1–8, June 2008.
- [12] Daniel Costa, Matthew Collins, Samuel Perl, Michael Albrethsen, George Silowash, and Derrick Spooner. An ontology for insider threat indicators: Development and application. In *Proceedings of the Ninth Conference on Semantic Technology for Intelligence, Defense, and Security*, STIDS 2014, pages 48–53. CEUR Workshop Proceedings, 2014.
- [13] Heather Hedden. *The Accidental Taxonomist*. Information Today, Inc., Medford, New Jersey, 2010. ISBN: 978-1-57387-397-0.
- [14] Gilles Falquet, Claudine MÃltral, Jacques Teller, and Christopher Tweed. *Ontologies in Urban Development Projects*. Advanced Information and Knowledge Processing 1. Springer-Verlag London Limited, 2011.
- [15] Natalya F. Noy and Deborah L. McGuinness. Ontology development 101: A guide to creating your first ontology.
http://protege.stanford.edu/publications/ontology_development/ontology101.pdf. Accessed: 2015-06-15.
- [16] Merriam-Webster.com. Merriam-webster.
<http://www.merriam-webster.com/dictionary/term>. Accessed: 2015-06-10.
- [17] Malek Ben Salem and Salvatore J. Stolfo. Modeling user search behavior for masquerade detection. In *Proceedings of the 14th International Conference on Recent Advances in Intrusion Detection*, RAID'11, pages 181–200, Berlin, Heidelberg, 2011. Springer-Verlag.
- [18] Lucas Davi, Alexandra Dmitrienko, Ahmad-Reza Sadeghi, and Marcel Winandy. Privilege escalation attacks on android. In *Information Security*, pages 346–360. Springer, 2011.
- [19] Ameya M Sanzgiri and Shambhu J Upadhyaya. Feasibility of attacks: What is possible in the real world—a framework for threat modeling. In *The 2011 International Conference on Security and Management, SAM*, 2011.