

Developing Systems for Cyber Situational Awareness*

James S. Okolica, J. Todd McDonald, Gilbert L. Peterson, Robert F. Mills, and Michael W. Haas

Abstract—In both military and commercial settings, the awareness of Cyber attacks and the effect of those attacks on the mission space of an organization has become a targeted information goal for leaders and commanders at all levels. We present in this paper a defining framework to understand situational awareness (SA)—especially as it pertains to the Cyber domain—and propose a methodology for populating the cognitive domain model for this realm based on adversarial knowledge involved with Cyber attacks. We conclude with considerations for developing Cyber SA systems of the future.

Index Terms—Cyber attacks, network defense, situational awareness, business continuity planning

I. INTRODUCTION

On February 18th, 2001, Robert Hanssen was arrested for selling American secrets to Moscow for a period of 22 years [1].

On April 28th, 2007, distributed denial of service (DDOS) attacks began on media website in Estonia. These DDOS attacks would later spread to attacks on Estonia's critical infrastructure including banks, ministries, and police.

On August 8th, 2008, scant hours after shooting began between Russian and Georgian forces in South Ossetia, cyber attacks began on Georgia's government and bank websites.

THE Department of Defense (DoD) NetOps strategic vision states that commanders, users, and operators (at all levels) need accurate and timely information when accessing the global information grid (GiG). Of course, the understanding of the health and mission readiness of the GiG remains vital for this goal to be achieved. At every level of the mission space, we need a coherent framework which translates events that occur in time and space to their (possible) deleterious effects on mission success.

What all of the above incidents have in common is that information was available that might have led to earlier detection and mitigation. Robert Hanssen had a password breaker program on his work computer [1]. Network probes and DDOS attacks were performed on Georgia's critical infrastructure as early as July 20, 2008¹. What is needed is a

means to increase awareness of what is happening in cyberspace—particularly from the viewpoint of attackers and malicious adversaries. What is needed is *Cyber Situational Awareness*.

With the advent of Cyber as a prominent operational concern and even a defined domain of operations in the U.S. Air Force, the DoD as a whole has come to realize that Cyber-based effects and defensive operations are integral to the overall success of air, land, naval, and space operations. Industry has also realized that vulnerabilities in this realm, including targeted malicious attacks, have huge monetary consequences and carry losses in both productivity and public trust.

In this article we offer a definition for situational awareness for the Cyber domain and present an overview of the problem space within which it resides. We show how traditional definitions of SA may be adapted for Cyber specifically in a sense/evaluate/assess loop which provides correlation between real events, key system components, and their corresponding business/mission impact. We propose a notion of the adversarial narrative, which provides a ground truth view of SA which knowledge and data discovery techniques ultimately attempt to replicate and refine. We also propose a methodology for building an automated discovery engine that can build a useful, actionable Cyber SA picture for commanders at various levels.

II. DEFINING CYBER SITUATIONAL AWARENESS

While there are several definitions of what is meant by situational awareness, one of the most accepted is by Dr. Mica Endsley [2]. It defines SA as "the **perception** of elements in the environment within a volume of time and space, the **comprehension** of their meaning, and the **projection** of their status in the near future". Endsley then extends his concept of SA to include a memory component and a decision/ action taken as a result of the SA. The decision / action is then considered to act upon the environment which produces a circular loop as SA begins again with a perception of the new environment (Figure 1).

Using Endsley's definition, there are three functions *any* SA system must perform: (1) it must sense its environment, (2) it must take its raw sense data and assemble it into a meaningful understanding of its environment, and (3) it must use its current understanding to predict the future. Figure 2 provides a specific Cyber example based on an attacker with inside

*The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

¹John Markoff, Aug 12, 2008 NY Times, "Before the Gunfire, Cyberattacks", <http://www.nytimes.com/2008/08/13/technology/13cyber.html>




Fig. 1. Endsley's situational awareness model [2]. SA leads to decisions and actions which affect the environment itself. SA captures the environment state through perception, comprehension, and project (predictive analysis), forming a loop.

knowledge and access to an organization (an insider threat).

First, the SA system senses elements of an individual's environment. Using an insider threat example, these sensors include emails sent and received by the individual as well as transaction logs from the applications the individual uses for his day-to-day activities. The SA system then assembles this information into a concept which matches its already known concept of "insider threat". At this point, the SA system has a suspicion that the individual might constitute an insider threat. The SA system then predicts that if the individual is an insider, he may (1) send information to computers outside of the local network and (2) possess password cracker programs. The SA system then decides to activate packet traffic and file locator sensors to determine if it is correct. When the results are positive, the SA system then combines the packet traffic and firewall information to determine what data vulnerabilities exist. The concept observed during this second pass is "data exfiltration." However, the SA system still only understands this concept in terms of data.

The final step is to incorporate an understanding of the relationship between business processes and data elements to determine the mission impact of the projected data exfiltration. It is this final step that is missing from many of the Cyber SA efforts to date. High level business processes must be broken down into detailed workflow steps performed by individuals within different organizations. Users and applications must then be associated with each functional responsibility and action respectively within each of the workflows. Once this association has been made, it is possible to relate data concepts to operational concepts. Then, when sensors extract user and application data and feed correlation tools that assemble it into a comprehensible picture of the data environment, business health assessment tools can then

translate the data environment into an operational environment. This complete process (Figure 3) then provides a holistic Cyber Situational Awareness.

Before proceeding, Endsley's term *comprehension* needs to be better framed. Specifically, we need a distinction between local comprehension and global comprehension. If it is possible for a single host to determine a concept, e.g., "I am under a DDoS attack", then we define that knowledge as a local concept. If the only way to determine a concept is to collect information from several hosts, e.g., "a worm is spreading across the network", then we define that knowledge as a non-local concept. For instance, a domain name server being singled out for a DDoS attack is a local concept.

Now, consider a non-Cyber example of this distinction. When a homeowner considers his water system, he thinks about the individual pipes, which rooms have faucets and whether the toilets are working. He also may give some thought to the water entering and leaving his home. However, when a city engineer considers his water system, the only parts of an individual's home that the engineer thinks about is the water entering and leaving a home. Not only doesn't the engineer care about the specific conditions in an individual home, he may not even use the same vocabulary, e.g. faucets and toilets. This implies that the vocabulary used to describe local perceptions may not be needed to describe global perceptions. Furthermore, the vocabulary used to describe global data environment perceptions may not be used to describe global operational perceptions. At each level, the transformation from perception to comprehension changes the language used to describe the environment (Figure 4).




Fig. 2. Insider threat Cyber SA example. Sensors at lower levels on individual devices focus on specific information/data elements. Evaluation matches activities and patterns of data to known threat categories which spawn additional sensor / data collection activities. Determination of particular offensive operations and associated vulnerabilities that support the operations are distilled. The health of the overall mission and plans that mitigate effects of the projected evaluation are assessed.

III. DEFINING THE CYBER SA PROBLEM SPACE

Developing an infrastructure that provides operational cyberspace situational awareness requires successfully solving multiple problems. As Figure 5 illustrates, in addition to developing sensors (problem 1), correlation tools (problem 2) and visualization tools (problem 3), there are several embedded issues that must be resolved.

First, detecting a non-local (i.e., distributed) attack requires correlating information from multiple types of multiple sensors. For instance, there are sensors that track network traffic on a specific host and there are sensors that track program executions on a specific host. Only by combining the information from both types of sensors across multiple hosts can a “low-and-slow” attack be detected. At the heart of this issue is the need to evaluate information from multiple types of sensors that both view and describe the network environment in different ways. Developing an infrastructure for describing information from disparate sources in a unified way is defined as the environment description language (EDL) problem (Figure 5-P4). One subset of the EDL problem is describing data information that is either: (a) local to the Host (i.e. Host Data EDL (HDEDL) problem) or (b) descriptive of the entire network (i.e., Network Data EDL (NDEDL) problem).

Second, in addition to minimizing sensors’ processor time on each individual host, correlating multiple sensors across

hosts requires minimizing network traffic between hosts. If all sensor information from each host is transmitted across the network, the result would be a self-inflicted denial of service attack. Instead, some sensor fusion at the local (i.e. individual host) level needs to occur before transmitting a more abstracted state to other hosts. Determining methods for summarizing local data and transmitting it efficiently is defined as the scalability problem (Figure 5-P5).

Third, an issue that emerges naturally from the first and second issues is identifying *what* to look at. Time and again, in the wake of an attack (cyber or otherwise), signs are uncovered that if they had been noticed and acted upon in a timely manner would have prevented the attack. Security professionals are left with the uncomfortable task of answering why they hadn’t been looking for that particular sign. Unfortunately, the reality is that it is impossible, even in a cyber environment, to look at and evaluate everything. Instead, security personnel must select a subset of the data to collect and analyze. Determining what to look at is defined as the feature extraction problem (Figure 5-P6).

Our fourth concern deals with single points of failure. If correlation occurs in a central location and the adversary is able to neutralize that target, the security of the network is significantly degraded. In addition, if an adversary is able to subvert a host and cause it to send out erroneous sensor information, the security of the network will also be compromised. Addressing these twin issues of single point of failure and sensor corruption is defined as the resiliency




Fig. 3. Cyber situational awareness (SA) model. Business continuity planning (BCP) based on workflow processes and models allow top-down mapping of mission, operational, and systems functions/organizations/equipment to the overall business goals and activities. Data at various levels capture both BCP and Cyber SA data. Sensors and correlation tools provide bottom-up knowledge synthesis, filtering and fusing data to provide top-level business process health.

problem (Figure 5-P7).

Finally, once we address these issues, it becomes possible to develop correlation tools to determine when the network, or its hosts, is/are under attack and what the implications of this attack are to the health of the data network.

While the four embedded problems listed above address determining whether the network is under attack, the issue still remains whether we can adequately communicate this information to security professions and senior management. While the problem of visualization is more of a human effects issue than a technological one, unless this problem is solved, efforts on the problems above are wasted. Related to the Visualization problem, is the “so what?” factor. While a good visualization tool can provide a Chief Information Officer with the relative health of her network, it does not address the Chief Operations Officer’s (COO) question of “can the operation fulfill its mission?” To answer this question, network health must be translated into business process health. In the same way that the data EDLs addresses disparate types of sensor information, an operational environment description language (OEDL) would allow business process engineers to describe the relationship between the data environment and the operational environment. Thus the EDL problem has three sub-problems: HDEDL, NDEDL, and OEDL (seen in Figure 4). With this information, visualization tools can be developed to provide the COO with the answers to her questions.

While IT specialists think of visualization tools as red light/green light monitors, this awareness represents only one type of visualization. Another involves providing a narrative description of the offensive operations being perpetrated on the organization. Consider the following example: several network sensors identify that Bob’s machine has a rootkit on it. The tools further identify what the rootkit is trying to hide.

What senior management wants to know is who installed the rootkit and for what purpose. A successful Cyber SA monitoring system might provide senior management with parts of the real-life story that involves Mallory, the employ who actually perpetrated several malicious actions that led to the rootkit installation and operation. Though the true, real-life narrative of the events would detail the underlying social, political, or personal motivations (i.e., Mallory targeted Bob out of personal vendetta related to a work-place affair), the Cyber SA narrative would determine that Mallory used social engineering. The awareness would include pertinent pre-exploitation details such as the fact that Mallory sent Bob an email with a link to a website that has a cross site scripting (CSS) vulnerability (which he clicked on), subsequently giving Mallory administrator privileges on his machine. She then used those privileges to install a rootkit and a backdoor so that she could access his machine. She started small by




Fig. 4. The Cyber SA environment. Environment description languages exist at three different levels, providing both local and global SA comprehension and expression of Cyber SA events.




Fig. 5. The Cyber SA problem space. Six different problem areas are delineated, capturing the primary research space for accomplishing successful Cyber SA.

changing his Outlook schedule and removing important meetings; however, she quickly moved on to sending emails (and trying to remove the evidence so Bob wouldn't notice) to other employees with racial and sexist jokes in order to have him fired for inappropriate conduct – she was extremely vindictive in her dismissal. Although this sort of narrative may be considered science fiction today, by defining the vocabulary and languages and developing the appropriate sensors and correlation tools, this sort of visualization may be commonplace in the business world of tomorrow.

Once all of the problems described in the problem space are addressed, there is still the issue of obtaining realistic data in order to test the Cyber SA systems and build confidence that projection accurately characterizes threats, offensive activities, and vulnerabilities.

IV. DEVELOPING A CYBER SA SYSTEM

We believe three overlapping activities are needed to develop a Cyber SA system: (1) developing a test environment that provides sensor data that can be correlated and fused, (2) developing one or more languages that can describe the cyber environment at different levels of abstraction, and (3) integrating the adversarial narrative into the abstraction space.

A. Developing a Cyber SA Test Environment

The purpose of a Cyber Situational Awareness system is to report on the health of an operational network. Therefore, an ideal dataset would provide data that duplicates an operational network. Some of the desirable characteristics include:

- (1) “Real” data including normal baseline traffic and attempted/successful malicious attacks. While the percentage of normal to malicious data may be modified to provide sufficient exemplar data, the percentages should be explicitly stated so that a realistic baseline can be defined.
- (2) “Timely” data from a time period long enough to model all activity expected on the operational network. This includes (a) peak usage data as well as off-peak (e.g., nighttime and weekend) data; (b) end of month/quarter/year usage data as well as day-to-day usage data;

- (3) “Functional” data of many different types of users including technical, clerical, operational, and management users.
- (4) “Scaled” data for an operational network of appropriate size. While this varies depending on where the operational Cyber SA system is intended, it is likely that the network data should include data from several hundred, if not several thousand, hosts.
- (5) “Heterogeneous” data that covers all of the possible inputs that an IDS might desire. While it is impossible to enumerate all possible inputs, representative data includes network traffic, operating system logs, application transaction data, and temporal operating system process data.

Unfortunately there is currently no publicly available dataset that satisfies all of these requirements. However, there are several datasets available that satisfy some of them. In 1998, MIT Lincoln Laboratories under Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratories developed the first dataset for evaluating intrusion detection systems [3]. They added to this dataset with additional datasets in 1999 and 2000 [4]. Although there have been several criticisms of the representativeness of the data [5], they still remain one of the most used datasets.

While DARPA has since sponsored a 2002 Cyber Panel Correlation Technology Validation effort, the datasets used are no longer publicly available. Instead, there are several datasets from other competitions that have been made available for public use. For instance, DEFCON is an annual convention for security professional and hackers. One of the principal events at DEFCON is its 72 hour Capture the Flag (CtF) contest where teams attempt to protect their own network while invading other teams (thus capturing their flag). The event traffic from DEFCON 8 CtF and DEFCON 10 CtF was recorded and made available by the Shmoo Group at <http://cctf.shmoo.com/>. Lastly, the 3rd International Knowledge Discovery and Data Mining Tools Competition focused on network intrusion and it has made its dataset available as well [6]. Unfortunately, what all of these datasets have in common is a lack of a baseline. While the DEFCON and KDD Cup data are real data, they were developed in an artificial contest environment and consequently contain unrealistic amounts of attack data with little or no baseline data.

Recognizing the issues inherent in synthesized IDS data, several organizations have developed testbeds as more realistic environments for measuring the success of intrusion detection systems. We describe four such environments which have representative features consistent with Cyber SA and development and sensor data analysis.

Originally built from Utah's EMULAB software, the cyber-Defense Technology Experimental Research (DETER) testbed has been configured to “provide stronger assurances for isolation and containment” [7]. Its goal is to specifically test network defense against attacks including distributed denial of service attacks, worms and viruses. DETER was developed to provide a medium-scale (approximately 300

nodes in two clusters) environment for “safe, repeatable, security-related experimentation to validate theory and simulation”. It is run by Information Sciences Institute, University of California at Berkeley funded by the National Science Foundation and the Department of Homeland Security. More information can be found at <http://www.isi.deterlab.net/>.

Netbed, also a descendant of EMULAB, is “a software system that provides a time- and space-shared platform for research education, or development in distributed systems and networks” [8]. It uses both local, dedicated nodes, geographically-distributed shared nodes and emulated Dummynet nodes. Researchers access these nodes via a virtual topology which causes Netbed to configure a physical topology. Netbed provides an experimentation facility that integrates these approaches, allowing researchers to configure and access networks composed of emulated, simulated, and wide-area nodes and links. Netbed’s primary goals are “ease of use, control, and realism, achieved through consistent use of virtualization and abstraction”. Netbed is run by The Flux Group, School of Computing, University of Utah. More information can be found at <http://www.emulan.net/>.

The Protected Repository for the Defense of Infrastructure against Cyber Threats (PREDICT) is a repository for current computer and network operational data accessible through a secure web-based portal and is made available to qualified cyber defense researchers located in the United States [9]. It is run By RTI International, a not-for-profit research institute funded by the Department of Homeland Security. More information can be found at <https://www.predict.org/>.

Finally, System Administrator Simulation Trainer (SAST) is a software simulator which artificially generates internet/network traffic and superimposes actual exploits on it. SAST provides a safe simulator for DoD security and personnel and system administrators to hone their capabilities by providing thousands of real world exploits and an environment that can mimic an organization’s information infrastructure. It is run by the National Center for Advanced Security Systems Research. More information can be found at <http://www.ncassr.org/project/>.

B. Describing the Cyber Environment

Language is “a systematic means of communicating by the use of sounds or conventional symbols” [10]. It must, at a minimum, contain names of items (e.g. John, George, Andrew, hit, smack, beat) and may also contain classifications of items (e.g., person, president, attack, and strike). Additionally, adding grammar enables communication of relationship between items (e.g., without a grammar {George beat Bill}, {Bill beat George} and {Bill George beat} are equivalent). As a result, language is generally considered to be composed of vocabulary (possibly containing classifiers) and the elements to manipulate them.

In order to (1) describe data that a Cyber SA system senses and (2) fuse that data into comprehensible concepts, a Cyber SA system requires a language. Relevant vocabulary may

include (1) devices connected to a network, (2), users of the network, (3) application software run on the network (4) user missions/operations enabled by the network, (5) actions performed by devices, users, and applications (6) communications between devices, users, and applications, (7) actions performed on devices, users and applications.

There are two distinct ways of communicating relationships. The first, and most obvious, is via grammar (e.g., “George beat Bill”). The second defines vocabulary such that a single item contains this information (e.g., attack(source=George, target=Bill, time=12-Jan-09;21:23:00, method=stick)). There are benefits to each technique. Formal deductive methods, e.g., predicate logic, benefit greatly from the explicit relationships between objects that grammars provide. On the other hand, since the formalization of relationships limits the expressiveness of language, knowledge from data discovery can benefit from the lack of grammars, allowing for unconsidered relationships to emerge.

Two primary application areas that are related to Cyber SA are intrusion detection and cyber forensics. While the authors know of no Cyber SA-specific language, there are several languages related to intrusion detection and cyber forensics that apply. The Intrusion Detection Message Exchange Format (IDMEF) was developed by an Internet Engineering Task Force (IETF) working group and sent out with a Request for Comments (RFC) in March, 2007. IDMEF uses extensible markup language (XML) to facilitate the multitude of sensor vendors. It provides for sensor input from network devices (e.g., switches and routers), O/S audit logs, and application transaction logs as well as alerts to be sent back to operators and actions to be taken in response to sensor input.

The IDMEF data model (RFC4765) shown in Figure 6 is an object-oriented representation of a space which includes source data with very little information (e.g., origin, destination, time, and name/description) and source data with too much information (e.g., application transaction logs with hundreds of fields in them). The IDMEF-Message entity is the top level class. All other entities are sub-classes of it. Currently the two subclasses of IDMEF-Messages are alerts and heartbeats. Alerts correspond to analyzer (i.e., sensor) alerts or events and occur asynchronously. There are several sub-classes within the alert class including tool alerts (to describe attack tools), correlation alerts (to describe previously grouped and correlated alerts), and overflow alerts (to describe buffer overflow attacks). The heartbeat class defines messages sent out at regular intervals from analyzers to managers (centralized tools used by operators to configure sensors, analyzers, data consolidators, etc.). Lastly, the object-oriented representation provides both flexibility and extensibility.

While the IDMEF model requires the implementer to define the relationships between classes, Pinkston *et al.* [11] have developed ontology, shown in Figure 7, which defines both the classes and the relationships between them. Although as described, TCO focuses on network attacks but might be easily extended to incorporate exfiltration or modification of host data. Furthermore, despite the fact that TCO cannot




Fig. 6. IDMEF data model. Alerts and heartbeats define all sub-classes of IDMEF messages, covering both asynchronous and continuous monitoring data.




Fig. 7. Target Centric Ontology (TCO).

describe distributed attacks affecting multiple hosts, it can detect them through the use of generic queries.

In addition to IDMEF and TCO, the National Center for Forensic Science and the University of Central Florida Department Of Engineering Technology have proposed the digital evidence markup language (DEML) as a method to model digital evidence [12]. Unlike IDMEF and TCO, DEML is more focused on characteristics of a specific device, e.g., hard disk model, partition size, O/S revision and uptime, etc. While DEML may not be expressive enough to be used to describe a large scale network-wide environment, its specificity makes it a good choice for describing a detailed host-level environment.

Although not specifically a language, MITRE has compiled the common vulnerabilities and exposures (CVE) list [13] to provide standardized names for different attacks and vulnerabilities. CVE has since received widespread adoption by a number of organizations and individuals.

C. Measuring Aggressor Cyber SA

A crucial final element needing integration into Cyber SA systems is the ability to accurately describe or measure what is actually happening in reality. We consider that for the most basic of Cyber SA questions (whether a Cyber attack is underway, imminent, or in preparation stages), only the attacker possesses ground truth situational awareness and only the attacker can define the ground truth narrative which describes who, what, why, when, and where. Unless an attacker acts for no reason at all (purely psychopathic motivations), the underlying reasons and goals of an attack can help us identify patterns of behavior. Likewise, the actual steps taken in a malicious attack are known by the attacker perfectly, though execution of them may not be perfect. This perspective helps shape the way we design and test systems for Cyber SA.

One way to describe Cyber SA then is how close assessment may come to the attacker's ground truth SA. Successful detection, identification, and differentiation of various malicious activities may be compared only rightly to the actual activities. Our methodology for resolving this question also forms a basis for refining a domain model that supports information fusion from bottom data/correlation tools to high-level Cyber SA abstractions (using environment descriptions and ontology). We envision test environments that involve use of real-world attacks (ARP cache poisoning, data exfiltration, social engineering, malware deployment, etc.) executed in the backdrop of configured sensors and data correlation tools. Such attacks give the bottom-layer data elements which may be fed to correlation tools and engines.

What prevents accurate, high-level Cyber SA in many cases is not knowing which data elements to look for and which data elements to keep. It is those missing data elements and correlation hints that prevent the high-level picture from being adequately created. By executing known attacks in an iterative manner, we expect that candidate domain models may be refined that capture a "middle" layer of knowledge conducive for populating our high level SA expressions. Our current research efforts focus on developing this middle layer of domain ontology and finding appropriate fusion algorithms with favorable predictive behaviors.

V. CONCLUSIONS AND FUTURE WORK

While the above steps bound the work of developing Cyber SA systems, we expect continued progress by researchers in the problem space areas will help candidate systems mature over the next decade. The co-problem of adequately defining the business mission space remains an open problem with a different and active research community. Without this fuller context of how Cyber may affect business process health and lower levels of correlation, Cyber SA systems may not find prominence in operational use. Our future work aims at developing adequate intermediary domain models that facilitate generalized fusion of lower-level correlation data with higher level SA statements.

ACKNOWLEDGMENT

This material is based upon work supported in part by the U.S. Air Force Office of Scientific Research under grant number F1ATA09048G001.

REFERENCES

- [1] Wise, David (2003). *Spy: The Inside Story of How the FBI's Robert Hanssen Betrayed America*, Random House Publishers, ISBN 0375758941.
- [2] Endsley, Mica (1995). "Toward a theory of situation awareness in dynamic systems". *Human Factors* 37(1), 32-64.
- [3] Lippmann, R., Fried, D., Graf, I., Haines, J., Kristopher, J., *et al.* (2000). "Evaluating Intrusion Detection Systems: The 1998 DARPA Off-line Intrusion Detection Evaluation," *DARPA Information Survivability Conference & Exposition - Vol 2*, pp.1012.
- [4] Haines, J., Rossey, L., Lippman, R., and Cunningham, R. (2001). "Extending the 1999 Evaluation", In the *Proceedings of DISCEX 2001*, June 11-12, Anaheim, CA. Datasets available at <http://www.ll.mit.edu/mission/communications/ist/corpora/ideval/data/index.html>.
- [5] McHugh, J. (2000). "Testing intrusion detection systems: a critique of the 1998 and 1999 DARPA intrusion detection system evaluations as performed by Lincoln Laboratory." *ACM Trans. Information System Security* 3(4), 262-294.
- [6] Available online: <http://kdd.ics.uci.edu/databases/kddcup99/kddcup99.html>.
- [7] Benzel, T., Braden, R., Kim, D., Joseph, A., Neuman, C., Ostrenga, R., Schwab, S., and Sklower, K. (2007). "Design, Deployment, and Use of the DETER Testbed". In *Proceedings of the DETER Community Workshop on Cyber Security Experimentation and Test*, August 2007.
- [8] White, B., Lepreau, J., Stoller, L., Ricci, R., Guruprasad, S., *et al.* (2002). "An Integrated Experimental Environment for Distributed Systems and Networks". *Proceedings of the Fifth Symposium on Operating System Design and Implementation*, Dec 2002, 255 - 270.
- [9] Available online: https://www.predict.org/Portals/o0/files/Documentation/MANUAL%20OF%20OPERATIONS/PREDICT_Overview_final.pdf.
- [10] Available online: wordnetweb.princeton.edu.
- [11] Undercoffer, J., Pinkston, J., Joshi, A., Finin, T. (2003). "Target-Centric Ontology for Intrusion Detection," *IJCAI Workshop on Ontologies and Distributed Systems (IJCAI'03)*, August, 2003.
- [12] Available online: http://www.ncfs.org/digital_evd.html.
- [13] Available online: <http://www.cve.mitre.org/cve/cve.html>.