

Mitigating Cloud Computing Security Risks using a Self-Monitoring Defensive Scheme

Steven Mazur
AFRL/RIGG
Air Force Research
Laboratory
Rome NY, 13441

Erik Blasch
AFRL/RYAAX
Air Force Research
Laboratory
Rome NY, 13441

Yu Chen
Department of Electrical
and Computer Engineering
Binghamton University
Binghamton NY, 13902

Victor Skormin
Department of Electrical
and Computer Engineering
Binghamton University
Binghamton NY, 13902

Abstract – Cloud Computing (CC) promises to offer seamless provisioning of on-demand services and data through the network. These services are made discoverable to the user in the form of simple abstractions via virtualized resources. These services offer unprecedented dynamic, primarily hardware dependent, scalability. Although CC promises to make life much easier for the user, it comes with significant security issues. Because on-demand service provisioning for applications and data will be used by hundreds of thousands (if not millions) of users simultaneously, a successful intrusion would not only expose sensitive data, but it could also completely cut users off from both applications and/or data. This paper examines the underlying security risks inherent to the CC paradigm, compares approaches to mitigate known security risks, and offers a solution that leverages intelligent multi-agent systems and network data ontologies to provide automated defense for both known and unknown malware security risks. We describe a mechanism whereby a dynamic ontology can be self-enriched over time to provide for some protection against unknown security risks.

Keywords: Cloud Computing, intelligent multi-agent system, dynamic ontology

I. INTRODUCTION

Cloud Computing (CC) has been described a number of ways: network virtualization, the re-birth of application service providers (ASP), or a lot of hype for a rebirth of the services that mainframe computers have already been providing for many years. For the purposes of this paper, we will use CC as described in the recently released Draft Special Publication 800-146, “DRAFT Cloud Computing Synopsis and Recommendations”[1].

“Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models.” The elements of a cloud are shown Figure 1.



Cloud Computing

Having secure access to all your applications and data from any network device

Fig. 1. Elements of a Cloud [2].

Although there is no set definition for CC, what is certain, is that this paradigm has caught the attention of many key players in information technology (IT) infrastructure and support, to include all the Branches of Service in the US DOD, National Institutes of Standards (NIST) [1], Defense Information System Agency (DISA) [3], as well as leading IT providers like IBM [4], HP [5], Microsoft [6], Oracle [7], and Google [8]. Despite all of this attention, it is not clear if Cloud will be able to withstand the attack from malware without a significant change to current infrastructure. This paper briefly describes various configurations that are considered to be part of the CC environment in Section 2. It identifies several key security risks that this paradigm introduces due to its open nature in Section 3. Section 4 reviews related work and in section 5 we offer a general solution to help mitigate some of the risks using intelligent agents, ontologies and Computational Intelligence (agent-based systems using ontologies and course of action reasoning). Our solution provides for an internal (to the cloud) semi-autonomous defensive security mechanism that leverages intelligent multi-agent systems and network data ontologies to provide automated defense for both known and some unknown malware security risks. Section 5 describes the overall self-defensive monitoring scheme to include a mechanism whereby a dynamic ontology can be self-enriched over time to provide for some protection against these security risks. Discussion and conclusions are drawn in Section 6.

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE JUL 2011	2. REPORT TYPE	3. DATES COVERED 00-00-2011 to 00-00-2011			
4. TITLE AND SUBTITLE Mitigating Cloud Computing Security Risks Using A Self-Monitoring Defensive Scheme		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Research Laboratory, Rome, NY, 13441		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Presented at the 2011 National Aerospace & Electronics Conference (NAECON'11), Dayton, OH, July 20 - 22, 2011, Government or Federal Purpose Rights License					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

II. THE CLOUD

CC can trace its roots back to the mainframe, client-server model where almost all of the actual computing was done at a central data aggregation point using expensive central processing, memory, and data storage that was shared by all users. Figure 2 shows a general cloud and subscriber view.




Fig. 2. General Cloud and subscriber view [1].

Over time, the cost of central processors, memory and storage significantly decreased, encouraging the use of dedicated personal computers at each work station, and only shared data on servers. Centralized approaches made the business case for highly distributed local node processing a logical investment decision. Within the past five years a near insatiable appetite for consumer access to streaming content has caused data service (bandwidth) providers to rapidly expand their enterprises and in the process have provided the key component that allows for a cloud construct to exist. Recent trends in advanced computing models bring enterprise processing power to the user in the form of *services*. These services are much more than just a Service Oriented Architecture [9] (which is really an attitude and not an architecture). Some of the components include the *Cloud* itself (a distributed collection of computing resources), *Cloud Services* (applications, systems software and hardware), *Cloud Technology* (dynamically scalable, virtualized resources), and the *Cloud Ecosystem* (users, developers, managers, datacenters, service providers, integrators, aggregators, infrastructure vendors and content providers) [10].

The Cloud is not comprised of a single technology, system or architecture - it is based on a number of technologies, hardware and software configurations, and is realized in the form of various services and deployment models [11]. The recently released NIST Special Publications 800-145 [12] describes CC as “a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. The NIST cloud model promotes availability and is composed of five essential characteristics, three service models, and

146 [1] further explains what some of the known about using the cloud as a tool and gives some indication of the limitations and applicability of cloud technologies. Variations in size and complexity of each instantiation of a cloud are determined based on these attributes, as determined by the customers’ needs. The trade space that will determine the final configuration is primarily based on the customer’s operational and security requirements, as well as economic factors.

Based on this information, and quite a bit of hind sight on the part of vendors, many organizations are coming forward, claiming that (even though they did not know it at the time) they have been running in a cloud environment for up to a decade. These include IBM [4] and Hotmail (Microsoft) [6]. One could argue that in hind sight they are simply binning technologies that they developed to fit them in with the new CC paradigm.

The cloud is here to stay. To drive this point home, the United States Defense Information Systems Agency has released its 2011-2012 campaign plan [13] and will focus on the maturation and migration to a cloud service models in an effort to provide both garrisoned users and those at the “edge” with access to decision quality information anywhere, anytime. The campaign plan will include technology enhancements such as improvements in information assurance capabilities, application monitoring, automated provisioning and automated infrastructure orchestration [14].

Clearly, one size does not fit all and the solution that is chosen needs to take into consideration the reliability and stability of the available (current and future) bandwidth. To a certain degree, CC has a significant dependence upon *bandwidth* which is the backbone of the Cloud concept. What the internal and external bus’s are to the typical motherboard, the network is to CC. Just as sound cannot travel in a vacuum, a Cloud cannot exist without on-demand high bandwidth. Detailing the benefits of any one Cloud instantiation over another is beyond the scope of this paper; however, suffice it to say the one security size does not fit all. As such, the concepts developed in this paper are not necessarily applicable to all cloud security approaches, but are presented in a fashion that the researcher might take what they need and modify it to suit their needs.

III. SECURITY IN THE CLOUD

Because of the large number of ways in which CC technologies can be implemented (different architectures, service and deployment models) and inter-operate with disparate technologies and software designs, security is especially challenging. The challenge is especially true for public clouds where the infrastructure, hardware and software are owned and operated by a third party selling these services to multiple subscribers. It is ultimately the responsibility of the data owner/service subscriber to ensure that the appropriate integrated, reliable, and repeatable security measures have been put in place to

provide adequate protection for their data. Depending on the level of service agreement, subscribers may shift significant responsibility into the hands of the Cloud providers to ensure, not only uninterrupted access to data and services – but to also provide an agreed to level of security as part of the quality of service. Part of this security is tied to where the data is located.

In several configurations, the data may exist in many locations – and never be one single complete data set at any one of those locations (or country). Not only company policy but also legal concerns related to data segregation need to be examined. In some cases encryption can be used to help provide the necessary protection, however, this could affect availability. Managing user access and data privileges can be very difficult because the subscriber does not have an ability to screen the service providers’ administrators. User access (e.g. authentication and authorization) is usually an internal activity that in-house security administrators control and there must be an adequate amount of trust and openness between the parties. In the case that data is used inappropriately/illegally there must be an audit mechanism in place to allow accurate forensic reconstruction of nefarious activities. Unfortunately all of the tools do not yet exist, and time may prove it an impossible task to identify and track all breaches. Disaster recovery efforts will see significant gains using Cloud technologies, however, how and where the data is backed up needs to be analyzed in detail to ensure risks are properly addressed. The short term advantages of CC may not be viable in the long term.

For the most part, current security solutions for CC are based on using the same solutions that have existed for a number of years. One virtual machine appliance (Catbird vSecurity™) for the *VMware* hypervisor has been advertised as a zero-cost security appliance solution – however hypervisors have significant security issues as shown in [15] and [16]. Looking deeper into the security aspects of CC, we must look at some of the internal security issues and how we can increase the subscribers’ confidence in the service providers’ ability to protect both data and applications. *NIST SP 800-39* [17] describes the need for a risk management framework to promote the concept of near-real-time risk management through the implementation of a continuous monitoring process – however that process has not yet been defined. Risk management must not only provide for continuous situational awareness [18] and control of the organizations system’s security, but also for continuous knowledge of threats and vulnerabilities [19]. This level of additional tracking and management will not be inexpensive and will require new methods for risk mitigation.

Previous research that built frameworks for tracking and managing security related data [20]-[24] do not go deep enough. We need to research and develop new methods, not only to track cyber security related data [25] and various metrics of cyber security [26], but also to analyze

cyber threats, coupled with course of action reasoning will allow for near real time threat detection and automated mitigation [27]-[30]. Some research in cyber situation awareness (CSA) has shown positive results for identifying certain types of malware [31, 32]; however, additional research is necessary to expand CSA to other instruction sets and to incorporate additional automated reasoning on data streams.

Simply inspecting packet dumps and host data files is not enough. Current network-based intrusion detection systems, host-based intrusion detection systems, and intrusion prevention systems (IPSs) based on signature-based or anomaly-based will not be enough in the cloud environment, and without automated self monitoring capabilities that can also generate and execute defensive courses of action, CC will be vulnerable to massive failures.

With the advent of malware such as the *stuxnet*, we have seen a small picture of things to come. Initial analysis indicates that this threat was able to exploit four 0-day vulnerabilities, compromise two digital certificates and inject code into supervisory control and data acquisition (SCADA) systems without the operator or IDS detecting it [33]. In a cloud construct, the subscriber is completely helpless when trying to trace data inconsistencies that result from such an attack, and would be at the mercy of the service provider to solve the problem. If the data trace inconsistency problem went unnoticed for several weeks or months the damage could be irreparable.

IV. RELATED WORK

A. Intelligent Agents

Over the last decade a number of researchers describe the use of intelligent agents [34]-[37] that in theory could effectively and automatically provide some basic levels of computer and network defense that could be applied to CC. The intelligent agents offer various solutions using agents to control access and authentication, distributed trust management, audit and intrusion detection, attack vector pursuit, and diagnostic and system restoration.

Most agent-based applications/systems have three common properties – *distributed* data knowledge and management mechanisms; a community of *autonomous* cooperative components (cognitive); and *inheritable* components as described in [35]. Agent technologies have given us tools for a number of advanced system configurations and management protocols. Modeling and Simulation has also benefited significantly from the use of multi-agent systems. Agent-oriented modeling and simulation for various aspects of computer network defense have been introduced [36]; however, the utility of agents does not end with modeling and simulation. These same technologies can also be applied to runtime systems.

Because agents most often exist as a service or daemon even when they are not actively executing, they remain in memory and take up precious resources. To be effective,

these agents must be relatively compact and provide for low level elementary actions. Higher level actions can be created using them in a self-configuring fashion [38] require significant memory. Examples where some of the earlier proposed agents have been replaced by policy and technologies include public key infrastructure (PKI) and secure sockets layer (SSL). In theory this makes sense, however, some of the more advanced agents that need to be developed could provide scalable services that are tailor made and created on demand.

An *Agent-based Adaptive Dynamic Semantic Web Service Selection* (AADSS) framework was proposed in [39]. Based on real-time conditions, this AADSS service can dynamically select the “right” service, and adaptively change the bound services as the needs of the subscriber change. A consumer-agent is used to maintain a list of candidate services and allows for the reassignment of assets as necessary.

In [40], a gateway is proposed that would allow agents to interact with services. The gateway would perform a mediation function between web service requests from (virtual) organizations of agents to enable gateway managed web service access.

Despite limited applicability of services similar to those mentioned above, without some additional cognitive capabilities many of these agents could take several months to develop – and a system effectively employing them could take years. In their original venue many of them would not only be coupled with the central processing unit instruction set, but also with the operating system itself. Because both of these can change frequently, this effectively limits the period that they would be effective. What is needed is a more general framework that provides for the realization of these actions at a high level, but allows for a dynamic environment that can reach down below the service level using an inferred ontology to effect the automated creation of a desired course of action.

B. Ontologies

An ontology can be defined in many ways. In Artificial Intelligence (AI), and in computer science in general, an ontology refers to an engineering artifact, made up of a specific vocabulary used to describe a certain reality. The ontology definition also includes a set of explicit assumptions regarding agreement on the intended meaning of the vocabulary words. A number of ontologies have been developed for use in Information Systems [41]-[43], [25]; however they were primarily concerned with characterizing and tracking data associated with the systems’ assets and not the content of the data and streaming information. Work in the cyber forensics domain such as the The Cybersecurity Information Exchange (CYBEX) format [44]-[46], [42] has generated ontologies that are very useful and could in part contribute to a security solution for CC. Initial research in this area

concentrated on individual devices and components that were related and interacted with one another.

Promising research in unsupervised ontology induction from text [47] has shown that semantic content can be extracted from unstructured using a method that induces and populates a probabilistic ontology using dependency-parsed text as input. The unsupervised ontology induction approach was found to be effective in extracting a medical knowledge base and could easily be extended for use in the cyber domain.

Another area of research that can help in the area is the hard-soft fusion (HSF) [48]-[51] that allows for heterogeneous data to be aligned semantically, in order to provide a pragmatic course of action. HSF can be extremely helpful when data has been generated by a number of disparate sources and a cross alignment of “like” data is needed to help prepare data to be ingested by a reasoning engine.

C. Computational Intelligence

Computational Intelligence (CI) is a relatively new term, although there are numerous conflicting definitions. For the purposes of this paper, we will describe it as advanced low level artificial intelligence that uses computational adaptation to mimic human logic and reasoning. CI attempts to model brain mechanisms and then applies these models to developing cognitive algorithms in an attempt to bring computer system performance closer to the brain-mind [52]. CI can be realized in a number of different ways. In this paper we suggest a paradigm that uses intelligent agents not only to collect data about system resources, and processor/memory states, but also to deliver and execute byte code modifications to mal-data that has been found to be present in a stream or depending upon the nature of the data, might also halt or kill a process (or data stream) depending upon the severity of the potential mal-activity.

CI core methods such as artificial neural networks, fuzzy systems, evolutionary computation, artificial immune systems, swarm intelligence, and soft computing are described in [53]. The authors analyze a significant number of research works and provide useful insights into how CI might be used in an intrusion detection system. Many of these CI technologies could easily be adapted to be used as mechanisms that can deal with the numerous security challenges that CC creates.

V. SOLUTION DESCRIPTION

For many years Computer network defense has been dependent upon automated signature-based intrusion detection capabilities as the primary defense mechanism. Unless a new paradigm for secure computing is designed, we will not be able to secure many of the CC constructs and realize the full potential that CC offers. The bottle neck in current systems is the fact that the speed with which these signatures are developed and implemented

determines whether or not a particular piece of malware will be successful in its mission. By not publishing the existence of new malware and immediately bringing multiple technologies to bear on creating solutions, new malware we will put significant amounts of data and application integrity at risk. We paper propose a self-monitoring defensive mechanism for CC which integrates intelligent agents, computational intelligence, and ontologies.

As depicted in Figure 3, distributed intelligent agents collect data within the cloud by monitoring devices, data streams (to include inter communication between devices and network processors) and code execution. In some cases this information is handled by brokering agents that can resolve some issues locally, and in other cases correlation or reasoning engines must be utilized to determine the appropriate course of action.




Fig. 3. Concept of Autonomous Cloud Security.

Brokering is similar to the way that we keep our balance when walking down a hallway while engaged in a conversation with a friend. For the most part, we don't think about maintaining our balance, it just happens. While we are walking we are usually thinking about something completely different. In a similar fashion, many (if not most) of the actions that would take place in our system would not require advanced analytics or fusion engines, but would require situational analysis metrics for brokering arbitration [54]. These brokering engines can be realized in the form of ontology-based fusion engines that continually change the behavior by regenerating the low level ontologies that describe the system, based on agent input and game-theoretic threat assessment. By also using an ontology to describe malware, we can then compare the two ontologies [55] and if there is a close enough match, we declare the presence of that malware and based on our policies (course of action) take steps to eliminate the threat. Elimination may be simply launching agents to terminate the code execution, freezing a process, alerting the system administrator, or something more involved such as initiating a significant system cleansing or system fail over operation.

For our purposes these ontologies are not huge monolithic structures that describe the entire cloud world within a single ontology. They are small entities, each with a very specific purpose for existing. We can think of them like organs in the body – built for a specific purpose (i.e. lungs to breath air, heart to pump blood, etc.). Within the system these ontologies can be operated on individually or brought together by mid-level and upper-level ontologies to create different instantiations of more advanced views. In some cases these ontologies will be used and thrown away, in other cases they will need to be persisted for advance long term analysis.

The concept of *automated ontology generation* (AOG) is not new. We propose that a new method be developed that can enrich each generation of an ontology during its creation by using related, known data to seed the process. Current methods of generating ontologies in an unsupervised fashion do not take into account the significant amount of supporting information that already exists. In [56] a knowledge-based, ontology-centric security management system is created to support the process of driving technical controls based on informal, high-level policy statements, but this does not go far enough. We propose a significantly different mechanism that takes advantage of the fact that in the network defense domain we already know many of the questions and policy issues *a priori*, and we postulate that by using semantic analysis we can help guide the ontology induction process, and increase the ontology quality in the AOG process. For example if we use policy, security assumptions, and the known queries (questions) that security analysts use to monitor system integrity, we might improve the process of using unsupervised ontology induction from using these and other corpora to seed the AOG. Using an iterative process we should be able to increase the quality of the ontologies and increase the level of situational awareness within the CC environment.

We can also automatically synthesize ontologies that accurately represent the desired state of the system, and via comparative analysis (and through the use of causal agents) bring the system in line with this view, thereby enabling automated configuration management. In a similar fashion we could also auto generate ontologies for other items of interest such as malware signatures that might not yet have been seen and we could then use them to detect the presence of that malware using ontology-based comparative analysis.

The AOG capability will allow us to bring a significant improvement to the time critical aspect of cloud self-monitoring mechanisms. The emergence an AOG-based autonomic cloud defense will enable a system that is flexible and adaptive by leveraging the content and context rich information that can be manipulated using ontologies. It will also allow the system to rapidly and dynamically react to new situations by integrating a number of additional, possibly heterogeneous resources.

VI. CONCLUSIONS

Failure to protect the cloud will not result in millions, but in billions of dollars in damage. To protect against this threat we need to act quickly to start designing intelligent defensive systems that can minimize these damages.

In this paper, we identified several key security risks that the CC paradigm introduces and proposed a general solution to help mitigate some of these risks. Our solution includes using the concept of computational intelligence in the form of an automated intelligent agent-based data collection mechanism that monitors data streams, services, network devices, byte code and machine language execution coupled with inference/ontology-based real time analysis which creates a novel autonomous course of action reasoning capability. Moreover, using the ability to auto re-generate the ontologies that form the basis of the analytical process, the system will be able to automatically adjust its defense posture so that it can immediately reason over new data as it is being processed within the Cloud. Because an ontology is like a person's frame of reference – shaped by the data and experiences that created it, we will be able to go much farther and faster to provide adequate security for CC if we increase our efforts to build and share libraries of ontologies.

REFERENCES

- [1] L. Badger et al., "NIST Draft Special Publication 800-146, DRAFT Cloud Computing Synopsis and Recommendations," Gaithersburg, MD: National Institute of Standards and Technology May 2011.
- [2] Cloud Computing. (n.d.), Retrieved Jul. 1, 2011, from http://en.wikipedia.org/wiki/Cloud_computing
- [3] Cloud Computing and Enterprise Services. (n.d.). Retrieved Jun. 26, 2011, from <http://www.disa.mil/computing/cloud/index.html>
- [4] IBM Smart Cloud, (n.d.), Retrieved Jul. 1, 2011 from www.ibm.com/cloud-computing
- [5] Cloud Computing. (n.d.), Retrieved Jul. 10, 2011 from <http://www8.hp.com/us/en/solutions/solutions-detail.html?compURI=tcm:245-300983>
- [6] Microsoft on Cloud Computing. (n.d.), Retrieved Jul. 1, 2011 from www.microsoft.com/presspass/presskits/cloud
- [7] Oracle Cloud Computing. (n.d.), Retrieved Jul. 1, 2011 from <http://www.oracle.com/us/technologies/cloud/index.html>
- [8] Google Apps for Business. (n.d.), Retrieved Jul. 1, 2011, from www.google.com
- [9] E. Blasch, G. Chen, D. Shen, H. Chen and K. Pham, "Services Oriented Architecture (SOA) based Persistent ISR Simulation System," *Proc. of SPIE*, vol. 7694, Apr. 2010.
- [10] W. Chang, H. Abu-Amara and J. Sanford, *Transforming Enterprise Cloud Services*, Springer, 2010.
- [11] H. Takabi, J.B.D. Joshi and G. Ahn., "Security and Privacy Challenges in Cloud Computing Environments," *IEEE Security and Privacy*, vol. 8, issue 6, pp. 24–31, Nov. 2010.
- [12] P. Mell and T. Grace, NIST Special Publication 800-145, "NIST Draft Special Publication 800-146 - The NIST Definition of Cloud Computing (Draft)," Gaithersburg, MD, National Institute of Standards and Technology, Jan. 2011.
- [13] Defense Information Systems Agency Campaign Plan 2011-2012. (n.d.). Retrieved Jul. 1, 2011, from http://www.disa.mil/campaignplan/DISA_CP_2011-12_ExecSum.pdf
- [14] Cisco Cloud Enablement Services. (n.d.). Retrieved Jul. 1, 2011, from http://www.cisco.com/en/US/services/ps2961/ps10364/ps10370/ps11104/services_cloud_enablement_overview_service_provider.pdf
- [15] J. Reuben, "A survey on virtual machine security," Tech. report, Helsinki University of Technology, Oct. 2007.
- [16] M. Christodorescu, R. Sailer, D. Schales, D. Sgandurra and D. Zamboni, "Cloud security is not (just) virtualization security," *Proc. ACM workshop on Cloud computing security (CCSW)*, Nov. 2009.
- [17] R. Ross et al., "NIST Special Publication 800-39, Managing Information Security Risk: Organization, Mission, and Information System View," Gaithersburg, MD, National Institute of Standards and Technology, Mar. 2011.
- [18] E. Blasch, I. Kadar, J. Salerno, M. M. Kokar, S. Das, G. M. Powell, D. D. Corkill, and E. H. Ruspini, "Issues and challenges of knowledge representation and reasoning methods in situation assessment (Level 2 Fusion)," *J. of Advances in Inform. Fusion*, vol. 1, no. 2, pp. 122-139, Dec. 2006.
- [19] G. Chen, D. Shen, C. Kwan, J. Cruz, M. Kruger and E. Blasch, "Game Theoretic Approach to Threat Prediction and Situation Awareness," *J. of Advances in Information Fusion*, vol. 2, no. 1, 1-14, Jun. 2007.
- [20] I. Kotenko, "Framework for Integrated Proactive Network Worm Detection and Response", *17th Euromicro Int. Conf. on Parallel, Distributed, and Network-Based Process.(PDP 2009)*, Feb. 2009.
- [21] L. Flagg et al., "Bringing Knowledge to Network Defense", Retrieved Jun. 23, 2011 from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.161.2958&rep=rep1&type=pdf>
- [22] N. Santos, K.P. Gummadi and R. Rodrigues, "Towards trusted cloud computing," *Proc. Hot topics in cloud computing (HotCloud 2009)*, Jun. 2009.
- [23] D. Raoui, S. Benhadou and H. Medromi, "New distributed platform for intrusion detection based on multi-agents system," *J. of Eng. and Tech. Research*, vol. 2 (10), pp. 200-206, Oct. 2010.
- [24] J. Feng, Y. Chen, W.-S. Ku and P. Liu, "Analysis of Integrity Vulnerabilities and a Non-repudiation Protocol for Cloud Data Storage Platforms," *2nd International Workshop on Security in Cloud Computing (SCC 2010)*, San Diego, California, USA, Sep. 14, 2010.
- [25] T. Takahashi, Y. Kadobayashi and H. Fujiwara, "Ontological approach toward cybersecurity in cloud computing," *3rd Int. Conf. on Security of Inform. and Networks (SIN 2010)*, Taganrog, Rostov region, Russia, Sep. 2010.
- [26] E. Blasch, J. Salerno and G. Tadda, "Measuring the worthiness in Situation Assessment," *IEEE NAECON*, Jul. 2011.
- [27] R. Ribeiro de Azevedo et al., "An Autonomic Ontology-Based Multiagent System for Intrusion Detection in Computing Environments inside Computing Environments," *International Journal for Infonomics (IJI)*, vol. 3, iss. 1, Mar. 2010.
- [28] D. Shen, G. Chen, J. B. Cruz, Jr., E. Blasch and K. Pham, "An Adaptive Markov Game Model for Cyber Threat Intent Inference", invited Ch. 21 in *Theory and Novel Applications of Machine Learning*, M. J. Er and Y. Zhou. (Eds.), IN-TECH, 2009.
- [29] D. Shen, G. Chen, L. Haynes, E. Blasch, and G. Tadda, "Adaptive Markov Game Theoretic Data Fusion Approach for Cyber Network Defense," *IEEE MILCOM 2007*, Oct. 2007.
- [30] D. Shen, G. Chen, L. Haynes, and E. Blasch, "Strategies Comparison for Game Theoretic Cyber Situational Awareness and Impact Assessment," *Int. Conf. on Info Fusion - Fusion07*, Jul. 2007.
- [31] A. Volynkin, V. Skormin, D. Summerville, J. Moronski, "Evaluation of Run-Time Detection of Self-Replication in Binary Executable Malware," *7th Annual IEEE Information Assurance Workshop*, West Point, NY, Jun. 21-23, 2006.
- [32] A. Tokhtabayev, V. Skormin, A. Dolgikh, "Dynamic, Resilient Detection of Complex Malicious Functionalities in the System Call Domain," *Military Commun. Conf.*, Nov. 2010.

- [33] W32.Stuxnet Dossier, Version 1.4. (Feb 2011). Retrieved Jun. 5, 2011 from http://www.wired.com/images_blogs/threatlevel/2011/02/Symantec-Stuxnet-Update-Feb-2011.pdf
- [34] L. Kagal, T. Finin, and A. Joshi, "Developing Secure Agent Systems Using Delegation Based Trust Management," *Security of Mobile MultiAgent Systems (SEMAS 02) held at Autonomous Agents and MultiAgent Systems (AAMAS 02)*, Jul. 2002.
- [35] I. Kotenko and A. Ulanov, "Agent Teams in Cyberspace: Security Guards in the Global Internet," *Proc. of the Int. Conf. on Cyberworlds (CW06)*, Nov. 2006.
- [36] I. Kotenko, "Active vulnerability assessment of computer networks by simulation of complex remote attacks," *Int. Conf. on Comput. Networks and Mobile Computing (ICCNMC)*, pp. 40–47, Nov. 2003.
- [37] I. Kotenko, "Multi-agent modelling and simulation of cyber-attacks and cyber-defense for homeland security," *IEEE Int. Workshop on Intelligent Data Acquisition and Advanced Computing Syst. (IDAACA 2007)*, pp. 614–619, Sep. 2007.
- [38] F. Sheldon, T. Potok, M. Langston, A. Krings and P. Oman, "Autonomic Approach to Survivable Cyber-Secure Infrastructures", *IEEE Int. Conf. on Web Services (ICWS 2004)*, California, USA, Jul. 2004.
- [39] J. Li, D. Ma; L. Li and H. Zhu; "AADSS: Agent-based Adaptive Dynamic Semantic Web Service Selection", *Int. Conf. on Next Generation Web Services Practices NweSP 08*, Oct. 2008.
- [40] B. J. Overeinder, P. D. Verkaik and F. M. T. Brazier, "Web service access management for integration with agent systems," *Proc. of the 23rd Annual ACM Symp. on Appl. Computing, Mobile Agents and Syst. Track*, Mar. 2008.
- [41] H. Assadi, "Construction of a regional ontology from text and its use within a documentary system," *Proc. of the Int. Conf. Formal Ontology and Inform. Syst., (FOIS 98)*, Trento, Italy, 1998.
- [42] A. Kim, J. Luo, and M. Kang, "Security Ontology for Annotating Resources," *Proc. of 4th Int. Conf. on Ontologies, Databases, and Applications of Semantics (ODBASE'05)*, Agia Napa, Cyprus, Nov. 2005.
- [43] S. Fenz and A. Ekelhart, "Formalizing information security knowledge," *Proc. of the 2009 ACM Symp. on Inform., Comput. and Commun. Security (ASIACCS 2009)*, Sydney, Australia, New York, Mar. 2009.
- [44] A. Rutkowski, Y. Kadobayashi, I. Furey, D. Rajnovic, R. Martin and T. Takahashi; "CYBEX – The Cybersecurity Information Exchange format," *ACM SIGCOMM Comput. Commun. Review*, vol. 40, no. 5, pp. 59-64, Oct. 2010.
- [45] A. Brinson, A. Robinson and M. Rogers, "A cyber forensics ontology: Creating a new approach to studying cyber forensics," *J. of Digital Investigation*, 3S, pp. 37-43, 2006.
- [46] J.W. Ji, "Holistic Network Defense: Fusing Host and Network Features for Attack Classification," M.S. thesis, Dept. Elect. and Comput. Eng., Air Force Institute of Technology, WPAFB, OH, 2011.
- [47] H. Poon and P. Domingos; "Unsupervised Ontology Induction from Text", *Proc. of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 296–305, Jul. 2010.
- [48] E. P. Blasch, "Ontological Issues in Higher Levels of Information Fusion: User Refinement of the Fusion Process," *Int. Conf. on Info Fusion - Fusion 03*, Jul. 2003.
- [49] E. P. Blasch, É. Dorion, P. Valin, and E. Bossé, "Ontology Alignment using Relative Entropy for Semantic Uncertainty Analysis," *Proc. IEEE Nat. Aerospace Electronics Conf (NAECON)*, 2010.
- [50] E. P. Blasch, É. Dorion, P. Valin, E. Bossé, and J. Roy, "Ontology Alignment in Geographical Hard-Soft Information Fusion Systems," *Int. Conf. on Info Fusion - Fusion10*, Jul. 2010.
- [51] E. Blasch, "Characterizing the Semantic Information Loss between Geospatial Sensors and Geospatial Information Systems (GIS)," *Proc. SPIE 8053*, Apr. 2011.
- [52] L. Perlovsky, "Computational Intell. Applications for Defense," *IEEE Computational Intell. Mag.*, Feb. 2011.
- [53] S. X. Wu and W. Banzhaf, "The use of computational intelligence in intrusion detection systems: A review," *Applied Soft Computing*, vol. 10, pp. 1–35, 2010.
- [54] E. Blasch, M. Pribilski, B. Daughtery, B. Roscoe and J. Gunsett, "Fusion Metrics for Dynamic Situation Analysis," *Proc. of SPIE*, vol. 5429, Apr. 2004.
- [55] Y. Jean-Mary, E. Shironoshita, and M. Kabuka, "Ontology matching with semantic verification," *J. of Web Semantics*, vol. 7, iss. 3, pp. 235-251, Sep. 2009.
- [56] B. Tsoumas and D. Gritzalis, "Towards an Ontology-based Security Management," *20th Int. Conf. on Advanced Inform. Networking and Applicat. (AINA'06)*, vol. 1, pp. 985-992, Apr. 2006.