Bringing Knowledge to Network Defense

L. Flagg G. Streeter A. Potter
Sentar, Inc.
4900 University Square, Suite 8
Huntsville, AL USA
+1 256.430.0860
{lflagg, gstreeter, apotter}@sentar.com

Abstract

Security managers must scan through multiple continuous data streams issuing from diverse sources in an effort to defend computer networks from attack. However, manual aggregation of this information is not achievable for vital decision-making within a narrow timeframe if security managers are not well-educated in current attack vectors. Thus, extensive and periodic training in attack methods, situation awareness and decision-making strategy should be required. However, it is challenging to provide training environments that can properly simulate multi-stage attacks effectively. Security managers are also impeded by the lack of dynamic feedback afforded by traditional training. This can result in false positive or negative readings of their preparedness. In this paper we discuss strategies to provide effective simulation and training of computer network defense for security managers through the integration of knowledge, intelligent agents, and proven network defense technologies.

Keywords

Network-centric warfare, Situational Awareness, Decision Support, Cyber Defense Training

1 Introduction

As computer network systems used by military, government, and business organizations become ever more vital to the organizational mission, they also grow larger, faster, more complex, more heterogeneous, and more difficult to protect [4]. Nowhere is the nation's reliance on network systems more critical than in network-centric warfare (NCW). As a well established military doctrine, NCW translates information advantage into warfighting advantage through robust networking of sensors, decision makers, and shooters to achieve shared situation awareness, increased speed of command, increased pace of operations, greater lethality, increased survivability, and self-synchronization. Successful conduct of NCW requires constant vigilance and an almost predictive awareness of the security posture of the Network Battlespace. For a security manager to be properly prepared to conduct NCW, several broad requirements must be met:

- They must understand the actual state of the network.
- They must understand the policies that define the desired state of the network
- They must have the necessary reasoning capability for determining what actions would be required to maintain actual conditions as closely as possible to the desired state of the network

Among the many challenges to successful NCW providing for these requirements necessitates the integration of many varied network defense technologies to provide all the correct network posture information.

However, currently available network defense technologies provide security managers with more information than can be assimilated in a time-critical environment. They must scan through multiple, continuous data streams issuing from heterogeneous network defense, monitoring and management tools in an attempt to dynamically gauge the current security posture of their organization's network. High-stress conditions compound the difficulty of situation awareness. Aggregation of time-sensitive information may not be achievable for vital decision-making within a narrow timeframe. Additionally, latency in human processing will easily exceed acceptable limits for effective response. A structure of (and a prior understanding of) what those data streams indicate is imperative to maintain a level of acceptable risk for large-scale, complex, critical information systems.

In order to prepare system administrators for security manager positions, extensive and periodic training in attack methods, situation awareness and decision-making strategy is essential. It is those decisions that

involve ambiguity, such as discovering and reacting to an attack before it has succeeded, for which it is most difficult to prepare and carry out remediation without interrupting service. There are many factors that affect a training program for network intrusion detection. It is challenging to provide training environments that can properly simulate multi-stage network attack circumstances effectively. Security managers are also impeded by the lack of dynamic feedback afforded by traditional training. This can produce inaccurate assessments of their preparedness. Thus, the need for network security system simulation and training technology is readily apparent.

In this paper, we present how knowledge can be integrated into a network defense environment using the dynamism of agents to provide an appropriate platform on which to build intrusion detection simulation and training capabilities. Incorporating intelligent situation awareness enhancement for dynamic feedback as well as proven intrusion sensor technology affords a powerful and realistic medium on which to base training. Through this automated and goal-oriented reasoning system, facilities can be integrated to dynamically challenge a trainee using evolving, real-world network defense scenarios, producing decision-oriented security managers that can identify attacks before they succeed.

2 Context

A wide variety of threats may be used in a network attack. [1] identifies the following broad categories: snooping, modification, spoofing, repudiation of origin, denial of receipt, delay, and denial of service. Snooping is unauthorized access to information. Snooping may result in a loss of confidentiality, or the information obtained through snooping may be used for subsequent attacks of a more destructive nature. For example, an attacker will frequently perform a port scan of a target system as part of the planning for an attack. By scanning the ports and examining the information returned, the attacker obtains valuable information for determining the system's vulnerabilities [8].

In a sophisticated attack, multiple exploits may be combined. Such attacks are known as blended attacks. For example, an attack may include the insertion of a backdoor or Trojan program on the target machine for use in later nefarious undertakings. Self-propagating worms commonly implement multiple exploits in order to increase their potency and longevity.

Due to the volume of attacks being produced daily and the growing complexity of such attacks, maintaining confidentiality, integrity, and availability for security managers is a daunting challenge, as threats manifest themselves in complex, varied, and adversarial ways. Cyber attacks occur continuously, and stories of resulting widespread network outages and system compromises have become a part of daily news. A significant denial of service attack concerning the Web infrastructure company Akamai caused a severe ripple effect, resulting in a temporary loss of availability for Google, Microsoft, Yahoo!, and Apple [5]. Intruders even accessed Cisco Systems corporate network and stole an undetermined amount of the company's Internetwork Operating System (IOS) source code [16]. This loss of confidentiality not only raises questions about the effectiveness of Cisco's security measures, but it raises the prospect that hackers may study the code to identify its weaknesses. This being the case, it follows that organizations must invest in a parallel education for their network administrators in order to level the playing field. Likewise, providing a high fidelity simulation through which attack identification and policy handling can be carried out is paramount in order to garner the greatest gain.

The next section will present our The Work Centered Interface for Computer Network Defense (WCI-CND) as a foundational platform on which such simulation and training infrastructures may be built.

3 Simulating Network Defense

The Work Centered Interface for Computer Network Defense (WCI-CND) is an extensible technology for integrating best-of-breed cyber-defense technologies and performing intelligent information fusion, correlation, and policy-based decision support. We initially developed the Work Centered Interface (WCI) technology in response to significant challenges for situation awareness posed in information rich environments. It is designed to provide organizations an effective tool for aggregating, correlating cyber-defense knowledge using multi-agent knowledge based technology [14, 20]. For groups of autonomous or semi-autonomous software agents to support achievement of human goals, the contributions of individual agents must be coordinated and optimized to support system objectives. By allowing for such collaboration through the decoupled architecture afforded by intelligent agents, new information can be seamlessly correlated, and contribute to the goal driven achievement of tasks.

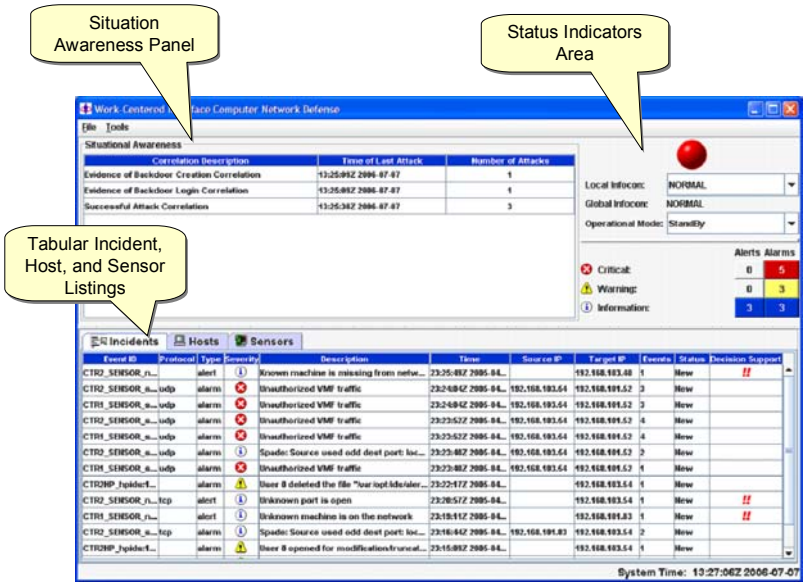

Figure 1: WCI-CND Console

The result of integrating network vulnerability, exploit, intrusion, attack and defense data with attack pattern and network policy information and using rule-based reasoning to reflect near-real time network posture and health is an advanced capability for Computer Network Defense (CND) situation awareness and decision-making. Situation awareness and decision support resulting from the synthesis of the aforementioned data into knowledge are reflected in the graphical user interface (GUI). The GUI, as shown in Figure 1 gives the security manager the high level and the detail level views of CND activity. From the vantage point of the WCI-CND console, the security manager can take in the overall network state at a glance.

The top part of the window is organized in two major areas, the Situational Awareness Panel and the Status Indicators Area. In the Status Indicators Area, there are summary indicators as to overall alert and alarm levels of varying degrees, as well environment parameters such as INFOCON and Operation Mode. The bottom half of the console is organized as three tabs: Incidents, Hosts, and Sensors. In the Situation Awareness Panel correlations are listed. These items are selectable to explicate all information available, supporting research and analysis.

The Incidents tab displays a listing of current and active alarms and warnings. Information provided for each incident includes the incident description, severity, source and target IP address, status, decision support, and other salient details associated with the incident. The user may open a window showing event details, and if the incident represents a consolidation of multiple events, the security manager may access a listing of these events. The hosts tab lists the target platforms in the given network, and includes information such as host name, IP and MAC address, operating system, location, and the number of alarms and warnings associated with the host. The sensors tab lists all sensors in the network. Information for each sensor includes sensor name, version, sensor type, location, host, and scan control (for sensors that have a scanning capability).

Global security status indicators enable the security manager to focus on the task of maintaining the security of the network instead of sifting through huge volumes of event data to locate and identify the important pieces while switching between user interfaces, products, and machines. This security status is depicted graphically in the WCI-CND with color coded indicators corresponding to network status (see Figure 1).

Visual indication of network security issues may be augmented by audible alarms to ensure a security manager does not miss a threatening network condition. These visual and audible indicators are backed by textual information summarizing the situation awareness picture of the network and consolidated sensor event information. The security manager can access the details from consolidated event information as needed. The WCI-CND GUI also provides security managers with decision support recommendations that may be taken to resolve network security issues.

This single integrated view of information is collected from across an entire network and integrated with policy and system configuration data. Processing this information results in a higher level of network security situation awareness because the underlying agent architecture supports the consolidation and correlation of active network events to construct a complete view of relevant network vulnerabilities and threat conditions.

The WCI-CND is implemented using Sentar's KnoWeb® multi-agent technology (see Section 3.2 for greater detail). From an implementation perspective, the WCI-CND is a multi-agent knowledge based system for highly interactive task-focused decision support and dynamic decision-making. As shown in Figure 2, the WCI-CND architecture consists of the KnoWeb core agents, a set of workflow agents, a set of rule-based agents, an active response agent, an extensible collection of sensors agents and sensor wrappers, the console agent, a database agent, and a set of auxiliary agents.

The sensor wrappers communicate directly with the sensors and report CND-related events to the sensor agents. The sensor wrappers use an industry standard messaging format for reporting events. This format, called the Intrusion Detection Message Exchange Format (IDMEF) was developed by the Internet Engineering Task Force (IETF). IDMEF is used by a variety of sensors, including Snort and Nessus, to facilitate XML-based interchange of cyber-defense information across the network. The sensor agents map this information into the system domain and report it to the WCI-CND. From there, the core agents marshal incident processing, calling upon the knowledge synthesis workflow agents to invoke appropriate rule agents for event consolidation, correlation, decision support, and knowledge storage. The console agent then requests updates from the knowledge store using a publish-subscribe messaging protocol.

Several network security technologies have been wrapped for inclusion to the framework, including IDSs, vulnerability scanners, network mappers, anomaly detectors, and integrity checkers. This infrastructure provides a critical mass of security posture information to be manipulated in the platform during operation.

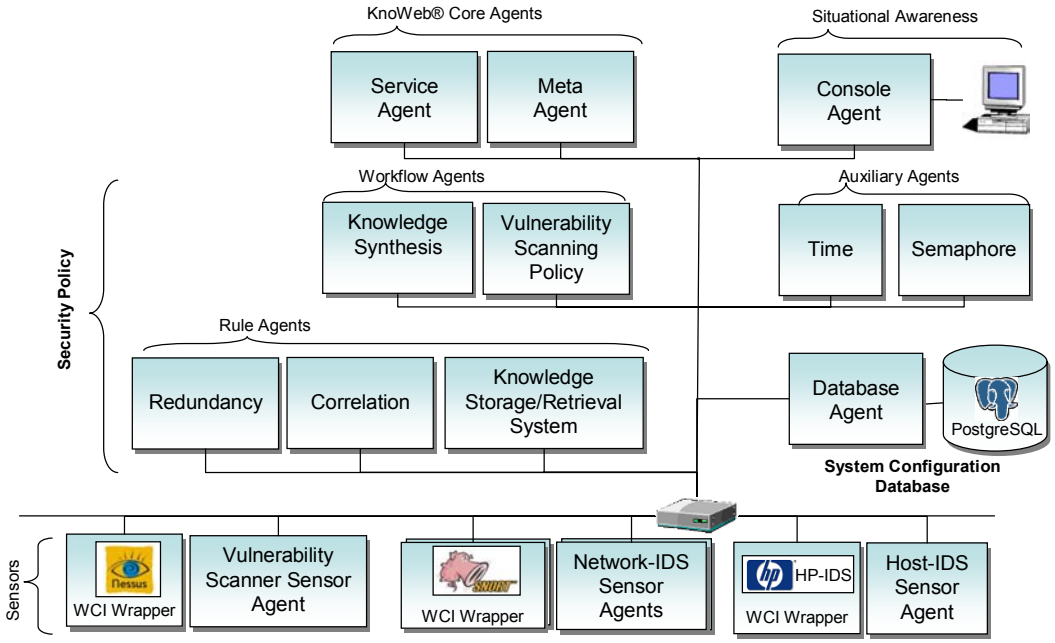

Figure 2: WCI-CND Architecture

3.1 Underlying Architecture

The WCI-CND is implemented as a CND-specific application of the more general WCI technology. The WCI was developed in response to challenges posed by distributed multi-agent systems for human-computer collaboration. The software solutions made feasible by multi-agent technology, particularly when integrated with distributed knowledge-based systems, were scarcely conceivable just a few years ago. Taking advantage of high-speed global networks, these systems can address a range of problems previously beyond the reach of any technology. Intelligent agents can perform cooperative problem solving in a manner that is best described as technological teamwork. And yet, for this to be of value, these capabilities must be harnessed to serve human needs.

The challenges addressed by WCI arise with the realization that as the functions of intelligent agents

become increasingly integral to human decision-making, the relationship between humans and computers becomes increasingly collaborative. For groups of autonomous or semi-autonomous software agents to support achievement of human goals, these goals must be understood at a system level, and the contributions of individual agents must be coordinated and optimized to support system objectives. Under such conditions, the discourse of interaction is not merely a matter of graphical design, mental models, or metaphor selection. As systems become more intelligent, the scope of human-computer collaboration deepens.

Thus the focus of the WCI is on providing intelligent problem-solving systems using advanced multi-agent technology. Systems employing work-centered technology are able to harness loosely coupled agents and coalesce their capabilities in support of task-oriented goals. This is accomplished through a novel integration of technologies including mediated reasoning, agent workflow management, and ontology-driven knowledge representation. The framework used to implement these capabilities is the KnoWeb multi-agent architecture.

3.2 KnoWeb®

KnoWeb® is a general-purpose multi-agent system architecture that uses mediated reasoning to perform dynamic decision-making and event driven decision support [14, 20]. As shown in Figure 2, KnoWeb® employs a small group of core agents to implement its reasoning model. These core agents engage a loosely coupled collection of specialist agents to carry out goal-driven and event-driven tasks. The core agents use the results of these tasks to synergistically coalesce and maintain higher order situation awareness and decision support. The core agents consist of a Meta Agent, a service agent, and one or more domain advisors. The Meta Agent provides domain-neutral mediation and conflict resolution. By centralizing reusable intelligence, redundant complexity in domain and ancillary components is reduced, allowing domain dependent information to be fully encapsulated. The service agent maintains a registry of agent capabilities used to provide brokerage services. The domain advisor provides conflict resolution and planning strategies used by the Meta Agent. Agents communicate using Sentar's Knowledge Agent Mediation Language (KNAML) [19].

KnoWeb mediated reasoning supports subsumption, unification, and binding. Subsumption uses existential conjunctive logic to establish the truth value of one graph based on the known value of another. Graph and subgraph unification support discovery by indicating how one graph is subsumed by another. Binding is used for joining graphs to produce knowledge synthesis. The binding operation also supports backtracking. This is necessary to assure that graphs are recoverable in the event of binding failure.

The reasoning process used by the Meta Agent is straightforward. Throughout the process, the Meta Agent uses an agenda to keep track of what it is doing and why, and it maintains a context of asserted propositions. The process consists of several phases: initiation, alliance, marshalling, resolution, and response. The process begins when an agent initiates a request. The agent does so by sending the request to the Meta Agent. When the Meta Agent accepts a request, it first checks to see if the answer is already in the context or if the request is already on its agenda. If the answer is already in the context, the Meta Agent uses it to generate a response. Otherwise, it proceeds with problem-solving.

The KnoWeb mediated reasoning process and agent encapsulation technology can be used to integrate the capabilities of heterogeneous information resources in support of high order reasoning to achieve a common set of goals. These advanced features also make KnoWeb highly extensible, so that it can be readily integrated without re-engineering established capabilities.

Agent communication and knowledge representation in a KnoWeb® application are ontology-based. That is, all application concept and relation types are specified a priori. Ontologies are also used to define the knowledge modalities used by the agents. For example, workflows are structured using a workflow ontology and rule bases are compliant with a rule ontology. Furthermore, goals of the problem solving system are specified ontologically. There is also an ontology of ontologies for use in defining new ontologies.

KnoWeb® applications use the Knowledge Agent Mediation Language (KNAML) developed by Sentar [19]. KNAML is based on conceptual graphs and represents knowledge using concepts, relations, and graphs. Concepts and relations may be linked together to form graphs, and graphs may be nested within other graphs or concepts. All concepts and relations are ontologically typed. KNAML supports the semantic description of knowledge in various forms, including rules, frames, decision trees, and databases.

KNAML may be viewed as an alternative to semantic content languages such as RDF, KIF, and OWL. In some cases, ontologies can be shared between KNAML and other representation languages; however, interoperability with more expressive languages may be limited. A benefit of KNAML is that it was designed specifically for use in multi-agent reasoning systems and is streamlined for agent communications and efficient machine interpretation.

The next section discusses the unique characteristics that make WCI-CND appropriate for integration into a simulation infrastructure.

3.3 Applying Simulation

Because it manifests both intelligent agents and a highly decoupled configuration, the WCI-CND provides an effective, flexible, and extensible integration platform. The use of intelligent agents gives the system advanced capabilities for accomplishing automated correlation and decision support. The use of rule-based processing enables the system to adapt to evolving policies and technological capabilities. The modular nature of the multi-agent system provides flexibility in resource deployment—the agents may be installed all on a single host, or they may be distributed across multiple hosts. Importantly, the use of a wrapper technology means that new components can be introduced with minor cost and minimum perturbation to the system. The platform also manifests a script injection capability that executes structured events to which the system can react. Thus, it can be termed "simulation-ready" in that, if fed events, it will execute according to the sensor and knowledge-based component configurations.

The key to utilizing WCI-CND as a platform for intrusion detection simulation and training resides in its flexible architecture. The integration of traffic generators and attack model servers require only the specification of an agent representative. Considerable research, definition and development of such components have been conducted in the area of network simulation as it relates to characteristic measurement and training. Fundamental research in modeling attacks, including trees, patterns, and discrete event representations, provide a foundation from which simulated intrusion events and detection measurement and training can be built. Likewise, network simulators and situational awareness tutors have been developed to replicate network operations for security analysis and education [2, 3, 6, 7, 9-12, 17, 18]. This research provides building blocks that can be integrated into a broad infrastructure like WCI-CND to provide a comprehensive and realistic simulated network defense infrastructure and training platform.

The discrete-event Scalable Simulation Framework (SSF) provides a C++ and Java API as well as protocol models to develop a network simulation platform. Several instantiations have been implemented for experimentation in real-time processing, performance, routing and intrusion detection (http://www.ssfnet.org/homePage.html). One C++ implementation, iSSF, uses model abstraction and parallelism to provide a network simulation that can pre-compute situational outcomes to address faster than real-time challenges [13]. This would include systems that use simulation within their regular processing to provide information to real-time decision-making algorithms.

The PRIME project uses iSSF as their network simulation platform as well as the Real-time Immersive Network Simulation Environment (RINSE) to simulate large-scale, distributed real-time network operations [9-11, 13]. Multiple areas of research are conducted using PRIME, including hybrid network traffic modeling, and network emulation. RINSE itself was developed in conjunction with the Univversity of Illinois – Urbana and utilized to provide network intrusion simulation platform [9, 10]. Attack models in RINSE, which include DoS and worms, concern network resource level assets rather than detailed traffic dynamics.

The Rapid Model Parameterization (RAMP) network traffic model is a network simulation tool that takes traffic traces and generates models off-line to be executed in network traffic experiments [7]. RAMP has also been augmented to simulate real-time traffic patterns through parallelism of the RAMP framework. The real-time traffic generated can then be visualized to provide operators with a live snapshot on which traffic engineering tasks can be executed.

Network defense training platforms have also been developed [2, 3, 12, 17, 18]. The US Naval Academy investigated the pedagogical effectiveness of using an attack simulator combined with multiple network simulators and viewers to educate students on attack recognition, differentiation and mitigation [3]. Training with the platform was introduced both as an immersive and as a passive mechanism in the classroom and produced a 31% increase in knowledge and retention.

A different but complimentary approach to simulation for network security education is Military Academy Attack/Defense Network (MAADNET). Security [2]. MAADNET manifests a client-server architecture where users can build, test, and submit network designs. MAADNET then grades the performance and security of the design through event and attack simulation.

NetSim is a fluid flow simulation model [12] combining a simulation engine with internet and database technologies developed to support the TOPOFF2 and TOPOFF3 cyber exercises as well as the Livewire national cyber exercise in 2003. The goal of the platform is to simulate attacks at a level appropriate for Network Operations Center personnel to detect them and report their findings upward.

Akin the network sensors integrated into WCI-CND, technologies like RINSE could easily be integrated with WCI-CND to proffer both volumes of coarse-grained simulated network traffic and according to realistic traffic intensive attack scripts. Likewise, RAMP could provide a complimentary, trace driven study of malicious traffic within WCI-CND. The combination of traffic generators and attack models with WCI-CND's visualization could also serve as a classroom based training platform like that used at the Naval Academy. However, WCI-CND provides a layer of situation awareness for students such that they could study how they should be interpreting events and can then take appropriate action. The decision support capabilities would subsequently provide them with a gauge of the correctness of their action during the intrusion detection exercise.

4 Conclusions

Research blending intelligent agents and training and simulation techniques indicate agents provide a suitable platform to manage and participate in the simulation [15]. Using the dynamism of agents, while incorporating an intelligent situation awareness enhancement mechanism for dynamic feedback as well as proven intrusion sensor technology, provides a powerful and realistic medium on which to base training. Sentar's KnoWeb®-based WCI-CND technology presents such a framework on which to build an intelligent agent training simulator. KnoWeb is a multi-agent integration platform that uses mediated reasoning to perform dynamic decision-making. WCI-CND, in turn, leverages this technology to integrate a critical mass of best of breed network security technologies to carry out goal-driven and event-driven situational awareness and decision support. This serves as an infrastructure through which varied and unlimited network, attack, training, and validation knowledge sources and programs may be integrated to form an immersive and realistic simulated network defense infrastructure.

5 Acknowledgements

The research and development presented in this paper was supported by the Missile Defense Agency through multiple SBIR grants at the Phase I, II and III level.

6 References

[1] M. Bishop, *Computer security: art and science*. Boston, MA: Addison-Wesley, 2003.

[2] C. Carver, J. Surdu, J. Hill, D. Ragsdale, S. Lathrop, and T. Presby, "Military Academy Attack/Defense Network Simulation," presented at 3rd Annual Information Assurance Workshop, United States Military Academy, West Point, New York, 2002.

[3] L. L. DeLooze, P. McKean, J. R. Mostow, and C. Graig, "Incorporating Simulation into the Computer Security Classroom," presented at 34th Annual Frontiers in Education (FIE 2004), 2004.

[4] V. D. Gligor, T. Haigh, D. Kemmerer, C. Landwehr, S. Lipner, and J. McLean, "Information Assurance Technology Forecast 2005," *IEEE Security and Privacy*, vol. 4, pp. 62-69, 2006.

[5] J. Hu, "Denial-of-service attack causes web blackout," in *silicon.com*, June 16 ed, 2004.

[6] S. Krasser, G. Conti, J. Grizzard, J. Gribschaw, and H. Owen, "Real-time and forensic network data analysis using animated and coordinated visualization," presented at 6th Annual Systems, Man and Cybernetics (SMC) Information Assurance Workshop, 2005.

[7] K. Lan and J. Heidemann, "A tool for RApid model parameterization and its applications," presented at ACM SIGCOMM workshop on Models, methods and tools for reproducible network research, Karlsruhe, Germany, 2003.

[8] C. B. Lee, C. Roedel, and E. Silenok, "Detection and characterization of port scan attacks," vol. 2004. San Diego, CA, 2003.

[9] M. Liljenstam, J. Liu, D. M. Nicol, Y. Yuan, G. Yan, and C. Grier, "RINSE: The Real-Time Immersive Network Simulation Environment for Network Security Exercises (Extended Version)," *SIMULATION*, vol. 82, pp. 43-59, 2006.

[10] M. Liljenstam, D. M. Nicol, V. H. Berk, and R. S. Gray, *Simulating realistic network worm traffic for worm warning system design and testing*. Washington, DC, USA: ACM Press, 2003.

[11] J. Liu, "Packet-level integration of fluid TCP models in real-time network simulation," presented at 2006 Winter Simulation Conference (WSC'06), Monterey, CA, 2006.

[12] D. McGrath, D. Hill, A. Hunt, M. Ryan, and T. Smith, "NetSim: A Distributed Network Simulation to Support Cyber Exercises," presented at Huntsville Simulation Conference, Huntsville, AL, 2004.

[13] D. M. Nicol, J. Liu, M. Liljenstam, and Y. Guanhua, "Simulation of large scale networks using SSF," presented at 2003

Winter Simulation Conference, New Orleans, LA, 2003.

[14] A. Potter and G. Streeter, "Work-centered services for the semantic Web," presented at 3rd International Symposium on Multi-Agent Systems, Large Complex Systems, and E-Businesses (MALCEB'2002), Erfurt/Thuringia, Germany, 2002.

[15] J. Rickel and W. L. Johnson, "Extending Virtual Humans to Support Team Training in Virtual Reality," in *Exploring Artificial Intelligence in the New Millennium*, G. Lakemeyer and B. Nebel, Eds. San Francisco: Morgan Kaufmann Publishers, 2002, pp. 217-238.

[16] P. Roberts, "FBI investigating Cisco source code leak," in *ComputerWorld*, May 15 ed, 2004.

[17] N. C. Rowe and S. Schiavo, "An intelligent tutor for intrusion detection on computer systems," *Computer Educ.*, vol. 31, pp. 395-404, 1998.

[18] J. Saunders, "Simulation Approaches in Information Security Education," presented at The 6th Nat'l Colloquium for Information Systems Security Education, Redmond, Washington, 2002.

[19] G. Streeter and A. Potter, "KNAML: A knowledge representation language for distributed reasoning," in *Conceptual Structures at Work*, K. E. Wolff, H. D. Pfeiffer, and H. S. Delugach, Eds. Berlin: Springer-Verlag, 2004, pp. 361-374.

[20] G. Streeter, A. Potter, and T. Flores, "A mediated architecture for multi-agent systems," presented at Seventeenth International Joint Conference on Artificial Intelligence: Workshop on E-Business and the Intelligent Web, Seattle, WA, 2001.

Biographies

Dr. Leigh Flagg is a Senior Research Scientist at Sentar, Inc., serving as a coordinator and technical contributor on many software architecture, computer security, and knowledge-based systems research and development projects. She is certified by the Committee on National Security Systems (CNSS) as an information systems security professional. Her technical areas of expertise include software integration, distributed systems, formal modeling and verification, secure system administration and audit, policy verification, web service enterprise information systems, and middleware.

Gordon Streeter is currently the Director of Technology for Sentar, Inc., where he is the lead designer and developer of KnoWeb®, a multi-agent technology for distributed, dynamic problem solving, and a co-inventor of KNAML, a graph-based language and environment for knowledge representation and reasoning. Mr. Streeter has broad interest in Cognitive Science, particularly in knowledge representation and cognitive processing.

Andrew Potter is the Chief Scientist at Sentar, Inc. and has been with the company for seven years, where he has led numerous research and development projects in the areas of multi-agent systems, knowledge acquisition, and knowledge representation. His interests include ontology, reasoning, and rhetorical structure theory. Mr. Potter is currently completing his PhD in Information Science from Nova Southeastern University.