



A Scenario-Based Approach to Mitigating the Insider Threat

By Robert F. Mills, Michael R. Grimala – ISSA member, Dayton, USA Chapter
Gilbert L. Peterson, Jonathan W. Butts

Join the Discussion
Connect

This article presents a risk management strategy tailored specifically for the insider threat.

Abstract

This article presents a risk management strategy tailored specifically for the insider threat. Key components include the emphasis on people, process, and technology, organizational introspection into critical information resources, and a scenario-based approach to determining how insiders might attack those resources. Risks can then be mitigated through a variety of control mechanisms in a manner that does not introduce unnecessary costs.

Background

The insider threat is perhaps the most significant challenge facing information technology (IT) and security practitioners. For the purpose of this article, the insider threat is defined as “...intentionally disruptive, unethical, or illegal behavior enacted by individuals who possess substantial internal access to the organization’s information assets.”¹ This includes current or former employees, contractors, or other trusted business partners. Theft of intellectual property and proprietary information by employees remains a top form of financial loss and other damages, such as loss in reputation. The recent WikiLeaks scandal – in which volumes of sensitive documents were leaked by a trusted insider and ultimately published on an open website – has caused

much embarrassment to the United States and other nations and represents the ultimate nightmare scenario when considering the insider threat problem.²

The insider threat is complex due to a variety of contributing factors and bears characteristics of a *wicked problem*³ because attempts to “solve” the problem may actually exacerbate it or introduce other problems. Malicious insider activity may be indistinguishable from normal actions, and attacks are difficult to detect until after damage has occurred. Most insider attacks are planned, however, and a window of opportunity exists during which people can intervene and prevent the attack, or at least limit the amount of damage. However, with the focus on lean management, supervisors have less time and are likely to overlook potential warning signs.

A holistic approach that blends people, process, and technology would be extremely useful in helping managers focus more on the behaviors and activities that appear to be risky if not outright malicious. The process outlined in this article

**Most insider
attacks are
planned.**

1 J. M. Stanton et al., “Analysis of end user security behaviors,” *Computers & Security*, March 2005.

2 K. Coleman (2011, January 13), “Wikileaks scandal raises many questions,” <http://defensesystems.com/articles/2011/01/24/digital-conflict-wikileaks-raises-questions.aspx>.

3 H. Rittel et al., “Dilemmas in a general theory of planning,” *Policy Sciences*, June 1973.

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 MAY 2011		2. REPORT TYPE		3. DATES COVERED 00-00-2011 to 00-00-2011	
4. TITLE AND SUBTITLE A Scenario-Based Approach To Mitigating The Insider Threat			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) Information Systems Security Association, Inc.,9220 SW Barbur Blvd #119-333,Portland,OR,97219			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 ISSA Journal, May 2011					
14. ABSTRACT This article presents a risk management strategy tailored specifically for the insider threat. Key components include the emphasis on people, process, and technology, organizational introspection into critical information resources, and a scenario-based approach to determining how insiders might attack those resources. Risks can then be mitigated through a variety of control mechanisms in a manner that does not introduce unnecessary costs.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			




Figure 1 – Observables for mitigating the insider threat

is based on a combination of conventional risk management, functional analysis of insider behaviors to develop threat scenarios, and evaluation and selection of control measures to mitigate specific insider risks. Cost/benefit metrics are discussed to illustrate potential trade-offs and to help managers decide how a particular scenario might be mitigated in the most efficient or cost-effective manner. While each of these concepts is not new or unique, this article presents an approach for how they can be used by any organization to address the insider threat in a practical manner.

Observable behaviors

Figure 1 illustrates some potential observables that could indicate malicious insider behavior.⁴ Security clearances and background checks are routinely used by government agencies and some businesses to determine an individual's trustworthiness and personal character prior to granting that person access to sensitive information. These background checks may also provide a deterrent for current employees, because they will be less likely to engage in undesirable behaviors when they know they are subject to increased scrutiny.

The lower half of the figure contains many "cyber actions" that can be observed and fused to develop a picture of employee behavior. Although the cyber actions provide limited insight into one's intent and character, they are much easier to automatically collect, process, and correlate than non-cyber ones. Some of these activities can be collected via standard system auditing, while others will require specific sensors. The ease with which these cyber actions can be monitored is a dual-edged sword, because the amount of raw data that can be collected may actually exacerbate the problem.

Some cyber actions can easily be tied to malicious behavior, such as installing unauthorized software or violating established acceptable use and security policies. In other cases,

it is difficult to separate a malicious act from normal everyday business, such as printing a document, browsing the Web and searching databases. Someone who has been assigned to a new project will likely be gathering information to learn more about the project and would probably be doing many of these activities. Thus, a change in behavior is not necessarily an indicator of malicious intent. Likewise, system and network administrators routinely perform network scans, install new software,

change file permissions, and manipulate log files. Writing data to removable media may also be normal and expected behavior, even though it represents a significant risk factor, as was the case with the WikiLeaks scandal.⁵

While system logs are clearly useful for insider threat detection, they were not designed and optimized for this purpose. Raw log data is typically limited to information such as user name, workstation identifier, IP address, time/date, and a brief description of the event. Significant processing and fusing are required to establish the contextual linkages between low-level system activities (keystrokes, file access, system usage) and higher-level concepts such as motive and intent. There will rarely be a specific event in a system log that indicates malicious behavior. For example, many log events (possibly from multiple sources over a long time period) would be required to determine that an individual has recently printed large numbers of documents using multiple printers throughout the facility.

While there has been much research into the technical and behavioral aspects related to the insider threat, there has been little research to address linkage between the two areas, leaving a "semantic gap." Low-level system audit data lacks contextual factors: Is the person working after hours, and if so, is there a good reason? Is the individual working on an important project with a deadline, such as a contract proposal? What is the person's job function (research engineer, budget analyst, systems administrator)? Do the observed cyber events indicate a violation of organizational policy? These contextual factors are necessary in determining whether further investigation is needed. Human analysis, to include cooperation among various business functions (security, human resources, legal affairs, supervisors, IT management, etc.), will be required to validate whether suspicious activities identified are truly malicious in nature.

4 Adapted from M. Maybury, "Detecting malicious insiders in military networks," MITRE Corporation, 2006.

5 Coleman, 2001.

Scenario-based insider risk management

Effectively dealing with the insider threat is complicated by many factors: decision makers must act with imperfect knowledge; everything cannot be protected all the time; it is not known what to look for until after something has happened; and it is impossible to plan for and deal with every possible contingency. We contend that a scenario-based risk management approach is an effective way to deal with them.

Too many organizations focus on security solutions without a clear understanding of what they are protecting and why.

Risk management is an analytical methodology used to evaluate trade-offs in protection strategies when mitigating risks subject to organizational constraints.⁶ The primary function of risk management is to assign protective measures to assure the ability of the organization to conduct its mission. Risk management includes risk assessment, risk mitigation, and evaluation/assessment.⁷ Collectively, these processes enable managers to identify and evaluate

risks so they can make informed decisions on how to mitigate those risks. In general, the insider risk mitigation process does not differ substantially from traditional risks associated with external attacks and/or environmental disasters and is discussed in the following steps.

Identify critical information resources

Understanding the mission and how information is used to support it is essential to an effective risk management process. This sounds obvious, but it is the authors' experience that too many organizations focus on security solutions without a clear understanding of what they are protecting and why. The criticality of a resource ultimately depends on the value it provides in the context of the overall mission. Everything cannot be afforded the same level of protection; given limited resources, information must be protected according to its assessed value (e.g., monetary worth, competitive advantage, consequences of loss or destruction, etc.).

Develop attack scenarios and identify observables

After a critical information resource has been identified, the organization must reason through one or more scenarios in detail, considering what the attacker would try to achieve (and why), and how the attacks could be carried out. This bounds the problem and allows focus on a smaller subset of behaviors, rather than trying to defend against everything. It is possible that a different type of attack may occur, but

the organization will be better off having gone through the exercise than if it had not done any analysis at all. Over time, the organization can build its knowledge base and refine scenarios based upon real world events within, and external to, the organization.

Scenarios should be objective-based (steal information, impede/disrupt business operations, etc.). Critical events that must occur within the scenario for the user to achieve the desired goal are identified and assessed, based on their detectability via one or more sensors (e.g., event logs or other cyber sensor). Finally, if one or more events are logged that indicate potential malicious activity, a subsequent, more rigorous investigation can be conducted to reveal if the event was related to malicious behavior or simply resulted from accidental or normal user behavior. To minimize false positives, each scenario should consider both the normal user and malicious user behavior to be detected. To the extent possible, non-cyber observables should be included.

There is no one "best" way to develop the attack scenarios. Documented insider threat cases represent a wealth of information and can be tailored as necessary. Lacking any pre-existing knowledge, an organization can develop attack scenarios using a functional decomposition process. This facilitates grouping similar activities in such a way as to minimize redundancy when identifying observables and placing security measures. The concept is similar to a modified attack tree whose purpose is to identify an attacker's goals and then specify the different ways in which those goals might be achieved. Attack trees alone are of limited use here because they do not address the distinction between legitimate and undesirable activities if an attacker uses authorized permissions.⁸

For example, insider threat behaviors could be categorized into four general types of activities:

- **Alteration:** The insider modifies data or system parameters in an unauthorized manner (e.g., deletes a file from the system or adds a covert user account)
- **Distribution:** Insider transfers rights to an unauthorized party (e.g., sends a restricted document to someone without legitimate access to the information)
- **Elevation:** User obtains unauthorized rights in the system (e.g., elevates privileges to system administrator access)
- **Snooping:** Seeks access to unauthorized information (e.g., browsing shared file systems or databases)

These threat classes are derived from Pfleeger and Pfleeger,⁹ but other categorization schemes can be used with successful results. Threat-actions are then decomposed step-by-step, beginning with the top-level category and continuing with

6 T. Finne, "Information systems risk management: Key concepts and business processes," *Computers & Security*, January 2000.

7 *Risk Management Guide for Information Technology Systems*, National Institute of Standards and Technology Special Publication 800-30, July 2002.

8 J. Butts et al., "Developing an insider threat model using functional decomposition," *Proceedings of the 2005 Mathematical Methods, Models, and Architecture for Computer Network Security Workshop*.

9 C. Pfleeger and S. Pfleeger, *Security in Computing: Fourth Edition*. Prentice Hall PTR: Indianapolis, IN, 2006.

can you control who has access to what?

Finding ways to easily and securely control your IT environments — physical, virtual and cloud — while also addressing your compliance requirements is crucial to your business success.

You can get that level of control from CA Technologies Content-Aware Identity and Access Management. It goes further than traditional security solutions by giving you control all the way down to the data level.

It gives you the ability to take control of your users, their access and their information use so you can easily answer the question: “Who has access to what?”

Take control of your IT security today. Start here: ca.com/security



you can



intermediate levels to final leaf nodes. A leaf node is the lowest level of abstraction and represents the tool or technique that a malicious insider might use. The process is illustrated in Figure 2, which shows activities associated with disseminating a sensitive file (*Distribution*). The figure shows only a partial decomposition and does not include complex attacks. For example, an insider could print a document and then scan/email it.

Another example might be to use a screen capture and then paste the image into another document which is emailed to an external party. As stated earlier, introspection will be required for any organization to understand and manage its insider threat risks.

This process provides insight into gaps and overlaps between organizational and system policies, which in turn allows development of appropriate mitigation strategies and audit sensors. The strength of this approach is that insider behaviors are systematically analyzed in the context of the organization's mission and business processes. Further, emphasis is placed on observable behaviors – i.e., something happened and was detected in one or more logs.

Analyze risk factors

The third step is to evaluate the risk from insider activity to the organization's critical information and key processes. Risk mitigation is an analytical process that involves prioritizing, evaluating, and implementing controls to mitigate risks to an "acceptable" level. There is no universally established accept-




Figure 2 – Partial scenario tree and placement of security mechanisms

able level of risk – different organizations may have different tolerance levels based upon their risk preference, historical events, resources, and/or other priorities. What is acceptable for one organization may be completely unacceptable for another organization.

Risk is often depicted as the product of likelihood and mission impact. For example, critical risks are those that are quite likely to occur and will have a significant impact – loss in revenue, damaged reputation, mission failure, etc. Once the various risks have been determined, they can be rank-ordered in level of severity. Control measures can then be evaluated on the basis of their ability to prevent, detect, or mitigate threats so they can be implemented to address the behaviors that are deemed most dangerous.

ISSA Connect – The discussion begins...

Join the Discussion
Connect



Pete
Lindstrom

Benevolent Interdiction? Hi, I'm from the government, and I'm here to help.

Did you miss the news that the FBI was granted permission to take over the Coreflood botnet and temporarily disable the software on 2 million unsuspecting users (purportedly all in the U.S.). Here's the rundown:

The government 1) sued some John Does; 2) took over domain names; 3) replaced command and control servers; 4) issued 'stop' commands; and 5) provided IP addresses to applicable ISPs, all while guaranteeing that this wouldn't hurt the end user PC. I have a mixed reaction: In a specific sense, I am glad they did it because this particular malware is stealing passwords and private information from its victims. And I suspect most of the victims will be glad they did this as well. But in a general sense, I really don't like the idea of the government getting involved in stopping processes since it is easy to see how this can escalate. At some point, somebody is going to ask why they are only 'temporarily' stopping the bots (they restart on boot) and then someone else is going to sue because the restarted service actually *did* compromise their privacy and another person is going to sue when the government *doesn't* disable their particular botnet...



Carl Staab

Re: Benevolent Interdiction? Hi, I'm from the government, and I'm here to help.

One very basic question is who do you trust more: the government or those that created the botnet? Once the government is in control of millions of PCs and see what power they hold, are they going to relinquish it?

...and continues with you.

Establish control measures

Finally, selected risks are mitigated through a combination of control measures (or mechanisms), that include anticipation, prevention, deterrence, detection, and response. A mechanism is any countermeasure or process employed to detect or counter a violation to the protection state. This includes detection methods (e.g., auditing and intrusion detection systems), automatic blocking of websites, and disabling specific network/host services or functions (e.g., administrative privileges are required to install software on a host computer). The dashed lines in Figure 2 depict possible placement of control measures.

The actual function performed by a mechanism could range from outright prohibition or denial of that capability to content filtering to increased auditing. It depends on the perceived threat and the extent to which the organization wants to minimize the risk. For example, Mechanism B addresses web browsing via the HTTP network protocol. Web browsing could be prevented outright by blocking HTTP traffic, or a proxy server could be used to limit browsing to designated individuals and monitor them closely. White lists (authorized websites) and black lists (unauthorized sites) could also be used. Finally, focused monitoring and detailed logging could be employed for individuals deemed at risk, perhaps because of suspicious behavior or critical job function. These solutions have costs and benefits, as will be discussed in the next section.

Control measures should include a mix of people, process, and technology. Everyone – employees, managers, and security professionals – has a role in mitigating the insider threat. Procedural controls, such as separation of privilege and compartmentalization based on “need to know,” can minimize potential damage by limiting what a single individual can do with the access that has been granted. While technology solutions are important, they alone are not sufficient.

Cost versus benefit analysis

Organizations should invest the time and energy into fully understanding their risks, what they are trying to protect (and why), and what it “costs” to protect critical resources. Managers must make informed decisions when striking a balance between the costs of protective measures (prevention, deterrence, detection, and response) and the benefits provided by protecting the organizational mission.

Security software and hardware cost money. Reviewing audit logs and following up on events of interest take time away from primary business functions. There is a cost associated with collecting too much data (storage and analysis), and there is a cost with not collecting enough (missed detections). Intangible costs include the human response – some employees might resent being under the watchful eye of “Big Brother” and may resort to passive-aggressive behavior. Others might attempt to work around security measures that are perceived as being unnecessary or counter-productive; while their intent might not be malicious, their actions could intro-

duce other undesirable risks

With many potential trade-offs involved, a “best” solution is open to interpretation. Difficult decisions are required to achieve a balance between being “secure enough” and being able to get the job done. A cost-benefit analysis approach could be used to derive an overall metric, which would then provide management with enough information to make these decisions. For example,

$$FOM_{XY} = \sum_j W_{Bj} B_{XYj} - \sum_k W_{Ck} C_{XYk}$$

where FOM_{XY} represents a figure of merit (FOM) for a given control solution X and scenario Y . B_{XYj} and C_{XYk} represent benefit and cost factors, respectively, associated with each X and Y – there may be multiple cost and benefit factors associated with each control/scenario pair. Optional weighting terms, W , can also be used to emphasize some costs or benefits more




Figure 3 - Calculating figures of merit



Design Your Future – Security 2020

October 20-21, 2011

Baltimore Convention Center

We look forward to seeing all members at the 2011 ISSA International Conference in Baltimore, Maryland, USA. Mark your calendars now and plan on joining us.

- Chapter Leaders Summit October 19*
- International Conference October 20-21
- 4th Quarter CISO Forum October 22*

Registration will be open soon. [Click here](#) for details and to be notified as soon as registration opens.

**Open to qualified attendees only.*

www.issaconference.com

than others. Benefits would include the ability to prevent, detect or otherwise mitigate a specific insider attack scenario. Cost factors might include storage space, time to review logs, and missing an attack pattern of interest.

As shown in Figure 3, FOMs for each control can then be calculated over all scenarios, and then compared to determine which control “best” addresses the list of scenarios (e.g., highest FOM). The determination of the cost factors is left to the organization and depends on the desired detail of analysis. If the cost/benefit factors are not quantifiable, non-numeric metrics (e.g., *high*, *medium*, and *low* scores) could be used.

As discussed earlier, auditing and other forms of oversight may disturb employees and might even cause unnecessary negative feelings. When implementing security controls, the organization must be cognizant of the human element. In some situations, high levels of scrutiny and oversight are to be expected, such as in the financial industry and government agencies. Other organizations may find that such intrusive oversight drives employees away. The extent to which its employees are antagonized or threatened by internal auditing is indeed one of the “costs” that must be evaluated.

Application

In this section we discuss how the scenario-based approach can be used to mitigate the insider threat. The example is somewhat generic and discusses various considerations and cost trade-offs for a business concerned with sensitive information being disseminated outside the organization. Assume a business is developing a new product and wishes to restrict access to proprietary information to a select few individuals. The information is stored on an internal server that may or may not be isolated from the Internet. Protection of the server includes a combination of internal firewalls and access controls, but a risk analysis indicates there is a potential for leakage using electronic file sharing by insiders not directly associated with the project.

Figure 2 shows an example application of the functional analysis and placement of security mechanisms. This analysis was performed from the perspective of the computer/server that holds the critical information being protected. Several mechanisms have been placed to prevent or monitor undesirable activity. Mechanism A addresses the connection to a local area network file system, while B covers actions using HTTP protocols. Mechanisms C and D address the risk of data being shared via physical means, such as printouts and writing to removable media. Finally mechanisms E and F address non-HTTP risks. The example shown is not meant to be all-inclusive and highlights the need to identify potential communications channels for the information to leak out, and then implement appropriate controls to block and/or monitor those channels. This diagram should also be updated on a regular basis.

The purpose and implementation of each mechanism will vary. Mechanism A could range from simple user authentication to denial of access to network file services. Mechanism B could involve blocking of all HTTP transactions, web proxy filtering, or content filtering – or some combination of these. Depending on the analysis of the HTTP risks, B could be separated into multiple mechanisms which perform these tasks separately. Mechanisms C and D might involve physically disabling (or removing) printer connections and external media devices. If an operational requirement exists to read/write to external media, then D could also include user authentication, two-person integrity, or increased auditing. The writable CD-ROM is left intact to allow for data backup and recovery. Mechanism E covers a broad range of non-HTTP methods of sharing information. Depending on the perceived risks and organizational mission requirements, E might be separated into a subset of more sophisticated measures addressing each of the nodes in the tree. Implementation could involve removal of client software, packet filtering, and routine scans to ensure compliance with organizational

Connect Learn Advance...Join Today!



Information Systems Security Association

For less than \$10 a month become an ISSA Member and take your career to the next level through:

Local Chapter Meetings • Face-to-Face Networking

The ISSA Journal • ISSA Web Conferences

Trusted Online Member Community

Discounts to Industry Conferences • Certification Study Courses

Continuing Professional Education (CPE) Credits

The Preeminent Trusted Global Information Security Community

www.ISSA.org

and system policy. Mechanism F mitigates the email risk – it seems unlikely that email would be allowed given the risk of a data spill, but F could be configured to allow email under certain situations and with heightened control measures.

Using Figure 3, FOMs could be derived for each control solution (combination of mechanisms), and overall metrics are tallied across the scenarios of interest. For example, when addressing the HTTP threat, each solution being considered (white listing, black listing, proxy filtering, and outright blocking traffic) would have its own column in the figure. Each FOM_{xy} would be calculated using the appropriate cost and benefit factors for how well each control addresses the risk posed by a given scenario.

The preceding discussion focused on the actual host or server that contained the information being protected. A similar process could be performed at the network level, or even focused on a specific group of people. For example, some may not actually need the ability to browse the Internet to perform their jobs, in which case limiting their access would eliminate a number of associated vulnerabilities. Likewise, another group may require the ability to read/write to removable media. In this case, training on proper procedures would be necessary, and procedures would be developed to minimize the potential for data leakage or theft.

Conclusions

In this article, we discussed the complexity of the insider threat and how it is best mitigated through a combination of people, process, and technology. The primary difficulty in dealing with insider threats is that by definition insiders are trusted, so they possess elevated privileges and insider knowledge when compared to external users. This makes it difficult to concisely characterize all of the activities that are malicious. While there are common motivational factors, such as greed and revenge, people who have been caught performing malicious acts do not fit a standard profile. The insider threat cannot be solved through technology alone, because security is at its very core a people problem.

The insider threat problem *can*, however, be approached in a straightforward manner using a combination of technology and standard security management practices, such as risk management, oversight, policy, and procedures. To do this effectively, an organization must take the time to determine what it is trying to protect, most likely and/or dangerous threats, and how much it is willing to invest to protect against those threats.

We have presented a systematic and repeatable process to help IT and security managers effectively manage the insider threat. The process begins with the development of threat scenarios that provide focus. Countermeasures, such as auditing or dynamic defense operations, can then be developed to mitigate specific risks. The process is scalable and has built-in flexibility for adapting to different organizational requirements. Security administrators can use our approach to focus on the behaviors presenting the greatest risk, as opposed to

simply trying to monitor, collect, and defend everything. The method can also be a valuable tool for executives faced with decisions about investments and balancing operational and security requirements.

Disclaimer

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 United States Government.

About the Authors

Robert F. Mills is an Associate Professor of Electrical Engineering at the Air Force Institute of Technology. He received his PhD in Electrical Engineering from the University of Kansas in 1994. His research interests are in communication systems, network management and security, information warfare, and systems engineering. He is a member of Eta Kappa Nu and Tau Beta Pi, and is a Senior Member of the IEEE. He may be contacted at Robert.Mills@AFIT.edu.



Michael R. Grimaila, CISM, CISSP, is an Associate Professor of Information Resource Management at the Air Force Institute of Technology. He received his PhD in Computer Engineering from Texas A&M University in 1999. Dr. Grimaila's research interests include cyber incident detection, mission assurance, network management and security, information warfare, and systems engineering. He is a member of the ACM, Eta Kappa Nu, ISACA, ISC2, ISSA, Tau Beta Pi, and he is a Senior Member of the IEEE. He may be contacted at Michael.Grimaila@AFIT.edu.



Gilbert L. Peterson is an Associate Professor of Computer Science at the Air Force Institute of Technology. He received his PhD in Computer Science from the University of Texas Arlington in 2001 and a BS in Architecture from the University of Texas Arlington in 1995. His research interests include digital forensics, insider threat mitigation, and artificial intelligence. He may be contacted at Gilbert.Peterson@AFIT.edu.



Dr. Jonathan Butts is an Assistant Professor of Computer Science at the Air Force Institute of Technology. Jonathan is an active duty captain in the Air Force. His research interests include critical infrastructure protection, insider threat mitigation, and national cyberspace policy. He may be contacted at Jonathan.Butts@AFIT.edu.

