Presented at the University of Louisville Cyber Securitys Day, October 2006

Solving the Insider Threat Problem

Dr. Bruce Gabrielson, NCE
Booz, Allen, Hamilton
CND R&T PMO

Abstract

The insider threat is significant and real within both the Department of Defense (DoD) and commercial sector. Many previous studies have looked at the broad scope of the problem without any real attempt to identify a solution. This presentation provides visibility into how this or any other broadly defined technology gap can be decomposed such that partial solutions are identifiable using a formal investigative process. In particular, it describes the approach taken by the DoD's Enterprise-wide Information Assurance (IA)/Computer Network Defense (CND) Solutions Steering Group's (ESSG's) Insider Threat Technology Advisory Group (TAG).

Insider Threat Status and Issues

As a general DoD definition, the "insider" is anyone who is or has been authorized access to a DoD information system, whether a military member, a civilian employee, employee of another Federal agency or the private sector. Some definitions, however, address the broader scope of "system components" or "computer software code" inserted inside a system and intended to carry out a malicious act. Of interest regarding the many broad descriptions of insider is that the definition proposed is often dependent on the perspective of the individual defining the problem. The real question arises, is the perpetrator simply someone exhibiting bad behavior or is this person representing a serious threat to our nation.

Regardless of the definition used, we do know the insider threat is significant and real. A recent DoD Inspector General (IG) report indicates that, for one set of investigations, 87 percent of identified intruders into DoD information systems were either employees or others internal to the organization.

Insider Threat Details

The definition of insider threat should encompass two main threat actor categories and five general categories of activities. The first actor category, the "true insider," is defined as any entity (person, system, or code) authorized by command and control elements to access network, system, or data. The second actor category, the "pseudo-insider," is someone who, by policy, is not authorized the accesses, roles, and/or permissions they currently have but may have gotten them inadvertently or through malicious activities.

The activities of both fall into five general categories: 1) exceeds given network, system or data permissions; 2) conducts malicious activity against or across the network, system or data; 3) provided unapproved access to the network, system or data; 4) circumvents security controls or exploits security weaknesses to exceed authorized permitted activity or disguise identify; or 5) non-maliciously or unintentionally damages resources (network, system or data) by destruction, corruption, denial of access, or disclosure.

Some investigators have cited four categories of the insider problem: traitor, zealot, browser, and well intentioned. The traitor category includes persons who have a malevolent intent to damage, destroy, or sell out their organization. The zealot category involves an insider who believes strongly in the correctness of one position or feels the organization is not on the right side of a certain issue. The browser category consists of persons who are overly curious in nature (often a violation of the need-to-know principle), while the well-intentioned insider commits violations through ignorance. Downloading shareware, disabling virus protection software, using unapproved CDs can all provide the assistance a hacker needs to penetrate a system. The well-intended user can become the unwitting and unknowing associate.

Because insider threat is a heterogeneous problem with many component parts, the solution becomes too complex a problem for anyone to expect a “silver bullet” type solution to handle it all. Managing a architecture consisting of a set of point solutions with multiple data gathering needs and potentially distributed stakeholders, each with their own data sharing or further investigative requirements, impacts the overall solution set architecture needed. An insider threat conceptual architecture should leverage an array of network and host-based sensors along with existing networked systems that provide network analysis or access controls. Figure 1 depicts the overall architecture that can support this integrated insider threat solution approach.

Note in Figure 1 that law enforcement and counter-intelligence (LE/CI) are special purpose legally authorized organizations that have formal investigation authority. Their potential solution set can include the use of specialized monitoring sensors to provide data to investigators. However, while these organizations can collect data from multiple sources, based on legal access requirements, no data can be further disseminated from the collecting authority to outside organizations.

Another issue that should be pointed out in Figure 1 is that any practical response approach must focus on a selection of component parts, each developed based on the unique needs of their user community. A related issue with point solutions is that not all approaches to identify and mitigate an insider are unique to insider threat mitigation. Mitigation techniques that have to be implemented for external threats often have overlapping capabilities to mitigate against both internal and external threats. Therefore, the tools commonly used by system administrators, network analysts, and/or criminal investigators can be integrated into a comprehensive insider threat mitigation toolset. Point solutions are most useful when they focus primarily on technology gaps where no solution from any other source exists.

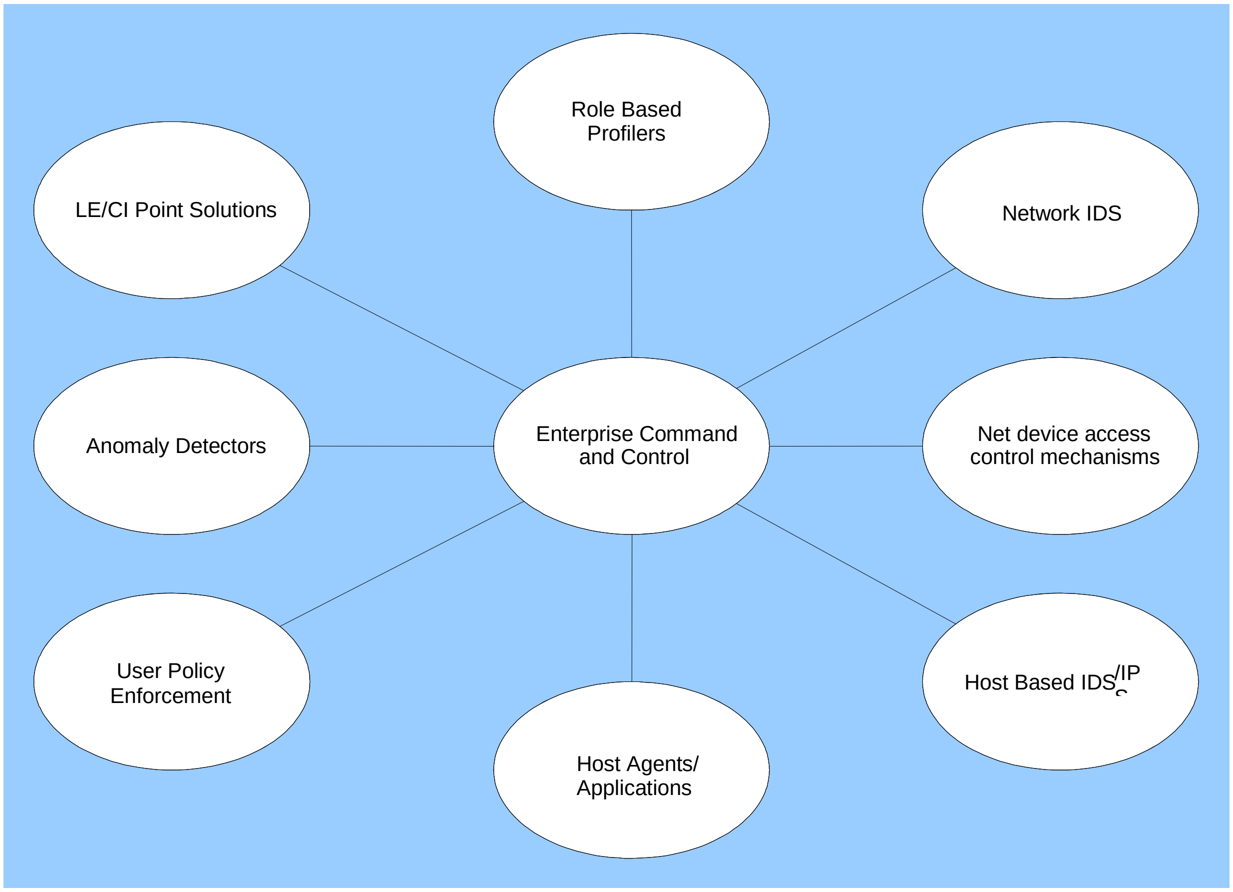

Figure 1 - Insider Threat Notional Architecture

Solution History

Organizations have been trying to solve the computer based insider threat problem for several years, most recently though network based means. It wasn't until a few years ago that the DoD formally attempted to actually identify solutions. However, because of the nature of the problem, the solution process has moved slowly, involving extensive learning and collaboration. Some significant milestones in solving the insider threat problem are listed below.

- Various workshops and working groups developed an initial set of requirements – through 2004.
- Technology Advisory Group is formed to address the problem – April 2004.
- USSTRATCOM creates an initial DoD vetted set of insider threat requirements, the Insider Threat Required Capabilities Document (RCD), September 2004.
- Government off the Shelf (GOTS) and Commercial off the Shelf (COTS) Days address current solutions and research activities –2004/2005.
- US Strategic Command Insider Threat Notional CONOPS developed – May 2005.
- Revised set of insider threat technical requirements developed in the fall/winter of 2005.

- Formal development of four part insider thread solution set definitions – March 2006
- Fully DoD vetted set of "testable" requirements completed – June 2006

The Insider Threat TAG currently consists of subject matter expert representatives from at least twenty-two services and agencies. This group has been charged with generating formal technical requirements and identifying solutions that will provide a baseline insider threat mitigation capability on the DoD enterprise. The group has reviewed current research activities, developed a vetted requirements specification, published Insider Threat RFI's, and reviewed existing COTS and GOTS solutions that address the insider threat problem space.

Solving a Difficult Problem

Government organizations face several challenges in stimulating research within identified capability gap areas. Many of these gaps represent protection issues that are simply too broad and contain too many sub-problems to be readily categorized into a focused research area and then resolved by one solution or a set of partial solutions. Further, specific solutions to most CND gaps are driven both by mission requirements and the ability of a solution to meet operational, functional, and information assurance concerns. For the insider threat problem, all these driving factors had not been well defined previously. To solve this problem, four steps were undertaken, led primarily by the NSA CND Research and Technology (R&T) Program Management Office (PMO) and the Insider Threat TAG:

1. Technology Decomposition:
    - Decomposes the needed CND technology into the basic functional components
2. Solution Mapping:
    - Mapping the various point solutions against this decomposition
3. Solution Evaluation:
    - Evaluate potential solutions and overlapping capabilities and then recommend those that offer the "greatest bang for the buck" and/or address the most pressing operational needs
4. Focus Research:
    - Focus research thrusts on to gap areas that are not fully addressed by those existing solutions evaluated
    - Research would also address the need for technology transition

For insider threat, the solution/operational need matching process has now been completed and an enterprise-wide baseline solution to the insider threat problem will soon be acquired.

Decomposing a Technology

Several organizations and considerable coordination is necessary when decomposing a broad technology research area. This step involves clearly defining the problem space based on determining every functional need or required capability. In this case, the problem set was an identified and sometimes ill-defined technology area known as insider threat. Initially, a set of required capabilities was developed based on discussions among subject matter experts from many different organizations. Additionally, a combination of commercial product descriptions, operational requests, and functional test results of related mature products was used to enhance the capabilities identified.

Unfortunately, every operational organization, vendor, researcher, or industry "expert" had a slightly different concept of what the technology means. Therefore, the first consolidation addressed the entire landscape of capabilities, and further iterations were necessary as the problem set became better understood and defined. The goal is to express each functional need at the lowest level so testable requirements can be described in a straightforward manner and the comparison of a solution against each requirement is possible.

Defining the Problem Space in Terms of Threats and Mission Needs

The previous definition of an insider threat identified five threat actors. Using the term "bad behavior" to cover the actions of both browsers and the well-intentioned, the inter-relationships between those who detect threat actor activities and those responsible for the investigation and mitigation of these activities can be mapped for various threat levels as shown in Table 1. The term "Computer Emergency Response Team" (CERT)[1] is used here to indicate any organization with network level intrusion detection monitoring capabilities and computer

Threat Actor	Threat Level	Initial Detection	Initial Detection Level	Solution Set Deployment	Report to Responsible Authority
Bad Behavior	Minimal	SA and/or security administrator	Enterprise and/or Enclave Network Sensors	Significant Deployment	Local Service/Agency
Outsider - Pseudo-Insider	Severe	Analyst	Enterprise and/or Enclave Network Sensors	Moderate Deployment	Enterprise and/or Enclave/Service CERT
Active Malicious Insider	Most Severe	Analyst	Enterprise and/or Enclave Network Sensors	Limited Deployment	Misuse Detection Investigator and/or LE/CI
Passive Malicious Insider	Most Severe	Intelligence Investigator and/or CI Analyst	Host Sensors	Limited Deployment	Intelligence Investigator and/or LE/CI

Table 1 - Integrated Detection and Investigative Capabilities

[1] Some CERTs now use the term CND Service Provider (CNDSP) as defined in DoDD O-8530.1 and DoDI O-8530.2 At the DoD level the former DoD CERT is now known as "NetDefense." In the Services, they are now known as NOSCs (AFNOSC, MCNOSC), TOCs (A2TOC), or CIRTs (NAVCIRT).

Solution Mapping

The combination of the functional capability breakdown at the basic need level and the set of specific operational requirements that have been generated enables a mapping of each available commercial solution or research activity against specific operational requirements. This process creates a clearer picture what is available and might soon be available. It also helps identify those capabilities needed that aren't being met, the "gaps" in available technologies.

For the insider threat case, solutions generally fall into two types, those that look for unauthorized activity (which includes improper behavior) and those that look for anomalous behavior that may indicate malicious activity. Either solution type will alert when the behavior of interest is identified. Additionally, some solutions concentrate on network-based activities while others concentrate on host-based activities. While most existing insider threat solutions are network based, actual case studies have shown that the most critical need, and the most common insider threat problem, relates to host-based rather then network-based monitoring and analysis.

Figure 2 depicts how mapping was used to superimpose solution functionality as an aid in identifying a possible solution set as well as highlighting where focus area technology gaps exist. Some interesting conditions begin to emerge once mapping is accomplished. Once the entire scope of the technology is understood and mapped, vendors using this approach can quickly ascertain their product's limitations and where improvements can be made. Researchers start to have a much clearer feel for where technology gaps exist. Users can identify limitations in their installed solution set. Planners start to see their "way forward" for budgeting and focusing research activities.

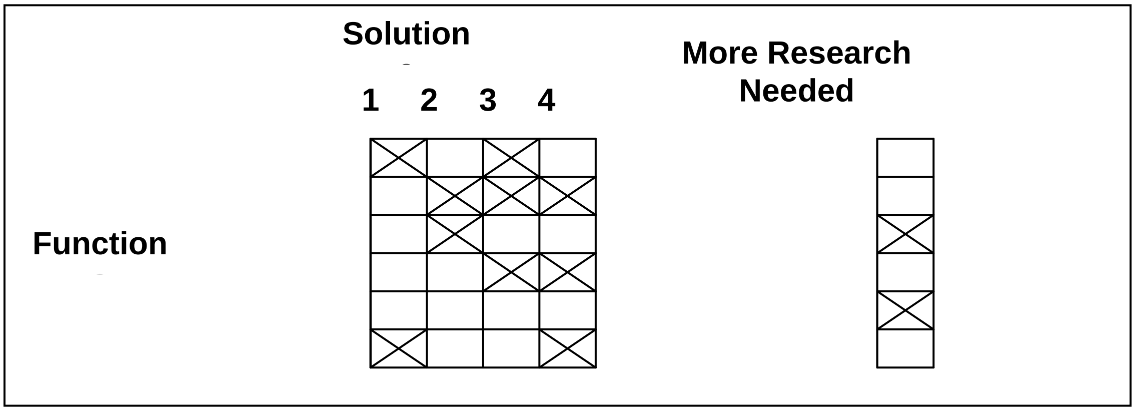

Figure 2 - Solution Sets and Gaps Highlighted by Mapping

Scope of the Problem Space

Based on capability mapping and current GOTS or COTS capabilities available, the insider threat solution is presently envisioned as encompassing four parts, each necessary to provide a complete solution to the problem. Parts 1 through 3 combined provide the baseline capability for initial misuse detection.

- Part 1 - Host-Based Anomaly Detector

    - Host-based insider threat sensors target per user activity. Using an installed agent, the sensor focuses on suspicious activity by authorized users performing actions. These actions, when correlated, may be determined to be bad behavior, accidental actions, or potentially an insider with malicious intent.

- Part 2 - Network-Based Anomaly Detector
    - Network-based insider threat tools tie suspicious, but normally non-actively malicious, behaviors to specific users. These tools, which include network behavior modeling tools, are also useful in detecting outsiders who have successfully penetrated the network and are acting as insiders.
- Part 3 - Correlator – Correlates data from the previous two insider threat parts (as well as log files and other feeds)
    - Correlators are specifically designed to enable the identification of authorized user behavior consistent with profiled insider threat activities.
- Part 4 – Network and Host-Based Focused Observation Tool
    - This is placed on an end-user's box, or on the network, to gather more data about a specific user, particularly the content of their activities and transmissions. While some of the tools in this space can be deployed to monitor actions without end-user knowledge, this may not be considered a critical feature.

Figure 3 below helps explain their relationship.

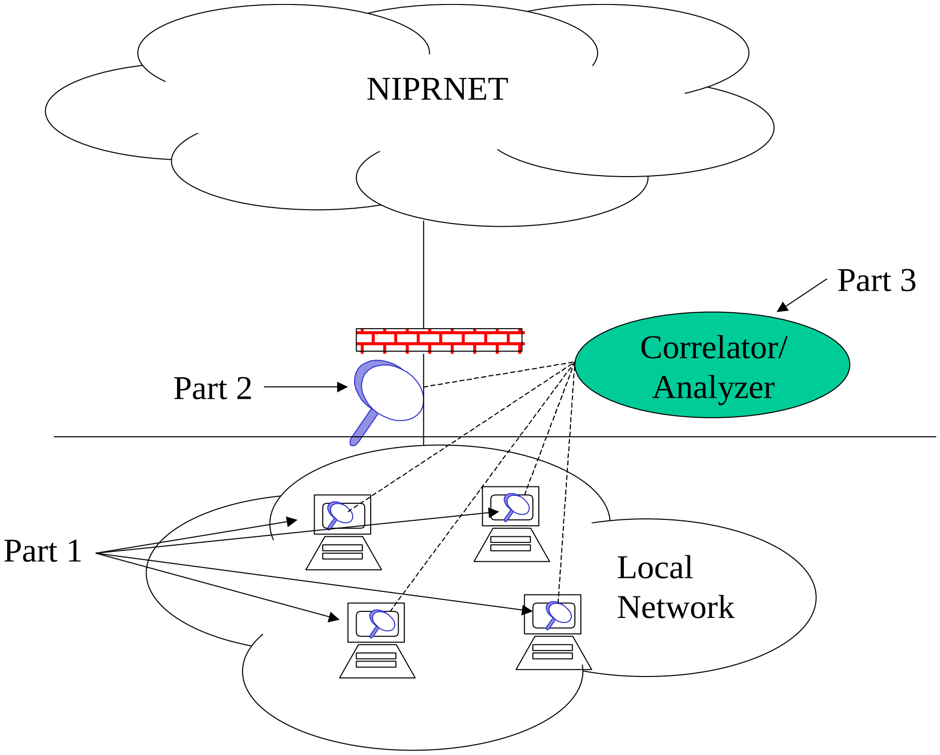

Figure 3 - Insider Threat Solution Architecture

Solution Selection Criteria and Issues

Having accomplished a significant understanding of the technology needs and solutions available using mappings, the actual selection and evaluation of a solution is still difficult. The most damaging insider threat to deal with relates to the individual on the inside who captures and then exfiltrates information in some manner while avoiding intrusion detection alarm conditions. Unfortunately, these malicious individuals attempt to mask their activities by operating within normal or abnormal but acceptable behavior. Since behavior can vary widely among individual users, the initial internal detection mechanism must be both lightweight and generate relatively few false alarms. This type of activity cannot be addressed by boundary solutions.

Existing behavior monitoring tools are computationally intensive and are not easily scalable. They usually require training and a learning period to avoid false alarms. It is also a key concern that an imposter could access the hidden profile used by the detection system and exploit its vulnerabilities before the system could validate the user. If normal activity on the system is regularly delayed until proper identification can be achieved, this delay could become a significant nuisance to the user. Additionally, this could result in a series of user actions that actively attempt to circumvent these controls since they add no user perceived benefit. These actions will increase the false positive rate of the system.

Another problem is the need for evidence preservation for some solution users. There are significant issues of evidence preservation, prosecution and damage assessment. These issues can be helped or hindered by other technical means, but the actual identification and mitigation of a true insider threat must include attention to the prosecution and administrative actions that are fundamental to deterrence of the malicious activity. Without legally admissible evidence neither of these can happen.

Solution Recommendation

Once potential candidate solutions are identified, the order of solution acquisition depends on both budgetary consideration and the organization’s ability or difficulty to integrate the solution into their existing infrastructure and operational model. Considering these problems, the last thing users want is to field a product that's complex and costly to use, doesn’t address the projected users’ unique needs, and doesn't catch the highest priority insider behavior. However, deploying a tool that simply produces an abundance of alerts on anomalous behavior for system administrators and doesn’t support evidence gathering would be nearly as useless. Many, and often the most dangerous, insiders are sophisticated and their activities are likely difficult detect and prove.

Since the threat is recognized and significant, it is also imperative that if the proper solution is identified, the acquisition process move forward rapidly with an aggressive series of pilots. These pilots would include point solutions to test and support solutions to the existing categories of insider threat. It is important to configure and deploy a combination of host-based activity monitors and intrusion detection system daemons or agents that are lightweight and have a low probability of detection. This initial alert capability will initiate the deployment of more robust tools that can be used by investigators in the actual identification, attribution, evidence preservation and prosecution gathering stages of an investigation. The initial solution should be compatible with, and not hamper the deployment of additional tools or agents, particularly if and when other alert mechanisms are identified. These products would also address those tools or agents that provide for the collection of evidence that can be used by the law enforcement and counterintelligence communities.

Focusing Research

As previously mentioned, a complete insider threat solution set does not exist for all needed capabilities. Research is necessary to fill the technology "gap" areas. Some initial technology gaps can be solved with existing tools and integrated quickly once a baseline solution is implemented, while other gaps will be solved near-term or are considered "Grand Canyons" that may take focused research over a much longer term to solve. However, simply having a good idea that results in a gap solution isn't enough. Research that develops an initial proof of concept or an emerging partial solution is not the same as having a transitioned product ready for deployment.

Historically, researchers have an idea, canvas the community as to what already exists and what problem set they should specifically go after, and then seek to develop their attempted solution. Unfortunately, while a customer needing the solution may exist, unless the researcher can locate a customer or venture capital source for commercialization, the research often goes no further than perhaps a proof of concept phase.

To help mitigate this problem, part of the research and development cycle should be focused on first creating a business case for the research and projected solution, and then mitigating any risks during the transition to operation. The rapid cycle of CND requirement identification, discovery of research and technology shortfalls, development of research and technology solutions, and the integration and acquisition of the solutions into deployable systems is an iterative process of mitigating risk at every step in the program life-cycle.

Currently Identified Gaps

The following table summarizes needs, gaps, and areas for exploration:

Table 2 - Insider Threat Research Gaps

Need	Gaps	Areas for Exploration
Insider Characterization and Modeling	Typology / taxonomy of insiders	Typology with respect to DoD and IC and significant assets Human characteristics, both individual and group; psychological profiling; examination of motivations and intentions
	Models of insider adversary behavior	Informal modeling Statistical modeling
	Validation of insider adversary behaviors and models	Empirical studies Experiments Simulations
Preventative countermeasures against the insider	Accountability for insider actions, particularly in heterogeneous	Multiple and coordinated forms of authentication across security domains or organizations

Need	Gaps	Areas for Exploration
	environments	
		Watermarking, fingerprinting, and other forms of marking data to provide a deterrent to or a detection of unauthorized actions (disclosure, modification)
	Access control mechanisms sensitive to insider threats	Differential access controls depending on roles, rights, privileges, access context, and history
Monitoring and detection of adversarial insider behavior	Effective modeling / profiling of adversarial insiders	Social network analysis
	Monitoring techniques for different classes of insiders	Monitoring and analysis of system administrators Application-based monitoring and analysis Correlation across multiple monitoring mechanisms Differential and adaptive monitoring
Reactive countermeasures for the insider adversary	Analysis capabilities	Tools for analyzing and correlating monitoring data and audit records Forensic tools on machines and storage devices Evidence collection and preservation
	Automated response capabilities	Dynamic determination of the need for, and implementation of, restricting access, initiating additional data collection or monitoring, compartmentalizing the organization’s network

Achieving Future Needs

With a comprehensive approach available, critical CND needs are being addressed. This approach must understand our operational challenges, be able to develop functional requirements, support the evaluation of solutions and emerging technologies against these requirements, and finally ensure that focused solutions are made available in a timely manner throughout the DoD. When this approach is coupled with focused approaches for ensuring research activities address emerging needs, the prospects for meeting the challenges of CND now, and the future, are promising.