

AFRL-SN-WP-TR-2002-1101

**INSIDER ANOMALY MEASUREMENT
PROCESSING SYSTEM (IAMPS)**

Dennis H. McCallam

**Northrop Grumman Information Technology
Defense Enterprise Systems
1813 Weihle Avenue
Reston, VA 20190**



MARCH 2002

Final Report for 19 December 2000 – 31 March 2002

Approved for public release; distribution is unlimited.

20020830 084

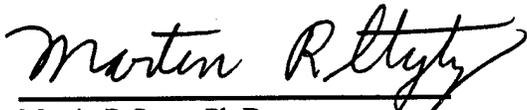
**SENSORS DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7318**

NOTICE

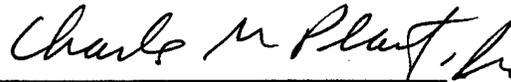
USING GOVERNMENT DRAWINGS, SPECIFICATIONS, OR OTHER DATA INCLUDED IN THIS DOCUMENT FOR ANY PURPOSE OTHER THAN GOVERNMENT PROCUREMENT DOES NOT IN ANY WAY OBLIGATE THE US GOVERNMENT. THE FACT THAT THE GOVERNMENT FORMULATED OR SUPPLIED THE DRAWINGS, SPECIFICATIONS, OR OTHER DATA DOES NOT LICENSE THE HOLDER OR ANY OTHER PERSON OR CORPORATION; OR CONVEY ANY RIGHTS OR PERMISSION TO MANUFACTURE, USE, OR SELL ANY PATENTED INVENTION THAT MAY RELATE TO THEM.

THIS REPORT HAS BEEN REVIEWED BY THE OFFICE OF PUBLIC AFFAIRS (ASC/PA) AND IS RELEASABLE TO THE NATIONAL TECHNICAL INFORMATION SERVICE (NTIS). AT NTIS, IT WILL BE AVAILABLE TO THE GENERAL PUBLIC, INCLUDING FOREIGN NATIONS.

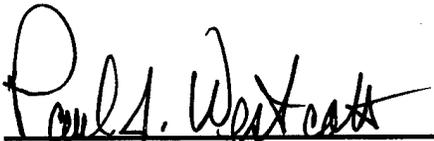
THIS TECHNICAL REPORT HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION.



Martin R Stytz, Ph.D.
Project Engineer
Electronic Warfare Branch
Sensor Applications & Demonstrations Division



Charles M. Plant, Jr.
Branch Chief
Electronic Warfare Branch
Sensor Applications & Demonstrations Division



Paul J. Westcott
Division Chief
Sensor Applications & Demonstrations Division

Do not return copies of this report unless contractual obligations or notice on a specific document require its return.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

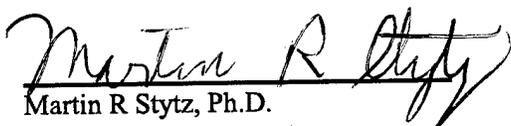
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 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 any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YY) March 2002		2. REPORT TYPE Final		3. DATES COVERED (From - To) 12/19/2000 – 03/31/2002	
4. TITLE AND SUBTITLE INSIDER ANOMALY MEASUREMENT PROCESSING SYSTEM (IAMPS)				5a. CONTRACT NUMBER F33615-01-C-1806	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 62301E	
6. AUTHOR(S) Dennis H. McCallam				5d. PROJECT NUMBER ARPS	
				5e. TASK NUMBER NZ	
				5f. WORK UNIT NUMBER 0H	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Northrop Grumman Information Technology Defense Enterprise Systems 1813 Weihle Avenue Reston, VA 20190				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Sensors Directorate Air Force Research Laboratory Air Force Materiel Command Wright-Patterson Air Force Base, OH 45433-7318				10. SPONSORING/MONITORING AGENCY ACRONYM(S) AFRL/SNZW	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S) AFRL-SN-WP-TR-2002-1101	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The solution detailed in this report is based on the fusion of information from a variety of cyber sensors, all looking for authentication consistency. In the event that authentication inconsistency is developed, the user holding the presented credentials is denied further access to the system. Several forms of authentication information and types of sensors were considered as part of the IAMPS suite, with the goal of using common COTS sensors to enhance transfer of IAMPS technology into real world systems. An additional sensor was conceptually developed to profile users based on the more hardware-related parameters all specifying computer usage. The basis of the IAMPS solution is in the application of sensor fusion approaches. While initially seeking to use only one form of fusion (i.e., Bayesian Networks, Dempster-Schaeffer, etc.), it was decided that a hybrid approach would work best. This avoids the problem of methods targeted to defeat certain forms of fusion if the detection fusion approach is known. A hybrid approach preserves algorithmic integrity. The research developed success criteria for evaluation of alternatives and then applied those criteria to the IAMPS solution. In summary, IAMPS directly addresses one of Sherlock Holmes' major concerns, as stated in the novel <i>A Study in Scarlet</i> : "There is nothing like first hand evidence." IAMPS leverages all sources of authentication information to develop aspects of first hand evidence.					
15. SUBJECT TERMS user profiling, access control, intrusion detection, learning systems, intrusion detection systems, user identification, authentication consistency					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT: SAR	18. NUMBER OF PAGES 60	19a. NAME OF RESPONSIBLE PERSON (Monitor) Martin R. Stytz, Ph.D. 19b. TELEPHONE NUMBER (Include Area Code) (937) 255-2811 x4380
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			

NOTICE

USING GOVERNMENT DRAWINGS, SPECIFICATIONS, OR OTHER DATA INCLUDED IN THIS DOCUMENT FOR ANY PURPOSE OTHER THAN GOVERNMENT PROCUREMENT DOES NOT IN ANY WAY OBLIGATE THE US GOVERNMENT. THE FACT THAT THE GOVERNMENT FORMULATED OR SUPPLIED THE DRAWINGS, SPECIFICATIONS, OR OTHER DATA DOES NOT LICENSE THE HOLDER OR ANY OTHER PERSON OR CORPORATION; OR CONVEY ANY RIGHTS OR PERMISSION TO MANUFACTURE, USE, OR SELL ANY PATENTED INVENTION THAT MAY RELATE TO THEM.

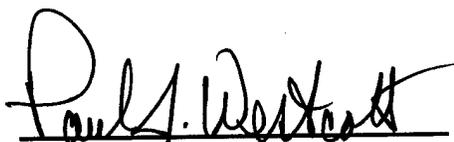
THIS TECHNICAL REPORT HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION.



Martin R Stytz, Ph.D.
Project Engineer
Electronic Warfare Branch
Sensor Applications & Demonstrations Division



Charles M. Plant, Jr.
Branch Chief
Electronic Warfare Branch
Sensor Applications & Demonstrations Division



Paul J. Westcott
Division Chief
Sensor Applications & Demonstrations Division

Do not return copies of this report unless contractual obligations or notice on a specific document require its return.

TABLE OF CONTENTS

SECTION NUMBER AND TITLE	PAGE
1.0 INTRODUCTION AND OVERVIEW	1
2.0 BACKGROUND	5
3.0 METHODS, ASSUMPTIONS, AND PROCEDURES	11
4.0 ALTERNATIVE SOLUTIONS AND EVALUATION	13
5.0 PROPOSED SOLUTION	14
6.0 RESULTS AND DISCUSSION	37
7.0 CONCLUSIONS AND SUGGESTIONS FOR FUTURE WORK	41
REFERENCES	44
GLOSSARY	46

FIGURE NUMBER	LIST OF FIGURES FIGURE TITLE	PAGE
1	IAMPS Process Showing Fusion of Authentication Information	4
2	Sample Physical Layout	19
3	Initial Concept of IAMPS Solution	21
4	Example of a Timecard System	22
5	Typical Premises Entry Devices	23
6	Typical Logon Authentication Devices	25
7	Computer Environment Analysis – Nominal Case Example	28
8	Computer Environment Analysis – Bad Case Example	28

LIST OF TABLES

Table Number and Title	Page
1 Literature Search Parameters and Results	6
2 Developed Evaluation Criteria	13
3 Relative Measures of Half Life Concept	20
4 Authentication Sensor Parametric Summary and Fusion Impacts	32
5 Sample Testing Scenario Generation	35
6 Summary IAMPS Solution and Evaluation	38

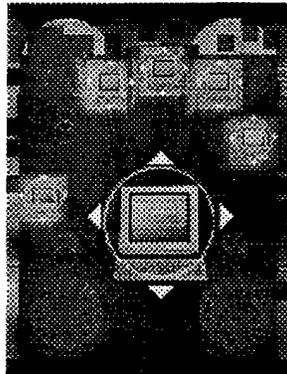
ABSTRACT / EXECUTIVE SUMMARY

The research detailed in this final report is the result of a DARPA sponsored program to “develop active (real time) techniques for profiling users and applications that will discriminate between normal and anomalous behavior for a given user, be able to discriminate among users, and identify new (unobserved) insider-initiated misuse.” There is no question that research into detection of insider-initiated misuse is of paramount importance considering that the insider still presents the greatest single threat to computer security in particular. Given this, the research set out to develop a process denoted as the Insider Anomaly Measurement Processing System, or IAMPS. The foundation of the IAMPS research considered the insider misuse to be from user(s) who are masquerading and are not who they claim to be, even though they have passed some level of authentication in the system. Furthermore by looking at several pieces of data (evidence), that inconsistency can be empirically supported. The belief was that users who attempted to gain access under these false pretenses are intentionally malicious.

The solution detailed in this report is based on the fusion of information from a variety of cyber sensors, all looking for authentication consistency. In the event that authentication inconsistency is developed, the user holding the presented credentials is denied further access to the system. Several forms of authentication information and types of sensors were considered as part of the IAMPS suite, with the goal of using common COTS sensors to enhance transfer of IAMPS technology into real world systems. An additional sensor was conceptually developed to profile users based on the more hardware-related parameters all specifying computer usage. The basis of the IAMPS solution is in the application of sensor fusion approaches. While initially seeking to use only one form of fusion (i.e., Bayesian Networks, Dempster-Schaeffer, etc.) it was decided that a hybrid approach would work best. This avoids the problem of methods targeted to defeat certain forms of fusion if the detection fusion approach is known. A hybrid approach preserves algorithmic integrity. The research developed success criteria for evaluation of alternatives and then applied those criteria to the IAMPS solution. The salient features of IAMPS are:

- A real time detection schema,
- Can operate with minimal or no operator interference,

- Can maintain forensic evidence,
- Is open to a wide range of cyber sensors, and
- Identifies any masquerading user(s) in spite of other system authentication.



The IAMPS Concept illustrating fusion of different forms of authentication credentials

IAMPS has some clear applications to current needs for both data and system security and terrorist identification. In fact several concepts for further work include:

- The Implementation and Testing of IAMPS in an alpha test environment to employ rapid, focused experiments to determine/refine the necessary sensors, sensor combinations, and intent models while simultaneously working to improve the technical quality and breadth of each of these component technologies. The data gathering for the cyber sensors and verification of the declaration process will occur in a real environment, such as one of the buildings at the Northrop Grumman Information Technology (NGIT) Sector. An extension to either variant would be to implement the IAMPS solution across a distributed network, and then possibly a network of networks to evaluate the impacts in real time, long distance authentication.
- Terrorist identification, an application of IAMPS technology with facial and image recognition. Events of September 11, 2001 caused a parallel line of applications thinking in how to apply non-repudiated authentication to the problem of terrorist identification. The result is the description of a program that using facial and other source image recognition to identify known terrorists. A variant of this would be to compare the images in a source issuer's database to look for multiple images against the same name or multiple names against the same images.
- Comparative testing for non-repudiated authentication. – This would develop a standardized testing set for evaluation of insider security solutions
- Integration of policy and IAMPS tools with measured testing. – Some of the types of security monitors look only for violations in the security policy. A broadening of that base would include the authentication portions of IAMPS with the policy verification portions of other commercially available tools. The standard test case would then be used to measure the integrated set

In summary IAMPS directly addresses one of Sherlock Holmes' major concerns, as stated in the novel *A Study in Scarlet*, "There is nothing like first hand evidence". IAMPS leverages all sources of authentication information to develop aspects of first hand evidence.

1. INTRODUCTION AND OVERVIEW

As a result of the rapid growth in computer technology, the government and private sector, has become extremely dependent on automated information systems. Access to these systems may be wanted by individuals or organizations seeking monetary gain, political blackmail or those dedicated to causing damage. As a consequence, access to systems is restricted and controlled to those who are trusted and approved. But what happens when those trusted and approved users begin to use the system for illicit, illegal, or seditious purposes? The number one threat for information warfare attacks for any system, military or otherwise, is the insider attack. The premise has been that adequate perimeter defenses can keep unauthorized users from entering the system through the IP connection and that means no possible intrusion. Current intrusion detection schemes rely on software checking software for abnormal behavior. This technology has had limited success in the IP based systems environment, but offers no additional protection for non-IP based command and control systems. All this falls short considering the insider attacks. Insider attacks transcend and redefine access control and intruder detection. Addressing the issue of the insider attack, has been more of a post-mortem analysis conducted after the damage has been done. Once a hacker gains access to a system (particularly root access), they have now become a legitimate user and are able to bypass all the conventional intrusion detection concepts. To some degree, this implies that all attacks in some form are really insider attacks. The issue centers on finding ways to decide, **in real time**, if there is an insider attack in progress.

The IAMPS program researched a methodology into detecting, in real time, insider misuse. This program considered the insider misuse to be from user(s) who are masquerading and are not who they claim to be, even though they have passed some level of authentication in the system. Furthermore by looking at several pieces of data (evidence), that inconsistency can be developed. The belief was that users who attempted to gain access under these false pretenses are intentionally malicious.

Succinctly stated, the goal of this research was to gain *non-repudiated authentication*, and to be able to prove that:

- A person is who they claim to be and we can support that claim empirically through electronic evidence, or,
- A person is not who they claim to be and we can discover the discrepancies using electronic evidence.

This research performed a detailed investigation concept study of an **insider attack detection** scheme using fusion concepts that integrates available corroborating information. The novelty and creativeness of this approach is fourfold:

INSIDER ANOMALY MEASUREMENT PROCESSING SYSTEM (IAMPS)

Author: Dennis H. McCallam

*Northrop Grumman Information Technology
Defense Enterprise Systems
1813 Weihle Avenue
Reston, Va. 20190*

ABSTRACT / EXECUTIVE SUMMARY

The research detailed in this final report is the result of a DARPA sponsored program to “develop active (real time) techniques for profiling users and applications that will discriminate between normal and anomalous behavior for a given user, be able to discriminate among users, and identify new (unobserved) insider-initiated misuse.” There is no question that research into detection of insider-initiated misuse is of paramount importance considering that the insider still presents the greatest single threat to computer security in particular. Given this, the research set out to develop a process denoted as the Insider Anomaly Measurement Processing System, or IAMPS. The foundation of the IAMPS research considered the insider misuse to be from user(s) who are masquerading and are not who they claim to be, even though they have passed some level of authentication in the system. Furthermore by looking at several pieces of data (evidence), that inconsistency can be empirically supported. The belief was that users who attempted to gain access under these false pretenses are intentionally malicious.

The solution detailed in this report is based on the fusion of information from a variety of cyber sensors, all looking for authentication consistency. In the event that authentication inconsistency is developed, the user holding the presented credentials is denied further access to the system. Several forms of authentication information and types of sensors were considered as part of the IAMPS suite, with the goal of using common COTS sensors to enhance transfer of IAMPS technology into real world systems. An additional sensor was conceptually developed to profile users based on the more hardware-related parameters all specifying computer usage. The basis of the IAMPS solution is in the application of sensor fusion approaches. While initially seeking to use only one form of fusion (i.e., Bayesian Networks, Dempster-Schaeffer, etc.) it was decided that a hybrid approach would work best. This avoids the problem of methods targeted to defeat certain forms of fusion if the detection fusion approach is known. A hybrid approach preserves algorithmic integrity. The research developed success criteria for evaluation of alternatives and then applied those criteria to the IAMPS solution. The salient features of IAMPS are:

- A real time detection schema,
- Can operate with minimal or no operator interference,

- First, the overall premise of this approach is that there is a measurable steady state of system performance parameters that a given system operates in terms of memory reads/writes, disk accesses, power used, internal ambient temperature, etc. While any one of these could fluctuate, it is our technical judgment that several of these being out of “steady state” ranges could indicate unauthorized access or use of the computer system.
- Second, we believe that there are other available measurements that exist that when viewed in their entirety, can potentially declare an insider attack and can perform that analysis in real time. The steady state analysis results can be fused *in real-time* with timecard, premises, and authentication information, also available in real time, to augment the declaration process and offer fidelity and certainty into the accuracy of the prediction
- Third, we postulate that the steady state measures of system use offer a new and creative method of profiling an attacker based on those measurements. This implies that an additional and empirical measure can be taken and used as “cyber evidence”.
- And finally, this technique detects the computer misuse in real-time without increasing the system overhead or workload.

“Circumstantial evidence is a very tricky thing. It may seem to point very straight to one thing, but if you shift your own point of view a little, you may find it pointing in an equally uncompromising manner to something entirely different.¹” Such are the issues raised as a result of researching the insider problem. First is to decide what is evidence and what is not. In fact, the approach taken by IAMPS is to consider any piece of information related to authentication as potential evidence. Second is to look at how to gather the evidence. IAMPS addresses this issue by using COTS sensors that already make available electronically the authentication-credential results. Third is a consideration of the privacy issues raised by using certain sensors that may develop profiles of users. IAMPS considers this in the application of the sensors to avoid the negative connotations of profiling. Finally there are the issues of solution impact on overall security in terms of people, process and technology. The IAMPS research was especially sensitive to this point. The IAMPS solution maintains a good balance between the people affected and practicing it, the security processes those people develop and adhere to, and the technologies selected and implemented to enforce and facilitate the security. Achieving that balance is difficult because changes in technology ripple through and affect previously good processes which in turn can affect the operability of the people.

Finding a solution to this problem is of extreme strategic and tactical importance to security personnel everywhere. Universally, by any measure, the largest single threat to security is the insider². Finding ways of detecting insider activity is a prime concern across the industry. But most of the initiatives have been focussed on misuse of computer resources in terms of policy violations, in short electronically breaking the rules. This falls short of the true definition of an insider as one who intentionally intends to compromise the computer resources or the data contained within those resources.

¹ From the website of Sherlock Holmes quotes, www.bakerstreet221b.de/canon/index.html, and the novel *The Boscombe Valley Mystery*

² One example is from the FBI year end reports on top security threats. In both 1999 and 2000, the insider was cited as the greatest threat to system security

In fact, there are several good COTS products on the market for security monitoring. Most promising are the one in the Intellitactics and e-Security genre. These solutions offered a great deal of promise in monitoring and managing a wide range of sensors and correlating the sensors to events. But the focus was on violations to policy in terms and the authentication information was accepted without verifying that the holder of the credentials is the actual owner of the credentials. In fact, we evaluated them against a set of criteria that addressed real time insider detection

Temporal structure and temporal analysis is a theory put forth by Victoria Koehler-Jones³. The two important elements of the temporal structure are in the pattern (or the way time and events move from the past, through the present and into the future) and the flow of those events. The point of the temporal analysis is to analyze and then to provide a quantitative assessment of the importance of the events. The postulation is that each person has a temporal signature or temporal identity. And just as important, the way those events are viewed can cause differences in the inferences. Similarly, IAMPS looked at events in time and viewed them as being related in the temporal sense and then looked for discontinuities in the authentication process.

The IAMPS solution developed an approach to authentication that involved the use of sensors that related all available information about a user. IAMPS shows that fusion of information yields better results in detecting an insider attack and is able to do so in real time. The fusion looks at information garnered from several sources and then draws inferences and conclusions from that set. By combining various inputs, the accuracy of those inferences increases over results achieved from using individual inputs. Conceptually, the process is shown in figure 1.

³ From "Disaster Responders Perception of Time", Victoria Koehler-Jones, University of Nevada Las Vegas http://www.library.ca.gov/CRB/96/05/over_8.html

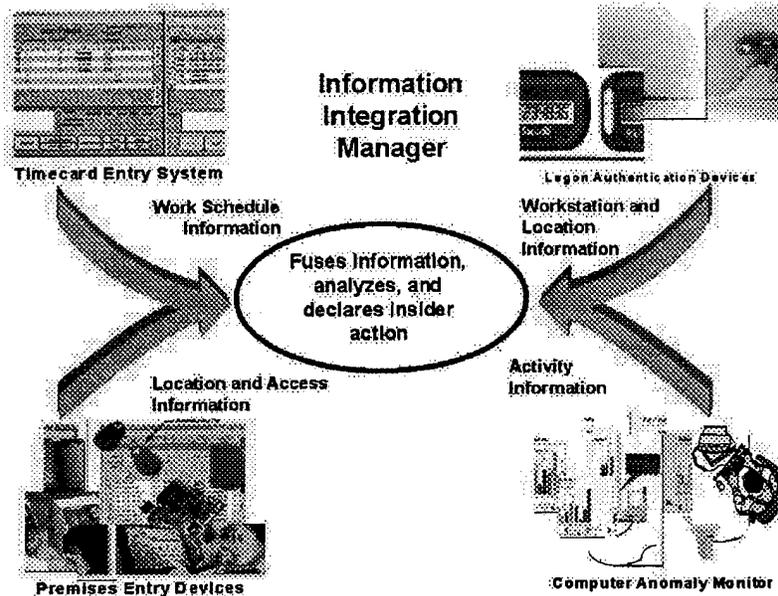


Figure 1
IAMPS Process Showing Fusion of Authentication Information

In essence, IAMPS provides a cutting edge solution to the insider problem. It is straightforward in the sense that it concentrates on the authentication portion of the user access process and continually adds new information to the aggregate and re-evaluates the new aggregate looking for profile consistencies and more importantly profile inconsistencies. IAMPS looks for the inconsistencies during the authentication process to be in the position of denying accesses as opposed to reacting to damage after the damage has been inflicted.

2. BACKGROUND

As mentioned before, the only methods of detecting the malicious insider attack was a posthumous reconstruction of the event. Recall that in the context of this study, the term malicious insider attack is an attack where the intent is to do damage to the data or operational code in a given system. This does not include the basic security infractions such as internet misuse. Those violations, while some consider to be insider attacks, are not of interest since they can be detected and adequately thwarted with firewalls. And in that reconstruction of the event standard techniques in cyber forensics are typically used [Mar01], [Ste99]. These techniques are in some ways a list of what to do, what not to do, and how to do it. The literature completely covers ways of maintaining evidence without contamination of the evidence and also covers ways to look at the evidence. In fact, according to Kruse and Heiser [Kru01], there are three basic steps to computer forensics that they refer to as the three As.

- Acquire the evidence without altering or damaging the original.
- Authenticate that your recovered evidence is the same as the originally seized data.
- Analyze the data without modifying it.

It is the third point that provides the both the focus and the dilemma. The posthumous discussion shows that the real nature of the solutions to date are confined to well after the fact and are using only the information (clues) left by the attacker. Recovery from the attack requires the reconstitution of the information within the system be restored, which as Omega Engineering found out can be almost impossible. And we are relying on the clues left by the attacker that might be clues the attacker *chose* to leave. Using this data to reconstruct the method and focus of the attack One of the possibilities in data that is left is the potential that the data and information may be

The initial step was to define what research areas we could look for that would have a bearing on the problem which is to detect an insider action in real-time. Clearly there are tools and processes available that detect misuse based on security policy, but that only solves the issues of a person trying to overstep policy directly. The concern of this project from the outset was to develop techniques and an approach for real time insider abuse detection to then be able to take a course of action to thwart any intended actions either during or preferably before those actions could damage the system

Part of the challenge to review previous or related research in the area of insider detection was finding the research. Initially, this was approached using two parallel methods. The first was to trace through government workshops on insider detection and the second was to perform a series of on-line searches. The point of the web search was to try and uncover any leads in the area of detecting an insider attack. Several different parameter combinations were used as search items for the online searches. The results are summarized in the table below. The search terms were selected after some time was spent doing more of a trial and error. The table only summarizes those search terms that were useful.

Table 1
Literature Search Parameters and Results

Search Engine	Search Parameters		Number of hits	Useful hits/Comments
	Include	Exclude		
Yahoo	Insider, detection	N/A	10600	Too many to evaluate, needed to prune tree and stick with specific phrase as opposed to word occurrence
Yahoo	"Insider abuse"	N/A	661	N/A
Yahoo	"Insider abuse"	trading	531	N/A
Yahoo	"Insider abuse"	market	339	N/A
Yahoo	"Insider abuse"	Trading, market, stock	293	N/A
Yahoo	"Insider abuse"	Trading, market, stock, news, key, article, fraud	48	Mostly articles or news, one on Profiler 2000 from AFRL
Yahoo	"Insider misuse"	Trading, market, stock, news, key, article, fraud	31	Neumann Publication, Research and Development Initiatives Focused on Preventing, Detecting, and Responding to Insider Misuse of Critical Defense Information Systems; Robert H. Anderson; RAND minutes publication
Yahoo	"Insider attack"	N/A	552	N/A
Yahoo	"Insider attack"	Market, trading	432	Similar results as insider misuse.

In essence, the references fit into a few predictive bins. First were stories and news articles about some insider attacks that were related to insider trading and insider loans. These references, after a while, were ignored, as they offered no direct connection to the problem at hand. The second group consisted of news articles that related to insider attack stories on computer systems mostly either historical in nature or instructive as to what insider attacks (misuse, etc.) really were. Finally, what was left referred to computer misuse or computer abuse in terms of how to look for it, results of research workshops, papers on approaches and some information on products. It is this last group that provided the background literature and any related work. We sought to look for insider solutions that did not fall into the categories of financial or internet policy misuse. We took that a few steps further to also

include any attempts to find solutions that encompassed sensor fusion. There were none that we were able to find.

The literature search provided five related efforts to solve the problem and all provide substantial background information. It is interesting to note that the technology for solving the insider problem is still in its infancy. What follows is a summary of the 5 most related approaches, a short description, and why they did not completely solve the problem at hand.

There are tools available that help with the insider problem, and EMERALD is a representative one [Por99]. The following is an excerpt from the EMERALD website¹. “*EMERALD's eXpert-BSM monitor is a host-based intrusion detection system that provides an unprecedented degree of real-time security monitoring for critical application servers and workstations in the SolarisTM Operating EnvironmentTM. It incorporates the most comprehensive knowledge base for detecting insider misuse, policy violations, privilege misuse or subversion, illegal resource manipulation, and other site policy violations upon operating systems. This fully packaged solution provides users with:*

- *a knowledge base of 39 host-oriented misuse-detection methods,*
- *extensive user ability to configure both the knowledge-base and surveillance policy,*
- *a graphical reporting console for managing sensor alerts,*
- *detailed response directives and human readable countermeasure recommendations,*
- *and real-time and batch data processing².”*

EMERALD was developed by SRI under a DARPA contract F30602-96-C-0294 and applied under contract F30602-98-C-0059. EMERALD provides a number of things, but the most relevant to this research is the ability to detect the misuse of system or network resources. EMERALD, and the family of tools like it, are all looking for IP based and higher level security policy violations. EMERALD and many of the tools in this class provide adequate detection in real or near real time of policy abuse, but fail to address the non-IP related abuse. Related to EMERALD, is the Production Based Expert System Toolset [Lin99], P-BEST. P-BEST looks at some types of attacks, notably syn flooding and buffer overruns.

Also found were tools performing types of integration and fusion. With several tools and tools suites, two vendors were selected that are representative of the integration and fusion tools: Intellitactics and e-Security.

Intellitactics has developed a suite called the *Network Security Manager* (NSM). According to the website information³, NSM has a distributed architecture that allows the monitoring of multiple sources of information from security assets dispersed across the network and/or the entire enterprise. It is a real time monitor and “By consolidating information from various best-of-breed devices, including firewalls, intrusion detection systems, security scanners, databases, operating system logs, and physical security devices, NSM can detect potential

¹ <http://www.sdl.sri.com/projects/emerald/news.html>

² <http://www.sdl.sri.com/projects/emerald/news.html>

³ http://www.intellitactics.com/pdfs/Supported_Devices_Cust.pdf

threats and attacks, occurring anywhere and anytime on your network from a single, unified view. Intellitactics' team of security professionals have spent over 20 person years of effort deciphering disparate code from various devices to make it easy for organizations to distinguish between a network broadcast, Denial of Service, SYN attack, and hundreds of other methods used by intruders.⁴

Specifically, NSM integrates the following sensors: Network based scanners, Vulnerability analysis and hardening tools, Intrusion Detection Log Watchers, Anti-viral products, Network based IDS, Perimeter defense (firewalls), VPN and Crypto Communications, Personal security (personal firewalls) and dial up authentication.

e-Security⁵ has developed a suite entitled Sentinel that also integrates and manages security sensors. The sensors supported by Sentinel include: information security devices such as firewalls and intrusion detection systems, security event network devices, applications and services (for operating systems, databases, and email), servers, and physical security devices (badge readers and process control devices). It is also a realtime system and provides correlation of events and inputs.

As far as the EMERALD and P-BEST are looking for security violations that are related to primarily Internet and/or equipment misuse. But the misuse is detected by looking at violations of security rules and policies and assumes that the authentication process was adequate and accurate. Both the eSecurity and Intellimatics are very robust systems, with real time operation and effective correlation. But this class of suites is primarily concerned with the view over different security events. Events that pass any of the security tests and rules do not appear to receive primarily consideration. They are much wider extensions of EMERALD and P-BEST, but are still 'security rule based' and do not directly address the malicious insider issue. In fact, both the eSecurity and Intellimatics solutions are discussed in section 4 of this report as candidate alternative solutions.

A second approach was to use pattern matching detailed in a recommendation from a University of Missouri proposal entitled "Insider Threat Detection for Robust Security Environment"⁶. The pattern matching approach was applied to securities and exchange transactions, which is vaguely similar to the previously mentioned policy based detectors. This approach is different because the application to the Wall Street Insider problem begins to get a solution that is closer to real time. After looking at this and analyzing, the difference in real time relates more to the process that securities and exchange go through and the ability to form some precursors of attack. These precursors are in the form of evaluating a list of potential insiders and any of those insider's actions as related to stock movement.

A third approach was focussed on profiling a theoretical attacker [Wod99] to model and understand the adversary.⁷ This approach characterizes the insider and provides a

⁴ http://www.intellitactics.com/html/nsm_overview.html

⁵ <http://www.esecurityinc.com/>

⁶ <http://www.umr.edu/~ff/Insider/proposal.pdf>

⁷ <http://www.rand.org/publications/CF/CF163/CF163.appb.pdf> provided a model for insider behavior as developed by Brad Wood of SRI International and this dealt primarily with profiling of the insider

preliminary trusted model of the characteristics of an insider. This in turn could spark “valuable insights and other observations that may lead to effective mitigation strategies for the Insider threat.”⁸ The point was that this effort was focussed on a *description* or *profile* of an attacker and this paper provides significant insight, but does not address solving the detection issues.

Another potential contribution came from a SANS Institute article⁹ by Terry Boston entitled “the Insider Threat”. This article was a good fundamental introduction to the insider problem and in addition described the social engineering aspect. At the conclusion of the article, it presented some solutions that were all process based¹⁰. The process approach, while part of a final solution set, does not address any technology issues. Process approaches are important, since good and effective security is composed of people, process, and technology.

The fifth approach analyzed [Neu99] analyzed presented a good problem decomposition and laid out requirements for solutions. In particular, Neumann called for “integrated detection that looks at hierarchical and distributed correlation using different sensors.” These requirements address directly the focus of the IAMPS research. Again, there were no direct solutions, but a good quantification of what problem a solution must address and what requirements the solution should have.

There are solid reasons why any previous attempts to solve the problem did not succeed. Foremost is that they did not address the real time nature of the problem from a true insider perspective. True insiders gain access and privilege and therefore “look” to the system like a legitimate user and in particular “look” to the system like a specific legitimate user. Even though subtly stated, there are two cases to look at as a result of that statement. The first case is an insider that already has access and has turned malicious. The second case is someone who, through any one of a number of means, has someone else’s cyber identity. Most systems do not authenticate beyond typical user id / password combinations so if someone can gain possession of a user id / password set, then as far as the system goes they are legitimate and from the cyber perspective are who they purport to be. In either event, there is someone with some level of access about to compromise part of the system.

The real time nature of the problem means the ability to catch the offender while the cyber crime is being committed. No solution found addressed this issue except in the policy

⁸ from the INTRODUCTION of the Brad Wood paper

⁹ http://rr.sans.org/securitybasics/insider_threat2.php

¹⁰ The following list is representative of those process recommendations and is found in the summary of the paper.

- Run backups and secure in a safe place.
- Install and execute appropriate anti-virus tools.
- Regularly check for viruses.
- Ensure that your computers have the most recent versions of anti-virus software.
- Update your anti-virus tools using vendor updates as they become available.
- Store copies of anti-virus tools offline, in a secure manner.
- Install software updates and patches.
- Routinely check for and update threat detection tools as needed, especially when new threats are discovered.
- Prevent unauthorized outgoing access at the firewall.
- Ensure that permissions are properly set.

determination. With this research being focussed on the act being committed due to authentication and verification issues, the policy based initiatives while interesting contributed minimally to the body of knowledge.

The second reason for previous attempts failing was related to single point checking. Verifications of users assumed that the real owner was using the user id/password combinations, as were access cards or tokens. No attempt was made to look across several pieces of data and look for inconsistencies.

Finally, there are forensic solutions that look at evidence after the effects of the intrusion. For the nature of posthumous evaluation, these are inadequate. The reasoning supporting this stems from the fact that a posthumous evaluation can only look at the final result of data and that data could have been manipulated as part of the attack. This could lead to drawing erroneous conclusions about the nature and perpetrator of the attack.

3. METHODS, ASSUMPTIONS, AND PROCEDURES

The methodology that was used to evaluate alternatives was identical to the methodology that was used to evaluate the IAMPS solution. Any methodology needs to be consistent across and between solutions, therefore it is important to set the standards so they would apply equally. Because of this, we did not see the need to either develop or consider alternative methodologies, this one suffices for all flavors of solutions. The solution set had to be consistent and any alternatives must adhere to that consistency. Primarily the methodology is concerned with five basic features of a solution:

- Paramount is real time nature of solution. The primary goal of this program is to detect insider attacks as they happen and use that detection to prevent any data or information compromise.
- Second was minimal if any interference by operators.
- Third is the ability to recognize, collect, and maintain forensic evidence that is germane to prosecution
- Fourth is the ability to select and implement courses of action for both immediate and longer term
- Fifth is the ability to select from and integrate a wide range of sensors.

The goal of this effort is to be able to identify insider attacks as they happen. In fact this translates into *non-repudiated authentication* that guarantees that the user is in fact who they purport to be. In effect, this is identical to determining that the user is not who they say they are. This in and of itself is a sufficient condition to declaring a security event and taking appropriate actions. The extension to this is to determine who that person really is which is beyond the scope of this effort. This step, a far more difficult problem at this time would declare Non-repudiated identification. This would determine the actual identification of who the person really is.

- There are four primary assumptions on the solution:
- Assumption #1 – We would seek to limit any need for a fixed solution set. This assumption is meant to guard against a solution that only works on a specific processor with a specific system. Technology transferability is one of the keys to this program and flexibility in the implementation is important.
- Assumption #2 – Use of commercially available sensor suites and develop a process and an algorithm concept that uses this available information in a new way. This is an additional support for the transfer of technology theme.
- Assumption #3 - Information from sensors could be made available electronically. The sensors being considered went far beyond the scope of just firewalls, intrusion detection system, etc. The sensors to be considered included physical access devices and external to the computing system sensors.
- Assumption #4 - Wanted to consider any piece of information that could be tied to an individual or the individuals' machine. Again this opens the fusion potential.

As stated before, the purpose of this research is to develop theory and recommendations to catching computer abuse / computer misuse in real time so any collateral damage could be

prevented. For the purposes of homing in on the solution set, we need to define up front the forms of insider attack we are looking for. The fundamental observation that can be made is that an outsider with access is an insider. What this means to the research is that we did not differentiate between the outsider and the insider. Once a person has access to the system, they are considered by the system to be an authorized user and hence an insider.

Considering all the various possibilities, there are four bins that users can be 'sorted' into:

- Case 1 - Users are who they say they are and that fact can be supported by several pieces of data. These users are not intentionally malicious.
- Case 2 - Users are who they say they are and that fact can be supported by several pieces of data. These users are intentionally malicious.
- Case 3 - User(s) are masquerading and are not who they claim to be, even though they have passed some level of authentication in the system, and that by looking at several pieces of data (evidence) we can develop that inconsistency. In addition, these users are not intentionally malicious and by carefully selecting and combining (fusing) data we can support that.
- Case 4 - User(s) are masquerading and are not who they claim to be, even though they have passed some level of authentication in the system, and that by looking at several pieces of data (evidence) we can develop that inconsistency. These users are intentionally malicious.

It is case 4 that is the focus of this research.

4. ALTERNATIVE SOLUTIONS AND EVALUATION

The fundamental requirement we placed on the solution was for the detection of the insider activity to be in real time. To put a finer point on that, the solution had to have the ability to detect in enough real-time to counteract any potential damage and be in a position to gather electronic evidence. Both portions of those requirements were crucial.

The evaluation of any alternative approaches was fairly straightforward. First, solutions employing other forms of sensor fusion were viewed as being the same as this one. The view was that an alternate fusion approach could be using different fusion engines or placing different weights on data or even using other data points. We did find related solution sets in the background research, eSecurity and Intellimatics. The criteria discussed in section 3 applies to these solutions and the results summarized in the Table 2.

One clear way to evaluate them adequately would be to compare the e-Security and Intellimatic suites with the implemented results of this research in a series of real tests. Both the implementation and any comparative testing was out of the scope of this effort, but could be an excellent mechanism for further research. Second, we rejected any solutions that were not real time in nature. This ruled out solutions that use the post-mortem approach. Post mortem solutions have a fundamental weakness in that the analysis is subject to only looking at the information and data residuals after the incident has occurred, and many times long after the incident. But more fundamental is the fact that the intruder could have left only the information that he/she wanted us to have. During the attack, the data could be manipulated in such a way that an analysis of only what remains could lead to erroneous conclusions.

Table 2
Developed Evaluation Criteria

Criteria	Comments and Evaluation	
	e-Security	Intellimatics
Solution operates in real time	Yes	Yes
Minimal operator interference	Yes, can involve operator if need be	Yes, can involve operator if need be
Maintain forensic evidence	Yes	Yes
Implements courses of action	Yes	Yes
Wide range of sensors	Yes	Yes
Technology transferability and portability	Yes	Yes
Identifies user(s) are masquerading and are not who they claim to be, even though they have passed some level of authentication in the system	No	No

5. PROPOSED SOLUTION

In looking at a solution, the goal was to gain *non-repudiated authentication*, or being able to prove two points:

- A person is who they claim to be and we can support that claim empirically through electronic evidence
- A person is not who they claim to be and we can discover the discrepancies using electronic evidence

Looking into past histories of authentication was a starting point for this research. Originally, people needed identification for social reasons. As interactions between people became more complex, proving who you were took on economic implications. Initially, names were the first form of identification with the appearance of surnames in Britain as early as 1066 [Cla94], [Fox06]. By the year 1538, parish priests began to keep registers of births, deaths and marriages for identifying purposes, many of which can still be found today. Interestingly enough, as early as 1300 in Britain, passports began to be issued. [Ehr66].

There is a variety of means for identifying a person's identity:

- name (what a person is called)
- possession (something or things that a person owns)
- descriptive appearance (how the person in terms of height, gender, weight)
- codes (what other names a person may have)
- knowledge (what the person knows)
- details of physiological features (fingerprints, facial characteristics)
- miscellaneous attributes (what the person is now, e.g. tags, collars, bracelets).

All authentication has the goal of protecting a system against unauthorized use. This means that there can be protection for legitimate users while taking actions that would deny or prohibit the impersonation of authorized, legitimate users. Authentication procedures are based on the following approaches [Woo99]:

- *Proof by Knowledge*. The verifier known information regarding the claimed identity that can only be known or produced by a principal with that identity (e.g. passport, password, personal identification number (PIN), questionnaire).
- *Proof by Possession*. The claimant will be authorized by the possession of an object (e.g. magnetic card, smart card, optical card).
- *Proof by Property*. The claimant directly measures certain claimant properties using human characteristics (e.g. biometrics).

Many of us carry issued identification cards, know passwords and PINs and use them in order to identify ourselves. In many areas, security is provided by requiring the use of badges, special badges or rules for visitors and possibly even the issuing of keys [Car95]. These are the most common means of identification since they have been the easiest to remember and the easiest to confirm.

The postulation that began this research centered on four claims summarized below:

- First, the overall premise of this approach is that there is a measurable steady state of system performance parameters that a given system operates in terms of memory reads/writes, disk accesses, power used, internal ambient temperature, etc. While any one of these could fluctuate, it is our technical judgment that several of these being out of “steady state” ranges could indicate unauthorized access or use of the computer system.
- Second, we believe that there are other available measurements that exist that when viewed in their entirety, can potentially declare an insider attack and can perform that analysis in real time. The steady state analysis results can be fused *in real-time* with timecard, premises, and authentication information, also available in real time, to augment the declaration process and offer fidelity and certainty into the accuracy of the prediction
- Third, we postulate that the steady state measures of system use offer a new and creative method of profiling an attacker based on those measurements. This implies that an additional and empirical measure can be taken and used as “cyber evidence”.
- And finally, this technique detects the computer misuse in real-time without increasing the system overhead or workload.

In the course of the research, we uncovered information that in some cases supported and others modified our initial postulation. Relating that back to the original claims, we found:

- Sensor fusion where various corroborating pieces of information could be ‘fused’ offers significant promise and should provide a solution space for detecting insider attacks in real time.
- Some sensors do not provide primary information, but have one or two attributes that can support some claims.

“ It is of the highest importance in the art of detection to be able to recognize out of a number of facts which are incidental and which vital. Otherwise your energy and attention must be dissipated instead of being concentrated”. – Sherlock Holmes, **The Reigate**.

The main postulation of this research activity was that fusion of information would yield better results in detecting an insider attack and is able to do so in real time. Most security around computer networks has forms of physical security, such as locks, electronic entry devices, etc. These are monitored by security personnel who look only at the entry/exit logs with any correlation done manually and well after any fact. For example, if something is stolen from a closed area, the logs can be examined to see who may have had opportunity and was present when the event occurred. Equally as important, the logs can be examined to see who wasn’t present, and they can be ruled out as suspects. This example taken with several sources of information is exactly the context of data fusion as it relates to this research. Data fusion looks at information garnered from several sources and then draws inferences and conclusions from that set. The expectation is that by combining various

¹ From the website of Sherlock Holmes quotes, www.bakerstreet221b.de/canon/index.html. There are a number of quotes across the works of Sir Arthur Conan Doyle that are applicable to the concept of fusion. In fact, Sherlock Holmes was probably the first literary detective to demonstrate the art of data fusion as applied to solving crimes.

inputs, the accuracy of those inferences would increase over results achieved from using individual inputs. The two questions we are trying to answer are:

- Are there authentication inconsistencies? and,
- How can we substantiate?

The first type of correlation looks at the authentication consistencies and the authentication inconsistencies to ascertain that a given person is in fact, who they claim to be and. Consider the case of a valid user id/password combination, a token for access into other parts of the system, devices (such as card readers) that allow physical access to computer equipment and information from the timekeeping system. Suppose the following case is considered: a Joe ID – Joe Password are valid users of the system, Joe Access Card, and Joe's record of his time for this week. One form of consistency is that Joe used his access card to gain entry to the area, then he used his Joe ID and Password to log on to the system the system, and Joe has time entered on his timecard for work on Project X. Taken one at a time, each of the data items are valid and would indicate that the person logged onto the system is in fact, Joe. This is an *authentication consistency*. But suppose the time card indicates Joe is on vacation. Again, each of the individual data items are valid and if taken one at a time, would indicate Joe could be on the system. But taken as a unit, they show an *authentication inconsistency* in Joe actually being on-site. At this point, there are courses of action that can be taken, the least of which would be to terminate the access from 'Joe' until a follow-up can be completed. In essence, we are developing a process where we are trying to prove that someone logged in to a computer is either who they claim to be or is impersonating a legitimate user, we have a good fit. In fact what we are trying to accomplish is to view the data, correlate where appropriate and present evidence.

Given that, we began to examine types of data fusion with intent to specify one over another. For example, is this the type of application that better suits Bayesian networks or should a Dempster-Schaeffer approach be used? The plan was to investigate both approaches, start to look for ways to potentially implement, but it was decided early in that process that hybrid approach would work best. The primary reason for that is there are methods available that can be employed to defeat certain forms of fusion if the detection fusion approach is known. If the fusion implemented is a hybrid of more than one approach, there is some degree of certainty that the algorithmic integrity can be preserved.

The concept of minimal essential [McC01] was developed under USAF Contract F-30602-97-C-0132 and related to the smallest subset of information that would be required, through various formulas and re-computations, to populate an entire data set. The concept was developed to support the real time recovery of systems so that operational continuity could be maintained. Fundamental to the continuity is the ability of the system to (in realtime) reconstitute data that may have been compromised as the result of an attack.

MEDS in the database context would be the smallest subset of the database that could be used to reconstitute the entire database. The subset could be a specific section of the database or could be a collection of data items across the database. It is a pro-active process where

data on a cyclic basis is gathered and then hidden² within the system and then made available in the event of a cyber attack. Once an event happens these trusted copies of information are retrieved and used in a series of pre-determined mathematical and logical computations to “fill out” the remainder of the data-base.

Consider, by way of an example, a ground based radar system that is either air defense or air traffic control. Both systems maintain target tracking, with the only differences being in the way the targets are actually processed³. Typically, targets are known or unknown, hostile or friendly and the databases that maintain information on these targets are usually fairly large. Should a cyber event wipe out portions or all of the database, we need to reconstitute the database in order to continue with the mission. In this case⁴, hostile targets were considered to be the minimal essential set. The reason for this is straightforward in the sense that a radar system gathers data on every sweep of its antenna. At that time, all targets ‘seen’ on that scan are correlated to the hostile targets first and then the friendly or unknown targets. If the hostile targets are the ones considered as the MEDS, then the database can be replenished with its hostile targets and then anything ‘seen by the radar’ that doesn’t correlate can be assumed to be friendly or unknown. In any event, the targets of specific interest and the main function of the system can continue without rebooting or restarting.

MEDS doesn’t care as much about the *value* of a piece of data, rather it cares about the *relationships* that a piece of data has. Those relationships if taken in the aggregate then define the transformation computations for all data in the system. MEDS, as a concept, has great similarities to the mathematical concept of basis vectors. Basis vectors are by definition the smallest set of linearly independent vectors that, taken in some combination, completely span the vector space. It is a similar relationship that MEDS exploits to span the data space of a given system.

The data half-life concept [McC01] also emerged under Contract F-30602-97-C-0132 and relates to the time currency of the value of a piece of information. Again using the example of a system given above, consider the relationship between time and target position and target velocity. Systems use time, position and velocity to extrapolate ahead to the next expected target position in order to maintain tracking filters on each target. Using a simple formula of:

Target Position_{t+1} = Target Position_t + Target Velocity_t * Δt , where Δt = t₂ - t₁

This formula works extremely well if the time intervals are close together, but the value of target position at time t_{1,2} might not be useful if the system is at t_{1,44}. This is the essence of the concept of half-life. Similar to the concept of radioactive half-life, information also has a period or window of validity. When combined with the MEDS concept, half-life helps

² possibly using steganographic approaches

³ An air traffic control system is the airborne eyes of civilian aviation. Its function is to maintain separation between civilian aircraft to avoid collisions. Air defense systems seek to bring targets together, so that friendly forces can eliminate hostile forces.

⁴ This case is exactly the one used in the Demonstrating Resiliency in Information Warfare (DRIW) program, USAF contract number F30602-99-C-0010. In fact, DRIW successfully demonstrated in real time the recovery of information compromised as the result of a cyber attack and recovered that information in real time with minimal system overhead penalties.

define how current a piece of MEDS really is. In fact, half-life helps define the intervals by which MEDS are gathered and stored.

As far as half life of information is concerned, there are situations where data is relevant and other situations where the data is stale. For example, there are situations where the time of a data input will have some parametric relationship with other information and there will be instances where the time associated with a piece of information has no relationship and in fact may even be parametrically indifferent to the entire computation. To better understand this, consider a site where the entrance to the site and all the buildings at that site require a card to open/access them (figure 2)

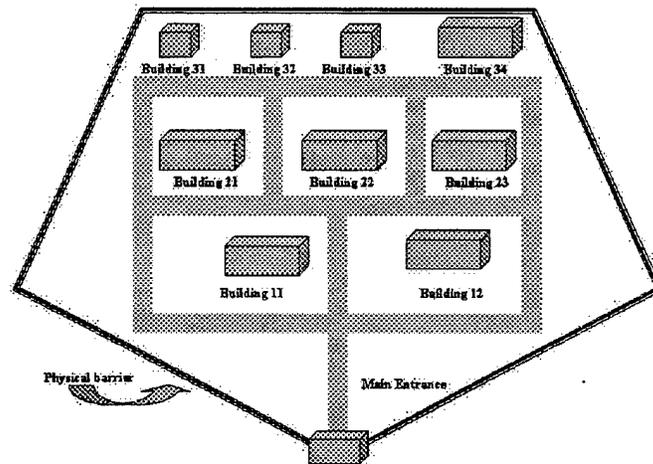


Figure 2
Sample Physical Layout

This raises an issue, namely can this information be of use in the fusion process and under what conditions is it useful. What might be relevant here is the time that a person went through the sensors because we could use that time to establish presence at a specific place in the compound and also to establish a conflict in location. To explain that, we can look at this case a little closer, through a hypothetical scenario. Arriving at the main gate, the time of entry (t_e) for employee n is noted from reading the access card swiped in the reader at the main gate. These access cards, similar to the one pictured below (figure bb) are connected to a main computer where all the accesses are recorded and saved. At some time, t_f , employee n enters building 22 and that time is noted. A third piece of information is the relative time it takes to travel from the entrance to building 22, t_{e22} . The table below summarizes the possibilities and relates them to relevance. A word about relevance in the case of looking at an insider attack, we are looking for those data entities that either *authenticate consistently* or *authenticate inconsistently*.

Table 3
Relative Measures of Half Life Concept

Case	Fusion Implication	Relationship to Relevance
$t_f < t_e$	Employee entered building 22 when there is no record of employee being on premises. Very dependent.	Very relevant. Chances are employee is not who the system thinks they are and is masquerading as someone else. Shows violation of security policy and potential for insider sabotage
$t_f > t_e$	Employee enters building 22 at some time after entering premises properly.	No conclusions can be drawn from this single time slice.
$t_f = t_e$	Employee enters building 22 at same time as initially entering the base.	Very relevant. There is a possibility of multiple copies of the same access card.
$ t_f - t_e > t_{e22}$	None. Employee didn't go directly to building 22 so any useful relationship between the data is minimal. Parametrically indifferent.	No conclusions can be drawn from this single time slice.
$ t_f - t_e < t_{e22}$	Employee has arrived at the building outside the bounds of possibility	Relevant in the sense that someone else is using a copy of the access card.

In the five cases illustrated above, relevancy was found in three of the cases and in all three cases there was a good notion of the half-life. In the other two, the information had lost its temporal relevancy and was outside the useful half-life window.

Far and away the most interesting portion of the research dealt with examining sensors for potential use in the fusion process. The initial departure point for the research postulated the use of four sensors as shown in figure 3 below.

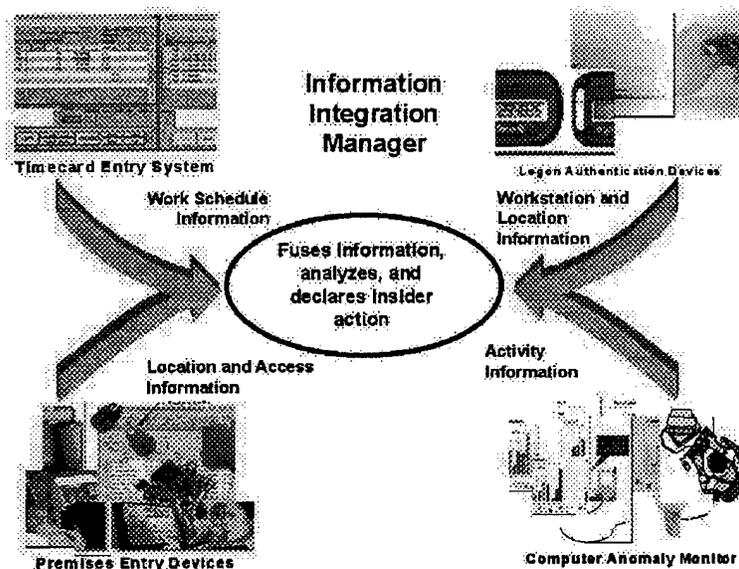
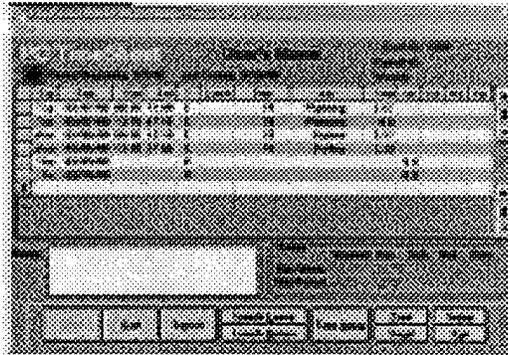


Figure 3
Initial Conception of IAMPS Solution

The initial concept used the information from each sensor, all postulated to be easily available. A separate discussion of candidate sensors follow, but the theory behind the fusion is best explained via an example. Consider the following case: the main server detects a logon from Joe Employee inside the building at 12 am. Joe's ID/password combination is legal and as far as the system is concerned they all belong to Joe who is authorized for the system. However, the premises entry device indicates that the person is not on the premises at this time and a check of the timecard system shows that Joe has signed for vacation. Clearly we have uncovered a highly suspicious access and can take a course of action such as to terminate the connection to the server.

What follows is a discussion of some generic sensors, how they could be implemented in the fusion scheme, and what types of information we could expect to deduce from the sensor and what types of information could not be deduced. The initial selection of sensors was done based on what ones were more commonly used and therefore the fusion approach could be retro-fitted. Biometric and facial recognition sensors were not initially part of the research, but the events of September 11, 2001 placed some additional emphasis on those types of sensors. There was some investigation into using biometric sensors, and what was garnered from the research is presented after the initial sensor discussions.

One of the initial sensors to be selected was from a timecard system, figure 4. Most businesses and organizations have some kind of automated method for tracking employee time and the information from that timekeeping system can be queried. One of the keys was to identify what information we would want and when in the fusion process did we want it. Some generic characteristics:



Timecard Entry System

Figure 4
Example of a Timecard System

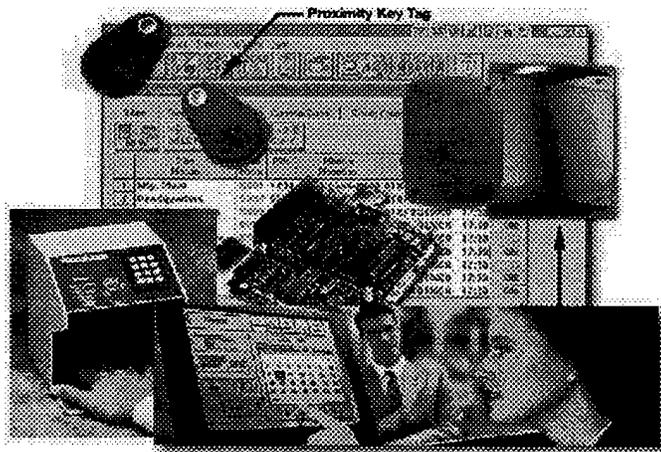
Timekeeping system can tell what kind of time a person has charged (project, overhead, vacation) with the vacation time being potentially more important than the others. For example, if a person is trying to log into the system when the timecard system indicates they are on vacation, then there is a discrepancy that needs to be addressed

There are many more items that the timecard cannot tell:

- The exact times a person was working,
- Where the person was working
- No guarantee that the information is in real time since time can be entered before worked or before vacation
- No guarantee that person charging time is even in plant, since project hours could be worked while on a trip

The information within a personnel system is usually highly proprietary and extremely sensitive, which imposes some constraints in the access to and use of data. To alleviate the concerns of data sensitivity, queries would be constructed that would essentially only ask if the time charged was vacation or not. Given the list of shortcomings versus the potential strengths, the reliance on timecard information was reduced.

Premises entry devices are a very good source of information that has a wide range of applicability in the insider process and then also for other forms of security in general. Most systems, such as the ones shown in figure 5, record all the times when the device is accessed.



Premises Entry Devices

Figure 5
Typical Premises Entry Devices

Typically, a user is assigned either a token or has an access code for the keypad or cypher locks. When the user employs the token or keypad to unlock the door, the device records user information and the time of access (or egress). The device itself or the system behind the device then looks into its database and validates the user (token) at this location and then if the user is authorized to enter this location, the device unlocks the door. The user information, or what can be garnered from the access device and its associated database, is the relationship between the token or keypad code and the specific user it is assigned to. In a perfect word, the user with a token is the user who was assigned that specific token.

What premises devices are able to supply to the fusion process are:

- What entry is being accessed in terms of physical location and possibly coordinates
- Whether it is entry to an area or egress from an area (which side of the door the device is on)
- What token or keypad sequence is being used
- Identification of the user assigned to that token or keypad
- Time of entry/egress

There are some pieces of information the premises devices cannot supply:

- What the user does in the area they accessed
- If the controlled area is large (a large building with the access control only on the main entrance) then the information may not be as meaningful. The user could be logged as entering the building but without any more information, an intrusion could still go unnoticed as the users movements within the building cannot be traced.
- Does not guarantee that the person using the token is in fact the person assigned the token

The main difficulty with using this information is the users themselves. The reality of this situation is that it is possible for a legitimate user to gain access to the facility without using his or her keycard by having someone hold the door open for them. It would reason that

having no record of someone's entry does not necessarily mean that they are not in the building. This relates heavily to the process portion of the people-process-technology triad. Process changes can be made in order to ensure that users verify their entrance to the area and no 'tailgating' occurs. There are some access systems that are only on one side of the entrance to an area and therefore do not register when a person leaves an area. In those cases it would mean that having a record of a person entering the building does not mean that they are still there.

Logon authentication is very similar to premises access systems except they are geared strictly towards the logon process with computers/workstations. There are wide ranges of devices, as in figure 6, and the information that becomes available is a function of the sophistication of the device. Some devices only require a password, some require that a token be used, and some are more complex biometrics. In addition, many of the same physical access mechanisms have been adapted to work as authentication devices for computer systems.

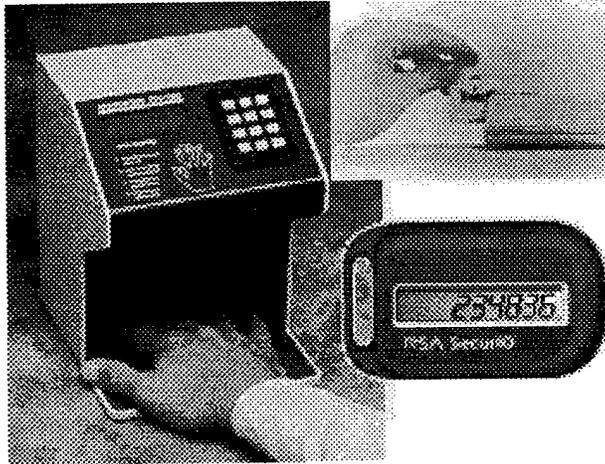


Figure 6
Typical Logon Authentication Devices

What logon authentication devices are able to supply to the fusion process are:

- The subnet address of the computer being logged onto.
- Time of the logon attempt
- What device is being used
- Identification of the user assigned to that device
- Possible location and coordinates of the login attempt.

There are some pieces of information the logon devices cannot supply:

- What the user does in the computer system they accessed
- Does not guarantee that the person using the token is in fact the person assigned the token, although a higher degree of certainty comes with using the biometric devices.

Sensors that *profile* individuals are useful in the sense that they can issue alerts if there are discrepancies between the expected profile and the observed action. Profiling sensors have been controversial due to issues raised over potential violation of privacy issues so tying a specific profile to a specific user, such as keystrokes, was not considered. The reasoning is as follows. If a tool were to monitor the keystrokes, then it could be possible for that tool to reconstruct other information from those keystrokes such as passwords or content. For this reason, any profile that was developed by looking at software was thought to be potentially controversial and in keeping with the desire to be able to transfer this technology easily, was ignored. This allowed for some research to be done on developing a profiling technique that did not monitor soft information. From this, a sensor was conceived that would measure hardware parameters and use those to develop a profile. Since nothing about information content is used, and information cannot be constructed, this approach should avoid all the pitfalls of more software-oriented sensors.

The way this sensor would work would be to have access to various spots in the hardware and register level instructions

A representative set would be:

- Number and frequency of disk accesses,
- Read/write accesses requested
- Average memory fetch cycles
- Average memory store cycles
- Internal bus traffic in terms of amount
- Disk /memory swaps
- Input / output
- Ambient temperature
- Power drawn
- Email
- Erases
- Copy
- Idle
- Loads/Stores

Given the above as an initial set to monitor, we can develop a profile of a user and then relate that profile to a period of time.

What the computer environment analysis devices are able to supply to the fusion process are:

- Provide a measure of the system being used (in hardware parametric terms) consistently against the measured profile.
- Provides an alert if the system is being used inconsistently with the user's everyday patterns
- Time of use

There are some pieces of information the computer environment analysis device cannot supply:

- This system cannot accurately identify the user. There are no authentication routines within this system. It will only detect a suspicious change in the user's activities which maybe be suspect and it may be an innocent coincidence.
- It does not actually authenticate the user by itself. Anomalies tracked by this system may be the work of the actual user. Simple mistakes on the users part could inadvertently trigger an alarm.

This type of monitoring offers some significant features:

- This system is very subtle and low-level. It would be very difficult to evade even if you knew of its existence. The only way to avoid triggering it is to try to mimic the legitimate user's activities on the system. This would make it difficult to accomplish any sabotage or unauthorized accessing of files.
- At more of a system level, it can alert if unauthorized attempts are made by a user to access files or applications.

This functionality could possibly hold the best chance of detecting abuse. Unfortunately it is also very subjective and possibly fraught with error. The information is subject to interpretation. What may appear to be suspicious could simply be an unusual task being undertaken by the user (requiring extra processing or such). The anomalies tracked by this system are less ambiguous than the System Behavior Profile (an increase in Memory Utilization may or may not be cause for alarm but repeated attempts to access a classified database would be).

Consider the example in figure 7. In this example, we are measuring disk accesses, the ambient temperature, read/write commands, memory swap to interleave and software, RAM usage, the power being drawn, EMAIL in and out activity¹, erases of memory, copy commands, and idle time. This is not a case where there is any concern, not necessarily because there is only one parameter being out of tolerance, but also because the most plausible explanation is that there is a Windows heap issue making the RAM usage measure out of balance. Of the ten parameters that are part of this measurement grouping, only 1 is out of expected tolerance. Here there would be no alert, however the computer environment monitor may want to keep on eye on this combination because 5 of the 10 measurements are rising. If all or part of them exceed the nominal limits, there is the potential of multiple copying and/or multiple emailing.

¹ No intention to perform any content analysis, just the number and sizes of primarily outgoing emails.

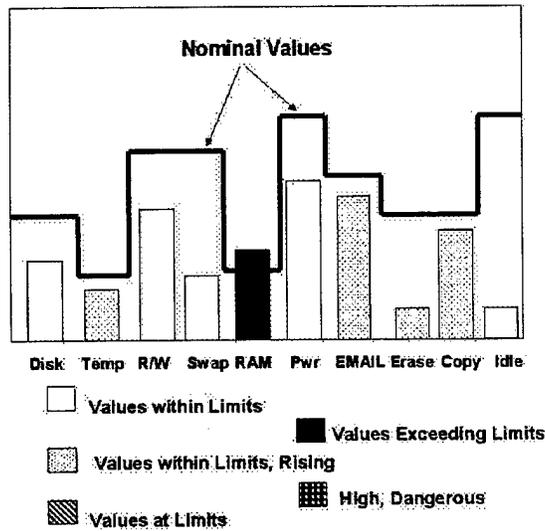


Figure 7
Computer Environment Analysis – Nominal Case Example

Next, consider the following case in figure 8. Here there are several measure that the monitor has made. There are excessive disk accesses coupled with many read/write commands. There are also an unusually high number of erasures possibly indicating a purge of memory. In all 5 of the 10 measured parameters indicate a problem

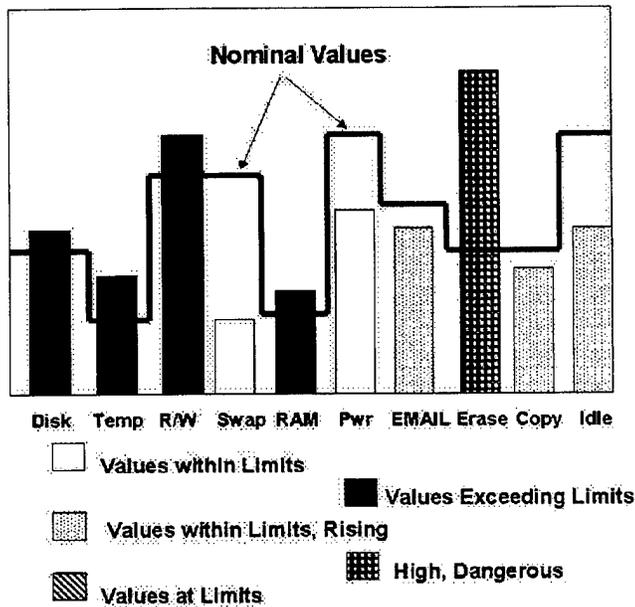


Figure 8
Computer Environment Analysis – Bad Case Example

Although not part of the original scope of the IAMPS effort, it was added in late in the research due to its strong authentication potential. According to the searchSecurity website definition², *“Biometrics is the science and technology of measuring and statistically analyzing biological data. In information technology, biometrics usually refers to technologies for measuring and analyzing human body characteristics such as fingerprints, eye retinas and irises, voice patterns, facial patterns, and hand measurements, especially for authentication purposes. Often seen in science-fiction action adventure movies, face pattern matchers and body scanners seem about to emerge as replacements for computer passwords. Fingerprint and other biometric devices consist of a reader or scanning device, software that converts the scanned information into digital form, and wherever the data is to be analyzed, a database that stores the biometric data for comparison with previous records. When converting the biometric input, the software identifies specific points of data as match points. The match points are processed using an algorithm into a value that can be compared with biometric data scanned when a user tries to gain access.*

Fingerprint, facial, or other biometric data can be placed on a smart card and users can present both the smartcard and their fingerprints or faces to merchants, banks, or telephones for an extra degree of authentication. IBM, Microsoft, Novell, and others are developing a standard, called BioAPI that will allow different manufacturers' biometric software to interact. There are privacy concerns about the gathering and sharing of biometric data, however. One suggestion to assuage those with privacy concerns is to encrypt biometric data when it's gathered and discard the original data to prevent identity theft.”

² http://searchsecurity.techtarget.com/sDefinition/0,,sid14_gci211666,00.html

Biometric devices offer a strong impact on authentication and are clearly an important component of the fusion and declaration process. Chances of false correlation decrease since the biometric information is inherently difficult to subvert. This would require access to the database to modify the comparing information. The biometric device is not sufficient alone to guarantee authentication, but it is expected that the information could be given larger weight in the fusion process. As far as evidence gathering is concerned, a link can be made between the biometric information from an access and a cyber event.

Facial recognition is a special case of biometrics, but operates in a similar manner as other bio information such as fingerprints, hand geometry, etc. In fact, the ensuing argument using facial information as the example works for any form of biometric sensor. The facial device provides information that can be used to make a positive identification of who the masquerading insider is, but this correlation requires a significant amount of database cross-cooperation. If there is a face from the sensor, a secondary form of identification could be requested (such as a license, passport, picture badge, military id, etc.) to make a comparison. In turn, this information could be verified through interactions with the issuing agency database to compare the holder of the credentials with the credentials as the issuing party has them. To accomplish this depth of image access and comparison a great deal of cooperation between issuing agencies must occur. Due to the international nature of some of this information and the likelihood that this is used at different spots around the world, the real time requirement is jeopardized under current communications limitations. This potential lag in accessing information from a remote site or from a site experiencing an overload of requests could cause problems with the fusion engine.

An alternate implementation approach is to have all the faces of interest already contained inside the system. These faces of interest could be users that are allowed into the system or conversely users who are of special interest (terrorists). The faces of interest should be a significantly smaller set than the general population, and the adverse timing impacts of querying for information around the world are eliminated. The comparison engine in the sensor can then indicate that a user is on a 'good' list and therefore allowed entry or if not found on this 'good' list, the user attempting access can then have that entry or access denied. Conversely if the facial sensor has found a person of interest, the system can take a course of action to prevent access and/or alert authorities. In section 7 of this report, we recommend a follow-on activity that uses this concept to identify terrorists in mass transportation settings.

The course of conduct of this research was to culminate with the theory for developing triggers in the declaration process. In fact, the argument for fusion of information is the theory behind the triggers and presents the only true approach to looking at the picture in its entirety. The expectation of the sensor process is that by fusing the information, the various outputs of the sensors will support one another for either a consistent authentication or an inconsistent authentication. A consistent authentication will occur when a user is who they claim to be and all forms of sensor inputs support that in terms of a continuous response. The inconsistent authentication occurs when there are one or more discrepancies between the information from the sensors and the response from the sensor is discontinuous.

In a perfect world, there is an expectation of information and conditions of the data taken from the sensor and used as part of the fusion process. The following table presents the sensor information that is to be considered as part of the fusion process and some comments on how the fusion engines need to process that information.

Table 4
Authentication Sensor Parametric Summary and Fusion Impacts

Sensor or Information Source	Expectation of sensor information	Comments
Password	Matches user Id, authenticates only user ID	Does not ensure that user is who the user appears to be. Password/Id could be compromised.
Secure ID Token	Ties holder to a specific password and user id combination. Authenticates that combination.	Imposter must have all 3 pieces of information to successfully pass this level of authentication.
Biometric ID including facial recognition	Ties issuer to a specific person in terms of something that person <i>is</i> . Usually used to authenticate a specific user id and password combination. Stronger authentication	Difficult to subvert, unless change in the database was made. Certain forms of biometric subject to tampering (false iris scans, fingerprint mask), however it is difficult. Doesn't indicate intent
Sub Net Location	Ties specific machine/node on the network to a specific fixed physical location.	Useful in the computations that show "being in two places at the same time". Involves not using dynamic addressing for network users. Fixes a network IP address to specific location through access tables.
Premises access, cypher lock, card access to areas	Authenticates a holder of the device is on premises and, potentially, where on premises they can be found.	Does not guarantee that holder of access token or password is the user assigned that token or password. Also if tailgating, or more than one person entering area on single access, could generate false positive. Need a policy (process improvement) to require every entry to use entry device.
Time Card	Provides idea of where user is applying time. Useful when indicating vacations or out of office time.	Weak authentication in long run, useful as additional piece of puzzle in terms of knowing that user is out of town or out of office when users credentials show up at other portions of system
Computer Environment Analysis	Ties a user to the way the user employs computer resources. Provides profile against facets of hardware parameters as opposed to software parametric profiling.	With many pieces of information being part of the monitored parameters, user profile is difficult to imitate. Providing the profile across the network allows for monitoring at any workstation in the network and is additional way of locating user through fixed IP address-physical location relationship.

Given that information, there are some observations on the fusion process that can be made:

- Consider some of the combinations, such as actions made on a users behalf with the user not having gone through the access devices. This case is suspect from either a violation of policy (tailgating through the access device) of a true masquerade.
- Facial and other forms of biometric information are clearly stronger indicators than password and password –token combinations.
- Using physics with range of motion of user could be beneficial in plotting movements and seeing if the movements disagree with access requests.

- Need to also look for non-matches of users entering premises and then never heard from in terms of any of the other devices. Situation, particularly in large physical areas or for protection or guarding of important assets could be significant.
- All anomalies must be treated as potential events
- Possible extensions include using a malicious code detector as another input to the sensor. This case specifically addresses the Omega Engineering situation in that a detection of malicious code could potentially be traced to a specific user providing not only evidence, but also could invoke a course of action to block execution of malicious code before any damage done.
- The time synchronized and normalized event data can be combined into logical steps correlating to progressive stages of an attack.
- During the fusion process, view the process as correlation of associated events and discrimination between like events.
- Provide a creation of and formulation of beliefs based on partial evidence relying only on what we've seen.
- In some cases, asking other sensors to provide verification, validation, and refinement of belief if secondary and timely (in regards to the half life) information available.
- Each correlation and fusion point provides additional knowledge to gain affirmation or manifestation of suspicions.
- States of the sensor outputs need to be considered. At any time, sensor information can be affected due to a variety of reasons, but the concern is due to sensor information being in one of four states. The first state is 'true positive', where the information is true, the reading by the sensor is a good with the combination indicating a true reading positive response. The second state is 'true negative' where the sensor indicates the credentials presented fail the test. The third step is 'false positive' where the sensor returns a positive result when the credentials are not bona fide. The final state is 'false negative' when the sensor indicates the credentials are faulty when they are actually bona fide.
- Setting and correcting the sensor tolerance limits through feedback loops is important. The entire fusion suite would contain a mechanism that programmatically provides performance "feedback" to these devices in order to reduce false positives. This would thereby enhance both the fidelity and the value of the COTS security component. Most COTS products have parameters and tables set at product installation time that are initialized by the network administrator and are based on site and user specific parameters. By considering adjustments based on other environmental factors, the fusion process can help refine those parameters.

Developing testing information for IAMPS is a process of considering all the various sensor states and then matching possible inputs together to view the impact on the fusion engine. The methodology for selecting the elements of the test matrix follows the concepts of minimal essential and is based more on characteristics of the information rather than on the value of the information itself. The time relevance of the information is of extreme importance to the fusion process because in some cases it is the time deltas that provide significant clues. Looking at the hierarchy of timing that describes any piece of sensor data or information in the system, there can be established a precedence in terms of timing computations between data elements of the system.

There are several points in selecting the parameters of the test matrix that need be taken into account:

- Examine the data for time of creation. It is important to understand when the data is initially created -- in terms of compile time, start-up/initialization time, etc.
- Data rate in terms of latency between the times the sensor acquires the information to the time the information is available to the IAMPS process.
- Data from the sensor is also examined for specifics on the refresh rate of that information. Some information has longer periods of time between updates. Knowledge of the refresh intervals becomes important when the period of validity of that information is considered.
- Data from different sensors arrives at random intervals.

The table below describes a partial-testing scenario. Developing a scenario, while outside the scope of this effort, was considered important in understanding the overall interplay of the fusion information.

Table 5
Sample Testing Scenario Generation

Timecard System	Premises Entry Control System	User Authentication and System Identification	System Behavior Profile	System Activity	Assessment
Entry for today	Users entry recorded	User logged in. System is located within the building	Normal level of activity	No unauthorized activity	Authorized User activity
No Entry	No entry recorded	User logged in. System is located within the building	Unusual level of activity	Unauthorized accesses recorded	Probable intruder masquerading as the user
Entry for today	User Entry recorded into the building	User logged in. System is located within this building	Unusual level of activity	No unauthorized activity	Possible intruder masquerading as the user (may require further investigation)
Entry for today	User Entry recorded into the building	User logged in. System is located within this building	Unusual level of activity	Unauthorized activity detected	Malicious activity detected by authorized user
*	Entry recorded at another site	User logged in. System is located within this building	*	*	Probable intruder. The user cannot be at one site and logged in at another.

From the outset, IAMPS as a process needed to collect information, time stamp that information and then maintain correlation results. This is in essence the type of information forensic investigators look for to uncover the perpetrator of the cyber event [Ste99], [Mar01]. IAMPS makes use of some related recovery technology [Mcc01], [Mcc01a] to use proven approaches for guarding and maintaining information in the face of a cyber event. What was specifically demonstrated was the selection of critical minimal essential data sets (MEDS) and then retaining that MEDS information. The methods of retention included fragmentation and encryption with the recovery processes later re-assembling the decrypted information. A similar approach can be followed for IAMPS. Once a piece of evidence has been developed, all the information surrounding that evidence (evidential MEDS) can be gathered and then sent to the evidence file. The process for collecting information in the file has to be well spelled out, and become a consistent process to withstand legal tests in the entire chain of evidence life. The conclusions reached are that the previously developed approaches are satisfactory and can be transferred from recovery processing to evidence gathering.

A simplified approach to this is as follows:

- Executes at pre-determined intervals to catch all evidence gathering occurring between major updating. Can also execute when significant evidence is found to ensure retention
- Ensure that all time information relative to the evidence has been accounted for and retained.
- Perform the evidential MEDS selection and parameter updating followed by the compression, encryption and storage.
- Perform log entry of evidence update

At the other side of the processing to reconstitute the evidence:

- Invoke special routines to access and de-encrypt, decompress evidential MEDS
- Correlate with log.
- Save information using secure electronic filing for report generation and off-line analysis.

6. RESULTS AND DISCUSSION

This section of the report presents an assessment of the IAMPS concept by looking at the efficacy of the solution. This includes an analysis of the technical and other implementation issues. The section concludes with a methodology for testing that includes a mechanism for modeling IAMPS.

In section 4 of this report, the evaluation criteria were applied to the alternate solutions to show the strong and weak points. To be consistent with that approach, the criteria were also applied to the IAMPS solution, there results are presented in table 6. The IAMPS solution provides for some additional capabilities and features that can be implemented. First is the ability to dynamically adjust tolerances. This is crucial to having IAMPS be adaptable to surroundings and environmental conditions. Having a feedback loop to automatically adjust the COTS tolerances allows IAMPS to operate under a variety of stress conditions. Another key feature is the ability to select courses of action that are directly a function of the presented credentials. Besides blocking accesses to physical areas or to computer resources, IAMPS can include a mechanism to contact proper authorities if the event warrants (for example, an authentication inconsistency results when a user presents credentials to access an area that contains ammunition or toxic chemical).

Table 6
Summary IAMPS Solution Evaluation

Criteria	Comments and Evaluation
Solution operates in real time	Yes, each piece of information related to the user in the system is considered to be part of the authentication process. As the user invokes more access, the system begins to correlate access attempts and can make decisions as events unfold.
Minimal operator interference	Yes, but can involve operator if need be. Pre-selected courses of action can be loose or tight in terms of denying or restricting access to resources based on results of authentication fusion. Depending on the critical nature of the asset being protected, adjustments in the tolerance can be made.
Maintain forensic evidence	Yes. IAMPS collects information in real time, time stamps, and logs processing. Using approaches demonstrated on recovery programs, IAMPS could maintain this information for later use. (For example, several USAF programs have developed real time recovery and reconstitution of information [Mcc01], [Mcc01a] with a demonstration of that technology available. The programs first postulated the use of steganography and then implemented an approach using steganography.)
Implements courses of action	Yes. Comment here is similar to above, where tolerance limits can be adjusted to have courses of action set as a result of other criteria (Threat levels, cyber interference, etc) and then alter the courses of action accordingly.
Wide range of sensors	Yes, but sensors that do not aid in the identification process of a user would not necessarily be part of the selection.
Technology transferability and portability	Yes. IAMPS architecture is open with the implementation being most likely a separate box integrated into the overall system and some specific software within each node of the system.
Identifies any masquerading user(s) in spite of other system authentication	Yes. IAMPS continually evaluates sensor and other data to look for inconsistencies in a user's authentication information.

In essence, IAMPS is a cutting edge solution to the insider problem. It is straightforward in the sense that it concentrates on the authentication portion of the user access process and continually adds new information to the aggregate and re-evaluates the new aggregate looking for profile consistencies and more importantly profile inconsistencies. IAMPS looks for the inconsistencies during the authentication process to be in the position of denying accesses as opposed to reacting to damage after the damage has been inflicted.

It is a well know adage that good security is a balance between the people affected and practicing it, the security processes those people develop and adhere to, and the technologies selected and implemented to enforce and facilitate the security. Achieving that balance is difficult and sometimes changes in technology ripple through and affect previously good processes which in turn can affect the operability of the people.

The IAMPS solution maintains a good balance between the people, the processes and the technology. As a technology, IAMPS provides a new set of algorithms operating on primarily existing data from sensors already in a system or well within the bounds of a system. For example, moving from a physical access that is a manual cypher lock to one that is an electronic card reader is not a large jump. Using commercially available sensor components as the foundation of IAMPS keeps the impacts of technology changes minimal, and thus maintain a balance in the people-process-technology triad. One new sensor was postulated as another identification sensor and this is the computer environment analysis sensor. This sensor would notice fluctuations in various hardware parameters and use expert reasoning to compare known user profiles against current measurements, thus providing an additional measure on a user. The implementation of this sensor could be on a chip/board resident in each system or more likely through an additional box on the network, similar to a firewall or security monitor. However, this is a non-intrusive sensor in that it does not interfere with or impact the user.

The implementation of IAMPS involves delineating a set of requirements and addressing some of the possible approaches to implementation. The requirements do not detail *how* to implement IAMPS, rather they specify what *features* the IAMPS implementation should have. The approaches to the implementation give guidance as to what options should be considered in the actual development.

The following list is designed to be a living list particularly through any follow on implementations and tests. It is the intent that the list be iterated and validated through measured testing and analysis.

- IAMPS shall be designed to ensure no loss of completeness, consistency, and meaning of the information in the system.
- IAMPS shall be transparent to the system operators.
- IAMPS shall not deny access to authorized users.
- IAMPS shall deny access to users whose presented qualifications do not correlate
- IAMPS shall not be accessible to unauthorized personnel.
- IAMPS shall be unable to be modified except by authorized and cleared personnel.
- IAMPS shall be tested prior to release using a standardized testing suite of conditions and responses.
- IAMPS shall be periodically re-tested to ensure the operability.
- An interface specification shall be developed between the cyber sensor components of the defensive suite to ensure the IAMPS receives sufficient and timely information on the attack.
- IAMPS shall develop a forensics mode based on steganographic technology
- IAMPS shall support unique and positive authentication of authorized users.
- IAMPS shall detect, log, analyze and report unauthorized access attempts.
- IAMPS shall have a policy for review of COTS sensor technology for security software to maintain the latest in COTS tools and procedures.
- IAMPS shall have the ability to disconnect from sensors or interfacing systems that are under IW attack.

- IAMPS will be developed using open architecture to guarantee portability

Implementing IAMPS allows for some latitude particularly in the availability of technologies and the ability to effectively implement technologies. It is expected that as technologies improve and different sensors are made available, the fidelity of the fusion will increase. The list that follows represents only a short-term view since it is expected that future work would be more immediate than long term, so relevant approaches can be recommended. The set of recommendations in the near term is for IAMPS implementations:

- Should involve the use of agents
- Should be based around and compatible with the major OS (Windows, Linux, UNIX)
- Be alpha tested in a real environment
- Use a separate server for all software and databases
- Fix the IP addresses so they can be correlated to specific physical locations

In a full operational implementation, there are still some non-technical issues that need to be addressed. The two most important ones are for human resources and legal. IAMPS is a monitor, and continuously monitors and provides some computations on any access attempt. From the legal standpoint, IAMPS should not be considered a toolset that profiles or isolates specific people. IAMPS does not process data that could be considered private, and only resides a persons computer as a 'cookie' (which is not illegal). Rather it treats everyone equally, essentially believing nothing until information is proved to the contrary. Some attention should be paid to legal issues and it is recommended that the implementation be discussed to evaluate any legal boundaries that might. Coincident with the legal issues are those that are from the human resources standpoint. Again, IAMPS should be reviewed from the users perspective with human resources (and possibly legal) personnel to ensure any final implementation could not be transitioned.

7. CONCLUSIONS AND SUGGESTIONS FOR FUTURE WORK

The baseline problem we set out to solve is one of ensuring that a user is who they appear to be, and that fact could be substantiated through a process of information fusion and extraction. This research went through a series of steps looking at alternative solutions and did uncover some related fusion approaches. Those were looked at in detail and it was found that they did not specifically address the authentication problem. Rather their focus was on correlating security events to look for contradictions in security policy and those events all assumed the authentication process was accurate.

We then evaluated some sensors to see what types of information could be extracted from them. This led to an analysis of true-false positive-negative data and how the declaration portion of the fusion process

This has been lacking as a concept because a great deal of trust has been placed in the something the user knows, something the user is, something the user has concept. The optimal solution is one that solves the key issues of insider or intrusion detection:

- The attack can be detected in real time as opposed to postmortem,
- The process to determine the attack is non-intrusive,
- The process returns a non-repudiated determination of an attack as it is going on,
- The process uses a minimal amount of information to determine the identification anomaly,
- The process can provide current situation reports instead of after action battle damage assessments.

The IAMPS program by all accounts was a success in proving fusion concepts as applied to the problem of non-repudiated authentication. There are five specific program concepts that could be explored to further the research. In each case, the emphasis is being put on testing in a real environment to generate meaningful results showing a clear path to transitioning the technology. In each case, there will be measurements against the following criteria:

- Achieve a false positive rate of less than .1% false alarm rate per distributed system.
- Bound the solution's computational overhead to no more than 2% of the CPU cycles.
- Detect 90% of true intrusions within 1 second of an authentication request and 99% of true intrusions within 3 seconds.
- A standardized test set that looks at sensors in normal operation, being inoperable, and giving versions of true-false positive-negative information to maintain consistency in the evaluation. This test set could be expanded to include other security events that involve other types of sensors (such as firewalls, etc.). This set could be developed from other recognized test sets, but the intent is to consider all types of sensor responses that could be related to various user states.

The potential follow-on and extension activities are as follows:

Program #1. The Implementation and Testing of IAMPS in an alpha test environment. – This activity would develop the IAMPS concepts and then implement in a real environment. The

approach will employ rapid, focused experiments to determine/refine the necessary sensors, sensor combinations, and intent models while simultaneously working to improve the technical quality and breadth of each of these component technologies. The data gathering for the cyber sensors and verification of the declaration process will occur in a real environment, such as one of the buildings at the Northrop Grumman Information Technology (NGIT) Sector. The real environment will be at a (NGIT) commercial location, thereby providing significant feedback for the user intent and cyber sensor threshold selection. The feedback and experiments run within both environment will be crucial factors in assessing the technologies, singularly and in combination. Part of this assessment will point to new types of cyber sensors that are needed to achieve the required authentication and detection objectives.

There are two variants of this program. The first would be to develop the computer environment profiler as an additional sensor and then integrate several other commercial sensors as part of the IAMPS suite. Testing in a real environment, as offered by NGIT, would ensue with measurements taken and refinements to the fusion process accomplished. The second variant would not develop the computer environment sensor but integrate other forms of user identification sensors. This program would run the standardized test set and then does an analysis of potential improvement that would result in a subsequent implementation of the environment sensor. This variant allows for incremental funding with measurement points in between to monitor progress and technical achievement.

An extension to either variant would be to implement the IAMPS solution across a distributed network, and then possibly a network of networks to evaluate the impacts in real time, long distance authentication.

Program #2. Terrorist identification, an application of IAMPS technology with facial and image recognition. Events of September 11, 2001 caused a parallel line of applications thinking in how to apply non-repudiated authentication to the problem of terrorist identification. The result is the description of a program that using facial and image source recognition to identify known terrorists. First is to populate the facial recognition system with only those images of wanted or suspected terrorists. In areas where photo identification is required, correlation can be made with the several forms of image identification that is available. First is the image as taken by the visual system. Second is the image from the photo id, and third are the images of terrorists where correlation is being sought. This system would seek to get a correlation between the visual or photo with the database. If there is a match, then a positive identification could be made. Extending this somewhat is going back to the source issuer of the photo identification (license or passport) and comparing the name / image pair on the presented identification with the name / image pair that the source issuer has in that system. If there is a non-match here, then the user is trying to use false identification and that would warrant further investigation. There is a significant challenge to doing this in real time since there are examples of dispersed, extended networks that operate almost around the world. This implies a significant number of long distance communications might need to be accomplished, potentially stretching the real time constraints. A variant of

this would be to compare the images in a source issuer's database to look for multiple images against the same name or multiple names against the same images.

Program #3. Comparative testing for non-repudiated authentication. – Using the standardized testing set, this program would look at comparing the results of an IAMPS implementation with the alternate solutions discussed in section 4 of this report (Intellitactics and e-Security).

Program #4. Integration of policy and IAMPS tools with measured testing. – Some of the types of security monitors look only for violations in the security policy. A broadening of that base would include the authentication portions of IAMPS with the policy verification portions of other commercially available tools. The standard test case would then be used to measure the integrated set.

REFERENCES

- [And00] Anderson, R., Bozek, T., Longstaff, T., Meitzler, W., Skroch, M., and Van Wyk, K., "Research on Mitigating the Insider Threat to Information Systems - #2: *Proceedings of a Workshop held August 2000*", Arlington, Va.
- [Car95] Carback, R. "Reducing Manpower intensive tasks through automation of security technologies" IEEE Annual International Carnahan Conference on Security Technology, Proceedings 1995, pp.331-339.
- [Cas01] Casey, E. (2001), *Handbook of Computer Crime Investigation*, Academic Press, Bath, Great Britain
- [Cla94] Clarke, R. "Human Identification in Information Systems: Management Challenges and Public Policy Issues" *Information & People*, vol.7, no.4 (December 1994) pp 6-37.
- [Ehr66] Ehrlich, T "Passports" *Stanford L. Rev.*, v.19, pp.129--149, 1966-67.
- [Fox06] Fox-Davies A.C. and Carlyon-Britton P.W.P. "A treatise on the law concerning names and changes of name", Elliot Stock, London 1906.
- [Kru01] Kruse, W., Heiser J.(2001), *Computer Forensics Incident Response Essentials* , Addison-Wesley, Boston, Mass.
- [Lin99] Lindqvist, U. and Porras, P., "Detecting Computer and Network Misuse Through the Production Based Expert System Toolset (P-BEST), Proceedings of the 1999 IEEE Symposium on Security and Privacy, Oakland, C., May 9 – 12, 1999
- [Mcc01] McCallam D., Piggott C., Newland R., *Rapid Recovery of Information for Real Time Intruded Systems*, AFRL-IF-RS-TR-2001-55, Rome, New York
- [Mcc01a] McCallam D., Newland R., Jajodia, S., Wheeler, A. *Demonstrating Information Resiliency*, AFRL-IF-RS-TR-2001-156, Rome, New York
- [Mar01] Marcella, A and Greenfield, R (2001), *Cyber Forensics: A Field Manual for Collecting, Examining, and Preserving Evidence of Computer Crimes*, CRC Press LLC, Boca Raton, Fl.
- [Neu99] Nuemann, P., "The Challenges of Insider Misuse", Prepared for Workshop on Research and Development Initiatives Focused on Preventing, detecting and Responding to Insider Misuse of Critical Defense Information Systems, Santa Monica, Ca., 16-18 August 1999

[Por99] Porras R. and Nuemann, P., "Experience with EMERALD to DATE", 1st USENIX Workshop on Intrusion Detection and Network Monitoring Santa Clara, California, 11-12 April 1999, pages 73--80

[Res99] "Research and Development Initiatives Focused on Preventing, detecting and Responding to Insider Misuse of Critical Defense Information Systems *Results of a Three Day Workshop*, Santa Monica, Ca., 16-18 August 1999

[Sha00] Upadhyaya, S., Chinchani R., and Kwiat, K. (2000), "A Comprehensive Reasoning Framework for Information Survivability", *Submitted Paper to the 2001 IEEE Man Systems and Cybernetics Information Assurance Workshop*, Baltimore, Md, 5-6 June,

[Son01] Song, D., Wagner, D., and Tian, X., "Timing Analysis of Keystrokes and Timing Attacks on SSH", *Presentation to the 10th Usenix Security Symposium*, Washington D.C., August 2001

[Ste99] Stephanson, P. (1999) *Investigating Computer Related Crime*. CRC Press LLC, Boca Raton, FL.

[Woo99], Wood, H.M. "The use of passwords for controlled access to computer resources" National Bureau of Standards Special Publication 500-9, US Dept. of Commerce/NBS

[Wod99] Wood, B. "An Insider Threat Model for Adversary Simulation", Appendix B from the Workshop on Research and Development Initiatives Focused on Preventing, detecting and Responding to Insider Misuse of Critical Defense Information Systems, Santa Monica, Ca., 16-18 August 1999

GLOSSARY

Authentication consistency: Each piece of information in the authentication stream supports the premis that the credentials presented by the user indicate the user is who they claim to be.

Authentication inconsistency: Each piece of information in the authentication stream supports the premis that the credentials presented by the user indicate the user is NOT who they claim to be.

Bayesian Networks: a joint probability distribution structure for representing knowledge about uncertain variables and computing impact of evidence on beliefs

Biometrics: automated methods of authenticating or verifying an individual based upon a physical or behavioral characteristic.

Biometric sensor: A device for processing and authenticating biometric information

Computer forensics: Science of using electronic evidence to reconstruct a computer crime

COTS: Commercial Off-The-Shelf

Cyber evidence: Electronic evidence left as a result of a cyber event

DARPA: Defense Advanced Research Projects Agency

Dempster-Schaeffer: A technique of evidential reasoning

EMAIL: Electronic mail

Evidential MEDS: Minimum essential set of evidence used to support an authentication

IAMPS: Insider Anomaly Measurement Processing System

IDS: Intrusion Detection System

IP: Internet Protocol

MEDS: Minimal Essential Data Set

NGIT: Northrop Grumman Information Technology

Non-repudiated authentication: being able to prove two points: A person is who they claim to be and we can support that claim empirically through electronic evidence and a person is not who they claim to be and we can discover the discrepancies using electronic evidence.

NSM: Network Security Manager, a product by Intellitactics

OS: Operating System

Post-mortem analysis: The analysis done on a computer system after the event is over and discovered. Operated only on data residuals.

RAM: Random Access memory

Sensor fusion: The combining of inputs from several sources to gain a more comprehensive idea of the situation in the world.