Who Really Did It? Controlling Malicious Insiders by Merging Biometric Behavior With Detection and Automated Responses

Dr. Bruce Gabrielson
Booz Allen Hamilton
Gabrielson_bruce@bah.com

Abstract

This ongoing research and development activity addresses aspects of a potential capability to detect credential misuse and a suggested alerting approach based on known attack conditions to support automated mitigation techniques. This research is based on the assumption that the audit data and human-computer activity characteristics extracted from networked components contain the footprint(s) of those trying to breach network security. It takes advantage of the combination of near-real-time suspicious activity detection with biometric behavior profiling to reduce profiling false positives and network access controls that enable faster and more focused responses to detected suspicious activities.

1. Introduction

During the past few years, increasing emphasis has been placed by vendor and research communities on triggers or precursor events that are indicators of insider threat activities. A trigger is the initial indicator that something out of the ordinary has taken place, such as the detection of a prohibited activity, an anomaly, or some heuristic detected difference. Today, event triggers need to evolve in more efficient and sophisticated ways to detect the subtle differences between the actions of a trusted individual and those of someone using the same credentials—because of the sophistication of hackers whose primary goal has been to acquire trusted credentials.

This paper focuses on narrowing the technology gap analyst's face in identifying and differentiating the trusted individual from an external hacker who is performing possibly malicious activities using stolen trusted credentials. It also proposes a path toward control of insider threat activity and continuous monitoring or mitigation of the actions of those who abuse credentials that will be more efficient than current manual approaches, which are expensive and time-consuming, and require significant resources. The advantage of this approach will be the near-real-time detection and automated control of malicious activities, which will reduce the amount of damage the insider can accomplish. The primary drawback of the approach proposed will be the risk that false positives may overburden local enclave resources. However, initial bandwidth research indicates that at a 1000 events/day/node detection rate, the impact would be 0.02% of 100MBPS or 0.002% of 1 GPS with our approach. Therefore, minimal impact on bandwidth is expected [1].

Ongoing and future research steps include:

1. Enhance the current capabilities of two existing prototype tools to incorporate remote biometric activity identification and malicious remote login identification.
2. Perform integration of the relevant aspects of the tools so detection can take place.
3. Investigate the use of an identified data standard framework for describing and alerting discoverable malicious events associated with attacks in an Extensible Markup Language (XML). The malicious events have been previously researched and documented by the Department of Defense (DoD) [2].
4. Based on the machine readable alerts, integrate a network access control (NAC) capability at the local enclave level to move platforms into more limited environments based on mission roles.
5. Use the emerging National Institute of Standards and Technology (NIST) Enterprise Remediation Automation Framework to develop standardized Common Remediation Enumeration (CRE) actions that must be taken to accomplish the distinct remediation required for each identified attack.
6. After remediation (or other actions) use the NAC capability to automatically move the platform back to the operational network.

2. The Trusted Credential Problem

What actions do malicious insiders perform? They have wide latitude from data exfiltration to impacting network operations. Because one of their major activities is stealing information in a stealthy low-profile manner, the installation of additional network defense controls would be most useful. Below are a few of the common activities a malicious insider or outsider could perform. Note that there are instances in which a malicious user hopes to accomplish his or her objectives regardless of whether those actions cause internal alerts; however, in most cases, malicious users hope to keep their activities undiscovered as long as possible by performing more "stealthy" activities. If the attacker chooses to perform an attack and remove targeted data quickly with an overt, easily detectable action, this could be considered a "less than stealthy" attack. Attempts to hide their actions would be performed as a short term delay so objectives could be completed. This might be the case of an attack launched from a stepping stone where the actual attacker has already protected his or her activities from trace back.

- The malicious user steals a user's trusted credentials, connects as the trusted insider, and accesses and exfiltrates that individual's information before removing all trace of his or her access and leaving. This would be a more covert attack.
- The malicious user steals a user's trusted credentials, connects as the trusted insider, and attempts to access a different user's workstation from which to carry out malicious activities. This may be a more overt attack.
- The malicious user actively takes control of a user's workstation, uses that user's access to collect and exfiltrate information, and removes all trace of his or her access and leaving. This is likely a more overt attack.
- The malicious user actively takes control of a user's workstation; expands the user's privileges to access, collect, and exfiltrate information; and removes all trace of his or her access and leaving. This is also likely a more overt attack.

Based on workshops and market research by the Department of Defense (DoD) [3][4][5][6], most recent malicious user threat research has focused on developing special-purpose collection tools, identifying relevant data to collect, incorporating common naming standards and formats, creating collection repositories, integrating multiple sensor data and collection components, looking for anomalies, and using other security and analytical monitoring tools that require labor intensive close investigative coordination. Additionally, a literature review concerning insider threat taxonomies [7] and compromised credentials also shows that existing taxonomies do not sufficiently address the class of attack or type of activity addressed in this work. That absence also suggests that there are insufficient tools for detecting or mitigating these types of insider threats.

Traditional internal detection tools model normal behavior and then look for deviations or focus on outsider activities aimed at internal systems. Current techniques for modeling the actions of legitimate users to watch for changes in their behavior over time (dynamic deltas) produce many false positives. Additionally, continuous monitoring solutions often require significant resources. Trained analysts are required to use multiple sets of detection tools that may be bandwidth intensive and may also require archive capacity.

The DoD, counterintelligence (CI), and law enforcement (LE) communities recognize the importance of attributing suspicious or malicious activity to the actual individual performing the activity. This type of analysis and use of audit data overall has been impeded by:

- The very large volumes of heterogeneously formatted text data
- The lack of tools and techniques for analysts to explore audit data in a readily intelligible format
- The lack of common terminology for similar audit events on different systems, resulting in the inability to associate, aggregate, and correlate such events ("audit normalization")
- The ability to configure auditing systems to most effectively identify unauthorized activity at the lowest cost in terms of the number of audit records and computational resources used
- The difficulty involved for the analyst to distinguish between the user performing the potentially malicious act and the trusted user whose credentials have been compromised
- The difficulty in distinguishing a user that is just doing his or her job but may also be doing something that resembles an attack
- The limited constraints on a trusted insider or outsider who gains trusted access to network platforms due to the absence of deployed

detection capabilities, particularly at the root level.

To reduce the information assurance (IA) computer network defense (CND) detection difficulties and the investigative loads imposed on analysts, detection capabilities are being developed that will not only reduce the amount of audit data required for detection and forensics analysis but also enable correlation of identifiable attributes of the individual performing the malicious activity, if known, in near real time. These alerting capabilities are referred to by the National Insider Threat Working Group (NISTWG) as "triggers". Near-real-time alerting could involve email, automatic phone calls, or even machine-readable xml code.

3. Available Detection Technologies

This section provides an overview of existing technologies that could be integrated and enhanced to develop a more robust detection and response capability. More specific details of the particular capabilities selected for the integration effort are described later.

3.1. Network Access Control

Automated network access control (NAC) reduces vulnerability problems by allowing for conditional network access, host isolation into specialized networks, and captive remediation. NAC host authentication supports fast, direct communication with network administration when policy violations are detected as well as automated placement of a user's system into a role with network access appropriate to its owners [8].

3.2. Host Policy Enforcement

For host detection and remediation, a specialized client-installed software agent (see also host-based intrusion detection system/host-based intrusion prevention system [HIDS/HIPS] below) automatically runs checks and sends the results back to a central location. The central repository compares the results against the expected policy and can reconfigure if a change is detected and indicated based on response policy. This approach allows for highly granular checks, including anti-virus definition dates and the presence of certain registry settings, which external scans cannot detect. Unfortunately, the approach has a scalability issue, making version control difficult to manage. For continuous policy enforcement, hosts are constantly scanned, thus creating increased network traffic. In addition there are other concerns including the cost associated with maintenance and licensing for large enterprises.

The author makes an assumption here that automated host policy enforcement would be better suited for the Non-classified Internet Protocol Router Network (NIPRNET) than for the Secret Internet Protocol Router Network (SIPRNET). SIPRNET is not always the operational network of choice. However, there are generally immediate operational needs required on SIPRNET. In particular, mission directed operational needs have priority over automated actions that may impact those needs. Therefore, further investigation of the potential for limiting NAC on an operational SIPRNET platform is needed.

3.2.1. Host-based Intrusion Detection System/Intrusion Prevention System (HIDS/ HIPS). A HIDS resembles a network intrusion detection system (IDS) with the advantage of working close or within the operating system protected kernel layer to detect root-level attacks that might be missed by network-based IDS. An agent based HIDS identifies intrusions by analyzing system calls, application logs, file-system modifications (binaries, password files, capability databases, access control lists, etc.). and other host activities and states. Unrecognized incoming malware that causes a configuration change is recorded and reported. As with IDSs, HIDSs are prone to false positives and require continuous signature and white listing updates.

HIPS combines a signature and anomaly-based IDS with a packet filtering firewall plus the capability to intercept and examine specific system calls and application programming interface (API) calls, which are used by applications to request services from the operating system. The HIPS can stop the attack itself by changing the security environment (e.g., reconfiguring a firewall), or changing the attack's content. HIPS can also take defensive actions such as blocking intruder addresses without causing a denial of service possibility on the network. Unfortunately, the HIPS requires constant management updating of its white list because of continuously changing trusted software applications.

3.3. Host Audit Detection and Correlation

As with policy detection agents, collection agents can collect and send audit information back to a central location for correlation and event detection. In some cases, an installed agent can immediately detect a known pre-defined common security policy violation and send an alarm or activity indication to the appropriate authority for response actions. What is considered a potential pre-defined common security violation is addressed later in this paper (4.3). The majority of these commercial capabilities require queries against large external repositories of collected audit data. Unfortunately, when large-scale deployments are involved, there are issues that may be problematic such as agent tuning and network management to create profiles, reduce false positives, and detect policy violations, plus the added bandwidth impact from audit data.

3.4. Network Behavior Profilers

Network behavior profilers passively collect statistical information from deployed sensors and perform continuous real-time monitoring of network behavior to better understand network usage, discover assets, identify malfunctioning devices, and detect activity trends. From this information, they create behavior profiles of hosts and expected normal network activity.

3.5. Host Biometric Behavior Profilers

Biometrics is not a new technology, but evolving biometric techniques provide for a more effective security model of constant verification based on physical activity. Keystroke biometric analysis that addresses the problem of user authentication, verification, and identification (attribution) has made significant progress in recent years [9][10]. Unfortunately, these profilers still have problems with false positive alerting. However, the integration of keystroke biometrics could transform security models from ones of "authenticate and trust" to "constant vigilance." Employing the proposed biometrics introduces a transparent method of constant security, which would enable much more effective discovery and attribution of unauthorized access. This is possible because of what is actually being used to identify the user. Keystroke biometrics analyzes the user's own actions as the basis for measurement. The fact that he or she uses the system gives the technology the data it needs to enforce security.

4. Integration of Technologies

To reduce the IA CND detection needs and the investigative loads for analysts, new approaches must be developed for near-real-time detection that will reduce the amount of audit data required for detection and enable the identification of the individual performing the malicious activity, if known, in near real time. The new approaches also will create the machine-readable alerting information that can direct an NAC-like response and enable response actions based on directing alerts to the proper response organization. Attributing suspicious or malicious activity to the actual individual performing the activity as quickly as possible will enable a response to be initiated, before the attacker has covered his or her tracks.

The approach taxonomy applied herein is to use static detection of potentially malicious or suspicious activity at the platform level while concurrently determining within an acceptable margin that the actor is either unknown, known but not the individual whose credentials are provided, or known and the same user performing the activity. Knowing any of these three metrics can enable either automated mitigation or directed action by an analyst to take place as indicated in Table 1. Note that remediation action or mitigation activities are driven both by risk and by local CI policy.

Table 1. Decision points

Security Event	Decision Point	Action/Mitigation
Policy Violation (Severe)	Automated	Quarantine
Policy Violation (Moderate)	CND/CI Alert	Per Policy
Outside Attacker (Unknown)	CND/LE Alert	Per Policy
Inside Attacker (Unknown)	CND/CI Alert	Per Policy
Inside Attacker (Known)	CI Alert	Per Policy

4.1. Near-Real-Time Activity Detection

Log and log-like analysis of user activities can be used to detect potential malicious sequences. For this to happen, the installed agent's reader must separate, parse, and then extract the desired information from each audit record or indicator location. Because

different operating systems define this data in differing ways, the extracted data must be mapped into named, normalized data elements. This ensures that collected data is normalized and that alerts are machine readable regardless of their source.

A sequence-of-events detection capability that enables specific sequences to be detected immediately with or without the transport of raw event data to central correlator is depicted in the block diagram in Figure 1.

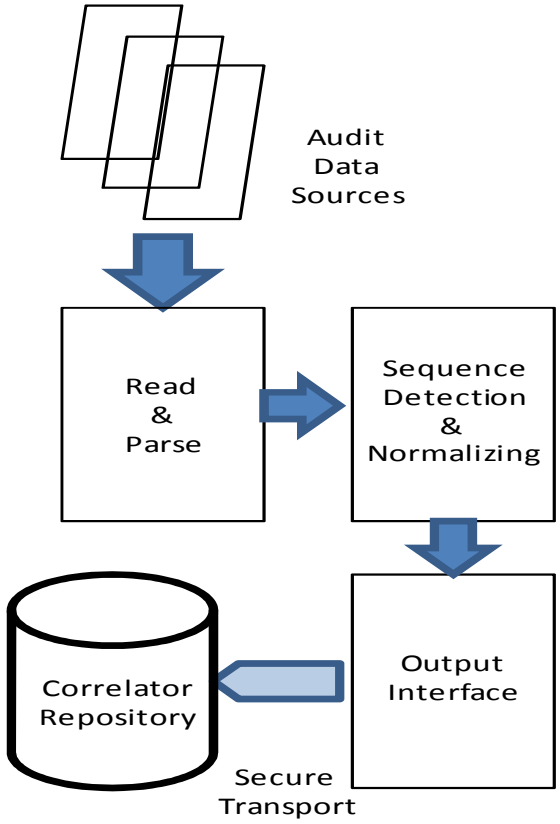

Figure 1. Existing Audit Data Extraction Utility (ADEU) sequence detection block diagram

4.2. Biometric Profiling Analysis

Keystroke biometric analysis is one technique for biometric profiling. Some other techniques include mouse movements and how applications are used. Keystroke analysis consists of measuring flight and dwell time, as well as key-to-key combinations during free-text keyboard activity. The analysis produces a unique signature that is dynamic, producing increasingly accurate results with time. If the credential for a known trusted user behavior profile were previously profiled and stored in an encrypted manner at a location isolated from the host and then quickly compared to the biometric profile of the individual who is using the credential, an immediate alert and quarantine could be initiated. Quarantine in this application could mean restricted access to a location with specific data sources and where additional monitoring tools are deployed rather than total removal from the network.

4.3. NAC and CAPEC

Another emerging standard is the Common Attack Pattern Enumeration Classification (CAPEC). CAPEC is a way of describing common methods for exploiting software. Our research concentrates on exploitation attacks from outside rather than malicious insider activity; however, the insider who performs a malicious act on another trusted user's workstation could use well-formed requests to an application, service, or device that results in the inadvertent disclosure of sensitive information by exploiting weaknesses in the design or configuration of the target. In this case, CAPEC, which could be considered in terms of a common policy violation, would provide the common attack activity description in a machine-readable format.

While a small subset of the CAPEC mechanisms of attack are applicable to the common insider threat attack sequences, the structure can cover most attacks. CAPEC mechanisms of attack currently address:

- Data Leakage Attacks
- Resource Depletion
- Injecting Content
- Spoofing
- Time and State Attacks
- Abuse of Functionality
- Probabilistic Techniques
- Exploitation of Authentication
- Exploitation of Privilege/Trust
- Data Structure Attacks
- Physical Security Attacks
- Network Reconnaissance
- Social Engineering Attacks
- Supply Chain Attacks.

For example, the summary of a CAPEC Data Leakage Attack from the Mitre [11] website reads:

"An attacker uses well-formed requests to an application, service, or device that results in the inadvertent disclosure of sensitive information by exploiting weaknesses in the design or configuration of the target resulting in the target revealing more information to an attacker than intended. The attacker may collect this information through a variety of methods including active querying as well as passive observation. Information may include details regarding the configuration or capabilities of the target, clues as to the timing or nature of activities, or otherwise sensitive information. Often this sort of attack is undertaken in preparation for some other type of attack, although the collection of information may be the end goal of the attacker in some cases. Information retrieved may aid the attacker in making inferences about potential weaknesses, vulnerabilities, or techniques that assist the attacker's objectives. Data leaks may come in various forms, including confidential information stored in insecure directories, or via services that provide rich error or diagnostic messages in response to normal queries."

CAPEC could be applicable in the detection and mitigation of known exfiltration attacks. In the proposed architecture, an attack indication would be mapped to a known CAPEC attack within the correlator. This would trigger the "to be developed" machine-readable xml alert for a severe CAPEC-described policy violation. For a workstation, this would trigger NAC-like quarantine on the platform, moving it to a controlled but undetectable environment where controlled monitoring can or would take place. If the detection were on a database, there could be an automated command to the database (or some other device such as a firewall) to stop the data exfiltration communications immediately.

Note that there may be many mitigation strategies for a particular attack event. The methodology for determining which strategy would be more effective with less impact on the organization is under investigation. An example would be deciding whether to remove a server from the network to do the reimaging, or to perform real-time reimaging.

If a policy violation is detected from someone who is possibly not a trusted insider, CND and LE/CI would be notified and the appropriate automated or non-automated response would be initiated. If the policy violation is detected from a trusted insider, CI can determine appropriate action based on what the user did. If there is suspected malicious activity from a user, but no actual violation has been detected, there may not be an automated response, but all three organizations may become involved, based on user actions.

Some means must be available to reconfigure and enable a compromised host to be reattached to the network. The Security Content Automation Protocol (SCAP) combines a number of open standards that are used to enumerate software flaws and configuration issues related to security. They measure systems to find vulnerabilities and offer methods to score those findings to evaluate the possible impact. It is a method for using open standards for automated vulnerability management, measurement, and policy compliance evaluation. In the proposed architecture, SCAP could be used to test and reconfigure the quarantined workstation to determine its suitability for reconnection to the network. Additional research is needed to resolve this question.

4.4. The Combined Architecture

Figure 2 illustrates a local enclave level integrated architecture with the red line indicating the NAC control initiation. This solution is best suited for the local enclave because it has direct control over its own environment and platforms. The figure assumes that latency and timing deltas between keystrokes from a user will be nearly uniform if the user is typing from the actual monitored platform or from a remote platform. The possibility of keyboard-related biometrics being applicable for remote users is discussed in Section 6.

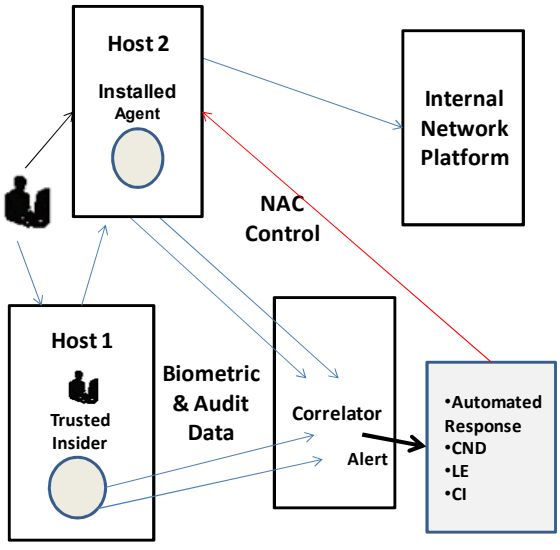

Figure 2. Integrated architecture

Figure 3 depicts the decision process inside the correlator.

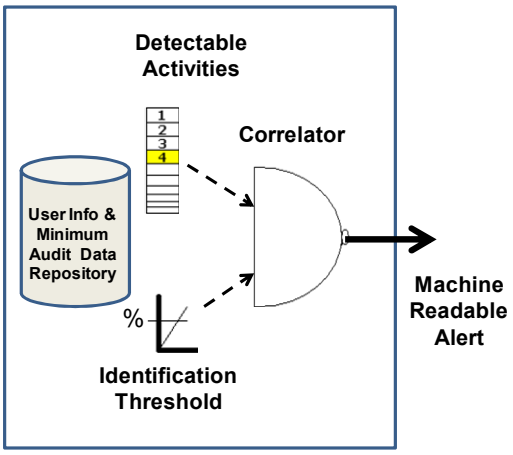

Figure 3. Decision process in correlator

5. Existing Components

Tools and techniques currently available for integration include the Digital Biometric System (DiBiS) [12][13], the Audit Data Extraction Utility (ADEU) [14], and any of the available NAC solutions, plus the CAPEC data standard for describing exfiltration attacks. For this effort, code will be developed to address automatic NAC quarantine.

DiBiS is a tool developed at the Air Force Research Lab (AFRL) that analyzes several of the unique characteristics of individuals and how they interact with computers, not simply keyboard biometric analysis. It can be used for the following:

- Biometric keyboard authentication
- Web and application patterns of activity based on application context and hardware profile data
- Fusion of multiple sensor and system output algorithms
- Interactive response mechanisms
- Anomaly detection
- Mouse movement identification (future).

The ADEU is a capability developed by DoD for extracting, aggregating, correlating, and reporting audit, log, and log-like data within a service-oriented publish/subscribe architecture. ADEU is capable of identifying potential abuse of trust or intrusion into a DoD network platform through near-real-time activity pattern matching rather than anomaly detection and profile tuning. The ADEU exists in four parts. The platform-deployed Tap agent forms the basic collection construct, while the ADEU Bridge, Correlator, and correlation database form the basic short-term storage and publication function for the audit management architecture. Although ADEU can parse, normalize, and extract all malicious activity-related log and log-like files from a network platform, only those data elements necessary to support defined use cases are extracted. ADEU can accommodate CAPEC-encoded event detection.

5.1. Trigger Points

Of importance with DiBiS is that it can produce a reasonable metric estimate of user identification within 10 seconds of monitoring a user. This assumes the user is on a workstation where a DiBiS agent is located. The issue with remote logins is discussed in Section 6 below.

DiBiS currently provides a metric of 0 to 1 with 1 being an almost 100 percent certainty of user identification. ADEU alerts from a malicious activity within a short time, normally less than 20 seconds. If the local user is vigorously typing, detection and recognition are almost immediate. If the hacker is performing a manual attack remotely, a longer period of time is needed to provide a metric related to whether the hacker is known or unknown, particularly because keyboard biometrics is less applicable. This is an uncertainty area. Which indicators would provide the best recommendations in the shortest time? At this point, a 60 percent or better estimate would allow for some reactions to be triggered.

The block diagram in Figure 4 represents a future enhanced situational awareness and digital management architecture based on the availability of standardized alerting and configuration languages and protocols for communicating between blocks. The vision would be to express, enumerate, measure, and interact with event data from heterogeneous sources to streamline event management. Common descriptions of events and data elements would be used to match signature patterns and reduce the volume of log data to be collected. The architecture also fits within the continuous monitoring environment envisioned for the DoD enterprise.

Pattern-matching alerts are sent to Tier 3 log data analysts based on the incident response plan of the detecting organization. As indicated by the communications channels of Figure 4, the protocol

identifying which events are to be provided to Tier 1, and the responsibility for review and action need further research. The intent is that only events representing enterprise-level significance would be reported to Tier 1, but all questionable activities would be reported to lower level tiers as directed.

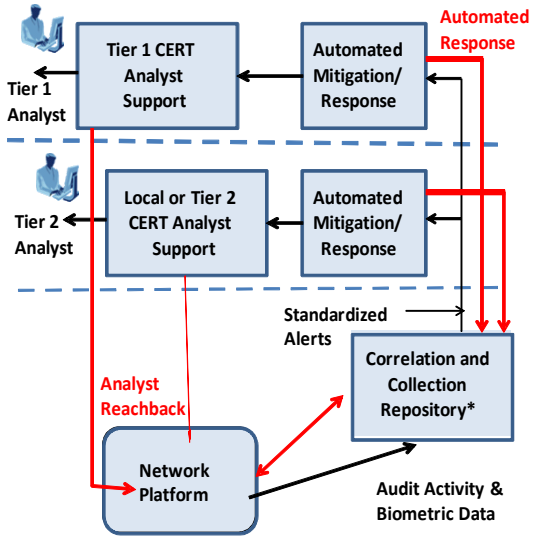

Figure 4. Block diagram transport overview

6. Areas of Additional Research

One of the current areas of interest is the type and limits of biometric detection that can be applied for a remote desktop logon when trusted credentials are used. If a remote logon is scripted and no biometric activity is detected, the user is likely a hacker and an immediate alert could be generated. Research will focus on remote logon and activity using a Transmission Control Protocol (TCP), typical of remote protocols such as Secure Shell (SSH), Telnet, Remote Desktop Protocol (RDP), and rlogin. If the login is with User Datagram Protocol (UDP), the login is connectionless and provides only best effort delivery. However, there is nothing to prevent connection management and login functionality being coded at the application layer with an underlying UDP connection.

RDP is a Microsoft protocol that provides the user with a graphical interface to another computer. Clients exist for most versions of Microsoft Windows, Linux, UNIX, Mac OS X, Android, and other operating systems. If the malicious remote user uses manual controls with RDP, some biometric indicators may also be useful.

The impact of TCP packet assembly and latency in larger networks on biometric identification is not known. Uniform latency should have minimal impact on biometric patterns, but packet assembly may be more problematic. Experiments are planned to assess these impacts, beginning with simple, small-scale networks and progressing to more complex and geographically diverse networks. In addition, a "virtual human interface adapter" will be developed to intercept keyboard, mouse, and other human interface activity from the remote protocol.

Another area of additional research is the length of time necessary to determine a high probability of identification (known or unknown) while reducing false positives. A 100 percent positive identification is not necessary when a suspicious activity is detected. If a trusted user is not performing malicious activities and his or her workstation is placed in quarantine, the user would normally contact the system administrator if he or she feels there is a connection problem. This happens so often under normal conditions that it is believed there would be very little impact on a typical user while tuning to reduce false positives takes place.

7. Conclusions

Near-real-time policy violation and malicious audit activity detection have made significant progress as viable solutions in recent years. Host-level biometric detection has likewise become an efficient means for identifying unique user attributes. The integration of both capabilities offers the potential for identifying cyber insider activities in near real time. While a CAPEC described attack is an immediate trigger, biometrics provides an alternative trigger to identify a loss of ongoing trust in a host attached to a network.

In a similar manner, developing a NAC-like capability to quarantine a suspicious workstation into an enclave with limited access until the suspicious activity is analyzed would provide another layer of protection, one that allows close monitoring and control of potentially malicious activities. Finally, with the addition of machine-readable alerts and the emergence of a digital reconfiguration capability, near-real-time detection and automated mitigation become an achievable possibility in the control of insider threats.

In summary, this evolving insider threat integration and development effort provides:

- A new capability that associates known users with their actions
- A new capability that quickly identifies a user that is unknown or using stolen credentials
- A capability that detects malicious activities before attackers can hide their actions
- A new capability that normalizes collected data to enable machine-readable alerting
- A capability that can be used to automatically trigger NACs and/or more intensive analysis and monitoring tools
- A capability that can provide DoD and government automated detection and response actions based on risk and mission needs.
- The possibility that this work could offer an extension to the Hansman and Hunt taxonomies [7], which in turn builds on previous taxonomies of threats and attacks.

8. References

[1] Wood, R. "Initial AEM Scaling Test, Revision 1," ITT Corporation, Rome, NY, 1 February 2011.

[2] Gabrielson, B., "Compiled Malicious Activity and Attack Use Cases and Examples (FOUO)," CND R&T Team, NSA I411, 29 August 2011

[3] Ferguson, R., "Enterprise-wide Information Assurance and Computer Network Defense Solutions Steering Group (ESSG), Insider Threat Technology Advisory Group (TAG), Market Survey," August 2011.

[4] Gabrielson, B., "ESSG Insider Threat Detection Technology Advisory Group Gap Analysis (FOUO), V2," STRATCOM, 6 May 2010.

[5] Gabrielson, B., "ESSG Market Survey, ADEU/AEM Key Features vs. Selected COTS Products," NSA CND R&T PMO (I71), 23 August 2010.

[6] "The Insider Threat to Information Systems, State-of-the-Art Report (FOUO)," Information Assurance Technology Analysis Center (IATAC), February 2009

[7] Hansman, S., and R. Hunt, "A Taxonomy of Network and Computer Attack Methodologies," Department of Computer Science and Software Engineering, University of Canterbury, Christchurch, New Zealand, 7 November 2003.

[8] Report: "Symantec Network Access Control," http://www.symantec.com/business/network-access-control. Accessed 4 April 2011.

[9] Ordal, P., D. Ganzhorn, D. Lu, W., Fong, and J. Norwood, "Continuous Identity Verification through Keyboard Biometrics," Department of Computer Science, University of Rochester, http://sa.rochester.edu/jur/issues/fall2005/ordal.pdf. Accessed 5 April 2011.

[10] Student Research Staff, University of Victoria, "Biometric Profiling System (Mouse and Keyboard) for Network Intrusion Detection," Biotracker Project, www.isot.ece.uvic.ca/projects/biotracker/. Accessed 4 April 2011.

[11] "Common Attack Pattern Enumeration and Classification (CAPEC)," http://capec.mitre.org/.

[12] Spink, B., "Digital Biometrics Whitepaper," 24th AF, AFRL, 27 May 2010.

[13] Fitzgibbon, J and C. Duy, "Digital Biometrics System: Middleware Solution," Booz Allen Hamilton, 15 May 2011.

[14] Gabrielson, B. and D. Wood, "Audit Data Extraction Utility (ADEU) Audit Management Concept of Operations," NSA I71, 13 November 2010.