WHITE PAPER

Report to help readers understand insider threat detection

INSIDER THREAT DETECTION

behavior profiling and anomaly isolation

PERSONAM has developed an advanced new technology that detects insider threats. We use the science of Machine Learning and Advanced Data Analytics to construct behavior profiles for the humans and machines on a computer network, and generate alerts when suspicious behavior is identified. The technology will detect an insider threat at the first sign of unusual behavior, even if the computer network itself is not the instrument of attack or exfiltration. The technology provides security where everyone is most vulnerable – inside.

personam

INTENTIONALLY B L A N K

INSIDER THREAT DETECTION

WHEN PEOPLE WITH ACCESS "TURN"

PREPARED BY PERSONAM, INC.

McLean, Virginia

September 2013

personam™

www.PersonamInc.com

On June 5, 2013 Edward Snowden took center stage in yet another high-profile case of an insider leaking US classified information. Snowden worked in a trusted position as a contractor at the National Security Agency (NSA) PRISM program[1] where he was granted sensitive "need to know" access to classified data. He underwent rigorous background investigations, signed foreboding legal agreements with harsh consequences for unauthorized disclosure (both for himself and for the security of his country), and worked in a highly monitored environment with advanced physical and cyber-security controls. None of these measures made a difference when he decided to leave the country with several laptops full of classified material and turn it over to the press in a public display of activism. The media quickly labeled Snowden the "Ultimate Insider"[2].

EDWARD Snowden, NSA Leaker ("Ultimate Insider")

Prior to Snowden that title might have been held by PFC Bradley Manning, who in a similar scenario turned over several hundred gigabytes of classified information to the activist website WikiLeaks.

Insider threats and rogue employees are by no means limited to acts of espionage or political activism. Insider threats are a growing and significant problem for private industry, with enormous losses due to internal fraud, intellectual property theft, and sabotage. In a recent survey[3], half of employees admit to taking intellectual property with them to new companies, and 40% plan to use the information in their new job. In a 2011 survey conducted by Carnegie Mellon's CERT, 46% of respondents claimed losses caused by insider attacks were more damaging than external attacks. A quick scan of recently filed cases on the Trade Secrets Institute website reveals dozens of court cases involving issues like industrial espionage and theft.

1 http://www.guardian.co.uk/world/2013/jun/09/edward-snowden-nsa-whistleblower-surveillance/

2 http://www.wired.com/threatlevel/2013/06/nsa-leaker-ultimate-insider/

3 http://www.cio.com/article/728520/Is_Stolen_IP_Walking_in_the_Door_With_New_Employees_

Numerous initiatives have been undertaken to mitigate the insider threat problem, but none have been effective at general-purpose early warning and detection judging by the continued insider attacks. The Defense Advanced Research Agency (DARPA) initiated the Cyber Insider Threat (CINDER) project in 2010, and in 2011 added the Anomaly Detection at Multiple Scales (ADAMS) project. In a separate effort, the FBI renewed its standing effort to address Insider Threat vulnerabilities following the espionage case of Robert Hanssen[1], though admits that the science of insider threat detection and deterrence is in its infancy[2]. Following the Bradley Manning incident, the White House recommended federal agencies use psychiatrists and sociologists[3] to assess workforce threats. The NSA's response to its own high-profile leak was the implementation of a two-person rule for all of its 1,000 system administrators, similar to how strategic missile crews operate[4].

While background investigations and best-practices for hiring are useful, many see the insider threat vulnerability as a cyber-security problem. However, the cyber-security industry is already fully energized by ubiquitous and ongoing cyber attacks against federal and corporate assets. Breaches by external parties are reported almost daily and carry the high cost of stolen credit cards, breaches of individual privacy, and stolen trade secrets. Advanced Persistent Threats (APTs) consume the attention of the cyber-security industry, leaving no real effort or investment directed at mitigating the Insider Threat problem.

In a week of headlines about Edward Snowden, two prestigious cyber-security venues failed to highlight insider threat detection. Conspicuously absent from the **Gartner Security and Risk Management Summit** in Washington, D.C. was any mention in the tracks, sessions, or research papers of insider threat vulnerability and/or mitigation technology.

Even at the **Suits & Spooks** conference in La Jolla, a highly-focused cyber-security event that brings together private industry with the intelligence community, there was almost no discussion of Snowden or other insider threat scenarios.

NUMEROUS PREVIOUS ATTEMPTS TO DETECT ROGUE INSIDERS HAVE FAILED

1 http://en.wikipedia.org/wiki/Robert_Hanssen

2 http://www.darkreading.com/insider-threat/5-lessons-from-the-fbi-insider-threat-pr/240149745/

3 http://www.bbc.co.uk/news/world-us-canada-12120850

4 http://www.usatoday.com/story/cybertruth/2013/06/24/nsa-two-man-rule-stop-the-next-edward-snowden/2453895/

SIEM & DLP
THE TOOLS WE HAVE

NOT THE TOOLS WE NEED

Where does that leave cyber-security technology in the fight against insider threats? When asked what tools are available to address insider threats specifically, many vendors recommend re-purposing existing technologies like SIEM (Security Information and Event Management) and DLP (Data Loss Prevention)[1]. SIEM tools were developed to address compliance and governance requirements in markets such as finance, banking and healthcare. SIEM tools aggregate and analyze network infrastructure log files then report anomalous discoveries and breaches in policy, such as a user attempting to access information they have not been granted permission to. DLP tools protect data assets themselves by tracking accesses to every asset and detecting movement of protected assets over networks. Both technologies are advanced and mature, but despite the sophistication they have serious shortcomings in their ability to mitigate insider threat scenarios: (1) they depend heavily on pre-defined rules and learned heuristics that limit detection sensitivity to known threat vectors; and (2) they are unable to discern between proper use versus malicious use of digital resources by an employee or contractor who has been granted legitimate access privileges.

1 http://www.darkreading.com/attacks-breaches/nsa-leak-ushers-in-new-era-of-the-inside/240156599

The media has used the recent leaks of NSA Surveillance Programs to sum up the current state of technological solutions to the insider threat vulnerability:

> "It's possible that somewhere within the thousands of U.S. federal government databases, information can - or perhaps will - be found that tracks Edward Snowden's activities leading up to his leak of top secret information. But it appears that government systems are inadequate to alert authorities in real time to potential leaks." – Eric Chabrow, Data Breach Today[2]

Without a robust solution on the market, the question remains: how many more Edward Snowden's are out there ready to steal or publish your sensitive information?

2 http://www.databreachtoday.eu/tools-available-to-stop-nsa-type-leaks-a-5826

If background checks and traditional cyber-security technologies don't solve the problem, what's left? There are actually a number of industry sources that agree on the answer: **Behavioral Analysis**. Unlike rules-based systems where the rules are often well-known to those intent on subverting them, behavior analysis takes advantage of the fact that insider threats behave differently than normal loyal employees. Behavioral analysis can detect indicators of *intent*, not just the act, of malicious behavior. The technology is based on advanced analytics and unsupervised machine learning techniques. It learns and profiles the behaviors of a workforce then compares every individual's behavior to their past behavior and to the behaviors of their most similar co-workers. This provides the basis for detecting abnormal or threatening behavior. Because the technique uses advanced analytics with many simultaneous feature dimensions, it is virtually impossible for a rogue insider to know how their behavioral patterns are being profiled and compared to others, so it is very difficult for anyone to hide "stealth behaviors" from such a detection system. Because the approach does not rely on pre-defined rules or signatures, it can identify new forms of rogue behavior and malicious intent. Behavior profiling is an important layer of defense in a highly connected digital environment, especially with increasingly utilized tools like DropBox (for file-sharing) and BYOD (Bring Your Own Device).

BEHAVIOR ANALYSIS and profiling

Carnegie Mellon's CERT program specifically calls out #17 in their Common Sense Guide recommending the use of **behavioral analysis** as an important tool in the detection of insider threats. At a lessons learned briefing given by the FBI following its insider threat program (where a predictive analytics approach failed spectacularly) one conclusion stated that "detection of insider threats has to use behavioral-based techniques"[1]. Gartner, the technology research firm, also recognizes the importance of behavioral analytics in the fight against insider fraud[2]. There are at least two government programs that have begun implementations of the behavioral analysis approach. They are DARPA's (Defense Advanced Research Projects Agency) ADAMS program (Anomaly Detection At Multiple Sales), and the Oak Ridge Cyber Analytics program at the Department of Energy's Oak Ridge National Laboratory.

1 http://www.darkreading.com/insider-threat/5-lessons-from-the-fbi-insider-threat-pr/240149745/

2 "The Five Layers of Fraud Prevention and Using Them to Beat Malware", Avivah Litan, Gartner Research, 21 April 2011.

FALSE POSITIVES ARE A KEY TECHNICAL CHALLENGE IN ANY ANOMALY DETECTION SYSTEM

Previous attempts to mitigate the Insider Threat or Rogue Employee vulnerability have failed to produce a robust behavioral analysis solution for threat detection, with efforts often frustrated by daunting technological hurdles. There are three primary challenges to the insider threat detection using behavioral analysis: The first is creating **accurate behavioral profiles** suitable to the needs of insider threat detection. How one defines thresholds for threat behavior (e.g. when does a new habit turn into malicious intent), and what features to consider are key components. Local area networks produce an enormous amount of raw data, including but by no means limited to Internet browsing habits, email, social media usage, file access, and work schedule patterns. Knowing which data to use, and which data not to use, is critical. The second challenge is a consequence of the first, i.e. **Big Data.** This requires a sophisticated capability in Big Data analytics that few companies have achieved. Not only is the volume and variety of data exceptionally large, but early threat detection requires data be analyzed by advanced algorithms at near real-time speeds. The third challenge, and perhaps the most written about, is the challenge of **false positives**.

False positives are a key technical challenge for any anomaly detection system that leverages unsupervised machine learning. To avoid rules-based approaches and their shortcomings, behavioral analysis detects anomalies in the background behaviors of everyday work activities. It's a process that detects "patterns of life" from data transmitted over computer networks. Unfortunately, human behavior is highly anomalous, even when it is normal. Raising alerts each time an employee visits a new website or accesses new files will quickly irritate a security administrator, rendering such technology useless. To be viable a solution must account for the high anomaly rate of everyday work life and still be sensitive to odd behavior.

BREAKTHROUGHS in technology

Two years ago, following the high-profile arrest of PFC Bradley Manning for leaking classified information, we embarked on an aggressive R&D program to invent new technology capable of detecting Bradley Manning and other "rogue high-trust employee" threats that evade existing technologies and security procedures. In the early stages of development, the three critical challenges were identified and became the focus of innovation involving new algorithms and engineering.

WikiLeaks exfiltrator, Bradley Manning, mostly stole information he was trusted and authorized to access. In many cases his access for theft was difficult to distinguish from his normal duty function. Detecting someone like Bradley Manning isn't about catching them in the act of downloading, an intervention requires a multi-dimensional profile of their behavior from which we can detect subtle differences.

PFC Bradley Manning

BEHAVIORAL MODELING 1

We achieved our first breakthrough by scrapping the conventional approach to behavioral analysis. Existing SIEM and DLP technologies use sensors to monitor, collect and collate user logins, file access and dozens of other activities across a computer network. The technologies then apply a set of rules and statistical analysis to generate alerts when a user or autonomous device attempts to perform an unsanctioned action. Some of the more advanced vendors have started to apply data analytics to assess behaviors of the data streams, such as how often a file is typically accessed and when. The analytics approach used by these products adds value above the traditional rules engine because it can detect when something like bandwidth usage suddenly increases above a baseline norm. However, such application of analytics still misses a critical element in the detection of Insider Threats. To detect an insider threat, one must define and assess the behaviors of the individual, not the behavior of a network metric like bandwidth usage or frequency of file access. It is also essential to compare a person's behavior to other people's behavior, and detect behavioral anomalies.

This screen-capture shows network users profiled by behavior. The closeness of the boxes indicates the degree of similarity in overall behavior. With this technology it is quite easy to discern humans from non-humans. When viewing human users the resulting graph will normally parallel the official organization structure (org-chart) of the business, even though the algorithms are given no clue as to the division of labor. Behaviors are sampled from a multi-dimensional feature-set, making it difficult or impossible to conceal or forge a behavior profile. Boxes shown in red are considered "suspicious".

Our first major breakthrough on this effort was the real-time extraction of observable features that represent a person's behavior patterns (i.e. their "patterns of life") well enough to distinguish individual behaviors, resilient enough to handle a dynamic work environment, and sensitive enough to recognize an anomalous behavior that could be threatening. A true test of our success came in the form of a machine-generated dendrogram of every behavior in a working office environment. The dendrogram closely resembled the company's organization chart, confirming that subtle behavior differences were being accurately detected. The dendrogram was created by computing the similarity of each individual's behavior to that of every other individual in the organization, then clustering the individuals with the most similar behaviors together. One could clearly identify groups of software developers, accountants, sales staff, and so on.

Previous unsuccessful attempts have been made to apply predictive modeling to insider threat detection. One such project was an FBI initiative described in a presentation at the 2013 RSA Conference. The FBI's predictive models, derived from supervised machine learning, proved even less accurate than random for detecting insider threats. The example highlights the shortcomings of a predictive modeling approach: generating sufficiently accurate models would likely require a different model for each protected environment and a large number of actual insider threat profiles from the environment to train the models.

BIG DATA ANALYTICS 2

Generating human behavioral profiles requires a sensor capable of collecting multiple simultaneous features in a large continuous data stream to cover all aspects of a person's normal and abnormal behavior in a work environment. Computer networks are an ideal target for such a sensor. The challenge is that even a small office produces enormous amounts of network traffic every day. Data analytics technology of just a few years ago could never have kept up with the tonnage of data.

Most Big Data technologies which can simultaneously analyze multiple feature dimensions on gigabytes of data require the data to be at rest, at least temporarily, in a data repository. Relying on map reduce frameworks and other Big Data technologies would require big infrastructure to store and triage data, making deployable behavior profiling instruments impractical. Worse, it would mean that the analytics of identifying suspicious behavior would have to be done in off-line batches. For a reliable insider threat detection capability, the behavioral profiles must be computed in near real-time on data "in flight".

To address this need, we invented several advanced technologies, including new machine learning algorithms that profile behaviors and identify suspicious activity using real-time streaming data at extremely high volumes and velocities.

HUMANS PRODUCE **A LOT** OF NOISY DATA

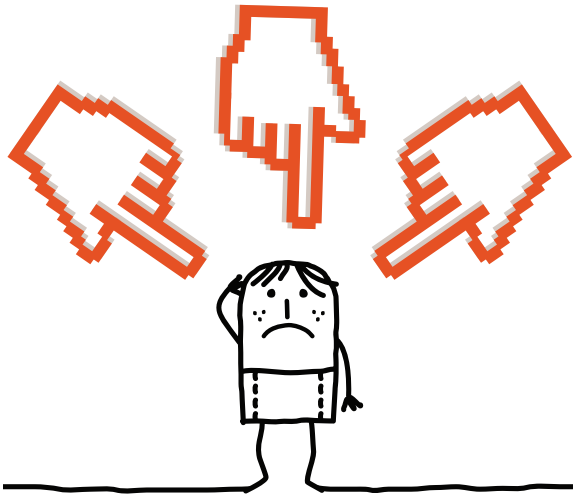

FALSE POSITIVES 3

Perhaps the most cited reason for the dearth of behavioral analysis technology is the tendency for such approaches to generate a frustrating number of false positives; a.k.a. "false accusations". The traditional approach is to establish a baseline model for an individual's normal computer use. When outlier activity is detected, the resulting alert would trigger an investigation to confirm or reject the suspicion. Unfortunately, with humans, normal behaviors are so highly anomalous that any statistical threshold worth using is constantly exceeded, resulting in a flood of alerts. On a given day, an employee might view dozens of websites or internal files they never previously accessed. The specific destinations and targets vary greatly from day to day, even within a highly correlated behavioral profile. A practical solution must embrace anomalies as part of normal baseline behavior.

The approach we invented is to incorporate low-level anomalies into a statistical behavioral profile. Individuals that express similar behaviors are grouped closely together, thus providing a local baseline from which to evaluate each member of the group. Groups are not only related by their network usage behavior, but also in the way they generate anomalies. Anomaly type and frequency are considered when comparing one individual's behavior to another.

This approach not only mitigates the problem of flooding with false positives, it also increases the sensitivity of the instrument for detecting insider threats.

DETECTING INSIDER THREATS

Having solved the three key challenges in behavioral analysis, we have successfully demonstrated a viable method for detecting insider threats. The mechanics of putting it together are straightforward, but within each sub-process there exists a sophisticated set of statistical models for determining baseline behaviors and identifying the outliers.

The five processing steps at the core of a real-time insider threat detection capability are:

1. Extract behavioral features
2. Compare behaviors across a population
3. Cluster individuals by behavioral similarity
4. Construct profiles of each group
5. Discover outliers within groups

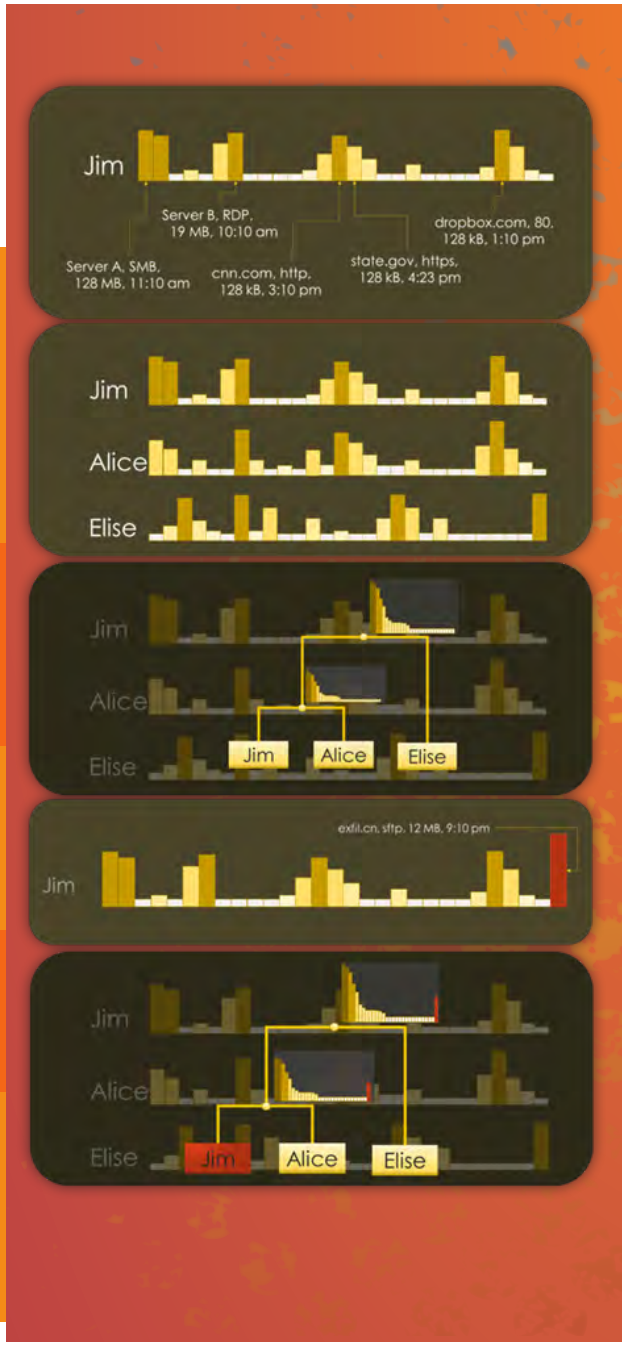

1 **Extract Behavioral Features**. Mirror port taps are strategically positioned on the network to capture all packets transmitted between monitored devices such as desktops, file servers and application servers, as well as external traffic to and from the Internet. From "big pipe" feeds, our sensors extract a "small pipe" of feature vectors that define an individual's behavior on the computer network.

2 **Compare behaviors across a population**. Using the feature vectors collected from raw network traffic, behavioral profiles are constructed for every individual and device. Each behavioral profile is then compared to all other profiles present on the network.

3 **Cluster individuals by behavioral similarity**. The system places individuals with similar behavior profiles into groups. Thresholds for group boundaries are determined with non-parametric statistics and are continuously evaluated.

4 **Construct profiles of each group.** Each group is characterized with its own multidimensional statistical distribution of network behavior. Baseline activity is determined, continuously monitored, and autonomously adjusted.

5 **Discover outliers within groups.** New data ingested by network sensors are evaluated within the context of a group's distribution function. If a feature falls outside a computed range then an anomaly is identified. An alert is raised if the anomaly falls outside the group's anomaly baseline.

Raw data captured by network mirror ports is exceptionally noisy and the features used in statistical analysis are not independent. Most traditional statistical methods operate on an assumption of independence[1], but the solution for insider threat detection requires a more sophisticated approach; one without the use of well-known distributions (e.g. Gaussian, Poisson).

Another reality of behavior to consider: organizations, teams and individuals change behaviors all the time. Individual behavior patterns change when the company adopts a new business model, adds new clients, or reacts to a strong market event (e.g. the Recession of 2008). Teams adjust behaviors as they transition from old projects to a new ones. Individuals change behavior when they get promoted or assigned new tasks and teams. Simply going between summer months and winter months causes behavior changes. The point is, behavior patterns change continuously and the statistical models used to determine baseline behaviors and identify anomalies must adjust with them. As such, any practical detector of rogue insider behavior must implement an appropriate unsupervised learning mechanism and a sliding temporal window. The statistical sample set needs to be constantly moving in time and never stop; something traditional "at rest" data analysis does not handle well.

1 Many traditional statistical methods are built on key assumptions, such as the assumption of independence. Failure to recognize when the assumptions do not hold true can have drastic consequences. Some credit the incorrect use of the assumption of independence as a key cause to the Sub-prime Mortgage Crisis. Information on the crisis and assumption can be found here http://en.wikipedia.org/wiki/Causes_of_the_Great_Recession; information on statistical assumptions can be found here: http://en.wikipedia.org/wiki/Statistical_assumption.

PRACTICAL | FIELDABLE | SCALEABLE

INSIDER THREAT DETECTION TECHNOLOGY

PERSONAM's Insider Threat Detection technology consists of a central processing device and enough deployable sensor units to cover critical network components, such as enterprise servers, local hosts, and WAN uplinks to the Internet. Sensors reduce the data by extracting key features and forwarding just the features to a central processing device where aggregation and behavioral analysis is performed.

A security professional operates the insider threat detection instrument through a web-based experience hosted by the central processing device. Once signed in, a user can access three primary views: Situational Awareness, Alerts, and an Anomaly Detail & Forensics Toolkit. There is also a system dashboard that summarizes the systems health, status, and performance.

The Situational Awareness function provides a sweeping view into the behavioral profiles across the protected organization. At the heart of the monitor is a dynamic graphic showing how each individual fits within the organization based on their behavior profile. Individuals with similar behaviors will be clustered together, forming cohesive groups. The real-time visualization also labels ongoing activity for each individual and group, including internal and external file transfers, chat sessions, web access, streaming content, and more. When an individual begins to behave in an unexpected way, the visualization immediately brings the suspicious activity to the attention of the security professional. Alerts are logged and notifications are sent via email or SMS to predefined recipients.

With the Alerts view, a security professional can examine past and current alerts to help assess each suspected threat. The source and severity of each alert is provided in a summary view. The instrument provides easy access to details regarding suspicious individuals and equipment, with convenient access to forensic tools. From the Alert screen, a security professional can adjudicate alerts as benign, suspicious, or threatening.

The Anomaly Detail and Forensics toolkit is accessible from both the Situational Awareness function and the Alerts view. It allows the user to dig deep into an individual's baseline profile and to see why current activity might have been considered suspicious and potentially threatening.

TYPICAL DEPLOYMENT

Most deployments of Insider Threat Detection are delivered as a subscriber service with onsite hardware for monitoring. The first step is for a specialist to review the organization's threat concerns and digital infrastructure. The specialist will provide the subscriber with an assessment of coverage that is possible, what live streams to mirror, and how many sensors are required. In many cases it is sufficient to monitor external web traffic only and not tap the LAN at all.

Other situations require monitoring "host to host" traffic, i.e. traffic between desktops, thin clients, file servers, application servers, and IP enabled appliances. The number of deployed sensors will depend on the desired level of coverage, the size of the organization, and the digital infrastructure.

Once sensors are installed the system is then professionally configured and calibrated. Threat Alerting can be adjusted between high-sensitivity and moderate-sensitivity. Alert "notify lists" and special "watch lists" will be created.

Once activated, the system enters a short calibration phase to establish baseline behavior profiles for the organization. During this phase there are a large number of false-positive alerts as the initial calibration is refined. Professional installers work with the subscriber to quickly adjudicate the initial findings. In this critical stage it is important not to dismiss alerts too quickly. Our advisors often find a number of unexpected infrastructure and configuration "surprises" during calibration, including back-office servers and network components that may be mis-configured or off-specification.

The calibration phase does not last long, and within a few days the administration and monitoring of the subscriber's network can be moved offsite to a remote security operations center. Our professional threat analysts continuously watch over the network in search of insider threat behavior and rogue employees. The subscriber can enjoy a heightened sense of security knowing their valuable assets are being monitored.

Although Insider Threat Detection is a powerful behavior profiling technology, it does not encroach on individual privacy. For a large majority of deployments the technology will never need to examine the packet data itself and therefore cannot be used to "snoop" into private communications. The Insider Threat Detector monitors behavior without opening people's mail or syntactically analyzing communications. The forensic reports it generates are very similar to reports already produced by common enterprise security infrastructure.

ESSENTIAL PROTECTION

Businesses and government agencies already appreciate the importance of protecting their digital perimeter against outsiders because daily attacks come from aggressive adversaries intent on gaining access to intellectual property, customer information, and even state secrets. External attacks attract most of our attention because they are persistent and somewhat detectable, making advanced defensive measures, such as next-generation firewalls, an everyday aspect of life.

However, the real threat is **inside** where defensive options are much more limited against the threats that exist within the secure perimeter, be it a person, an appliance, or malware. Prior to developing this technology, there was little if anything one could do to systematically protect against human behavior. Behavior analysis delivers peace of mind and finds "bad guys" when trusted people act in an untrustworthy way.

Despite the substantial money and effort invested in defensive cyber-security, the inside is left almost completely undefended. Anti-virus technologies only defend against malicious machine-executable threats, they do not detect humans with ill intent. Even the newest whitelisting solutions that secure network endpoints cannot detect when someone exploits a legitimately authorized access for illegitimate purposes.

personam™

www.PersonamInc.com

Contact us:
Personam, Inc.
1420 Spring Hill Rd., Suite 525
McLean, VA 22102

email: Insider@PersonamInc.com
tel: (571) 297-9371

Personam's Insider Threat Detection technology is currently the most advanced behavior detection system in the world for early warning and is able to quickly pinpoint insider threats such as rogue employees. The technology is designed specifically to identify malicious human behavior. The technology also alerts to off-specification hardware and other more traditional cyber attack vectors.