

Enhanced Security and Compliance through Neural Data Analysis

A white paper by Privacyware

Published by:

Privacyware
68 White Street, 2nd Floor
Red Bank, NJ 07701
Email: info@privacyware.com
URL: <http://www.privacyware.com>
July 2007

The Neural Security Layer

A steady and seemingly endless stream of high-profile security breaches, affecting some of the world's largest and most respected businesses and information technology systems, has been reported during the past several years. The list of personal information stolen or otherwise exposed to skilled and often well-funded criminals includes the social security numbers, credit histories, bank account and credit card numbers of literally millions of individuals. In other cases, information and assets pertaining to customers, product designs, critical infrastructure, financial and other sensitive information are of primary interest to attackers. The risks associated with unauthorized access to critical systems and data have never been higher and the impact of a serious incident can be devastating.

Organizations possessing the most valuable assets are logically the prime and most frequent targets. Ironically, these entities also possess the greatest ability to defend against Internet and similar cyber-attacks. The take-away of this fact is that realistically, it will never be possible to eliminate such attacks altogether and organizations of any size that handle sensitive data that expect to remain strategically competitive, demonstrate compliance with the increasingly explicit regulations, and maintain the confidence of the respective communities they server must consider every viable technique to demonstrate as proactive and effective a defense as possible.

Each day, the security, network and compliance professionals on the front lines in the pursuit of this objective face increasingly diverse and dynamic challenges to maintain the critical systems on which their enterprise relies as well as to detect and investigate possible breaches or policy violations. Attaining the upper hand in this process is dramatically enhanced through the review

and analysis of logs and other data generated by the various systems, applications, firewalls and other devices deployed across the enterprise. Advances in analytic technologies are helping administrators do a better job of identifying threats and policy violations amidst a rapidly expanding volume of data.

Among the emerging technologies demonstrating compelling value include software applications that leverage the latest neural technologies to enhance security and other related data analysis. Neural software applications use complex mathematical algorithms that scour vast amounts of data and categorize them in much the same fashion as would a human. What's the difference? Neural applications can examine far more data in less time and more comprehensively than a human can, highlighting those events that appear suspicious enough to warrant human or automated attention. As security administrators deal with these events, the actions they take are added to the knowledge base, enabling the neural system to continuously "learn" more about its operating environment. The expert security or compliance professional, coupled with neural data analysis technologies, form what Privacyware calls the "**Neural Security Layer**" within a complete information assurance and enterprise risk mitigation framework.

Privacyware is a leading innovator in the use of neural-based technologies that enhance data analysis and threat prevention. The company's Adaptive Security Analyzer (ASA) software applies techniques such as **fuzzy clustering** to define normal sets of activity and **kernel classifiers** to deal with events that don't neatly fall into any predefined cluster. Adaptive Security Analyzer can work entirely on its own, defining a baseline of normal activity and then reporting on events outside of that norm, but is optimized when mentored through interaction with the security or compliance practitioner.

A Human Touch

Adaptive Security Analyzer (“ASA”) classifies new events and highlights those that appear most threatening, allowing the security expert to be the final arbiter of what is and is not an actual threat or defer judgment to the system itself. In the process, the system continuously updates itself, learning more about its environment.

In that respect, ASA mimics the steps a human security professional would logically follow to analyze network and application traffic data. Consider the example of a company that hires a new security professional to work with existing security tools. The normal course of action would have this individual begin to learn the network configuration and become familiar with normal traffic data patterns. Early on the employee may be notified of an alert where a certain DLL is being accessed frequently. Since it is the first time the person has seen this alert, they will do some research and eventually determine that it’s happening because the company’s antivirus software was just updated. So, this represents an unusual event but not a harmful event. The next time this alert is received they will know what is occurring and will quickly dismiss it. The event and its unique attributes have been added to the security professional’s body of knowledge. Adaptive Security Analyzer processes information in a similar way, considering events with evolving levels of knowledge and scrutiny and thereby producing the increasingly accurate classifications that you would expect from a capable professional over time given experience and training.

Technologies including fuzzy clustering and kernel classifiers allow Adaptive Security Analyzer to automatically and via interaction with the administrator learn about the environment in which it functions. ASA will identify events that are out of the ordinary and that have not been classified as benign.

In that fashion, ASA allows the security professional to quickly *hone in* on events that may be harmful, out of the thousands of events that occur each day or even each hour. The security professional can thereby prioritize his workload based on the events which reflect the highest severity and potential risk to the organization.

Six Steps to Producing Security Intelligence

1. Designate Data: Data can be system log entries or any other raw or formatted measure of activity in an IT environment.

2. Model Analyst Expertise: Variables, weights, model centers, and pertinent event knowledge comprise the analytic or data mining model and are configured based on the specific analysis requirements and the unique attributes of the particular environment.

3. Train Model: This is the process of organizing the designated security data into multi-dimensional “event vectors” within the context of the analytic models. In this way, a baseline of activity that is typical for the particular environment during a specific period of time can be established.

4. Generate Knowledge: Live or offline data is compared against the training baseline and is classified or grouped accordingly. Reports list in order of severity abnormalities which represent the most unusual and potentially threatening activity within the system. The specific elements of each event will be ranked in relation to the degree they contributed to the event’s classification, thereby guiding the analyst to an understanding of the root cause of the event.

5. Tune Baseline: User-supervision and infusion of expert knowledge is essential to accurate event classification and system baseline maintenance. This will help to filter out non-threatening anomalous activity. Administrators can override the automated classification of single or multiple events and manually initiate retraining and re-querying of the training baseline. Over time, these interactions will lead to a significant reduction of output or “noise” that while unusual, may not represent threats.

6. Leverage Knowledge: The system output is invaluable for real-time or offline analysis, along with the desired detection and prevention of potential internal or external criminal activity or system misuse. One now has the ability to document, prioritize and address security events rapidly. This intelligence can be used to expand or enhance existing device and system policies. Output can also be exported to other applications for additional reporting or analysis.

Fuzzy clustering

Fuzzy clustering is one of the neural technologies at the heart of Adaptive Security Analyzer. The technology works by "training" itself, creating a baseline profile of the network, database, or device in various states to determine what happens under normal conditions. For example, it can determine what different system users do -- the resources they typically request, what types of files they transfer and so on. All those routine events are then grouped into clusters that represent normal activity.

For example, it may be sensible to define models that focus on different sorts of users, such as administrators, marketing employees, and anonymous ("un-trusted") end users. For each type of user, Adaptive Security Analyzer will determine which events are considered normal and group them into a cluster. The goal is not to determine an exact profile of what any given type of user does but rather to establish patterns. "**Fuzzy-logic**" uses algorithms that identify these patterns and separates clusters accordingly.

Kernel classifiers

Kernel classifiers add additional analysis as to when an event or group of events comes along that can't be neatly classified into an existing cluster. The classifiers use algorithms that allow Adaptive Security Analyzer to determine which cluster the event most likely belongs. The algorithms are based on non-linear distribution laws, which use statistics to track what happens over extended periods of time.

In essence, the profile that Adaptive Security Analyzer provides when it is first installed amounts to a series of distribution laws, which can be thought of as averages. For example, a marketing employee typically logs in at 9 a.m., looks at the CNN Web site for 15 minutes, then logs into the sales system, then the CRM extranet and so on. That series of

events represents a typical, average day for that employee and is defined in a distribution law. If this employee's routine would deviate significantly from this pattern, ASA would detect this activity and quantify the extent of the deviation. For example, if this same employee accessed the network at 2am on a Sunday morning and generated large file transfer volumes, ASA would alert and provide substantive evidence of the activity to the administrator.

ASA utilizes special metrics that "measure" how far any given event is from any existing, known cluster. In this way, the system can determine which cluster a new event should most likely belong to, a process known as dynamic clustering.

Putting ASA to work

Enterprise Security Professionals can easily configure Adaptive Security Analyzer to examine data derived from almost any network resource or information system asset by creating a model for that resource through the use of ASA's graphical user interface. Implemented as a Snap-In to the Microsoft Management Console (MMC), Adaptive Security Analyzer has a wizard-driven interface that makes it a straightforward exercise to define the data source to be analyzed, the variables to consider and any special considerations related to this data and the intended purpose and objectives of the analysis.

For example, perhaps management has compliance requirements to monitor activity in a messaging environment like Microsoft Exchange. There is a general requirement to be notified of any abnormal activity but also very specific events such as first mailbox access. The event variables for this model could be UserName, UserDomain, DayWeek, and ObjectName. Within the model you could also assign weights to the model variables, perhaps specifying that one variable is more

important in the total analysis than others. In addition, special filters or rules pertinent to events that should carry custom metrics values or be filtered out can be defined. All of these attributes add up to form a model or framework in which the data can be organized and compared in a manner that emulates the thinking process of the human analyst.

Next, you specify a time period, perhaps one week or one month, which ASA will use to train itself and create the baseline of activity. Once ASA completes its training cycle, it creates a series of clusters that represent normal (trusted or untrusted) activity.

This knowledge base will be continuously updated based on results of day-to-day activities, such as when administrators reclassify events. Perhaps ASA flags an event as suspicious and the security administrator knows this event is benign, as in the previous example suggesting the anti-virus software update. The administrator will reclassify that event, and the knowledge base will be updated accordingly.

The knowledge base can also be updated using data from third-party sources, such as Intrusion Detection Systems (IDS), or information compiled by internal IT resources such as list of untrusted IPs or URLs. Ideally, these rules are applied before the initial training, so ASA can flag known threat events within the training set and classify them accordingly.

This ability to take advantage of data from existing security tools is one reason that ASA is an extraordinary complement to existing IT investments. For example, Adaptive Security Analyzer can enhance the power of a Security Information Management (SIM) tool that collects and aggregates disparate device and system data into a single massive repository. ASA can reduce millions of records collected by such systems to a top ten or top 100 list of events reflecting the highest potential threat in order of priority.

ASA also includes pre-built analytic models for a variety of common log types and behaviors of interest including unusual logons or privileged access, web/http traffic, firewall events, Web Proxy, database transactions, application level activity and others.

Once ASA has established the analytic baseline, users must specify ranges of data (typically by date or some other period) to compare against the baseline. The ASA administrator employs MMC to monitor the reporting output that ASA produces. Reports are displayed in tabular and graphic formats. In addition, OLAP-compliant tools, such as Microsoft Analysis Services can be utilized to perform additional report manipulation and presentation. ASA lists unusual events for any period of time, listed in order of severity or by date, enabling users to quickly identify and focus on the most critical vulnerabilities amidst perhaps millions of events.

Adaptive Security Analyzer System Requirements

Hardware

- 1 GHz Pentium® III or faster
- 1 GB RAM
- 5 MB of free disk space (for ASA Pro software)

Software

- One of the following operating systems:
 - Windows® 2000 Server
 - Windows® 2003 Server
 - Windows® 2000 Professional
 - Windows® 2000 Advanced Server
- Microsoft® Analysis Services (for SQL-based OLAP)
- Service Pack 3 (or later) for Windows® 2000
- TCP/IP network

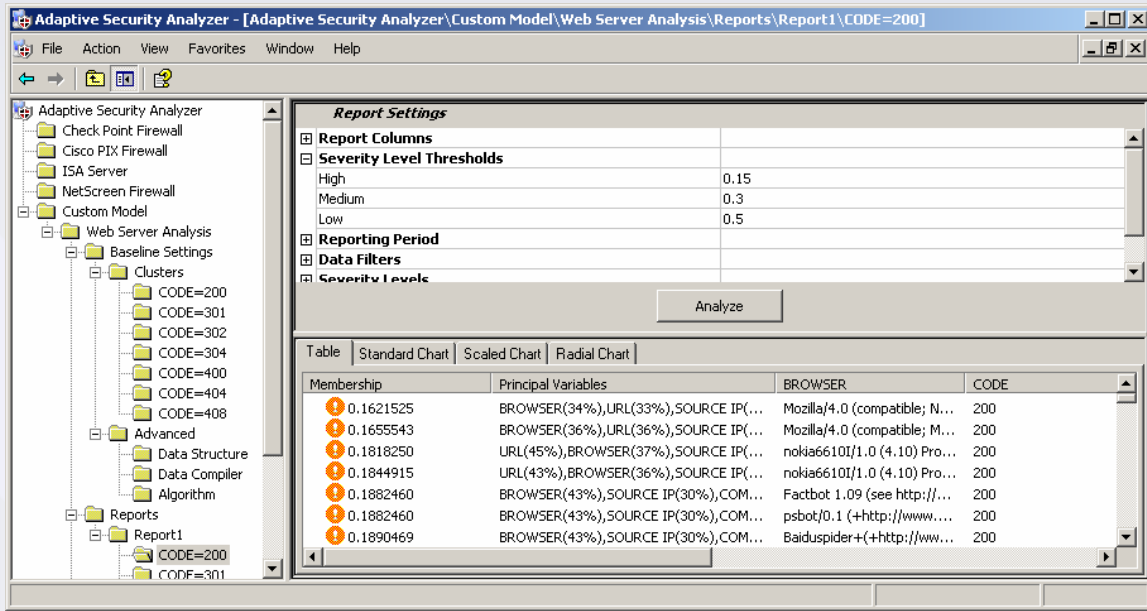
Implementation:

Microsoft Management Console (MMC) Snap-In

Data formats and Analytic Models:

Custom analysis can be performed using most standard data formats such as syslog, .csv, .txt and databases including MS SQL, MySQL, and Oracle.

ASA supports leading SIM/SIEM solutions from vendors such as Cisco, LogLogic, netForensics, Sensage, and Quest Software and includes pre-built models designed specifically to analyze logs from Juniper, Cisco, Check Point and Microsoft ISA Server devices and applications.



Adaptive Security Analyzer generates reports in MMC that identify all events that depart from an established normal baseline. It measures the extent of the event abnormality and provides information regarding the factors that most influenced the event's classification. ASA also utilizes rules, filters and "schemas" (expert rules and pattern detection) to identify events of interest. This information enables security staff to identify the most critical vulnerabilities immediately and address threats more efficiently and effectively.

Conclusion

Adaptive Security Analyzer combines expert rules with concept-based artificial intelligence technology to identify and provide insight about known threats and activity of interest as well as new or elusive threats. ASA empowers administrators and analysts challenged with monitoring and securing computing systems, managing compliance and assuring the integrity of database or other critical infrastructures. ASA's neural core emulates the cognitive and self-learning process of a human analyst and can sift through massive volumes of logs to quickly reveal and prioritize the most critical events and advise users of the factors of highest influence to event classification.

Adaptive Security Analyzer's capacity vastly surpasses its most capable human counterpart or advanced correlation, log or security event analysis tool. For security and compliance practitioners responsible for the review and analysis of security related information, ASA delivers an early warning capability to alert, assess and aid in the remediation of known and new security and compliance-related threats.

ASA complements and enhances the effectiveness of related systems and personnel, helps overcome overwhelming information overload challenges, improves

network security and systems management, and strengthens the ability to meet the analytic aspects of regulatory compliance demands.

Adaptive Security Analyzer performs the needed analytic heavy-lifting, enabling network and security personnel to devote more time to the critical tasks of eliminating vulnerabilities, improving and optimizing network performance and serving the needs of end users.

The benefits of implementing a "Neural Security Layer" using Adaptive Security Analyzer are direct and measurable: Higher levels of network and system availability, decreased overall risk, increased productivity and maximized return on security and other IT asset investments. Neural security applications dramatically reduce the resources required for security data management and analysis and generate intelligence indispensable to defending against digital saboteurs, network trespassers or any other type of cyber-criminal.

The technology to deliver these advantages is emerging today, and as neural security applications and solutions continue to evolve, deploying a Neural Security Layer will become recognized as a best practice within the overall enterprise security framework.

About Privacyware

Privacyware is an innovative developer of security data analysis and threat prevention software. Leveraging advanced competencies in non-linear mathematics, neural networks and self-learning systems, and proficiency in complex software and systems, we deliver security data analysis products that enable enterprise security and compliance personnel to more thoroughly understand the environments for which they are responsible and to more effectively identify and comprehend malicious and/or deviant activity. Privacyware also develops award-winning host and desktop defense offerings that increase the level of protection from new and known malware and intrusions in individual, small business and large enterprise computing environments.

For more information about Privacyware visit: <http://www.privacyware.com>, call +1-732-212-8110 x235, or email info@privacyware.com

About the Authors

Gregory J. Salvato

Mr. Salvato serves as Chairman & CEO of PWI, Inc., a leading provider of advanced technology products and solutions in IT Security, Enterprise Messaging, Analytics and Custom Software Development, and serves as a Director of electronic transaction and funds management services provider, Internet Transaction Solutions, Inc. Mr. Salvato is a published author and a graduate of The Ohio State University with a degree in English. He is a member of the ISSA and IEEE.

In 1999, with Dr. Konstantin Malkov, he conceived of and led the launch of Privacyware, to provide direction and legitimacy to the practical use of neural analytics in the monitoring, detection and prevention of harmful computing activity, information assurance and risk management. In addition to his day-to-day operations responsibilities, Mr. Salvato drives the company's thought leadership activities and serves as the company's chief spokesperson.

Konstantin V. Malkov, Ph.D.

Based in the US, Dr. Malkov is a recognized, international scientist and the Chief Technology Officer for Privacyware. He is a leading expert in non-linear mathematics and is co-founder of the Department of Non-linear Dynamic Analysis and The I&C Laboratory at Moscow State University. Dr. Malkov has managed a number of high-level mathematical modeling project in physics, mechanics, and control theory in the United States and Europe dealing with a wide spectrum of requirements including: seismological inverse problems, numerical analysis, and flight control. He has overseen the development of dozens of commercial software packages and authored more than 30 scientific articles.

He is a former full professor of Applied Mathematics and Computer Science at Moscow State University and received Ph.D.s in Mathematics and Computer Science from that institution in 1986.

References

1. Herbrich, R. "Learning Kernel Classifiers." The MIT Press, 2002, Cambridge, MA, USA.
2. Cooley, R. Mobasher, B., Srivastava, J. "Web Mining: Information and Pattern Discovery on the World Wide Web (A Survey Paper)." Proceedings of the 9th IEEE International Conference on Tools with Artificial Intelligence, (ICTAI'97).
3. Yager, R. "Indistinct Sets and Theory of Capabilities." Radio i sviaz, Moscow Russia, 1986.
4. Averkin, A.N., Batyrshin, I.Z., Blishun, A.F., Silov, V.B., Tarasov, V.B. "Indistinct Sets in Models of Management and Artificial Intelligence." Nauka Publishers, 1986.
5. Kofman, A. "Introduction in the Theory of Indistinct Sets." M., Radio i Sviaz, 1982.
6. J. Gray, J., Bosworth, A., Layman, A., Pirahesh, H. "Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals." Proceedings of the IEEE International Conference on Data Engineering, pp. 152-159, New Orleans, February, 1996.
7. Rogerson, Dale E. "Inside COM.", Microsoft Press, 1997.
8. "Report to Russian Foundation of Basic Researches". RFBR Grant # 97-01-00140-a, 1997.
9. Ben-Hur, A., Horn, D., Siegelmann, H.T., Vapnik, V. "Support Vector clustering." Journal of Machine learning Research, Vol. 2, (2001), 125-137.
10. Krishnapuram, R., Keller, J. M. "A Possibilistic Approach to Clustering." IEEE Trans. Fuzzy Systems. Vol. 1. No. 1 (1993), 98-110.
11. Scholkopf, B., Smola, A., J. "Learning with Kernels: Support Vector Machines, Regularization, Optimization and Beyond." The MIT Press Cambridge, Massachusetts (2002).
12. Eskin, E., Arnold, A., Prerau, M., Portnoy, L., Stolfo, S. "A Geometric Framework for Unsupervised Anomaly Detection: Detecting Intrusions in Unlabeled Data." Applications of Data Mining in Computer Security, Kluwer, (2002).
13. Bottomley, L. "Dataset: A day of HTTP logs from the EPA WWW Server." Duke University, (1995) <http://ita.ee.lbl.gov/html/contrib/EPA-HTTP.html>.
14. Girolami, M. "Mercer Kernel Based Clustering in Feature Space." I.E.E.E Transactions on Neural Networks, 13 (4), (2001), 780-784.