

SANS

ANALYST PROGRAM

Sponsored by Sourcefire

Real-Time Adaptive Security

A SANS Whitepaper – December 2008

Written by: Dave Shackleford

Event Data in Context

**Additional Real-Time
Threat Management
Benefits**





Introduction

In today's dynamic threat and networking environments, standalone Intrusion Detection/Prevention Systems (IDS/IPS) cannot protect against ever-changing attacks and vulnerabilities. The reason: Standing alone, IDS/IPS lacks the context it needs to reliably distinguish an event from a non-event and prioritize protection based on business-critical rules. Context can be helpful in determining when an event indicates a security incident, such as a deliberate remote buffer overflow exploit attempt, as well as when events are nothing more than false positives, such as poorly configured applications sending out broadcast packets.

IDS/IPS also needs context to adapt to changes occurring in enterprise networks. New user-demanded applications such as Voice over IP (VoIP), Web 2.0, virtualization, and other infrastructure applications open new attack surfaces. In these cases, attacks against the Session Initiation Protocol (SIP) and H.323, as well as Web application attacks including Cross-Site Request Forgery (CSRF) and SQL Injection have become commonplace.

Context can also help IDS/IPS recognize new types of attacks and vulnerabilities. Commonly-used exploit frameworks and toolkits such as Metasploit allow for obfuscation and modification of attack processes specifically to prevent signature- and behavior-based detection. Meanwhile, the mean time between when a vulnerability is published and when the exploit code for that vulnerability is released has narrowed significantly. According to Symantec's Internet Threat Report, Volume XI, 25% of vulnerabilities had exploit code available within one (1) day of release, and 31% had exploit code available within six (6) days.¹ In the second half of 2007, nine (9) zero-day exploits were found in the wild.²

For organizations to adapt in real-time to "unknown" or "unknowable" threats being presented in this dynamic network and threat environment, their intrusion systems should tap into security information generated by users, devices and systems throughout the network. This includes data from passive traffic monitoring correlated against data from vulnerability assessments, Network Access Controls (NAC), user data and other system data. Figure 1 depicts a simple evolution in network security strategy that leads from our current state (point solutions that perform specific functions in silos) to a more real-time adaptive approach to intrusion detection and prevention.

¹ eval.symantec.com/mktginfo/enterprise/white_papers/ent-whitepaper_internet_security_threat_report_xi_03_2007.en-us.pdf

² eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xiii_04-2008.en-us.pdf





Figure 1: Moving to continuous compliance enforcement

Here's an example of how adaptive security works: A behavioral-based rule triggers IPS alerts for multiple malformed packets. Instead of sounding an alarm immediately, the intrusion analysis system checks the most recent scanning results on the server under attack. Those results show the system is missing several recent patches. Passive traffic analysis reports then reveals that the server has been attempting to communicate with unusual ports on local systems. With this additional information, it becomes clear that the organization has a zero-day attack occurring inside its network.

With security actions based on context, intrusion systems can adapt to real-time threats like these while giving visibility into what to investigate, where to investigate, and even take or recommend action based on preset rules.





Event Data in Context

Over time, many enterprise networks have become cluttered with point solutions and security controls that can only perform a single function. Each of these controls provides only one piece of the overall puzzle—a puzzle that is growing larger and more complex by the day. Figure 2 depicts a simple network diagram with security controls as they're commonly employed today. First-generation, signature-based intrusion detection and prevention systems can detect and block well-defined attacks, but they have no knowledge of applications in use, normal traffic patterns or user activity in the context of a network's normal behavior patterns. Standing alone, each of these controls cannot tell an actual attack from a non-attack, nor can they assess impact, track down users and affected systems, set priorities, or take defensive and offensive action. To do any of these things takes integration and correlation between security controls.




Figure 2: Traditional network security controls





See What You're Protecting

Today, organizations can build or find IDS/IPS correlation engines to add this context for security that can adapt to new threats and changes in real-time. This starts with visibility into today's complex networks in which systems and applications may be improperly configured or have known vulnerabilities, and users might be accessing systems they shouldn't even have access to.

Previous generations of intrusion detection and prevention technology have only looked at attacks as individual events, with no consideration given to the surrounding network's state at any given time. An IPS could be configured to understand that certain types of services were available on certain subnets (for example, a Web server DMZ consisting of the subnet 10.10.10.0/24 with HTTP and HTTPS services available).

There is a lot more data available, though! Passive monitoring technologies can provide system and application identification data; local system logs and user directory services can provide information on which users are logged into systems and applications at any given time; active scanning tools can provide data about current system vulnerabilities; and vulnerability databases can correlate this with known vulnerabilities in the wild. An intrusion prevention system could easily leverage this data to gain a better understanding of what systems and applications it's protecting, where they are in the network, and what they may be susceptible to.



Understand Behaviors

Network Behavior Anomaly Detection (NBAD) techniques were originally developed to monitor network traffic thresholds for shifts or changes that could indicate Denial of Service (DoS) attacks or signal that troubleshooting was needed to address specific issues with systems. Over time, NBAD has evolved into Network Behavior Analysis (NBA), which does not rely as much on traffic thresholds and, instead, focuses on the establishment of a comprehensive network baseline. This overall baseline is then continually monitored for deviations or exceptions. There are three major components of a modern network behavior monitoring strategy for use in information security and network operations:



1. Traffic flow patterns and data: Many NBA products today have focused on network flow data such as NetFlow, sFlow, and Jflow. Although some formats of flow data are specific to one vendor or another, most include a number of traffic attributes that provide information about what systems are communicating, where the communications are coming from and headed to, and in what direction the traffic is moving. Common attributes include source and destination IP addresses, source and destination ports, and device interfaces. Recent versions of NetFlow contain a number of new attributes, such as MultiProtocol Label Switching (MPLS) and IPv6 data. An example of NetFlow records is shown in Figure 3, including interfaces, source and destination IP addresses, protocol in use (TCP), source and destination ports, packet counts, and start and end times of flow data capture:

Sif	SrcIPAddress	Dif	DstIPAddress	Pr	SrcP	DstP	Pkts	Octets	StartTime
EndTime		Active	B/Pk Ts Fl						
0059	127.0.0.1	005b	219.140.194.174	06	50	4f3	1	40	0721.21:58:00.593
0721.21:58:00.593		0.000	40 00 14						
0059	127.0.0.1	005b	219.148.205.228	06	50	6ef	1	40	0721.21:57:56.533
0721.21:57:56.533		0.000	40 00 14						

Figure 3: Sample NetFlow records

2. Network Performance Data: Traditional network monitoring tools have focused on Simple Network Management Protocol (SNMP) events and Quality of Service (QoS) data for system and network performance. This data is most often used for performance and uptime monitoring and maintenance. It can also be leveraged for security purposes. For example, availability of Voice over IP (VoIP) networks is critical, because any interruptions may cripple telephone service in a business. Spikes in network traffic can also indicate malware or Denial of Service (DoS) attacks underway and may be the first signs of a zero-day worm or other fast-spreading malicious code.

3. Passive Traffic Analysis: Passive traffic analysis tools are similar to sniffers in the way they operate, employing network interfaces in promiscuous mode to intercept and analyze traffic as it passes by. The focus of these tools is very different, however. Passive analysis tools can continually monitor traffic for protocol anomalies, tunneled protocols such as IRC commands within ICMP traffic, use of non-standard ports and protocol field values, etc. All of these are indicative of different security events and are subject to further correlation. Because these tools are also capable of inspecting application-layer traffic, security professionals can observe unique application attributes and behaviors to identify operating platforms and their potential vulnerabilities, based on the application's traffic characteristics. Ultimately, these tools can also provide a much higher degree of visibility into what systems and applications are communicating with one another and where they are on the network, which provides intrusion prevention systems with much-needed environmental context.





Know Users

Insider attacks are on the rise. According to the 2008 CSI Computer Crime and Security Survey, an average of 50% of responding organizations claimed insider abuse had occurred in the past year³. To monitor and act on internal abuse, as well as comply with data protection regulations, organizations need to tie network security events to local systems and user credentials.

Correlating threat information from intrusion prevention systems with actual user identities (logged on to local systems) allows security professionals to identify breaches of policy and fraudulent activity more accurately within the internal network. They can use NetFlow to gather MAC and IP addresses and other device data as it traverses the network. Tying this location data to logon information from Lightweight Directory Access Protocol (LDAP) stores (or other centralized account data repositories) allows security teams to track down specific users and workstations to remediate vulnerabilities or stop attacks much more quickly. This can also help security personnel avoid wasting time on false-positives because they can contact a user to investigate behavior instead of having to track down systems manually.



Determine Impact

It's a known fact that a lot of time is wasted on analyzing false positives generated by first generation, stand-alone intrusion systems. Depending on the environment, false positives can often be numerous and very difficult to verify, costing analysts a valuable time determining whether or not something is an event the analyst should be worried about.

Real-time Adaptive Security begins with passive traffic analysis. Once a threat is determined, the intrusion system leverages tools like Nessus to scan impacted systems and dive deeper into the vulnerability detected. For more environmental context, vulnerability assessment data can then be analyzed in conjunction with network behavioral data to provide a true real-time snapshot of what attacks are happening and assess the likely impact on the target system(s).

As an example, suppose that a number of high-severity alerts are generated for a series of Windows-focused attacks against DMZ Web servers. If these servers are scanned for vulnerabilities to correlate with the alert data, analysts may find that all of the systems have been patched for the vulnerability or have a non-vulnerable Service Pack level, thus reducing the overall impact of the attack and setting an appropriately lower level of priority. In another scenario, these systems may all be running a Linux or UNIX operating system, thus rendering the attack entirely ineffective, so no alert would be processed.

³ www.gocsi.com





Respond Appropriately

Once systems have been given a criticality rating (for example, 1 for mission-critical database servers storing sensitive information and 5 for Web servers with less important static content) and a continuous stream of contextual environment data is gathered, then intrusion prevention solutions can begin to dynamically and proactively take actions to reduce operational overhead. For example, intrusion prevention rules that are not applicable to certain systems and applications in a particular IP range can be disabled, significantly reducing false positives. If new data indicates that a certain system has become vulnerable to known attacks, rules can be re-enabled.

Additionally, changes in system or network behavior can be linked with defined system criticality to elevate the need for additional investigation, helping analysts look for possible zero-day attacks or unidentified malware. For truly dynamic response, IPS sensors can trigger short-term network access control modifications to block malicious traffic immediately.



Enforce and Tune

Security teams have less time than ever to monitor and maintain a number of distinct network security controls. However, auditors assessing these controls to ensure compliance with regulatory, industry or internal policy-driven mandates will want to see that attacks are detected, incident response procedures are initiated, and log and audit trail data is generated for later analysis. In addition, regulators will want to see improvement. So, in addition to real-time protection, the ultimate goal of leveraging real-time adaptive security is getting to a state of continuous monitoring, assessment and network security improvement.





Additional Real-Time Threat Management Benefits

Because network security technology has evolved, security teams can adapt policies and processes to adapt to new threats posed by new technologies and architectures. The following are some examples:



Data Loss Prevention (DLP)

Since 2005, over 245 million sensitive data records have been involved in security breaches in the U.S., according to the PrivacyRights.org Web site.⁴ So, it's no wonder that security minds are starting to shift their focus from network protection toward data centric protections.⁵ How can network security technologies help with this problem? The answer is: Baby DLP. Instead of spending significant dollars right away on DLP, users can leverage real-time adaptive security to monitor data and behavior in several ways:

- **IPS** can *alert* on signatures with specific *keywords* related to sensitive internal data. If this data crosses the network, the IPS will drop the packets or send alerts.
- **Network behavioral analysis** can determine when *unusual patterns* of computing activity are occurring, possibly indicating fraudulent behavior.
- **Access control systems** can reveal *suspicious use* of *User IDs*. For example, a User ID logged onto a system containing sensitive information and then sending data to other systems may be indicative of stolen credentials or other illicit behavior.



Centralized Monitoring and Correlation

As more security and network components become integrated into a real-time threat management infrastructure, the ability to manage and monitor incidents and threat reports from a single console is paramount. Analysts cannot work with multiple consoles for each of their devices and platforms without creating additional confusion and correlation work for security and network teams. In large, complex environments, ease of use and single-console management can help to ease the administrative burden of analyzing events, reviewing alerts and generating reports.

⁴ www.privacyrights.org/ar/ChronDataBreaches.htm

⁵ www.networkworld.com/supp/2008/100908-trendwatch-information-protection.html





Virtual Environments

New security questions arise just as quickly as virtualization spreads in the enterprise. One of the biggest issues facing security professionals is “virtual sprawl,” or lack of visibility into what is running. Given the ease of deployment associated with virtual machines (VMs), many IT departments are realizing that they no longer have a way to maintain up-to-date system inventories. Additional concerns focus on visibility into the inter-VM traffic on a virtual host platform, because the entire switch is virtualized as well. Traditional physical monitoring solutions will not be able to inspect this network traffic. Finally, the new virtualization platforms themselves may have vulnerabilities that could lead to significant compromise, so these systems must be monitored and assessed for patches, configuration errors, etc.

Context-aware network security technologies can play a role in protecting virtual assets in the following ways:

- **IDS/IPS** can be adapted to *monitor virtual switches* via software integration or proper credentials on virtual host platforms.
- **Host platform vulnerabilities** can still be protected from network-based exploit with *intrusion prevention signatures*.
- **Passive traffic analysis** can be leveraged to *combat virtual sprawl*, because all virtualization platform vendors have assigned Organizationally Unique Identifiers (OUIs) for MAC addresses. As these systems come online, they can be detected and assessed.





Conclusion

Although most organizations currently employ some form of network intrusion detection or prevention, they're typically using first-generation tools that lack the context needed to react and adapt in real-time. Without context, IDS/IPS will continue in its notoriety for an over-abundance of false-positives that keep administrators tracking down unimportant issues while missing those that are important.

With context, intrusion sensors can dynamically react and respond to changing networks and threats, leveraging integral data from other sources on the network. With context, intrusion systems become more adaptive, real-time and accurate than their first generation predecessors. To gain context, intrusion systems leverage traffic monitoring and behavior analysis to determine what systems and applications are online and what machines and ports they're trying to communicate with.

Real-time adaptive intrusion systems integrate with network access controls and user data repositories for tracing events to systems and specific users; dynamically monitor traffic patterns to mitigate threats; leverage vulnerability assessment data to correlate with alerts to reduce false positives down to actionable alerts; and can be used to continuously tune sensors and rules.

All of these technologies and processes working together provides context that IDS/IPS needs to adapt to new threats in today's ever-changing network and threat environments. Real-time adaptive intrusion systems represent the next generation of IDS/IPS—adaptive, real-time, and accurately determining events, dropping non-events, and setting priorities.





About the Author

Dave Shackelford, Director of Configuresoft's Center for Policy & Compliance, is an instructor and course author for the SANS Institute, where he also serves as a GIAC Technical Director. He is the co-author of "Hands-On Information Security" from Course Technology, as well as the 'Managing Incident Response' chapter in the Course Technology book, "Readings and Cases in the Management of Information Security."

Previously, Shackelford worked as CTO for the Center for Internet Security, as well as for a security consulting firm in Atlanta. He has also worked as a security architect, analyst, and manager for several Fortune 500 companies. He has consulted with hundreds of organizations in the areas of regulatory compliance, security, and network architecture and engineering. His specialties include incident handling and response, intrusion detection and traffic analysis, and vulnerability assessment and penetration testing.



SANS would like to thank the sponsor of this paper:

SOURCEfire[®]

