InformationWeek
:: analytics

Analytics.InformationWeek.com

September 2010
$199

Strategy Session

Presented in conjunction with

SECURITY dark READING
Protect The Business · Enable Access

What’s Going On?

Monitor Networks to Thwart Intrusions

Your networks may be under attack as you read this, but unless your security personnel are analyzing logs and leveraging common tools that are well known to your network operations teams, you may not find out until it is too late. In this Dark Reading Tech Center report, we explain how your security and network teams can cooperate and use common tools to detect threats before your databases are compromised.

By John Sawyer

Report ID: S1810910

TABLE OF CONTENTS

ABOUT US | ***InformationWeek Analytics'*** experienced analysts arm business technology decision-makers with real-world perspective based on a combination of qualitative and quantitative research, business and technology assessment and planning tools, and technology adoption best practices gleaned from experience.

If you'd like to contact us, write to managing director **Art Wittmann** at *awittmann@techweb.com,* executive editor **Lorna Garey** at *lgarey@techweb.com* and research managing editor **Heather Vallis** at *hvallis@techweb.com.* Find all of our reports at *www.analytics.informationweek.com.*

John Sawyer
University of Florida

John Sawyer is a senior security engineer with the University of Florida, Gainesville, and a Dark Reading, *Network Computing* and *InformationWeek* contributor and blogger. Sawyer's current duties include network and Web application penetration testing, intrusion analysis, incident response and digital forensics. He was recently awarded a 2010 Superior Accomplishment Award from the University of Florida for his work as part of the UF Office of Information Security and Compliance.

Sawyer is a member of team 1@stplace, a small group of righteous hackers that won the electronic Capture the Flag computer hacking competition at DEFCON in Las Vegas in 2006 and 2007. His certifications include Certified Information Systems Security Professional and GIAC Certified Web Application Penetration Tester, Incident Handler, Firewall Analyst and Forensic Analyst. He is a member of the SANS Advisory Board and has spoken to numerous groups, including the Florida Department of Law Enforcement and Florida Association of Educational Data Systems (FAEDS), on network attacks, incident response and malware analysis.

He holds a bachelor of science in decision and information science from the University of Florida.

Executive Summary

It is not uncommon to find network monitoring devices and logging mechanisms set up, only to be abandoned and forgotten. They are casualties of the information overload that modern IT professionals face daily. When problems arise, be they network or security, someone has to dust off the log documentation, if it exists, and start digging in to figure out what's going on. By then, it is often too late; sensitive data is already in the hands of the attacker.

The complaint often heard in IT shops is that there are simply too many logs, so monitoring suffers. To make matter worse, in many cases, no one knows what they should be looking for or how their data could be useful to various groups, such as security, applications and network operations. The key to solving these problems often requires cooperation, since each group holds a piece of the puzzle; without collection, management and correlation, effective network monitoring is nearly impossible.

As attacks get more sophisticated, the different teams, so often used to working independently, will have to work together and share data to effectively protect their network. As a result, the network operations team will have a better understanding of the network and applications residing on it; the applications team will know how network problems can affect application performance, and security will benefit from being able to correlate data and detect attacks that would have gone undetected.

Sharing Common Information

Security is challenged to gain the cooperation to leverage existing network monitoring efforts because it is often seen as the bad guy by users and other teams that are working to make the network faster and keep applications running smoothly. Security is perceived as a business inhibitor, trying to keep an ironclad fist on the network with choke points for monitoring and firewalls to keep the bad things out.

It's a delicate balance of security and business functionality, but working with each other to share monitoring systems and implement projects that cross groups, such as network access control, make the teams more efficient so they can catch issues and solve problems.

Figure 1

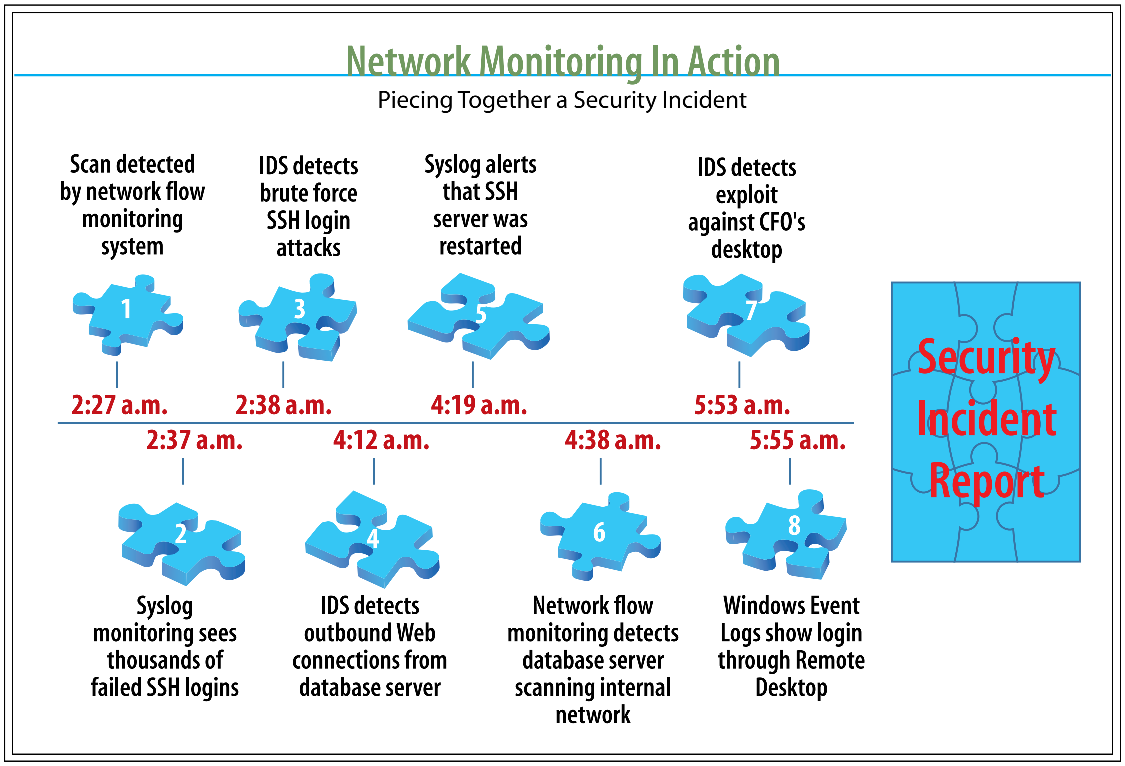
Network Monitoring In Action
Piecing Together a Security Incident
Scan detected by network flow monitoring system
1
2:27 a.m.
IDS detects brute force SSH login attacks
3
2:38 a.m.
Syslog alerts that SSH server was restarted
5
4:19 a.m.
IDS detects exploit against CFO's desktop
7
5:53 a.m.
2:37 a.m.
2
Syslog monitoring sees thousands of failed SSH logins
4:12 a.m.
4
IDS detects outbound Web connections from database server
4:38 a.m.
6
Network flow monitoring detects database server scanning internal network
5:55 a.m.
8
Windows Event Logs show login through Remote Desktop
Security Incident Report

Having all logs at your fingertips makes it easier to identify how the network scanning detected through network flow monitoring was followed by brute-force SSH attacks and the attacker downloading tools like a Trojanized SSH server. Additional data from the IDS and Windows Event Logs show the attacker inside the network exploiting the CFO's desktop and then logging in via Remote Desktop. Only having access to one or two sets of logs would make it hard to piece together the entire puzzle surrounding what happened leading up to the data breach.

Commercial network flow monitoring products can help by providing one location for network flow data, which is important for both networking and security teams. Seeing the opportunity to address network and security team needs, vendors started including multiple dashboards that let the network operations team drill down into performance metrics that indicate, for example, that a WAN link is saturated. Meanwhile, the security team has its own dashboard that indicates multiple malware infections at the remote site as the root cause of the network scanning saturating the link.

Centralize Logs

Leveraging existing network and application monitoring should be a priority for any organization to increase its security posture and detect security incidents before they become major breaches. The recent Verizon Business 2010 Data Breach Investigations Report *(http://www.verizonbusiness.com/resources/reports/rp_2010-data-breach-report_en_xg.pdf)* reminds us that there's a wealth of information contained in the logs, but it is rarely used properly. Verizon reports that it "consistently find that nearly 90% of the time logs are available but discovery via log analysis remains under 5%."

Logs are not going completely unnoticed. Network operations staff monitor router performance and SNMP traps to ensure the network is running smoothly. The question is, why aren't other groups doing the same so a security incident doesn't get missed? Simple: Too many logs and not enough time.

One way to help reduce the impact of having so many logs is to centralize them to one or two indexed, searchable locations. This gives analysts a fighting chance to spot patterns, compared with attempting to pore through dozens to hundreds of systems with their own logs. Centralizing logs also provide the opportunity to take advantage of power log management and security information and event management (SIEM) solutions, which we'll discuss later.

According to the Verizon report, more than 60% of 2009 breach investigations uncovered that attackers took several days or more (even months) to compromise data. The report goes on to say that "defenders can take heart that they will likely get more than one chance at detection. If real-time monitoring fails to sound an alarm, perhaps log analysis will be able to spot it." Unfortunately, Verizon's investigations show that organizations fail to take this chance to reuse the logs and find the breach through log analysis.

Detecting sophisticated attacks often requires correlation of events across multiple logs, including those from an IDS/IPS, network flow monitor, SSH server and antivirus. The difficulty in detecting these attacks as they happen is that different logs are the responsibility of different teams. By centralizing all of these logs into one place, the security team can, in a sense, recycle the logs from other teams into a security data source to get a better picture of what's going on across multiple systems and layers of the corporate infrastructure.

For example, the security team may have IDS alerts from the previous night that show SSH brute force attacks that stopped after four hours, but they don't know if the attacks were successful or not. Network flow data may show a 30-minute SSH session after the attacks during which several outbound FTP and HTTP connections were made so the attacker could download tools. Or, SSH server logs would show thousands of failed logins followed by one successful one. Without having all the logs at hand, the security team would have to guess what happened and contact the server admins for confirmation (see Figure 1,"Network Monitoring in Action," page 5).

Using a Hammer for More Than Nails

When applications are running sluggishly and Web sites are loading at a snail's pace, it's always the network's fault. For that very reason, network operations teams usually have all the tools they need to diagnose problems quickly and get things back on track. But these tools can also be used to supplement the security team's efforts to keep the corporate network and data secure.

We touched on network flow data tools above as one area where we're already seeing a convergence in which one tool can be used by network and security operations teams. There are other valuable network data sources that can be used by security to proactively catch and stop attacks before they get out of hand. We'll examine some of those areas and how they can be extended to help secure the network.

Network flow data is essentially network traffic metadata. It describes the network traffic using time stamps, the protocol used, source and destination IP addresses and ports, amount of data transmitted, and similar information. The difference between it and a full packet capture is that network flow data does not include the contents of the packet. As an analogy, think of a phone bill that includes the phone numbers and call length but not content of the conversation. Don't be fooled by the lack of content; a lot can be learned from network flow monitoring, and it's become an incredibly useful tool for security teams.

Network flow data is produced by network routers and some Layer 3 switches. It can come in several forms, such as NetFlow and sFlow, depending on the vendor, but they all typically provide the same basic details. The simplest way to take advantage of network flow data is by setting up a Linux-based system, install flow tools and configure the routers to export flow data to the system. For more information on flow tools and their usage, No Starch Press published "Network Flow Analysis" in June (*http://nostarch.com/networkflow.htm*), which details ways to use flow tools for network and security monitoring.

Two of the simplest uses of network flow data by security teams is detection of network scanning and potentially infected hosts. Network scanning is easy to spot because no traffic content is needed for detection. If a host attempts a certain number of connections to another host or series of hosts within a certain timeframe, then it is likely to be scanning. Or, for example, if hosts within the corporate network are connecting to external hosts that are listed as known Zeus malware hosts in the Zeus Tracker (*https://zeustracker.abuse.ch/*) and downloading data, those internal hosts are likely to be infected, and a first responder should be sent to determine whether they are infected.

Vendors like Lancope, Riverbed Technology and Arbor Networks have developed advanced network behavior anomaly (NBA) systems that use network flow data to not only monitor network performance but to also baseline hosts' behavior and alert on anomalous activity. For example, using the baseline, if a host that has traditionally been a network client begins acting as a server, an alert will be sent. Or, if an IP that's never been used begins communicating on the network, administrators can be notified. Likewise, similar alerts can be configured to detect scanning, excessive bandwidth and traffic seen outside of normal operating hours.

In addition, a number of free and open source network flow data monitoring tools have been released over the years. Two of the more popular are argus and flow tools. While they require more manual analysis and do not include behavior-based detection capabilities like commercial NBA solutions, the free network flow monitoring tools are very powerful and can help quickly pinpoint suspicious connections.

SNMP from network devices contains a variety of useful data from informational alerts. For example, you may see that a device was just connected to a network port to high CPU utilization in a router. From a networking perspective, those messages may not be interesting, but to security it may mean a rogue device was connected to an unused network port in an office supply room or an attack is causing the CPU usage to spike.

SNMP traps can provide informational alerts to show when MAC addresses have changed on a network port or when more than one MAC address is on a port. MAC change messages could indicate a rogue device has been placed on the network in place of the original device or that a network hub or switch has been plugged in and a rogue device is now connected alongside the original device.

Having detailed, accurate network and physical infrastructure documentation is important in being able to act on information from these systems in a timely manner before an attack has the chance to successfully compromise data. For example, if an attacker plugs a wireless access point into a network port in a conference room or a multifunction copy machine is compromised and used to scan the internal network, it is critical to have diagrams indicating where those physical network ports are located, so security can grab the system immediately for analysis.

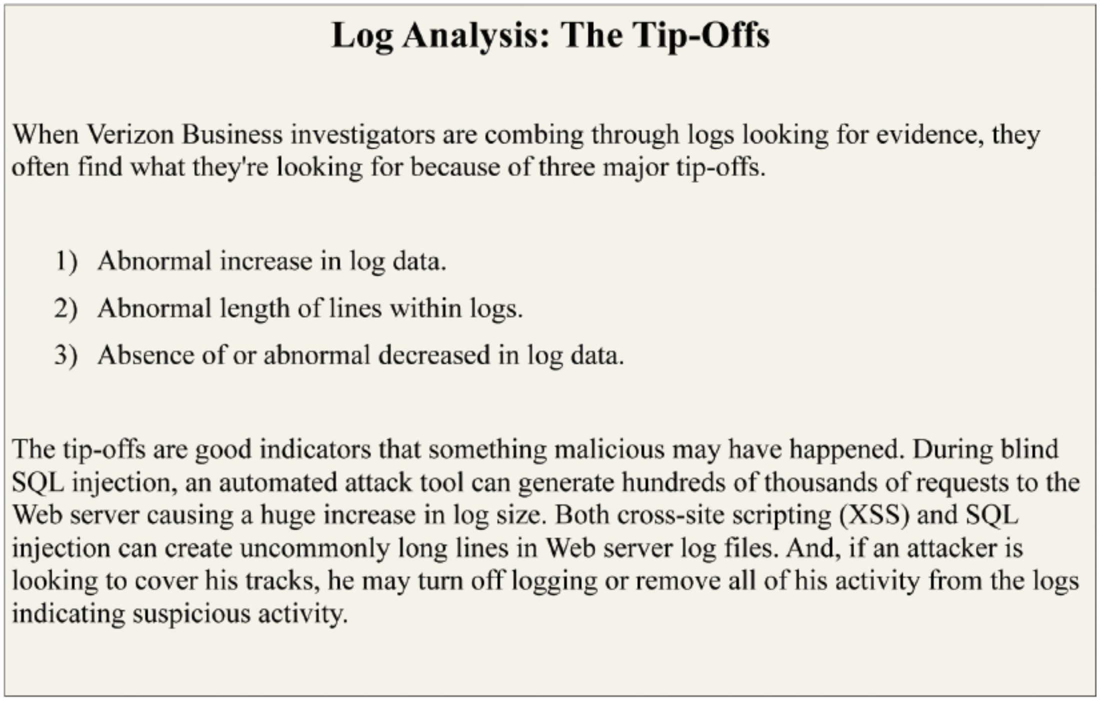

Log Analysis: The Tip-Offs

When Verizon Business investigators are combing through logs looking for evidence, they often find what they're looking for because of three major tip-offs.

1) Abnormal increase in log data.
2) Abnormal length of lines within logs.
3) Absence of or abnormal decreased in log data.

The tip-offs are good indicators that something malicious may have happened. During blind SQL injection, an automated attack tool can generate hundreds of thousands of requests to the Web server causing a huge increase in log size. Both cross-site scripting (XSS) and SQL injection can create uncommonly long lines in Web server log files. And, if an attacker is looking to cover his tracks, he may turn off logging or remove all of his activity from the logs indicating suspicious activity.

Source: Verizon Business 2010 Data Breach Investigations Report

Syslog has been the standard remote logging mechanism for Unix-based systems for many years. By having a central syslog server, organizations have been able to centralize logs on Unix-based systems for years while Microsoft IT shops were left in the dark. Free and commercial tools were eventually released that allow Windows system administrators to ship their proprietary Windows Event logs to a centralized syslog server. Some examples include Snare, Lasso and EventReporter.

One of the great benefits of syslog is that it has been around for many years, so there are several tools designed to monitor syslog to generate real-time alerts and provide customized reporting. Modern log management and SIEM solutions typically leverage syslog to centralize logs for analysis, although some have custom agents that use proprietary protocols. Splunk is an example of a popular solution for indexing and searching through centralized syslog data from systems of all types.

Firewalls and their logs are extremely valuable. Some organizations rely on the network engineers to manage and monitor the firewall, while others place that role on the security team. No matter who manages them, the logs from firewalls can give network operations insight into what may be causing performance issues, and security teams may use the same data to detect scans and attempted attacks.

Network operations is likely to have found on more than one occasion that a performance problem was related to some poorly maintained, or misunderstood, firewall rule. During the course of investigating the problem, it's often found that a firewall rule had an adverse effect on production traffic because the traffic behaved differently than documented by the vendor.

For security, it's easy to detect scans, as the attacker generates large numbers of denied log entries as she scans hosts and networks looking for live hosts with accessible ports.

Next-generation firewalls like those available from Palo Alto have evolved from the classic static filtering, stateful filtering and deep packet inspection firewalls to have a clearer understanding of the applications traversing them. They are capable of identifying the applications and filtering based on rules regarding application content such as instant messaging running on TCP port 80 and SSH traffic disguised as HTTPS traffic.

The advanced nature of next-generation firewalls gives the network and security teams a better picture of the traffic traversing in and out of their network border, which is why some enterprises are leveraging them as data leakage prevention (DLP) devices. Since they are already inspecting all the traffic going out, it's a natural next step to use them to monitor for Social Security numbers, credit card numbers or other sensitive data that should not be leaving the corporate network.

DHCP servers and their logs do not necessarily seem particularly important in terms of network security, but monitoring the logs can help detect subtle changes that could indicate an attack. For example, when a host connects to the network, the DHCP server will log the name of the system being given an IP address. A script could be written to record the mappings of all hostnames with their respective MAC address in a database and do a comparison every time a new DHCP request was made.

When a hostname or MAC address changes, it could indicate a rogue device is impersonating the authorized device. Or, if the name comes up as Backtrack, a malicious user may have booted his workstation with the popular penetration testing bootable CD.

It's also possible that the DHCP server could be used to block malicious systems by setting short DHCP lease times and handing out unroutable IPs and nonexistent DNS servers to the malicious system when it checks in to refresh its lease. However, this is only particularly useful in environments where most of the switches are not managed, making it difficult to pinpoint the location of the malicious system. It is not effective against a skilled, determined attacker and requires correlation from another system (e.g., IDS or NBA) to identify a host as malicious.

Network access control (NAC) is another network system that could be managed by either network operations or security, but is likely to be co-managed. The obvious value to security is it requires authentication of users prior to their systems being admitted into the corporate network. Based on user and groups, the computers can even be put on specific network segments specific to their role in the organization.

NAC also provides security teams the ability to perform endpoint posture assessments on computers before they are allowed access to the network. Computers can be blocked, quarantined or given limited network access if they do not meet certain criteria, such as having the latest patches installed, updated antivirus or banned software, such as peer-to-peer applications.

Because computers can fall out of compliance after they've passed the posture assessment and are on the network, many NAC systems continuously monitor the hosts. In the event of a computer becoming noncompliant or compromised, NAC can automatically quarantine the computer to prevent, or at least, reduce damage to the rest of the network.

Network taps and mirror ports allow all traffic or a portion of the traffic traversing a network segment or switch to be duplicated for capture and/or analysis. Network operations staff will often have a computer system connected to a network tap or mirror port to capture traffic on demand to help troubleshoot problems. The network operations systems usually do not capture traffic all the time and are only used as problems arise.

Security teams use the same taps and mirror ports to connect their IDS and/or DLP tools to monitor for malicious traffic at all times.

Additionally, network and security teams may use a network recorder such as those from NetWitness or Network Instruments to capture all traffic to disk to be used later for network performance analysis and security investigations.

Network recorders capture everything, so it's easy to step back and see what activity occurred at a specific time. Implementations often include an interface for network and security teams to extract files from Web, e-mail and VoIP sessions. Network recorders are very powerful but come with hefty hardware and storage requirements to keep up with busy networks.

Bringing It All Together

All of the different network monitoring tools and logs are useless if the output is not collected and reviewed. Doing so is difficult due to the sheer volume of data, making log management tools and security information and event management (SIEM) systems a compelling option, as compared with the countless man-hours if attempted manually.

Log management solutions focus on aggregation of logs from disparate sources across the enterprise into one location. Once collected, the logs are indexed and go through a life cycle where the data may be archived to a less expensive storage media and eventually deleted. The indexed logs are searchable, and real-time monitoring is achieved by setting up queries for events of interests that generate alerts when triggered. There are often little to no correlation capabilities in standard log management tools.

SIEM products are notoriously known as big, clunky and difficult to implement because of the level of integration needed within the enterprise in order to make them useful. But, once up and running, they can be invaluable from a real-time alerting standpoint and the ability to correlate multiple data sources into actionable data that's linked by time, user, IP, country, etc.

Security data visualization is a topic of much interest over the past several years with a couple of books dedicated to the subject. In this area, SIEM products excel because they provide rich dashboards to help security teams pinpoint patterns that may only be seen by a human eye.

Mi Herramientas Estan Su Herramientas (My Tools Are Your Tools)

By leveraging existing network monitoring tools, security can become more efficient and have much better visibility into the network with little to no additional cost. The examples above paint the picture of how both teams can use the same data while achieving different goals. Often, a little glue is needed to bring it all together, but that's where log management and SIEMs rise to the occasion.

Can security do everything they need using the network operations' tools? No, but combined with their current toolset, network monitoring can be provide security with the advantage it needs to get the job done faster and more effectively.