RSA White Paper

SIEM: Moving Beyond Compliance

Dr. Anton Chuvakin @
Security Warrior Consulting

Executive Summary

This whitepaper covers strategies and tips on effectively using SIEM and log management tools beyond regulatory compliance. Many organizations acquire SIEM solutions for PCI DSS and then slowly start using the tools for other security and operational concerns. This paper will help jumpstart this process and highlight common SIEM usage scenarios for organizations of all sizes.

It will also explain how to operationalize the SIEM tool and utilize it for many security use cases and scenarios, from Web site threats to security incident response. Specific examples from RSA's enVision platform are used to illustrate the concepts in the paper.

RSA®
The Security Division of EMC

Contents

Introduction

Security information and event management tools first appeared on the market around 1997. Their original purpose was for reducing network intrusion detection system (IDS) "false positives" that plagued IDS systems of that long-gone era of security. The tools were only used by the largest organizations with the most mature security programs and often with 24/7 Security Operations Centers. Over the years, the critical requirements for SIEM have evolved to compliance, user tracking, application monitoring and even fraud detection. The tools have matured significantly and are close to becoming usable by a wider range of smaller organizations. In addition, dedicated log management tools have emerged to address broad log retention and log review requirement across IT, beyond the traditional security space.

Security, compliance and IT operations are three drivers for SIEM

Over the years, the following areas where SIEM and log management tools can deliver value have emerged.

Security, detective, and investigative: sometimes also called threat management, this focuses on detecting and responding to attacks, malware infection, data theft and other security issues.

Compliance, regulatory (global) and policy (local): this focuses on satisfying the requirement of various laws, mandates and frameworks.

Operational, system and network troubleshooting and normal operations: specific mostly to log management, this use case has to do with investigating system problems as well as monitoring the availability of systems and applications.

SIEM for compliance

Recent research indicates that up to 70 or 80% of SIEM deployments are driven by PCI DSS or other regulations. The following table shows a few example regulations that affect SIEM and log management.

Regulation	SIEM and Logging Relevance
PCI DSS	The Payment Card Industry Data Security Standard (PCI DSS) applies to all organizations that handle credit card transactions. PCI mandates logging specific details and log review procedures to prevent credit card fraud within companies that store, process or transmit credit card data.
ISO27001	ISO27001 is a direct descendant of ISO17799 and British Standard 7799. ISO specifies requirements for managing the security of information systems. Audit logging and review of audit logs as well as their retentions are prescribed.
NERC	North American Electric Reliability Council (NERC) publishes Critical Infrastructure Protection (CIP) standards that contains important information security requirements. These standards affect utility companies in U.S. and Canada. Among them are requirements about logging, alerting, log review as well as broader security monitoring.
US State Data Breach and Data Protection Laws	CA SB 1386 started the trend of data breach disclosure laws in 2002. Since that time similar laws have spread to 44 of the states and a few countries as well. While not prescribing logging directly, the provisions to notify those whose confidential information has been stolen leads to access auditing and granular data logging requirements.
HIPAA/HITECH	The Health Insurance Portability and Accountability Act of 1996 (HIPAA) outlines relevant security standards for health information. NIST HIPAA documents detailed log management requirements for the securing of electronic protected health information such as the need for regular review of information system activity, such as audit logs, access reports and security incident-tracking reports.

Path beyond compliance

The easiest way to expand the use of log management or SIEM tools beyond compliance is to actually start using them for compliance, but using them well. Based on this and other examples from the author's recent consulting practice, we can formulate the following success criteria for moving beyond compliance.

First, the path to effective operational use of SIEM tools starts from operationalizing compliance practices. Few people remember that PCI DSS prescribes a large set of periodic tasks, from annual to daily (log review being the most well-known example of a daily practice).

Second, an incident response capability must exist – the personnel operating the SIEM tool should know what to do if a high risk alert is triggered. This is due to the fact that the easiest and most common security use for log management and SIEM tools is related to incident response and forensics.

Third, a certain degree of security practice maturity has to exist if an organization falls under the mistaken perception that buying the tool is enough to make them compliant, the tool likely will become "shelf-ware". SIEM operators have to follow a particular workflow to accomplish their goals.

Fourth, the concept of monitoring – whether for regular availability or threats – should exist. Simply buying a tool that is capable of enabling such monitoring does not create a monitoring capability. Such capability combines skilled personnel and effective SIEM tools. Fortunately, most organizations have monitoring tools for operational visibility – uptime monitoring. Full Security Operations Center (SOC) is not required; however, the organization must have or start to build security monitoring capabilities such as dedicating a person or team to ongoing periodic security monitoring.

Fifth, an organization must be able to integrate data sources as well as asset data sources into their SIEM tool. This will enables them to review alerts and then respond to them in the context of their organization. Feeding the SIEM tool with logs, vulnerability scan data, asset information, and security configuration management information will enable it to perform its mission with high efficiency and thus solve more business problems. The organization must also accept the responsibility for tuning and customizing their deployed SIEM tool.

Know your key SIEM requirements and use cases before deploying

Detailed Common Use Cases

While it is desirable for the organization to come up with their own requirements for a SIEM and their own use cases, we can try to help by outlining the most common SIEM use cases that are addressed by today's SIEM tools and that are successfully implemented at many organizations.

Our discussion of common use cases is structured as follows:

- **Use case name and description:** what the use case is and what business and security problems are solved by using a SIEM tool in this manner.
- **Collection:** what logs need to be collected in order to be able to address this use case, and what other context information – such as vulnerability data – is needed to successfully solve this problem. Log collection methods are also discussed.
- **Reporting and dashboards:** how the collected data will be presented and summarized via reports and monitoring dashboards, what reports need to be created and run – and how often and by whom they need to be reviewed.
- **Correlation and alerting:** what correlation rules must be defined, tuned and enabled in order to solve the problems, who should receive the alerts and what they should do about them.
- **What else:** what processes and procedures need to be in place for successful SIEM implementation for this purpose.

This framework should allow us to build practical guidance and tips for using a SIEM for the following scenarios.

- Server user activity monitoring
- Tracking user actions across disparate systems
- Comprehensive firewall monitoring
- Malware protection
- Web server attack detection
- Incident response enablement
- Everything together: SOC operation

Let's proceed to the detailed use case review.

Server user activity monitoring

Organizations that deploy thousands of servers with various operating systems, such as Linux, Solaris, or Windows have a challenge tracking who is logging in to all those servers. While centrally collecting all the login and other authentication logs from thousands of servers presents a challenge, intelligently analyzing all the authentication data is even more difficult.

Typically, a company would like to know whether people who are accessing the servers are doing it legitimately and with business purpose in mind. Also, organizations would like to know whether anybody is trying to compromise a server by trying multiple usernames and passwords, possibly in an automated fashion.

Being able to know that access by a particular user to a particular server is suspicious or malicious allows companies to detect possible hacking and insider abuse incidents at an early stage, before most of the damage is done.

On top of this, multiple regulations prescribe login monitoring to servers with sensitive data such as health records or payment card information. This is why the server monitoring use case is often one of the first to be implemented after purchasing a SIEM tool.

Collection: in order to use a SIEM tool for server access monitoring, logs from all servers of all platforms that include authentication records need to be collected. It is very important to collect both successful and failed access attempts. Only logging or only collecting failed access attempts will completely undermine this monitoring effort!

While no additional asset and context information is needed for the successful implementation of this use case, the following types of information would be extremely helpful while prioritizing responses to detected events:

- Server function and importance
- User identity and role within the organization

Below is an example of a successful Unix login message:

Mar 13 16:26:09 combo sshd[8714]: Accepted password for anton from 10.120.2.133 port 57019 ssh2

Here is an example Windows login message from Windows 7:

Event ID 4624 The event is generated when a logon session is created.

It should be noted that Windows contains multiple types of login messages and only select types can be obtained from domain controllers or active directory servers; local login types may be tracked by collecting logs from individual servers.

Log collection methods: server log collection methods differ dramatically between platforms.

Unix and Linux syslog are the easiest to collect and centralize. Since there is no need to configure anything to make sure that user access attempts are logged (on by default), the only configuration change needed is to make sure that such log entries are sent to a SIEM server.

Older Windows platform offer a plethora of choices. One can use an agent to convert Windows event logs into syslog which is then sent to a SIEM. Another option is remote collection of Windows event logs using Windows APIs directly, which it has its own challenges. Only modern Windows versions such as Vista and Windows 2008 have native XML-based log centralization options that can be utilized for wide scale log collection.

Mainframe and midrange servers that also record user authentication events present an additional challenge. However, in many cases collection can be resolved by using a syslog or text converter as well that will render SMF into readable text.

Reporting and dashboards: collected data can be used for security, operational and compliance purposes as well. Here are some example reports that help to serve security purpose:

First, a simple authentication failures report can be used to quickly check who is trying to get to various Unix servers:

A similar report can be run across other types or servers or all of them.

On the compliance side, PCI DSS mandates access monitoring. “Successful Connection” report that shows all successfully authenticated sessions is useful to jumpstart your server access monitoring efforts by focusing on the actions of privileges users.

Operationally, knowing which users access what servers is useful for defining access policies and possibly provisioning additional servers if needed.

If periodic reporting is not sufficient and near-real-time monitoring is required, the dashboard “Top Failed Login Accounts” can be used to track user authentication events across servers.

However, the most value in this use case can be obtained from automated rule-based correlation of login success and failure events.

Correlation and alerting: a rule can be used to automate nearly the entire server access monitoring process. One of the most useful correlation rules is the following:

- 10 login failures on any server in 1 minute
- followed by
- successful login within 1 minute to the same server

Another useful rule is the one that tracks increases of failed logins over a daily baseline. A large increase over an accumulated baseline is almost always worth investigating.

These correlation rules might need to be tuned for the environment in order to only produce alerts in cases of real security incidents. Here is how this rule can be tuned:

- If the rule will fire too frequently in case of legitimate server access, increase the count of failed login events
- If you would like to make the rule more sensitive to slow password guessing, change the timeframe for the rule.

The alerts can then be reliably triggered without producing "false positives."

What else: in addition to deploying SIEM technology, collecting logs, running reports and using correlation to trigger alerts, operational procedures need to be in place to have an effective server access monitoring process.

These should include:

- Periodically reviewing all server access successes and failures using reports. PCI DSS mandates daily log review; in other cases weekly server access log review is sufficient.
- Notifying the server administrators in case of an abuse or attack. Multiple mechanisms such as email, SMS, or SNMP traps are available for notification.
- Notifying the business owner in case of a serious risk of server downtime or compromise.
- An incident response process should be in place in order to define what happens if a server is compromised by the attackers.

These processes need to be put in place by the security team, but other parties such as system administrators need to be involved as well.

Tracking user actions across disparate systems

Security incident response, compliance as well as Human Resources (HR) requirements call for investigating user activities across multiple information systems. Log management and SIEM tools are ideal for that since they contain traces of user behavior across possibly every system in the organization. Recently, investigation of insider fraud cases has increased the need for efficient, quick and comprehensive user activity investigation across servers, network access devices and applications. In addition to this, individual user activity monitoring can be used when suspicion exists that the user is "up to no good."

Collection: for this use case, the collection effort covers most every log source that records user name information. That will exclude some network devices – such as routers and firewalls – but will include most every piece of information technology deployed at an organization.

Collection methods: log collection methods vary dramatically for this use case. Starting from syslog across Unix and Linux, WMI or syslog agent for Windows, database table pool across databases, file logs for Web proxies and VPNs, and file download for many application logs, collection of all user activity logs presents a challenge for SIEM implementers.

One approach to solve this challenge is to prioritize your collection efforts based on ease of gathering the data and priority. For example, Unix syslog will be the easiest to collect while SAP application logs are more difficult, yet extremely valuable.

Reporting and dashboards: reports that show user activities across multiple systems might present the username time and date as well as the nature of activity. Also useful are summary counts for each activity.

If looking for a particular user, information review starts from entering the username into the report condition or search filter and tuning the time that report covers.

In many cases, the approach will be to run a report across 24 hrs of user activity logs and then expand coverage to weeks or even months.

If near real time monitoring of user activity is desired the dashboard "PCI Windows Failed Logins" can be useful. A customized version with select users can be created during the incident response process in order to watch the suspicious users.

This allows security analysts to review everything that this particular user does as it happens and allows to take action over anomalous or malicious activities.

Correlation and alerting: in most cases, user activity tracking is performed over historical data, collected over days, weeks or months of activity. Correlation rules are of limited use for this scenario. However one can use simple filtering rules to alert in case a particular user commits a particular activity or transgression. For example, a security administrator might want to be alerted if a user he is monitoring suddenly starts to upload large amounts of data outside the company.

This can be accomplished by the following rule:

- If user on watch list and destination=external and size of file transfer=large, send an alert

In particular this rule allows detection of suspicious file transfers by watched users. It might indicate that the user is trying to send intellectual property or other regulated or valuable data outside the company. It requires collection of Web proxy logs or logs from a firewall that records file transfer sizes and types.

Other user monitoring and alerting rules include:

- Watched user connects to critical, sensitive or regulated server
- Watched user creates new user accounts on servers
- Watched user performs financial operations with high amounts (if financial application logs are collected)

Thus, correlation rules are very useful for watching for user activities after initial suspicions are established and the users are added to a watch list.

What else: while performing the user monitoring and investigation as well as live detection of user activities, a policy that allows such action is absolutely mandatory. In addition to policy, the organization must document user investigation and monitoring procedures that need to be followed in case of suspicions.

In many countries, such policies will be affected by privacy regulations and other laws that protect employees from unreasonable snooping. For example, the policy might state that "user activity logging might be increased if suspicions exist that the user is violating corporate policy or relevant laws."

Another critical success component that is required for the effective use of SIEM for user monitoring is having a response policy and process. Specifically, if a user is detected transferring corporate data outside or performing other illegal acts or violations, HR action needs to be part of the plan.

Comprehensive firewall monitoring (security + network)

Since the early days of SIEM technology, firewall log data has been considered as one of the most useful and commonly collected information sources.

Apart from allowing and denying connections to and from the network, firewalls allow recording or logging of every single connection denied or allowed by the firewall. An example would be connections from the outside world to the DMZ Web server, or connections by users inside the company to their favorite social media Web site.

Analysis of such logs is extremely useful for security, compliance and even operational purposes such as network management, bandwidth management, etc. For example, on the compliance side, PCI DSS, HIPAA, NERC/FERC all have firewall logging implications. Firewall logs are also extremely useful for incident response and forensics since they can help identify the connectivity pattern and serve as "poor man netflow." On top of this, firewall logs can be used to assess the health of the firewall itself and to optimize the ruleset performance.

Collection: comprehensive firewall log collection is mandatory for this use case, and it is important to remember that firewalls can record both failed and successful connections through the firewall – both types are essential for SIEM.

Some firewalls log two messages per each allowed network connection: the first when it is initiated and the second when it is terminated. The latter message often contains connection duration and the number of bytes transferred.

Traffic logs from a firewall are further subdivided into inbound and outbound messages – the former are connections from the Internet to the company while the latter are connections from the company systems to outside. While in the past inbound attempts were seen as valuable firewall data, today it is more useful to collect outbound connection messages since these can be used for detecting malware infections inside the company. It should be noted that the separation into inbound and outbound logs is overly simplistic: most organizations operate firewalls with multiple segments and multiple network interfaces.

In addition, most firewalls will record administrative and other actions on the firewall itself; these logs also need to be collected for analysis. Non-traffic logs cover various firewall performance and administration messages, access to the firewall system itself, as well as logs from other components of a multi-function firewall device.

Additional context information which is useful for firewall monitoring is information that maps internal IP addresses to asset information. This allows security personnel to not only identify the offending IP address but also the function and the ownership of the system that initiated suspicious traffic or launched an attack against third parties.

For external systems, the ability to resolve the DNS name and look up a WHOIS record is useful for assessing the impact of a possible incident.

Collection methods: most popular firewall types such as Cisco and Juniper use syslog to transfer messages from a firewall to SIEM. Collecting those logs is as simple as configuring all firewalls to send their messages to the SIEM collector. Some firewall logging settings need to be adjusted in order to record connections established through the firewall and not only blocked connection attempts. Syslog is also used by other firewall manufacturers common in smaller businesses.

Other firewalls, such as Checkpoint, use proprietary logging mechanisms. Typically a SIEM product will have a collector to pull the logs from individual firewalls or from firewall management server. Such periodic pulling is not real time but typically is often good enough for most firewall log analysis.

Reporting and dashboards: rich firewall log data set, including both traffic and non-traffic logs, as well as a wide range of uses for firewall data allows for many visual representations and reports.

The following report is useful for reviewing connectivity denial patterns across the firewall:

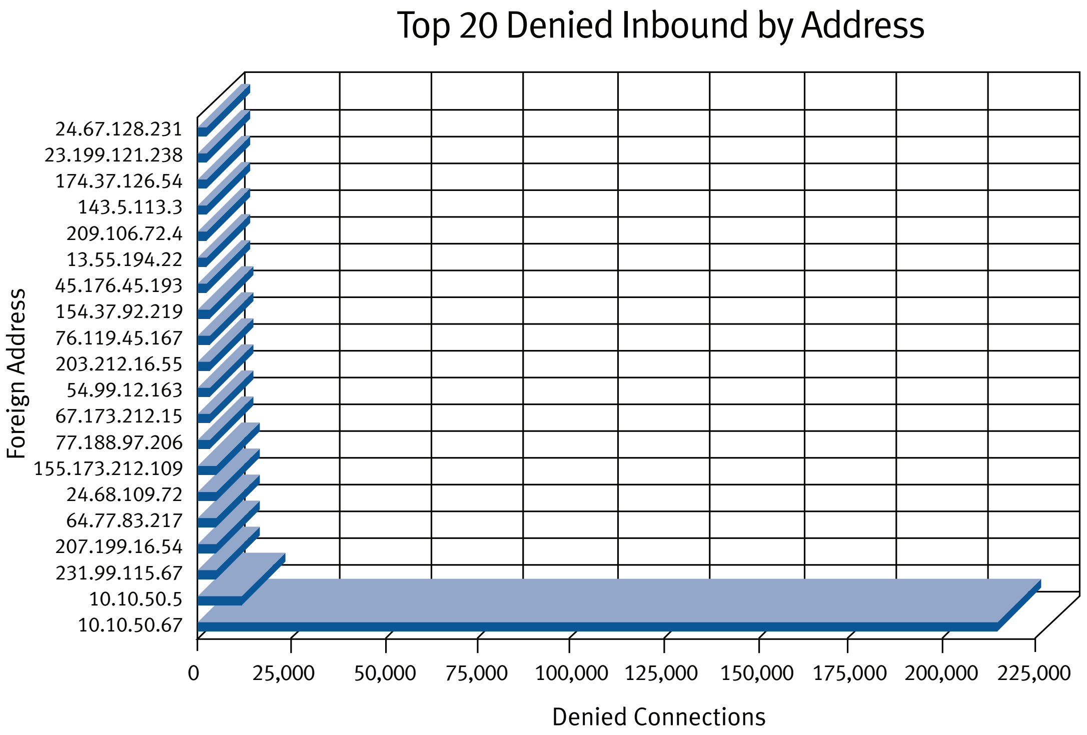

This graph displays the top 20 foreign addresses that were denied inbound access

Time range: Wed Jun 16 13:15:04 EDT 2010 to Wed Jun 16 17:15:04 EDT 2010

Reviewing successful access from internal systems towards the internet is even more useful than reviewing denied access from external system to inside. For example, “Outbound HTTP Traffic” report can be used to see all such access over port 80.

These reports are useful for tracking suspicious internal systems that generate too much traffic across the firewall.

The latter report, if limited to internal systems, can be used to track not only legitimate bandwidth use, but also malware trying to connect to outside from the company environment.

Finally, these reports are useful for tracking connectivity to and from regulated environment such as PCI DSS payment networks. “All Outbound Traffic from Regulated Environment” and “All Traffic Allowed Into Regulated Environment” are examples of such reports.

Overall, looking through multiple firewall reports allows firewall administrators to solve both security issues as well as connectivity flaws and other network problems.

Correlation and alerting: firewall data is useful for correlating with other types of data (such as network IDS or IPS) and also for rules that make use of data from multiple firewalls.

One of the most useful correlation rules detects an internal system that attempts to connect to many systems outside the company, indicating possible malware:

– If internal system connects OR tries to connect to outside systems at 10x normal volume, trigger an alert

Baselining firewall connectivity also renders useful results such as knowing when internal systems start misbehaving.

Finally, while looking for port scans from the internet is difficult, a rule that shows that the scan was followed by a successful established connection to your systems is still fairly useful.

Alerts from such rules should be going to security analysts and firewall administrators.

What else: as with other use cases we discussed, correlating firewall log data calls for a response process in case a compromised system is discovered. Similarly, if analysis of firewall logs reveals that a system is misconfigured, a notification process for the administrator and possibly business owner of that system is essential.

Malware protection

It is well-known that signature based antivirus technologies are losing their efficiency as a primary weapon in the war against malicious software. Detection and clean rates have been dropping dramatically over the last few years.

To detect modern commercial malware, desktop and gateway anti-virus tools need to be reinforced with network traffic analysis and log analysis. In addition, scenarios where anti-virus technology detects the threat but is unable to delete it are not uncommon. Using SIEM for detecting and highlighting such situations is within the capabilities of most organizations.

Another useful scenario for malicious software analysis using SIEM presents itself during a major malware outbreak. In this case, using correlation technology allows organizations to track which systems are infected and spread malware. Detecting systems that attempt to connect to other systems in order to spread malware presents one of the effective ways to curb the outbreak.

Finally, as botnets and other modern commercial malware become even bigger threats, SIEM presents the way to analyze diverse sources of information thus making it possible to detect advanced malicious software missed by antivirus solutions.

Collection: in order to address this group of use cases aimed at detecting malicious software, multiple types of devices should be logging into a SIEM. First, antivirus and other anti-malware logs must be collected in order to analyze events where antivirus protection fails or is disabled by malicious software. Antivirus log analysis also enables the security analysts to know when a virus is detected but not cleaned. On top of this, multiple compliance regulations prescribe regular antivirus updates. SIEM alerting allows organizations to monitor for failed anti-malware updates and take rapid action in order to restore security defenses and compliance status. PCI DSS mandates generation and collection of antivirus logs.

Firewall logs present another useful information source for detecting malware. Specifically, such logs are useful for detecting connection attempts by infected systems to their command and control (C&C) botnets.

IPS logs are also useful since modern intrusion prevention systems have signatures for detecting network behavior characteristics of malware.

Additional data useful for detecting malware consists of vulnerability data and asset data. Vulnerability data is used to qualify systems that might be infected by a malware. What is more important, while addressing malicious software incidents, information about infected system owners, systems business roles and other asset parameters is of high value and should be integrated into a SIEM.

Collection methods: given a breadth of log sources, collection methods vary. Most antivirus tools log into a windows event log or into their own proprietary logging mechanisms, typically text files. Firewall and intrusion prevention systems might be using syslog which makes collection of such logs simple and effective.

Reporting and dashboards: antivirus log data, firewall and IDS/IPS log data related to infections as well as anti-malware safeguard failures can be represented in multiple reports.

Some of the useful reports are “Top infected systems” that help prioritize which infected systems to clean first. Also, a useful report to run is the one to show all systems with clean-only events. It helps reveal systems that are currently infected by malicious software which cannot be cleaned by existing antivirus defenses.

Similarly, a report on systems with anti-virus failure and anti-virus update failures is a critical report that highlights systems where malicious software managed to take down antivirus tools. For example, “Anti-Virus Update Procedures” is based on PCI DSS Requirement 5, which mandates that anti-malware updates are operational and also capable of producing logs.

These reports should be executed at daily to weekly intervals. Reports related to live infections should be run daily while reports indicating antivirus failures and signature update failures should be run weekly or, preferably, daily as well.

It is advisable to use real time dashboards to monitor for infected internal systems.

Correlation and alerting: ideally, organizations should be equipped to deal with infected systems immediately. Thus defining alerts on select antivirus logs malware network activity is critical.

All anti-virus logs that indicate that malware was detected but not cleaned need to trigger real time alerts sent to system owners and security administrators. Infections may be contained before further damage is done by malicious software.

Repeated anti-virus failures, especially across multiple systems, also need to trigger alerts in real time. Such failures indicate a spreading infection by malicious software capable of terminating the antivirus process. It is also possible that a critical bug in antivirus software is disabling the protection – which also needs to be brought up to the security team’s attention as soon as possible.

Events indicating attempts to send virus and other malware from inside the company network to the outside also call for immediate action. Define correlating rules such as the following:

- internal system attempts to connect to many external systems
- internal system generates a large volume of connection attempts
- internal system attempts to connect to systems on the malware blacklist

The latter alert needs to be acted on immediately due to its high reliability.

Overall, modern commercial malware presents a significant risk to organization’s regulated data, sensitive information and other IT resources. Real time response is highly desirable (even if difficult due to the stealth properties of said malware) and SIEM correlation enables automation of such response.

What else: response to malware events often requires collaboration between security analysts, network administrators and system or desktop administrators. While security tools can detect malicious software, it is often required that a desktop management team perform cleaning and other operations on desktops. Alerts of critical malware events should be going to both security and desktop administrators. In case of critical infections –it also makes sense to notify the business owner of the infected system especially if there is a chance that the system needs to be taken offline.

Web server attack detection

Web application attacks have increased in recent years by a huge margin. Research indicates that a majority of Fortune 1000 companies' Websites are vulnerable to various Web application attacks, from cross site scripting to SQL injection as well as business logic flaws. Also it was discovered that the vast majority of credit card data theft occurs through Web attacks, at least as one of the stages that leads to data theft.

What makes Web application attacks so prominent is a multitude of factors. Web servers are always exposed to the Internet in order to engage in e-commerce and partner transactions. Many Web applications – including those that handle regulated data – are written by companies internally or by outsourced developers. This prevents an organization from patching it when the vulnerability is discovered since development of such security patch requires cooperation from the application developers.

In light of this, Web site security monitoring and reporting presents a critical requirement that is also increasing in importance. SIEM allows organizations to collect and analyze Web server logs in order to detect possible Web site compromise, thus saving the company from direct losses and embarrassment.

Collection: Web server logs present the primary information source about Web application attacks. It must be noted that all types of Web server logs need to be collected – this commonly means collecting both access and error logs. Ideally, middleware application server and back and database server logs also need to be collected.

For a Microsoft IIS Web server, the Windows event log must also be collected and filtered for log messages to be enabled. This is due to the fact that critical server operation messages are logged to the Windows event log, while legitimate access to the Web site is recorded in dedicated plain text logs.

If an organization deploys a Web application firewalls (WAF), its logs are also very useful and need to be collected.

Critical context information for Web attacks is the result of Web vulnerability scan data. Dedicated Web vulnerability scanners can detect issues with custom applications as well as Web application platforms. Such data can be correlated with Web application firewall logs in order to provide reliable and effective attack detection.

Finally, security use of Web server logs is overshadowed by the operational and business use of Web logs. For example, the commerce Web server logs are frequently analyzed in order to determine customer behavior and make the Web site more accessible to customers. Such use cases are outside of the scope of this paper.

Collection methods: most Web server logs can be collected as plain text files stored on the Web server. As mentioned above, windows event logs also need to be collected. Web application firewall logs can be in syslog or plain text format and need to be collected accordingly.

In many cases the following architectural challenge presents itself. Web servers are deployed in the DMZ, on the public network or at the outsourced, hosted location – while SIEM is deployed on the internal network. The challenge is in moving logs from DMZ or other public network to the internal network. Since logs are often collected by file transfers, such as SCP, direct access from the internal network to the Web server is required. It is not recommended to store Web server logs in public Web server directories since such logs may occasionally contain passwords and other sensitive data.

Reporting and dashboards: Web server access patterns need to be reviewed for normalized activity using reports. Similarly, Web server error logs must be looked at in order to determine unusual errors, Web application failures and malicious access attempts.

Another useful Web server report is a trend of errors by type over time. Even reviewing access to Web server pages over time might reveal the pattern of malicious activities which increases the volume of access as well as the volume of errors generated in Web server logs.

Web server log reports that show file types served off the Web server can help detect injection attacks against the Web server. For example, if your Web server shows that it served the *.exe file to a browser, this indicates that malware injection has taken place. Checking all the served extensions is very useful for detecting such attacks.

If a Web server is used for authenticating the access to a Web application, reviewing authentication status records is also important – HTTP error 401 indicates failed password while HTTP code 300 indicates successful authentication.

In some cases, reviewing the trend of failed access is useful for detecting Web security scanning and exploitation; use report "Microsoft IIS / Microsoft IIS – Top 20 Page not Found (404)" for that purpose. Another useful report of the same type is "Microsoft IIS / Microsoft IIS – Top 20 Script Errors(501)."

Run the Web server access summary report every week or less frequently if your business does not rely on Web servers as a key business activity. Review Web server error reports weekly or even daily in order to detect compromise attempts promptly.

Given the volume of Web server logs, real time monitoring is unlikely for most organizations.

Correlation and alerting: real-time detection of suspicious activities across Web server logs is entirely possible using correlation technology.

First, it is useful to create a rule that triggers when an unusual number of Web access errors is registered:

– Too many failure types off same source IP

Use hourly baseline functionality for this.

Baseline functionality can also be used even though any large number of authentication failures need to be detected.

If a Web server serves a malicious file that means that your Web server is compromised and is now serving malware. Use the following rule to detect such occurrences.

Malicious attempts to proxy connections through a Web server must also be detected; trigger alerts if successful CONNECT requests are executed at your Web server.

The first common mistake is storing logs for too short a time

Also, check for large files being downloaded off your Web server – especially files that you didn't put there – this activity calls for an alert.

More advanced correlation rules might use both database and Web server logs for advanced attack detection, provided that external Web server messages can themselves be correlated with database audit logs.

Alert both Web admin and security when the above issues are discovered. If the Web server is used for credit card transactions, its compromise will lead to immediate PCI compliance status violation.

What else: when analyzing Web server log data for Web application attack detection, it is essential to be prepared to conduct incident response without putting the server offline. For many organizations, the security team will not be able to disconnect the server, even in case of a severe compromise, or malware infection. This puts extra emphasis on early attack detection using correlation technology.

Incident response enablement

SIEM and log management tools that can collect massive volumes of diverse log data without issues are hugely valuable for incident response. Having a single repository for all activity records, audit logs, alerts and other log types allows incident responders to quickly assess what was going on during an incident and what led to a compromise or insider abuse. Incident response is the only unavoidable part of information security.

The 2009 Verizon breach report indicates that a majority of system compromises are discovered by third parties and not by organization's security teams. In light of this, incident response process might need to be activated at any moment when notification of a possible incident arrives. From this point onward, the security team will try to contain the damage and investigate the reason for the attack or abuse based on initial clues.

Collection: the scope of log collection for incident response is very simple – any and all logs from networks, hosts, applications and other information systems can be useful for a response to an incident. If you' re missing one piece of the puzzle, you may not be able to pinpoint the root cause.

The same applies to context data – information about users, assets, and vulnerabilities will come handy during the panic of incident response.

Overall, having as much data as possible will allow your organization to both effectively investigate what happened and to prevent its recurrence in the future.

Reporting and dashboards: incident response often calls for ad hoc reports or message review based on keywords and other criteria, such as a particular user name, IP address or a type of an attack.

Overall, many reports can be useful during incident response, depending upon the exact nature of a reported incident.

Correlation and alerting: many types of rules can be used during an incident to perform ongoing detection of future similar attacks. Almost any rule that has high confidence of triggering on a real incident can be used for automatically opening an incident task.

What else: incident responders rely on both process and technology in order to accomplish their difficult mission. Having an incident response plan, and having it tested in a simulation, will ensure that when a real incident strikes, the team will be prepared from both a procedural point of view as well as a skilled tool operational point of view.

While responding to an incident, it is important to keep track of all the findings and notes.

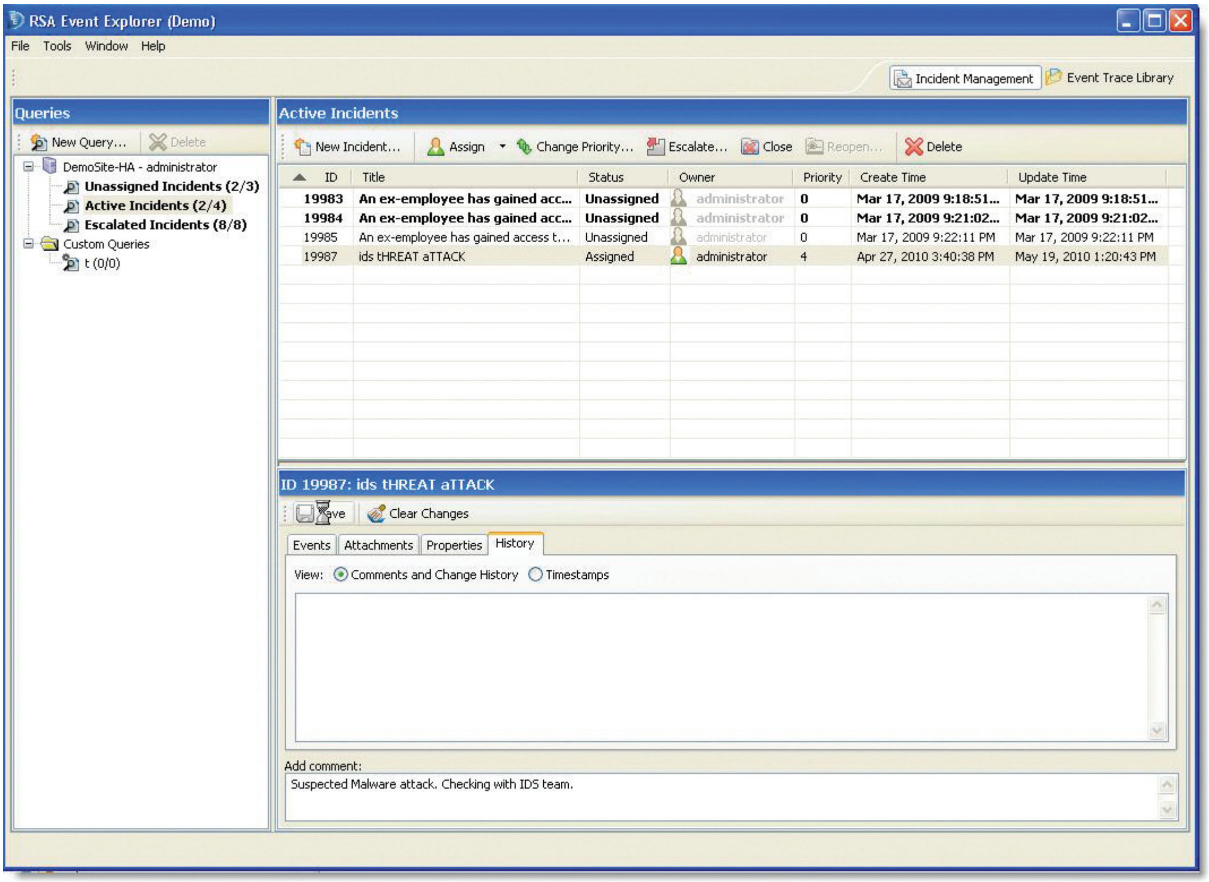

ID	Title	Status	Owner	Priority	Create Time	Update Time
19983	An ex-employee has gained acc…	Unassigned	administrator	0	Mar 17, 2009 9:18:51…	Mar 17, 2009 9:18:51…
19984	An ex-employee has gained acc…	Unassigned	administrator	0	Mar 17, 2009 9:21:02…	Mar 17, 2009 9:21:02…
19985	An ex-employee has gained access t…	Unassigned	administrator	0	Mar 17, 2009 9:22:11 PM	Mar 17, 2009 9:22:11 PM
19987	ids tHREAT aTTACK	Assigned	administrator	4	Apr 27, 2010 3:40:38 PM	May 19, 2010 1:20:43 PM

View: ● Comments and Change History ○ Timestamps

Add comment:
Suspected Malware attack. Checking with IDS team.

> Trying to use advanced SIEM features before mastering log collection and simple reporting will likely not work well

Pitfalls, mistakes, errors and "worst practices"

Given that effective SIEM deployments that go beyond regulatory compliance are often challenging for organizations, let us review a few common mistakes that will help you make your SIEM deployment and operation more pleasant.

The first common mistake is storing logs for too short a time. This makes the security or IT operations team think they have all the logs needed for monitoring and investigation or troubleshooting and then they have the horrible realization after the incident that all logs are gone due to their shortsighted retention policy. It often happens (especially in the case of insider attacks) that the incident is discovered a long time – sometimes many months – after the crime or abuse has been committed. One might save some money on storage hardware, but lose a large amount due to theft or fines. The incident response use cases discussed in this paper call for longer term retention, especially if insider abuse incidents are concerned.

In fact, organizations that are just starting on their journey to SIEM and log management should use longer retention times since they are less likely to rely on well tuned correlation rules and other near real time alerting mechanisms.

The next mistake is related to log record prioritization. While people need a sense of priority to better organize their log analysis efforts, the common mistake nowadays is in prioritizing the log records before collection. In fact, even some "best practice" documents recommend only collecting "the important stuff." But what is important? This is where the guidance documents fall short – by not specifying it in any useful form. While there are some approaches to the problem, all that I am aware of can lead to glaring holes in security posture or even undermine the regulatory compliance efforts.

The third mistake is in trying to use advanced SIEM features before mastering log collection and simple reporting. More than a few organizations ended up with failed SIEM deployments due to the fact that they tried to use advanced functionality and collect unusual logs in the first phase. This led to disillusionment and inability to achieve success using SIEM technology. On a similar note, not focusing on basics and fundamental requirements of log collection and reporting, organizations can lose their chances to eventually graduate to using advanced features.

Conclusions

To conclude, recent challenges with SIEM and log management that affected some organizations frequently stem from the fact that the powerful and advanced SIEM technology is purchased to address a narrow compliance mandate. Expanding the use of a SIEM beyond compliance to security and operational use cases happens slowly, if at all.

However, benefits from using SIEM go much beyond "checking a compliance box." Being able to leverage the power of SIEM takes some determination, knowledge of your environment, awareness of your business priorities –and some trial and error time with your SIEM tool.

Further, while SIEM presents a layer above security point solutions such as firewall, IDS, antivirus, Web proxy and others, many tools can enhance the SIEM mission as well as expand its use for security and compliance. Some of the key technologies that enhance the value of security information and event management are:

- Security configuration management (SCM) – combining configuration data with SIEM allows additional awareness of changes as well as other system issues
- Data leak prevention (DLP) – tightly integrating DLP and SIEM helps organizations improve efficiency and better prioritize incidents by correlating infrastructure risks with sensitive information, addressing both security and compliance problems.

> Trying to figure out what logs to collect before actually collecting them will fail; collecting 100% is important

- Finally, using Governance, Risk and Compliance (GRC) as a layer above SIEM allows some organizations to integrate their security management practices and process with higher level corporate concerns.

As your organization learns to operate SIEM for many of the use cases described in this paper, it makes sense to start exploring additional SIEM integrations with the above technologies. Finally, a few additional SIEM deployment and operation success tips are presented below:

- Always deploy and operationalize SIEM in phases; such phases will apply to both the scope of log collection and the utilization of SIEM features. Go from traditional server and firewall logs to advanced application logs, similarly, advance from collection and simple reporting to correlation, real time alerts and analysis.
- Think about the use cases discussed in this paper while deploying and using SIEM. Even if compliance is a primary driver, focusing on achieving outcomes useful for your business will give you more success on your journey to information security.
- Solidify your success for each use case before advancing to more log collection, context data collection and using advanced features.
- If building and running a SOC is your ultimate goal, make sure to familiarize yourself with your SIEM technology by successfully implementing and operating simpler use cases. SOC operation integrates all other use cases together in a coherent blend of technology, process and people.

About the Author

Dr. Anton Chuvakin is a recognized security expert in the field of log management and PCI DSS compliance. He is an author of books “Security Warrior” and “PCI Compliance” and a contributor to “Know Your Enemy II”, “Information Security Management Handbook” and others. Anton has published dozens of papers on log management, SIEM, correlation, security data analysis, PCI DSS, security management. His blog “Security Warrior” is one of the most popular in the industry.

In addition, Anton teaches classes and presents at many security conferences across the world; he recently addressed audiences in United States, UK, Singapore, Spain, Russia and other countries. He works on emerging security standards and serves on advisory boards of several security start-ups.

Currently, Anton is developing his security consulting practice, focusing on logging and PCI DSS compliance for security vendors and Fortune 500 organizations. Dr. Anton Chuvakin was formerly a Director of PCI Compliance Solutions at Qualys. Previously, Anton worked at LogLogic as a Chief Logging Evangelist, tasked with educating the world about the importance of logging for security, compliance and operations. Before LogLogic, Anton was employed by a security vendor in a strategic product management role. Anton earned his Ph.D. degree from Stony Brook University.

About RSA

RSA, the Security Division of EMC, is the expert in information-centric security, enabling the protection of information throughout its lifecycle. RSA enables customers to cost-effectively secure critical information assets and online identities wherever they live and at every step of the way, and manage security information and events to ease the burden of compliance.

RSA offers industry-leading solutions in identity assurance & access control; encryption & key management; governance & risk management; compliance & security information management and fraud protection. These solutions bring trust to millions of user identities, the transactions that they perform, and the data that is generated. For more information, please visit www.RSA.com and www.EMC.com.

The Security Division of EMC

RSA Security LLC
RSA Security Ireland Limited
www.rsa.com