Security information and event management

Introduction:

Security Information and Event Management (SIEM) automates incident identification and resolution based on built in business rules to help improve compliance and alert staff to critical intrusions. IT audits, standards and regulatory requirements have now become an important part of most enterprises' day-to-day responsibilities. As part of that burden, organizations are spending significant time and energy scrutinizing their security and event logs to track which systems have been accessed, by whom, what activity took place and whether it was appropriate. Organizations are increasingly looking towards data-driven automation to help ease the burden. As a result, the SIEM has taken form and has provided focused solutions to the problem. The security information and event management market is driven by an extremely increasing need for customers to meet compliance requirements as well as continued need for real-time awareness of external and internal threats. Customers need to analyze security event data in real time (for threat management) and to analyze and report on log data and primarily this has made security information and event management market more demanding. The market remains fragmented, with no dominant vendor.

This report entitled ' Security Information and Event Management (SIEM) Solutions' gives a clear view of the SIEM solutions and whether they can help to improve intrusion detection and response. Following this introduction is the background section which deeply analyzes the evolution of the SIEM, its architecture, its relationship with the log management and the need for SIEM products. In the analysis section, I have analyzed the SIEM functions in detail

along with real world examples. Finally the conclusion section summarizes the paper.

Background:

What is SIEM?

Security Information and Event Management solutions are a combination of two different products namely, SIM (security information management) and SEM (security event management). SIEM technology provides real-time analysis of security alerts generated by network hardware and applications. The objective of SIEM is to help companies respond to attacks faster and to organize mountains of log data. SIEM solutions come as software, appliances or managed services. Increasingly, SIEM solutions are being used to log security data and generate reports for compliance purposes. Though Security Information and Event Management and log management tools have been complementary for years, the technologies are expected to merge.

Evolution of SIEM:

SIEM emerged as companies found themselves spending a lot of money on intrusion detection/prevention systems (IDS/IPS). These systems were helpful in detecting external attacks, but because of the reliance on signature-based engines, a large number of false positives were generated. The first-generation SIEM technology was designed to reduce this signal-to-noise ratio and helped to capture the most critical external threats. Using rule-based correlation, SIEM helped IT detect real attacks by focusing on a subset of firewall and IDS/IPS events that were in violation of policy. Traditionally, SIEM solutions have been expensive and time-intensive to maintain and tweak,

but they solve the big headache of sorting through excessive false alerts and they effectively protect companies from external threats. While that was a step in the right direction, the world got more complicated when new regulations such as the Sarbanes-Oxley Act and the Payment Card Industry Data Security Standard followed much stricter internal IT controls and assessment. To satisfy these requirements, organizations are required to collect, analyze, report on and archive all logs to monitor activities inside their IT infrastructures.

The idea is not only to detect external threats, but also to provide periodic reports of user activities and create forensics reports surrounding a given incident. Though SIEM technologies collect logs, they process only a subset of data related to security breaches. They weren’t designed to handle the sheer volume of log data generated from all IT components, such as applications, switches, routers, databases, firewalls, operating systems, IDS/IPS and Web proxies. With an idea to monitor user activities rather than external threats, log management entered the market as a technology with architecture to handle much larger volumes of data and with the ability to extend to meet the demands of the largest enterprises. Companies implement log management and SIEM solutions to satisfy different business requirements, and they have also find out that the two technologies work well together. Log management tools are designed to collect report and archive a large volume and breadth of log data, whereas SIEM solutions are designed to correlate a subset of log data to point out the most critical security events.

On looking at an enterprise IT arsenal, it is likely to see both log management and SIEM. Log management tools often assume the role of a log data warehouse that filters and forwards the necessary log data to SIEM solutions for correlation. This combination helps in optimizing the return on investment while also reducing the cost for implementing SIEM. In these tough economic times it is likely to see IT trying to stretch its logging technologies to solve even more problems. It will expect its log management and SIEM technologies to work closer together and reduce overlapping functionalities.

Relation between SIEM and log management:

Like many things in the IT industry, there’s a lot of market positioning and buzz coming around regarding how the original term of SIM (Security Information Management), the subsequent marketing term SEM (Security Event Management), the newer combined term of SIEM (Security Information and Event Management) relate to the long standing process of log management. The basics of log management are not new. Operating systems, devices and applications all generate logs of some sort that contain system-specific events and notifications. The information in logs may vary in overall usefulness, but before one can derive much value

out of them, they first need to be enabled, then transported and eventually stored. Therefore the way that one does gather this data from an often distributed range of systems and get it into a centralized (or at least semi-centralized) location is the first challenge of log management that counts. There are varying techniques to accomplish centralization, ranging from standardizing on the syslog mechanism and then deploying centralized

syslog servers, to using commercial products to address the log data acquisition, transport and storage issues.

Some of the other issues in log management include working around network bottlenecks, establishing reliable event transport (such as syslog over UDP), setting requirements around encryption, and managing the raw data storage issues. So the first steps in this process are figuring out what type of log and event information is in need to gather, how to transport it, and where to store it. But that leads to another major consideration about what should one person want to do with all those data. It is at this point where the basic log management ends and the higher-level functions associated with SIEM begins. SIEM products typically provide many of the features that remain essential for log management but add event-reduction, alerting and real-time analysis capabilities. They provide the layer of technology that allows one to say with confidence that not only are logs being gathered but they are also being reviewed. SIEM also allows for the importation of data that isn't necessarily event-driven (such as vulnerability scanning reports) and it is known as the " Information" portion of SIEM.

SIEM architecture:

Long term log management and forensic queries need a database built for capacity, with file management and compression tools. Short term threat analysis and correlation need real time data, CPU and RAM. The solution for this is as follows:

> Split the feeds to two concurrent engines.

> Optimize one for real time and storage up to 30 days of data. (100-300GB)

> Optimize the second for log compression, retention, and query functions. (1TB+)

The block diagram showing the architecture of the SIEM is as follows: [Source: Reference 2]

A collector is a process that gathers data. Collectors are produced in many shapes and sizes from agents that run on the monitored device, to centralized logging devices with pre-processors to split stream the data. These can be simple REGEX file parsing applications, or complex agents for OPSEC, LEA, for . Net/WMI, SDEE/RDEP, or ODBC/SQL queries. Not all security devices are kind enough to forward data, and multiple input methods, including active pull capabilities, are very essential. Also, since SYSLOG data is not encrypted, it may need a collector to provide encrypted transport.

A threat analysis engine will need to run in real time, continuously processing and correlating events of interest passed to it by the collector, and reporting to a console or presentation layer application about the threats found. Typically reporting events that has happened for 30 days are sufficient for operational considerations. A log manager will need to store a great deal of data, and may take either raw logs or filtered events of interest, and need to compress store and index the data for long term forensic analysis and compliance reporting. Capacity for 18 months or more of data is likely to be required. Year end closing of books and the arrival of the auditors often necessitate the need for 12 months of historic data plus padding of several months while books are finalized and an audit to be completed.

At the presentation layer a console will present the events to the security staff and managers. This is the primary interface to the system for day to day operations, and should efficiently prioritize and present the events with a full history and correlation rationale.

SIEM functions:

With some subtle differences, there are four major functions of SIEM solutions. They are as follows:

1. Log Consolidation – centralized logging to a server

2. Threat Correlation – the artificial intelligence used to sort through multiple logs and log entries to identify attackers

3. Incident Management – workflow – What happens once a threat is identified? (link from identification to containment and eradication).

Notification – email, pagers, informs to enterprise managers (MOM, HP Openview…)

Trouble Ticket Creation

Automated responses – execution of scripts (instrumentation)

Response and Remediation logging

4. Reporting

Operational Efficiency/Effectiveness

Compliance / SOX, HIPPA, FISMA….

Ad Hoc / Forensic Investigations

Coming to the business case for SIEM, all engineers are perpetually drawn to new technology, but purchasing decisions should by necessity be based on need and practicality. Even though the functions provided by SIEM are impressive they must be chosen only if they fit an enterprise’s needs.

Why use a SIEM?

There are two branches on the SIEM tree namely, operational efficiency and effectiveness, and log management/compliance. Both are achievable with a good SIEM tool. However since there is a large body of work on log management, and compliance has multiple branches, this coursework will focus only on using a SIEM tool effectively to point out the real attackers, and the worst threats to improve security operations efficiency and effectiveness. It can be believed that the most compelling reason for a SIEM tool from an operational perspective is to reduce the number of security events on any given day to a manageable, actionable list, and to automate analysis such that real attacks and intruders can be discerned. As a whole, the number of IT professionals, and security focused individuals at any given company has decreased relative to the complexity and capabilities demanded by an increasingly inter networked web. While one solution may have dozens of highly skilled security engineers on staff pouring through individual event logs to identify threats, SIEM attempts to automate that process and can achieve a legitimate reduction of 99. 9+% of security event data while it actually increases the effective detection over traditional human driven monitoring. This is why SIEM is preferred by most of the companies.

Reasons to use a SIEM:

To know the need for a SIEM tool in an organization is very important. A defense in depth strategy (industry best practice) utilizes multiple devices: Firewalls, IDS, AV, AAA, VPN, User Events – LDAP/NDS/NIS/X. 500, Operating System Logs… which can easily generate hundreds of thousands of events per day, in some cases, even millions. No matter how good a security engineer is, about 1, 000 events per day is a practical maximum that a security engineer is about to deal with. So if the security team is to remain small they will need to be equipped with a good SIEM tool. No matter how good an individual device is, if not monitored and correlated, each device can be bypassed individually, and the total security capabilities of a system will not exceed its weakest link. When monitored as a whole, with cross device correlation, each device will signal an alert as it is attacked raising awareness and threat indications at each point allowing for additional defences to be brought into play, and incident response proportional to the total threat. Even some of the small and medium businesses with just a few devices are seeing over 100, 000 events per day. This has become usual in most of the companies says the internet.

Real world examples:

Below are event and threat alert numbers from two different sites currently running with 99. xx% correlation efficiency on over 100, 000 events per day, among which one industry expert referred to as “ amateur” level, stating that 99. 99 or 99. 999+% efficiency on well in excess of 1, 000, 000 events per day is more common.

Manufacturing Company Central USA – 24 hour average, un-tuned SIEM day of deployment

Alarms Generated 3722

Correlation

Efficiency 99. 06%

Critical / Major

Level Alerts 170

Effective Efficiency 99. 96%

[Source: Reference 2]

In this case, using a SIEM allows the company’s security team (2 people in an IT staff of 5), to respond to 170 critical and major alerts per day (likely to decrease as the worst offenders are firewalled out, and the worst offenses dealt with), rather than nearly 400, 000.

Financial Services Organization – 94, 600 events – 153 actionable alerts – 99. 83% reduction.

[Source: Reference 2]

The company above deals with a very large volume of financial transactions, and a missed threat can mean real monetary losses.

With respect to the Business Case, a good SIEM tool can provide the analytics, and the knowledge of a good security engineer can be automated

and repeated against a mountain of events from a range of devices. Instead of 1, 000 events per day, an engineer with a SIEM tool can handle 100, 000 events per day (or more). And a SIEM does not leave at night, find another job, take a break or take vacations. It will be working always.

SIEM Selection Criteria:

The first thing one should look at is the goal. (i. e.) what should the SIEM do for them. If you just need log management then make the vendor can import data from ALL of the available log sources. Not all events are sent via SYSLOG. Some may be sent through:

Checkpoint – LEA

Cisco IDS – RDEP/SDEE encryption

Vulnerability Scanner Databases – Nessus, Eeye, ISS…

AS/400 & Mainframes – flat files

Databases – ODBC/SQL queries

Microsoft . Net/WMI

Consider a product that has a defined data collection process that can pull data (queries, retrieve files, WMI api calls…), as well as accept input sent to it. And it is essential to be aware that logs, standards, and formats change, several (but not all), vendors can adapt by parsing files with REGEX and importing if one can get them a file. However log management itself is not usually an end goal. It matters about for what purpose these logs are used for. They may be used for threat identification, compliance reporting or

forensics. It is also essential to know whether the data captured is in real-time. If threat identification is the primary goal, 99+% correlation/consolidation/aggregation is easily achievable, and when properly tuned, 99. 99+% efficiency is within reach (1-10 actionable threat alerts / 100, 000 events).

If compliance reporting is the primary goal, then consider what regulations one is subject to. Frequently a company is subject to multiple compliance requirements. Consider a fortune 500 company like General Electrics. As a publicly traded company GE is subject to SOX, as a vendor of medical equipment and software they are subject to HIPPA, as a vendor to the Department of Defense, they are subject to FISMA. In point of fact, GE must produce compliance reports for at least one corporate division for nearly each and every regulation. Two brief notes on compliance, and one should look at architecture: Beware of vendors with canned reports. While they may be very appealing, and sound like a solution, valid compliance and auditing is about matching output to one’s stated policies, and must be customized to match each company’s published policies. Any SIEM that can collect all of the required data, meet ISO 177999, and provide timely monitoring can be used to aid in compliance. Compliance is a complex issue with many management, and financial process requirements, it is not just a function or report IT can provide.

Advanced SIEM Topics:

Risk Based Correlation / Risk Profiling

Correlation based on risk can dramatically reduce the number of rules required for effective threat identification. The threat and target profiles do most of the work. If the attacks are risk profiled, three relatively simple correlation rules can identify 99%+ of the attacks. They are as follows:

IP Attacker – repeat offenders

IP Target – repeat targets

Vulnerability Scan + IDS Signature match – Single Packet of Doom

Risk Based Threat Identification is one of the more effective and interesting correlation methods, but has several requirements:

> A Metabase of Signatures – Cisco calls the attack X, ISS calls it Y, Snort calls it Z – Cross Reference the data

> Requires automated method to keep up to date.

> Threats must be compiled and threat weightings applied to each signature/event. Reconnaissance events are low weighting – but aggregate and report on the persistent (low and slow) attacker

Finger Printing – a bit more specific, a bit higher weighting

Failed User Login events – a medium weighting, could be an unauthorized attempt to access a resource, or a forgotten password.

Buffer Overflows, Worms and Viruses -high weighting -potentially destructive – events one need to respond to unless one has already patched/protected the system.

> The ability to learn or adjust to one’s network Input or auto-discover which systems, are business critical vs. which are peripherals, desktops, and non-essential

> Risk Profiling:

Proper application of trust weightings to reporting devices (NIST 800-42 best practice), can also help to lower “ cry wolf” issues with current security management

Next-generation SIEM and log management:

One area where the tools can provide the most needed help is in compliance. Corporations increasingly face the challenge of staying accountable to customers, employees and shareholders, and that means protecting IT infrastructure, customer and corporate data, and complying with rules and regulations as defined by the government and industry. Regulatory compliance is here to stay, and under the Obama administration, corporate accountability requirements are likely to grow. Log management and SIEM correlation technologies can work together to provide more comprehensive views to help companies satisfy their regulatory compliance requirements, make their IT and business processes more efficient and reduce management and technology costs in the process.

IT organizations also will expect log management and intelligence technologies to provide more value to business activity monitoring and business intelligence. Though SIEM will continue to capture security-related data, its correlation engine can be re-appropriated to correlate business processes and monitor internal events related to performance, uptime,

capability utilization and service-level management. We will see the combined solutions provide deeper insight into not just IT operations but also business processes. For example, we can monitor business processes from step A to Z and, if a step gets missed, we'll see where and when. In short, by integrating SIEM and log management, it is easy to see how companies can save by de-duplicating efforts and functionality. The functions of collecting, archiving, indexing and correlating log data can be collapsed. That will also lead to savings in the resources required and in the maintenance of the tools.

CONCLUSION:

SIEM is a complex technology, and the market segment remains in flux. SIEM solutions require a high level of technical expertise and SIEM vendors require extensive partner training and certification. SIEM gets more exciting when one can apply log-based activity data and security-event-inspired correlation to other business problems. Regulatory compliance, business activity monitoring and business intelligence are just the tip of the iceberg. Leading-edge customers are already using the tools to increase visibility and the security of composite Web 2. 0 applications, cloud-based services and mobile devices. The key is to start with a central record of user and system activity and build an open architecture that lets different business users access the information to solve different business problems. So there is no doubt in SIEM solutions helping the intrusion detection and response to improve.

References:

1. Nicolett. M., Williams. A. T., Proctor. P. E. (2006) ' Magic Quadrant for Security Information and Event Management, 1H06' RA3 1192006.

2. Swift. D. (2006) ‘ A Practical Application of SIM/SEM/SIEM Automating Threat Identification’

3. ‘ SIEM: A Market Snapshot’ (2007) from http://www. crn. com/security/197002909; jsessionid= BVQXTH11HH14JQE1GHPSKH4ATMY32JVN [Date Accessed: 20th November, 2009].

4. ‘ WHAT IS SIEM’ (2008) from http://www. exploresiem. com/resource-center. html [Date Accessed: 24th November, 2009].

5. ‘ Securing and Managing Your Enterprise: An Integrated Approach’ (2008) from

http://www. exploresiem. com/images/WP-Securing-and-Managing-Your-Enterprise. pdf [Date Accessed: 25th November, 2009].

6. Shipley . G.(2008) ‘ Are SIEM and log management the same thing?’ from http://www. networkworld. com/reviews/2008/063008-test-siem-log-integration. html [Date Accessed: 26th November, 2009]

7. Levin. D. (2009) ‘ The convergence of SIEM and log management’ from http://www. networkworld. com/news/tech/2009/031909-tech-update. html [Date Accessed: 26th November, 2009]