



Interested in learning more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

A Practical Application of SIM/SEM/SIEM Automating Threat Identification

The goal of this paper is to explain how to use a SIEM effectively to identify and respond to security threats. The paper begins with level set information including definitions, capabilities requirements, architecture and a business case. Later I will cover aggregation and correlation design concepts, with real world examples including architectural design, risk based profiling, finite state engines, and merging traditional network operations data into security operations tools for improved detection.

Copyright SANS Institute
Author Retains Full Rights

AD

Build your business'
breach action plan.

START NOW

 **LifeLock**
BUSINESS SOLUTIONS

No one can prevent all identity theft. © 2016 LifeLock, Inc. All rights reserved. LifeLock and the LockMan logo are registered trademarks of LifeLock, Inc.

A Practical Application of SIM/SEM/SIEM
Automating Threat Identification

By David Swift

December 23, 2006

© SANS Institute 2007, Author retains full rights.

Table of Contents

Introduction.....	2
SIEM Functions.....	4
Reasons to use a SIEM.....	6
Real world examples.....	7
Alternatives.....	9
Taxonomy of an Attack.....	11
Correlation Example 1.....	14
SIEM Selection Criteria.....	16
SIEM Architecture.....	18
Aggregation, Normalization, and Correlation.....	20
Basic Correlation Principles.....	24
Correlation Example 2.....	25
Notification and Event Response.....	26
Advanced SIEM Topics.....	27
<i>Risk Based Correlation / Risk Profiling</i>	27
Correlation Example 3.....	29
<i>Finite State Engine</i>	31
NOC meets SOC.....	32
Correlation Example 4.....	33
Summary.....	35
Acknowledgements.....	36
Appendix A - SIEM Vendors.....	37
Request for Input.....	38

Introduction

In the world of IT security, if you've ever felt like Davey Crocket at the Alamo, it's you and a few good men versus thousands of heavily armed attackers, and you still hope to win the battle, then read on. Proper deployment of a SEM tool prior to an incident can radically increase one's effectiveness at identifying an incident in progress. In GCIH terms SEM tools are part of preparation and identification, and can be invaluable in forensics as well. In this paper I will try and stay focused on general capabilities of all SIEM tools, though your mileage may vary depending on the tool you choose to implement. A large focus will be on proper correlation techniques that could be applied manually, but can be made over 1000 times more effective when automated. Many of the real world examples, and capabilities discussed will be from my experience with two SIEM products - OpenService's Security Management Center (SMC), and EIQ's Network Security Analyzer. While there may be other better tools, I want to share real world practical advice based on experience rather than extrapolation. For those of you using other tools, I would appreciate additional feedback, and have listed my email and specific areas for additional requests for information from the broader community at the end of the paper.

The goal of this paper is to explain how to use a SIEM effectively to identify and respond to security threats. The paper begins with level set information including definitions, capabilities requirements, architecture and a business case. Later I will cover aggregation and correlation design concepts, with real world examples including architectural design, risk based profiling, finite state engines, and merging traditional network operations data into security operations tools for improved detection.

First let's define a rather broad, widely misused term.

What is SIM/SEM/SIEM?

Security Event Management - SEM

Security Information Management - SIM

Security Information and Event Management - SIEM

For purposes of this paper the acronym SIEM will be used generically to refer to tools with the capabilities outlined below.

No SIEM tool is an island. To function effectively, a SIEM tool will require pre-deployment and integration with several security devices. For optimum effectiveness, reporting data from a firewall, and IDS sensor, an authentication service (AAA, LDAP, AD, etc..), and vulnerability scan data will need to be integrated during the incident handling preparation phase. Correlations, and operational efficiency gains are directly related to the identification phase.

In addition, for forensic identification and prosecution the data capture and correlated can be invaluable. For auditing and compliance, proper reporting can go a long way towards proving compliance.

© SANS Institute 2007. All rights reserved.

SIEM Functions

With some subtle differences, there are four major functions of SIEM solutions:

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

© SANS Institute 2007. Author retains full rights.

Next, we'll analyze the business case for SIEM.

As an engineer I'm perpetually drawn to new technology, but purchasing decisions should by necessity be based on need and practicality. Given this as a basic business tenant, this paper will also attempt to build a valid business case for cost justification.

Why use a SIEM?

There are two branches on the SIEM tree - 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 paper will focus on using a SIEM tool effectively to ferret out the real attackers, and the worst threats to improve security operations efficiency and effectiveness. I'm continually asked "Who do I bop on the head?" by clients when deploying security tools. SIEM can answer that question better than any other tool I've seen.

I believe 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 actually increasing effective detection over traditional human driven monitoring.

Reasons to use a SIEM

It is not uncommon for management to fail to see the need for such tools, I'd like to cover a few basics.

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 to deal with. So if the security team is to remain small they'll need to be equipped with a good SIEM tool.

No matter how good an individual device, 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 defenses 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 I've worked with are seeing over 100,000 events per day.

© SANS Institute 2007. All rights reserved. This document is for personal use only. All other rights reserved. No part of this document may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or by any information storage and retrieval system, without the prior written permission of SANS Institute.

Real world examples: (company names removed due to NDA conflicts, only industry listed)

Below are event and threat alert numbers from two sites currently running with 99.xx% correlation efficiency on over 100,000 events per day, 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.

Real world examples:

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


397471	Events	Events Per X
24	Hours	16561
1440	Minutes	276
86400	Seconds	5

Alarms Generated	3722
Correlation	
Efficiency	99.06%
Critical / Major	
Level Alerts	170
Effective Efficiency	99.96%

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.

Real World Example

Financial Services Organization - 94,600 events - 153 actionable alerts - 99.83% reduction.



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

*

With respect to our Business Case:

A good SIEM tool can provide the analytics and 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 doesn't leave at night, find another job, or take vacations.

Alternatives

One alternative is the use of a *Unified Threat Management (UTM)* security device - a single do all correlate all security device. In a "green field" environment, or a small to medium size business these may represent a "good enough" solution that affects the greatest increase in overall security. In a larger organization with significant security assets already in place, the replacement of those devices may be cost prohibitive. This is not to say that a UTM may be able to augment an existing security infrastructure, by adding one or more functions needed, without replacing existing systems. In some cases if a UTM can fill two or more needs, it may be justified by just those features alone. As with all decisions, there are tradeoffs.

The Upside to a UTM:

May ease the burden of administration and change control.
May reduce training costs, and increase your staffs' abilities to affect tighter security.

The Downside to a UTM:

Cost - you may have to "Rip and Replace" you're entire security infrastructure.

Capabilities - A UTM is unlikely to do everything you need well. You will loose best of breed options. Correlation, if possible, is by device, rather than across the enterprise.

Other practical applications of a SIEM may also make the decision much easier. Though not covered in this paper, a SIEM can dramatically improve a company's ability to meet compliance regulations and industry best practices. The ISO 17799/27001/BS7799 standard is an underpinning to any good SIEM solution, and an integral part to meeting Sarbanes Oxley, PCI, GLBA, HIPPA, FISMA, or any internal or external compliance goal you may need to achieve.

SOX (HR 3763 / 107 section 404) requires:

"Timely monitoring and auditing of systems used to track financial data with an annual review and statement of the effectiveness of the tools used to do so."

Whether the audit is for internal compliance, or one of the external regulations timely monitoring and response is a reoccurring theme.

© SANS Institute 2007, Author retains full rights.

Taxonomy of an Attack

Just for clarity, we'll examine an attack, with and without a SIEM tool.

Basic Steps a Determined Evasive Attacker Takes

Discovery Phase

1. Attacker Scans the Firewall (NMAP, Firewalker, HPING, etc...) Which IP addresses respond? Which ports are open? Low and Slow to avoid triggering automatic protections

2. Finger Printing

Continued, targeted scan (NMAP, HPING, etc...) What operating system is running on discovered hosts? On the discovered hosts, what applications are running?

Targeting with IDS Evasion

Send targeted attacks of known vulnerabilities (buffer overflows, With Fragmented packets (fragroute, nemesis) With signature evading patterns (admutate, metasploit)

Compromise

System Crash, Denial of Service, or Data theft
Install sniffers, backdoors or rootkits for ongoing access


DSNIFF. Ettercap, Ethereal

Netcat, VNC

BackOrifice, LRK, AFK, KIS

David Swift

Basic Network Diagram - Single Firewall, IPS, and Vulnerability



In each phase an attack can be crafted to bypass a single individual protection, having learned how to penetrate the preceding device.

The same attack with a SIEM

Discovery Phase

The Router or Firewall sends events to SIEM indicating port scans and an alert is built at minor/warning level.

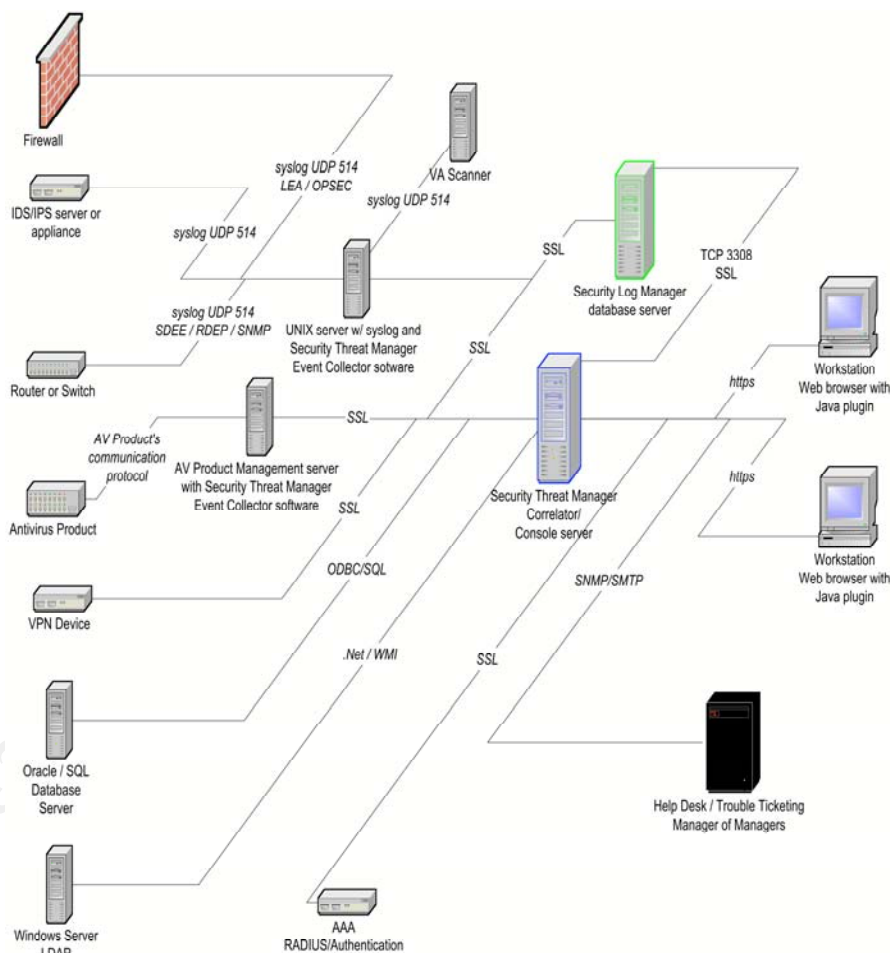
Finger Printing

The IDS/IPS reports system scans, and other possible signature matches, and the alert is raised to an elevated level. Security staff is notified (email, pager, etc...).

Targeting with Evasion

The firewall reports fragmented packets, the IDS may report certain signatures, and the alert level is raised to high. If the IDS sees an event and the vulnerability scanner knows the event can compromise a system, the alert is escalated to critical. Security staff is notified of a high probability threat and automated responses (firewall ACL, system shutdown, etc...), are taken.

Compromise



Under worst case scenarios, the download of exploit code will be detected by most IDS/IPS systems and anti-virus software.

Security staff is notified of a critical event and automated responses (firewall ACL, system shutdown, etc...), are taken.

Even if the individual events do succeed in bypassing the firewall, and evading the IDS and any anti-virus software, the total number of questionable packets should trigger a major threat. Instead of being bypassed in succession, each device is reporting on the unsuccessful events and raising the threat level of an alert and with it the defensive posture and responses of an alerted security staff. A single evade all compromise without failed tries would be a legendary accomplishment, and cannot be mitigated.

Visual Walkthrough Link:

<http://www.openservice.com/hacker1.php>

© SANS Institute 2007, Author retains full rights.

Correlation Example 1 (International Banking Company)

Attacker compromises an account. The attacker attempts repeated logins (brute force), on a known user account against a custom database application using a "low and slow attack" to locate and penetrate a custom database.

Custom Database application (Oracle Financials), does not report via SYSLOG, nor log either source or destination IP in event logs making forensics and correlation challenging. In order to enable correlation and SEM, REGEX filters are used to normalize the Oracle Event logs and pulled into the SEM application via a log parsing script. The destination IP is inferred to be the Oracle database Server's IP and appended to the data as it is parsed.

By using low/slow (2 login attempts per hour), the attacker avoids IPS detection and application/OS account lockout. The attack can continue indefinitely until the account password is discovered, and is only thwarted by a SEM Application, or the user changing their password (not commonly required on commercial sites for customers). All events are recorded in each correlation instance they could apply to.

Count	Event Thresholds Before Alert by Type
10	User
100	Source
500	Destination
1000	Port
1 Hour	Decay Rate (50% lower priority if no new events received on a given correlation)

The thresholds are set to dampen noise, and avoid frequent false positives.

Correlation Example (International Banking Company)

Event	Event Description	Correlations Triggered		
		Source (1)	Destination (2)	Port (3)
1-50	Dropped connections on a firewall – invalid port/dest. Destination, Port and Source			
	Each instance < alert threshold of 100, all 50 events added to each correlation			
51-53	Failed Login to Database	User (4)	Destination (2)	
	database reports only User ID, and Destination IP, source may be spoofed			
	events added to destination correlation already instantiated			
	< 2 failed attempts / hour fails to trigger account lockout, HIDS, or IPS			
54-60	Failed Login to Database	User (4)	Destination (2)	
	Additional failed logins "Low and Slow"			
	< 2 failed attempts / hour fails to trigger account lockout, HIDS, or IPS			

User correlation (4) triggers an alert at event 60 (10 failed logins), starting incident response/handling. Port correlation (3) will have aged out of the system. User Correlation (4) and Destination correlation (2), will be elevated and create threat alerts. Additional future events will continue to escalate the alerts. Source correlation (1) will have aged out of the system. User correlation (4) will have only User ID and Destination IP.

However, Destination correlation (2) will have a complete history of the attack, with a full listing of all events, and can provide user ID, source IP, ports, and destination IP.

SIEM Selection Criteria

The first thing one should look at is the goal.

What do you want your SIEM to do?

If you just need log management then make sure you're vendor can import data from ALL of your log sources.

Not all events are sent via SYSLOG, consider:

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 be aware that logs, standards, and formats change, several (but not all), vendors can adapt by parsing files with REGEX and importing if you can get them a file.

However log management itself is not usually an end goal.

What are you going to use the logs for? Threat Identification? Compliance Reporting? Forensics? Does it have to be real-time? Is next day OK?

If threat identification is your 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 your primary goal, then consider what regulations you're subject too. 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 every regulation.

Two brief notes on compliance, and we'll 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 your 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 17799, and provide timely monitoring can be used to aid in compliance. Compliance is a complex issue with many management, and financial process requirements, not just a function or report IT can provide.

© SANS Institute 2007. Author retains full rights.

SIEM Architecture:

Two birds, one stone - Split Architecture / Dual Data Streams

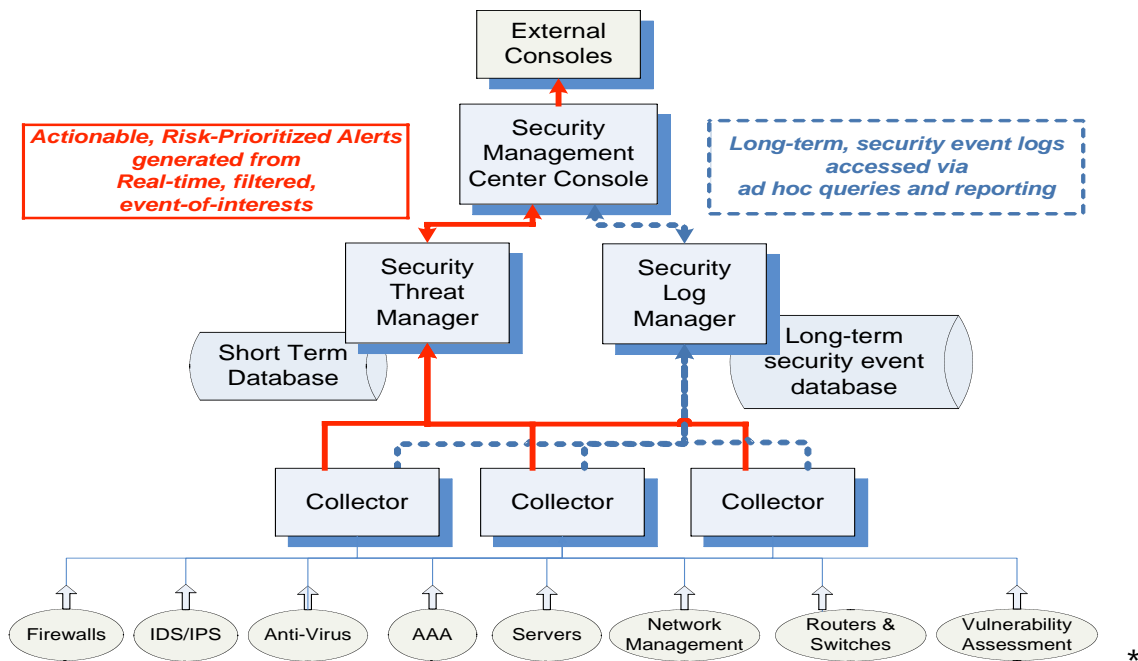
Design Issues:

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.

Solution:

Split the feeds to two concurrent engines. Optimize one for real time and up to 30 days of data. (100-300GB) Optimize the second for log compression, retention, and query functions. (1TB+)



Functionally:

A collector is a process that gathers data. Collectors come 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 essential. Also, since SYSLOG data is not encrypted, you may need a collector to provide encrypted transport.

A threat analysis engine will need to run in real time, processing and correlating events of interest passed to it by the collector, and reporting to a console or presentation layer application the threats found. Typically reporting events 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 needs 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 complete an audit.

At the presentation layer a console will present the events to 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.

Aggregation, Normalization, and Correlation

Assuming your goal is increased security effectiveness and efficiency, let's drill down a little on how to achieve it.

Basic Premise:

To thwart a threat, you must first identify it. If we can find the threat, we can put additional controls in place to prevent it (i.e. additional firewall rules, patch a system, take a system offline before infection, and drop malicious content [IPS and/or AV]).

Threat Identification can be broken into three parts:

Aggregation, Correlation, and Normalization

Aggregation - Everything counts in large amounts.

Useful Daily Reports

Top 10 Attackers - Where are my attacks coming from?

Top 10 Destinations - Which systems are under attack?

Top 10 Attacks - What are my most common threats?

Even aggregating and consolidating the data from multiple devices isn't as easy as it sounds. An event can come in with one or all of the possible event IDs and signature descriptions. Running a search for a given attack would require extensive OR IF clauses, and be impractical, if the data weren't normalized first. So first we have to normalize the data.

Consider the Sasser Worm from an IDS Perspective (a small portion shown)

IDS Vendor	Identifier	Description
Cisco Systems	3030/0	IDS Signature TCP SYN Host Sweep
Cisco Systems	3338/0	IDS Signature Windows LSASS RPC Overflow
Cisco Systems	3142/0	IDS Signature Sasser Worm Activity
SNORT	2507	NETBIOS DCERPC LSASS bind attempt
SNORT	2508	NETBIOS DCERPC LSASS DsRolerUpgradeDownlevelServer Exploit attempt
SNORT	2509	NETBIOS SMB DCERPC LSASS unicode bind attempt
SNORT	2510	NETBIOS SMB DCERPC LSASS bind attempt
SNORT	2511	NETBIOS SMB DCERPC LSASS DsRoler UpgradeDownlevelServer exploit attempt
SNORT	2512	NETBIOS SMB-DS DCERPCLSASS bind attempt
SNORT	2513	NETBIOS SMB-DS DCERPC LSASS unicode bind attempt
SNORT	2514	NETBIOS SMB-DS DCERPC LSASS DsRoler UpgradeDownlevelServer exploit attempt
SNORT	2524	NETBIOS DCERPC LSASS direct bind
SNORT	2525	NETBIOS SMB DCERPC LSASS direct bind
SNORT	2526	NETBIOS SMB-DS DCERPC LSASS direct bind
ISS Real Secure	15699	Microsoft Windows LSASS buffer overflow

Consider the Sasser Worm from a Vulnerability Perspective:

VA Vendor	Identifier	Description
Mitre (CVE)	CAN-2003-0533	Stack-based buffer overflow in certain Active Directory service functions in LSASRV.DLL
SecurityFocus	BID-10108	Microsoft Windows LSASS Buffer Overrun Vulnerability
nCircle IP360	3643	Worm: Sasser
Nessus	12219	Sasser Worm Infecting
Nessus	12220	Sasser Worm Infection
ISS/XForce	15818	Microsoft Windows MS04-011 patch is not installed

Normalization

Normalization is the process of cross referencing and enriching event data such that regardless of the source, the event ID, or the description, a common value can be derived. Normalized events can then be used to dampen repeat events from a single device, or multiple devices repeating the same event. Cross referencing and enriching event data with BugTrak and/or CVE vulnerability and threat databases are a suggested starting point. Either database is cross linked to vendor published vulnerabilities and patches for remediation.

A key element to consider in SIEM is whether the vendor keeps a metabase or is capable of normalizing and keeping multiple vendor signatures up to date.

Filter for events of interest (NIST 800-92)
<http://csrc.nist.gov/publications/drafts/DRAFT-SP800-92.pdf>
All events are not created equal. Logging everything results in data overload.

Even the Department Of Justice (DOJ), doesn't require it to be admissible as real evidence (DOJ business practice computer records reference)
http://www.usdoj.gov/criminal/cybercrime/usamarch2001_4.htm
Rules of evidence "If a business routinely relies on a record, that record may be used as evidence."

© SANS Institute 2007, All rights reserved.

Correlations

Correlation - A single packet can kill a host. Watch for the magic bullets.

Event Based

You're IDS reports a Signature X targeted at Host Y, Your VA scanner knows that Y is vulnerable = BINGO! We have a Winner.

Rules Based

If X + Y + Z then do A, or If X repeats more than 3 times in interval Y then do Z

Anomaly Based

If the traffic on port X exceeds the standard deviation of historic traffic patterns then there may be a problem (i.e. new worm, bot, or application)

Risk Based

If attack type = destructive (i.e. Buffer Overflow vs. SYN Scan), and target = critical asset (server vs. workstation), and reporting device = trustworthy (SourceFire RNA vs. untuned Snort), then open a threat alert or escalate a threat instance

How do I find the effective hacker? Who do I Bop On the Head?!

© SANS Institute 2007, Author retains full rights.

Basic Correlation Principles:

Watch for repeated events over a long period of time. Low and Slow Scans

Complex rules are not required.

In theory an attacker is going to follow a progressive attack, and we can watch for reconnaissance events followed by finger printing, followed by targeted exploits. Even a simple pair of aggregation rules that say if the events from source X or the events to destination Y exceed a noise threshold then Alert will work (and work well - 99% correlation from real-world users).

Start watching at the first line of defense (firewall or border router), and build a history of an attack.

Don't let each device be an Island.

Firewalker and NMAP can find the openings in a firewall - it's just a filter.

Fragroute, and morphed attacks (AdMutate) can evade IDS and AV - pattern matching can be fooled.

Don't just watch the IDS/IPS and AV Event logs. Too often security engineers review only the "best" log(s).

Who rang the doorbell? Did he go any further?

An attacker isn't going to know how to get through all of your defenses without first probing them.

Tying the devices together to look for an attack with a SIEM product means the attacker has to evade All of the Devices All at once, or get caught.

Create match/correlation rules for critical single packet (or single session), attacks that could compromise a system. Such as...an IDS signature match to a vulnerability found by your Vulnerability Scanner.

Correlation Example 2

A vulnerability scan discovers a security vulnerability, followed by an IDS sensor alert of an inbound exploit of that vulnerability targeted at the vulnerable device.

Events	Event Description	Correlations Triggered		
1	IDS Alert	Source IP(1) Attacker	Destination IP (2)	IDS / VA Match (3)

Microsoft Exchange Vulnerability

<http://national.uscert.org.au/render.html?it=6285&cid=>

Nessus Vulnerability Scan finds vulnerability CVE-2006-0027 on IP address X - primary Exchange server for corporation Y.

IPS sensor Z reports vCal/iCal possible Exchange exploit.

In this case a single event drives a critical level correlation creating an immediate alert, and incident handling procedures must be undertaken ASAP to avoid further exploits.

Since the Exchange server X could now be compromised, proper containment phase procedures should include isolation, and eradication procedures may require a complete rebuild of the server, patching to avoid a repeated compromise, and restoration of data. Additionally, the source IP may merit blocking at the border firewall, and forensic investigation and prosecution.

Correlations 1 and 2 will not yet have reached a critical threshold and may go unnoticed. However correlation 3 will have triggered an instant critical alert based on a single event.

SIEM is effectively tying the existing resources already deployed in the network into a cohesive synergistic defense.

Notification and Event Response

Consider using Multi-Factor Notification Trees

Time of Day - Anyone who's ever carried a pager, knows the value in NOT paging for non-critical events during off hours.

Type of Threat - Notify the party with the authority and knowledge to remediate the issue found.

Examples:

For Active Directory - repeated user failed login, or sensitive file access - notify the AD Admin

When a repeated attacker is found notify the firewall admin and consider blocking the source.

When repeated attacks against a target are correlated, notify the admin for the targeted system.

Severity

Place appropriate thresholds to dampen spastic IDS engines and Firewalls with constantly repeating non-destructive reconnaissance events, such that the events are categorized as Low/ Informational, or Guarded/ Minor. Consider notifying only during business hours, or when the quantity indicates a determined attacker rather than the daily webcrawler.

Elevated / Warning level for issues of concern, but not yet destructive threats (i.e. multiple failed logins to a sensitive account during the day).

For events that could indicate a threat, escalate the threat to High/Major status and notify appropriately (i.e. multiple repeated root failures on a Unix host during off production hours).

Severe/Critical status should be used for events known to be capable of causing immediate harm (i.e. when an IDS sensor sees an inbound back to a device your vulnerability scanner has reported as vulnerable to that attack).

Levels and color coded notification as suggested by the Department of Homeland Security.

* *Public Domain graphic from Department of Homeland Security.*



Many SIEMs can also automate creation of trouble tickets (via application informs, XLM/XLST, or SNMP alerts, and even feed events up to a Manager of Manger application. This is an advanced topic and feature that frequently requires customization.

Workflow and incident management, are features I expect to see expanded in SIEM products over the next few years.

Automated Response

SIEM tools frequently include the ability to execute external scripts, or automate rule additions to existing security devices (though OPSEC and other common APIs have failed to gain acceptance, there are still some integrations that allow automated rule insertion). In most cases hardening of SNMP v3 managed devices and systems that have an API or scriptable changes can be automated. While this sounds appealing, it is often impractical, human intervention by a trained incident handler is often more appropriate. There are times for forensic reasons or business needs a compromised system may be kept online.

Note to Management: (forgive the overstatement of the obvious [again], but...)

It is impossible to eliminate all security knowledgeable humans from the security process no matter how much you spend, if you want your defenses to work properly.

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.

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 - I don't want BlackHat1 to know what type of system/software I'm running.

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 I need to respond to...unless I've already patched/protected the system (see IDS and VA correlation)

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.




Diagram from NIST 800-42 Guideline on Network Security Testing p 12

A more complete implementation of risk profiling extends the NIST 800-42 best practice by applying similar weightings to attacks and targets in addition to reporting devices.

Correlation Example 3 (Profiling to limit Alerting)

Frequent IDS alerts against non-critical assets can be dampened to avoid "cry wolf" syndrome overtaxing security responders.

IP assets are initially assigned a weighting. Desktop network subnet X is given a weighting of 2 (low). Routers, and border devices are given a weighting of 1 (lowest, and as border guards are subject to the highest frequency of events). Back Office servers (mail, web, database, file servers), are given a weighting of 5 (high).

Count	Event Thresholds Before Alert by Type
10	Device Event Threshold – Warning
50	Device Event Threshold - Critical
Subnet X	Non-Critical Systems
20:00-05:00	Notification Delay Interval for Non-Critical Systems (log, and report, do not page)

Attacks are multiplied by their respective weighting during threat calculations such that two events against the same critical target will trigger an immediate notification, however, 5 events against the same desktop would be required to trigger a notification, which will be logged during non-production hours.

Correlation Example (Profiling to limit Alerting)

Events	Event Description	Correlations Triggered		
1	SMTP Debug Wiz	Source (1) Attacker	Destination (2) Desktop	Signature (3)
2	Sasser Worm	Source (1) Attacker	Destination (2) Desktop	Signature (4)
3	TFN2K	Source (1) Attacker	Destination (2) Desktop	Signature (5)
6	SMTP Debug Wiz	Source (1) Attacker	Destination (6) Server	Signature (3)
7	Sasser Worm	Source (1) Attacker	Destination (6) Server	Signature (4)
8	TFN2K	Source (1) Attacker	Destination (6) Server	Signature (5)

- Assumes none of the events are an IDS Signature/VA Scan match.

Correlation 1 will reveal a repeated (but ineffective), attacker that may be investigated, and possibly blocked by rule at the firewall to avoid future attacks.

Correlation 2 will remain at informational level only, not rising to the warning level threshold set for dampening.

Correlation 6 will be escalated to warning level at event 7 (2 x 5 = Warning Threshold of 10), and security responders will be notified.

A critical threshold (set to 50), will not be reached in the series, and depending on notification rules, need not be sent out during off hours.

Correlations 3, 4, and 5 will record most common signatures at informational level to allow prioritization of patch and prevention efforts.

Correlations 1 and 6 may indicate the most common targets and allow prioritization of patch and prevention efforts at a system level.

Finite State Engine


Consider the use of a Finite State Based engine rather than a rules based engine.

Rules Based vs. Finite State Comparison

In the first diagram (*) 20 TCP SYN Scans are reported by an IDS Sensor to a rules based SIEM.

A rules based SIEM will dampen the events based on a rule watching for 5 events to occur in a 20 second interval.


The result is three separate Alarms of the same severity with 5 events each.



In the next sequence (*), the same 20 events are received by a finite state based engine.

With a finite state based machine, a correlation instance is instantiated into memory when the first TCP SYN Scan event is received.

Each event is added to the same instantiation, and the threat level of a single alert is raised and escalated based on system or user defined thresholds as the events are seen. The single alert would include the full history of an event.



In the example above, a rules based engine reduces 20 events to 3 alerts with a resulting 85% correlation efficiency, dropping 5 events for 75% accuracy.

The finite state based engine reduces those same 20 events to 1 alert with a resulting 95% correlation efficiency and 100% accuracy.

While both engines can dramatically increase security efficiency, finite state has clear advantages.

NOC meets SOC

A network operations center (NOC), traditionally monitors SNMP events for traffic flow analysis and infrastructure support. Over the past few years security operations centers (SOCs), are being deployed to monitor and respond to security threats.

Network Monitoring Tools (HP OpenView, OpenService NerveCenter, etc...), can be used to tie the NOC and SOC together. Advanced SIEM tools often integrate with other SNMP management tools using to collect and act on SNMP events. Traditional SNMP events can provide additional insight into internal security threats.

Consider the following examples for correlation:

At 2 am, 5 failed logins are detected, followed by a successful login, and a configuration change to a router or VPN concentrator.

While this may be legitimate, I would want it logged, and checked.

One of your security devices sends out an SNMP power supply or fan failure notification, and the system has only one power supply or fan, and fails open.

Visual Walkthrough: <http://www.openservice.com/hacker2.php> (link should be available soon)

Other common NOC tools for Denial of Service Prevention and Network Flow analysis can also be integrated to provide network visibility and anomaly detection.

A Denial of Service prevention device (Arbor Peakflow, Mazu, etc...), or packet shaping engine (A LOT, Sitara, Packeteer), can be used to gather data (and filter or rate limit), on anomalous traffic on a given port, often detecting a 0 day virus or worm before signatures are released and vulnerabilities become public.

Correlation Example 4

(Network Monitoring feeding Security)

A compromised router/VPN device is used to steal key intellectual property.

A network engineer would change otherwise effective Access Control Lists (ACLs), connect via an intermediate VPN tunnel hiding his actual source IP, and download to an offsite system key intellectual property. ACLs on a border router/VPN device were set such that FTP was blocked in production rules. In order to avoid detection, SNMP (port 161/162), and SYSLOG (port 514), would be blocked at the first router configuration change, and when data transfer was complete, the router configuration would be reset.

Count	Event Thresholds Before Alert by Type
1	Enable Access Granted
1	SNMP Configuration Change
1	FTP during unusual hours (6 pm – 6 am)

All configuration change events outside of production hours were monitored and logged to a centralized SYSLOG collection point. SIEM analytics were applied with correlation algorithms to include SNMP event data. Local IPS data was unavailable during the event (SYSLOG was blocked, and had no local storage), however the packet shaper/QoS device did store and forward SYSLOG data.

Correlation Example (Network Monitoring feeding Security)

Events	Event Description	Correlations Triggered		
1	Router – Enable Access Granted	Source (1) Attacker	Destination (2)	
2	SNMP Configuration Change Alert	Source(3) Router		
3	Unusual port activity (FTP), and unusual time	Source(4) Target	Destination (2) Router/VPN tunnel	User (6)
4	SNMP Configuration Change Alert	Source(3) Router		

Correlation (2) will have a full picture of the attack, having picked up the true attacking IP from event 1, and the user and compromising application from event 3. Individual router events did not raise concern, as they were common, and deemed to be valid, and made by a valid (generic), user. Port activity was detected in Netflow data from another source (bandwidth shaper/QoS device). When reviewed with production rules, FTP was seen to be blocked, and believed to be a false positive on the bandwidth shaper. The destination IP of the VPN/Router was misleading, and the enable password did not point to any specific user. User data was generic, and did not point out the true attacker. Only the initial event (1) provided a source to pursue leading to identification of the attacker.

An employee who had already resigned, but not yet left the company was found to be the culprit.

Summary

As an incident handling tool, SIEM can be highly effective at increasing a security staff's ability to identify and handle a large number of events while simultaneously making them more effective. By consolidating and correlating events, a SIEM product can spot attacks that would otherwise go unnoticed. Cross correlation of multiple devices can improve the accuracy of threat identification while effectively joining devices together in a conflagration that must be defeated as a whole rather than serially. A SIEM can aid in integrating traditional network management tools from one's NOC into a SOC increasing the security team's effectiveness in detecting and responding proactively to internal security threats. With the push for compliance, the need for centralized logging, and the pressures on small security staffs to deal with an avalanche of events, a SIEM tool should be part of any enterprise security solution design.

Sincerely,

David Swift dgswift@verizon.net
GSNA, GCIH, CISSP, MCSE, MCNE, AIX-CSA, SUN-CSA, CCNA

© SANS Institute 2007, Author retains full rights

Acknowledgements

Editing, suggestions and comments

Geoffrey Coulter CTO OpenService

Correlation Techniques

Jason Baroush Systems Engineer OpenService

Finite State Graphics

Scott Kappel Systems Engineer OpenService

Graphics and Release

Mike Schmitt CEO OpenService,

and. OpenService Marketing documents

Thank you each for your contributions.

© SANS Institute 2007, Author retains full rights.

Appendix A - SIEM Vendors

Arcsight	www.arcsight.com
EQ	www.eiqnetworks.com
E-Security	www.novell.com
Intellitactics	www.intellitactics.com
Network Intelligence	www.network-intelligence.com
OpenService	www.openservice.com
Sensage	www.sensage.com
Symantec	www.symantec.com
TriGeo	www.trigeo.com

Be prepared for a starting price at around \$900 device you want to take log input from (\$20,000 is a realistic entry level). For an enterprise deployment, the budget needs to start at \$75,000, and may be in excess of \$500,000 or more for a fortune 500 company.

* beside multiple graphics denotes images borrowed with permission from OpenService

© SANS Institute 2007, Author retains full rights.

Request for Input

I would like to see this body of work grow. SANS is a community of highly intelligent security focused professionals, and I'd like you're input.

Is anyone using a free SIEM tool that you like and would recommend?

In the spirit of SANS, I'd like to be able to point to free tools for student CDs.

Help complete the list of vendors - if you know of a good SIEM tool please submit it for addition.

Please list strengths and weaknesses.
Please include a current URL

Related topics that could use input:

Windows / Active Directory events of interest for security
i.e. Using AD objects and events 560 & 567 for Compliance reporting resolving SIDs in AD events to machines using DDNS, and MAC addresses to locate DHCP systems, and the user logged in at the time of an event
Unix - Linux/AIX/HP-UX/Solaris - events of interest and what to correlate

Effective Correlations discovered/created in the field.



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS 2018	Orlando, FLUS	Apr 03, 2018 - Apr 10, 2018	Live Event
SANS Abu Dhabi 2018	Abu Dhabi, AE	Apr 07, 2018 - Apr 12, 2018	Live Event
Pre-RSA® Conference Training	San Francisco, CAUS	Apr 11, 2018 - Apr 16, 2018	Live Event
SANS London April 2018	London, GB	Apr 16, 2018 - Apr 21, 2018	Live Event
SANS Zurich 2018	Zurich, CH	Apr 16, 2018 - Apr 21, 2018	Live Event
SANS Baltimore Spring 2018	Baltimore, MDUS	Apr 21, 2018 - Apr 28, 2018	Live Event
SANS Seattle Spring 2018	Seattle, WAUS	Apr 23, 2018 - Apr 28, 2018	Live Event
Blue Team Summit & Training 2018	Louisville, KYUS	Apr 23, 2018 - Apr 30, 2018	Live Event
SANS Riyadh April 2018	Riyadh, SA	Apr 28, 2018 - May 03, 2018	Live Event
SANS Doha 2018	Doha, QA	Apr 28, 2018 - May 03, 2018	Live Event
SANS SEC460: Enterprise Threat Beta Two	Crystal City, VAUS	Apr 30, 2018 - May 05, 2018	Live Event
Automotive Cybersecurity Summit & Training 2018	Chicago, ILUS	May 01, 2018 - May 08, 2018	Live Event
SANS SEC504 in Thai 2018	Bangkok, TH	May 07, 2018 - May 12, 2018	Live Event
SANS Security West 2018	San Diego, CAUS	May 11, 2018 - May 18, 2018	Live Event
SANS Melbourne 2018	Melbourne, AU	May 14, 2018 - May 26, 2018	Live Event
SANS Northern VA Reston Spring 2018	Reston, VAUS	May 20, 2018 - May 25, 2018	Live Event
SANS Amsterdam May 2018	Amsterdam, NL	May 28, 2018 - Jun 02, 2018	Live Event
SANS Atlanta 2018	Atlanta, GAUS	May 29, 2018 - Jun 03, 2018	Live Event
SANS London June 2018	London, GB	Jun 04, 2018 - Jun 12, 2018	Live Event
SANS Rocky Mountain 2018	Denver, COUS	Jun 04, 2018 - Jun 09, 2018	Live Event
DFIR Summit & Training 2018	Austin, TXUS	Jun 07, 2018 - Jun 14, 2018	Live Event
SANS Milan June 2018	Milan, IT	Jun 11, 2018 - Jun 16, 2018	Live Event
SANS ICS Europe Summit and Training 2018	Munich, DE	Jun 18, 2018 - Jun 23, 2018	Live Event
SANS Crystal City 2018	Arlington, VAUS	Jun 18, 2018 - Jun 23, 2018	Live Event
SANS Oslo June 2018	Oslo, NO	Jun 18, 2018 - Jun 23, 2018	Live Event
SANS Cyber Defence Japan 2018	Tokyo, JP	Jun 18, 2018 - Jun 30, 2018	Live Event
SANS Philippines 2018	Manila, PH	Jun 18, 2018 - Jun 23, 2018	Live Event
SANS Boston Spring 2018	OnlineMAUS	Mar 25, 2018 - Mar 30, 2018	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced