



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

A Process for Continuous Improvement Using Log Analysis

GIAC GSEC Gold Certification

Author: David Swift dgs swift@verizon.net

Advisor: Johannes Ullrich

Accepted: 9/12/2011

Abstract

Good security is a moving target. Walls and castles were once good defenses against attackers, but they stand as little chance of preventing an attack by a modern army. Like all defenses, if left unattended, any information security strategy will become obsolete and fail. The problem with building or improving a defense strategy is where to start. Our knowledge and defenses are seldom perfect. More often than not the task of securing a network is gargantuan, and daunting. A good logging and analysis strategy can point the way. By accepting that defenses and configurations are never perfect and ever changing and by analyzing input from the event sources we already have, we can detect threats, direct responses, and tune our defenses. In the paper that follows, a repeatable process for continuously improving security and an outline of log analysis with case studies and sample output based on actual data will be detailed. The process is broadly applicable, and does not require a Security Information and Event Management (SIEM) or centralized log management (LM) system, though they do make the process easier.

1. Introduction

A great deal of money has been spent by organizations on security technology, with only moderate success. Technology is often installed, but often left untuned and unmonitored. Though vendors have touted self-defending networks (Gleichauf, 2005), and claimed their products are impervious, reality teaches otherwise. The existing state of the art security requires people and processes, and if left untended will become practically useless.

We are in a virtual arms race to stay ahead of our attackers and the permeability of our security defenses is directly related to the skill and effort applied to them. New threats, exploits and evasion techniques are constantly being developed. Signature and firmware updates are not enough. With any given new signature, the standard default action is to alert, not to block. We must continue to analyze and tune, as well as patch and update, or security will fail.

And even then, there is no guarantee that you will detect or block every threat no matter how good your defenses are, but by applying the process detailed in the paper that follows you will detect a higher percentage of threats, and you will decrease the number of systems exposed or compromised. This is not conjecture. The reports reviewed in this paper are not theoretical; they are examples from real world results at multiple real world organizations.

Many organizations including SANS document what to log, list top 10 reports, and provide analysis based on specific threats, however there are few papers that document the process of log analysis for general and emerging threats. In the paper that follows, general principles that can be applied to detect both existing and emerging threats are detailed and real world examples and responses are presented to show how defenses can be improved and responses can be targeted.

Alternative log management processes are available from the National Institute for Standards and Technology in the 800-92 publication (NIST, 2006), and several respected authors including Anton Chuvakin, Raffael Marty and Lenny Zeltzer are publically available. SANS also offers a course and annual Log Management Summit meetings on the topic as well.

The following basic and repeatable process can direct your security efforts to the highest threats, and continuously improve your defenses.

The steps are:

- A. Run monthly reports
 - a. Simple top 10 reports can provide a starting place.
 - b. Top 10 events grouped by unique signature/username/source.
- B. Analyze the monthly reports
 - a. Get to know the events and sources, and build that knowledge into the reporting tool your using.
 - b. Build white lists for known normal approved traffic. (see [Appendix A](#))
 - c. Build black lists for known attackers.
- C. Improve your defenses
 - a. Scan, update, and clean suspect hosts.
 - b. Add Blocking Rules to firewalls, IDS, HIPS and other active defenses for known offenders, and signatures with low false positive rates.
- D. Repeat

2. Reports

With only two basic reports, and a bit of applied logic and analysis, security professionals can continuously improve defenses.

For each device logging events in your network, generate a report that:

1. Counts (SQL SUM) by signature (or address, or user name), how many events occurred per month (or period of time).
2. Counts by unique sources or destinations how many threatening events occurred per month (or period of time).

Applying limits to show the top 10 events is a common way to bring the scope down to a reasonable actionable list, and can be used for management summary reports. One should not infer only the top 10 matter, the limit is arbitrary, and is just to provide a minimum starting point and reasonably digestible report for management.

These same two basic reports can be applied to nearly any log source, and several simple derivative reports can provide even more targeted analysis.

The type of report, and SQL “group by” options will be determined by the log source.

It is possible to create reports in a variety of tools. Most IDS/IPS products, Anti-Virus frameworks (i.e. McAfee EPO, Symantec SEP...), and other point solutions can be used to generate product specific versions of each report. With centralized log management tools (Splunk, Syslog-NG, LogLogic...), or a SIEM (ArcSight, Qradar, NitroView, EnVision...), reporting can be done for multiple devices in the enterprise and across multiple disparate point solutions much easier.

The tool and level of effort will vary, but the process is the same.

David Swift dgswift@verizon.net

2.1. Common Log Sources and Corresponding Reports

1. Firewalls

Firewall logs typically include only source and destination IP addresses and corresponding port pairs. These tuples by themselves provide limited data. However by grouping the data appropriately we can look for several potential threats.

Reports:

- a. Group by unique IP source address
to identify top talkers, possible servers, scanners, or attackers.

Response:

Investigate, add white lists for approved servers to remove common normal traffic from future reports, and disable services on unapproved servers.

1. It is not uncommon to find hosts serving traffic for applications that are not approved by systems that haven't been effectively hardened.
2. A common misconfiguration and vulnerability is to allow connection to HTTP (TCP 80), or SQL (TCP 1433), from any remote host on a system not intended to be a server.

Add blocking rules for scanners, and port probes.

3. Source IP addresses attempting connection to incrementing destination addresses on the same IP address are easily distinguishable as hosts performing reconnaissance.
4. Source IP addresses attempting connections to incrementing (or common ports), on the same IP address are also easily distinguishable as host scanners.

- b. Group by unique IP destination address
 - i. Top destinations, these may be servers, external business partners or compromised hosts.

Response:

1. Investigate, and add white lists for known valid servers.
2. Disable services, or block access to unauthorized servers.

David Swift dgswift@verizon.net

When possible, overlay GEO IP data, and consider denying traffic to untrusted countries. Also consider comparing against common known malware lists of command and control, botnets, and known attackers.

2. IDS/IPS/HIPS/AV Devices

Intrusion detection systems and intrusion prevention systems, are often poorly configured. By analyzing the volume and variety of signatures and hosts reporting in in a given month, we can identify signatures that could be blocked, and signatures that are not being triggered.

One proactive approach is to block all signatures that have not triggered in the past 30 days. These may add only minor value, but if we know in a given network that the signatures would not have blocked legitimate traffic, blocking going forward can provide incrementally better security with little risk.

Reports:

- a. Group by signature
 - i. A simple analysis of how many of each signature occurred can show which threats are most prevalent and need corrective actions.
 - ii. When a given signature count is exceedingly high, it can indicate a poorly written signature that can be disabled.
 - iii. When a given signature count is low, it can indicate a candidate signature for blocking going forward.
- b. Group by signature and unique source address
 - i. Signatures with a high number of unique IP addresses often indicate common normal network traffic causing false alerts.
 - ii. Signatures with a low number of unique sources and threatening behavior should be investigated as possible infections. A common false positive may be to discover servers of specific types (Mail, Database, File Service...).
 1. If the service is unauthorized on the host, it can be disabled or blocked.
 2. If the service is authorized, an asset tag/identification label can be added and the specific signature for that specific IP can be filtered or white listed.

A bit of basic analysis can yield a short list of signatures that should be blocked to prevent infection (few unique host IPs trigger the signature, and one or more was confirmed to have malware). When done each month in a systematic approach, blocking more each month with predictive results yields improved defenses.

3. Authentication Sources (Active Directory, TACACS, RADIUS...)

While we often have many authentication sources, watching for hosts with repeat failed login attempts can find service misconfigurations and possible brute force attacks. Alternative analysis of target accounts with a high number of failed logins can help identify accounts that have been harvested and are being targeted.

Reports:

- a. Group by unique target user name
 - i. These can be further analyzed by limiting the scope to only:
 1. Default User Accounts (admin, root, guest...)
 2. Administrator accounts (all admin equivalents)
 3. Service accounts
- b. Group by unique target user name and authentication type
 - i. These reports can be used to track inactive accounts
 1. Any active user ID that has not logged in during the previous period can be disabled.
 2. Consider running the report for 90 days, or over a time period to match your policy on inactive accounts.
 - ii. Providing these reports to the respective administrators for each authentication source can be used to check for
 1. Active accounts that should be disabled (terminated employees, contractor accounts active after the contract expires...).
 2. These reports are of particular importance for VPN authentication sources.
- c. Group by unique target user name, authentication type, and source address.
 - i. These reports can be used to track admin account, and service account usage on unauthorized hosts.

4. Flow Analysis

In situations where flow data is available (Netflow, JFlow, SFlow, QFlow, VFlow), host profiling, by port and application can provide another vector to watch for anomalous behavior and misconfigured systems.

A simple report listing all unique IP addresses where the source port is one of the known common ports can identify unauthorized and/or compromised hosts.

Reports:

- a. All hosts by unique IP address where source port is a common well known port (20,21, 22, 23, 25, 110, 2049, 445...). Ports under 1024 are reserved by most operating systems for administrator/root level services.
- b. Traffic from internal sources not on the approved white list matching common server ports.

Desktops and servers not under IT's control and approval, should not be serving applications (HTTP/TCP 80, HTTPS/TCP 443, File Service/TCP 445 or TCP 2049, Mail/TCP 25 or 110, FTP TCP 20 or 21, Telnet/TCP 23, SSH/TCP 22), and monitoring for the corresponding port activity internally can identify misconfigured or poorly configured hosts, and potentially compromised systems.

Often times exfiltration of data will occur on encrypted ports (SSH/22, HTTPS/443), and look as if the compromised system is acting as a server. Monitoring for unapproved use of server side traffic (Source Port), matching common services can help identify misuse and misconfiguration of systems.

3. Analysis

Taking the reports we've collected, a bit of analysis, intelligence, and investigation will need to be applied. Consider the volume and variety of events.

Consider the total volume of events.

- If the event count is higher than normal, why? And when did it start? (review volume charted by day).
- If the even count is low, check the source device.
 - Is it working properly, and are we receiving events into the reporting tool?
 - Is the logging policy inclusive enough?

Consider the variety of events.

- If the number of unique event types is low, you may want to increase the logging policy of the source (turn on more signatures/logging).

Consider how many of each event occurred in a given month.

- If the volume is high and the number of unique sources is high, perhaps the event is prone to false positives, and logging for that event should be disabled, or excluded from the report.
- If the event volume is high, but the number of unique sources is low, it may be a legitimate indication of malware, or may be a sign of normal network behavior that should be white listed. Investigate the source.
 - DNS lookup
 - WHOIS lookup
 - Determine where the device lives in your network
 - Traceroute
 - Show ARP
 - Review network diagrams
 - Contact the asset owner, and discuss what you're seeing.
 - Consider full forensics if it's serious enough.
 - Update Anti-Virus signatures and patches, and rescan.
 - Reboot the host from a bootable image, and scan the host for malware.

As you research any given event, augment your tools with the information. Don't repeat your work. Add labels with easily identifiable names overlaid on network and asset objects. Add filters for known attackers (black list), and normal traffic (white list).

David Swift dgs swift@verizon.net

4. Log Review

Analyze the volume and variety of events, sources and signatures.

How many of each event occurred this month?

- If the number is small, that event/signature can be blocked without detrimental impact to the network.
 - If the number is high, consider disabling the signature, it may be prone to excessive false positives in your network.
1. Analyze the reports.
 - a. Which events are most prevalent? (volume)
 - b. Which events are most threatening? (where the signature name contains WORM, VIRUS, BOT, MALWARE, W32, Buffer Overflow...)
 - c. How many unique source addresses triggered this event? (who to clean)
 2. Investigate
 - a. Check your conclusions.
 - i. Verify the given malware detected, or unauthorized traffic detected is emanating from the source.
 - ii. If not, the signature may be a false positive.
 - b. If in fact you do find malware, expand your search and report to show any IP address with the same symptom, or communicating with the same destination (other possible infected hosts).
 3. Tighten your defenses.
 - a. Implement a control (Blocking rule on a firewall, IPS or HIPS...), to block low volume threats.
 - b. Disable unnecessary/unauthorized services.
 4. Improve your reporting.
 - a. White list normal traffic that has been identified by reports.
 - i. In a SIEM, this can be applied by source IP and port to make the white list very specific for only known allowed traffic to/from a specific host.
 - ii. In an IPS, HIPS, or other device consider turning off logging for signatures with excessive false positives.
 - iii. For firewall events, filter out accepts from known source addresses on approved ports.

Review [Appendix B](#) for a list of suggested reports.


David Swift dgswift@verizon.net

5. Case Studies

Customer IP Addresses have been changed to protect the guilty, though some public IP addresses, and RFC1918 NAT addresses were left unaltered. Customer names have been removed, however the data used is real, and these reports were presented and reviewed with management. Most of the data collected for the reports below was gathered a SIEM tool, exported as CSV data and then manipulated into graphs and tables using a spread sheet and formatted and annotated for presentation a presentation tool.

5.1. Firewall & SIEM

Top 10 Source Countries



Top 10 Foreign Attackers

109.72.146.154	Europe	15393914
188.72.213.59	Belarus	2977444
91.212.135.186	Russian Federation	2699576
91.212.135.136	Russian Federation	2249550
91.205.41.235	United Kingdom	1758810
193.104.12.102	Panama	1375574
64.71.246.28	Canada	1056001
91.205.41.164	United Kingdom	1042430
24.153.22.142	Canada	996563
109.72.146.155	Europe	962600

A large number of the attacks, are targeting DNS (port 53), and may have been exploiting previous weaknesses now patched with Windows 2008 upgrades to DNS Servers and Domain Controllers.

Figure 1: Monthly Summary Report – Foreign Attackers


This report summarizes the traffic to and from foreign countries. While not relevant to every organization, foreign attacker reports are usually popular and generate interest at all levels and offer a good ice breaker and discussion point for a monthly status report. If your reporting devices support GEO IP, or some auto-overlay for country, consider including this report.

At one financial services organization only doing business in the United States, this report, with corroborating data from netflow events, showed a large volume of data was actually being

David Swift dgswift@verizon.net

transferred to foreign countries. A project to replace the organizations firewalls was immediately launched.

Infrastructure Security



4

Figure 2: Top 5 Unexpected Sources and Destinations

This report lists the top five sources and destinations based on volume of events. The host list was highly filtered with white lists for known servers and part of a mature SIEM deployment. By looking for high volumes from unknown sources, both unauthorized servers, and compromised hosts can be found.

In the top five sources, three proved to be infected systems. At the time this report was generated, Conficker and Virut N still had no IPS or AV signatures available. Detection through firewall events helped the security team build a defense response that included adding firewall rules and emergency patching via WSUS, and kept the spread of the infections to less than 30 systems on a 50,000 node network.

David Swift dgswift@verizon.net

© 2011 SANS Institute, Author retains full rights.




Figure 3: Monthly Summary Report – SIEM Overview

Each month, the analysis tool itself must be analyzed for effectiveness and consistency and should be reported on as in the above slide.

In this report by watching for anomalous event rates we can detect when data collection failed (in this case the SIEM was just installed on the 22nd of the previous month, and there is no data from the 12th through the 22nd).

Additionally, key correlation rules are seen to fire which is a good indication events are flowing and the rules are working, however, “Foreign Source in Network” is firing so frequently as to drown out other more threatening events and needs to be tuned. A white list of allowed networks and better perimeter ACLs to prevent access by foreign sources are good next steps.

The chart above showing events per day and correlated attacks can be difficult to generate without a SIEM or log management solution. However the essential Top 10 Sources and Destinations can be produced by most devices.

David Swift dgswift@verizon.net


A high volume of DNS traffic to 4.2.2.2 pointed out an opportunity to install a caching DNS server and improve overall web response for this customer.

A review of top 10 sources at one healthcare customer showed a spike of traffic from a medical records desktop in Houston. The host had been compromised and was serving pornography for a site hosted out of the Netherlands.

David Swift dgswift@verizon.net

5.2. Anti-Virus/Anti-Spyware

Infrastructure Security



7

Figure 4: Top 10 Report – Anti-Virus/Anti-Spyware

For each log source we're responsible for a monthly summary, such as the one above, can provide oversight to ensure the tool was effective, and offer insight into other areas for improvement.

This report was a follow up the month after firewall example 1, and showed the results of new signatures applied the customer's anti-virus client. The spread of conficker, sality, and virut n, were contained. No specific tuning was applied as a result, however this report helped management understand why emergency changes had been done two months prior, and justified the investment and continued push to the anti-virus client installed on every desktop and server.

David Swift dgswift@verizon.net

Anti-Virus (AV)




Figure 5: Monthly Report – Anti-Virus/Anti-Spyware

The report above is supplied to show a normal month. In this report there are no anomalies to investigate. Malware incident rates are low, events per day mirror a typical seven day pattern for this customer's network. In this instance both volume, event rates and unique IP addresses for which there were events, and variety, the number of unique signatures, were as expected for the network being reported on. A good operational report can be generated using the top 10 threats using a SQL "group by" option on the signature name, "group by" on the source IP address, and a "unique" function applied to the source IP address. A list of the few hosts that triggered the top threats should result allowing remediation efforts to focus on just those hosts that were infected.

David Swift dgswift@verizon.net

5.3. HIPS

Host Intrusion Prevention (HIPS)

HIPS is currently running less than 200 events/day.

>99% of all events are when a user is prevented from running a forbidden application (Lotus Notes).

Turning on additional HIPS signatures and capabilities (IPS), should be investigated, and reporting added as volume increases

Observed HIPS Events of concern:

Adobe Reader Plug-in Cross-Site Scripting Vulnerability

Sticky Keys File Replacement Backdoor

Sun Java WebStart JNLP Stack Buffer Overflow Vulnerability

These events are currently **permitted**.

Recommendation: Build/Enable a HIPS blocking policy for identified threats

Step 1: Put all events in Log only mode for some period (1 week – 30 days)

Step 2: Monitor events for frequent false positives

(ArcSight reports by Device Event Class ID and Name)

Step 3: Turn on additional blocking monthly with top identified threats (Repeat Monthly)

Figure 6: Monthly Report – Host Intrusion Prevention System


In the HIPS report above, both volume and variety were significantly outside of expectations. In terms of volume, fewer than 200 host IPs in a network of over 3000 hosts logged a single event. In terms of variety, of a possible signature set of over 4,000 signatures fewer than 20 had any events for the month. For this financial services customer, a report showing that very few hosts were actively using the Host Intrusion Prevention System (HIPS), module was unexpected. Even more concerning was the detection, but failure to block of multiple threats. A project was launched shortly after this report to improve policies and push the HIPS agent to all desktops, with servers to follow.

Log analysis at multiple customers has shown this failure to deploy HIPS/HIDS agents is more common than not.

David Swift dgswift@verizon.net

Infrastructure Security

Top 10 Blocked Events



- Vulnerability in Server Service Could Allow Remote Code Execution
- svchost Buffer Overflow (RPC DCOM)
- MSSQL Resolution Service Buffer Overflow (Slammer)
- COM Object Instantiation Memory Corruption Vulnerability
- Double File Extension Execution
- RPC DCOM Stack Buffer Overflow (Blaster, Nachi)
- Windows ASN.1 Heap Overflow Vulnerability
- Generic Buffer Overflow
- LSASS Dcpromo Log File Buffer Overflow (Sasser)
- IIS IPP .printer Buffer Overflow

Description: A list of the top most active HIPS signatures with a short description.

An increasing number of patterns and corresponding events are blocked each month with ongoing HIPS tuning.

HIPS Blocked Events

Vulnerability in Server Service Could Allow Remote Code Execution	57,064
svchost Buffer Overflow (RPC DCOM)	1,708
MSSQL Resolution Service Buffer Overflow (Slammer)	1,586
COM Object Instantiation Memory Corruption Vulnerability	1,022
Double File Extension Execution	227
RPC DCOM Stack Buffer Overflow (Blaster, Nachi)	87
Windows ASN.1 Heap Overflow Vulnerability	57
Generic Buffer Overflow	30
LSASS Dcpromo Log File Buffer Overflow (Sasser)	14
IIS IPP .printer Buffer Overflow	5
Total	61,800

6

Source: SIEM HIPS Connector (counts reflect the number of unique hosts reporting events in a 24-hour period)

Figure 7: Monthly Report - Host Intrusion Prevention Systems

The report above was from the third month of a HIPS deployment, and proved useful in tuning, by pointing out signatures that could be blocked without major network disruptions. In this instance, a healthcare customer, took a very conservative approach, as deployment of the HIPS agent was widespread (>30,000 hosts). Signatures with fewer than 2,000 events/month and reported by fewer than 10 hosts were modified from log to block. Each month this report was used to iteratively identify and block additional signatures improving the overall preventative capabilities of the HIPS deployment.

David Swift dgswift@verizon.net

5.4. IDS/IPS




Figure 8: Monthly Summary Report - NIPS

It's a maddening exercise to summarize millions of events into consumable chunks, but we should strive to produce summary reports that are meaningful and digestible by management. To that end, one should consider using graphs and summarizations of no more than top 10 events. Though busy, the slide above condenses a great deal of information into a single slide.

Follow up investigation of the Network Intrusion Prevention System (NIPS), top blocked signature for the month revealed "SHELLCODE: *" as a common false positive, and the signature was disabled. The signatures "NetBIOS-SS: Microsoft Server Service Remote Code Execution Vulnerability" and "SMB:NTLMSSP_AUTHUnauthorized..." were also noted as prone to false positives and alerting was filtered to ignore these events when the source and destination of the event were both internal. These signatures did prove useful in discovering previously undocumented internal networks, and were alerting and blocking when the destination was not an internal address.

David Swift dgswift@verizon.net

Top 10 IPS Events

Blocked Attacks	Count
WORM: W32/Netsky@MM Worm	357
WORM: W32/Netsky@MM Worm Variants II	154
WORM: W32/Zafi@MM Worm	92
SMTP: Incorrect MIME Header with Executable Attachment Found	47
SHELLCODE: Shellcode Detectedfor HP PA-RISC Family CPUs	32
DCERPC: Microsoft RPCSS Heap Overflow I	32
HTTP: Microsoft Frontpage fp30reg.dll Buffer Overflow	21
DCERPC: Microsoft RPC DCOM Buffer Overflow	12
WORM: W32/Mydoom@MM Worm Variants IV	7
DCERPC: Microsoft Workstation Service Buffer Overflow	6
DCERPC: Microsoft Windows LSASS Buffer Overflow	5
BACKDOOR: Web Serve CT Backdoor	4
HTTP: WebDAV Method URL Overly Long	3
SMTP: TURN Command	2
SMB: NLTMSSP_AUTH Unauthorized ChangeServiceConfigW Request	2
DCERPC: Microsoft Plug and Play Service Buffer Overflow	1
PCANYWHERE: Host Logon Engine Buffer Overflow	1
Total:	778

Figure 9: IDS/IPS Data




Figure 10: IDS/IPS Pie Chart

When using data analysis tools, you may also need to work around their reporting limitations. I've found many of the best SIEM and LM tools have poor visualization tools. However, in each case, there has been a way to output the text to CSV, XLS, or another raw text format.

By manipulating the data to produce graphs in common desktop tools (spreadsheets, and presentation tools), one can then produce more readily consumable reports and in a format that executives can cut and paste from or manipulate on their own. Output to HTML, PDF or other forms that are difficult for most people to manipulate often leave the audience frustrated. By converting to common desktop application format the data will be much more accessible to the audience.

The data in figures 9 and 10 were generated using a SIEM tool, output to CSV and manipulated in a spread sheet, then imported into a presentation tool to generate a monthly dashboard for executive review. The same data was extracted via the NIPS vendor supplied reporting tool, and used to validate accurate and complete data was being collected.

During initial setup of a new SIEM/LM and reporting, as a best practice, I create a spread sheet and presentation tool template that will auto generate the graphs to be used. Each month the raw data can be cut and pasted into the template and presentation with relatively minimal effort.

It is important when using any secondary analysis tool (SIEM, Log Management...), that may have filtering, aggregation or other manipulations applied that you validate the data against the raw source data until you are confident you are receiving all relevant events.

By documenting and validating reporting and monitoring of blocked events, management gave approval to increase the number of signatures in IPS/Blocking mode.

For reasons that should be obvious, I cannot share detailed reports showing actual systems. But suffice it to say that anomalous or concerning events for each report were followed up on with detailed reports to show which systems were showing symptoms, and then investigated.

5.1. Authentication

Infrastructure Security




Figure 11: Monthly Report – Failed Authentication

Authentication reports such as the one above are intended to identify the targets of attacks, and the sources of those attacks. By reporting on the top sources of failed logins, we can identify brute force attacks, and misconfigured services. By reporting on the top target accounts, we can identify accounts that may have been previously harvested, or those that are the target of brute force attacks.

In the example above the (redacted) System account had thousands of failed logins as a result of conficker, causing the system account to be locked out repeatedly. Tracking those sources, we were able to clean multiple hosts. The (redacted) cerner account failures were due to a service misconfiguration, and the account on the source host was adjusted correcting the issue and bringing a failed business critical service back online.

David Swift dgswift@verizon.net




Figure 12: Monthly Report - Authentication

The report above was generated in a new installation at a banking customer's network.

In this report, the use of the guest account had to be traced. The guest account was believed to have been disabled.

According to Microsoft's [Well-known security identifiers in Windows operating systems](#) (Microsoft, 2010), S-1-15-18 is the local system account, and the events indicate devices that were not properly joined to the domain.

The high number of SYSTEM and Administrator failed logins are also concerning. No details are available regarding the findings.

David Swift dgswift@verizon.net

6. Measuring Value

Multiple formulas and models exist for measuring the cost of malware, including Time Based Security (TBS), Annual Loss Expectancy (ALE), Single Loss Expectancy (SLE) formulas. An interpretation of ALE and SLE formulas, (Gregg, 2005), are summarized below.

Annual Lost Expectancy Formula:

$$\text{ALE} = \text{ARO} \times \text{SLE}$$

- ALE (Annual Loss Expectancy)
 - ARO (Annual Rate of Occurrence)
 - SLE (Single Loss expectancy)
-

$$\text{SLE} = \text{EF} \times \text{AV}$$

- EF (Exposure Factor)
 - AV (Asset Value)
-

Unfortunately a good deal of subjective value must be placed on the asset with ALE formulas, and often times the asset is not of value in and of itself, but rather is a PC required for someone to be able to do their job.

As a more accurate measure of actual cost impact, consider substituting the cost of remediation in actual man hours and dollars associated those hours.

Ideally an infected host should be reformatted and a clean hardened patched image installed on the host. In many corporate networks, in order to reduce downtime a more abbreviated process is often followed.

For simple issues on machines that do not contain sensitive data, the basic process is:

1. Boot the host from a clean CD and run an AV scan and Rootkit scan. Remove any infections found.
2. Update the OS and Application Patches on the System
3. Update Anti-Virus Signatures
4. Ensure the system is participating in patch management processes (EPO, WSUS, Bigfix...).
5. Check for and re-apply any system hardening/configuration standards.

Typically this means the host is out of service for around two hours while a technician with some security experience is testing and checking the system. I substitute a standard rate of \$100/hour

David Swift dgs swift@verizon.net

for the IT Tech's time, and ignore the lost productivity for the employee since the actual measure is typically quite subjective (can that person do other work while the system is down?).

I've also found the Exposure Factor (EF), is subjective, and usually substitute a threat multiplier (TM), where $TM = \text{the likelihood that the infection will spread} * \text{the number of systems likely to be infected as the infection spreads}$.

This too is of course somewhat subjective, and requires experience and judgment.

Consider:

- A worm could spread rapidly, particularly if it had infested a system inside an organization that only has an internet perimeter firewall, and no IPS.
- If the worm spreads via email and email is scanned and has an appropriate signature for the threat, it may not spread.
- A virus would spread more slowly if on a local host, and more rapidly if on a network share.

Time based security principles (TBS), directly indicate the need to continuously improve defenses. From time based security, (SANS, 2006), we can derive the following:

Time Based Security (TBS)

- Basic principle: Effective security measures are those where protections last longer than the time to detect a threat plus the time to remediate that threat.
-

$MTP > MTD + MTR$

- MTP (Mean Time to Protect)
 - MTD (Mean Time to Detect)
 - MTR (Mean Time to Repair)
-

In practice, TBS can be applied to help track the results of log analysis and remediation efforts.

- Establish a baseline for time to detect and remediate threats.
- Monitor and track the mean time to detect, and the number of systems (multiply by your mean time to repair in hours or dollars).

Over time, your MTD, and MTR should decrease.

David Swift dgswift@verizon.net

I've gone as far as calculating ROI per month for preventative controls, however these values are frequently met with skepticism.

ROI Calculations

	NIPS	HIPS	SurfControl	AV
Monthly Blocked Counts	690	82,844	104,971	270,889
Probability of Infection	1%	3%	1%	5%
Spread Multiplier	15	1	1	3
False Positive Rate	35%	75%	90%	95%
Repair Cost (2 hours@\$100)	200	200	200	200
Lost Productivity (Hours)	2	2	2	2
Value Per Block or Clean				
\$	19.5	1.5	0.2	1.5
Hours	0.195	0.015	0.002	0.015

Calculated Value				
\$	13,455	124,266	20,994	406,334
Hours	135	1,243	210	4,063

Figure 13: Monthly ROI Calculations

The probability that any given log event represents an infection must be reduced to eliminate duplicate log events, low quality events (i.e. PUP: cookie deleted), and threats that though valid are targeting non-vulnerable systems. When accounting for an infected host, each host should only be counted once per unique infection. Since each event source will have a differing quality of signatures, and varying levels of repetitive events each blocking source is reduced by a false positive rate to discern actual infection rates.

In addition, each infected host may in turn infect other hosts, and an appropriate threat multiplier should be applied to account for secondary infections. The threat multiplier will vary by event type and source. For network based malware such as worms that would be detected and blocked by NIPS, the infection multiplier is high, as these types of mobile malicious code attempt to spread via network means. For other sources such as anti-virus, a single file may be infected, and

David Swift dgs swift@verizon.net

unless that infected file is opened by another user, only the single file and host will be affected and the corresponding threat multiplier is low.

Adjust these values (false positive rate, spread multiplier, cost/hour, and the probability of infection), to your network. When the signature quality is high, or the patch rate is low, the probability of infection will be higher. When the number of duplicate events for a log source is low, or few benign meaningless events clutter the logs, the false positive rate will be low.

One application of less subjective IT only costs to valuation:

During a two year time frame at one customer, using this process, the security team reduced the average number of hosts infected per incident from over 300 (1% of the total population) to less than a single subnet (<15 hosts/incident). The average time to detect malware went to near zero day, typically 3-7 days (down from nearly 30 days), post rumor and questioning on the internet about new symptoms, and from one to four weeks sooner than IDS an AV signatures were released.

Considering the cost per infection involved two hours of an engineer's time at \$100/hour, the average cost dropped from \$60,000 to \$3,000. And in the reports above you can see that Conficker, Sality, and Virut N all hit in the same year.

David Swift dgswift@verizon.net

7. Conclusion

Our network defenses are unlikely to be perfect, and our knowledge of the network and assets we protect is often less than complete, but you can always make improvements.

You may be surprised at how far your overall defense posture will improve using some simple log reporting and analysis, and acting upon it each month.

In a best case scenario, should you fail to find anything to tune, you will have a report to show to management (and the auditors), that your defenses are effective and well deployed, and may even be able to calculate the value of security services for management and budgeting.

The average network is in a constant state of flux. Change is the only constant. To adapt and keep up, we must continuously teach our system about those things we discover, and modify our defense accordingly.

David Swift dgswift@verizon.net

8. Disclaimer

Data in reports has been intentionally modified to remove site specific information, and does not represent any individual or organization, and none should be inferred.

9. References

- Chuvakin, Anton. (2010), *Open source and free log analysis and log management tools*
Retrieved from <http://www.securitywarriorconsulting.com/logtools/>
- Cisco. (2008). *Self Defending Networks*. Retrieved from
http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns171/net_brochure0900aecd800efd71.pdf
- Clausius, R. (1865). *The Mechanical Theory of Heat – with its Applications to the Steam Engine and to Physical Properties of Bodies*. London: John van Voorst, 1 Paternoster Row. MDCCCLXVII.
- Gleichauf, Bob. (2005). Cisco Press, *Core Elements of the Cisco Self-Defending Network Strategy*. Retrieved from <http://www.ciscopress.com/articles/article.asp?p=379750>
- Gregg, Michael. (2005). *CISSP Security-Management Practices*. Retrieved from
<http://www.informit.com/articles/article.aspx?p=418007&seqNum=4>
- Ken, Karent & Souppaya, Murugiah. (2006), *NIST (800-92) Guide to Computer Security Log Management* Retrieved from <http://csrc.nist.gov/publications/nistpubs/800-92/SP800-92.pdf>
- Marty, Raffael. (2010), *Maturity Scale for Log Management and Analysis* Retrieved from
<http://raffy.ch/blog/2010/06/07/maturity-scale-for-log-management-and-analysis/>,
- Microsoft. (2010). *Well-known security identifiers in Windows operating systems*. Retrieved from
<http://support.microsoft.com/kb/243330>
- NIST. (2006). *Guide to Computer Security Log Management* Retrieved from
<http://csrc.nist.gov/publications/nistpubs/800-92/SP800-92.pdf>
- SANS. (2006). *Audit 507 - Auditing Networks, Perimeters and Systems*
- Shimonski, Robert. (2004). *Risk Assessment and Threat Identification*. Retrieved from
http://www.windowsecurity.com/articles/Risk_Assessment_and_Threat_Identification.html

David Swift dgswift@verizon.net

Smith, George. (2002). *Newton's Philosophiae Naturalis Principia Mathematica*. Retrieved from <http://plato.stanford.edu/entries/newton-principia/>

Swift, David. (2009). *Successful SIEM and Log Management Strategies for Compliance*. Retrieved from http://www.sans.org/reading_room/whitepapers/auditing/successful-siem-log-management-strategies-audit-compliance_33528

Swift, David. (2009). *A Compliance Primer for IT Professionals*. Retrieved from http://www.sans.org/reading_room/whitepapers/compliance/compliance-primer-professionals_33538

Swift, David. (2006). *A Practical Application of SIM/SEM/SIEM, Automating Threat Identification*. Retrieved from http://www.sans.org/reading_room/whitepapers/logging/practical-application-sim-sem-siem-automating-threat-identification_1781

Timm, Kevin. (2002). *Justifying the Expense of IDS, Part Two: Calculating ROI for IDS*. Retrieved from <http://www.securityfocus.com/infocus/1621>

Zeltzer, Lenny. (2010), *Official Page for Critical Log Review Checklist* Retrieved from <http://www.securitywarriorconsulting.com/logchecklist/>.

David Swift dgswift@verizon.net

10. Acknowledgements

Thank you to [Accuvant](#) (my employer), for allowing me to share the included content, having developed and/or refined much of it “on the job” at numerous client installations.

I owe general credit to SANS courses for information and processes that have become part of my standard way of thinking, if not specifically quoted in the paper.

David Swift dgswift@verizon.net

Appendix A - White Listing

Make white list entries specific.

Include both the source address and port, or source/destination address and signature name.

Example: Windows Domain Controller

Source IP: 10.1.1.1

Source or Destination Ports: 389 (LDAP), 636 (LDAPS), 3268 (MSGCS)

Example: NFS Server

Source IP 10.1.1.2

Destination Ports: 111 (Portmapper), 2049 (NFS),

LOCKD, STATD (can be static or dynamically assigned).

The exception to the rule: common false positives.

Consider creating a “benign events” filter for common false positive signatures, or disabling logging of those events all together.

These may be application specific, and may be a join of two or more filters.

Example:

Mail Server Filter (all mail server ips, destination ports 110, 25)

AND IDS Mail Events (alert on SMTP TO, alert on SMTP FROM...) command detected.

Example:

FTP Servers (all ftp server source addresses with destination port 20, or source port 21)

IDS FTP Events (alert on FTP USER command, alert on FTP PUT...)

Note: In ArcSight these can be applied using filters. In Qradar, these can be applied through building blocks.

David Swift dgswift@verizon.net

Appendix B - Reports

Consider producing and reviewing the follow reports monthly, or on demand as needed.

User Activity Reports

All Active User Accounts (any valid/successful login by account name in the past 30 days)

Active User List by Authentication type

VPN Users

Active Directory Users

Infrastructure Device Access (Firewalls, Routers, Switches, IDS)

To be reviewed/certified as valid by the administrator/manager for each authentication source.

Unused accounts should be disabled.

Any unexpected account usage (default accounts, terminated employees...), should be investigated and explained.

User Creation, Deletion and Modification

A list of all user accounts created, deleted or modified by authentication type, to include the date, time, and User ID that made the change.

Active Directory

RADIUS/TACACS

Local (Unix, Windows Server...SSH, LSAS...)

Access by any Default Account

Guest, Root, Administrator, or other vendor default account usage

Access by any terminated employee, expired contractor, or other expired account

Access by Privileged Accounts (root, administrator...)

Noting time, date, source IP, and where possible source user name that became admin 'su' log

Correlation of "Privilege Escalation for User X" followed by Privileged Execution (Server Reboot, Log Clear, File Deletion) on windows

Service Account Usage

A list of all service accounts grouped by target address

This report should be reviewed by the service account user owner and validated monthly.

Any unexpected use of a service account, or use on an unexpected address should be investigated and explained.

David Swift dgswift@verizon.net

Configuration Change Reports

Configuration Change Report

Any configuration changes on monitored devices

Date, Time, and User ID that made the change

Access Reports

Access to any protected/monitored device by an untrusted network

VPN Access to Protected Network

Wireless Access to Protected Network

Access by a Foreign Network to a Protected Network

Foreign Country

Any Non-Internal/Company/Intranet Source Address

Access to a Higher Security Network by a Lower Security Network

Internet Usage by Protected Device

Any traffic from a protected device to a network other than the protected and trusted networks.

Incident Tracking

Current Open Ticket List

A list of all incidents not yet closed.

Closed Ticket Report

A list of all tickets closed in the past X days (from 1-90).

Details must include the total time the ticket was open (Time to Resolution), the root cause if found, and the person who opened and closed the ticket.

Time to Resolution by Ticket Type

For each ticket type (See Attachment A – Required Correlations), the minimum, maximum, and average time to resolution.

On Demand Operational Reports

User Login Tracking

All Logins for a User ID for the past 30 days - Group by User Name and Source Address

Use: Identify any Hosts a Terminated/Suspicious employee has logged on to for secondary investigation of those hosts.

Host Login Tracking

All logins on a given host for the past 30 days

Use: Identify who may have compromised a host that is misbehaving.

Malware Source Report

A list of host addresses for any identified malware or attack – group by malware name or attack name

Use: Identify the source IP addresses of any given malware or attack for targeted removal/remediation.

Malware Occurrence Report

A count of any given malware (group by IDS signature/Anti-Virus Signature/Attack Name), over the past 30 days.

Use: Defense Tuning – if <100 occurrences of a signature, block, If >100,000 blocking could disrupt the network. Log only mode on new signatures for 30 days, monitor, and block as appropriate.

Monthly Summary Reports

These reports are produced and reviewed monthly.

Top Sources & Destinations
Filtered to remove known servers.
Total Correlated Events/Events of Interest
Grouped and totaled by Event of Interest/Correlation name
Total Events / Day / Log Source
Top 10 Events Per Log Source (Anti-Virus, IDS...)
Top 10 Failed Logins
Grouped by Source IP (Top 10 sources of failed logins)
Grouped by Target User Name (Top 10 accounts with failed logins)
Web Content Filter Summary
Top 10 Destinations by Domain Name
Top 10 Blocked Sources by IP Address
Top 10 Blocked Sources grouped by Network (subnet)
Foreign Attacker Report
Top 10 Source Countries involved in Attacks (FW, IDS, AV, AUTH...)
Top 10 Sources IPs of Foreign Attacks
Top 10 Destinations of Foreign Attacks

David Swift dgswift@verizon.net