



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.

Correlating Event Data for Vulnerability Detection and Remediation

Copyright SANS Institute
Author Retains Full Rights



Sponsored by McAfee

Correlating Event Data for Vulnerability Detection and Remediation

October 2013

A SANS Analyst Whitepaper

Written by Jacob Williams

SIEM Integration with Endpoint Data for Accuracy and Speed *PAGE 2*

Asking Questions of the Data *PAGE 4*

Case Studies *PAGE 7*

Introduction

The 2012 Saudi Aramco “spearphishing” attacks are an excellent example of how to misuse a security information and event management (SIEM) platform for incident response. To recap, attackers installed malware on office computers at the company, which is responsible for a tenth of the world’s oil production, compromised several hosts and gained control of the network.¹

From there, malware spread to all machines in the domain, affecting business systems, but not affecting the vital oilfield, pipeline and storage control systems, which were presumably on a separate network, as is common with supervisory control and data acquisition (SCADA) systems.² The breach event cost Saudi Aramco an estimated \$15 million or more, simply for incident response.

The attack was so detrimental that the company took the drastic step of disconnecting its systems from the Internet for almost two weeks.³ Yet it could have been detected, contained and eradicated at a number of points if the data had been properly analyzed and the results acted upon. These missed opportunities include:

- The initial infection by the spearfishers
- Attackers’ scanning of additional systems for vulnerabilities
- Widespread infection propagation, involving some 30,000 hosts
- Communications among devices that should never be talking to one another (e.g., workstations communicating via admin shares)

The company’s SIEM system had the data about these activities, so why weren’t the issues detected sooner? In essence, the SIEM and detection systems weren’t connecting related events. Organizations need visibility into associated events for better identification, containment and remediation.

In this paper, we examine how this attack could have been thwarted with the help of a SIEM platform that combines the power of historical data with real-time data from network data sources and security policies. Such a system provides the context around application usage, user behaviors and other operations for better, more accurate reporting.

¹ “Saudi Aramco Says Cyberattack Was Aimed at Production,” www.nytimes.com/2012/12/10/business/global/saudi-aramco-says-hackers-took-aim-at-its-production.html?_r=0

² “Why wasn’t Saudi Aramco’s Oil Production Targeted?,” Digital Dao, <http://jeffreycarr.blogspot.com/2012/09/why-wasnt-saudi-aramcos-oil-production.html>

³ “Shamoon was an external attack on Saudi oil production,” www.infosecurity-magazine.com/view/29750/shamoon-was-an-external-attack-on-saudi-oil-production

SIEM Integration with Endpoint Data for Accuracy and Speed

Attackers running malicious code on your machines is bad enough, but real losses begin when they exfiltrate data from your network. The need to detect attacks before this stage is, therefore, paramount—and time is of the essence. The 2013 Verizon Data Breach Investigations Report (DBIR) concluded that the majority of attackers begin exfiltrating data within *hours* of the initial compromise.⁴ The Verizon report also reveals that although the majority of attacks take days to contain, almost a quarter of victims take months to identify breaches. That doesn't include the time it takes for them to determine the impact and remediate the vulnerabilities being exploited.

Endpoint Management and SIEM: A Happy Marriage

For visibility—the first step in analysis and detection—organizations need near-real-time alerts and searchable data. They need to identify all infected nodes, but that's just the start: They also need to be able to analyze seemingly normal data.

Crafty attackers disguise their tracks to make their traffic look like normal activity, avoid early detection and enable deeper entrenchment in the network. This slow and stealthy approach, mimicking normal user behavior, is where combining data from SIEM and endpoint management can speed up detection and assist with remediation.

Let's take an example where attackers are using SSH to perform their exfiltration under the guise of a routine administrative activity. The containment efforts go something like this:

- First, the incident responder obtains a list of all endpoint IP addresses that have used SSH within a certain window of time.
- Then he or she queries the endpoint management system to determine which endpoints and/or users were authorized to use SSH.
- The incident responder removes the corresponding entries from the next query, this one against SIEM data.
- The remaining addresses become the starting point for containment efforts.

Without integration between the SIEM system and endpoint management tools, duplicate query results must be removed manually (or with some script-driven assistance) to provide maximum value. The problem becomes even more difficult if endpoints are configured with dynamic IP addresses, as is often the case with desktops and wireless devices on IPv4 networks, making more work for analysts trying to track down the infected endpoints.

⁴ http://verizonenterprise.com/resources/reports/rp_data-breach-investigations-report-2013_en_xg.pdf

Make Reporting Easy

Responders need to focus on the incident at hand instead of fighting their data aggregation systems for answers. They need to be alerted to incidents they don't—or can't—catch manually.

Truly integrated solutions must allow for easily constructed queries, preferably using natural language. Instead of merely searching historical data in a SIEM system, responders need to go directly to the source and correlate on the fly, so a system should make it easy to find the infected endpoints in the first place. Finding compromised machines can be as easy as looking for machines that don't exhibit normal behavior. The SIEM data contains the history of what processes (or other artifacts) are normal for the organization's endpoints, thus forming the baseline. Then the analyst can quickly focus the investigation by examining machines that deviate from this baseline, using real-time queries to obtain the data quickly.

To correlate against data from a SIEM, endpoint data must be quickly queried and correlated against historic SIEM data and vice versa. Basic requirements should include the following:

- SIEM and endpoint management integration for easy correlation
- Network log data from firewalls, IPSs and “netflow” monitoring systems
- Real-time query of endpoints provided through the SIEM system or middleware that interoperates with the SIEM system
- Plain-language queries, preferably supported by standards so that other tools can interact without technicians having to learn new languages for each tool
- Preconfigured queries for common requirements
- Immediate access to current and historical security data

The Meaning of Integration

True integration is about much more than having a slick front end. Tools should be integrated to make the user's life easier, rather than fill a marketing-driven checklist; they can actually change the way an analyst thinks about a problem, making it more likely that they'll more quickly identify similar security events in the future.

Without the SIEM information, particularly the historical data the SIEM stores, the ability to query endpoints is only half as useful as it can be. Information stored in the SIEM can provide short-term and long-term snapshots of events related to endpoints. It can provide data that shows, over time, what is normal behavior; this can help detect abnormal behavior indicative of security-related events.

A holistic integration in which information on endpoints and from the SIEM are easily acquired will have these characteristics:

- Alerts are generated from all sources.
- Thresholds reflect asset value and risk.
- Initial analysis is automated.

This combination allows IT professionals to focus on critical operational and response tasks during and after an incident.

Asking Questions of the Data

Having information from every endpoint is one thing, but making sense of it is another. The driving requirement for such a chore is the ability to correlate real-time information from an endpoint manager with the historical data from the SIEM platform to understand the processes, files, countermeasures and other attributes that deviate from the baseline. Simply examining the real-time query responses and SIEM data is not enough; integration with the policy management system is necessary to provide the context critical to the meaning of the data.

Query for Compliance

One use of real-time data is verifying compliance with security policies. The SIEM repository may contain historical information, but systems fall out of compliance all the time. New vulnerabilities are discovered, new patches are released and unauthorized configuration changes are made with alarming regularity. Attackers rely on the vulnerabilities left open because of patching delays and failures. The mere existence of a patch can point the way to the automatic creation of an exploit, so the exploit clock starts ticking as soon the patch is released to the public.⁵ IT professionals know that just because an endpoint is configured for centralized patch management, it won't necessarily get all required patches. These failures leave the door open to attackers using "one-day" exploits that are reverse-engineered from patches, making independent verification of machine state critical to endpoint security.

Even verifying that security software is operating per policy can be a challenge, because attackers frequently disable a host's antivirus and firewall software after compromising it; the really smart ones also go after the security reporting tools.

If the SIEM platform supplies the only warning on these activities, the incident responder must wait for the next scheduled audit interval to discover disabled security software. If the attacker disables a real-time query agent on the machine, a null response from that machine (which should be reachable on the network), tells the analyst that trouble is afoot. This is why we also need correlation against endpoint data.

Query for Unknown Vulnerabilities

Another security threat is that of unapproved, user-installed applications with vulnerabilities that attackers can exploit. A good application whitelisting policy may prevent malware from executing, but crafty attackers have found ways around that as well, perhaps by using a tool that is already on the whitelist, such as Windows PowerShell. An approved application may also become dangerous if a new vulnerability is discovered. A method to mitigate risk based on new breach potential (e.g., blocking execution from a particular directory used by new malware) is a valuable tool, especially if it can achieve results in moments.

⁵ <http://bitblaze.cs.berkeley.edu/papers/apeg.pdf>

Visibility into new applications is as important as visibility into changes made to approved applications and must be part of the endpoint monitoring, alerting and querying strategy. The combined data can help IT and security professionals see into new and changed applications, such as those that evade traditional monitoring tools. For example, user-installed (and prohibited) applications such as file-sharing tools can get around protections in at least two ways:

- Installers place the binaries in the user's profile directory, bypassing Windows' application installation and execution controls. These unapproved applications offer a perfect vehicle for attackers to covertly exfiltrate data and establish command and control channels without detection.
- Alternatively, defiant users or attackers may utilize so-called "portable apps" on USB sticks that run on Windows without actually installing software.

Further, some malware executes only in memory, leaving no file-based artifacts on the disk.⁶ Relying on standard audit reporting intervals used by some SIEM solutions isn't sufficient when malicious processes leave few physical traces and are only present in memory during their execution.

Real-time queries address this problem by reporting the current machine state, detecting changes by comparing with historic data and searching for unapproved applications, misconfigurations and similar violations of the security policy. This step can filter systems that are not affected by attacks in progress, while showing what systems remain vulnerable to the attacks.

Query Your Wi-Fi

It's hard to purchase a laptop that doesn't have Wi-Fi as well as wired Ethernet. In most corporate environments, the laptop normally uses the wired connection for greater throughput, but Wi-Fi access is often available for the convenience of guests, or in flexible "hoteling" environments. (Many desktop computers also have Wi-Fi built in, just to make things interesting.)

Security experts have warned for years of attacks where wireless adapters on compromised machines are reconfigured, bypassing corporate network protections by connecting victims to a—sometimes phony—"guest" network or rogue access point.⁷ Such connections are often short-lived and can be missed by intermittent SIEM reports. Figure 1 shows an example of this sort of attack.

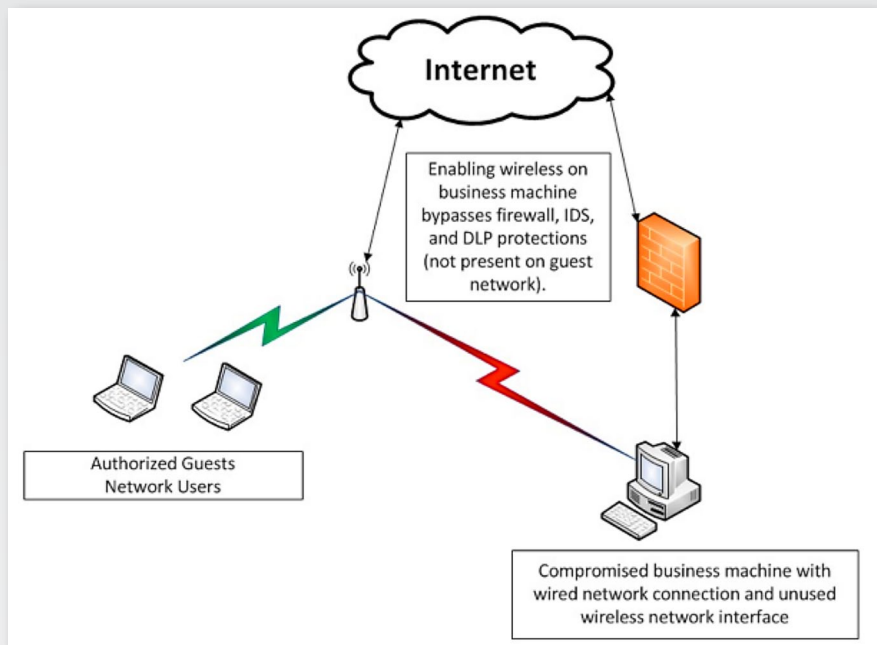


Figure 1. Typical Wi-Fi Hijacking

⁶ www.infoworld.com/d/security/java-based-web-attack-installs-hard-detect-malware-in-ram-188960

⁷ [https://files.sans.org/summit/pentest09/PDFs/Josh Wright - Wireless Weaponry - SANS PenTest Summit09.pdf](https://files.sans.org/summit/pentest09/PDFs/Josh%20Wright%20-%20Wireless%20Weaponry%20-%20SANS%20PenTest%20Summit09.pdf)

Asking Questions of the Data (CONTINUED)

Other examples of wireless chicanery include power users configuring rogue access points or connecting unauthorized wireless devices. All of these make the ability to query nodes for statuses in real time vital to a security response.

As with queries of historical SIEM data, an endpoint manager's real-time queries must be intuitive, fast and accurate. Although the ability to construct ad hoc queries using plain language is essential, an extensive array of preconfigured queries should be available. Ideally, the data returned from the real-time query can be automatically correlated against data stored in the SIEM, highlighting changes in the environment.

Case Studies

To illustrate how integrating SIEM and endpoint management can simplify incident response, let's examine two case studies, both based on actual events.

Case Study 1: Endpoint Remediation in a Regulated Environment

This case study takes place in a major health care provider's network that supports both patient care and research. Attackers targeting the health care industry are typically looking for research data (intellectual property theft) or personal information (identity theft). In the United States, personal health care data is specifically regulated by HIPAA. A HIPAA-defined data breach can lead to regulatory actions and costly fines.

Although the details in this case are hypothetical, such attacks and the sanctions resulting from them are all too real. Between October 2009 and March 2010, WellPoint, the largest managed health care company in the Blue Cross and Blue Shield association, fell victim to such an attack, resulting in a \$1.7 million fine for the company.⁸ If incident responders can stop attackers from obtaining access to regulated data, they justify their work and protect their organizations from unfavorable compliance rulings.

Traditional SIEM-Driven Investigation

The health care provider's SIEM system alerts on inbound HTTP requests on all machines that are not providing web services. Although HTTP traffic in an enterprise network is not itself suspicious, legitimate users' startup connections normally travel *outbound* to Internet web servers, making this a good test of suspicious activity.

In our example, a security analyst receives an alert from the SIEM system indicating HTTP traffic to and from a machine listed in configuration management as a database server. It could be that a developer set up a web server without adding it to the known-good list, but there's no way to know this without correlating an IP address to a device's function. The organization would greatly benefit from additional queries to the endpoint to find out more. It must be able to make such queries without the manual correlation, scripts, foreknowledge and tools upon which most such investigations rely.

Then, to derive additional information after an alert is received from the SIEM platform, administrators use the SIEM data to search another system—typically, an endpoint management system of some kind—to find misconfigurations related to the alert the SIEM system sent. Although determining whether the alerting machine is in a group authorized to listen for network services in general (and HTTP in particular) sounds easy enough, it won't work if there's a transcription error, outdated configuration information or a lack of communication among the sysadmins. This lack of information fusion can create a nightmare for incident responders.

⁸ "HHS Fines Wellpoint \$1.7 Million for a Major Breach of PHI," www.healthdatamanagement.com/news/breach-notification-hipaa-privacy-security-wellpoint-ocr-46377-1.html

SIEM and Endpoint Data Working Together

However, with integration of SIEM and endpoint management tools, the incident responder can determine whether the machine in question is authorized to have an HTTP server running in the first place. If not, our responder can look into other relevant aspects of the machine. The questions the responder should be asking and the information available from SIEM and endpoint management systems are summarized in Table 1.

Table 1. Comparison of Data from SIEM and Endpoint Management Systems

Question	Data Available from SIEM System Query	Data Available from Endpoint Management System Query
Is the security software still running on the machine?	Historical data from security tools includes versions and updates, plus the latest check-in times for agent-based security software.	Immediate, real-time data on the state of security tools and systems being monitored (e.g., host-based intrusion prevention system [HIPS] logs) is available.
Is the machine listening on TCP port 80?	Historical data of TCP port 80 use and policy for that usage is available.	If sudden spikes in TCP port 80 traffic are detected on the endpoint, or if new forms of encrypted traffic at unusual times of day are detected, this would be an immediate alert.
What process is associated with the listening port?	Historical data only lists what was running at a point in time and may not catch the malicious process. It is, however, good for baseline information.	Real-time data is queried as needed, enabling identification of active processes before the attacker can terminate them.
Are there any additional active processes on the machine that may be suspicious?	Data provided is good for baseline historical data and includes a list of processes previously running on the machine.	Immediate real-time data is essential: What is running now? The attacker may terminate these processes after use, making on-demand, real-time data critical. Attackers often use the same malware processes on multiple hosts during a compromise. By building an attack fingerprint, multiple machines involved in a compromise can be quickly identified with suitable tools.
Does the machine have up-to-date antivirus software, host-based IPS, etc., including definitions?	Data includes update logs for definitions and when the agent-based software last checked in.	Endpoint management systems query the antivirus/host-based IPS software for current operational status and last definition update.
Has the machine received scheduled updates to its OS and third-party software?	SIEM data is historical and is based on update logs.	Systems query for patch levels in real time and verify that patches reported installed by third-party patch management solutions actually completed installation correctly.

Case Studies (CONTINUED)

Although some of this information may be in the SIEM platform, the information is historical and, at this phase, incident responders need real-time information. With the combination of data from their endpoints collected in real time, and historic data stored in the SIEM system, the responding analyst almost instantly discovers that:

- The security software is still running.
- The machine is still listening on TCP port 80.
- The process associated with the listening port is **winservice.exe** (a portmanteau of two legitimate system processes, **winlogon.exe** and **services.exe**; such disguises are meant to fool casual or inexperienced observers).
- There are no other immediately obvious malicious processes on the machine.
- The machine has current antivirus and HIPS definitions.
- The machine has up-to-date patches for its OS and most third-party software, but it is missing a recent critical patch for Adobe Flash.

At this point it appears that the machine has been compromised, evidenced by the presence of the suspicious **winservice.exe** process. The security software, OS and application software are all up to date, except for Adobe Flash, so a publicly available Flash exploit may be to blame.

Our responder can use the historical SIEM data to view the source IP address of all inbound connections to TCP port 80 on the exploited machine, then query the SIEM for machines using those IP addresses to identify other machines to investigate. Then, our analyst can scan devices on the network in real time to find machines running **winservice.exe**.

Containing the Damage

Initially, containment for this incident could start with configuring a “black hole” route for traffic to known malicious IP addresses. During incident response, such routings can easily log data that would otherwise be sent to an attacker, where dropping an IP address might signal the malware that the connections are failing and lead the attacker to change tactics. For example, an attacker and his malware would simply contact the other compromised hosts from another IP address or act under a different process name. Adaptive tactics such as this are a core component of the “persistence” in advanced persistent threats (APTs).

Sending traffic to a black hole gives the administrator time to run real-time queries on every managed machine, identify other anomalous processes and artifacts useful in developing indicators of compromise and follow leads all the way back to the initial compromise.

The investigation doesn't stop there. With this query for connections to the suspect IP address, the analyst finds another machine talking to that address, only this time using the process **lsasss.exe** (note the extra “s”; this purposeful misspelling is another example of a misleading reference to a common Windows process). This process-level granularity is only possible (and feasible) with a real-time endpoint query.

Merely watching the attacker exfiltrate HIPAA-protected data would be reckless. Armed with tools that allow quick correlation of data, the incident responder is assured that all primary attack components have been identified and can begin remediation with a minimal possibility of additional, unknown malware.

Neutralizing the Threat

With the appropriate tools at their disposal, analysts can also neutralize the threat from a central console. Our analyst can send specific real-time commands to force-kill the known malicious processes on the infected machines, and then either create an action to delete the malware from the machine or move it to a quarantine folder for later analysis.

The job still isn't done; some malware may launch "monitor" routines, either running another process or injecting code into a neighboring process. These routines can restart the malware if it is killed, downloading a new copy from a malicious server if needed. This is why routing the attacker to the black hole is important: Analysts need time to find the infected machines, repair them and, thus, prevent malware from "rising from the dead."

In addition to executing the force-kill, our analyst would be wise to refresh process lists on affected machines afterward to see if malicious processes have re-spawned; this is a trivial task with real-time queries. Additionally, the analyst might schedule periodic process lists across endpoints to see if any new known malicious processes have started up since the initial remediation and configure the endpoint manager to auto-kill any reincarnation of the malicious process.

Fixing the Vulnerability

The vulnerabilities that allowed the malware to propagate must also be repaired. This is not only critical in keeping the same malware family from propagating again, it is also a key tenet of the Critical Security Controls—vulnerabilities must be regularly monitored and remediated.⁹

Remember that centralized visibility allowed the analyst to quickly establish that "Patient Zero"—the initially infected system—was missing a critical third-party software patch.

Without changing consoles, our analyst can see if other infected machines were also missing this patch, through historical SIEM data or real-time queries. Because the analyst has already established the correlation between the presence of the malware and the absence of the patch, he or she searches enterprisewide to discover other machines requiring the patch and then deploys it through the patch management system.

⁹ www.sans.org/critical-security-controls

Case Study 2: Analyzing a Financial Services Breach

The second case study involves a capital escrow firm that attackers believe is vulnerable. Once they locate a machine with the necessary permissions (by examining browser cookies and histories) and find two-factor authentication software, they will move swiftly to move funds offshore using automated clearinghouse (ACH) transfers.

In January 2013, such an attack closed down an escrow firm in California when attackers stole \$1.5 million via ACH transfers.¹⁰ A few months later, it took just three ACH transfers to steal more than \$1 million from a hospital in Washington State. The transfers were executed just before the bank closed for the weekend, suggesting that the thieves laid in wait before transferring funds.¹¹

Detection *before* the money was stolen was a critical missed opportunity, but it represents only half of the equation.

Setting the Stage

Although specific details surrounding such losses are rarely reported, let's examine how the attack may have unfolded:

1. An endpoint antivirus package using heuristic detection fires off an alert, which the analyst receives in real time via the endpoint management system.
2. The analyst checks to see that the machine has the latest signature definition files, which it does.
3. He or she notes that the detection was based on heuristics rather than a signature match.
4. The heuristic monitor alerted on a process called **svchost.exe**, which could be legitimate.
5. Although Windows machines normally have many **svchost.exe** processes running, they should all be running from `C:\windows\system32`.
6. This process, however, is using a nonstandard path, triggering the alert.

Figure 2 shows the process our analyst could take upon receiving the alert.




Figure 2. Analyst Discovery Process

Without the endpoint information, investigators would likely have to physically locate the suspect machine to continue the investigation, or alternatively use another product to query the endpoint. Neither option is ideal, and both provide a window of opportunity for attackers to complete the theft.

¹⁰ Although specific attack results have not been disclosed, the scenario is based on the end result described at: “\$1.5 Million Cyberheist Ruins Escrow Firm,” <http://krebsonsecurity.com/2013/08/1-5-million-cyberheist-ruins-escrow-firm>

¹¹ “Cybertheft heists \$1 million from Leavenworth hospital,” www.wenatcheeworld.com/news/2013/apr/26/cybertheft-heists-1-million-from-leavenworth

SIEM and Endpoint Data

Using a combined SIEM/endpoint management dashboard, our analyst immediately discovers the following:

- The security software is running normally.
- The destination IP address is 10.10.0.4 with a destination TCP port of 443; attackers often use common ports to blend with legitimate traffic. In this case, TCP 443 is associated with HTTPS.¹²
- No other processes are associated with the destination address or port, allowing the analyst to conclude that the anomalous `svchost.exe` is the source of the network traffic.
- The malware's parent process has exited, but the malware is running as the local user, with no child processes; because the parent process has exited, the malware can't be tied to a particular application.
- The machine has current antivirus definitions.
- The machine has up-to-date OS patches.

Determining Scope

Although it isn't yet obvious how the machine was compromised, our analyst searches the endpoint manager for other machines with connections to 10.10.0.4, and examines them. Armed with heuristic detection information about the anomalous `svchost.exe` process, the analyst creates a new detection policy to alert on any processes by that name running from a location other than the `C:\windows\system32` directory.

Additionally, the new policy will alert on machines with network communications to 10.10.0.4. Although this could be accomplished at the perimeter firewall, it often requires the help of another team; networking specialists generally manage firewalls, whereas host security specialists usually manage antivirus software. Deriving the answers from real-time queries facilitates efficiencies not possible when the actions of multiple independent teams must be coordinated.

Two additional machines are found communicating with 10.10.0.4. Once they receive the new heuristic detection policy, they also alert on communications from `svchost.exe` processes in nonstandard locations. Our analyst now has three machines to examine in the hope of finding a common infection vector.

More importantly, the company's endpoint security policy has been updated to alert on the threat heuristic. The analyst will be immediately notified when additional machines behave in the same manner as the infected systems.

Our analyst then uses the endpoint manager's real-time query to determine the timestamps of the `svchost.exe` files. They could indicate the probable infection time for the machines that did not initially alert on the activity. Understanding when the infection happened is critical when creating search parameters for the SIEM query. Investigators might then use the query results to reconstruct the network communications immediately before the alerts.

In this case, the analyst discovers that all three machines communicated with the IP address 172.16.0.8 using TCP port 443 within minutes of their infections; this may be the infection point.

¹² In a real attack, this will be a public IP address; our example uses private, RFC 1918-defined addresses to avoid inadvertently defaming any legitimate user of a public address.

Taking Action

The analyst might now block the suspicious IP address, but he or she should first determine if other machines have communicated with 172.16.0.8, in order to protect and remediate them. The analyst checks the SIEM data and determines that there are *thousands* of connections to that address over the last month, from more than a hundred endpoints.

The analyst can then use data gathered from the SIEM with the endpoint manager's reports to immediately query those endpoints that had been communicating with 172.16.0.8 for suspicious activity (i.e., `svchost.exe` processes in nonstandard paths).

Once the analyst determines the scope of affected systems, he or she blocks communications to the suspicious address, and within minutes the helpdesk is barraged by users complaining that they can't reach a popular website for marketing leads; the suspicious address corresponds to this site. The analyst now recognizes this as a potential "watering hole" attack, where victims are infected by visiting legitimate sites that have been compromised for the purposes of distributing malware to a specific group (those who use the compromised site).

Using real-time queries from the endpoint manager, our incident responder obtains the contents of the prefetch directory on infected machines. Using this information, the analyst is able to determine that `java.exe` was executed at the time of the infection and concludes the watering hole infected computers by way of a malicious Java applet targeting the marketing site's servers.

Our incident manager again queries the endpoint manager, finding that the machines are current on Java patches, indicating the possible use of a zero-day vulnerability, such as the one used to compromise UK kitchenware retailer Lakeland in July 2013.¹³ Legal counsel and the public relations team are tasked with contacting the compromised marketing site, leaving the analyst to focus on securing the network. (Removing Java from endpoint machines until a patch becomes available would be the safest approach in this example, but may not be possible in all cases.)

Zero-day vulnerabilities can be extremely challenging to remediate; by their very nature, no patch exists, but blocking access to the delivery website may fix the problem at hand. However, this doesn't help if the exploit is installed on business-related sites frequented by users for legitimate purposes.

Making Improvements

The integration of real-time query data, endpoint security configurations and information in the SIEM repository makes answering the questions posed by auditors and executives easier, and it can change the way security teams interact with business units.

Armed with this information, our analyst can present a case to business leaders to remove a risky application from managed endpoints until a patch is available. If the application were not a requirement for work, the easiest approach would be to blacklist the application on all machines that do not require it.

¹³ "Lakeland Kitchenware Hacked with Java 0-Day," www.infosecurity-us.com/view/33615/lakeland-kitchenware-hacked-with-java-0day

Conclusion

Monolithic security tools fail at correlating data from disparate sources. Even SIEM products that specialize in correlation are failing at catching the attacks occurring in our networks. Detecting a compromise is only half the problem. Security analysts and incident responders must have tools with the capabilities to respond to threats quickly—before data can be exfiltrated.

The integration of SIEM-based data correlation with real-time query and response data from endpoints is critical to tackling today's security challenges.

With the breadth and depth of integrated endpoint and SIEM orchestration, organizations have the security intelligence needed to react to events in progress, run real-time queries against the endpoints and check histories in the SIEM repository to determine scope and impact. This level of visibility is critical in today's threat landscape, where advanced attacks rapidly reach the stage of exfiltration and regulators are demanding increasingly detailed audit and incident reports.

About the Author

Jacob Williams is the chief scientist at CSRgroup computer security consultants and has more than a decade of experience in secure network design, penetration testing, incident response, forensics and malware reverse engineering. Before joining CSRgroup, he worked with various government agencies in information security roles. Jake is a two-time victor at the annual DC3 Digital Forensics Challenge.

SANS would like to thank its sponsor:





Upcoming SANS Training

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

SANS Tampa - Clearwater 2017	Clearwater, FLUS	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS Network Security 2017	Las Vegas, NVUS	Sep 10, 2017 - Sep 17, 2017	Live Event
SANS Dublin 2017	Dublin, IE	Sep 11, 2017 - Sep 16, 2017	Live Event
SANS Baltimore Fall 2017	Baltimore, MDUS	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Copenhagen 2017	Copenhagen, DK	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS London September 2017	London, GB	Sep 25, 2017 - Sep 30, 2017	Live Event
Data Breach Summit & Training	Chicago, ILUS	Sep 25, 2017 - Oct 02, 2017	Live Event
Rocky Mountain Fall 2017	Denver, COUS	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS SEC504 at Cyber Security Week 2017	The Hague, NL	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Oslo Autumn 2017	Oslo, NO	Oct 02, 2017 - Oct 07, 2017	Live Event
SANS DFIR Prague 2017	Prague, CZ	Oct 02, 2017 - Oct 08, 2017	Live Event
SANS Phoenix-Mesa 2017	Mesa, AZUS	Oct 09, 2017 - Oct 14, 2017	Live Event
SANS October Singapore 2017	Singapore, SG	Oct 09, 2017 - Oct 28, 2017	Live Event
SANS AUD507 (GSNA) @ Canberra 2017	Canberra, AU	Oct 09, 2017 - Oct 14, 2017	Live Event
Secure DevOps Summit & Training	Denver, COUS	Oct 10, 2017 - Oct 17, 2017	Live Event
SANS Tysons Corner Fall 2017	McLean, VAUS	Oct 14, 2017 - Oct 21, 2017	Live Event
SANS Tokyo Autumn 2017	Tokyo, JP	Oct 16, 2017 - Oct 28, 2017	Live Event
SANS Brussels Autumn 2017	Brussels, BE	Oct 16, 2017 - Oct 21, 2017	Live Event
SANS Berlin 2017	Berlin, DE	Oct 23, 2017 - Oct 28, 2017	Live Event
SANS San Diego 2017	San Diego, CAUS	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS Seattle 2017	Seattle, WAUS	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS Gulf Region 2017	Dubai, AE	Nov 04, 2017 - Nov 16, 2017	Live Event
SANS Miami 2017	Miami, FLUS	Nov 06, 2017 - Nov 11, 2017	Live Event
SANS Amsterdam 2017	Amsterdam, NL	Nov 06, 2017 - Nov 11, 2017	Live Event
SANS Milan November 2017	Milan, IT	Nov 06, 2017 - Nov 11, 2017	Live Event
Pen Test Hackfest Summit & Training 2017	Bethesda, MDUS	Nov 13, 2017 - Nov 20, 2017	Live Event
SANS Paris November 2017	Paris, FR	Nov 13, 2017 - Nov 18, 2017	Live Event
SANS Sydney 2017	Sydney, AU	Nov 13, 2017 - Nov 25, 2017	Live Event
SANS San Francisco Fall 2017	OnlineCAUS	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced