



Sponsored by Bit9

**Network and Endpoint Security
“Get Hitched”
for Better Visibility and Response**

July 2013

A SANS Whitepaper

Written by Jerry Shenk

On the Network PAGE 2

On the Endpoint PAGE 4

Network and Endpoint Security: Working Together PAGE 6

Executive Summary

It seems as if the bad guys are winning. If you follow the technology news, and even the mainstream news, stories about security breaches in large companies appear to be on the rise. Those are just the reported breaches. According to the latest Verizon Data Breach Investigations Report, the ratio among organization sizes that have been breached “is fairly even.”¹ In other words, organizations of all types are being broken into—quite often by attacks of low to medium difficulty, according to the Verizon report and another report, the Microsoft Security Intelligence Report for the last half of 2012.²

Why? It’s a matter of visibility. To begin with, IT security and operations teams need to pay more attention to the information available to them on their networks and endpoints, including laptops, workstations, kiosks and servers. Unfortunately, the ability to make sense of their security information is also something organizations are having trouble with. According to the SANS Eighth Annual Log and Event Management Survey published in May 2012,³ respondents are collecting more event data from their endpoints and on the network than ever before.

What they’re missing is the context with which to analyze the information they are collecting. And, to get that context, they need more than just raw data from logs and security sensors scanning device traffic on the network. They need visibility into what is installed and running on the endpoints/servers, as well as a history of activity and changes to that endpoint in order to make decisions during perceived events and for use in forensics. In addition to endpoint scan information for device state, IT groups also need to know if untrusted applications have been installed, which means checking memory, files, registries and other computing processes for malware imprints, and comparing them to the policies and network monitoring data related to that device.

This paper will discuss how endpoint visibility, coordinated with network intelligence, can help identify threats that were not discovered by other means, determine the level of threat, recognize previously unknown threats and follow up with more accurate information for regulators and investigators.

1 www.verizonenterprise.com/DBIR/2013, p. 5

2 www.microsoft.com/security/sir/default.aspx, p. 20 and www.verizonenterprise.com/DBIR/2013, p. 6

3 www.sans.org/reading_room/analysts_program/SortingThruNoise.pdf

On the Network

There are many tools already in use on today's networks to catch intruders and assist with forensics. And yet, intruders continue to advance into systems due to lack of visibility, intelligence and coordination between endpoint and network activity. For example, organizations have deployed the following tools with limited success:

- Advanced firewalls to filter inbound traffic and even watch outbound traffic for signs of trouble
- Intrusion detection systems/intrusion prevention systems (IDS/IPS) to watch inbound and outbound traffic to compare against signatures of known activity that is deemed malicious, such as an email attachment that exploits an Adobe vulnerability or a website that is designed to hijack the browser
- Network monitoring tools that watch activity coming from applications, such as large file transfers that aren't expected or a high number of blocked ports
- Log managers and SIEM systems that can collect information from all the devices on the network and watch for signs that systems are under attack or have been compromised
- Next-generation network devices that can do dynamic execution and analysis of inbound files before they get to a workstation

The question is: How do organizations leverage that data to differentiate between real events and events that don't apply to them? False positives waste valuable time in efforts to track them down. And, often organizations expect to plug in a magic SIEM box and have it all work, but the reality is that these tools need tuning—and that means knowing the types of actions attacks and malware will take, such as:

- Large data transfers
- Long connection times
- Connections to known hostile IP addresses
- Connections to countries that an organization does not deal with
- A dramatic increase in inbound or outbound traffic being blocked by a firewall
- A spike in failed DNS requests

All of these are indicative of a real threat in the environment. Unfortunately, this isn't the type of data organizations are typically collecting off their firewalls, IDS and other network devices. Without this level of context, there are many ways attackers get around current security controls (while leaving patterns that the tools themselves cannot detect):

- **Firewalls.** Although they are good at filtering, firewalls are still just filtering devices for inbound and outbound packets. They must also be properly programmed to catch outbound exfiltration and commands, which get around the firewall's filtering processes by hiding in accepted ports and services.
- **IDS and IPS.** IDS/IPS has become very good at identifying hostile IP addresses, indicators of compromise (IOC) and attacks. However, if an attacker wants to hide traffic from an IDS, it may be possible to encrypt the data so the IDS doesn't see it in the packets. Some encrypted traffic has significant patterns that can allow it to be identified as belonging to specific malware families, but in general, encrypted data is a problem for IDS/IPS technologies. For example, if attackers use SSL on port 443, the traffic appears to be normal encrypted traffic and may not be detected.
- **Network inspection.** One definition of network inspection is a network security appliance that takes IDS/IPS one step further by conducting deep inspection on web, email and FTP traffic looking for executable files, comparing their file signatures (hashes) with known-good and known-bad files, as well as running the unknown program files in a sandbox or virtualized execution engine to see what it does. Like IDS/IPS, network inspection appliances can also miss encrypted attacks or attacks that don't come directly across the network. For example, if a web server is using SSL encryption to deliver malware, network monitoring usually doesn't see what is being transported.

Likewise, if an email message gets delivered with a password-protected zip file, chances are that network monitoring won't be able to detect a hostile file inside the zip file. It is also possible that the malware could have come from another computer that is sitting inside the network, in which case, network monitoring may miss the traffic. USB drives or CD-ROMs and DVDs can also be used to deliver malware to a victim computer, and the appliance won't catch it.

Gaps in visibility can be managed through correlation with network and endpoint intelligence. Now the question becomes: How do organizations better leverage their network monitoring systems to identify hostile traffic? Furthermore, what information do they need about their endpoints to correlate with network data in order to make actionable decisions and forensics analysis?

On the Endpoint

The most common protections running on the endpoint involve antimalware/antivirus, or signature-based technologies. These technologies, which rely on signatures and known behaviors (heuristics), do well at protecting against known threats. The key problem with any signature-based detection system is the zero-day (or 0-day attack). A zero-day is an attack that is not known by the antimalware system.

Zero-Day Attacks

Zero-day attacks are hard to defend against because they have never been seen before. For example, in the physical world, flying a huge plane loaded with fuel into a skyscraper was something that most people had not considered before the morning of Sept. 11, 2001. An hour later, though, it was a known attack, and the passengers in Flight 93 above southwestern Pennsylvania responded by crashing the plane to avoid further damage. If we know what an attack looks like, we can respond. But, what if we don't know what it looks like?

Another problem with signature-based antimalware is protecting against attacks that are specifically designed to evade detection. Sometimes, attackers modify malware specifically to avoid the known signatures being used by detection systems.

Once they infiltrate the endpoints, malicious files are propagated across the network, searching out and infecting any computer they can find that is vulnerable to attack.

Table 1 shows some common examples of how endpoints are compromised—each of which would leave a trail that usually goes undetected.

Endpoint	Threat
<i>Endpoints: mobile devices, desktop and notebook PCs, and so on</i>	Users click malicious links or attachments, or are infected by drive-by downloads.
<i>Endpoints and servers</i>	In a practice known as <i>privilege escalation</i> , attackers attempt to upgrade a low-access account to allow administrative access.
<i>Servers and security devices</i>	Attackers seek server space to store tools and data that can be picked up later from outside the organization. These could stay quiet for a long time.
<i>Network boundary</i>	Edge servers and routers are often the last point to defend against an advanced attack. Criminals know this and tune their outbound traffic to leave the network beneath the radar.

Table 1. How Endpoints Are Compromised

On the Endpoint (CONTINUED)

These activities leave evidence on the endpoint and network that could be leveraged to detect the event. For example, suppose changes in a registry occur on one machine being monitored for system state changes. Shortly later, that machine begins making calls to another machine it wasn't previously authorized to connect to. Coincidence?

Network and endpoint data can be leveraged to determine that this sequence of events is suspicious and should be examined further, which will be discussed in the next section.

Endpoint Monitoring

Endpoint security such as antimalware, and monitoring tools such as MDM (mobile device monitoring) and NAC (network access control), are also useful prevention and discovery tools, but they only know about the devices they are running on or scanning. The security process or monitoring device doesn't know where the malicious file came from or how widespread the attack is.

Therefore, the next step for endpoint security is to have the endpoints and network security solutions communicate with a controller to share information about the attack, including when it arrived and where it came from. If there is malware on one computer, you might respond in one way. But, if there is a virus on multiple computers, a different response is warranted.

Network and Endpoint Security: Working Together

With the right coordination between network and endpoint security, administrators can quickly determine whether malware ran, where it ran, how it got on the endpoint and whether it was previously known or is new.

By integrating network and endpoint security, it is possible to have better visibility into the operation of the entire network. We can also learn where files came from, where they attempted to run and where they are currently running. This information is important for responders, forensics practitioners and auditors who need answers to the following questions:

- Did it execute?
- What did it do/what happened?
- Did anything else happen in association to this?
- What machines are impacted and how?
- What further action is needed?

So, how do network and endpoint security work together for this level of visibility? It starts with analysis. And, once malware is known through endpoint analysis, the network security systems can start blocking new instances of that malware. See Figure 1 for an example.

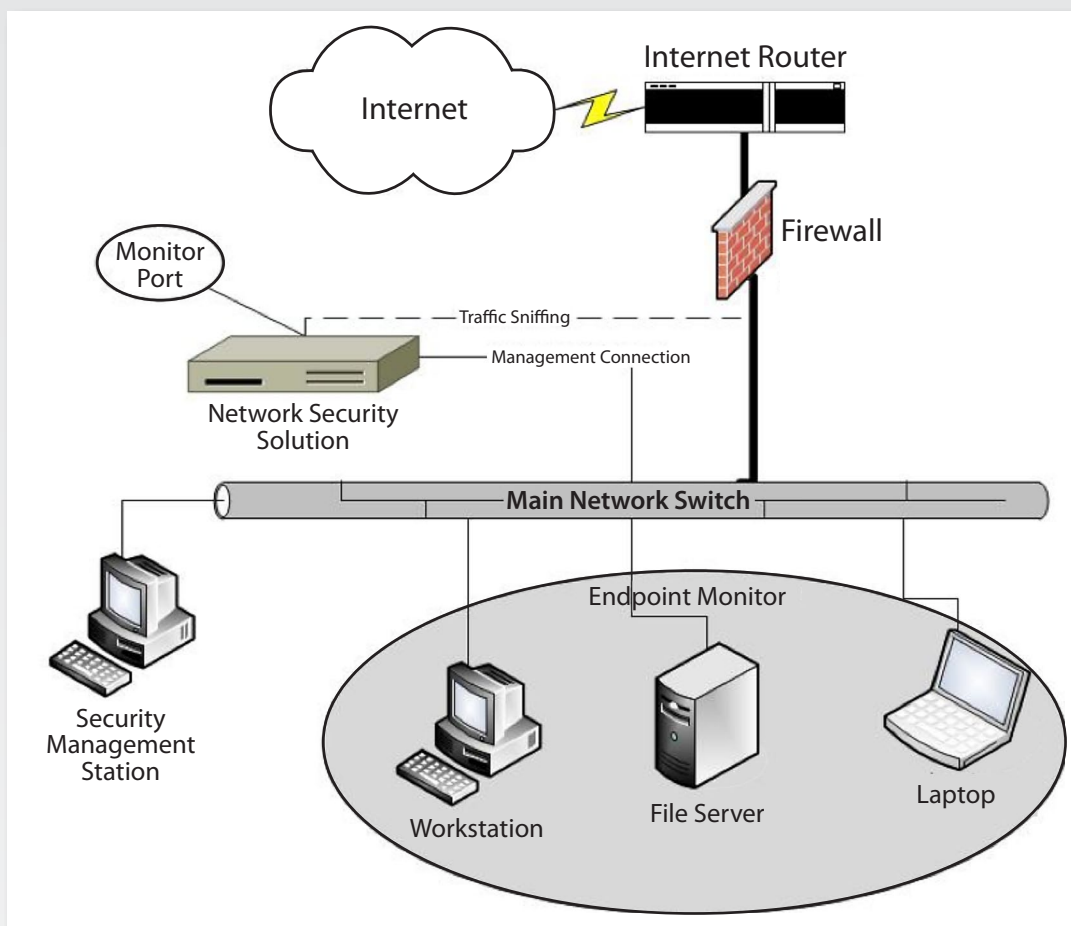


Figure 1. Network + Endpoint Monitoring Infrastructure

Network and Endpoint Security: Working Together (CONTINUED)

In this figure, endpoint monitoring software is running on all servers, workstations, kiosks, laptops and remote devices. These devices report activity including what processes are running, unknown executables and any detected malware to the security management station. If discovered, unknown executable files can also be sent to the management station by the endpoints for analysis.

The network security solution monitors the network for file transfers and executes any executable files in a protected sandbox. It can also execute any unknown files that it receives from the security management station for deeper analysis, identification and assignation of reputation.

Identify Known and Unknown Malware

Whether or not there is an antivirus signature for them, all attacks have a specific pattern of infection that can be analyzed. They may start with the insertion of a malware-laden USB stick on a back-end system or through an endpoint (a user clicking a malicious email, a drive-by download or other such event). Once into the target system, the attack loads an executable, which will attempt to hide in registries, memory and files. This executable may attempt to find passwords and data of interest to the attacker. Some malware will attempt to raise its privileges to get access to additional areas of the system. Once it has achieved its goal, it may pull additional software onto the system, or it may attempt to compromise other machines it finds in the network or on other networks accessing it.

Each of these activities leaves a trail on the endpoint and network that can be correlated. Table 2 provides an introduction to the indicators of an attack on both endpoints and networks.

Threat	Endpoint Indicators	Network Indicators
Users click links or attachments, or are infected by drive-by downloads.	Changes in registries, processes, memory executions, suspicious user/access activity	Installation traffic to the endpoint, confirmation traffic from the endpoint, unauthorized applications
Attackers seek out privileged accounts to gain a deeper foothold into the organization (privilege escalation).	Suspicious user activity, changes in user access credentials, misuse of access credentials/VPN	Access requests and traffic between systems that normally don't talk to each other, unusual VPN traffic
Attackers seek server space to store tools and park data bound for exfiltration.	Changes in registries, processes, memory executions, suspicious user/access activity	Traffic between devices that usually don't talk to one another, pings to the device from the Internet
Edge servers and routers are often the last point to defend against an advanced attack. Criminals know this and tune their outbound traffic to leave the network beneath the radar.	Activity from the egress point should be detected at this stage, but with encryption, and by hiding in authorized traffic like HTTP, this can be difficult to detect.	Packets sent low and slow or in unpredictable bursts to IP addresses outside the network. Heavy outbound traffic or a regular outbound traffic that does not match normal business processes can also be indicators.

Table 2. Endpoint and Network Threat Indicators

If some type of malware actually runs on one of your endpoints, you want to know about that immediately. It's possible that, at the end of the attack, egress filtering is still blocking the results of the attack getting back to the attacker, but this is the end game. This security breach needs to be detected and fixed at the endpoint, which means providing additional layers of security at the endpoint (such as whitelisting, reputational services, advanced intelligence and scans for changes in the registries and files that are not authorized in policy).

Catch Them at Their Own Game

Advanced attack intelligence can provide needed context to correlate between network security and endpoint events to determine, in near real time, if an alert needs to be issued for an attack in progress. If determined to be malicious, the program and its originator (if discovered) can be blacklisted, and the hash of the file can be stored in a database to be shared with all other endpoints and network analysis appliances for the future. After this analysis, if the program tries to transmit itself across the network or ends up on another workstation, it will be a known and detected. Ideally, this process would be automatic and supplementary to a larger database that combines third-party intelligence with intelligence gathered on events discovered internally.

Block and Tackle

Policies should be predetermined to correlate activities on endpoints, including servers, so that specified programs are allowed based on the actions of those files in addition to a list of "signatures" of bad and good files. Specified programs and actions taken by programs and files, then, come under scrutiny rather than just the applications themselves.

The strictest configuration can be most useful for installations in which there is no deviation from normal operation, such as a PC running a gas pump or a computer running a point-of-sale terminal (cash register). Server and kiosk endpoints often also run a rather strict policy because they are usually running well-defined services, such as DNS or database processes on the servers, and kiosks would have limited functionality.

Policy on what programs and files are allowed to execute should be flexible. Workstations often need to have additional leeway to run programs common to the user population. So, for example, when a new executable shows up on a workstation, it can be checked against known files (both known-good and known-bad files), the locally-stored database of approved files and/or the approved files listing from a cloud-based service. The network analysis appliance can then execute the program in a sandbox or virtualized execution engine.

If the program passes all these tests, then execution on the workstation can be allowed—but would still be monitored for malware-like activity. If the unknown file is determined to be malware, then all other detection points could be updated with the hash of this new file. All of these processes should be managed by policies and customized for different use cases and policies within the organization.

Network and Endpoint Security: Working Together (CONTINUED)

This system might operate as illustrated in Figure 2.

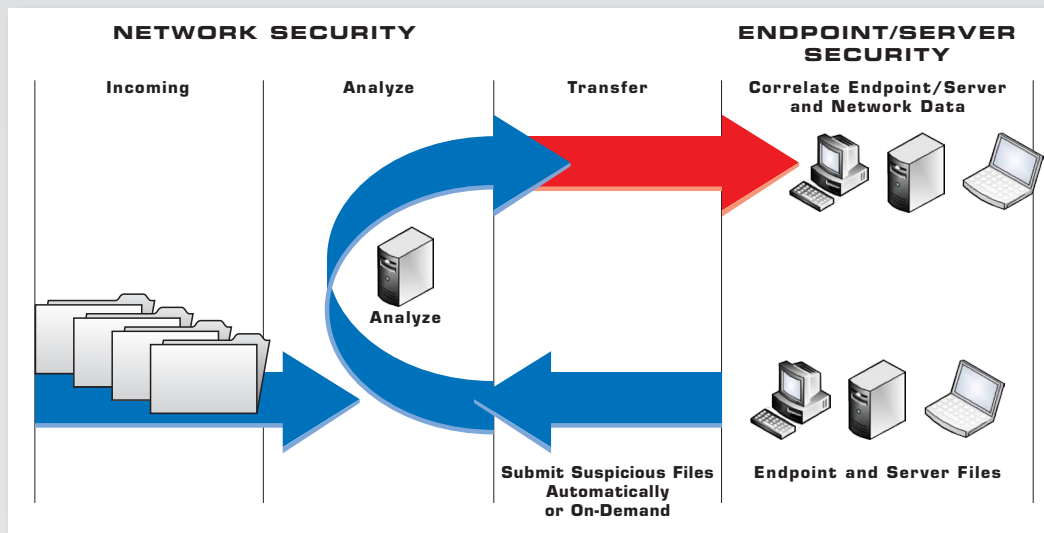


Figure 2. Example of Coordinated Network and Endpoint Security

Endpoint security can give us quick insight on new attacks if it is able to share that information with a central management console. Moreover, if that management console can be integrated with network security, additional attacks can be stopped before they ever get to the endpoint. This insight can also identify which legitimate process started the malicious process, and all other processes started by the legitimate process can be given a more thorough examination to verify that they were not part of the infection.

Satisfy the Critical Security Controls Guidelines

One of the key recommendations of the Critical Security Controls for Effective Cyber Defense⁴ is automation. Assess and inventory your devices and applications to know what applies to you; then automate defenses, alerts and responses.

Most organizations will want to focus initially on automating alerts. Once an organization has tuned the system properly and knows that it is reliable, it can begin to enable automated actions (e.g., auto bans).

The following individual controls listed in the Critical Security Controls (CSCs) apply:

- Inventory of Authorized and Unauthorized Software (Control 2)
- Malware Defenses (Control 5)
- Controlled Use of Administrative Privileges (Control 12)
- Boundary Defense (Control 13)
- Account Monitoring and Control (Control 16)
- Incident Response and Management (Control 18)

Efforts to secure endpoints may touch on other controls in specific instances. For example, Control 11, "Limitation and Control of Network Ports, Protocols, and Services," is related if a DNS server were to be attacked and had launched a shell. With that said, however, the situation is more specifically addressed by Controls 5 and 18. Each organization attempting to implement the CSCs will need to determine how it is affected.

⁴ www.sans.org/critical-security-controls

Conclusion

Many organizations are struggling with security issues. Both the Verizon and Microsoft studies mentioned at the beginning of this report indicate that correlation and visibility into attacks is a real issue across all sizes and types of organizations. Typically, organizations do not find out about security problems for weeks, months and sometimes even years—and when they do, it is usually by third parties alerting them, according to many reports. During that time, attackers may have full access to steal trade secrets, destroy or corrupt research, and siphon protected customer information (as in the famous Heartland Payment Systems breach in 2009).⁵

Some organizations are coordinating their endpoint security with network security information to help detect security events before they become bigger problems. With that coordination, network administrators can quickly determine whether malware ran, where it ran, how it got on the endpoint, and whether it was previously known or if it is new.

By combining this information—through the use of self-developed or better-managed third-party tools—organizations can help satisfy the demands for better response and automation detailed in the CSCs. Integration of endpoint monitoring and network security enables organizations to shorten their response time, target their resources better, stop new threats that don't have a known signature, ban those new threats throughout the organization, and build better awareness of what is happening between devices and on the network that could and should be correlated for a real view of events in progress.

5 http://usatoday30.usatoday.com/money/perfi/credit/2009-01-20-heartland-credit-card-security-breach_N.htm

About the Author

Jerry Shenk currently serves as a senior analyst for the SANS Institute and is senior security analyst for Windstream Communications, working out of the company's Ephrata, Pa., location. Since 1984, he has consulted with companies and financial and educational institutions on issues of network design, security, forensic analysis and penetration testing. His experience spans networks of all sizes, from small home-office systems to global networks. Along with some vendor-specific certifications, Jerry holds six Global Information Assurance Certifications (GIACs), all completed with honors: GIAC-Certified Intrusion Analyst (GCIA), GIAC-Certified Incident Handler (GCIH), GIAC-Certified Firewall Analyst (GCFW), GIAC Systems and Network Auditor (GSNA), GIAC Penetration Tester (GPEN) and GIAC-Certified Forensic Analyst (GCFA). Five of his certifications are Gold certifications. He also holds the CISSP certification.

SANS would like to thank its sponsor:

