IT@Intel White Paper

Intel Information Technology
Security
December 2009

Getting Ahead of Malware

Since implementing our security event monitor and detection processes two years ago, we have seen a 40 percent decrease in the number of formal incident response events.

Greg L. Bassett
Senior Information Security Specialist, Intel IT

Jeff Boerio
Senior Information Security Specialist, Intel IT

Executive Overview

To minimize the threat posed by malicious software, or malware, making its way into the enterprise, Intel IT has established a process that actively seeks to identify and take action against the malware before it reaches Intel's user base. This process focuses on real-time monitoring and interpretation of security events on the network and taking immediate action against any identified threats.

Our security event monitor (SEM), a tool that provides a centralized framework for collecting these events, coordinates our proactive approach, which includes:

- **Early detection of malware.** Through the use of strategically placed network intrusion detection sensors, we continually look for and research new threats. We also deploy signatures or patterns based on malware samples and monitor for systems that trigger the signature.
- **Rapid analysis of threats.** If an intrusion occurs, we immediately take steps to investigate and identify it. Our malware analysis team determines the best course of action.
- **Updates to the detection system.** We apply what we learn from any malware intrusions by continually updating the SEM, which refines our process and better defends against future attacks.

Our proactive approach is helping to reduce the number of formal events and limit the damage caused in the environment. Since implementing our SEM and detection processes two years ago, we have seen a 40 percent decrease in the number of formal incident response events.

Minimizing the number of malware intrusions has reduced the amount of time we need to spend in a reactive mode, allowing us more opportunities to increase our effectiveness at detecting new malware.

Contents

Executive Overview 1
Background 2
SECURITY EVENT MONITOR 2
  Detecting Malware 3
  Analyzing the Threat 4
  Remediating Infected Systems 5
  Results 6
Next Steps 6
Conclusion 7
Acronyms 8

IT@INTEL

IT@Intel is a resource that enables IT professionals, managers, and executives to engage with peers in the Intel IT organization—and with thousands of other industry IT leaders—so you can gain insights into the tools, methods, strategies, and best practices that are proving most successful in addressing today's tough IT challenges. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BACKGROUND

The volume of new malicious software, or malware, released every day increases the likelihood that some of this code will make its way into the enterprise and disrupt normal operations. Often by the time a user calls to report that a computer has been infected with malware, the computer's performance has already degraded, or it is exhibiting unusual behavior.

Malware can make its way to a computer in various ways, including:

- Uninstalled security updates and patches
- A zero-day attack, in which malware exploits a previously undisclosed vulnerability in a computer application or the OS
- Visits to infected Web sites, either through direct browsing or clicking links from e-mail or instant messages
- Attacks to a universal serial bus device
- Non-secure file share access

In the past, our IT Emergency Response Process (ITERP) dealt with malware after it had already taken hold and was actively proliferating in the environment. These events required multiple teams and large amounts of time to identify, contain, mitigate, and remediate these threats.

To help us stay ahead of malware, our security operations team has a charter to actively seek to identify the malware and take action before it reaches Intel's user base. Our approach to proactive discovery of malware, which requires fewer resources and less time than before, includes:

- **Early detection.** We continually look for and research new threats, deploy signatures or patterns based on malware samples, and monitor for systems that trigger the signature.
- **Rapid analysis.** If a suspected intrusion occurs, we immediately take steps to investigate and identify it. Our malware analysis team, which is under the umbrella of ITERP, decides how to respond to possible intrusions.
- **Updates to the detection system.** If an intrusion occurs, we take a malware sample from the infected system to run on systems in our malware analysis lab. The goal is to obtain intelligence data that can help us refine the network signature to detect the presence of the malware and, through a deeper analysis, identify the threat's intent, whether to steal passwords, collect e-mail addresses, or export intellectual property. We apply what we learn from any malware intrusions to refine our detection process and better defend against future attacks.

Our security event monitor (SEM) coordinates our proactive methods. This tool provides a centralized framework to monitor and analyze security events on the network.

SECURITY EVENT MONITOR

The SEM, the key to our proactive approach to dealing with malware, acts as the first line of defense in threat protection.

As part of daily security operations, we use various data collection methods to stay up to date with the changing malware landscape. Collecting data not only helps us monitor for known malware, but it is also an essential part of researching new threats of which we're still unaware.

In addition to using resources that are specifically looking for external security news, we rely on intelligence sources for data to construct lists of potentially malicious sites. These sites and addresses are conveyed through malware reports, security mailing lists, security Web sites, and other sources.

Detecting Malware

The SEM's network intrusion detection sensors (NIDS) and other strategically placed network detection devices monitor for malware attacks. These sensors study packets on the network, looking for indications of the presence of malware or potential intrusion attempts. The process is shown in Figure 1.

- The sensors use available signatures or patterns to determine whether the contents of a network packet meet the criteria of an attack and whether the signature of the potential malware matches any of the given signatures derived from malware we already know about.
- If a match is determined, the SEM logs the event and notifies our malware analysis team; the team then analyzes the event and determines whether to take any action.
- The SEM may also send an automated notification to the user. This notification alerts the user to the malware and outlines the immediate steps the user must take to confirm the detection. These steps include updating the system's antivirus (AV) software and performing a system scan, as well as collecting data to assist in the malware investigation.

A critical component in identifying and containing potential malware is finding a trigger to detect newly infected systems. Triggers can come from known malware or a possible variant of it.

- **Known malware.** Once we identify new malware, it becomes "known" malware. We identify known malware by confirming an infection on a source system and removing the infection with known remediation steps, usually an AV scan. We extract the malware if we suspect it is a new malware variant (NMV).
- **New malware variant.** This is a type of malware that our AV system doesn't detect, because it is new and its signature is not yet known. An NMV may have the same characteristics as known malware, but has undergone minor changes to its code base enabling it to get around the defenses.

As part of the detection process, we proactively search for sites that may be hosting malware and then investigate those sites to get a sample of the malware for analysis. Once we have a sample, we can create a signature and monitor for a variant with a pattern that matches that signature. We continually update the SEM with the information we collect.

As part of this early detection process that involves monitoring and confirming suspicious sites, we can discover false positive sites that may not be malicious or are hosting legitimate content. We may need to fine-tune our detection rules to reduce the false positive triggers. The ongoing process of building more effective SEM detection rules reduces the potential for false positives. Any false positive detection is analyzed and fed back into the detection signatures to remove further detection.

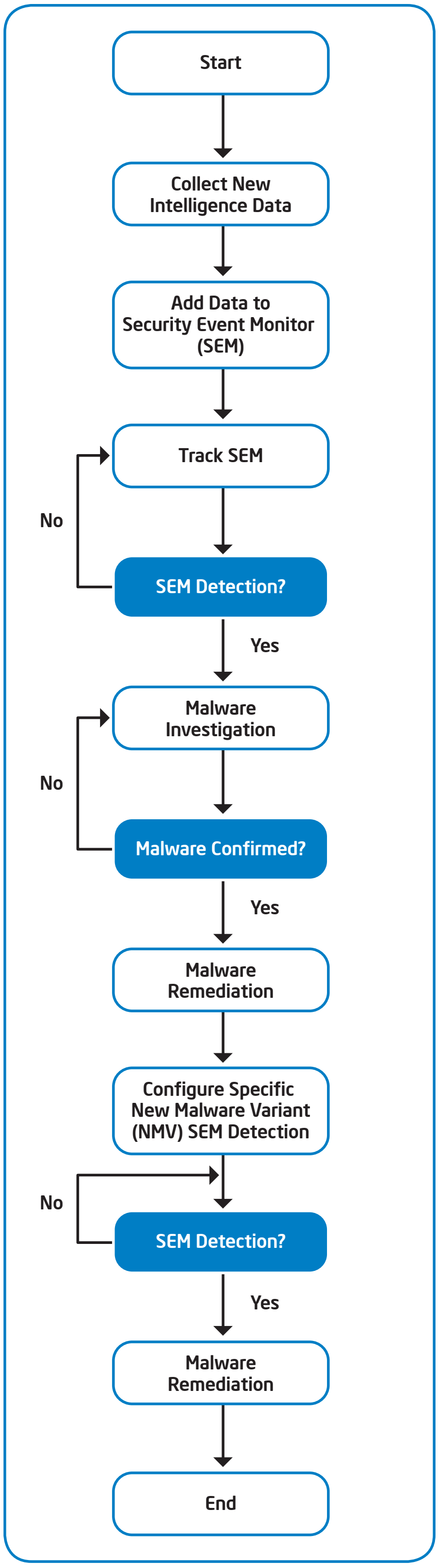

Figure 1. The security event monitor (SEM) coordinates our proactive process for dealing with malware.

Our detection process produces two types of events:

- **Systems that are attempting to contact addresses that are on our suspected malware hosting list.** This type of event indicates potential malware, and the systems need to be investigated to verify its presence. If the detected malware is from a known family for which we already have an existing signature, we update our rule set so that the SEM can report this detection as part of the same malware family.
- **Systems that trigger a known malware signature, generated through prior analysis.** Once we obtain and analyze a sample, we can develop secondary signatures to detect these known pieces of malware and then actively monitor for variants of this malware. We also make sure that the AV vendor knows about it, so their database gets updated and the information gets into the detection system.

If we identify that an intrusion has occurred, we initiate an investigation. We send users automated e-mail messages requesting they take immediate action and letting them know we need to investigate their systems for possible new malware. If numerous machines are involved, we may initiate our formal incident response team. We then formulate a plan to determine the severity of the intrusion and whether the affected machines need to be removed from the network.

Analyzing the Threat

If the SEM detects a possible NMV and isolates it to a particular computer, we analyze the system as quickly as possible to validate and contain the threat. Determining the function and purpose of the malware is also an essential part of this process; otherwise, subsequent remediation efforts may not meet our expectations.

Our analysis provides empirical data that includes the malware signature, behavioral characteristics, network traces, and other information that we use to properly identify and remediate the malware. As we acquire new data about a particular variant, particularly the network traces and the signature data, we add this information to the SEM to improve its ability to detect early infections.

IDENTIFYING THE NEW MALWARE VARIANT

Although we always hope to find evidence of the initial activity that caused the infection, this isn't always possible. Instead, we can check for suspicious file names, services, and other resources, and attempt to identify the process that generated the detected activity.

- In certain cases, we conduct a hands-on analysis of the infected system to identify the malware. However, this is a time-consuming process that requires access to the system, either physically or remotely.
- To automate the data collection process, we use Rapid Assessment and Potential Incident Examination Report (RAPIER), a security tool Intel developed and then released as open-source. RAPIER, which can be run on systems that are offline, gathers logs and registries and runs security programs to get a snapshot of the system's status (Table 1).

Once we collect the relevant data, we can review it to determine next steps. We also give this information to the AV software vendors so they can update the DAT files in their software to accurately detect new threats.

COLLECTING MALWARE SAMPLES

We also take a malware sample from the infected system to run on systems in our lab. A preliminary analysis of the sample can provide intelligence data that can help us refine the signature and improve the process for detecting the NMV.

Network traces of the packet communications can provide unique packet characteristics that we can use to create custom SEM signatures. The sample analysis may reveal target port information, additional target addresses, or regular expression strings of a specific packet that can ultimately be used to improve the initial detection process.

- **Target port information.** We can expand the detection rule set to look for a target IP address on a target port with which the malware is communicating.
- **Additional target addresses.** If the malware also communicates to another target address, we can add this target IP address to the detection rule set as well.
- **Regular expression strings of a specific packet.** If the malware has a specific packet, we can search for its regular expression strings instead of searching for generic target addresses.

A deeper analysis of the sample can help us characterize the malware and determine its intent. For example, its purpose might be to steal passwords, log keystrokes,

Table 1. Some Data Collected by the Rapid Assessment and Potential Incident Examination Report (RAPIER) Security Tool

TOPIC	DESCRIPTION
Alternate Data Streams (ADSs)	The file system's ADSs
AuditPolicy	Microsoft Windows* Audit Policy status
Checksums	Microsoft Windows system file checksums
Dump Users	Local system users
File Handles	Open file handles
File Time Stamps	Modified, Access, Created times of files on system drive
Hidden Files	Existence of hidden files on the system drives
Browser Activity	Temporary files and cookies
List Dynamic Link Libraries (DLLs)	Associated DLLs of running processes
Logs	System, application, and security logs
Network	Interface configuration, Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) activity, ports opened by processes
Registry Dumps	Local registry hive infection
Rootkit	Rootkit detection
Services	Services that are running on the system
System Information	Generate system information about the hardware, OS, and installed software

attempt secondary downloads, collect e-mail addresses, or export intellectual property. Acquiring this information usually involves the use of static, behavioral, or code analysis.

- Static analysis can provide basic file characteristics, such as checksums, ASCII strings embedded in the file, and file compression or packing techniques, which help to establish a profile of the attacker and the techniques used for the attack.
- Observation of the malicious code can provide a behavioral analysis of files, processes, and network traffic and registry entries, furthering helping to refine the malware detection signature. This process often requires repeated testing to reveal what the malware requires to run and what it does once it is running.
- Code analysis may involve reverse engineering or the debugging of suspicious samples to provide specific details on the actual attack. This effort can also reveal target addresses or URLs that we can use for monitoring purposes.

Remediating Infected Systems

Once we have a thorough understanding of the extent of the malware infection and have identified the components and signature of the malware, we can focus on removing it from the computer. The malware analysis team must decide whether remediating the computer can safely return it to service. The SEM triggers infected systems. The automated trigger provides the user with details on remediation steps and where to get help.

In some instances, identifying all of the malware and removing it may not be possible. In that case, the malware analysis team may decide that rebuilding the computer or restoring the OS image is necessary to remove the malicious files and the related registry keys.

For example, rootkits are a type of malware that can hide their existence from analysis tools, making identification and remediation more difficult. Kernel-level rootkits, in particular, are extremely difficult to remove and may require the computer to be rebuilt if we cannot reveal and disable their hidden components.

Results

In the two years that have passed since implementing our SEM and detection processes, we have been able to reduce the number of formal incident response events by 40 percent (Figure 2). This has reduced the amount of time we spend reacting to malware in the environment and allowed us to focus efforts on increasing our effectiveness at detecting and analyzing new malware.

NEXT STEPS

Our process is constantly evolving. We are always working on new ways to detect, contain, and remediate malware.

Some of the ways in which we are improving our detection capabilities include:

- **Expanding the capabilities of current tools.** The AV vendor may add new features and capabilities with upgrades and version changes. Taking advantage of these new tool features is one method we can use to enhance our detection and response capabilities.
- **Integrating new data sources into the SEM.** Log sources from across the enterprise provide a rich source of event traces that enhance the correlation of detection rules. Understanding the data source events to determine appropriate triggers is a key element in successfully detecting malware intrusions.
- **Enhancing detection of proxy-aware malware.** Malware authors are always improving their tactics, finding new ways to evade detection and attempting to secure a foothold in the environment. We continually seek out methods to improve the way we detect botnet command-and-control traffic, update traffic, and other malicious communications.

While not the focus of this paper, part of our process for continual improvement includes modifications to the containment and remediation aspects of our malware response.

Containment. The containment of infected systems within a global enterprise can be challenging. Going forward, we plan

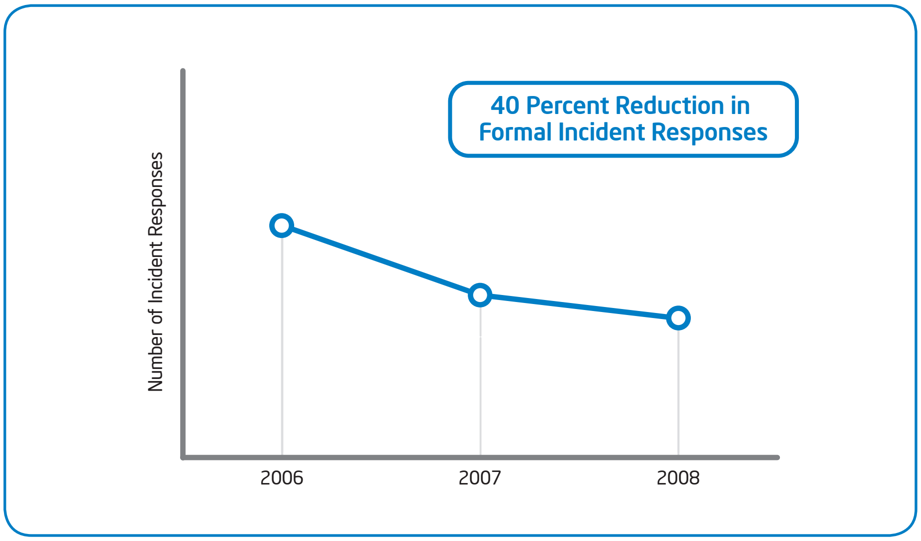

Figure 2. We have seen a 40 percent reduction in formal incident responses since implementing the security event monitor.

to investigate the use of Intel® vPro™ technology to provide faster containment capabilities, while also providing the framework for early malware detection and remote imaging capabilities.

We can also restrict network access of infected systems to contain the system, which allows further investigation and analysis while limiting the risk to enterprise operations. Only the services required to allow analysis and remediation can be provided to the infected system.

Remediation. Systems face a wide array of malware families and variants during daily attacks through common user activities. Some variants are more difficult to clean than others, often requiring the external development of specific cleaner tools to remove the infection.

To help reduce the impact and loss of productivity experienced by the end user, efforts are underway to streamline our systematic approach of understanding the variant, the remediation tasks, and when to rebuild or reimage the system. In particular, we are looking at ways to streamline the rebuilding and reimaging processes using Intel vPro technology.

CONCLUSION

Implementing the SEM has helped us centralize and streamline our security event collection and investigation process. It has also helped us shift to a proactive approach that enables us to stay ahead of malware.

Strategically placed NIDs that monitor for network anomalies that may indicate possible malware attacks are one component of our proactive approach. We also rely on our malware analysis team to immediately respond to and investigate any possible intrusions. Should a suspected intrusion occur, we collect data from the computer to identify any malicious files and also take a sample for lab analysis. We apply what we learn from any malware intrusions to refine our process, further improving our ability to detect new malware and better defend against future attacks.

We also continue to update the SEM with additional intelligence as it becomes available by monitoring known malware sites and implementing zero-day signature detection.

ACRONYMS

ADS	alternate data streams
AV	antivirus
DLL	dynamic link library
ITERP	IT Emergency Response Process
NID	network intrusion detection
NMV	new malware variant
RAPIER	Rapid Assessment and Potential Incident Examination Report
SEM	security event monitor
TCP	Transmission Control Protocol
UDP	User Datagram Protocol

For more straight talk on current topics from Intel's IT leaders, visit www.intel.com/it.

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel, the Intel logo, and Intel vPro are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2009 Intel Corporation. All rights reserved.

Printed in USA
1209/JLG/KC/PDF

♻ Please Recycle
322029-003US