InformationWeek :: reports

Next

Reports.InformationWeek.com May 2012

Rooting Out Sophisticated Malware

As malware gets increasingly sophisticated, so, too, must the technology and strategies we use to detect and eradicate it (or, better yet, stop it before it ever makes it onto network systems). There is no one product or product category that can do the job alone. Instead, security professionals must become familiar with—and adept at using—a combination of technologies. Security pros must also become skilled at connecting the dots among sometimes innocuous-seeming events to root out trouble. In this report, we examine the tools, technologies and strategies that can ease some of the burden.

By John H. Sawyer

Presented in conjunction with

SECURITY dark READING
Protect The Business Enable Access

Report ID: S4880512

TABLE OF CONTENTS

ABOUT US

InformationWeek Reports' analysts arm business technology decision-makers with real-world perspective based on qualitative and quantitative research, business and technology assessment and planning tools, and adoption best practices gleaned from experience. To contact us, write to managing director **Art Wittmann** at *awittmann@techweb.com,* content director **Lorna Garey** at *lgarey@techweb.com,* editor-at-large **Andrew Conry-Murray** at *acmurray@techweb.com,* and research managing editor **Heather Vallis** at *hvallis@techweb.com.* Find all of our reports at *reports.informationweek.com.*

John H. Sawyer
InformationWeek Reports

John H. Sawyer is a senior security analyst with InGuardians, specializing in Web, mobile and network penetration testing. His experience in IT enterprise security includes penetration testing, system and network hardening, intrusion analysis and digital forensics. He was formerly a senior security engineer with the University of Florida, Gainesville, and is a Dark Reading, Network Computing and *InformationWeek* contributor and blogger.

John was a member of team 1@stplace, a small group of righteous hackers that won the electronic Capture the Flag computer hacking competition at DEFCON in Las Vegas in 2006 and 2007. He has consulted with federal, state and local law enforcement agencies on malware analysis, hacker attacks and digital forensics. His certifications include GCIH, GCFA, GCFW, GWAPT and CISSP.

EXECUTIVE SUMMARY

Every week there's a new piece of malware or botnet in the news threatening enterprise and home users. Indeed, it seems we can't go a day without hearing about malware affecting a co-worker's or family member's computer, or about some botnet wreaking worldwide havoc. In addition, malware is getting more sophisticated—it often combines worms, Trojans and bots, and can morph automatically to prevent detection. Malware is everywhere, and because of the low barrier to entry for cybercriminals buying point-and-click malware kits, not to mention the potential for substantial reward, the situation is likely to get worse before it gets better.

During the last five years, it's become painfully obvious that traditional antivirus products on their own are not able to keep up with the onslaught of new malware. Antivirus vendors are supplementing their products with collective analysis and powerful processing afforded by the cloud, but they are still challenged to adapt quickly enough to meet today's threats.

A new category of products has emerged to help deal with unidentified malware. Network-based malware-detection systems and malware sandboxes go beyond the traditional signature-based and limited heuristics capabilities offered by antivirus vendors. These systems can run suspicious files through virtual machines and monitor for malicious behavior at the network, file system and process layers. Their goal is to detect and prevent malicious files from ever making it into the target network. After all, if the files don't reach the desktops or servers, they can't compromise them.

However, no one product—no matter how sophisticated—can detect all malware, nor can it replace a layered security system. A combination of technologies and best practices can aid enterprises in the fight to detect and stop advanced malware attacks before a serious breach occurs.

Rooting Out Sophisticated Malware

Malware authors are developing new malware variants at a breakneck pace. Not so long ago, malware defense meant recognizing a virus or a Trojan horse and eradicating it. But today's advanced malware is designed to be resistant to detection and removal. Malware authors also have developed many new techniques for hiding malware or making it appear benign by tunneling its command-and-control traffic as part of standard HTTP or encrypted HTTPS traffic. In this special report, we offer a look at some methods for recognizing advanced malware and mitigating the effects of malware should it make it past your organization's defenses.

The goal of enterprise malware-prevention efforts should be to stop malware from ever getting to the desktop. To do that, analysis, detection and prevention need to take place at the network layer. Starting at the perimeter, content filtering gateways, next-generation firewalls and new network-based malware-detection appliances provide the first layer of defense. They have the ability to analyze traffic, detect malicious files and prevent malware from ever getting to its intended target. The concern, of course, is whether these systems can keep up with the ever-increasing number of new malware specimens being released

Figure 1

Layers of Security

Shown is an example of a layered approach to security, leveraging desktop antivirus, network malware-analysis appliances, content security gateways, and cloud-based intelligence and scanning services.

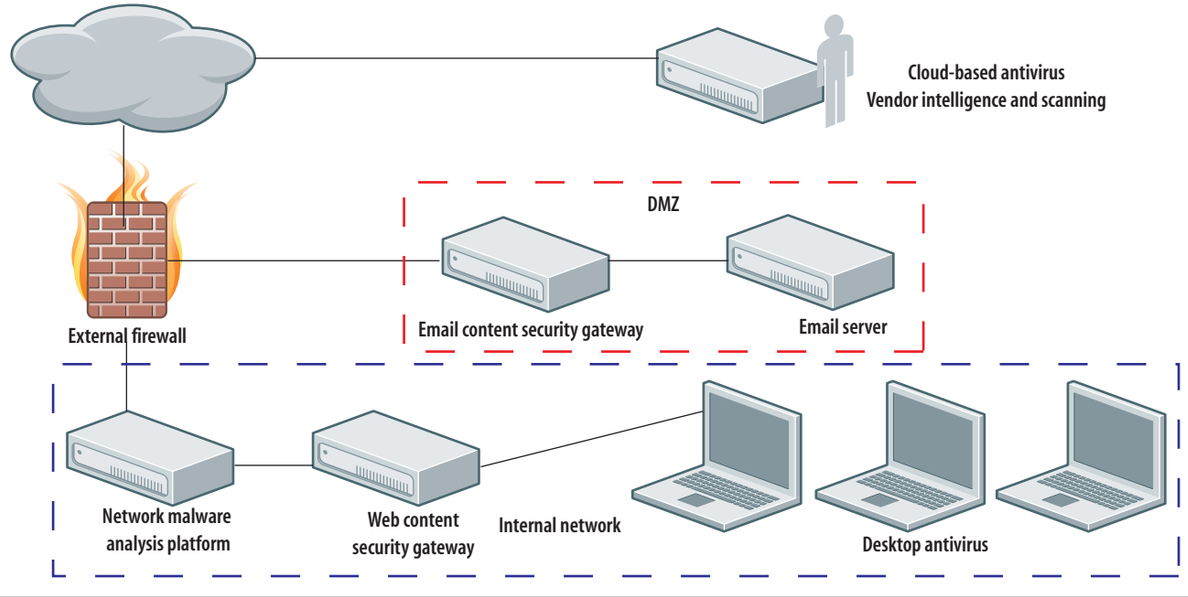

S4880512/1

daily, and whether they can efficiently deal with increasing network throughput demands.

To supplement network-based malware-detection products, many vendors are turning to cloud-based services to off-load analysis and computing capabilities. Cloud-based computing services provide more computing power so more malware samples can be analyzed, and they serve as a collective analysis resource.

And while we would prefer to stop all malware at the network level so it never reaches the desktop, we know that's an unrealistic goal. Desktop antivirus still has a place, and many desktop antivirus vendors are using the same cloud-based services for file comparison and reputation lookup as the network detection systems.

No matter what technology you're using, traditional best practices should be implemented and followed to provide the best opportunity for detection and prevention. These include applying the principle of least privilege, aggressive patch management of Internet-facing hosts and workstations allowed to browse the Internet, separation of privileges, change management to detect operating system changes, and log monitoring for detecting anomalous events.

Figure 2

Industries Being Targeted by Advanced Attackers

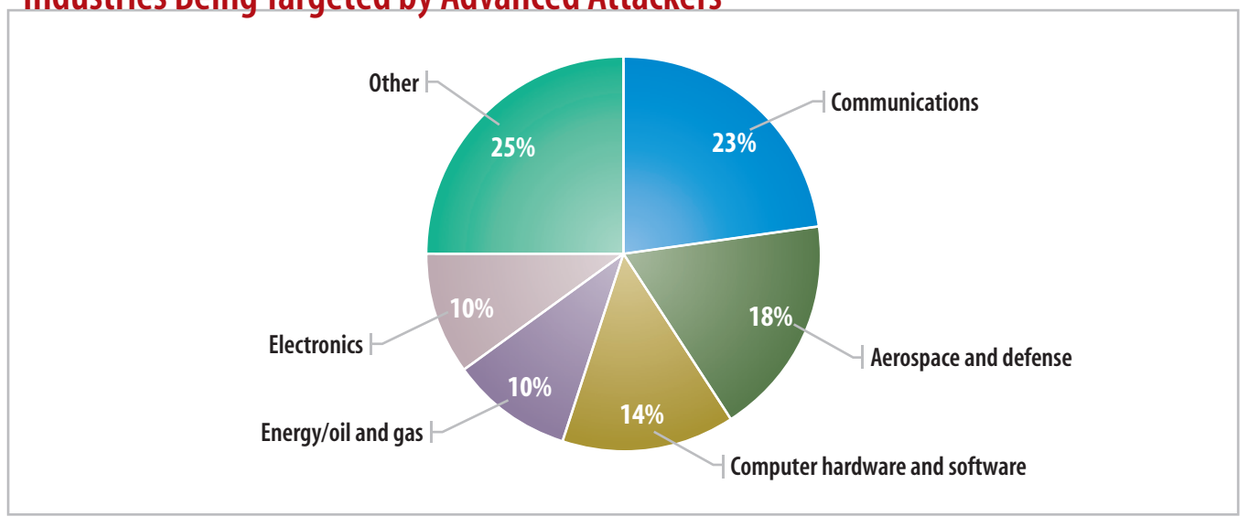

Mandiant M-Trends 2012 Report S4880512/2

The Woes of Content Inspection

One of the biggest network security challenges is content inspection of network traffic at wire speed. At the network perimeter, content security gateways and next-generation firewalls inspect network traffic in transit and block malicious content. These systems have traditionally relied on signature-based antivirus to detect and block known bad files or websites.

The antivirus engine used within these products is often an OEM system, and the ability to block known files is only as good as the

Strategy: Threat Intelligence: What You Really Need to Know

If there was ever a time when threat intelligence could be put on autopilot, that time is over. With the increase in advanced, multidimensional threats, organizations can no longer depend solely on existing gateway tools to weed out nefarious activity. More and more organizations are considering development of an in-house threat intelligence program, dedicating staff and other resources to deep inspection and correlation of network and application data and activity. In this report, we will examine the drivers for implementing an in-house threat intelligence program, the issues around staffing and costs, and the tools necessary to do the job effectively.

Download

included signatures. New or targeted malware can easily bypass these devices. To adapt, vendors are leveraging cloud-based services to supplement their detection capabilities for file comparison and reputation lookups.

This makes sense: Consider the large number of files that need to be analyzed and the amount of network traffic within an enterprise that needs to be inspected. Leveraging the collective knowledge gained through cloud-based intelligence and analysis supplements any weaknesses in antivirus scanning engines. Content security gateways can consider the reputation of the source of the file, the content of the file, and any other information that can be gleaned from the file. That data can then be compared with data collected through the analysis of files from other customers.

Next-generation firewalls were created to bridge the gap between firewalls and content inspection. They provide some of the same capabilities as security gateways because they are application-aware and can understand protocols such as HTTP, SMTP and FTP, as well as application-specific traffic like instant messaging and file sharing.

The really big benefit of next-generation firewalls is that rules can be written to prevent application-specific traffic, rather than traffic based on IP addresses and port numbers. The drawback to performing so much content inspection on a next-generation firewall is that the process can impact the firewall's core functionality. The more analysis performed on traffic as it passes through, the more likely that latency that impacts performance will be introduced. However, even with those concerns, these firewalls are providing better tracking and blocking capabilities out of the box than traditional firewalls. Add to that the fact they can be used to leverage reputation-based information and block known bad domains and IP addresses to prevent malware from coming in, and it's clear that next-generation firewalls have promise.

Content security gateways and next-generation firewalls provide a starting layer of protection at the perimeter, but they are limited in the ability to fight against sophisticated malware and targeted attacks. They don't, for example, provide dynamic analysis to detect malicious behavior during the execution of suspicious files. This is where newer network-based malware-detection devices step up to try to fill the gap.

Dynamic Analysis on the Fly

Network malware-detection appliances such as those from FireEye and ValidEdge have the ability to dynamically analyze suspicious files to determine whether they are malicious. All of this is done at the network layer using on-board virtual machine and emulation technologies. The appliances open or execute files within a virtual machine, or emulator, and profile the files' behaviors to determine if they are malicious. The appliances then decide whether to allow the file through or sound some kind of alert.

During analysis, the products look for outbound connections to known malicious command-and-control servers, modifications to the Windows registry, creation of new services, code injection into running processes, and other suspicious activities.

These are the same kinds of things that a malware analyst or incident responder would look for during the investigation of a possible malware infection. The process is intense and laborious when done manually; the appliances attempt to automate the process in real time to make quick decisions as network traffic passes through.

Network malware analysis appliances evolved from standalone malware sandboxes, which are used to perform analysis on an on-demand basis with little to no impact on the local system or network—often, completely separate from their network because they are Web-based services. Some of the more well-known sandboxes are Anubis, GFI SandBox, Joebox, Norman and Cuckoo. These types of sandboxes emulate or completely virtualize a Windows system to monitor what impact a suspicious file would have on a system. Configuration changes such as modifications to the registry or the addition of new services are reported, along with attempts to start new processes and perform network communications.

Figure 3

Malware Infection Vectors by Percent of Breaches Within Malware

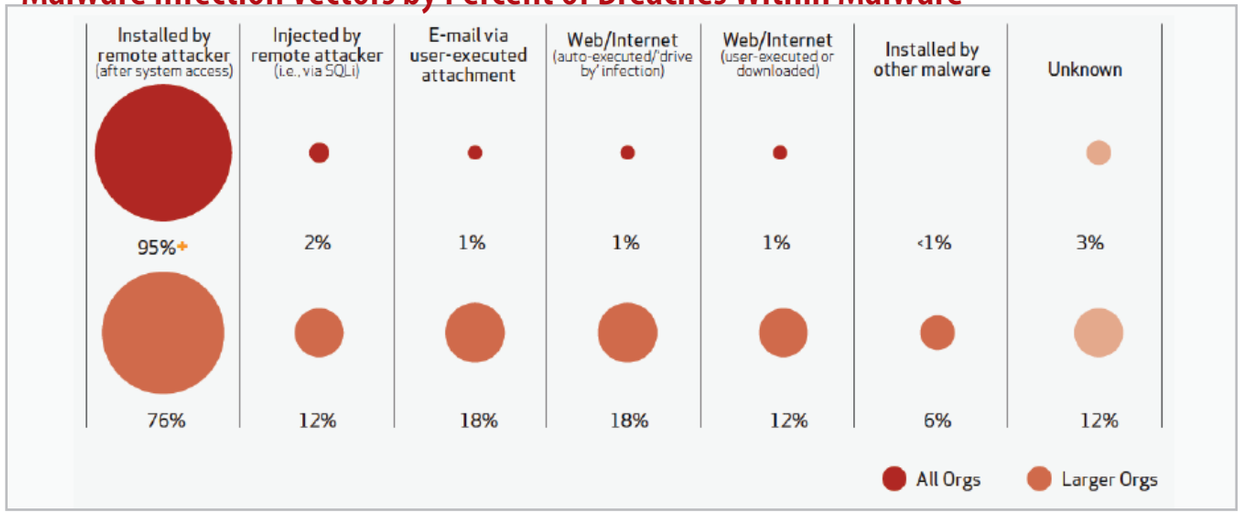

Verizon Data Breach Investigation Report 2012

S4880512/3

One of the great benefits of sandbox systems is that they are typically fast and can provide quick intelligence that can be utilized for creating new intrusion-detection systems and firewall rules. The downside is that they could pose a risk if their protection mechanisms are bypassed, leading to a live malware infection within your enterprise network. Attackers are certainly aware of sandbox technologies and have been working to defeat them and obfuscate their malware's behavior to avoid detection.

Automated malware-analysis systems alone are not an easy fix for detecting advanced malware attacks, nor are they inexpensive. To perform real-time analysis of network traffic, a beefy box is necessary, and with that comes an equally beefy price tag. There are several

very capable free and/or open source sandbox tools available (including Cuckoo and Zero Wine), but their ease of use and configuration complexity vary greatly. In addition, some organizations may be fearful of submitting files from their environment, possibly revealing that they were the victim of a targeted malware attack.

The Pains of Manual Labor

Network-based malware detection and sandboxes do a very good job of automating dynamic analysis of suspected malware, but

Figure 4

Security Breaches Over Past Year

Which types of security breaches or espionage have occurred in your organization in the past year?

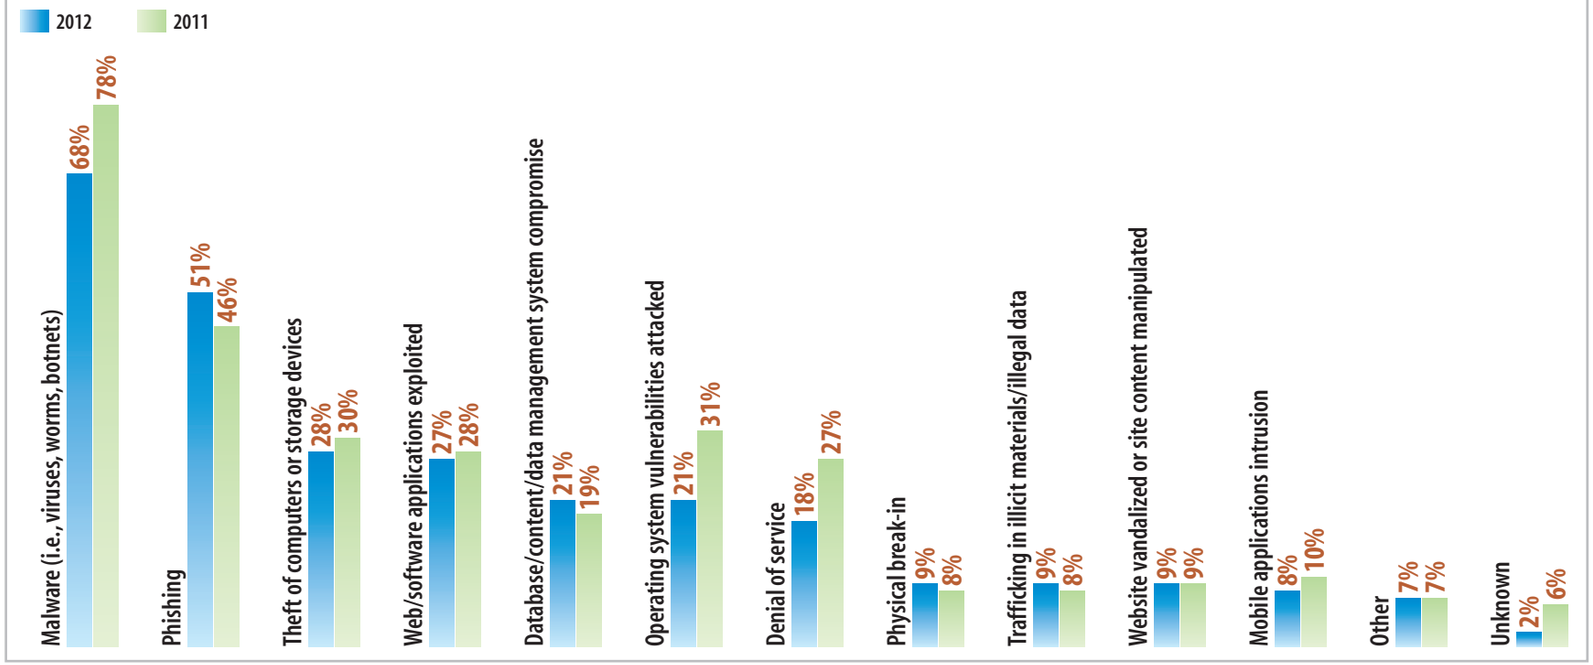

Note: Multiple responses allowed
Base: 183 respondents in March 2012 and 219 in March 2011 experiencing a security breach within the past year
Data: *InformationWeek* Strategic Security Survey of business technology and security professionals at organizations with 100 or more employees

R4670512/5

they do not negate the need for manual analysis. There will be cases in which automated technologies simply cannot adequately analyze certain types of malware sepals. To make matters worse, malware analysis can be a long, tedious process that involves specialized skills to reverse engineer the suspicious file, determine its features and functions, and perform dynamic analysis by executing the suspicious file on a live system and monitoring the changes made to the system.

The goal of enterprise malware-prevention efforts should be to stop malware from ever getting to the desktop.

The reverse engineering process may include disassembling the malicious binaries using Ida Pro disassembler, live execution analysis through a debugger and use of some of the newer malware-specific reverse engineering tools. These include HBGary Responder and AccessData's new Cerberus malware triage tool, which helps automate part of the reverse engineering and classification of unknown files. They understand what common attributes of malware are and can help identify malware based on their internal functions. All that information is used to develop a profile for malware to stop it or detect it once it has already made it onto systems.

When dealing with a targeted attack, it is critical to understand the capabilities of the malware used, including its propagation methods, how it persists on systems once it infects them, its purpose and any communication channels it may use for interaction by the attacker. Developing an accurate profile for malware is critical in order to identify systems throughout the enterprise that have been compromised. The attributes within the profile—often referred to as indicators of compromise—consist of processes, file names, Windows registry entries, event logs, network traffic and any other bits of information that can uniquely identify a piece of malware. The profile can be developed through several methods, including sandbox analysis, dynamic analysis and forensic analysis of known infected systems.

The key here, of course, is to have as accurate a profile as possible. Understanding how the malware propagates can help identify systems that could potentially be infected. This information could be cross-referenced through patch management systems and vulnerability scanners. Any information about exfiltration data or network traffic used by the malware to communicate back to the attacker can also be useful. New IDS and firewall rules can be put in place to detect and stop the communications. Monitoring DNS lookups and even inserting bogus entries for known bad domain names can help with detection and prevention.

Rooting Out Compromised Systems

During the last few years, we've seen an emergence of new enterprise incident response tools, including AccessData Enterprise, Carbon Black, F-Response Enterprise, Encase Enterprise and Mandiant Intelligent Response. The capabilities of these products vary, but the goal of each is to empower incident responders and enterprise security

professionals to perform incident response procedures on a large scale via agents on the desktop or centralized collection of data.

This is where creating that profile of malware becomes extremely important, because those attributes from the malware can be used to search across many systems or in the logs to determine which machines have already been compromised. They are extremely powerful, with features that include memory analysis, remote disk imaging and remote forensic analysis. If analysis of malware shows that certain files or services are created on a compromised system, these tools can enable security pros to quickly search all computers within the enterprise to determine if compromised files or services exist on any of the systems.

Taking the Fight to the Streets

Even with malware authors using new and ever-more-clever techniques to cover their tracks on compromised systems, there remains the simple truth that most malware wants to be persistent. Some malware is designed to be a downloader, whose sole purpose is to make it onto target systems and download additional malicious components. While downloaders are temporal and serve a singular purpose, the components they download will seek persistence on the target system to carry out their nefarious activities. When the malware takes residence on the system and sets itself up to survive system restarts, it gives itself away by making changes to the victim computer system.

How do we know what to look for? We've already discussed the indicators of compromise that can be used to profile and detect malware using enterprise incident response tools. But traditional best practices, including change management, log monitoring and the principle of least privilege, can also be effective in this battle.

The 2012 Verizon Data Breach Investigation Report states that 84% of the breaches Verizon investigated could have been identified in process had the victim organizations been monitoring their logs. The very logs that held the evidence that helped Verizon determine what happened during the breach were the same ones that the victim could have been monitoring to catch the intrusion. According to the report, useful compromise indicators include log file line count, log file line length, spikes in traffic types, country origin of the IP connection and email message sent/received.

Unfortunately, logging seems to be a difficult undertaking for many organizations, which is surprising considering that logs are generated by every operating system, network device and service out there. Centralizing logs and performing basic analysis for anomalies such as those mentioned earlier don't take much time, effort or money. The Verizon report goes on to say, "The really interesting thing about this type of monitoring is that it doesn't take a ton of cash to implement an effective solution. It can be done with a few commands on a Linux or Windows system."

The perception is that centralized logging requires an expensive enterprise logging platform with a specialized Web interface and

secret data correlation sauce under the hood. Instead, there are many low-cost and free options out there that can enable companies to start collecting, centralizing and analyzing log data today for malware activity. For example, all antivirus products create logs—most often, right there in the Windows Event Log. The logs can be centralized using free syslog tools and easily searched using command-line tools or automation.

Configuration management databases and change management monitoring tools could be used to detect new services that are created to allow persistence for malware. Some change management tools have the ability to perform file integrity checking. New files or changes to files within the Windows system directories outside of normal patch times might indicate malware attempting to establish a permanent home on the system.

Even with malware authors using new and ever-more-clever techniques to cover their tracks on compromised systems, there remains the simple truth that most malware wants to be persistent.

Abnormal network behavior and previously unseen network communications are other areas that can be monitored using existing tools—that is, of course, assuming that netflow-enabled routers, firewalls and intrusion-detection systems are already in place. These devices perform logging and are often logged to a central server, but their logs often go unchecked. Tools such as Tenable Security Center enable security pros to correlate that information to identify hosts offering network services for the first time, hosts communicating with known malicious servers and domain names, spikes in network traffic, and new network protocols—all potentially signs of an attack.

MORE LIKE THIS

Want More Like This?

InformationWeek creates more than 150 reports like this each year, and they're all free to registered users. We'll help you sort through vendor claims, justify IT projects and implement new systems by providing analysis and advice from IT professionals. Right now on our site you'll find:

Strategy: How to Boost Security via FFIEC Compliance: With just a smartphone, users can conduct nearly all their banking business at any time. However, all this flexibility and convenience opens up new avenues for fraud and cybercrime. Guidelines laid out by the FFIEC several years ago predate many of the capabilities—and vulnerabilities—that are in place today. In this report, we examine the latest guidelines and provide advice on how you can extend the work done to comply with FFIEC guidelines to strengthen your organization's overall security posture and keep customers and their data safe.

Research: 2012 Strategic Security Survey: When it comes to security and risk management, it's tempting to try to address everything. A more effective approach: Focus on the most likely threats. Our survey results show security and IT pros are concentrating on risks over which they have some control, such as implementing better access control, vetting cloud providers, safeguarding mobile devices, educating users and building more secure software. See what else should be on your list.

Strategy: Monitoring and Measuring Cloud Provider Performance: There is no ignoring the cloud, which means that IT pros must find ways to monitor and measure the performance of cloud providers. Just as security groups often struggle with managing security inside a company when in a governance role, we struggle with governing the security of assets that no longer sit within our own data centers. The challenge is to develop and implement a strong governance model for cloud offerings that ensures that security is part of the conversation.

PLUS: Find signature reports, such as the *InformationWeek* Salary Survey, *InformationWeek 500* and the annual State of Security report; full issues; and much more.