WHITE PAPER:
ENTERPRISE SECURITY

A comprehensive approach to unique and targeted attacks

Turning the Tables on Malware

Who should read this paper

Enterprise Information Security executives and teams can use this document to understand a new reputation-based security technology that automatically identifies and blocks even highly obfuscated malicious code with no user intervention and improves system performance.

Confidence in a connected world. Symantec.

Content

Introduction

Malicious code is in a constant state of malignant evolution, finding new forms, disguises, and vectors to reach, penetrate, and compromise its targets. Signature-based antivirus software is effective against viruses, Trojans, and worms for which signatures have been catalogued. It is also relatively efficient, although the performance impacts of scanning for signatures grow with file size, compression, complexity, and number of signatures that must be scanned. But signature-based defenses are useless in encounters with new malware.

Because of this vulnerability, signature-based defenses are typically backed by additional lines of defense—often tools that look for so-called generic signatures, and advanced heuristics that look for evidence of malicious behavior. These add substantial protection, but have much more impact on performance than scanning for signatures. Worse, because they are heuristic—that is, not empirically validated—they have lower accuracy, and generate a much higher rate of false positives, which reduce employee productivity and create perverse incentives for companies to dial back protection.

A tectonic shift

Data collected by Symantec for its annual Internet Security Threat Report document an overwhelming shift in the creation and distribution of malicious software, from mass distribution of a relatively small number of threats to narrow distribution of a large number of unique threats. The change threatens to upset the current balance between signature-based and heuristic defenses, making both approaches less effective, and leaving enterprises exposed to a barrage of unique attacks. Left unaddressed, it may force enterprises into unacceptable compromises between security and performance.

Some of the key findings in the report:

- **More attacks:** 5.5 billion attacks in 2011 represented an 81 percent increase over 2010.
- **More unique threats:** 403 million unique threats were seen in 2011; in 2010 this number was 286 million.
- **More Web-based attacks:** 36 percent increase in web-based attacks in 2011 versus 2010.

Mass mailings and other "traditional" attacks are still common, and years-old malware still finds thousands of unprotected computers to exploit. But the trends are clear: in the very near future, the online environment will contain *hundreds of millions* of threats that may overwhelm signature-based defenses. Relying on heuristics alone will compromise the efficiency and productivity of enterprise networks worldwide. The purpose of this paper is to identify new ways to detect malicious code, and outline requirements for more comprehensive protection.

Dimensions of detection

Signature-based detection searches for known patterns of data within executable code. If the code is new or rare, no signature may yet exist. "Generic" signatures depend on an executable's similarity to previously identified code, and are similar to signature-based defenses, although less accurate.

Advanced heuristics test software's behavior. This takes time, raising the probability of performance impacts on the host machine, or that the software itself will be overwhelmed in a sustained or high-volume attack. And because heuristics are less accurate than signature-based defenses, a barrage of new malicious code raises both false alarms and the probability of a miss.

We need a new line of defense to augment efficient but porous antivirus defenses and less-reliable, more resource-intensive heuristics. Such a defense would allow the antivirus layer to block known threats, while keeping even the majority of new threats from reaching the advanced heuristics.

Context is the key to this new line of defense. Context comprises all the information that can be collected about a file that is not embodied in the file itself, that is, information beyond its contents and behavior. Since the definition is open-ended, no list of context elements can be exhaustive, but it includes at a minimum:

- **Prevalence**—how many instances of the file exist?
- **Emergence**—when did the first instance of this file appear?
- **Origin**—where did it come from?
- **Connections**—with what other files is the file of interest connected or co-located?
- **Experience**—has it been scanned by other machines on the Internet—and with what results?

Context-based defenses exist that rely on whitelists of trusted files and domains, and blacklists of suspicious ones. These are incomplete and, like signatures, ineffective against unknown threats: incomplete because they exclude “grey files” not yet identified as malignant or benign, and ineffective because they have no information at all about new, unique, or targeted files, and must pass them through for heuristic analysis with its associated processor overhead.

Requirements

To address a threat environment with hundreds of millions of new threats, defenses must augment signature-based and heuristics defenses with efficient context-based defenses that use a much broader definition of context than white lists and blacklists. The challenges are steep, but the following requirements outline an approach that turns hackers’ greatest strength against them: the ability to create an unending stream of new, unique threats that may be targeted against individual companies.

1. Analyze and catalog every accessible executable file on the Internet

This approach requires identifying, logging, testing, and rating every file likely to contain a virus, Trojan, or worm, and assigning it a trust rating. Technically unfeasible only a few years ago, comprehensive analysis in full context is the only way to measure “prevalence” and “emergence” independently, and the only way to approve low-prevalence but (ultimately) trustworthy files.

Comprehensive analysis is also a key to performance. Once a file has been rated, defenses can skip known good and bad files, so they won't be swamped by proliferating copies of malware, waste time revalidating trustworthy files, or raise unnecessary false positive alarms.

This analysis must be done on a global basis, which requires an enormous network of sensors and a lot of computing power. The larger the network, the more accurate detection results will be, and the lower the burden on any individual sensor or server.

2. Determine prevalence and emergence data for every unique file

Again, the numbers are immense, but continuous cataloguing can keep prevalence and emergence information up to date for every file in the database.

Prevalence and emergence data are important because of their relationship to trustworthiness. Malicious code derives its advantage from stealth—it is rare, often targeting a single organization or even computer, and it relies on "zero-day" emergence to evade signature-based

detection. Malicious code that is widespread and mature is more likely to have been identified, and its signature published. In contrast, low-prevalence, recently-emerged files are the riskiest.

3. Determine the origin and connections of each file

This step is important to see through email or Web address aliases and identify files of suspect origin, or those associated with files of suspect origin. The technologies to this are all ready well-established—security companies already use them to identify:

- **Command-and-control servers and botnets**—servers and the networks of computers they recruit to propagate spam, organize denial-of-service attacks, and spread more malicious code
- **Malicious Web servers**—computers known to spread spam or malicious code, or implant such code on visitors through browser vulnerabilities
- **Devices**—USB “thumb” drives, mobile phones, and other unusual vectors for uploading executable code to an enterprise network endpoint
- **File-sharing executables**—applications used to exchange pirated media and software also at high risk to implant malicious code

These relationships may be subtle, but data-mining techniques originally developed to find hidden associations and dependencies in large databases are well suited to tease out associations between infections, websites, and files, even on a global scale.

4. Apply layered detection technologies to each file

These are the same technologies applied on enterprise networks, in corporations and on consumer devices. The advantage of applying them en masse is to offload processing burden from those networks and endpoints by performing massively parallel detection in the "cloud" of global networks. The technologies involved include:

- **Network analysis**, looking at individual packets to identify patterns of attacks
- Signature-based antivirus technologies, to determine whether the executable has been previously identified as malicious
- **Static behavioral profiling (heuristics),** to look for suspicious code inside files
- **Dynamic or real-time behavioral analysis**, testing code in an emulator or "sandbox" to look for suspicious activity and targets

Context isn't a standalone detection tool; instead, it improves the focus and performance of other detection technologies. Prevalence, emergence, origin, and connection data are correlated with malicious intent, so signature and heuristic scans can bypass files known to be safe, heuristics can scan suspect files more deeply, and firewalls can block connections known to be malicious without prompting users.

5. Deliver actionable data

Administrators and users need to understand the risks embedded in the executables they download and manage, and this approach delivers the information they need with full context, so they can make informed decisions.

A corollary benefit is the improvement in security awareness when administrators and users see, for example, that a downloaded file has never been seen before, comes from an unreliable or unknown source, and contains code that resembles—though it may not exactly match—known malicious software.

Users today are expected to make important security decisions without context: whether or not to visit sites, install files, or launch executables. Ideally, users would be given context and risk information before they click to install programs or to allow them to access their

data and connections. Every file download into a corporation should be accompanied by risk information so that administrators and users can understand the risks embedded in the executables they download and manage.

6. Set policies based on risk tolerance

In enterprise IT environments, policy management relies heavily on automation, to keep enforcement consistent across the organization and produce reports that document compliance to internal auditors and external regulators. IT Professionals must now extend their policies to incorporate new prevalence, emergence, and connection information. A hypothetical set might include:

- **IT Development Environment** (high risk tolerance)—permit installing low-reputation software
- **Financial data** (low risk tolerance)—do not permit access by applications with low reputation
- **Confidential data** (low risk tolerance)—do not permit access to payment card, customer database, or personnel files by applications with medium to high risk
- **Finance department** (low risk tolerance)—may only install high-reputation software

The additional granularity provided by the new information allows policies to track risk tolerance precisely—delivering security where it's required without compromising performance where it's most useful. And it allows IT professionals to establish endpoint, Web or network blocking policies based on risk tolerance. When they know, for example, that a file was recently created and has low prevalence, they can decide to increase protective measures based on the risk tolerance appropriate for a given department, network domain, or device.

7. Close the loop

An approach like this one also makes it easy to close the loop between security organizations and their enterprise clients, by reporting information on new files to a centralized database. With each contributing user and contributed file, the database will become more powerful, faster, and more useful. And because the database is global, there will be no enclaves where it is safe for malware to hide and propagate.

The growth of accuracy with sample size, and the shared motivation of all participants to keep malware under control, combine to create a virtuous spiral of improved malware identification, resulting in fewer misses and false alarms, resulting in faster scans and lower resource overhead, releasing more resources for the identification and evaluation of the remaining code.

How this approach changes the game

This new approach offers better malware detection with fewer false alarms and higher performance than today's security stack. Today, cybercrime has the upper hand, because criminals can quickly customize and execute targeted attacks. This new approach turns their model on its head: new and custom attacks with low emergence and prevalence ratings will be blocked by all but the most risk-tolerant environments. And "mature" malware will be discovered, fingerprinted, and then blocked with less scanning and heuristics overhead than today.

At the same time, the approach cuts false positives by confirming the legitimacy of files sent through reputable channels, and by checking the reputation of suspect software in quarantine before permanently removing it.

Finally, the approach reverses the effect of the malware explosion on performance, by moving the evaluation and rating overheads—and much of the scanning—off enterprise clients and endpoints. Because the new ratings offload work from both signature—and heuristics-based defenses, every one of the security technologies works less, and therefore better.

Conclusion

Custom, focused, and malicious code, targeting enterprises and created for financial gain, is a serious threat. But a comprehensive approach to establishing and rating file context promises to reverse the value of stealthy and targeted attacks, and create a virtuous spiral of increasing IT security with reduced overhead.

How to get started

Insight is available in several Symantec and Norton products today. These products include Symantec™ Endpoint Protection, Symantec™ Endpoint Protection Small Business Edition, Symantec™ Endpoint Protection.cloud, Symantec™ Web Gateway, Norton™ Antivirus, Norton™ Internet Security and Norton™ 360™.

Learn more about our detection technologies and enterprise security products from your Symantec representative, or at http://go.symantec.com/insight.

About Symantec

Symantec is a global leader in providing security, storage and systems management solutions to help consumers and organizations secure and manage their information-driven world. Our software and services protect against more risks at more points, more completely and efficiently, enabling confidence wherever information is used or stored.

For specific country offices and contact numbers, please visit our website.

Symantec World Headquarters
350 Ellis St.
Mountain View, CA 94043 USA
+1 (650) 527 8000
1 (800) 721 3934
www.symantec.com

Symantec helps organizations secure and manage their information-driven world with **security management**, **endpoint security**, **messaging security**, and **application security** solutions.

6/2012 21155056-1