Realtime
publishers

The Essentials Series: Increasing Performance in Enterprise Anti-Malware Software

Why is Traditional Anti-Malware So Slow?

sponsored by

GFI

by Eric Schmidt

Why Is Traditional Anti-Malware Software So Slow? 1

Traditional Anti-Virus Software Is Slow 2

Multiple Endpoint Products Performing the Same Function 3

Scan Methodologies 3

Real-Time Scans 4

Scheduled Scan 5

Client Reporting 6

Anti-Malware Shouldn't Be Slow 6

Copyright Statement

Why Is Traditional Anti-Malware Software So Slow?

Since the beginning of the Internet, anti-malware software has been an essential component of every business computer. As the Internet continues to grow in scope and sophistication over time, so does the intelligence built into malicious software. Today's malware has evolved from its origins of viruses as a mostly harmless demonstration of self-replicating code to a full-fledged underground industry with its own motives for profit.

To that end, the malware threat landscape has changed dramatically in recent years. Today's threats are being created for far more insidious goals, including cyber-terrorism, identity theft, and corporate espionage. This reality has been made easier with the introduction of shrink-wrapped malware toolkits that enable the point-and-click invocation of attacks. These toolkits are easily downloadable through online forums that facilitate the collaboration and development of such attack software.

The solutions that protect computers from these types of attacks were developed shortly after the first viruses. Many of those products are still around today. To combat the ever-changing tactics of malware writers, those early "anti-virus" products have also evolved to address new threats. Unlike with malware, however, this sophistication in protection products comes at a cost. That cost relates to the actual system resources necessary to run their protective processes.

In effect, the slow evolution of anti-malware software over the long term has created a spaghetti of detection and removal mechanisms. This aggregation of today's solutions atop yesterday's code bases creates a major performance problem for enterprise computers. That loss in performance impacts employees' ability to accomplish their needed tasks, reduces business processing agility, and increases the computing infrastructure's total cost of ownership.

One solution for this problem involves an entirely new approach to creating anti-malware solutions for business. This approach throws away the aging code base of the past in favor of truly recreating the wheel. By electing to break with the software of the past, new and highly-optimized solutions can be developed to protect from malware while ensuring the best possible user experience for your desktops, laptops, and servers.

This, the first of three articles in this series, will discuss why many traditional anti-malware solutions have a negative impact on system performance. It discusses the risks to business operations that are created as a result of this performance degradation. Continuing with this discussion are two subsequent articles that focus on the factors you must consider when selecting an anti-malware product as well as what can be done to optimize your selected solution's performance.

Traditional Anti-Virus Software Is Slow

Let's face facts. Traditional anti-malware solutions are a painful but historically necessary function of computing. The fact is there is a problem intrinsic to any software solution that has evolved over a long period of time. Most traditional anti-malware solutions in use today are the result of years of development, with much of their evolutionary updates done to the same code base originally created for the ancient operating systems (OSs) of yesteryear.

Like any software company, Microsoft releases new OSs every few years. Each of those OSs includes dramatic changes to their core kernel. Those changes mandate equivalent changes to protective software such as anti-malware solutions. At a high level, anti-malware solutions operate very "close" to the kernel, intercepting file system calls and monitoring processes and process threads. Architecturally, anti-malware's close proximity to the kernel itself requires it to evolve with the OS. Ultimately, as Microsoft releases new OSs, anti-malware solutions must change to support each new version. One problem is that such solutions must also support legacy OSs. As such, traditional anti-malware solutions grow heavier and heavier with each new OS release, making their use less optimized over time.

Further, the ways in which the "bad guys" write malware evolves over time as well. The growth in malware code base accommodates each OS release, addresses new attack vectors, and includes new features. The vast majority of commercial anti-malware software has grown from simple virus scanning to comprehensive suites that provide protection from viruses, malware, and spyware; some also include firewall features. The goal of these products is to provide total endpoint protection from every possible angle. Although they can provide greater protection than their predecessors offer, much functionality has been bolted on and shimmed into the products in a way that places much higher demands on system resources. The higher resource requirement forces companies to make critical decisions about how their systems will be protected.

The first option is to purchase more expensive hardware (more RAM, faster processors) to ensure that the system and anti-malware suite perform at acceptable levels. This approach may resolve the issue; however, IT budgets are tight. Dumping money into hardware just to accommodate slow software is always a poor business decision. The most common solution is to limit the components that are implemented or come up with creative solutions in order to maintain optimum performance. Ultimately, throwing hardware answers to what is really a software problem is not a smart solution. Needed are improvements to the software itself that reduce the performance impact of anti-malware solutions overall.

Multiple Endpoint Products Performing the Same Function

There are also circumstances in which an anti-malware suite may excel in one area but be deficient in another. To accommodate for deficiencies, an enterprise may choose to install multiple products on the same system to achieve the best level of protection. This overlap can have a huge impact on performance and potentially result in an unstable system or products that are ineffective.

A common issue that occurs when multiple products are installed is an overlap in functionality. For example, each installed product is performing real-time scans simultaneously. In this situation, each product requires every file to be scanned before it is available to the user. This delay in opening files and applications will be very visible to the user. In some cases, the conflicting products will create a situation in which files are not scanned at all or they are never allowed to be opened.

Running two or more products also complicates the task of troubleshooting application and performance issues. During the troubleshooting process, a determination has to be made as to which product is doing what. This will help determine the product in need of further investigation. Once the offending product has been identified, the issue may be resolved, but at what cost? What was the offending products' role? If it is uninstalled or disabled will that put the system at risk? Can the product be reconfigured so that the issue won't return?

Although it might be possible to run two products in a complementary manner, rare is the vendor that will recommend and/or support such a configuration.

Scan Methodologies

Another contributor to slow performance with anti-malware software involves scan methodologies and the engines that run them. Before examining the different scanning methodologies, it is important to discuss the processes and mechanisms that enable scans to work in the first place:

- Malware Signatures and Definitions—The first and most well-known method involves malware signatures or definitions. These definitions encompass a list of known viruses or malware and the marks they leave on an infected system. These breadcrumbs are used when a scan is performed. While performing a scan, the engine uses signatures as a reference to compare against the OS state and files. If something is found that matches a signature, it's marked as infected, blocked from execution, and quarantined or deleted.
- Heuristics—Heuristics are used to identify infected files where no signature is readily available. The benefit of an effective heuristics engine is that systems are no longer exclusively dependent on the receipt of signatures in order to stop an infection. Heuristic scanning is a very complicated process that infers certain behaviors as being malware-like. Due to this complexity, in some cases, legitimate files can be marked as infected.

- Behavioral Analysis—Like heuristics, behavioral scanning provides the ability to detect malware without relying on definitions. Behavior analysis examines the actions of a program in order to identify malicious activity. Some examples of suspicious program actions include writing to protected parts of the OS or registry. This is considered a malware-like activity because legitimate applications don't display these behaviors. Eliminating the dependency on definitions for threat detection has tremendous benefits; however, there can be a risk of false positives when legitimate but poorly written applications are used.
- Emulation—Another method that has been developed involves emulation. Here, a potentially malicious application is opened in a protected area to identify malicious behavior. This protected area emulates the OS so that files can be opened and analyzed for malicious activity in a protected and temporary environment. If the file is infected, it can be allowed to perform all its actions, which then allows heuristic and behavioral analysis to be fully performed without harming the actual system. Once the execution and analysis has completed, the emulated environment and all the changes that were made are safely removed. Good heuristics and the use of emulation have become increasingly important in confronting modern attacks, as the possibility of zero-day exploits has dramatically increased in recent years.

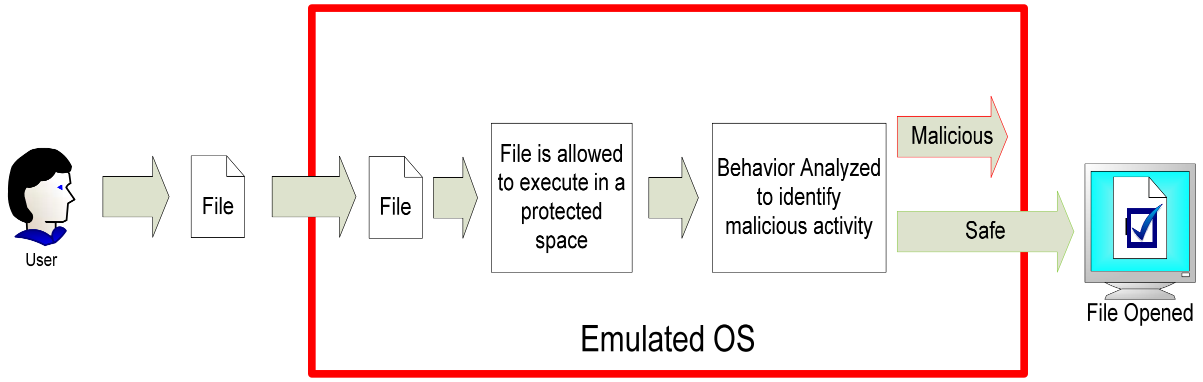
User
File
File
File is allowed to execute in a protected space
Behavior Analyzed to identify malicious activity
Malicious
Safe
File Opened
Emulated OS

Figure 1 – Emulation used to analyze file

Real-Time Scans

Real-time scans are often the first method of scanning to prevent systems from being infected with viruses and malware. Real-time scanning requires a client that runs continuously and monitors every file that is opened or executed. When a file is opened, it is first scanned by the engine and evaluated using one or more of the methodologies previously described. Real-time scans act as the first line of defense against malware because they monitor the OS and all attempts to change protected areas such as the system files, registry keys, and system services.

Traditional real-time scans can have a negative impact on performance because the file being opened must first be scanned. If the anti-virus software is slow or poorly written, files will take longer to open and programs will run slower because they must wait for the real-time scan to release its hold on files. This can have a direct impact on user perception of overall computer performance, resulting in decreased user productivity and unnecessary calls to the Service desk.

There are also risks to other computers on the network if real-time scanning is slow. It is possible for a system to become overloaded to the point where there are communication delays between the agent and its management server. In these situations, an overloaded agent may not be able to receive the latest definition files or provide status reports. If this were to occur during an active malware attack, delayed communication could result in an infected system. Another thing that will happen with more savvy users or support staff when performance is slow is that they can turn off their anti-virus software, which obviously then places the system at risk of being infected. Anti-malware solutions that do not incorporate a policy-based approach to defining client configurations are particularly at risk for these user behaviors.

Scheduled Scan

The purpose of scheduled scans is to proactively evaluate all the files on a computer—some of which may be dormant—to detect viruses, malware, spyware, and adware. This approach is as critical as a real-time scan because it can find and stop the propagation of a threat before it is opened and given a chance to execute. Scheduled scans also have a negative impact on performance if the anti-virus client is slow or demands a significant amount of resources while the scan is being performed. Disk I/O is one resource that is heavily impacted—in addition to significant processing and RAM consumption. Slow clients will take longer to perform the scan and, during this time, the overall system performance will be slower as the client examines each file. This can be of particular concern with today's large hard drives as well as the amount and type of data being stored (for example, virtual machines, email archives, images, documents, and spreadsheets).

Many organizations attempt to alleviate this performance impact by scheduling scans during off hours; however, this is only a band-aid approach to what is really a core software problem. The off-hour approach may work fine for desktops that never move and can be left on overnight. Yet more and more companies are moving to a mobile workforce—they are replacing desktops with laptops. Day-to-day use with laptops is very different than with desktops because laptops tend to be powered on only when they are being used. With laptops, scheduled scans can run while the user is trying to use the system, thereby making it critical that the anti-malware software is lean and efficient while performing scheduled scans.

When such is not the case, performance is impacted by scheduled scans. To resolve user performance complaints, the decision is often made not to perform scheduled scans and rely solely on real-time scanning. In some cases, users may also stop a scheduled scan in order to restore system performance. These actions eliminate an important method of virus and malware detection, which puts the system and infrastructure at risk of infection.

Client Reporting

For enterprises, administrators rely on the communication between the anti-malware clients and the servers that manage them. Clients are configured to communicate with management servers in a bi-directional manner for several reasons. The first is to enable the rapid distribution of malware signatures and client updates. The second aspect is the client reporting its status back to the server. Client status reporting is one of the most important aspects of limiting the impact of a malware infection; it relies on the ability to collect and analyze data from every client. The type of data that is needed includes the health of the client, which is determined by the version and/or date of the virus definitions on the client. Clients will also report back any infections that are found and the actions performed. This reporting enables administrators to assess the overall threat to their infrastructure and take appropriate action. Clients that are overloaded or slow to report their status limit an administrator's ability to properly manage and protect the infrastructure.

Anti-Malware Shouldn't Be Slow

To address the ever-increasing sophistication of threats, software vendors have created more sophisticated anti-virus and anti-malware solutions, but the cost of this development is decreased performance both from the system perspective and from the anti-virus software itself. Many factors contribute to this poor system performance, such as code bloat, the requirements of OS support, and products that were bolted together over time to provide a suite of solutions. Only through the use of new and specifically-targeted solutions for anti-malware will today's IT environments ensure the highest levels of protection while maintaining good performance in their computing infrastructures.