White Paper

Endpoint Security: It's Not Just Black or White

By
Christopher Beier, Senior Product Marketing Manager, McAfee
Barbara G. Kay, Principal Analyst, Secure By Design Group

Table of Contents

Executive Summary

IT staff are tired of patching, salvaging corrupted systems, nursing legacy systems, and trying to impose security standards on reluctant users and field systems. They demand more practical ways to control and secure the software running on their endpoints. A concrete example of this change is IT reasserting control over management of corporate-owned devices by using whitelisting to push back on the free-for-all of Bring Your Own Device (BYOD) programs.

Endpoint security strategies available today can be categorized as:

- *Block the known bad*—Comprehensively thwart attempts to infect or compromise applications and their execution environment: across every attack vector, throughout the stack (blacklisting based on signatures, behavior, and reputation).
- *Trust the known good*—Decide what software or applications you want to permit, validate them, and make them manageable (whitelisting with dynamic updating).

Each of these approaches has merit. Blacklisting continues to offer important protection against opportunistic, high-volume attacks. Signatures are just the first line of defense, backed up by more advanced blacklisting that applies behavioral protection; detection below the operating system (OS); and file, message, and IP reputations from the cloud. All of these options make some sort of decision about known risks.

Used primarily at the application level today, but also in techniques like signed drivers, whitelisting offers ways to prevent targeted attacks by reducing the attack surface. Dynamic whitelisting adds manageability and automation so that these techniques—and the IT administrators who use them daily—can survive real-world environments.

Together, these strategies represent a modern incarnation of defense-in-depth. Instead of layers of antivirus blacklists from different research labs, modern defenses use layers of blacklisting and whitelisting, enhanced with real-time analysis and reputation data, to protect the endpoint computing stack. This meshed defense can detect and block both malicious payloads and malicious actions that compromise a system and its network infrastructure. Without a persistent home in the endpoint, attackers can't exfiltrate data or use an infected endpoint as a gateway to other systems.

Secure the Stack

The endpoint attack surface is no longer just the OS and its vulnerabilities—nor just the browser or email. The attack surface includes the complete software computing stack: the BIOS, OS, applications, data, and cloud. As we look at securing the endpoint, we need to consider the different forms of attack at each layer. This view helps us see where different defensive approaches are most effective.

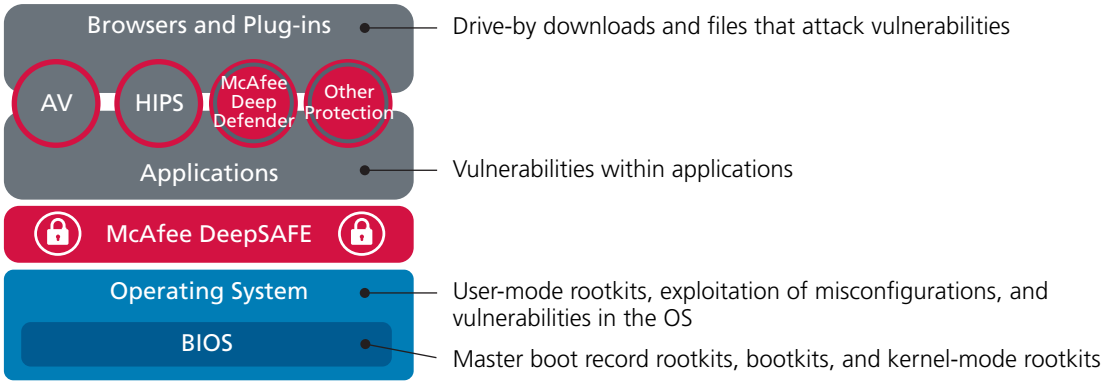

Figure 1. Securing the stack.

"One strain of malware targets a computer's master boot record (MBR)—an area that performs key startup operations. Compromising the MBR offers an attacker a wide variety of control, persistence, and deep penetration. These attacks, including mebroot, Tidserv, Cidox, and Shamoon, have rapidly increased their numbers and have set a new record high for two quarters running."

—*McAfee® Labs Threats Report*, First Quarter 2013.

The primary security defense today is blacklisting. It's necessary now and will remain a key defense for the foreseeable future, since malicious developers and their automation technologies have created a "long tail" for known attacks. Like tuberculosis in the physical world, an old virus will still crop up years after it has been presumed cured with a .DAT signature.

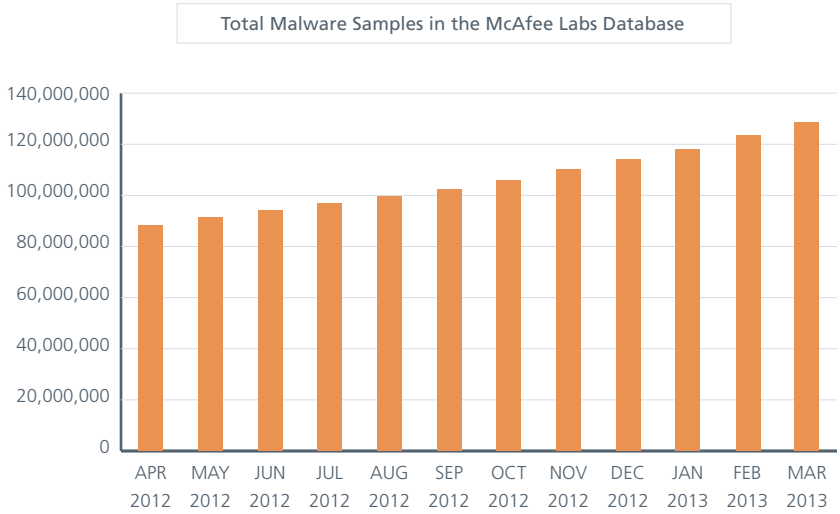

Figure 2. The malware zoo may reach 175 million samples by the end of 2013, up from 75 million samples at the end of 2011. (Source: *McAfee Labs Threats Report,* First Quarter 2013)

However, the relentlessly swelling volume of malware makes relying solely on blacklisting a risky strategy. Most security experts recommend that IT should presume some malware will get through. And then what?

If keeping up with all of the "known bad" stuff is too difficult, what about creating a moat around the "known good" through whitelisting? Whitelisting approaches have matured in recent years and can now be categorized in three different strengths based on sophistication and capability:

- Basic or manual application whitelisting, available in most enterprise endpoint security suites, is beneficial, but entry level. It requires ongoing input and maintenance of file names or wholesale fingerprinting of entire systems. This hands-on approach creates a management headache whenever any new application or patch is introduced. And users resent the constraints on their work environment.
- Dynamic whitelisting uses a "trust but verify" model to ensure the reputation of a new application or updated application is understood.
- Graylisting provides a temporary right to use until an application is assessed and either approved or blocked by IT. This option requires extra defenses on the endpoint (such as host intrusion prevention) to ensure that the graylisted application doesn't compromise the system.

Today's diverse endpoint systems require a well-designed blend of blacklisting and whitelisting technologies to achieve effective security and optimum productivity. No single solution provides the required techniques to both eradicate the growing threat of the malware landscape and keep up with the technological innovations and user software installations. A comprehensive approach to security that takes advantage of both blacklisting and whitelisting, incorporates real-time reputation, and works throughout the entire software stack—from the boot sequence up through the application plug-ins—provides the solution to this dilemma.

"IT managers worry about the increased sophistication of malware as part of low-volume, high-complexity targeted attacks. One increasingly popular technique signs malicious code with the certificates of good applications to evade traditional detection methods."

—McAfee Labs Threats Report First Quarter 2013

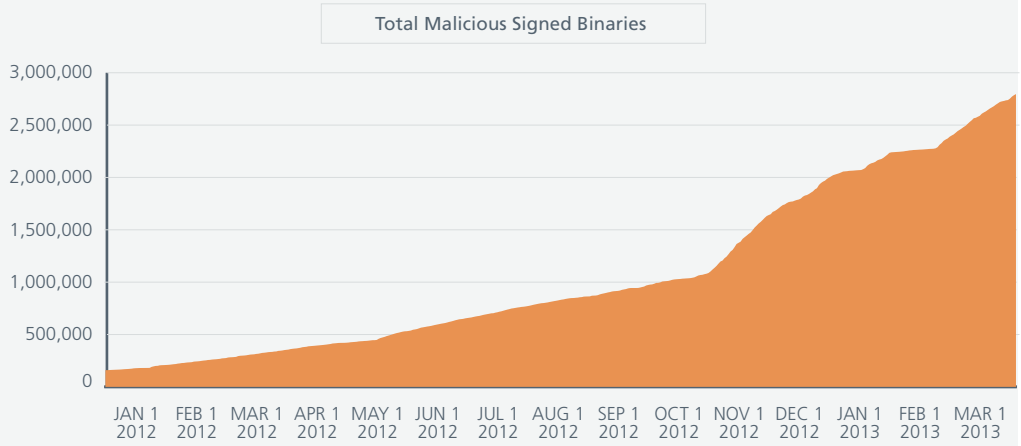

Basic Blacklisting: File Signatures

Basic blacklisting is familiar to those using antivirus and intrusion detection. Each suspicious code sample spawns release of a protective file known as a signature, which tells the security product to block or blacklist that image if it sees it. While potent, blacklisting alone is not powerful enough for all of today's malware. Despite the responsive nature of antivirus vendors, the operational constraints of most enterprises can result in it taking many hours or even days to roll out updated signature-based protection. Where threats travel the globe in seconds, the traditional delayed approach offers little or no protection against new malware. Bad actors never throw out a tool, so eliminating the known bad remains an essential part of endpoint security.

Blacklisting reduces the "known" noise in the environment.

Blacklisting benefits:

- Updates to virus lists are automatic and do not require time-consuming maintenance.
- It allows known malware to be identified and eliminated, reducing infections, storage costs, risks of propagation, and network bandwidth.

Blacklisting drawbacks:

- There's a need for regularly scheduled and on-demand real-time updates to virus and spyware definitions, which increase the load on hardware and network bandwidth.
- The solution requires viruses or spyware to be identified and added to the blacklist, leaving workstations and networks vulnerable to a zero-day attack.
- The scanning of all incoming and outgoing IP traffic results in slower workstation performance.
- Remote users may miss updates when they are offline or not active.
- It may only protect one layer of the software stack.

Although some have called these solutions dead, there is a measurable negative effect when signature antivirus is removed from the endpoint. Specifically, in its *Security Intelligence Report (SIR v14)*, Microsoft found that computers that had zero anti-malware protection were 5.5 times more likely, on average, to be infected with malicious code. The more recent the OS, the more important the malware protection. Compared to the same systems with active anti-malware protection:

- Older unprotected Microsoft Windows XP systems were 3.5 times more likely to be infected.
- Unprotected Windows 7 Service Pack 1 systems were 9.5 times more likely to be infected.
- Brand new unprotected Windows 8 systems were 14 times more likely to be infected than the same system with anti-malware software.[1]

Advanced Blacklisting: Beyond Signatures

Advanced blacklisting uses additional data points beyond virus signatures as well as real-time analytics to make blacklisting decisions. IP addresses of known botnets, bad URLs, file reputation, behavior, and other attributes are used heuristically to determine the disposition of an application or file. While advanced blacklisting provides additional views and determining factors, these solutions remain limited by what is "known bad." For example, a newly published website will not have a URL reputation but could still contain very malicious content.

Intrusion prevention systems (IPS) can use blacklisting to block files attempting to misbehave, such as malware that tries to change a registry. An IPS might also block attempts by a system to contact a known malicious IP address, such as a botnet control. But if the bot control center is a newly compromised and unknown host, a connection to that address would be allowed.

Advanced blacklisting benefits:

- It may potentially intercept zero-day attacks.
- It provides another layer of protection because it does not rely exclusively on file reputation.
- Updates to behavior rules are minor and less frequent than virus signatures.

Advanced blacklisting drawbacks:

- Legitimate computer usage and applications may also fall into the suspicious pattern, resulting in "false positives" and access denial.

Basic Whitelisting: Manual Lists

Most commonly, whitelisting helps to manage the overwhelming problem of how to determine a "good" application. It can also be used to constrain drivers, web downloads, or user changes to the operating system. The general concept is quite simple. Instead of attempting to block malicious files and activity, application rules only permit known good files. The model is changed from a "default allow" to a "default deny" for all files and executables.

Basic whitelisting limits content download or execution to approved content. For application whitelisting, an administrator sets up a list of applications to either always permit or always block, essentially defining rules. Any time a new application needs to be added or blocked, the administrator changes a rule.

This model works reasonably well for uniform groups of fixed-function devices, such as embedded systems or ATMs. It can work well for limited-function use cases, such as a call center or manufacturing floor kiosk where only a few drivers, operating systems, and applications are permitted. These environments tightly restrict changes to ensure availability. Uptime is much more important than user convenience.

However, manual whitelists fall short when you consider the usage models of full-featured endpoints and knowledge workers. Users want to install their own applications without waiting days for IT—and IT can't keep up with the range and volume of requested changes.

Some manual whitelisting implementations also have coverage limitations. Basic application whitelisting solutions only cover executables (EXEs) and dynamic-link libraries (DLLs). In addition, approved applications may also have newly discovered vulnerabilities that could be exploited, giving the attacker free rein if applications are not patched. And what about Java, ActiveX controls, scripts, and specialty code such as drivers and kernel components? Attackers go after whatever part of the software stack you leave unprotected, and basic application rules don't apply control to these routinely compromised software components.

Manual whitelisting benefits:

- No virus or spyware definition updates are needed; systems are always protected from known virus attacks.
- Constant scanning of incoming and outgoing IP traffic is not necessary; therefore, there is no decrease in performance.
- No unauthorized executable files, such as a chat program, P2P, spyware, or Trojan will run; staff productivity increases and downtime decreases.
- No illegal or unlicensed software can be executed on endpoints.
- Hardware and support budget costs are reduced because of a decrease in re-imaging PCs on a regular basis; organizations can re-channel resources to other activities.

Manual whitelisting drawbacks:

- The total cost of ownership (TCO) is impractical: no single whitelist will meet the needs and configurations of all users, applications, devices, and servers.
- It's too labor-intensive to scale for each new or updated application.
- Responding to legitimate exception requests requires an approval process and a manual change to the database of authorized software.
- As approved whitelists proliferate, they compound the resource burden until it becomes unsustainable.
- Approved applications may be subject to attack or manipulation.
- It may only protect one layer or a part of a layer of the stack.

Advanced Whitelisting: Tighter and More Manageable Control

To overcome the administrative overhead and "one size doesn't fit all" limitations of basic manual whitelisting, you can automate the process of creating an approved application list, then let the software maintain that approved list using a trusted, delegated update model. This approach is flexible, in that it can continually adapt to the patches, updates, and user requests of a typical business environment.

Active whitelisting can use a "trust but verify" model to ensure the reputation of a new or updated application is validated and approved before the update is installed. Each user starts with a trusted environment, usually defined based on roles tracked within Microsoft Active Directory: such as administrator, accountant, sales representative, or executive.

Sometimes users need access to applications that are not yet corporate-approved. For example, a user may need WebEx software to view a webinar. In these instances, users can install new software without waiting for an IT approval. IT can inspect these self-approvals (conditional applications) and create enterprise-wide policies to either ban the application or permit it on all systems. The self-approval options should be used on a limited basis to ensure timely and manageable audits and dispositions by administrators. Additional countermeasures, such as antivirus and host intrusion prevention (blacklisting), should be used to protect against malicious actions by the application until approved or removed.

For extra confidence, an "approved configuration" checklist can live in a separate location, so the system that is changing can double-check itself against a master list. A checksum function can verify that the code being installed on the user system is the same code that was housed on the server. Signed certificates validate changes coming from publishers, such as Patch Tuesday patches for Microsoft Office coming from Microsoft.

This dynamic control over the configuration ensures the health of the application while it is stored on the disk.

Advanced whitelisting benefits:

- No virus or spyware definition updates are needed.
- No constant scanning of incoming content is required, eliminating the user impact of scanning.
- Users can install unlisted but good applications, saving IT resources and user productivity (graylisting).
- Hardware and support budget costs are reduced from a decrease in re-imaging PCs on a regular basis; therefore, organizations can re-channel resources to other activities.
- Whitelists are automatically maintained, saving IT resources.

Advanced whitelisting drawbacks:

- Difficulty in managing and removing unapproved or unwanted applications.

Optimized Endpoint Security

Since each of these approaches has strengths and weaknesses, they work best together. Seamless integration of a range of blacklisting and whitelisting technologies and real-time analytics secures your systems and data against sophisticated malware and diverse compliance and user requirements. This combination of controls is ideal for modern enterprises whose demanding users expect freedom—as well as complete protection from the risks their mobility and flexibility entail.

McAfee Complete Endpoint Protection—Enterprise suites integrate multiple layers of anti-malware, and each layer includes one or more techniques for detecting and blocking threats, permitting a robust and complete anti-malware strategy:

Trust the good

Decide what applications you want to permit, validate them, and make them manageable.

- McAfee Application Control for Desktops ensures that only trusted applications run on endpoints, improving protection and application visibility and reducing patching cycles on standardized desktop operating systems.
- By verifying application reputation through integration with McAfee Global Threat Intelligence™, McAfee Application Control simplifies decisions about what is "good."
- McAfee Application Control for Desktops provides automatic acceptance of new software through your authorized processes and the option of user assisted self-approval for temporary whitelisting until IT can audit and verify against security policy.

Block the bad

Comprehensively thwart attempts to infect or compromise known applications and their operational environment—across every attack vector, throughout the stack.

- McAfee Host Intrusion Prevention shields your trusted applications by protecting the entire runtime environment against known vulnerabilities and specific attacks on the host. It includes special protections against common server and database attacks, such as directory traversal, DDoS, and SQL injection, and provides buffer overflow protection by monitoring applications and protecting critical memory address space.
- McAfee antivirus, antispam, and anti-spyware solutions will detect, clean, and kill known malware before it can attempt to install, keeping this code away from your applications.
- McAfee Deep Defender is a first-of-its-kind, hardware-enhanced endpoint security technology, enabled by McAfee DeepSAFE™ technology. It operates as the computer boots up and protects before the operating system is active so it can detect, block, and remediate advanced, hidden attacks. It removes low-level hidden threats that often evade traditional OS-based protection, lowering reimaging and remediation costs and enhancing overall security.

Figure 3. McAfee Application Control permits a variety of secure, automated, dynamic updates to enterprise whitelists.

Cloud-based Threat Intelligence

McAfee Global Threat Intelligence (McAfee GTI™) factors in behavior, reputation, vulnerability data, and McAfee experience to gauge the risk of new, unknown code. It keeps these protections up to date using real-time reputation assessments. McAfee GTI is an exclusive McAfee technology that tracks the reputations of files, messages, and senders in real time using millions of sensors worldwide. McAfee products use this reputation for real-time decision making on risk. For example, McAfee Application Control uses this cloud-based knowledge to determine the reputations of all files in your computing environment, classifying them as good, bad, and unknown. McAfee anti-malware can look up new and suspicious unknown files and block them based on the file's reputation.

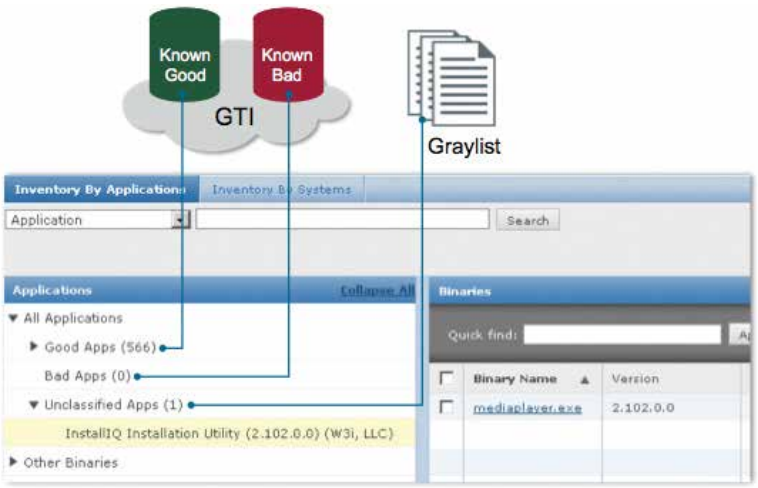

Figure 4. McAfee GTI constantly monitors the reputation of files and senders, allowing automatic blocking of known bad files and graylisting those with no known reputation.

Included in the suites is the integrated security management platform, McAfee® ePolicy Orchestrator® (McAfee ePO™) software, which brings together policy management, software deployment, maintenance, and reporting across these products. With just one management environment and integrated real-time questions, McAfee ePO software has been proven to help organizations reduce time spent managing endpoint security.

Here's how these protections team up to cover the entire computing stack:

- Below the OS, McAfee Deep Defender uses behavior and signatures to prevent stealthy rootkits.
- To protect the OS and application layers, McAfee suites include anti-malware protection in McAfee VirusScan® Enterprise, McAfee Host Intrusion Prevention System, and McAfee SiteAdvisor® Enterprise web filtering and content control. These systems guard endpoints with blacklisting and behavioral protections, enhanced by the McAfee Global Threat Intelligence network.
- At the application layer, McAfee Application Control complements other system security tactics with dynamic whitelisting that reinforces blacklisting and behavioral protections.
- McAfee Application Control also extends coverage to executable files, libraries, drivers, Java applications, ActiveX controls, scripts, and specialty code for greater control over application components.

Learn More

Most businesses have deployed antivirus and a firewall, which traditionally use blacklisting and rules to block known bad activities on the endpoint. You can reduce your attack surface and reduce your vulnerability to both opportunistic and targeted attacks by increasing the range of bad activities that can be blocked and by applying the real-time intelligence collected by a globally connected cloud of sensors. Today, you can also take advantage of new approaches that focus on known good, approved content.

By leveraging integrated solutions that include pretested and centrally managed suites of these capabilities, you can provide a secure endpoint environment for your business without the complexity of managing point products.

McAfee Complete Protection Enterprise unites industry-leading endpoint security and data protection with real-time security management for powerful, complete security that streamlines operations and eases compliance. For more information, visit www.mcafee.com/endpoint.

About McAfee

McAfee, a wholly owned subsidiary of Intel Corporation (NASDAQ: INTC), empowers businesses, the public sector, and home users to safely experience the benefits of the Internet. The company delivers proactive and proven security solutions and services for systems, networks, and mobile devices around the world. With its visionary Security Connected strategy, innovative approach to hardware-enhanced security, and unique global threat intelligence network, McAfee is relentlessly focused on keeping its customers safe. http://www.mcafee.com.

McAfee
An Intel Company

2821 Mission College Boulevard
Santa Clara, CA 95054
888 847 8766
www.mcafee.com

[1] *Microsoft Security Intelligence Report*, Volume 14, July 2012 to December 2012, pages 3-7.

60262wp_ep-security-bw_0713_fnl_ETMG