Technology Blueprint

McAfee
An Intel Company

FIND, FREEZE, AND FIX ADVANCED THREATS

Eliminate the malware behind zero-day and stealthy, evasive threats

SECURITY CONNECTED REFERENCE ARCHITECTURE

LEVEL 1 2 **3** 4 5

Security Connected

The Security Connected framework from McAfee enables integration of multiple products, services, and partnerships for centralized, efficient, and effective risk mitigation. Built on more than two decades of proven security practices, the Security Connected approach helps organizations of all sizes and segments—across all geographies—improve security postures, optimize security for greater cost effectiveness, and align security strategically with business initiatives. The Security Connected Reference Architecture provides a concrete path from ideas to implementation. Use it to adapt the Security Connected concepts to your unique risks, infrastructure, and business objectives. McAfee is relentlessly focused on finding new ways to keep our customers safe.

Eliminate the malware behind zero-day and stealthy, evasive threats

The Situation

Your management is asking for an advanced malware detection and response strategy that is as sophisticated and adaptive as the attackers your organization faces. You've read that sandboxing isn't a panacea. You don't have a stable of forensic investigators. Network traffic loads keep rising. You already have a headache. What's your next step?

Driving Concerns

Several design trends in advanced malware security affect the performance, efficacy, cost, and management complexity of an overall solution.

- **Rapidly maturing technologies**—While signature-based protection is still needed for timely blocking of known malware, security vendors are now transplanting the techniques of manual forensic analysis into automated defenses that can monitor traffic continuously and identify new threats when they appear. Market evolution including vendor consolidation makes it difficult to be sure how best to balance performance and protection.
- **Per-protocol deployment requirements**—Some solutions require a specific, separate analysis appliance for each protocol—file sharing (FTP), web (HTTP), and email (SMTP/IMAP)—in addition to the appliance used for standalone analysis. This model increases both equipment costs and administrative overhead.
- **Host-agent dependency**—Certain solutions delegate the task of separating files and executables from their message streams to client software running on host systems throughout the environment. Files are then forwarded to a central inspection service for analysis. This model offloads an initial increment of processing from the malware platform, but transfers that workload to client systems and adds a new agent to each host image, which can massively increase host management costs. All this handing off also increases the delay before detection and conviction of a malicious file.
- **Cloud-based analysis**—One option is to strip files and executables at the firewall and ship them to a cloud-based service for analysis. This design has the disadvantage of adding a decomposition workload to the firewall, where any impact on throughput performance may affect the entire environment. Some organizations may also be reluctant to allow large volumes of internal content and communications to be routed outside the network for security processing. And a pure cloud analytical model reduces your opportunity to understand the implications of an attack on your environment.
- **Automated detection, but manual response**—Many advanced malware detection solutions are capable of automatically finding advanced threats, but have no automated response mechanism to immediately freeze similar threats in the network, or to quickly identify and fix infected hosts. While the detection itself may alert incident response staff to an attack in progress, great damage will be done by the attack during the time consumed by all of the manual processes and hand offs required to halt and remediate the attack.
- **Dependence on manual forensics**—Some solutions provide sophisticated analytical tools, but remain dependent on security personnel with advanced forensic skills to assess the nature and impact of the malware or the breadth and depth of the actual infiltration. These approaches tend to have high labor costs and higher latency from detection to response than solutions that integrate forensic cycles and related contextual data with incident response processes.

Solution Description

The standard defense-in-depth idea of layers gets tricky when you have a business to support, and network traffic is its lifeblood. Your layers have to balance protection and performance. This is where your blueprint has to blend art and science.

Static and Dynamic Analysis: A Comparison

Two distinct analytical approaches account for much of the progress in detecting previously unknown and well-disguised threats: static code analysis and dynamic evaluation or sandboxing. Both tactics are important tools that should be integrated into a layered defense strategy.

Static analysis parses and examines object code without actually executing it, using emulation and heuristics to extrapolate runtime behavior, uncovering alternate execution paths, and revealing potential attacks, including those that may not execute immediately. Static analysis provides useful visibility into the intent of a file or executable, but it is limited in its ability to understand complex interactions between an executable and other elements in an environment.

Lightweight analysis may also be susceptible to the many types of code obfuscation that can be used to hide an incriminating string or conceal the true intent of a malicious instruction. Emulation, the common form of real-time static analysis, has a short window and cannot fully examine the code. Full static code analysis, which includes unpacking of often-obfuscated code, can perform thorough analysis of every code path. However, this deep analysis is computationally intensive and can take significantly longer than other methods to identify a threat.

Dynamic analysis places suspect code in a safely isolated runtime environment (a sandbox) to observe its behavior. This approach is more likely to identify an elaborately choreographed attack sequence, but it can be fooled by a programmed attack delay or by an executable with the ability to sense its environment, identify a virtual machine or sandbox, and behave innocently under observation. In addition, the available sandboxes may not match all the potential victim environments within an organization. Unsupported operating systems and browsers, for example, can still be exploited. This gap leaves a hole in the analysis: any code that isn't executed in the sandbox isn't analyzed.

In short, while static and dynamic analysis each identify malware attacks that are invisible to signature-based defenses, both technologies have strengths and weaknesses that are largely complementary. Any security product based on one of these technologies alone only provides partial protection and must be supplemented with other controls to provide visibility in the blind spots. A preferred approach is to combine signature-based inspection with static analysis and dynamic sandboxing to provide multilayered threat detection in a single service.

Reputation and the Cloud

The cloud offers data and analytical tools that will increasingly augment on-site systems. Today, reputation—of files, senders, IP addresses—can be very useful in both determining risk of traffic and shutting down active attacks. Cloud analytics can provide extra resources to characterize content. Designing these systems into your threat defenses can boost both real-time detection and incident response.

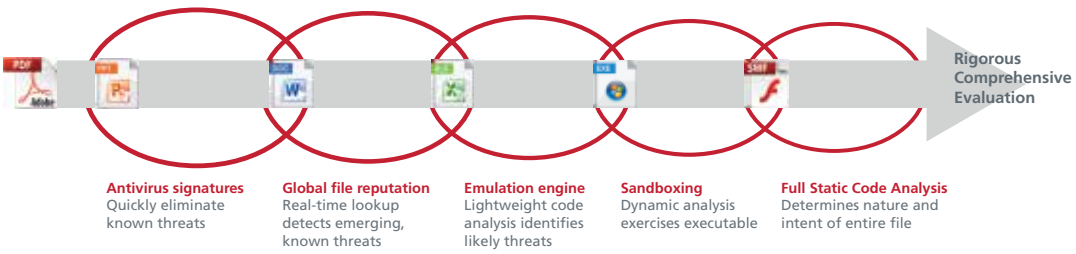

Figure 1. Balanced use of multiple techniques optimizes both protection and performance.

Decision Elements

These factors could influence your architecture:

- Do you prefer the architectural simplicity of a single, shared, network-based solution, or are you prepared to manage a more distributed solution?
- What are your existing endpoint and gateway solutions?
- Do you have a way to feed threat, malware, and countermeasure information to or accept this information from gateway and endpoint defenses?
- Do you have in-house forensic examiners?

Layers to "find" malicious files

Since no individual detection method penetrates all of advanced malware's varieties of camouflage and evasive behaviors, the most reasonable approach is to apply proven technologies as a coordinated system—signature, reputation, lightweight and deep static analysis, and dynamic sandboxing—in a unified, layered, and extensible architecture. To optimize processing efficiency, the inspection engines should be stacked in order of increasing computational intensity, so that samples can be identified with fast, simple processes and can be blocked and discarded from the funnel prior to more CPU-intensive stages.

This integrated service should be deployable in a way that utilizes its capacity efficiently, leverages the services of other security resources, isolates other operational hosts from its workloads, and supports simple, cost-effective capacity scaling. These operational considerations are especially important to large organizations and those with high-value assets that attract cutting-edge attackers.

For example, many of the inspection elements can or are already implemented inline, as gateways at the perimeter and IPS sensors monitoring internal network traffic. These systems operate in real time to block malicious traffic entering and traversing the network and reduce the number of suspect samples that need full static code analysis and dynamic sandboxing. These latter analytics, which take longer than a few micro-seconds, can be performed out of band by a forensic system.

For thorough yet expeditious evaluation, you can use dynamic and full static code analysis engines in parallel. The dynamic sandboxing testing can "detonate" some of the code in virtual victims—replica environments that should be configurable to match your enterprise's asset configurations: operating systems, productivity applications, browsers, and plug-ins such as Adobe PDF readers.

Simultaneously, full static code analysis can explore unexecuted code in the file to see its capabilities and similarities to known malware families. For instance, sophisticated attacks frequently include a variety of evasive tactics to compromise a victim: multiple browser exploits, exploit of zero-day vulnerabilities, and tools for deactivating anti-malware on the victim. Full static code analysis can reveal the scope and nature of the payload in detail, showing what it might do in another host.

This layered approach—high-efficiency "noise" filters followed by thorough, painstaking dissection of the remaining suspect files—minimizes the chance of malicious traffic on your network. This is an operationally efficient and complete environment for identifying advanced malware.

Depending on the forensic appliance, you may have the option of performing some or all of these analyses within the out-of-band system. Whichever design you use—inband, out-of-band, or a combination—to glean intelligence about advanced threats, the "find" effort should be integrated into your "freeze" and "fix" processes. The tighter the integration, the shorter the timeline for data exfiltration, network reconnaissance, or sabotage.

"Freeze" any infection

Once the malware has been confirmed, the freeze phase stops its spread around the network. Validated threat intelligence (in a standard form such as an MD5 or SHA-1 hash identifier) feeds into existing gateway defenses so they can immediately start blocking new infections at all traffic control points.

"Fix" the situation

Finally, while it is obvious that infected hosts should be remediated, your incident response may legitimately take two forms: covert or overt.

Option A: Covert

A covert response may be appropriate if you believe the malware could be part of a persistent attack, since once you quarantine or remediate any infected hosts, you lose visibility into the process activity and volatile memory information they contained. These are very important sources of forensic information and represent a key best practice advance from the standard "dead box" forensics of the past. Killing processes and disconnecting a host will also signal the attackers that you have detected their activities, which can let them cover their tracks or activate an alternate attack sequence.

To understand the attack, you will want to dig into the malware's available and demonstrated behavior, reconstruct the attack sequence, and find related events that could help you see the full scope of the attack. For example, you could use the MD5 hash of the newly found malware to identify and prioritize infected hosts, and then profile the activities of those hosts:

- What processes are running?
- What registry keys have been altered?
- Are there any other suspicious files or configurations on those hosts?

You may also be able to determine details such as the software profile and patch level of the host, active or recent connections, and if any other hosts are communicating with those IP addresses.

In collecting this data from running (and compromised) hosts, your goal is to capture the fullest possible understanding of the attack before revealing to the attacker that you have uncovered his trail.

Option B: Overt

Different hosts usually merit different intervention strategies. The nature of the malware and the asset value of the infected system should drive the response strategy. For example, laptops and end-user workstations may be fine to quarantine or take offline right away, whereas business-critical servers such as an email system must be remediated according to a careful schedule to avoid disrupting the business.

The threat analysis should provide a file identifier that can be used to pinpoint which systems are infected. Other malware information such as registry key changes and software exploited can provide guidance on what to do to repair the system or if a full reimage is necessary.

Technologies Used in the McAfee Solution

McAfee addresses the three key requirements of today's advanced malware problem, the abilities to find, freeze, and fix a previously undocumented attack. Used with McAfee® network security products, the McAfee Advanced Threat Defense forensic appliance provides the full set of analytics required to ***find*** advanced malware. McAfee Advanced Threat Defense then shares the fingerprint of the validated threat locally with other McAfee products. This data sharing can ***freeze*** the threat, preventing further infections within your networked infrastructure.

Finally, McAfee Advanced Threat Defense works with McAfee Real Time software and McAfee managed endpoints to establish the scope of infiltration by seeking out the malware fingerprint throughout the network and revealing important behaviors and attributes of the malware, then providing tools to begin to ***fix*** damaged systems.

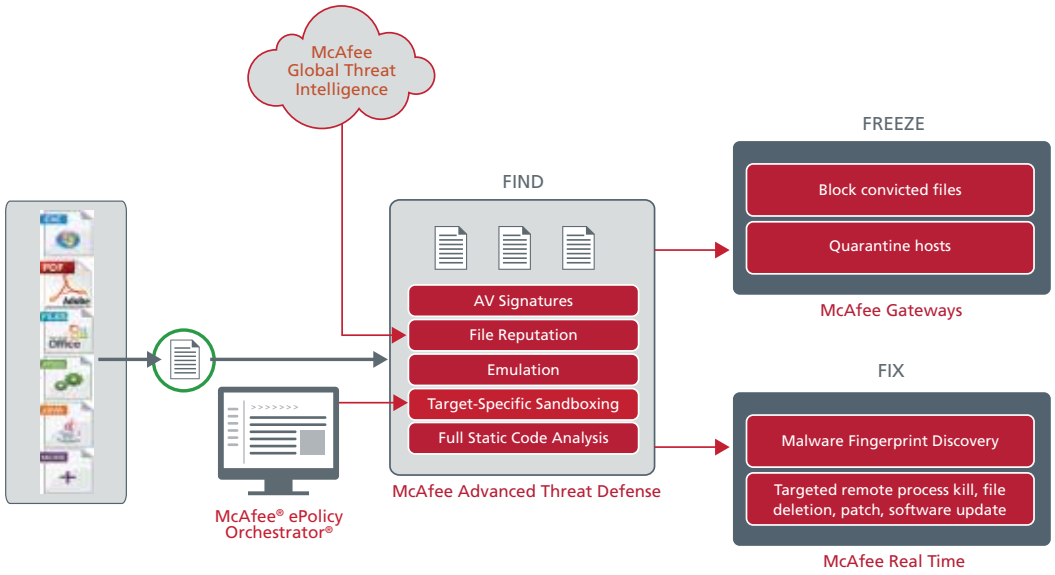

Figure 2. McAfee Advanced Threat Defense can integrate with other McAfee systems to improve the efficiency of detection and response.

Architectural options

McAfee gateway solutions including the McAfee Network Security Platform intrusion prevention system (IPS) and the McAfee Web Gateway can screen network traffic in real time to block malware through signatures, reputation lookups, and lightweight static analysis or emulation. This configuration skims off the bulk of malware without delaying traffic. The gateways can direct remaining suspect files through to McAfee Advanced Threat Defense for dynamic and full static code analysis.

Where you have these real-time, inline defenses in place, you can configure McAfee Advanced Threat Defense to perform only the more detailed and resource-intensive sandboxing and static analysis. As it detects malware, the Advanced Threat Defense can notify the McAfee gateways to block subsequent copies of the file.

Alternatively, the McAfee Advanced Threat Defense solution can perform all of these forms of analysis as a full-featured forensic system. Its detailed reports can feed manual efforts.

Finding advanced threats

McAfee Advanced Threat Defense is a multilayered malware detection solution that stacks an extensible series of inspection engines and analytical capabilities in a down-select, multitier sequence of increasing intensity.

- **Signature-based detection**—Detects viruses, worms, spyware, bots, Trojans, buffer overflows, and blended attacks. Its comprehensive knowledgebase is created and maintained by McAfee Labs, and currently includes close to 150 million signatures.
- **Reputation-based detection**—Looks up the reputation of files using the McAfee Global Threat Intelligence network to detect newly emerging threats.
- **Real-time static analysis and emulation**—Provides real-time static analysis and emulation to quickly find malware and zero-day threats not identified with signature-based techniques or reputation.
- **Full static code analysis**—Reverse engineers file code to assess all its attributes and instruction sets, and fully analyze the source code without execution. Comprehensive unpacking capabilities open all types of packed and compressed files to enable complete analysis and malware classification, helping organizations better understand the specific malware they are dealing with and the impact it has on their organization.

- **Dynamic sandbox analysis**—Executes the file code in a virtual run-time environment and observes the resulting behavior. Virtual run-time environments are configured to match the target host based on queries to McAfee ePO software. Analyzing execution behavior under the exact conditions of the intended host produces accurate results, quickly and efficiently. Without understanding the target environment, actual file behavior may not be accurately determined and execution may need to be run multiple times under varying conditions in an attempt to determine the actual file behavior on the intended target, lengthening analysis time. Since many advanced attacks are designed to evade detection while sandbox execution is attempted, McAfee Advanced Threat Defense includes innovative techniques to ensure code execution during dynamic analysis.

These techniques work together in coordination to efficiently identify many types of known and unknown malware. The full static and dynamic analysis reveal the obfuscated and advanced malware not positively identified through lighter-weight analysis engines.

The combination of full static code and dynamic analysis provides a complete assessment of threat potential that neither technique can achieve alone. Dynamic analysis identifies a single observed execution path, while static analysis reveals input-dependent behaviors and delayed or hidden execution paths—often not executed in a dynamic environment—through unpacking and full code analysis.

Freezing advanced threats

Unlike other malware detection solutions that can find stealth attacks but not intervene to stop them, McAfee Advanced Threat Defense can be used with McAfee network gateways to initiate an immediate and comprehensive response whenever a threat is identified. If the attack is new, McAfee Advanced Threat Defense generates and makes available a hash and other metadata for network gateways such as McAfee Network Security Platform and McAfee Web Gateway to enable instant protection. The network IPS can identify the users and applications involved in the download and quarantine the infected hosts or applications to prevent further infection and block communication with botnet command and control channels.

Fixing advanced threats

Finally, to begin remediation, McAfee Real Time can be used to help administrators discover which machines to take offline, understand the extent of the compromise, and gauge the need for a reimage. When the right remediation has been chosen, you can terminate processes and remove files on all (or a subset of) affected computers by simply clicking a button to target the action to those machines. You can even set up a blacklist to monitor for reinstallation of a file or application and automatically delete it when detected.

McAfee Real Time does on-demand checks of current status, rather than querying historical data, so it also gives you visibility into other activities on a compromised host. It can collect data from local log files and application, system, or security event logs to find particular strings, IDs, or regular expressions. This search option helps administrators determine the prevalence and location of specific errors in real time to decide if a newly discovered problem is common. McAfee Real Time can also enforce updates to DAT files, install or wake up security software to increase protections, or patch existing software to reduce the available attack surface on vulnerable hosts.

Optional Integrations

Since advanced malware is often just one part of a persistent attack, McAfee also facilitates the process of understanding the malware's role in the attack's "kill chain." Data collected by McAfee products can feed into the McAfee security information and event management (SIEM) system, called McAfee Enterprise Security Manager, where it is normalized and correlated against other endpoint, network, and third-party data to bring visibility and analytical tools to unfolding events.

Impact of the Solution

McAfee brings together the diverse security technologies that have evolved to meet the expanding capabilities and increasing stealth of advanced malware attacks. By combining multiple inspection engines that apply signature- and reputation-based inspection, real-time emulation, full static code analysis, and dynamic sandboxing in a single, easily deployed appliance, McAfee Advanced Threat Defense achieves highly effective detection of malware-based attacks. More importantly, it not only detects these attacks but initiates an immediate and comprehensive response that blocks future incidents and helps administrators and forensic examiners characterize the infection and remediate infected hosts.

Furthermore, because McAfee Advanced Threat Defense is tightly integrated with other McAfee security solutions, and deploys as a proxy service that supports multiple security controls in the network and at the perimeter, its protections are cost-effective, scalable, and manageable.

Additional Resources

www.mcafee.com/advancedthreatdefense

www.mcafee.com/ctp

www.mcafee.com/epo

www.mcafee.com/webgateway

www.mcafee.com/nsp

For more information about the Security Connected Reference Architecture, visit:
www.mcafee.com/securityconnected

About the Author

Michael Lawson, a systems engineer with the McAfee Advanced Technologies Group, has over 15 years of experience helping enterprises to understand and to cope effectively with the latest technologies and emerging threats. In his role at McAfee, Michael works closely with a wide range of organizations to understand how the threat landscape is evolving, and how to leverage Security Connected solutions to counter them. Michael served in The United States Navy for 10 years as an Information Technology Specialist and is a Certified Security Compliance Specialist.

McAfee
An Intel Company

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

60522bp_find-freeze-fix-L3_0913_fnl_ETMG