



Beyond the Sandbox: Strengthen Your Edge-to-Endpoint Security

What You Will Learn

Over the years we've all heard claims of simple, seemingly magical solutions to solve security problems, including the use of sandboxing technology alone to fight advanced malware and targeted threats.

This paper explores:

- Where sandboxing technology stands today
- Why it fails to meet the needs of organizations
- What's needed for effective malware analysis

Overview

Online threats are becoming increasingly sophisticated and elusive. Cybercriminals are launching attacks through a variety of vectors, including tools that users trust or view as benign. And targeted attacks are on the rise, creating a persistent, hidden presence within an organization and from which to execute their mission. Studies by Cisco find that 75 percent of all attacks take only minutes to begin data exfiltration but take much longer to detect. More than 50 percent of all attacks manage to persist without detection for months or even years before they are discovered. And even after they're discovered, several weeks can pass before full containment and remediation are achieved. During this time hackers are free to wreak havoc.

Why are data breaches taking so long to detect? Many organizations struggle with detection and incident response skills, processes, and technologies. Identifying the source and scope of an attack typically requires manual analysis of the breach. Adding more skilled security staff isn't a viable solution given a worldwide shortage of security professionals estimated at one million people according to the Cisco 2014 Annual Security Report.

To help address attacks that have infiltrated the network, dynamic malware analysis or 'sandboxing,' tools were developed to analyze suspicious malware samples in a safe environment. As the technology has matured we now see three ways that organizations typically deploy sandboxing solutions:

- A stand-alone solution that generally sits a network TAP/span port, designed to feed itself samples for analysis without dependency on other security products.

- A distributed feeding-sensor approach, with firewalls, an intrusion prevention system (IPS), or unified threat management (UTM) solutions having integrated sandboxing capabilities
- Built into secure content gateways, such as web or email gateways.

While each deployment option has its own set of pros and cons, traditional sandboxing technologies generally work in the same way and face similar limitations: they create a flood of alerts, can be evaded by environmentally-aware advanced malware, require additional manual effort to retrieve the reports and expertise to make sense of the analysis; and offer limited remediation capabilities. It's time for traditional sandboxing technologies to be improved upon with automation, context, and edge to endpoint deployment options.

A Closer Look at Today's Sandboxing Options

Over the years we've all heard claims of 'silver bullet' solutions to solve security problems including the use of sandboxing technology alone to fight advanced malware and targeted threats. Not surprisingly, a host of security vendors rushed into the sandboxing race and created a new market segment centered on detecting the more sophisticated attacks designed to evade traditional defenses. The goal of sandboxing is to provide a safe environment to improve the detection rates of threats that penetrate the network and to help speed the analysis to improve response and remediation, preventing serious damage. In reality, sandboxing is not always successful, and a number of breaches that have made headlines prove this. Why the disconnect?

- According to the Cisco 2015 Midyear Security Report, malware authors are increasingly using methods such as macros, BIOS names, and filename tricks to detect sandboxes and virtual machines. If the malware detects the presence of a sandbox it will take evasive action to dodge those defenses. Sandboxing technologies are guilty of generating a lot of alerts that are difficult to prioritize and can delay the response to critical threats.
- Not all approaches are the same. Some can be quite expensive and complex to scale. Moreover, their coverage of target endpoint operating systems and files vary.

- Sandboxing produces good raw data, but it is more of a research tool. It typically lacks user-friendly reporting capabilities. The data requires expert analysis and manual correlation if you are to gather details about the threat and its criticality.
- According to Ponemon Research, 41 percent of organizations do not have automated tools to capture intelligence and evaluate the true threat posed by malware.

Given this reality, security experts generally agree that sandboxing is not a silver bullet. In fact, some IT security vendors are realizing it is best to build sandboxing into their other existing security tools as an additional layer of analysis and detection after filtering by existing anti-malware, IPSs, secure web gateways, or firewalls. Some vendors offer sandboxing as just one of multiple methods in their advanced anti-malware offerings to determine whether an unknown file is malicious. Other vendors focus exclusively on sandboxing technology, offering customers yet another disparate product that won't, and can't, interoperate with already bloated volumes of IT security tools.

Whatever the approach, traditional sandboxing is a means to an end, but it doesn't provide adequate content in a consumable, accessible way that allows security staff to take swift action. Sandboxing shouldn't create more work but reduce it. To truly be effective, malware analysis must empower tier-1 analysts to deal with more threats themselves without having to escalate to tier-3 analysts. This requires an integrated approach that increases confidence through verification of data and offers automation to accelerate response.

Why the Current Approach Is Failing

The industrialization of hacking has spawned a new era of professional, entrepreneurial, and resourceful cybercriminals. Compensated to break in to specific organizations and complete a specific mission, they are persistent in their efforts and launch multifaceted attacks using multiple attack vectors. They choose the path of least resistance, but if thwarted they remain persistent and will continue to innovate to accomplish their mission.

Traditional sandboxing technology isn't designed for today's dynamic threat landscape. Instead, it typically operates in the following way:

- Detonates a file in a virtual environment that is easily detected by environment-aware malware
- Generates black or white responses and simply adds the threat to a blacklist or whitelist (that is, it reports whether a file is good or bad but does not provide context about why, or determine whether a suspicious file should continue to be tracked and monitored)
- Does not analyze files it has already seen
- Becomes easily defeated by malware that sleeps for extended periods of time before further execution, or malware that fingerprints the host and determines whether to run or not based on host- or network-level characteristics

Sandboxing technology also contributes to the complexity and fragmentation that IT organizations face today. Many organizations security departments are caught in a cycle of layering on the latest security tool, relying on as many as 40 to 60 disparate security solutions. Traditional sandboxing solutions add to the challenge by creating gaps in visibility and protection based on their design:

- Most sandboxes have self-serving architectures; as standalone products or only integrated as a feature on a company's existing products (such as a firewall) they offer little value to the broader security ecosystem.
- Some malware sandboxing approaches can be costly and complex to scale:
 - In order to provide comprehensive coverage, many vendors require that their appliances be deployed at every ingress point.
 - Because each appliance is threat-vector specific, it also requires the deployment of multiple appliances geared toward web, email, network, and endpoint. Customers essentially create an overlay network, which is very expensive to deploy and difficult to manage.
- Sandboxing technologies produce a limited amount of data, typically basic threat feeds and lists of bad IPs. They don't provide visibility into what is happening globally, so they can't identify trends and provide an early warning. And because they aren't integrated, they can't correlate local intelligence or an analysis of a company's infrastructure, hampering their ability to determine the scope of an attack and its criticality to the organization.


It's Time for a New Era in Malware Analysis

To overcome these increasingly detrimental limitations, what's needed is a more robust approach to malware analysis that takes the best of what traditional approaches offer and raises the bar to help organizations fight well-funded attackers that get better at what they do every day.

An Optimized Workflow

The workflow would look something like this, with a defense that:

Figure 1. Malware Analysis Workflow



1. Automates submission from multiple security monitoring points, including your network and endpoints.
2. Detonates suspicious files in a protected virtual or emulated environment to discern the file's true intent
3. Leverages an array of both static and dynamic analysis techniques and post-processing techniques to produce the highest possible fidelity of data.
4. Integrates into the existing security ecosystem with APIs to span a wide array of attack vectors, to deliver automation of response and a threat-centric approach to security
5. Easy retrieval, and consumption of data by other security monitoring and analysis tools
6. Balances fidelity and scalability of data, supporting an increasing number of analysis techniques, samples, and data without hindering processing time
7. Delivers analysis in an easily consumable way with context to prioritize response
8. Enables the correlation of current threats with historical data to develop more effective response plans and capabilities.

Context Is Critical

The value of context cannot be overstated. At a time when organizations are under siege by attackers, context is critical to understand where the real threats are and to accelerate response. Context is essential to helping you home in on malware that uses evasion techniques to bypass traditional sandboxes and can persist in an organization for months or even years. Traditional sandboxing isn't designed to take into account everything happening across the organization's infrastructure or in the outside world. It provides a sheltered place to test and analyze a suspicious file, which in and of itself does provide value. However, the addition of context expands the security team's ability to see more and take faster, more decisive action.

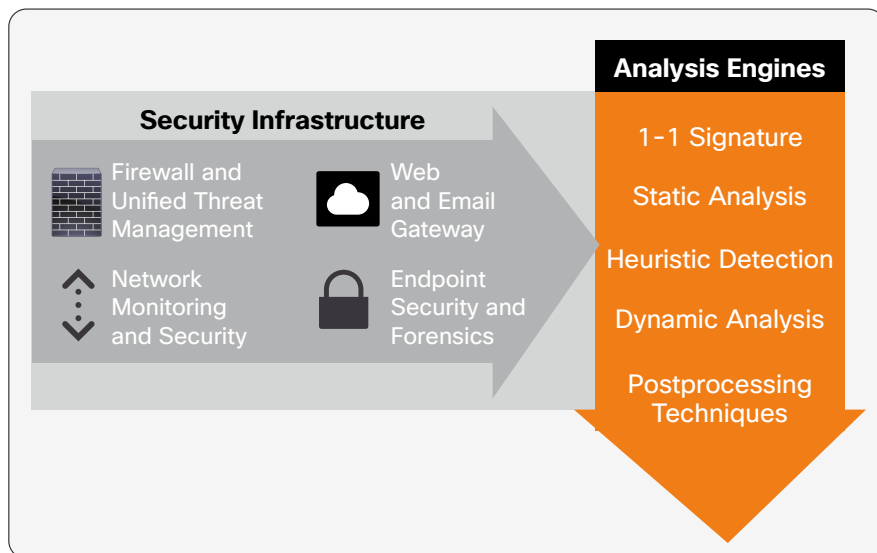
Malware analysis that is context-rich:

- Analyzes a file's code structure and runs it in a virtual environment
- Eliminates instrumentation within the virtual machine detectable to environment-aware malware
- Analyzes every sample, whether the checksum is known or not, allowing for the tracking of changes over time
- Provides contextual information around the sample based on region or similarity to other malware
- Delivers a threat score, a reflection of maliciousness, to determine severity relative to the specific environment
- Combines global crowd sourced intelligence, behavioral indicators of compromise, threat intelligence feeds, and other enrichment to determine whether a sample is malicious, suspicious, or benign
- Addresses threats holistically, across the full attack continuum—before, during, and after an attack

A Fourth Deployment Option: From Edge to Endpoint

An optimized workflow and context-rich malware analysis hinge on an integrated, threat-centric approach to security. In this scenario, malware analysis is an integrated component across the security architecture. The solution is deployed from the firewall to highly secure email and web gateways, to network security solutions, and endpoint security solutions. It conducts multiple checks as the file travels to its destination. It can be deployed either as a software-as-a-service option, as a licensed software option that can be turned on within an existing security appliance, or as a standalone malware analysis appliance. (See Figure 2.)

Figure 2. Edge-to-Endpoint Malware Analysis



At each point along the way, beginning with the firewall and ending with the endpoint, a file is checked and either identified as malicious and blocked, known safe and delivered to the end user, or unknown and passed through to the next checkpoint. If the file is still identified as unknown once it reaches the endpoint, it is automatically submitted for malware analysis by multiple engines. In addition, the capability to check behavioral indicators using global and historical data runs in the background. You are advised if a threat that was previously deemed safe is now identified as malicious.

A verdict is returned with a threat score to help determine the threat's severity and prioritize action. If it's identified as known bad, the file is automatically blocked and blacklisted. If the file passed through all previous checks and is later identified as malicious, an integrated, threat-centric approach to security lets you see the file's trajectory, quarantine any infected devices, and execute automated or hands-on remediation before reattaching the device to the network. This capability is called retrospective security and is critical to reducing time to detection (TTD) and accelerating remediation.

Malware Analysis in Action at the Center for Internet Security

Home to the Multi-State Information Sharing and Analysis Center (MS-ISAC), the Center for Internet Security (CIS) is a nonprofit cybersecurity firm that provides 24/7 cybersecurity services to its 19,000 members that include state, local, tribal, and territorial governments. With the increase in attacks from nation-state actors, CIS needed a solution that could help them provide automated malware analysis services in mass scale with the context to understand if a sample was malicious, suspicious, or benign—and why. CIS members needed to automatically analyze malware in a trusted environment without submitting samples in a public domain where attackers are lurking and can change their methods if they discover an investigation is under way.

The scalable infrastructure had to be able to handle malware submissions from thousands of entities across the nation. Using the feature-rich API, CIS built a customized portal to the Cisco AMP Threat Grid solution with robust access-control capabilities. They created communities so that incident response teams can gain a complete threat picture of the specific malware samples they are submitting without sharing the information more broadly. They can search for many different types of indicators and can pivot and pull all the samples that have the same indicator. Within a few minutes CIS members have the context they need to understand what the malware is doing or attempting to do, the scope of the threat it poses, and how to defend against it.

Malware Analysis Evaluation Checklist

When evaluating sandboxing technology, you should consider the following aspects to make sure you are getting a more robust approach to malware analysis that delivers the optimized workflow, context, and deployment flexibility required for effective protection.

- **Deployment options:** Malware sandboxing requires both cloud-based and on-premises analysis capabilities because of different sensitivities with regard to privacy concerns and compliance issues for more heavily regulated verticals. Look for solutions that offer the deployment flexibility to meet your organization's needs.
- **Scalability:** Enterprises kicking the tires on malware sandboxing technologies should evaluate each potential solution's capability to scale deployments. If it requires creating an overlay network, that approach could be costly and complex to deploy and manage.
- **Efficiency:** You should gain an understanding of what it takes operationally to best exploit the technology to speed the time between detection, analysis, and remediation.
- **Suitability:** Organizations should ensure that potential suppliers of the technology can replicate key components and endpoint configurations, and can examine file types that are a part of your organization's environment.
- **Integration:** Be sure you understand how the solution integrates with and shares information across other IT security solutions within your environment. Not only will this help accelerate detection and response during and after an attack, but also prevent similar attacks by automatically updating policies.

Conclusion

Just as attackers don't rely on one technique to accomplish their mission, so too must security professionals not rely on sandboxing, or dynamic malware analysis, as their single point of defense. But you can increase the effectiveness of every component of your defense. Each component must be part of a threat-centric approach to security.

Your defense should focus on visibility, context based on local and global threat intelligence, and control using analysis and automation to protect against threats whenever and wherever they are detected.

For More Information

For more information about Cisco Advanced Malware Protection, visit: <http://www.cisco.com/c/en/us/products/security/advanced-malware-protection/index.html>.