PATHFINDER REPORT

Next Generation Endpoint: Hype or Hope?

NOVEMBER 2016

COMMISSIONED BY

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

About 451 Research

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
20 West 37th Street
New York, NY 10018
+1 212 505 3030

SAN FRANCISCO
140 Geary Street
San Francisco, CA 94108
+1 415 989 1555

LONDON
Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 207 426 1050

BOSTON
One Liberty Square
Boston, MA 02109
+1 617 598 7200

Abstract

For years, the attacker/defender relationship has resembled a game of leapfrog. Defensive approaches have largely addressed or mitigated the challenge at hand, solving the problem for a short time. Yet the bad guys are typically not burdened with the things that slow down defenders – politics, legality and ethics, for example. This has traditionally tipped the scales in favor of attackers, especially in terms of time. It can take years to develop defenses, but only days or weeks for attackers to learn how to evade them.

In the world of endpoint security and antivirus, we've clearly seen this effect in play with the need to shift away from using signatures as the primary method for malware prevention. No less than 37 new endpoint security products have emerged in the past five years, each claimed to address the changing landscape. A closer look shows that most of these new products either focus on a narrow spectrum of attacks, or use techniques that can be easily evaded. A more reliable approach should include cross-product integration and several layers of resilient prevention and detection techniques. This paper addresses the trends that have emerged to deal with malware challenges on the endpoint and presents customers with the information necessary to evaluate the right approach for their organization.

The Importance of the Endpoint

Given that the endpoint is rarely the end goal for attackers and criminal campaigns, why is it still such a focus for attacks? The answer, overwhelmingly, is twofold: people are still easy, reliable targets; and while Windows has made great strides in terms of hardening and more secure defaults, Microsoft's ubiquitous operating system is often still an easy, reliable target as well.

For some time, the endpoint has served as a launching point for attacks against a business. Recently, with the rise of ransomware, we've seen endpoints become the more lucrative target. It is important to understand that criminals are typically motivated to go after the biggest paycheck for the least amount of effort. Launching a fully automated campaign that infects endpoints and demands ransoms is much easier than using an endpoint as a launching point for a complex and risky attack against a company's most lucrative assets. These assets are typically better protected and defended these days, and attackers see less risk, less work and more reward in ransomware.

It has become increasingly clear that the battle for the endpoint is one of the more important problems in the industry. It has also become clear, with mobile and portable devices making up a high percentage of corporate assets today, that whatever we do to address this problem must live on the endpoint, or travel with it.

NGAV, EDR and Threat Hunting: Hype or Hope?

Are these nothing more than buzzwords, or our best hope in the battle against endpoint attacks? First, let's take a stab at defining these terms:

NEXT-GENERATION ANTIVIRUS (NGAV)

No, AV isn't dead; it just needed a change in tactics. 451 Research defines NGAV as having the ability to detect and prevent threats without prior knowledge of them. Some of the technologies that have been popular among NGAV vendors include machine learning, exploit prevention, behavioral analysis and hardening. Individually, these techniques are not sufficient to provide protection across the spectrum of approaches used by attackers. Historically, when simplistic single-layer defenses are deployed, it doesn't take long before attackers find a way to bypass them. For example, an NGAV solution that relies solely on machine learning is not effective against 'fileless' malware that often exploits vulnerabilities in common software and uses macros and memory-based attacks to compromise the endpoint.

The NGAV market is almost entirely focused on building a better malware mousetrap, and these approaches are likely to merge with existing AV techniques. The bottom line? The more broad and layered a product's defenses are, the more effective they will be in protecting against a broad spectrum of attacks.

ENDPOINT DETECTION AND RESPONSE (EDR)

First and foremost, the EDR market is focused on gathering forensic data from endpoints. The data commonly collected includes changes to files and the registry, network connections, process details, and configuration changes. A notable benefit of EDR products is that the data collected can be used to detect attacks that don't use malware. Detecting and preventing attacks that don't use malware will become increasingly important as NGAV technology improves and succeeds at stopping most malware.

The most notable difference between EDR and NGAV is that EDR products generally require a lot more tuning, monitoring and management. In other words, EDR products are likely to require considerable investment in specialized security staff to produce results, or the outsourcing of this role to managed service providers with similar staff. NGAV products, on the other hand, should stop the vast majority of malware with little to no input from security staff. For this reason, we generally recommend implementing NGAV first, to cut down on the number of successful attacks and infections – especially in environments where commodity malware still often gets past existing antivirus products. With commodity malware under control, putting the effort into leveraging an EDR product makes more sense, since this will presumably free up staff resources, allowing them to spend more time on products that have to be tuned and monitored. EDR is most commonly used to catch the more targeted malware and attacks that make it past preventative controls. Due to the nature of EDR products, they are typically more effective as detective controls than preventative ones.

THREAT HUNTING

Threat hunting is a trend that emerged out of necessity in recent years. As threat intelligence became more common and easy to come by, the need to quickly and easily search all endpoints for indicators of compromise emerged. As such, threat hunting is more of a use case than a product category. We define it simply as the ability to search an environment (which could include network traffic, logs or data – not just endpoint artifacts) for signs of an attack or compromise. Nearly all EDR products support the endpoint threat hunting use case.

Advanced Is the New Normal

Years ago it became clear that attackers were working on automating the generation of malware in such a way that it would always yield a unique binary likely to frustrate signature-based antivirus. Today this approach is commonplace. Nothing has really changed with the malware itself – it is still the same malicious software; it's just dressed up in different 'clothes.' Thus, what was once considered something only seen in attacks labeled as 'advanced' or 'sophisticated' has now become the norm. Commodity malware can't be considered advanced, so it is only reasonable that the anti-malware products labeled as advanced also be thought of as the norm – not upgrades or premium offerings.

Exploit Prevention: Strengths and Cautions

Software vulnerabilities are still the primary method used by malware to infect endpoints. While software vendors are improving response speed, it can still take weeks or months to get from vulnerability discovery to patches deployed throughout organizations. Malware authors are well aware of these windows of opportunity, and take advantage of them. One technique some NGAV vendors have employed to counter such attacks is exploit prevention. An endpoint product monitors software processes for well-known exploit attempts in memory (e.g., heap sprays and buffer overflows) and suppresses them when detected.

To ensure compatibility, however, such a solution needs to be extensively tested in an IT environment. While NGAV approaches typically require less effort to tune and manage compared with EDR approaches, exploit prevention can be the exception, at least in the initial stages of implementation. The effort necessary to properly tune exploit prevention is also dependent on the complexity and diversity of end-user environments within an organization. We recommend customers consider the labor necessary to fine-tune exploit prevention to suit their organization's needs.

Machine Learning: Clever Trick or Long-Term Strategy?

Machine learning (ML) – sometimes referred to as artificial intelligence – appears to be effective in quickly identifying malware variants. Machine learning is leveraged in malware detection by identifying features and labels. In ML parlance, features refer to the data available for analysis, and labels refer to the desired results. In this case, the label will simply tell us if an executable file is good or bad. Once a model is trained and deemed accurate enough, it can be deployed to endpoints to provide near-real-time detection of malicious software. In some cases, several models are trained and even made available to the customer, who can choose based on how well a given model balances accuracy with false positives, which are an unavoidable side effect – legitimate software such as low-level drivers and some security products can look like malware.

First, feed the ML project with large amounts of good quality data.

There are two keys to ensuring a model is accurate. First, feed the ML project with large amounts of good quality data. The input here will be existing malware samples. Malware rarely undergoes major changes, with many new malware families and strains reusing code from previous malware and re-implementing existing techniques. To create entirely new malware that looks completely different from everything that's previously been seen would be a huge effort for most criminal operations, which allows defenders to have the economic advantage over attackers in this case. **Our conclusion is that the bigger and better the repository of files an ML-based NGAV solution has been trained with, the more effective it will be at identifying previously unseen malware**.

The second key to ensuring an accurate model is keeping up with malware and attacker trends.

The second key to ensuring an accurate model is keeping up with malware and attacker trends. Eventually attackers will develop their own ability to test malware against ML models, just as they've built their own versions of the Virus Total service, allowing malware to be tested against industry products prior to using it in campaigns. ML models need to follow these trends and retrain models as necessary in order to ensure that an increasing number of new malware samples aren't getting past existing models.

Beyond Machine Learning: Behavioral Analysis, Hardening and More

ML is designed to catch the majority of malware threats, not 100% of them. This is why we still advocate that an endpoint protection solution employ additional layers of controls designed to catch what gets through the initial layers of preventative controls.

Take ransomware, for example. While an ML model may stop the majority of malware, some may still get through, causing significant damage to a business. Here, **behavioral analysis** can be a useful layer. Researchers have noticed that instances of crypto-ransomware have some behaviors in common. Most ransomware seeks to first disable VSS (the Windows volume shadow copy service) and delete snapshots, to ensure victims can't avoid paying the ransom by restoring data from saved snapshots. Researchers can use this knowledge to implement behavioral rules that look for any untrusted process trying to disable VSS or delete snapshots. Any that are found can, by default, be assumed malicious and terminated. Behavioral analysis coupled with host integrity policies can be useful in protecting against such threats, likely to be missed by ML.

Application and device control can be effectively used to harden an endpoint to prevent malware from executing on the endpoint in the first place. While application and device control can be used to completely lock down a system, it can also be employed more selectively to maximize malware prevention while minimizing end-user disruption.

Host-based intrusion prevention, combined with a **host-based firewall** and behavioral analysis, can be particularly useful in recognizing malware communicating with command-and-control servers. These controls can be effective in preventing an attacker's lateral movement throughout the rest of the organization's environment.

Finally, a complete endpoint protection solution should offer robust methods for containing and remediating any detected malware. Attempting to manually remove malware and all related artifacts can be a painstaking job. If done incorrectly, the malware could return. For this reason, many organizations choose to wipe and reimage the affected computer's hard drive. In many cases, this could actually be more disruptive and costly to the business than any other option, particularly if the malware turns out to be benign. Many NGAV and EDR products can be limited in this area — they may be capable of detecting attacks, and perhaps isolating the malware or the infected system, but not capable of doing the cleanup work so that a production system can be returned to normal operation. (See the subsequent section on the true cost of securing the endpoint for estimates on the potential cost of reimaging infected systems.) We recommend that customers bear in mind the need for remedial features, and the comparative cost of using manual workflows instead.

Importance of Integration

Integration is also important — the security industry is entering the age of platforms, where security tools no longer need to 'go it alone' as point products. The fact of the matter is that not all endpoints will be guaranteed to have corporate anti-malware software installed, which is where network-based products can fill the gap. Also, not all corporate assets will be on the corporate network, which is where endpoint agents can fill in, sending suspicious files, sensor data and other intelligence off to the cloud for analysis. The integration of security products is a bit like using lighting in a photography studio to eliminate shadows — a single light source can't do it; multiple lights and reflective surfaces are necessary to eliminate them completely.

What's the True Cost of Securing the Endpoint?

451 Research's Voice of the Enterprise (VotE) service collects data on a quarterly basis from over 800 IT leaders and budget holders. In the first two quarters of 2016, malicious software held the number one spot among internal security pain points by a wide margin. Combined with endpoint security, which was listed as a separate pain point, this market is overwhelmingly — at 23.3% of respondents, vs. 'data loss/theft' at number two overall with 10.2% — the biggest pain point for the

enterprises we talk to. Additionally, when the group surveyed was asked about the top three InfoSec projects over the next 12 months, endpoint security also took the top spot, despite nearly 100% of respondents saying they already had endpoint security products in place.

Why the sudden popularity of endpoint security? Malware has always been a top pain point for enterprises, but in recent years we saw vendors explore other markets, such as appliance-based malware sandboxes, for relief. The answer is that the endpoint security market is now flooded with more than 80 vendors, over half of which are complementary offerings. They are complementary in the sense that these products are intended to be used in addition to an existing antivirus product, and don't replace them. They aim to temporarily fill a gap that enterprises are desperate to fill, especially with the popularity of ransomware. The problem is also that while enterprises have decent confidence in their endpoint security tools to detect and prevent known malware (75% and 68% of respondents, respectively), there is little confidence in these tools' ability to detect and prevent unknown malware (29% and 25%, respectively).

The result is that multi-product use is rampant in the enterprise. Customers are typically loath to install more than one agent for security, but our survey results show that only 13% rely on a single product. Two products are used by 27% of those surveyed, and, staggeringly, 59% use three or more products concurrently. In addition to layering on newer anti-malware products to fill a confidence gap, the majority of respondents also depend on features that are typically found only in mature endpoint security suites, like file integrity monitoring, full-disk encryption and configuration management.

This all means high cost and complexity. Complementary products cost anywhere from $20-150 per endpoint and represent additional management, maintenance and monitoring overhead. These products also have to work together. It is common for endpoint products to mistake other security products as malware and remove or disable them. Also, while a single security agent may claim to be lightweight, three or more agents increase the likelihood of slowing down users, regardless of how 'light' they might be individually. Additional considerations include:

- Are additional staff necessary to monitor/maintain products?
- Is training necessary for operational staff?
- Are managed services available? How much do they cost, in addition to the product?
- How does the product work? Preventative products (AV, NGAV) tend to have a low labor-to-value ratio – they either stop malware or they don't – whereas detective products (EDR) typically have a high labor-to-value ratio. Since detective products look for signs that malware has already infected systems, they're typically used as part of the incident response process, which is much more expensive in terms of labor compared with pre-execution detection and prevention.
- Can the product contain the threat?
- Who is cleaning up this mess? Most of the newer products on the market focus on finding the attack, but don't clean up the mess. 451 Research estimates that the cost of cleaning up a single infection ranges from $650 to over $2,800, when calculating the cost of employee labor and lost productivity, and can involve up to 38 work hours across all involved parties in a best-case example where the infection turns out to be benign.
- The endpoint security market is currently consolidating, meaning that some enterprises currently paying for these complementary technologies may start to find comparable next-gen functionality in existing primary endpoint security products for a much lower price, or even at no additional cost.

The ideal product would provide AV, NGAV and EDR functionality in a single agent, from a single vendor. Additionally, the ability to integrate with network threat detection and prevention products would be so much icing on the cake.

As a result of this consolidation, we anticipate that in just a few years the market will consist of two main categories. One will be a platform approach, or suite of products or modules, which will include threat detection/prevention capabilities, traditional EPP features and EDR features. The second category will be focused on hardening – making Windows more resilient in general against attacks, and will include things like privilege management, whitelisting, better process segmentation or application sandboxing, and other threat-agnostic approaches.

Conclusion: Defense in Depth Is Necessary to Win the Battle for the Endpoint

Each new defensive technology or technique tends to push adversaries down new or different paths. The more effective the technology, the more likely it is that attackers will be forced to try new things, and it is time to start anticipating and planning for these shifts in tactics and strategy. Even if the industry were to succeed in abolishing malware entirely, companies will still get hacked and breaches will still occur. In fact, credential theft is often the initial point of entry in some of the largest breaches that have been analyzed, where malware only plays a part.

It is time to look beyond point products and look at the big picture. Consider a layered approach and plan for failure. Ask the question, "What if this layer fails? What picks up where this control leaves off?" Even more importantly, ask the question, "What if this layer succeeds?" What path will our adversaries choose instead? Do we have what we need to follow them down each new potential path?