



Endpoint Advanced Protection

Version 1.3

Released: November 11, 2016

Author's Note

The content in this report was developed independently of any sponsors. It is based on material originally posted on [the Securosis blog](#), but has been enhanced, reviewed, and professionally edited.

Special thanks to Chris Pepper for editing and content support.

This report is licensed by Check Point Software Technologies Ltd..



www.checkpoint.com

Check Point Software Technologies Ltd. (www.checkpoint.com) is the largest pure-play security vendor globally, provides industry-leading solutions, and protects customers from cyberattacks with an unmatched catch rate of malware and other types of attacks. Check Point offers a complete security architecture defending enterprises' networks to mobile devices, in addition to the most comprehensive and intuitive security management. Check Point protects over 100,000 organizations of all sizes. At Check Point, we secure the future.

Copyright

This report is licensed under Creative Commons Attribution-Noncommercial-No Derivative Works 3.0.

<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>



Endpoint Advanced Protection

Table of Contents

The State of the Endpoint Security Union	4
The Endpoint Advanced Protection Lifecycle	7
The Evolution of Prevention	12
Detection and Response	19
Remediation and Deployment	24
Summary	28
About the Analyst	29
About Securosis	30

The State of the Endpoint Security Union

Innovation comes and goes in security. Back in 2007 network security had been stagnant for more than a few years. It was the same old same old. Firewall does this. IPS does that. Web proxy does a third thing. None of them did their jobs particularly well, all struggling to keep up with attacks encapsulated in common protocols. Then the next generation firewall emerged, and it turned out that regardless of what it was called, it was more than a firewall. It was the evolution of the network security gateway.

The same thing happened a few years ago in endpoint security. Organizations were paying boatloads of money to maintain endpoint protection, because PCI-DSS required it. It certainly wasn't because the software worked well. Inertia took root, and organizations continued to blindly renew their endpoint protection, mostly because they didn't have any other options.

Organizations were paying boatloads of money to maintain endpoint protection, because PCI-DSS required it. It certainly wasn't because the software worked well. Inertia took root, and organizations continued to blindly renew their endpoint protection, mostly because they didn't have any other options.

But in technology, inertia tends not to last more than a decade or so (yes, that's sarcasm). When there are billions of [name your favorite currency] in play entrepreneurs, investors, shysters, and plenty of other folks swarm in, trying to get some of that cash. So endpoint security is the new hotness. Not *only* because folks think they can make a buck displacing old and ineffective endpoint protection.

The sad fact is that adversaries continue to improve — both in the attacks they use, and the way they monetize compromised devices. One example is ransomware, which some organizations discover several times each week. We know of some organizations which tune their SIEM to watch for file systems being encrypted. Adversaries continue to get better at obfuscating attacks and exfiltration. As advanced malware detection technology matures, attackers have discovered many opportunities to evade detection. It's still a cat and mouse game, but both cats and mice have gotten much better at it. Finally, every organization still has to deal with employees, who are usually the path of least resistance. Regardless of how much you spend on security awareness training, knuckleheads with access to sensitive data will continue to enjoy clicking pictures of cute kittens... and other stuff.

So what about prevention? It has been security's holy grail for decades. To stop attacks before they compromise devices. Unfortunately prevention turns out to be hard, so the technologies don't work very well. Or they work but only in limited use cases. The challenge of prevention is also

The challenge of prevention is also compounded by the shysters I mentioned above, who claim nonsense like "Our products stop all zero-days!" — but of course there is no evidence, or it's completely bogus.

compounded by the shysters I mentioned above, who claim nonsense like "Our products stop all zero-days!" — but of course there is no evidence, or it's completely bogus. Obviously they have heard you never let truth get in the way of marketing. There has been incremental progress, and that's good news. But it's not enough.

On the detection side people realized more data could help detect attacks. Both close to the point of compromise, and afterwards during forensic investigation. So endpoint forensics is a thing now. It even has its own name: EDR (Endpoint Detection and

Response), named by the analysts in their ivory tower who label technology categories. The key is that as more organizations invest in incident response, they can leverage the granular telemetry offered by these solutions. But they don't really provide visibility for everyone because they require specialized security skills. For those who understand how malware really works, and can figure out how attacks manipulate kernels, these tools provide excellent visibility. Unfortunately these features are useless to most organizations.

But we have still been heartened to see a focus on more granular visibility, which provides skilled incident responders (who we call 'forensicators') more and richer data to figure out what happened during attacks. Meanwhile operating system vendors continue to improve their base technologies to be more secure and resilient. Offerings like Windows 10 and OS X 10.11 are far more secure, and popular applications (primarily office automation and browsers) have been locked down and/or re-architected for stronger security. We also have seen add-on tools to further lock down operating systems, such as [Microsoft's EMET](#).

State of the Union: Sadness

We have seen plenty of innovation. But the more things change, the more they stay the same. It's a new day, but security professionals will still spend a portion of it cleaning up compromised endpoints. That part hasn't changed.

The security industry also faces an intractable security skills shortage. As mentioned above, granular endpoint telemetry cannot help if you don't have staff who understand what the data means, or how similar attacks can be prevented. And most organizations don't have that skill set in-house.

The security industry also faces an intractable security skills shortage. As mentioned above, granular endpoint telemetry cannot help if you don't have staff who understand what the data means, or how similar attacks can be prevented.

It is really the best of times, and the worst of times. But if you ask most security folks, they'll tell you it's the worst.

Thinking Differently about Endpoint Protection

But it's not over. Remember "[Nothing is over until we say it is](#)," (hat tip to Animal House — though be aware that clip contains strong language). If something is not working, you need to think differently, unless you *want* to be having the same discussions in 10 years.

We need to isolate the fundamental reason it's so hard to protect endpoints. Is it that our ideas of *how* are wrong? Or is the technology not good enough? Or have adversaries changed so dramatically that all the existing ways to do endpoint security (or security in general) need to be tossed out? It's actually all of the above. We have used ineffective techniques far too long, technology that is long in the tooth, and we underestimate the innovation coming from our adversaries. It's a failure any way you look at it.

Fortunately technology which can help has existed for a few years. It's just that not enough organizations have embraced new endpoint protection methods. And many of the same organizations continue to be operationally challenged in security, which doesn't help — you are pretty well stuck if you cannot keep devices patched, or take too long to figure out someone is running a remote access trojan on your endpoints, on your networks.

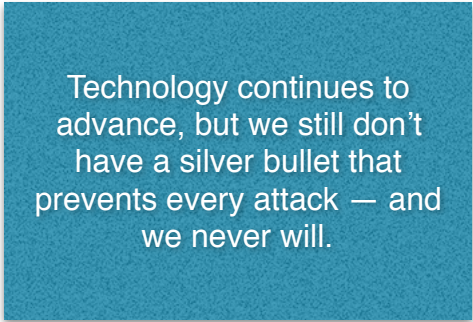
The Endpoint Advanced Protection Lifecycle

At some point you need to decide where endpoint protection starts, and where it ends. Figuring out how your endpoint security strategy integrates with the other security defenses in your environment is critical because today's attacks require more than just a single control — you need an *integrated system* to protect devices. The other caveat, before we jump into the lifecycle, is that we are actually trying to address the *security* problem here — not just compliance. We aim to actually protect devices from advanced attacks. We understand that's a very aggressive objective — some say crazy, given how fast adversaries learn and improve attacks. But we wouldn't be able to sleep at night if we merely accepted mediocrity for our defenses, and we figure you are similar... so let's aspire to this lofty goal.



Let's dig into the different aspects of the Endpoint Advanced Protection (EAP) lifecycle:

1. **Gain Visibility:** You cannot protect what you don't know about — that hasn't changed, and isn't about to. So the first step is to gain visibility into all devices that have access to sensitive data within your environment. It's not enough to just find them — you also need to assess and understand the risk they pose to your organization. We will focus on traditional computing devices here, but smartphones and tablets are increasingly used to access corporate networks.
2. **Reduce Attack Surface:** Once you know what's out there, you want to make it as difficult as possible for attackers to compromise it. That means practicing good hygiene on devices — making sure they are properly configured, patched, and monitored. We understand many organizations pretty much stink operationally, but endpoint protection is much more effective once you get rid of the low-hanging fruit making it easy for attackers to own your devices.
3. **Prevent Threats:** Next try to stop successful attacks. Unfortunately, despite continued investment and promises of better results, the results are still less than stellar. And with new attacks like ransomware making compromise even worse, stakes are getting higher. Technology continues to advance, but we still don't have a silver bullet that prevents every attack — and we never will.
4. **Detect Malicious Activity:** You cannot prevent every attack, so you need a way to detect attacks after they penetrate your defenses. There are a number of detection options. Most are based on watching for patterns indicating a compromised device, but there are many other indicators which can help you identify devices being attacked. The key is to reduce the time after a device is compromised but before you know about it.
5. **Investigate and Respond to Attacks:** Once you believe a device has been compromised, you need to verify the attack succeeded, determine your exposure, and take action to contain the damage as quickly as possible. This typically involves a triage effort, quarantining the device, and then moving to a formal investigation — including a structured process for gathering forensic data, establishing an attack timeline to help determine the attack's root cause, an initial determination of potential data loss, and a search to determine how widely the attack spread within your environment.
6. **Remediate:** Once the attack has been investigated you can put a plan in place to recover. This might involve fixing the device by rolling back the attack, or reimaging it and starting over again. This step can leverage ongoing hygiene tools such as patch and configuration management. There is no point reinventing the wheel, and tools to accomplish the necessary activities are already in use for day-to-day operations.



Technology continues to advance, but we still don't have a silver bullet that prevents every attack — and we never will.

Gaining Visibility

You need to know what you have, how vulnerable it is, and whether it can be exploited. With this information you can prioritize your exposure and design a set of security controls to protect your most important assets. Start by understanding what in your environment would interest an adversary. There is something of interest at every organization. It could be as simple as compromising devices to launch attacks on other sites, or as focused as gaining access to your environment to steal your crown jewels. When trying to understand what an advanced attacker is likely to come looking for, there is a fairly short list of asset types — including intellectual property, protected customer data, and business operational data (proposals, logistics, etc.).

Once you understand your potential targets, you can profile adversaries likely to be interested in them. The universe of likely attacker types hasn't changed much over the past few years. You face attacks from a number of groups across the continuum of sophistication. Starting with unsophisticated attackers (which can include a retiree with too much time on their hands... or possibly a 10-year-old wielding Metasploit), organized crime, competitors, and/or state-sponsored adversaries. Understanding likely attackers provides insight into probable tactics, so you can design and implement security controls to address those risks. But before you can design a security control set, you need to understand where the devices are, as well as their vulnerabilities.

Once you understand your potential targets, you can profile adversaries likely to be interested in them. The universe of likely attacker types hasn't changed much over the past few years. You face attacks from a number of groups across the continuum of sophistication.

Discovery

This process finds the devices accessing critical data and makes sure everything is accounted for. This simple step helps to avoid “oh crap” moments — it's no fun to stumble over a bunch of unknown devices with no idea what they are, what they have access to, or whether they are cesspools of malware.

A number of discovery techniques are available, including actively scanning your entire address space for devices and profiling what you find. This works well and is traditionally the main method of initial discovery. You can supplement with passive discovery, monitoring network traffic to identify new devices from their communications. Depending on the sophistication of the passive analysis, devices can be profiled and vulnerabilities can be identified, but the primary goal of passive monitoring is to discover unmanaged devices faster. Passive discovery is also useful for identifying devices hidden behind firewalls and on protected segments, which active discovery cannot reach.

As if you needed further complications, these cloud and mobility things everyone keeps jabbering about make discovery more challenging. Embracing software as a service (SaaS), as pretty much everyone has, means you might never have a chance to figure out exactly which devices are accessing critical resources. For devices which don't need to go through your monitored corporate networks, you need other ways to discover and protect them. That could involve a trigger on authentication to a SaaS service, or possibly having your endpoint protection capability leverage the cloud and phone home to relay device telemetry to a central management system.

Assessment

Once you know what's out there, you need to figure out how vulnerable it is. That typically requires some kind of vulnerability scan on discovered devices. Key features to expect from your assessment function include:

- **Device/Protocol Support:** Once you find an endpoint you need to determine its security posture. Compliance demands that we scan all devices with access to private/sensitive/protected data, so any scanner should assess all varieties of devices in your environment which have access to critical data.
- **External and Internal Scanning:** Don't assume adversaries are purely external or purely internal — you need to assess devices both inside and outside your network. Look for a scanner appliance (which might be virtualized) to scan from the inside. You will also want to monitor your IP space from the outside (either with a scanner outside your network, or a cloud service) to identify new Internet-facing devices, find open ports, etc.
- **Accuracy:** False positives waste your time, so accuracy of scan results is key. Also pay attention to the ability to prioritize results. Some vulnerabilities are more important than others, so being able to identify the ones truly posing risks to your organization is critical.
- **Threat Intelligence:** Adversaries move fast and come up with new attacks daily. You'll want to ensure you factor new indicators into your assessment of security posture quickly and without extensive manual effort.
- **Scale:** You likely have many endpoints. Today's large enterprises can have hundreds of thousands — if not millions — of devices requiring assessment. Also make sure your tool can assess devices that aren't always on the corporate network, smartphones & tablets, and hopefully cloud resources (such as desktop virtualization services).

The assessment provides insight into how each specific device is vulnerable, but that's not the same as risk. Presumably you have a bunch of network defenses in front of your endpoints, so attackers may not be able to reach a particular vulnerable device. That's the difference between something that is *vulnerable* and an asset that is *exploitable*. Consider this as part of risk prioritization.

It may not be as sexy as advanced detection or cool forensics technology, but these assessment tasks are necessary before you can even start thinking about building controls to prevent advanced attacks.

A Risk-Based Approach to Defending Endpoints

Security practitioners have an unfortunate tendency to miss the forest for the trees when discussing endpoint advanced protection. The reality is that each device contains a mixture of data types — some data types present great risk to the organization, but others don't. Keep in mind that some protection techniques are very disruptive to end users, and can be expensive to both procure and manage.

So we advocate a risk-based approach to protecting endpoints. This involves grouping endpoint devices into a handful (or less than a handful) of risk categories. Then determine the most effective means to protect the devices in each category. For example you might want to implement whitelisting on all kiosks in stores and warehouses. Or you might add an advanced exploit prevention agent to devices used by senior management, Human Resources, and Finance, and anyone else handling especially sensitive or attractive information. Finally you might just use free AV on devices which only have outbound access from common areas, because they don't have access to anything important on the corporate network.

So we advocate a risk-based approach to protecting endpoints. This involves grouping endpoint devices into a handful (or less than a handful) of risk categories. Then determine the most effective means to protect the devices in each category.

There are as many permutations as devices on your network. To scale this approach you need to categorize risk tiers effectively. But a one-size-fits-all approach won't work either, given the variety of different approaches available for detecting advanced malware.

The Evolution of Prevention

Once you know what you need to protect, and how vulnerable it is, you try to prevent attacks, right? Was that a snicker? You've been reading the trade press and security marketing telling you prevention is futile, so you're a bit skeptical. You have every right to be — time and again you have had to clean up ransomware attacks (hopefully before they encrypt entire file servers), and you detect Command and Control traffic indicating popped devices frequently. A sense of futility regarding actually preventing compromise is all too common.

So the key objective of any prevention strategy must be making sure you aren't the path of least resistance. That requires reducing attack surface and risk-based prevention.

But despite feelings of futility, we still see prevention as key to any endpoint protection strategy. It needs to be. Imagine how busy (and frustrated) you would be if you completely stopped trying to prevent attacks, and just left a bunch of unpatched Internet-accessible Windows XP devices on your network, figuring you'd just detect and clean up every compromise after the fact. That's about as silly as counting on stopping all attacks.

So the key objective of any prevention strategy must be making sure you aren't the path of least resistance. That requires reducing attack surface and risk-based prevention. Shame on us if devices are compromised by attacks which have been circulating in the wild for months. Ensuring proper device hygiene on endpoints is job one. Then it's a question of deciding which controls are appropriate for each specific employee (or, more likely, group of employees). There are plenty of options for blocking malware attacks, some more effective than others. Unfortunately the most effective controls are also highly disruptive. So you need to balance inconvenience against risk to determine which makes the most sense in which scenario. If you want to keep your job, that is.

'Legacy' Prevention Techniques

It is often said that you can never turn off a security control. We see the truth in that adage when we look at the technologies used to protect endpoints today. We carry around (and pay for) historical technologies and techniques, largely regardless of effectiveness, and that complicates actually defending against the attacks we see.

The good news is that many organizations use endpoint protection platforms (EPP), which over time reduces the use of less effective tactics — at least in theory. We cannot fully cover prevention tactics without mentioning legacy technologies. They are still in use, largely under the covers of whichever EPP you select.

- **Signatures (LOL):** Signature-based controls are all about maintaining a huge blacklist of known malicious files to prevent from executing. Free AV products currently on the market typically use *only* this strategy, but the broader commercial endpoint protection platforms have been supplementing traditional signature engines with additional heuristics and cloud-based file reputation for years. This technique is used primarily to detect known commodity attacks representing the lowest bar of attacks in the wild.
- **Advanced Heuristics:** Endpoint detection needed to evolve beyond what a file looks like (hash matching), paying much more attention to what malware does. The problem with early heuristics was a lack of sufficient context to know whether an executable was taking a legitimate action. Malicious actions were defined generically for each device based on operating system characteristics, so false positives (notably blocking a legitimate action) and false negatives (failing to block an attack) were both common — a lose/lose scenario. Fortunately heuristics have evolved to recognize normal application behavior, minimizing false positives. This dramatically improved accuracy by building and matching activity against application-specific rules. But this requires understanding all legitimate functions within a constrained universe of frequently targeted applications, and developing a detailed profile for each covered application. Any unapproved application action is blocked. This requires vendors to maintain a positive security model for each application — a tremendous amount of work.
- **AWL:** Application White Listing entails implementing a default deny posture on endpoint devices (and often servers as well). The process is straightforward: define a set of authorized executables which can run on a device, and block everything else. With a strong policy in place AWL provides true device lockdown: no executables (either malicious or legitimate) can execute without explicit authorization. But the impact to user experience is often unacceptable, so this technology is mostly restricted to very specific use cases, such as servers and fixed-function kiosks, which shouldn't run general-purpose applications.
- **Isolation:** A few years ago the concept of running apps in a “walled garden” or sandbox on each device came into vogue. This technique enables us to shield the rest of a device from a compromised application, greatly reducing the risk posed by malware. Like AWL, this technology continues to find success in particular niches and use cases, rather than as a general answer for endpoint prevention.

Advanced Techniques

We cannot afford to ignore old-school techniques, because a lot of commodity malware still in circulation can be stopped by signatures and advanced heuristics. Maybe it's 40%. Maybe it's 60%. Regardless, we need more to fully protect endpoints. So endpoint security innovation has focused on advanced prevention and detection, as well as optimizing for prevalent attacks such as ransomware.

Let's unpack the new techniques to make sense of all the security marketing hyperbole getting thrown around. You know, the calls you get and email flooding your inbox, telling you how these shiny new products can stop zero-day attacks, with no false positives and insignificant employee disruption. But we don't know of any foolproof tools or techniques, so we will focus the latter half of this paper on detection and investigation. In fairness, advanced techniques do dramatically increase the ability of endpoints to block attacks.

You need to make sense of all the security marketing hyperbole getting thrown around. You know, the calls you get and email flooding your inbox, telling you how these shiny new products can stop zero-day attacks, with no false positives and insignificant employee disruption.

Anti-Exploit/Exploit Prevention

The first major category of advanced prevention techniques focuses on blocking exploits before the device is compromised. Security research has revealed how malware actually compromises endpoints at a low level, so tools now look for those indicators. You can pull out our favorite healthcare analogy: by understanding the fundamental changes an attack causes within an organism, you learn what to look for generally, rather than focusing on a specific attack, which can morph in an infinite number of ways.

These tactics break down into a few buckets:

- **Profiling Exploit Behavior:** This takes the advanced heuristics approach described above deeper into the innards of the operating system. Where advanced heuristics focus on identifying anomalous application behavior, anti-exploit tools focus on what happens to an actual machine when malicious code takes over the device. The key insight is that there are a discrete and known number of ways to compromise a device, regardless of attack vector. Blocking those behaviors stops exploits.
- **Memory Analysis/Protection:** One of the latest waves of attack doesn't even deal with traditional malware files. Malicious code is inserted directly into a command line or other means of manipulating an operating system without ever hitting disk. Defending against these attacks requires analyzing device memory continually, preventing memory corruption and logic flaws. This technology is very sophisticated, and can have a severe impact on device operation, so it demands thorough testing to ensure there is no unacceptable impact on your devices.
- **Compromised System Processes:** Aside from hiding attacks in memory, attackers now use fundamental operating system features to defeat whitelisting and isolation techniques. The most frequently targeted OS services include WMI, PowerShell, and EMET. This is also known as a "malware-less" attack because it doesn't use traditional malicious software. These attacks are much more challenging to detect because these system processes are authorized by definition. To defend against these attacks advanced technologies need to monitor the behaviors of all processes to make sure an approved process hasn't been

hijacked. This requires profiling legitimate behavior of common system processes, then looking for anomalous activity.

All 'advanced' endpoint protection technology includes these techniques, though they may be branded differently. It is all largely the same approach of looking for anomalous behavior, focused on OS and device innards instead of userspace applications.

Endpoint Bot Detection

Pretty much every modern attack, whether it involves malware or not, involves communicating with a Command and Control network to download the attack payload and receive instructions. So endpoint network-based detection has evolved to look for Command and Control patterns, similar to detecting malware on the network with a purpose-built device (typically residing on the network perimeter).

This is important because endpoints aren't always on the corporate network, which you are presumably already scanning for Command and Control traffic. So recognizing when a device in a coffee shop or hotel is communicating with known malicious sites can help you detect a compromise before it reconnects to the corporate network. This requires integration with a threat intelligence source, to keep an updated list of known malicious sites on the endpoints.

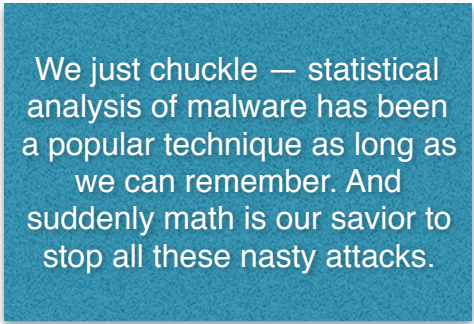
Dynamic File Testing

Many attacks still involve a compromised file executing code on a device, so network and cloud sandboxes are heavily used to dynamically execute inbound files and ensure they are not malicious. You have a number of options for where to test files, including the perimeter and/or email security gateway. But remote personnel remain challenging, because their network traffic doesn't run through the network's corporate defenses.

You can supplement those corporate controls with the ability to extract and test files on endpoints. The file will be checked to see if it has a known bad hash; if not it can be retested in the corporate sandbox. Some organizations now convert any easily compromised file (Office files) into a sanitized PDF to remove any code without impacting document appearance. If the original file is needed it can be routed to the recipient after clearing the sandbox.

Enabling Technologies

The first technology you have certainly been hearing a lot about is machine learning. It is used in many contexts other than endpoint protection, but the endpoint advanced security messaging has become very prominent. We just chuckle — statistical analysis of malware has been a popular technique as long as we can remember. And suddenly math is our savior to stop all these nasty attacks.



We just chuckle — statistical analysis of malware has been a popular technique as long as we can remember. And suddenly math is our savior to stop all these nasty attacks.

But the math really *is* better now. Combined with a much more detailed understanding of how malware actually compromises devices, more sophisticated static file analysis does help detect attacks. But we have to wonder whether these new techniques are really just the next generation of AV signatures.

We try to avoid getting wrapped up in vernacular or semantics. If these techniques help detect attacks more accurately at scale, whether they look like signatures or not isn't the important thing. It's not like we (or anyone else) believe machine learning is a perfect solution for endpoint protection. It's just another development in the never-ending arms race of malware vs. protection.

The other enabling technology that warrants mention is threat intelligence. Or security research, as endpoint protection vendors have been calling it for a decade. The fact is that whether you are adding new indicators to an endpoint agent or updating a list of known malicious sites for Command and Control detection, each endpoint agent needs to be updated frequently to stay current. Especially devices that don't sit behind the corporate network's perimeter defenses.

Protecting the Point of Attack

We should address the best place to place protection because you have a few options. The path of least resistance remains network-based solutions, which can be deployed without user impact. Of course these options cannot protect devices which aren't behind the corporate perimeter. Nor can network-based solutions provide context on individual user behavior, like something running on the device can.

You can run all traffic through a VPN or a cloud-based filtering service to provide some protection for remote devices. Running traffic through either enables you to gather telemetry and enforce

These new offerings represent an evolution of endpoint protection; either incumbents will add these capabilities to their existing offerings, or they won't survive.

corporate usage policies. On the downside this impacts traffic flow, and can be evaded by both savvy users and attackers. But it offers an option for addressing the limitations of filtering traffic through network defenses.

So should you add yet another agent to your endpoints, or use a plug-in into a common application such as a browser, to protect against the most common attack vector? If for some reason you cannot replace your existing endpoint agents, typically due to budget, politics

or a rogue assessor, looking at a bolt-in approach to provide additional protection for those devices that warrant it can certainly help as a stopgap.

But in case we haven't yet made it clear, these advanced endpoint security offerings are neither a long-term alternative, nor meant to run alongside an existing endpoint protection platform (EPP). These new offerings represent an evolution of endpoint protection; either incumbents will add these capabilities to their existing offerings, or they won't survive. And this is not just about prevention — it is also true for endpoint detection and response.

Anti-Ransomware

We don't normally call out specific attacks because they change so frequently. But ransomware demands some attention. The ability to so cleanly and quickly monetize compromised devices has rapidly made it the most visible attack strategy. And ransomware is not restricted to just one size or type of company, or of devices. We have seen ransomware targeting everyone and everything.

How can you combine advanced techniques to block ransomware? Fortunately, in technical terms ransomware is just another attack, so it can be profiled and blocked using advanced heuristics and exploit profiling. First look for attack patterns as they attempt to compromise the device; ransomware doesn't look too different from other attacks to detect with standard techniques and updated criteria.

Next look for clues within the endpoint's network stack — particularly Command and Control traffic — because attackers need to deliver their payload to lock down the machine. You can also look for anomalous searching of file shares, because ransomware often targets shared file systems for extra impact.

Additionally, because ransomware encrypts the local file system, you can monitor file I/O for anomalous activity. We suggest organizations more aggressively monitor storage networks and arrays for anomalous file activity. This can help shorten the detection window, and stop encryption before too much data is impacted.

And yes, these storage and network-centric techniques are out of the proper scope for this research, but device and data backup are essential for quick restoration of service in case of ransomware attack.

A Note on 'Effectiveness'

It is worth mentioning how to evaluate the effectiveness of these solutions. We refer back to our [Advanced Endpoint and Server Protection](#) research a few years ago, as this hasn't changed.

"As you start evaluating these advanced prevention offerings, don't be surprised to get a bunch of inconsistent data on the effectiveness of specific approaches. You are also likely to encounter many well-spoken evangelists spouting monumental amounts of hyperbole and religion in favor of their particular approach — whatever it may be — at the expense of all other options. This happens in every security market undergoing rapid innovation, as companies try to establish momentum for their approaches and products.

A lab test favoring one product or approach over another isn't much consolation when you need to clean up an attack your tools failed to prevent. And those evangelists are nowhere to be found when a security researcher shows how to evade their shiny technology at the latest Black Hat conference. We at Securosis try to float above the hyperbole and propaganda to keep you focused on what's really important — not claimed 1% effectiveness differences. If products or categories are within a few percent of each other across a variety of tests, we consider that a draw.

But if you look hard enough, you can find value in comparative tests. An outlier warrants investigation and a critical assessment of the test and methodology. Was it skewed toward one category? Was the test commissioned by a vendor or someone else with an agenda? Was real malware, freshly found in the wild, used in the test? All testing methodologies have issues and limitations — don't base a decision, or even a short list, around a magic chart or a product review/test."

Detection and Response

Despite all the cool innovation happening to prevent compromises on endpoints effectively, the fact remains that you cannot stop all attacks. So detecting the compromise quickly and effectively, and then figuring out how far the attack spread within your organization, continues to be critical.

Commercial endpoint detection tools were basically black boxes, not really providing visibility to security professionals. And the complexity of purpose-built forensics tools put this capability beyond the reach of most security practitioners.

Until fairly recently endpoint detection and forensics was a black art. Commercial endpoint detection tools were basically black boxes, not really providing visibility to security professionals. And the complexity of purpose-built forensics tools put this capability beyond the reach of most security practitioners. But a new generation of endpoint detection and response (EDR) tools is now available, providing much better visibility and more granular telemetry, along with a streamlined user experience to facilitate investigation, regardless of analyst capabilities.

Of course it is better to have a more-skilled analyst than a less-skilled one, but given the hard truth of the security skills gap, our industry needs to provide better tools to make those less-skilled analysts more productive, faster. Now let's dig into some key aspects of EDR.

Telemetry/Data Capture

In order to perform any kind of detection you need endpoint telemetry. This raises the question of how much to collect from each device, and how long to keep it. This is almost a religious question, but we remain firmly convinced that more data is better. Some tools can provide a literal playback of activity on the endpoint, like a DVR recording of everything that happened. Others focus on log events and other metadata to understand endpoint activity.

You need to decide whether to pull data from the kernel or from user space, or both. Again, we advocate for data, and there are definite advantages to pulling data from the kernel. Of course there are downsides as well, including potential device instability from kernel interference. As always, testing products to understand the impact on stability and user experience is a good idea when looking at any innovative technology.

We need to acknowledge that more data is great... up to a point. If you gather so much telemetry that you can no longer effectively move, store, or analyze it... you passed the point of diminishing returns. Your focus should remain on data quality, to ensure what you are gathering is useful in detecting and responding to attacks.

You should collect telemetry appropriate to the risk presented by the device. We recommend you take a risk-centric view on protecting endpoints. Some devices possess very sensitive information, so you should collect as much telemetry as possible. Other devices present less risk, and may only warrant log aggregation and periodic scans.

There are also competing ideas about where to store the telemetry captured from all these endpoint devices. Some technologies are predicated on aggregating the data in an on-premise repository, while others perform real-time searches using peer-to-peer technology, and a new model involves sending the data to a cloud-based repository for larger-scale analysis.

Again, we try not to get religious about any specific approach. Stay focused on the problem you are trying to solve. Depending on your organization's sensitivity, storing endpoint data in the cloud may be politically infeasible. On the other hand, in a highly distributed organization, it might be very expensive to centralize data. Understand there isn't right and wrong here, rather figuring out what's going to work best for your organization.

Threat Intel

It's not like threat intelligence is a new concept in endpoint protection. AV signatures are a form of threat intel — but the industry has never called it that. The difference is that now threat intelligence goes far beyond hashes of known bad files, additionally trying to capture behavioral patterns that indicate an exploit. Whether the patterns are called Indicators of Compromise (IoC), Indicators of Attack (IoA), or something else, you can watch for these indicators on endpoints in real time to detect and identify attacks.

The difference is that now threat intelligence goes far beyond hashes of known bad files, additionally trying to capture behavioral patterns that indicate an exploit.

This new generation of threat intelligence is clearly more robust than yesterday's signatures. But that underplays the impact of threat intel on EDR. New tools provide retrospection, searching the endpoint telemetry data store for new attack patterns. This enables you to see if a new attack has been seen in the recent past on your devices, before you even recognized it as an attack.

The goal of detection is to shorten the window between compromise and when you know you've been compromised. If you can search for indicators when you learn about them — regardless of when the attack occurs — you may be able to find compromised devices before they start misbehaving, and presumably trigger other network-based detection techniques. We are focused on endpoint-centric threat intelligence and indicators here, but there is leverage to be had from being able to analyze patterns within a broader data store which includes network, user, and application data.

One key to selecting any kind of Endpoint Advanced Protection (EAP) product is to ensure that the vendor's research team is well-staffed and capable of keeping up with the pace of emerging attacks. The more effective the security research team is, the more emerging attacks you will be able to detect before an adversary can compromise your devices. This is the true power of threat intelligence.

Analytics

Once you have all the data gathered, and have enriched it with external threat intelligence, you are ready to look for patterns which may indicate compromised devices. Analytics (or also commonly called machine learning) is now a very shiny term in security circles, which we find very amusing. Early SIEM products offered analytics — you just needed to tell them what to look for. But security marketers are going to market, so whatever the particular vernacular, more sophisticated analytics do enable more effective detection of sophisticated attacks.

Every company claims they use its to find zero-day attacks and all other badness with no false positives or latency. No, we don't believe the hype. But the advance of analytical techniques, harnessed by math ninja called "data scientists," enables detailed analysis of every attack to find commonalities and patterns.

But what does that even mean? First we should define probably machine learning within the context of security analytics, because every company claims they use its to find zero-day attacks and all other badness with no false positives or latency. No, we don't believe the hype. But the advance of analytical techniques, harnessed by math ninja called "data scientists," enables detailed analysis of every attack to find commonalities and patterns. New techniques, notably graph analysis, have dramatically opened up the kinds of patterns and indicators that data scientists and security researchers can derive from aggregated telemetry and malware analysis.

So these patterns and indicators can be leveraged for both static analysis (what the file looks like) and dynamic analysis (what the program does) to make detection faster and more accurate. Even if you object to the horribly overused machine learning label (as we do), these new techniques matter and are making a big difference in our ability to detect complicated attacks.

Response

Once you have detected a potentially compromised device, you need to engage your response process. We have written extensively about Incident Response (including [Using TI in Incident Response](#) and [Incident Response in the Cloud Age](#)), so we won't go through the details of the IR process again here. But as we have described, EAP tools now provide more granular telemetry, along with the ability to investigate attacks within the management console.

Additionally, these tools increasingly integrate with other security response tools in your environment. EAP products add several response capabilities, including:

1. **Attack Visualization:** In many cases being able to visualize the attack on a device is very instructive for understanding how the malware works and what it does to devices. The management consoles of some EAP products offer a visual map to follow the activity of malware — including the process the attack impacted, kernel-level activity, and/or API calls. This timeline (of sorts) must also specify the files involved in the attack, and track network connectivity.
2. **Understanding Outbreaks:** As discussed above, a key aspect of EAP products is their ability to aggregate telemetry and search after the fact to determine whether other devices have been attacked by similar malware. This provides invaluable insight into how the attack has proliferated through your environment, and identifies specific devices in need of remediation or quarantine.
3. **Forensics:** You also need your endpoint agent to gather raw telemetry from the device and provide tools to analyze it. At times, especially with skilled forensicators involved, you will need full data to really dig into what the malware did. A key aspect of forensic analysis is enforcement of chain of custody for collected data, especially if prosecution is an option.
4. **Blocking the Next One:** Once an attack is detected and investigated, you will have all the information you need to look for that attack moving forward. Shame on you if the same attack gets you again. So integrate the response data with preventative controls on both the network (including firewalls and web filters to block C&C sites and other malicious addresses) and endpoints.
5. **Ease of Use:** EAP tools have been built for general security practitioners — not only for forensics ninja — so user experience has been a focus, in order to help less experienced professionals be more productive. This requires a much easier workflow for drilling down into attacks, and then pivoting to find the root cause.
6. **Integration with Enterprise Tools:** Another key criteria for EAP products is making sure they play nice with tools already in use. You will want the ability to send data directly to a SIEM for further correlation and analysis. You will also want to integrate with a case management system to track investigations.

Hunting

Finally, we should acknowledge another very shiny concept in security circles: hunting. It seems every practitioner aspires to be a hunter nowadays. OK, maybe that's a little exaggerated, but it's a cool gig. Hunters go out and proactively look for adversary activity on networks and systems, as opposed to waiting for monitors to alert and *then* investigating.

Psychologically, hunting is great for security teams because it gives the team more control over their environment. Instead of waiting for a tool to tell you things are bad, you can go figure it out yourself.

But the reality is that hunting is primarily relevant to the most sophisticated and advanced security teams. Looking around requires staff time, and unfortunately most organizations are not sufficiently staffed to achieve core operational goals, so they don't have folks with time to proactively hunt for bad stuff.

Keep in mind that hunters' tools are largely the same ones used to validate attacks on endpoints. A hunter needs to be able to analyze granular telemetry from endpoints and other devices. They need to search through telemetry to find activity patterns that could be malicious. They need to forensically investigate devices when they find something suspicious. Hunters also need to retrospectively look for attack indicators to understand which devices have been targeted — all pretty much what EDR tools do.

We aren't maligning hunting. If your organization can devote resources to stand up a hunting function, that's awesome.

Remediation and Deployment

Now that we have gotten through 80% of the endpoint advanced protection lifecycle, we can get into actually fixing the issues you find, and then some finer points on deploying advanced endpoint security capabilities.

Remediation

Once you have detailed information from the investigation, what are the key decision points for figuring out how to return your environment to a working condition? As usual, to simplify we step back to *who*, *what*, *where*, *when*, and *how*. And yes, any time we can make something difficult feel like being back in grade school, we do.

1. **Who?** The first question is about organizational dynamics. In this new age, when advanced attackers seem to be the norm, who should take lead in remediation? Without delving into religion or politics, the considerations are really time and effectiveness. Traditionally IT Operations has tools and processes for broad updates, reimaging, or network-based workarounds. But for advanced malware or highly sensitive devices, or when law enforcement is involved, you might also want a small Security team which can remediate targeted devices with specialized techniques.
2. **What?** This question is less relevant because you already know you are remediating a device. There may be some question of whether to prevent further outbreaks at the network level by blocking certain sites, applications, users, or all of the above, but here we are talking about endpoints.
3. **Where?** One of the challenges of dealing with endpoints is that you often have no idea where a device will be at any point in time, so remote remediation is essential. Sometimes you need to reimage a machine, and often that is not feasible to do remotely. But having a number of different remediation options, depending on device location and the nature of the attack, can ensure minimal disruption to impacted employees.
4. **When?** This is one of the most challenging decisions, because there are usually reasonable points on both sides of the argument: whether to remediate devices immediately, or quarantine the device and observe the adversary a bit to gain intelligence. We generally favor quick and full eradication, which requires leveraging retrospection to find all impacted devices (even if they aren't currently participating in the attack) and cleaning devices as quickly as practical. But sometimes more measured remediation is called for.

5. **How?** The question is whether reimaging the device, or purging malware without reimaging, is the right approach. We favor reimaging because of all the sneaky ways attackers can remain persistent on a device. Even if you think a device has been cleaned... perhaps it really wasn't. But with the more granular telemetry gathered by today's endpoint investigation and forensics tools (think DVR playback), it is possible to reliably back out all changes made, even within OS innards. Ultimately the decision comes back to the risk posed by the device, as well as disruption to the employee. The abilities to both clean and reimage are central to remediation.

There is a broad range of available actions so we advocate flexibility in remediation... as in just about everything. We don't think there is any good one-size-fits-all approach anymore — each remediation needs to be planned out according to risk, attacker sophistication, and the skills and resources available between Security and Operations. Taking all that into account, you can choose the best approach.

We don't think there is any good one-size-fits-all approach anymore — each remediation needs to be planned out according to risk, attacker sophistication, and the skills and resources available between Security and Operations.

EPP Replacement?

One of the most frustrating aspects of doing security is having to spend money on things you know don't really work. Traditional endpoint protection platforms (EPP) fit into that category. Which begs the question: are these Endpoint Advanced Protection (EAP) products are discussing robust enough, effective enough, and broad enough to replace the EPP incumbents?

The answer depends on a few smaller questions. First, the main reason you renew your anti-malware subscription each year is to fill that box on a compliance checklist. So you need a sense of whether your assessor/auditor would give you a hard time if you proposed something that doesn't use signatures to detect malicious activity. If they are likely to push back, maybe find a new assessor. Kidding aside, we haven't seen much pushback lately, thanks to the overwhelming evidence that Endpoint Advanced Detection/Prevention is markedly more effective at blocking current attacks. That said, it would be foolish to sign a purchase order to swap out protection on 10,000 devices without at least putting a call in to your assessor and understanding whether there is precedent for them to accept a new type of agent.

You also need to look at your advanced endpoint offering for feature parity. Existing EPP offerings have been adding features (to maintain price points) for a decade. A lot of stuff you don't need has been added, but there may be some capabilities you would miss. Make sure replacing your EPP won't leave a gap you will just need to fill with another product.

Keep in mind that some EPP features are now bundled into operating systems. For example full disk encryption is now available free as part of the operating system. In some cases you need to manage these OS-level capabilities separately, which weighs against an expensive renewal which doesn't cover all the bases to effectively protect endpoints.

Finally, consider price. Pretty much every enterprise tells us they want to reduce the number of security solutions they need. And supporting multiple agents and management consoles to protect endpoints doesn't make much sense. In your drive to consolidate, play off aggressive new EAP vendors against desperate incumbents willing to perform unnatural acts to keep business.

Migration

Endpoint protection has been a zero-sum game for a while. Pretty much every company has some kind of endpoint protection strategy. So every deal that one vendor wins is lost by at least one competitor. Vendors make it very easy to migrate to their products by providing tools and services. Of course you need to verify what's involved in moving wholesale to a new product, but the odds are it will be reasonably straightforward.

Many new EAP tools are managed in the cloud. Typically that saves you from needing to test and deploy an onsite management server. This makes things much easier and facilitates migration — employees can connect to a cloud-based software installation/distribution engine without bringing devices to HQ for upgrades. Some organizations still resist cloud-based management — if this includes you, you will want to confirm your vendor can support on-premise installation.

Finally, when planning the migration you need to consider which security functions should be implemented on each category of devices, as defined by the risk they pose. Earlier in this paper we talked about categorizing devices into risk buckets, and implementing controls based on the risk they present. You can install or enable different EAP modules depending on employee or device needs.

The vendor may well make it worth your while to license all their capabilities on all your devices. There is nothing wrong with that if the price is right. But do not consider only purchase price — keep in mind the total cost of managing the various capabilities across all your devices; as well as the impact on employees, in terms of device performance and user experience. Not every device needs application whitelisting, for example. Likewise only a subset of your devices may warrant EDR, given the challenge of moving endpoint telemetry across the network.

Integration

Finally, any new EAP offering needs to play nice with existing enterprise security tools. Here are a few with their integration points.

- **Network Controls:** As discussed above, if you detect an attack on an endpoint and isolate the C&C (Command and Control) network it's connecting to, wouldn't it be great to automatically block that address so other devices cannot connect to that bot network? That's why many EAP vendors also offer network security devices, or at least partner with those players to offer an integrated experience.

- **Security Monitoring/Analytics:** An EAP product — especially when using EDR — generates a bunch of telemetry which can be useful within your security monitoring environment. So the ability to send it directly to a SIEM or security analytics program helps leverage it in any analyses you perform.
- **Forensics/Case Management:** If you can foresee a situation where you'll want to prosecute an attacker, you need the ability to integrate with an existing case management system. This helps protect the chain of custody for captured data and allows more sophisticated forensics tools to use endpoint data, to better determine what malware does to devices.
- **Operations Platform:** Finally, we need to highlight potential integration with an IT operations platform, especially as it relates to endpoint hygiene and asset management. An EAP product gathers detailed device data which can be very useful to Operations.

Security is too complicated for any tool to stand on its own, so any EAP offering's ability to send and receive data to and from your other security tools is a key selection criteria.

Summary

Enterprises seem to have finally concluded that existing Endpoint Protection Platforms (EPP) don't really protect their endpoints sufficiently. We feel that epiphany is better late than never. But we suspect the catalyst for this realization was that the new generation of tools simply does a better job.

The Endpoint Advanced Protection (EAP) concept entails integration of many capabilities previously only offered separately, including endpoint hygiene to reduce attack surface, prevention of advanced attacks including memory attacks and malware-less approaches, and much more granular collection and analysis of endpoint telemetry ('EDR' technology). The availability of much more detailed data and far better analytics to identify attack patterns have clearly benefited organizations' ability to protect their devices.

But endpoint protection cannot stand alone. Leveraging a broader threat intelligence function in use by other controls can help identify compromised devices which were attacked by malware you didn't know was malicious at the time (retrospection). You will also want to focus on integration with network-based controls and security monitoring environments to ensure you can block attacks earlier in the attack lifecycle, to shorten the window between compromise and detection.

Good progress has been made in protecting endpoints. We see legitimate alternatives to the ineffective EPP products which have been holding organizations hostage for years. But before jumping in with both feet you need to test the tool, plan and stage your migration, and implement a risk-based approach to protecting endpoints.

Of course adversaries will continue innovating, trying to stay a few steps ahead of defenders. But we are hopeful that these new EAP capabilities are balancing the scales, enabling organizations to focus on more strategic security projects instead of just reimagining the same devices day after day.

If you have any questions on this topic, or want to discuss your situation specifically, feel free to send us a note at info@securosis.com.

About the Analyst

Mike Rothman, Analyst and President

Mike's bold perspectives and irreverent style are invaluable as companies determine effective strategies to grapple with the dynamic security threatscape. Mike specializes in the sexy aspects of security — such as protecting networks and endpoints, security management, and compliance. Mike is one of the most sought-after speakers and commentators in the security business, and brings a deep background in information security. After 20 years in and around security, he's one of the guys who “knows where the bodies are buried” in the space.

Starting his career as a programmer and networking consultant, Mike joined META Group in 1993 and spearheaded META's initial foray into information security research. Mike left META in 1998 to found SHYM Technology, a pioneer in the PKI software market, and then held executive roles at CipherTrust and TruSecure. After getting fed up with vendor life, Mike started Security Incite in 2006 to provide a voice of reason in an over-hyped yet underwhelming security industry. After taking a short detour as Senior VP, Strategy at eIQnetworks to chase shiny objects in security and compliance management, Mike joined Securosis with a rejuvenated cynicism about the state of security and what it takes to survive as a security professional.

Mike published [The Pragmatic CSO](http://www.pragmaticcso.com/) <<http://www.pragmaticcso.com/>> in 2007 to introduce technically oriented security professionals to the nuances of what is required to be a senior security professional. He also possesses a very expensive engineering degree in Operations Research and Industrial Engineering from Cornell University. His folks are overjoyed that he uses literally zero percent of his education on a daily basis. He can be reached at mrothman (at) securosis (dot) com.

About Securosis

Securosis, LLC is an independent research and analysis firm dedicated to thought leadership, objectivity, and transparency. Our analysts have all held executive level positions and are dedicated to providing high-value, pragmatic advisory services. Our services include:

- **Primary research publishing:** We currently release the vast majority of our research for free through our blog, and archive it in our Research Library. Most of these research documents can be sponsored for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements and conform to our Totally Transparent Research policy.
- **Research products and strategic advisory services for end users:** Securosis will be introducing a line of research products and inquiry-based subscription services designed to assist end user organizations in accelerating project and program success. Additional advisory projects are also available, including product selection assistance, technology and architecture strategy, education, security management evaluations, and risk assessment.
- **Retainer services for vendors:** Although we will accept briefings from anyone, some vendors opt for a tighter, ongoing relationship. We offer a number of flexible retainer packages. Services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluation, and merger and acquisition assessment. Even with paid clients, we maintain our strict objectivity and confidentiality requirements. More information on our retainer services (PDF) is available.
- **External speaking and editorial:** Securosis analysts frequently speak at industry events, give online presentations, and write and speak for a variety of publications and media.
- **Other expert services:** Securosis analysts are available for other services as well, including Strategic Advisory Days, Strategy Consulting engagements, and Investor Services. These tend to be customized to meet a client's particular requirements.

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors.

Additionally, Securosis partners with security testing labs to provide unique product evaluations that combine in-depth technical analysis with high-level product, architecture, and market analysis. For more information about Securosis, visit our website: <<http://securosis.com/>>.