Endpoint Defense: Essential Practices

Version 1.4
Released: March 26, 2015

Securosis, L.L.C. 515 E. Carefree Highway Suite #766 Phoenix, AZ 85085 **T** 602-412-3051
info@securosis.com www.securosis.com

Table of Contents

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 and professionally edited.

Special thanks to Chris Pepper for editing and content support.

Copyright

Introduction

The area of security seeing the most increased focus lately is endpoint protection. Once you stop snickering it makes some sense. For years (or decades, depending on how cynical you want to be) endpoint security was the primary beneficiary of the compliance driver. Whether the technology actually protected anything was beside the point. Assessors would show up, and you needed to have AV. Then advanced attackers happened and the industry started innovating, starting with network security, leaving endpoints largely unprotected.

But that is no longer a defensible strategy. Endpoints are more often untethered than not, so these devices are no longer within the corporate perimeter. You could route all traffic through your corporate network, but that loses most of the benefit of the cloud and the Internet. We have seen an endpoint security renaissance of sorts, with lots of interesting technologies designed to protect endpoints. We covered many of these developments in our Advanced Endpoint and Server Protection paper.

But the fact remains that many organizations are not even prepared to deal with *unsophisticated* attackers. You know, that dude in the basement banging on your stuff with Metasploit. Those organizations don't really need advanced security now — their requirements are more basic. It's about understanding what really needs to get done — not the hot topic at industry conferences. They cannot do everything to fully protect endpoints, so they need to start with essentials.

But the fact remains that many organizations are not even prepared to deal with unsophisticated attackers. They cannot do everything to fully protect endpoints, so they need to start with essentials.

This paper is all about these Essential Practices of Endpoint Defense.

Securing Endpoints Is Hard

Why is this still a discussion? Endpoints have been around for decades, and organizations have spent tens of billions of {name your favorite currency} to protect these devices. But every minute more devices are compromised, breaches result, and your Board of Directors wants an explanation of why this keeps happening. Two issues underlie the difficulties of endpoint protection. First, let's be candid. It's a software issue: software has defects, which attackers exploit. Second, employees routinely fall for simple social engineering attacks, resulting in a software install or clicked link — the beginning of a successful attack.

And you are a target, whatever of the size of your organization. You have *something* someone else wants to steal, and they will try. To complicate the situation, adversaries continue to automate reconnaissance and attack efforts. You are not protected by resource constraints — the entire Internet can be scanned for common vulnerabilities daily.

The status quo doesn't work for our side. We need to take a step back and look at protecting endpoints with fresh eyes. This provides an opportunity to determine what's really essential.

Defending Endpoints

There are two major aspects to defending endpoints: hygiene and threat management. They are interdependent — you cannot address either alone and expect your endpoints to be protected.

Endpoint Defense Yin Yang

- **Endpoint Hygiene:** The operational aspects of reducing device attack surface are integral to endpoint security strategy. You need to ensure you have sufficient capabilities to manage patches and enforce security configuration policies. Additionally you should ensure employees have the least privilege necessary on each device, to prevent privilege escalation; and lock down device ports.
- **Endpoint Threat Management:** Advanced attackers are only as advanced as they need to be: they take the path of least resistance. But the converse is also true. When these adversaries need advanced techniques they use them. Traditional malware defenses such as antivirus don't stand much chance against a zero-day attack. An effective threat management process must incorporate people, processes, and technology.

Now let's dig into both aspects of endpoint defense.

Endpoint Hygiene

Consistent and effective hygiene practices are difficult to master, both personally (object lesson: your dentist's fancy car, paid for by your failure to floss) and within security. It is not a lack of desire — everyone wants their devices to be safe from compromise. The challenge is operational excellence. To be clear, effective hygiene cannot completely protect endpoints, but it does make them much harder targets.

It is not a lack of desire — everyone wants their devices to be safe from compromise. The challenge is operational excellence.

The essential practices we lump under hygiene include:

- Patch Management
- Configuration Management
- Device Control
- Least Privilege
- Patch Management

Patch managers install fixes from software vendors to address vulnerabilities. The most well-known patching process is Microsoft's monthly Patch Tuesday, when the company issues a variety of software fixes to address defects in its products — many of which could result in system exploitation. Other vendors have adopted similar approaches, combining a periodic patch cycle with out-of-cycle patches for more serious issues. Once a patch is issued your organization needs to assess it, figure out which devices need to be patched, and install it within the window specified by policy — typically a few days. A patch management product scans devices, installs patches, and reports on the success

or failure of the process. Our Patch Management Quant[1] research provides a detailed view of the patching process, so refer to it for more information.

Configuration Management

Configuration management enables an organization to define an authorized set of configurations for devices. Configurations can control pretty much everything that happens on a device, including: applications installed, device settings, running services, and on-device security controls. Another aspect of configuration management is assessment of configurations and identification of changes, useful because unauthorized configuration changes may indicate malware execution or an exploitable operational error. Additionally, configuration management can help ease the provisioning burden of setting up and re-imaging devices after infection, making it more palatable to nuke compromised devices to ensure total cleanup of a compromise.

Device Control

End users love the flexibility USB ports provide for 'productivity'. Unfortunately USB doesn't just enable employees to share music with their buddies — it also lets them download your entire customer database onto their phones. It all became much easier once the industry standardized on USB twenty years ago. The ability to easily share data facilitated employee collaboration, while also greatly increasing the risks of data leakage and malware proliferation. Device control technology enables you to enforce policy — both *who* can use USB ports and *how* — and capture whatever is copied to and from USB devices. As an active control, monitoring and controlling device usage addresses a major risk.

Least Privilege

Employees don't mean to mess up their devices, for the most part. But allowing them to install software, use new devices like printers, and change endpoint configurations can lead to device exploitation. So eliminating device owners' ability to manage devices can dramatically reduce attack surface. That said, a lot of endpoint changes are legitimate, so a key aspect of implementing least privilege is ensuring there is a clear process to allow employees to do their jobs. For instance trusted employees might be able to get a 24-hour grace period for a change, while less sophisticated employees may need to run through an approval process before installing new software.

Employees don't mean to mess up their devices, for the most part. But allowing them to install software, use new devices like printers, and change endpoint configurations can lead to device exploitation.

Endpoint Threat Management

We define threat management within the context of dealing with an attack, as a subset of a larger security program — typically the most visible capability. Now let's consider its components.

Assessment

You cannot protect what you don't know about — that hasn't changed and is not about to. So the first step is to gain visibility into all devices, data stores, and applications that present risk to your environment. Additionally you need to understand the security posture of anything you have to protect.

You need to know what you have, how vulnerable it is, and how exposed it is. With this information you can prioritize your exposure and design a set of security controls to protect your assets.

---

[1] https://securosis.com/Research/Publication/project-quant-survey-results-and-analysis

- **Mission Assessment:** As we described in our CISO's Guide to Advanced Attackers[2], you need to understand what attackers will try to access in your environment, and why. We call this Mission Assessment, and it involves figuring out what's important in your environment.
- **Discovery:** This process finds the endpoints and servers on your network and makes sure everything is accounted for. It includes an ongoing discovery process to shorten the window between something popping up on your network, you discovering it, and figuring out whether it has been compromised.
- **Determine Security Posture:** 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 the devices you discovered. There are many aspects to vulnerability scanning at the endpoint, server, and application layers. Check out our Vulnerability Management Evolution[3] research to understand how a vulnerability management platform can help prioritize operational security.

It may not be as sexy as a shiny malware sandbox or advanced detection technology, but these assessment tasks are necessary before you can even start thinking about building a set of controls to prevent attacks. Assessment needs to happen on an ongoing basis because your technology environment is dynamic, and the attacks you see are subject to change as well — sometimes daily.

Prevention

Next you try to prevent attacks from succeeding. This is where most of the effort in security has gone over the past decade, with mixed (okay, lousy) results. A number of new tactics and techniques are modestly increasing effectiveness, but the plain fact is that you *cannot* prevent every attack. It has become a question of reducing your attack surface as much as practical.

- **Traditional Signatures:** Signature-based controls are all about maintaining a huge blacklist of known malicious files to prevent from executing.
- **Advanced Heuristics:** You cannot depend on matching what a file looks like so you need to pay close attention to what it does, and profile typical patterns of successful attacks. These patterns are behind the advanced heuristics used to detect malware.
- **Application Control/Whitelisting:** Application control implies a default deny posture. You define a set of authorized executables that can run on each device, and block everything else. With a strong policy application control provides true device lockdown — no executables (either malicious or legitimate) can execute without explicit authorization. Check out our Application Control[4] research for a lot more detail on this approach.
- **Isolation:** In addition to better profiling malware and searching for indicators of compromise, another prevention technique growing in popularity is isolation of executables from the rest of the device, by running them in a sandbox. The idea is to spin up a walled garden for a limited set of applications, to shield the rest of the device from anything bad happening within those frequently targeted applications.

Now we need to tell you a hard truth. You cannot block all attacks. Adversaries have gotten much better, attack surface has increased dramatically, and you are not going to prevent every attack. Pwnage happens, so what you do next is critical — both to protecting critical information in your environment, and to your success as a security professional.

---

[2] https://securosis.com/research/publication/the-cisos-guide-to-advanced-attackers

[3] https://securosis.com/research/publication/vulnerability-management-evolution-from-tactical-scanner-to-strategic-platf

[4] https://securosis.com/research/publication/reducing-attack-surface-with-application-control

Detection

There are a number of different detection options — most based on watching for patterns which indicate a compromised device. The key is to shorten the time between when the device is compromised and when you *discover* it has been compromised.

In the broader sense, detection needs to include finding attacks you missed during execution because:

1. *You didn't know it was malware at the time* — This happens frequently, especially given how quickly attackers innovate. Advanced attackers have stockpiles of unknown exploits (0-days) which they use as needed. So your prevention technology could be working exactly as designed but still fail to recognize an attack. There is no shame in that, and is a reality of doing security.
2. *The prevention technology missed the attack* — This is common because advanced adversaries specialize in evading known preventative controls.

So how can you detect after compromise? Monitor other data sources for indicators that a device has been compromised. Very few organizations have the dubious distinction of being first to see a new 'advanced' attack, so you should be able to look for emerging attack indicators, IP and file reputation, etc. as a basis for detecting attacks. This kind of "threat intelligence" enables you to benefit from the misfortune of others, by looking for attacks you haven't seen yet.

Investigation

Once you identify a potentially compromised device you need to verify your suspicion. Verification involves scrutinizing what the endpoint has done recently for indicators of compromise, or other activity that confirms a successful attack. This typically involves a formal investigation — including a structured process to gather forensic data from devices, triage to determine the root cause, and a search to determine how widely the attack spread within your environment.

Verification involves scrutinizing what the endpoint has done recently for indicators of compromise, or other activity that confirms a successful attack. This typically involves a formal investigation

- **Data Capture:** To have the ability to thoroughly investigate a device, you need to systematically capture what's happening on all endpoints and servers at a very granular level and keep that data for as long as you can. Then you have the activity history of the device when you need to dig into it in order to find the compromise. This includes file activity, registry changes, privilege escalation, executed programs, network activity, and a variety of other activity.
- **Analytics:** Endpoints and servers generate a huge amount of data, so products needs to perform Big Data style analysis on telemetry data to identify patterns and develop relationships across data sources. Having the data is the first step. Supplementing it with external information to prioritize focus is second. Analyzing data to provide useful information to security practitioners and incident responders is the third leg of the device activity monitoring triangle.

Remediation

Once you understand what happened you establish a plan to recover. This might involve cleaning the machine, or more likely reimaging it and starting over again. This step can leverage ongoing hygiene activities (such as patch and configuration management), because you can and should use tools you already have to reimage compromised devices.

It also requires tight integration with the Operations team because most organizations separate out threat management functions from endpoint operations. This means integrating systems and ensuring that the handoffs between the security and Ops teams are well-structured and efficient.

Bringing It All Together

The key to making both sides of endpoint defense work well is a common data model. You should be able to integrate and analyze data about endpoints, without moving between systems or only looking at only half the story (either threat management or hygiene). For example if you detect a known malware file on an endpoint you know has been patched to protect it from that compromise, you can move on to other more pressing concerns.

On the other side of the coin, if a different device has known malware installed and recently escalated privileges (as recorded by policy), you know that's a serious problem; you can immediately quarantine the device by shutting down the network connectivity, then locking down what software it can execute by enforcing a whitelisting policy. Without hygiene and threat management consolidating data into a common view you cannot attain that level of integrated defense.

You do not need to use one solution for everything, but you must be able to integrate data to build a consistent end-to-end view.

You do not need to use one solution for everything, but you must be able to integrate data to build a consistent end-to-end view. This might involve sending data to a separate aggregation platform like a SIEM or security analytics product, or ensuring that both your hygiene and threat management vendors can export data to your integration point.

Summary

Perfectly defending against endpoint attacks is a pipe dream, so organizations need to shift away from ineffective legacy protection technologies and procedures. Endpoint security has two major components: hygiene and threat management. Neither is sufficient alone — you need to implement and ensure the effectiveness of both to adequately defend endpoints. It is tempting to focus on state-of-the-art defenses to protect against advanced attacks, but without a strong foundation to reduce attack surface and ensure endpoint hygiene, your devices may be compromised by unsophisticated attackers.

This is another situation where you need to walk before you can run. Get the essential pieces of the foundation in place, and then you can layer more advanced prevention and detection technologies onto it to protect against more sophisticated adversaries. That isn't what most practitioners want to hear, but it is necessary. If you can't get the basic functions right you have no chance against adversaries who know what they are doing.

About Securosis

About the Author

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. 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. Mike published The Pragmatic CSO <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.

After taking a short detour as Senior VP, Strategy at elQnetworks to chase shiny objects in security and compliance management, Mike joined Securosis in 2010 with a rejuvenated cynicism about the state of security and what it takes to survive as a security professional. 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.

We provide services in four main areas:

- **Publishing and speaking:** Includes publishing independent objective white papers, webcasts, and in-person presentations.
- **Strategic consulting for end users:** Includes product selection assistance, technology and architecture strategy, education, security management evaluations, and security program reviews.
- **Strategic consulting for vendors:** Includes market and product analysis and strategy, technology guidance, product evaluations, and merger and acquisition assessments.
- **Investor consulting:** Technical due diligence services including product and market evaluations, available in conjunction with deep product assessments with our research partners.

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. For more information about Securosis, visit our website: <http://securosis.com/>.