The Endpoint Security Management Buyer's Guide

Version 1.3
Released: September 12, 2012

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

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.

Licensed by Lumension Security, Inc.

Lumension® IT Secured. Success Optimized.™

Lumension Security, Inc., a global leader in endpoint management and security, develops, integrates and markets security software solutions that help businesses protect their vital information and manage critical risk across network and endpoint assets. Lumension enables more than 5,100 customers worldwide to achieve optimal security and IT success by delivering a proven and award-winning solution portfolio that includes Vulnerability Management, Endpoint Protection, Data Protection, Antivirus and Reporting and Compliance offerings. Lumension is known for providing world-class customer support and services 24x7, 365 days a year. Headquartered in Scottsdale, Arizona, Lumension has operations worldwide, including Texas, Florida, Washington D.C., Ireland, Luxembourg, Singapore, the United Kingdom, and Australia. Lumension: IT Secured. Success Optimized.™ More information can be found at www.lumension.com.

Copyright

Table of Contents

The Business Impact of Managing Endpoint Security

Keeping track of 10,000+ of *anything* is a management nightmare. With ongoing compliance oversight and evolving security attacks against vulnerable endpoint devices, getting a handle on managing endpoints becomes more important every day. Complicating matters is the fact that 'endpoints' now include all sorts of devices – including a variety of PCs, mobiles, and even kiosks and other fixed function devices. We detailed our thoughts on endpoint security fundamentals a few years back, and much of that is still very relevant. But we didn't take it to the next logical step: a deeper look at how to buy these technologies.

So we now introduce a new type of research paper, an "Endpoint Security Management Buyer's Guide", focused on helping you understand what features and functions are important – in the four critical areas of patch management, configuration management, device control, and file integrity monitoring. In this paper you won't see any mention of anti-malware. We have done a ton of research on that, including Malware Analysis Quant and Evolving Endpoint Malware Detection, so we will defer on an anti-malware Buyer's Guide for the time being. But now let's talk about the business drivers for endpoint security management.

Keeping track of 10,000+ of *anything* is a management nightmare.

Business Drivers

Regardless of what business you are in, the CIA (confidentiality, integrity, availability) triad is important. For example, if you deal with sophisticated intellectual property confidentiality is likely your primary driver. Or perhaps your organization sells a lot online so downtime is your enemy. Regardless of the business imperative, failing to protect devices with access to your corporate data won't turn out well. Of course there are an infinite number of attacks that can be launched against your company. But we have seen that most attackers go after the low-hanging fruit because it's the easiest way to get what they are looking for.

As we described in our recent Vulnerability Management Evolution research, a huge part of prioritizing operational activities is understanding what's vulnerable and/or configured poorly. But that only tells you *what* needs to get done – someone still has to **do** it. That's where endpoint security management comes into play. Before we get ahead of ourselves, let's dig a bit deeper into the threats and complexities your organization faces.

These so-called “advanced attackers” are only as advanced as they need to be. If you leave the front door open they don’t need to sneak in through the ventilation ducts.

Emerging Attack Vectors

You can’t pick up a technology trade publication without seeing terms like “Advanced Persistent Threat” and “Targeted Attacks”. We generally just laugh at all the attacker hyperbole thrown around by the media. You need to know one simple thing: these so-called “advanced attackers” are only as advanced as they need to be. If you leave the front door open they don’t need to sneak in through the ventilation ducts.

Many successful attacks today are caused by simple operational failures. Whether it’s an inability to patch in a timely fashion, or to maintain secure configurations, far too many people leave the proverbial doors open on their devices. Or attackers target users via sleight-of-hand and social engineering. Employees unknowingly open the doors for attackers – and enable data compromise.

We will not sugarcoat things. Attackers are getting better – and our technologies, processes, and personnel have not kept pace. It is increasingly hard to keep devices protected, so you need to take a different and more creative view of defensive tactics, while ensuring you execute flawlessly — because even the slightest opening provides opportunity for attackers.

Device Sprawl

Remember the good old days, when devices consisted of PCs and a few dumb terminals? Those days are *gone*. Now we have a variety of PC variants running numerous operating systems. Those PCs may be virtualized and they may connect in from anywhere in the world – including networks you do not control. Even better, many employees carry smartphones in their pockets, but ‘smartphones’ are really computers. Don’t forget tablet computers either – each having as much computing power as mainframes had a couple decades ago.

So any set of controls and processes you implement must be consistently enforced across the sprawl of all your devices. Every attack starts with *one* compromised device. More devices means more complexity, and a higher likelihood of something going wrong. Again, you need to execute endpoint security management flawlessly. But you already knew that.

Any set of controls and processes you implement must be consistently enforced across the sprawl of all your devices.

BYOD

As uplifting as dealing with these emerging attack vectors and device sprawl is, we are not done complicating things. The latest hot buzzword is BYOD (bring your own device), which means you need to protect not only corporate computer assets, but employees' personal devices as well. Most folks assume this just means dealing with those pesky Android phones and iPads, but we know many finance folks would just *love* to get all those PCs off the corporate books, and that means you need to eventually support any variety of PC or Mac any employee wants to use.

Of course the controls you put in place need to be consistent, whether your organization or an employee owns a device. The big difference is *granularity* of management. If a corporate device is compromised you just wipe it and move on – you know how hard it is to truly clean a modern malware infection, and how much harder it is to have confidence that it really *is* clean. But what about those pictures of Grandma on an employee's device? What about their personal email and address book? Blow *those* away and the uproar is likely to be much worse than just idling someone for a few hours while they wait to get their work desktop back.

So BYOD requires flawless execution, with an additional layer of granularity you haven't had to worry about before. Good times.

A More Strategic View of Endpoint Security Management

Between emerging malware, device sprawl, and BYOD, you have your work cut out for you. You need a much more strategic view of endpoint security management, as automation and integration become critical to dealing with these problems. That is what we will help you develop in this paper, with ideas on the management lifecycle for endpoint security — and we will talk about answers, not just issues.

The ESM Lifecycle

The world is complex and only getting more so. You need to deal with more devices, mobility, emerging attack vectors, and virtualization, among other things. So you need to graduate from a tactical view of endpoint security.

Thinking about how disparate operations teams manage endpoint security today, you probably have tools to manage change – functions such as patch and configuration management. You also have technology to control use of endpoints, such as device control and file integrity monitoring. So you might have 4 or more different consoles to manage a single endpoint device. We call that *swivel chair management* – you switch between consoles enough to wear out your chair. It's probably worth keeping a can of WD-40 handy to ensure your chair is in tip-top shape.

We call that *swivel chair management* – you switch between consoles enough to wear out your chair. It's probably worth keeping a can of WD-40 handy to ensure your chair is in tip-top shape.

Using all these disparate tools also creates challenges in discovery and reporting. Unless the tools integrate cleanly, if your configuration management system (for instance) detects a new set of instances in your virtualized data center, your patch management offering might not even know to scan them for missing patches. Likewise, if you don't control use of I/O ports (USB) on endpoints and a compromised thumb drive installs malware, you might not know the system files have been replaced unless you are specifically monitoring those files. Obviously, given ongoing constraints in funding, resources, and expertise, finding operational leverage *anywhere* is a corporate imperative.

So it's time to embrace a broader view of Endpoint Security Management and improve integration across the tools in use to fill these gaps. Let's describe endpoint security management — the foundation of an endpoint security management suite— and its component parts, and ultimately how these technologies fit into an enterprise management stack.

The Endpoint Security Management Lifecycle

As analyst types, the only thing we like better than quadrant diagrams are lifecycles. So of course we have an endpoint security management lifecycle. But none of these functions are mutually exclusive. Keep in mind that you can start anywhere, and most organizations already have some technologies in place to address these problems. It has become rare for organizations to manage endpoint security manually.

We push the lifecycle to highlight the importance of looking at endpoint security management **strategically**. A patch management product can solve *part* of the problem, tactically. And the same with each of the other functions. But handling endpoint security management as a *platform* can provide more value in terms of operational and cost leverage, compared to dealing with each function in isolation.

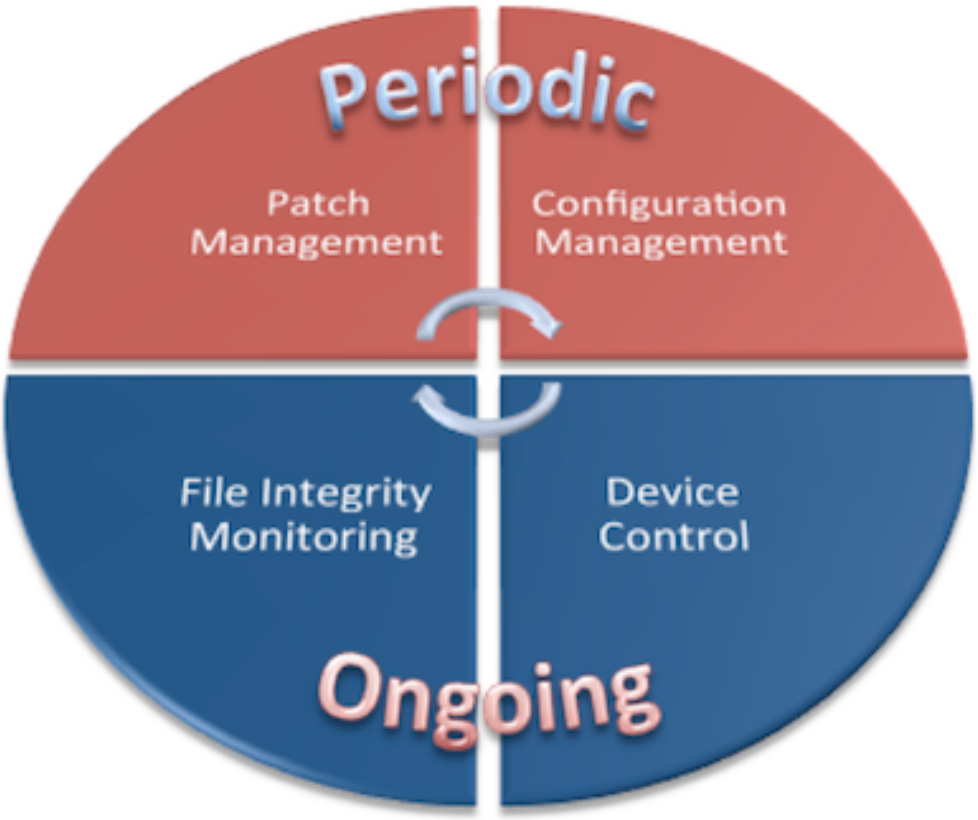

This picture illustrates the lifecycle. It shows periodic functions — patch and configuration management — which typically occur every day or two. It also shows ongoing activities — device control and file integrity monitoring — which need to run all the time – typically using agents deployed on each device.

Let's describe each part at a high level, and then dig into each function throughout the rest of this paper.

Configuration Management

Configuration management provides the ability for an organization to define an authorized set of configurations for devices in use within the environment. These configurations govern the applications installed, device settings, services running, and security controls in place. This capability is important because a changing configuration might indicate malware manipulation or an operational error causing an unauthorized change. Additionally configuration management can help ease the provisioning burden of setting up and/or reimaging devices. Configuration management enables your organization to define what *should* be running on each device based on entitlements, and to identify non-compliant devices.

Patch Management

Patch managers install fixes from software vendors to address software vulnerabilities. The best known patching process comes from Microsoft every month. On Patch Tuesday Microsoft issues a variety of software fixes to address defects that could result in exploitation of their systems. Once a patch is issued your organization needs to assess it, figure out which devices need to be patched, and ultimately install the patch within a window specified by policy – typically a few days. A patch management product scans devices, installs patches, and reports on the success and/or failure of the process. Patch Management Quant provides a very detailed view of the patching process, so check it out if you want more information.

Device Control

End users just love the flexibility their USB ports provide for their 'productivity'. You know – the ability to share music with buddies and download your entire customer database onto their phones – it all got much easier once the industry standardized on USB a decade ago. The ability to easily share data really has facilitated better collaboration between employees, but it also greatly increased the risks of data leakage and malware proliferation. Device control technology enables you to enforce policy for both who can use USB ports and for what; and also to capture what is copied to and from USB devices. As an active control, monitoring and enforcement of device usage eliminates a major risk on endpoint devices.

File Integrity Monitoring

The last control we will mention explicitly is file integrity monitoring, which watches for changes in critical files. Obviously many files do legitimately change over time – particularly during patch cycles. But most files are generally static, and changes to core functions (such as the IP stack and email client) often indicate some type of problem. This active control allows you to define a set of files (including both system and other files), gather a baseline for what they should look like, and watch for changes. Depending on the type of change, you might even *roll back* changes before more bad stuff happens.

The Foundation

The centerpiece of the ESM platform is an asset management capability and console to define policies, analyze data, and report. A platform should have the following capabilities:

The centerpiece of the ESM platform is an asset management capability and console to define policies, analyze data, and report.

- **Asset Management & Discovery:** Of course you can't manage what you can't see, so the first critical capability of an ESM platform is sophisticated discovery. When a new device appears on the network the ESM platform needs to know about it. That may happen via scanning the organization's IP address ranges, passively monitoring traffic, or integrating with other asset management repositories (CMDB, vulnerability management, etc). Regardless of how the platform is populated, without a current list of assets you cannot manage endpoint security.

- **Policy Interface:** The next key capability of the ESM platform is its ability to set policies – a very broad requirement. You must be able to set standard configurations, patch windows, device entitlements, etc., for groups of devices and users. Obviously you need to balance policy granularity against ease of use, but without an integrated policy encompassing all the platform's capabilities, you are stuck in your swivel chair.
- **Analytics:** Once policies are defined you need to analyze and alert on them. The key here is the ability to set rules and triggers across all functions of the ESM platform. For instance, if a configuration change occurs shortly after a patch fails, followed by a system file being changed, that might indicate a malware infection. We aren't talking about the sophisticated multivariate analysis available in enterprise-class SIEMs – just the ability to set alerts based on common attacks you are likely to see.
- **Reporting:** You just cannot get around compliance. Many security projects receive funding and resources from the compliance budget, which means your ESM platform needs to report on what is happening. This isn't novel, but it is important. The sooner you can provide the information to make your auditor happy, the sooner you can get back to the rest of your job.

The key is to look for integration across asset management, policies, analytics, and reporting, to provide the operational leverage you need.

Enterprise Integration Points

No platform really stands alone in an organization. You already have plenty of technology in place, and anything you buy to manage endpoint security needs to play nice with your other stuff. So keep the other enterprise management systems in mind as you look at ESM. Make sure your vendor can provide sufficient integration, or at a minimum an SDK or API to pull data from it for other systems.

- Operations Management – including device building/provisioning, software distribution/licensing, and other asset repositories
- Vulnerability Management – for discovery, vulnerabilities, and patch levels
- Endpoint Protection – including anti-malware and full disk encryption, potentially leveraging agents to simplify management and minimize performance impact
- SIEM/Log Management – for robust data aggregation, correlation, alerting, and reporting
- Backup/Recovery – many endpoints house valuable data, so make sure device failure doesn't risk intellectual property

Next we will dig into the periodic functions: patch and configuration management.

Periodic Controls: Patch Management

When Microsoft got religion about the security issues in Windows XP about a decade ago, they created a wide-ranging process called Trustworthy Computing to restore confidence in the integrity of the Windows operating system. That initiative included a monthly patch cycle to fix software defects that could cause security issues. Patch Tuesday was born, and since then almost every company in the world has had to patch every month.

Over the past decade, many software companies have instituted similar patch processes across many different applications and operating systems. Some vendors are trying to simplify the process by moving to a silent install process to avoid requiring any effort of customer organizations. But most security and operations personnel don't feel comfortable without control over what gets installed and when. So organizations needed to look beyond tactical software updates, handling patching as an operational discipline. Once a patch is issued each organization needs to assess it, figure out which devices need to be patched, and ultimately install the patch within the window specified by its policy – typically a few days.

Once a patch is issued each organization needs to assess it, figure out which devices need to be patched, and ultimately install the patch within the window specified by its policy – typically a few days.

Patching Process

Patching is an operational discipline, so an organization's patching process must first be defined and then automated appropriately. We documented a patch process in Patch Management Quant, and if you are looking for an over-arching process for all your patching we suggest you start there. That process map is detailed and highly granular – just use the parts that make sense in your environment.

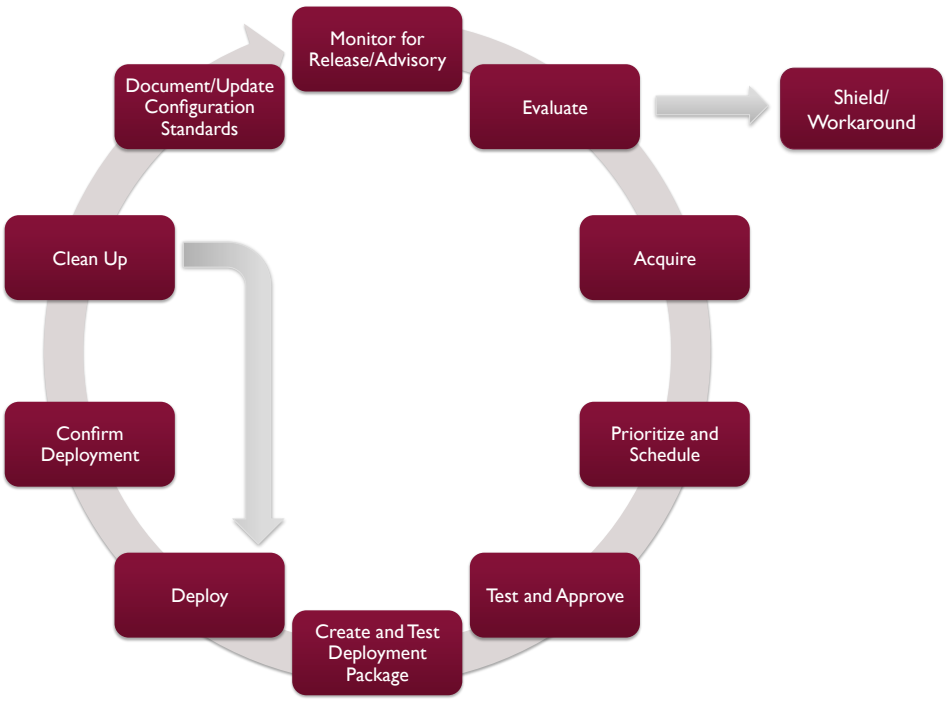

For simplicity's sake, we've recategorized the more details process as follows:

1. **Define targets:** Before you jump into the Patch Management process you need to define which devices will be included. Is it just endpoints, or do you also need to patch servers? These days you also need to think about cloud instances. The technology is largely the same, but increased numbers of devices have made execution more challenging. In this paper we largely restrict our discussion to endpoints, as server operations are different and more complicated.

2. **Obtain patches:** You need to monitor for release of relevant patches, and then figure out whether you need to patch or you can work around the issue.

3. **Prepare to patch:** Once the patch is obtained, you need to figure out how critical the issue is. Is it something you need to fix right now? Can it wait for the next maintenance window? Once priority is established, give the patch a final Q/A check to ensure it won't break anything important.

4. **Deploy the patch:** Once preparation is complete and your window has arrived, you can install.

5. **Confirm the patch:** Patches don't help if the install fails, so confirm that each patch is fully installed.

6. **Reporting:** Compliance requirements for timely patching make reporting an integral function.

Technology Considerations

The good news is that the tools (products and services) to automate patch management are reasonably mature and work fairly well. But there are important technology considerations to keep in mind:

The good news about transforming a function from a security problem to an operational discipline is that the tools (products and services) to automate operational disciplines are reasonably mature and work fairly well.

- **Coverage (OS and apps):** Obviously your patch management offering needs to support the operating systems and applications you need to keep current.
- **Discovery:** You can't patch what you don't know about, so you must ensure you have a way to identify new devices and get rid of deprecated devices – otherwise the process will fail. You can achieve this with a built-in discovery capability, bidirectional integration with asset management and inventory software, or (more likely) both.
- **Library of patches:** Another facet of coverage is accuracy and support of operating systems and applications. Just because something is 'supported' on a vendor's data sheet doesn't mean they support it *well*. So make sure to test the vendor's patch library and check on the timeliness of their updates. How long do they take to package and deploy the patches to their customers after a patch is released? Remember that the clock is ticking if a weaponized exploit is available as a result of the patch.
- **Reliable Deployment of patches:** If patches don't install consistently, which involves the updating, adding and/or removing software that makes more work for you. This can easily make a tool more trouble than it's worth.
- **Agent vs. agentless:** Does the patch vendor assess device via an agent, or do they perform an 'agentless' scan (typically using a non-persistent or 'dissolvable' agent), and then how to do they deploy patches? This is almost a religious dispute, but fortunately both models work. Patching is a periodic control so both models are valid.
- **Remote devices:** How does the patching process work for remote devices? This could be a field employee's laptop or a device in a remote location with limited bandwidth. What kind of features are built in to ensure the right patches get deployed regardless of location? And finally, can you be alerted when a device hasn't updated within a configurable window – perhaps because it hasn't connected?
- **Deployment architecture:** Some patches are hundreds of megabytes, so it is important to have some flexibility in patch distribution – especially for remote devices and locations. Architectures may include intermediate patch distribution points to minimize network bandwidth, and/or intelligent patch packaging to only install appropriate patches on each device.

- **Scheduling flexibility:** Of course it's essential that disruptive patching not impair productivity, so you should be able to schedule patches during off hours and when machines are idle.
- **Value-Add:** As you consider a patch management tool, make sure you fully understand the tool's value-add – what distinguishes it from low-end and low-cost (free) operating system based tools, such as Microsoft's WSUS. Make sure the tool supports your process and provides the capabilities you need.

Periodic Controls: Configuration Management

As we said earlier when discussing the ESM lifecycle: *Configuration Management provides the ability for an organization to define an authorized set of configurations for devices in use within the environment.* These configurations govern installed applications, device settings, running services, and security controls. Where patch management focuses on addressing vulnerabilities, configuration management protects against inadvertent changes (typically during operational updates or reconfiguration) by defining entitlements for what should run on each device, and identifies non-compliant devices. Configuration management also detects and remediates changes made by malicious software.

Configuration management protects against inadvertent changes (typically during operational updates or reconfiguration) by defining entitlements for what should run on each device, and identifies non-compliant devices.

Configuration Management Process

We haven't documented the configuration management process as thoroughly as patching, but let's take a high-level look to see how it works.

- **Establish configuration baselines and/or benchmarks:** First define acceptable secure configurations for each managed device type. Many organizations start with the benchmarks from CIS or NIST (PDF) for granular guidance on how devices should be configured.

- **Discovery:** Next find the devices that need to be managed. Ideally you can leverage an endpoint security management platform with an integrated asset management repository. You will also want to categorize and group assets to avoid unnecessary services. Engineering workstations, for example, require different configurations than Finance systems.

- **Assess, alert, and report changes:** Once devices are discovered and categorized, define a frequency for assessments. How often will you check them against policy? Some vendors use the term "continuous assessment", but their assessments aren't really continuous. Fortunately this isn't normally a problem – not least because most operational groups wouldn't be able to validate alerts and perform corrections in real time anyway.

- **Remediate:** Once a problem is identified, either it needs to be fixed or someone needs to grant an exception. You are likely to have too much work to handle it all immediately, so prioritization is a key success criterion. We offered some perspective on prioritization for vulnerability management, but the concepts are the same for configuration management. You will also probably need to verify that changes actually took place for the audit.

Technology Considerations

As with patching, configuration management tools have been around for a while and are reasonably mature. There are many similarities to the technology considerations presented above for patch management. There is significant leverage to be gained from a single platform which handles both periodic endpoint security management functions.

There is significant leverage to be gained from a single platform which handles both periodic endpoint security management functions.

- **Coverage (OS and apps):** Of course your configuration management offering must support your operating systems.
- **Discovery:** You can't manage configurations you don't know about, so you need to ensure you have a way to identify new devices. You also don't want to clutter your environment and should purge deprecated devices. This can be managed through a built-in discovery capability, bidirectional integration with asset management and inventory software, or more likely both.
- **Supported standards and benchmarks:** The more built-in standards and/or configuration benchmarks offered by the tool, the better your chance of finding something you can easily adapt to your own requirements.
- **Policy editing:** Policies generally require customization to satisfy your requirements. Your configuration management tool should offer a flexible policy editor to define policies and add new baseline configurations and/or benchmarks.
- **Scalability:** Scanning each device for configuration changes can be demanding of both the endpoints and the network, so understand how to distribute scanners effectively and make sure scanning frequency is flexible.
- **Dealing with remote devices:** How does assessment work for a remote device? This could be a field employee's laptop or a device in a remote location with limited bandwidth. What kind of recovery features are built in to ensure the correct remediations are implemented regardless of location? And finally, can you be alerted of devices which haven't been assessed recently – perhaps because they haven't connected?

- **Agent vs. agentless:** Does the configuration management vendor use an agent to perform assessment, or do they perform 'agentless' scans (typically using a non-persistent 'dissolvable' agent), and then how do they apply changes? This borders on religion, but fortunately both models work. Configuration management is a periodic control, so either model is valid here.
- **Integration with operational process:** Make sure any identified configuration issues are reported to the central help desk system to close the operational loop, to ensure a proper process for authorizing and applying changes. This may be managed within the endpoint security management platform, but integration with enterprise systems can make things easier.
- **Process to deal with policy exceptions:** As mentioned above, there may be situations where a configuration change represents an authorized exception. To make things more fun, authorization is often granted *after* configuration management detects (and perhaps reverses) the change. You must be able to handle these situations and without bogus alerts every time the device is assessed.
- **Value-Add:** As you consider a configuration management tool, make sure you fully understand the tool's value-add – what distinguishes it from low-end and low-cost (free) operating system based tools, such as Microsoft's SCCM. Make sure the tool supports your process and provides the capabilities you need.

Periodic Controls: Other Considerations

More Food for Thought

Despite the maturity of patch and configuration management, a number of emerging technologies impact these processes. For instance some organizations are considering virtual desktops (VDI) to improve endpoint management. You need to understand how your tool supports VDI and which virtualization infrastructure products it can manage. The goal is to avoid a separate management capability for the virtual environment. Cloud instances are similar but generally handled within the server context. We could write an entire paper just digging into that topic – and probably will – but for now we'll just note that managing cloud *devices* adds a layer of complexity for patching and configuration management.

We also need to recognize the impact of mobile devices on endpoint security management. For now mobile device management (MDM) offerings have emerged to address the management need, but we expect to gain leverage from managing *all* devices (both PC and post-PC) from a single platform.

Finally, consider alternative deployment models – which may include Software as a Service (SaaS). Obviously to assess and patch internal devices you need some kind of internal device (usually a virtual machine or physical appliance) as the inside control point, communicating results and taking direction from the cloud. This is another semi-religious decision, and we still see a majority of organizations deploying onsite platforms for endpoint security management. But that seems likely to change over the next few years.

Initial Buying Considerations

Patch and configuration management tools are mature, so how should you choose? We will go into detail later, but it depends on who is responsible for patching and configuration management. If it's Operations, an operations-oriented platform with broader data center and server management capabilities is probably the way to go. On the other hand, if the endpoint/device team is responsible, a tool or platform optimized for endpoints makes sense. If auditors are the driver, focus on assessment capabilities for validation and reporting. If different teams handle different functions, an integrated platform may not offer significant leverage.

With mature technologies, products rarely differ radically from each other. There are *always* differences in user experience and marginal features, but primary feature sets are consistent. First decide how you want to work, and then find a tool or platform to automate it. That's why we start with process and *then* automate as appropriate.

Ongoing Controls: Device Control

Device control technology provides the ability to enforce policy on what users can and can't do with devices. That includes locking down ports to prevent copying data (primarily via removable media), as well as ensuring any data allowed onto removable media is encrypted. Early in this technology's adoption cycle, we joked that the alternative to device control software was supergluing the USB ports shut. Obviously superglue doesn't provide sufficient granularity in the face of employees' increasing need to collaborate and share data using removable media, but it would at least prevent many breaches.

Device control technology provides the ability to enforce policy on what users can and can't do with devices.

So let's get a bit more specific with device control use cases:

- **Data leakage:** You want to prevent users from connecting their phones or USB sticks and copying your customer database. You would also like to allow them to connect USB sticks, but not copy email or databases, or perhaps limit them to copying a limited amount of data. Don't let your intellectual property escape on removable media.

- **Encryption:** Obviously there are real business needs for USB ports, or else we could all have saved serious money with superglue. If you need to support moving data to removable media, make sure it's encrypted. If you think losing a phone is easy, USB sticks are even easier – and if one has unencrypted and unprotected sensitive data, you will get a chance to dust off your customer notification process.

- **Malware proliferation:** The final use case gets back to the future. Remember how the first computer viruses spread via floppy disks? Back in the day sneakernet was a big problem, and this generation's sneakernet is the found USB stick that happens to carry malware. You will want to protect against that attack without the superglue.

Device Control Process

Implementing technology controls for endpoint security management without the proper underlying processes never works well, so let's offer a reasonable device control process:

1. **Define target devices:** Which devices pose a risk to your environment? Probably not all of them, so start by figuring out which devices need to be protected.

2. **Build threat models:** Next put on your attacker hat and figure out how those devices are likely to be attacked. Are you worried about data leakage? Malware? Build models to represent how *you* would attack your environment. Then take your threat models to the next level. Maybe the marketing folks should be able to share large files via their devices, but folks in engineering (with access to source code) shouldn't. You can get pretty granular with your policies, so you can do the same with threat models.

3. **Define policies:** With the threat models you can define policies. Any technology you select should be able to support the policies you need.

4. **Discovery:** Yes, you will need to keep an eye on your environment, checking for new devices and managing the devices you already know about. There is no reason to reinvent the wheel, so you are likely to rely on an existing asset repository (within the endpoint security management platform, or perhaps a CMDB).

5. **Enforcement:** Now we get to the operational part of endpoint security management: deploying agents and enforcing policies on devices.

6. **Reporting:** We security folks like to think we implement these controls to protect our environments, but don't forget that at least some of our tools are funded by compliance. So we need some reports to demonstrate that we're protecting data and compliant.

Technology Considerations

Now that you have the process in place, you need some technology to implement the controls. Here are some things to think about when looking at these tools:

- **Device support:** Obviously the first order of business is to make sure the vendor supports the devices you need to protect. That means ensuring operating system support, as well as the media types (removable storage, CDs & DVDs, tape drives, printers, etc.) you want to define policies for. And make sure the product supports all the ports on your devices, including USB, FireWire, serial, parallel, and Bluetooth. Some offerings can also implement policies on data sent via the network driver, though that begins to blur into endpoint DLP, which we will discuss later.

- **Policy granularly:** Make sure your product can support different policies by device. For example, this allows you to set a policy to let an employee download any data to an IronKey but only non-critical data onto an iPhone. You will also want to be able to set up different policies for different classes of users and groups, as well as by type of data (email vs. spreadsheets vs. databases). You may want to limit the

amount of data that can be copied by some users. This list isn't exhaustive, but make sure your product supports the policies you need.

- **Encryption algorithm support:** If you are going to encrypt data on removable media, make sure your product supports your preferred encryption algorithms and/or hooks into your central key management environment. You may also be interested in certifications such as EAL (Common Criteria), FIPS 140-2, etc.

- **Small footprint, secure agent:** To implement device control you will need an agent on each protected device. Besides making sure the agent isn't a pig, stealing massive amounts of compute power from the device, you'll also want to ensure some kind of tamper resistance to protect the agent. You don't want an attacker to turn off or compromise the agent's ability to enforce policies.

Besides making sure the agent isn't a pig, stealing massive amounts of compute power from the device, you'll also want to ensure some kind of tamper resistance to protect the agent.

- **Integration with endpoint security management platforms:** Don't reinvent the wheel – especially for cross-functional capabilities such as discovery, reporting, agentry, and agent deployment/updating/maintenance – so leverage your endpoint security management platform to streamline implementation and for operational leverage.

- **Offline support:** Devices aren't always connected to the network, so make sure policies are still enforced even when disconnected. You will also want to ensure you can configure alerts on policy violations for disconnected devices on reconnection.

- **Forensics:** In the event of data loss you will want forensics, so having the product log all user activity can be quite helpful. Some offerings also keep a copy of any files copied to any protected device ports, which can serve as a smoking gun in the event of data loss.

- **Override or temporary grace period:** Finally, there are times when a policy may simply need to be overridden. Like when your CEO is trying to get a deal done at the end of the quarter and needs to share an agreement with a customer. Having the ability to allow certain employees to override policies (with proper alerting and audit trails) can make deployment work much better.

Endpoint DLP Overlap

There are many areas of overlap between policies you can implement using device control and what endpoint DLP offers, and many device control offerings claim DLP capabilities as well. The good news is that Securosis offers a ton of research on DLP. Check out the following links for more on DLP:

- Understanding and Selecting a DLP Solution

- Low Hanging Fruit: Quick Wins with DLP
- Best Practices for Endpoint DLP

The relevant findings from our DLP research: DLP is really about *deep content analysis,* which offers much more granularity than a typical endpoint device control product. As an example, device control can build a policy based on file type (`.doc` vs. `.pdf`), whereas real endpoint DLP will enable a policy to detect and block mention of a research project within a large `.doc` file. DLP also adds significant capabilities for classification of sensitive data and enforcement beyond endpoints – it extends to the network (data in motion) and on storage devices (data at rest). Additionally, endpoint DLP agents can typically block functions such as Copy & Paste and Print Screen.

Another emerging issue for device control is the fact that “removable media” may not actually be removable anymore.

Endpoint DLP generally does not offer built-in encryption for removable media, nor can it normally assess or protect against malware attacks via removable media. But with the consolidation of many standalone DLP vendors, endpoint DLP agents now tend offer more than just content analysis. So device control vendors are also adding more DLP-style content analysis to their offerings, further blurring the line between product categories.

Finally, another emerging issue for device control is the fact that “removable media” may not actually be removable anymore. Cloud offerings such as Dropbox and Box.net have in many cases replaced removable media. So over time your device control offerings should be able to restrict and/or encrypt data going to cloud services, using the same kinds of policies as protect traditional removable media.

Ongoing Controls: File Integrity Monitoring

After hitting the first of the ongoing controls, device control, we now turn to File Integrity Monitoring (FIM). Also called change monitoring, this means monitoring files to see if and when they change. Here are a few scenarios where FIM is particularly useful:

- **Malware detection:** Malware does many bad things to your devices. It can load software and change configurations and registry settings. But another common technique is to change system files. For instance a compromised IP stack could be installed to direct all your traffic to a server in Eastern Europe, and you might never be the wiser.
- **Unauthorized changes:** These may not be malicious but can still cause serious problems. They can be caused by many things, including operational failure and bad patches, but ill intent is not necessary for exposure.
- **PCI compliance:** Requirement 11.5 in our favorite prescriptive regulatory mandate, the PCI-DSS, requires file integrity monitoring to alert personnel to unauthorized modification of critical system files, configuration files, or content files. So there you have it – you can justify this expenditure with the compliance hammer. But security is about more than checking the compliance box, so we will focus on getting value from the investment as well.

FIM Process

Again we start with a process that can be used to implement file integrity monitoring. Technology controls for endpoint security management don't work well without appropriate supporting processes.

1. **Set policy:** Start by defining your policy, identifying which files on which devices need to be monitored. But there are tens of millions of files in your environment so you need to be pretty savvy to limit monitoring to the most sensitive files on the most sensitive devices.
2. **Baseline files:** Then ensure the files you assess are in a known good state. This may involve evaluating version, creation and modification date, or any other file attribute to provide assurance that the file is legitimate. If you declare something malicious to be normal and allow it, things go downhill quickly. The good news is that FIM vendors have databases of these attributes for billions of known good and bad files, and that intelligence is a key part of their products.

3. **Monitor:** Next actually monitor changes. This is easier said than done because you may see hundreds of file changes on a normal day. So knowing a good change from a bad one is essential. You need a way to minimize false positives from legitimate changes to avoid wasting everyone's time.

4. **Alert:** When an unauthorized change is detected you need to let someone know.

5. **Report:** FIM is required for PCI compliance, and you will likely use that budget to buy it. So you need to substantiate effective usage for your assessor. That means generating reports. Good times.

Technology Considerations

Now that you have the process in place, you need some technology to implement FIM. Here are some things to think about when looking at these tools:

- **Device and application support:** Obviously the first order of business is to make sure the vendor supports the devices and applications you need to protect. We will talk about this more under research and intelligence, below.

- **Policy granularity:** You will want to make sure your product can support different policies by device. For example a POS device in a store (within PCI scope) needs to have certain files under control, while an information kiosk on a segmented Internet-only network in your lobby may not need the same level of oversight. You should also be able to set up those policies based on groups of users and device types (locking down Windows XP tighter, for example, as it lacks newer protections in Windows 7).

> Remember that to take full advantage of continuous monitoring, you must be able to respond immediately to every alert. Do you have 24/7 ops staff ready to pounce on every change notification?

- **Small footprint agent:** In order to implement FIM you will an agent on each protected device. Of course there are different definitions of what an 'agent' is, and whether it needs to be persistent or can be downloaded as needed to check the file system and then removed – a "dissolvable agent". You will need adequate platform support as well as some kind of tamper-proofing of the agent. You don't want an attacker to turn off or otherwise compromise the agent's ability to monitor files – or even worse, to return tampered results.

- **Frequency of monitoring:** Related to the persistent vs. dissolvable agent question, you need to determine whether you require true continuous monitoring of files, or whether batch assessment is acceptable. Before you respond "Duh! Of course we want to monitor files at all times!" remember that to take full advantage of continuous monitoring you must be able to respond immediately to alerts. Do you have 24/7 ops staff ready to pounce on every change notification? No? Then perhaps a batch process

could work. Though keep in mind having more forensic data about exactly when file changes occur can greatly assist with investigations.

- **Research & intelligence:** A large part of successful FIM is knowing a good change from a potentially bad change. That requires some kind of research and intelligence capability to do the legwork. The last thing you want your expensive and resource-constrained operations folks doing is assembling monthly lists of file changes for a patch cycle. Your vendor needs to do that. But it's a bit more complicated, so here are some other notes on detecting bad file changes.
    - **Change detection algorithm:** Is a change detected based on file hash, version, creation date, modification date, or privileges? Or all of the above? Understanding how the vendor determines a file has changed enables you to ensure all your threat models are factored in.
    - **Version control:** Remember that even a legitimate file may not be the right one. Let's say you are updating a system file, but an older legitimate version is installed. Is that a big deal? If the file is vulnerable to an attack it could be, so track file versions and integrate with patch information.
    - **Risk assessment:** It's helpful if the vendor can assess different kinds of changes for potential risk. Replacing the IP stack is a higher-risk change than updating an infrequently used configuration file. Either could be bad but you need to prioritize and a first cut from the vendor can be useful.
- **Forensics:** In the event of a data loss you will want a forensics capability such as a log of all file activity. Knowing when different files were accessed, by what programs, and what was done, can be very helpful for assessing the damage of an attack and nailing down the chain of events which resulted in data loss.
- **Closed loop reconciliation:** Thousands of file adds, deletes, and changes happen every – and most are authorized and legitimate. But for both compliance and operational reliability you should be able to reconcile the changes you *expect* against the changes that actually happened. During a patch cycle, a bunch of changes should have happened. Did all of them complete successfully? We mentioned verification in the patch management process, and FIM technology can provide that reconciliation as well.

Knowing when different files were accessed, by what programs, and what was done, can be very helpful for assessing the damage of an attack and nailing down the chain of events which resulted in data loss.

- **Platform integration:** There is no reason to reinvent the wheel – especially for cross-functional capabilities such as discovery, reporting, agentry, and agent deployment/updating/maintenance – so leverage your endpoint security management platform to streamline implementation and facilitates operations.

Platform Buying Considerations

We have alluded to the *platform* throughout this paper, but what exactly does that mean? What do you need it to do?

Platform Selection

As with most other technology categories (at least in security), the management console (or 'platform', as we like to call it) connects the sensors, agents, appliances, and any other security controls. Let's list the platform capabilities you need.

The management console (or 'platform', as we like to call it) connects the sensors, agents, appliances, and any other security controls.

- **Dashboard:** You will want user-selectable elements and defaults for technical and non-technical users. You should be able to only show certain elements, policies, and/or alerts to authorized users or groups, with entitlements typically stored in the enterprise directory. Nowadays with the state of widget-based interface design, you can expect a highly customizable environment, letting each user configure what they need and how they want to see it.
- **Discovery:** You can't protect an endpoint (or any other device, for that matter) if you don't know it exists. So once you get past the dashboard, the first key feature of the platform is discovery. Surprise is the enemy of the security professional, so make sure you know about new devices as quickly as possible – including mobile devices.
- **Asset repository integration:** Closely related to discovery is the ability to integrate with an enterprise asset management system/CMDB to get a heads-up whenever a new device is provisioned. This is essential for monitoring and enforcing policies. You can learn about new devices proactively via integration or reactively via discovery. But either way you need to know what's out there.
- **Policy creation and management:** Alerts are driven by the policies you implement in the system, so policy creation and management is also critical. We will delve further into this later.

- **Alert management:** A security team is only as good as its last incident response, so alert management is key. This allows administrators to monitor and manage policy violations which could represent a breach. Time is of the essence during any response, so the abilities to provide deeper detail via drill down and send information into an incident response process are critical. The interface should be concise, customizable, and easy to read at a glance – response time is key. When an administrator drills down into an alert the display should cleanly and concisely summarize the reason for the alert, the policy violated, the user(s) involved, and any other information helpful for assessing the criticality and severity of the situation. We will dig deeper later.

- **System administration:** You can expect the standard system status and administration capabilities within the platform, including user and group administration. Keep in mind that for a larger distributed environment, you will want some kind of role-based access control (RBAC) and hierarchical management to manage access and entitlements for a variety of administrators with varied responsibilities.

- **Reporting:** As we mentioned under specific controls, compliance tends to fund and drive these investments, so substantiating their efficacy is necessary. Look for a mixture of customizable pre-built reports and tools to facilitate *ad hoc* reporting – both at the specific control level and across the entire platform.

In light of the importance of managing your policy base and dealing with the resulting alerts – which could represent attacks and/or breaches – let's go deeper into those functions.

Every environment has its own unique characteristics but the platform vendor should provide out-of-the-box policies to make customization easier and faster.

Policy Creation and Management

Once you know what endpoint devices are out there, assessing their compliance (and remediating as necessary) is how the platform provides value. The resource cost to validate and assess each alert makes filtering relevant alerts critical. So policy creation and management can be the most difficult part of managing endpoint security. The policy creation interface should be accessible to both technical and non-technical users, although creation of heavily customized policies almost always requires technical skill.

For policy creation the system should provide baselines to get you started. For patching you might start with a list of common devices and then configure assessment and patching cycles accordingly. This works for the other controls as well. Every environment has its own unique characteristics but the platform vendor should provide out-of-the-box policies to make customization easier and faster. All policies should be usable as templates for new policies. We are big fans of wizards to walk administrators through this initial setup process, but more sophisticated users need an 'Advanced' tab or equivalent to set up more granular policies for more sophisticated requirements. Not all policies are created equal, so the platform should be able to grade the sensitivity of each alert and support severity thresholds.

Most administrators tend to prefer interfaces that use clear, graphical layouts for policies – preferably with an easy-to-read grid showing the relevant information for each policy. The more complex a policy, the easier it is to create internal discrepancies or accidentally define an incorrect remediation. Remember that every policy needs some level of tuning, and a good tool will enable you to create a policy in test mode to see how it would react in production, *without* firing all sorts of alerts or requiring remediation.

Alert Management

Security folks earn their keep when bad things happen. So you will want all your tools to accelerate and facilitate the triage, investigation, root cause analysis, and manage the process when you respond to alerts. On a day to day basis admins will spend most of their time working through the various alerts generated by the platform.

When assessing the alert management capabilities of any product or service, first evaluate them in terms of supporting your existing processes. Obviously you will need to adapt some process to get the most out of your tools, but you should not need to blow up your existing processes to take advantage of the platform.

Obviously you will need to adapt some process to get the most out of your tools, but you should not need to blow up your existing processes to take advantage of the platform.

In terms of what to look for, the first stop is the alert queue: a summary of all alerts, hopefully with the ability to assign each to a specific analyst for triage. Alert status should be clearly indicated with color-coded sensitivity (based on the policy violated) and severity (based on the volume of the violation or some other factor defined in the policy). Alerts should be sortable or filterable on any field. Policy violated, user, alert status (open, closed, assigned, unassigned, under investigation), and responsible party should also be indicated and easily changeable for instant disposition.

By default, closed alerts shouldn't clutter the interface – you should be able to treat the interface like an email inbox. Each administrator should be able to customize everything to suit his or her own work style. Alerts with either multiple policy violations or multiple violations of a single policy should only appear once in the incident queue to simplify things, but highlighted in some way to show multiple issues caused the alert (which may represent increased urgency).

When digging into an alert the platform should list all the relevant details – including *who, what, where, when,* and *how*. A summary of other recent violations by that user or device, and other violations involving that data (which might indicate a larger event), is particularly useful. The tool should allow administrators to make comments, assign additional resources, notify management, and upload any supporting documentation.

To Cloud or Not to Cloud

Now let's address everyone's favorite buzzword. In this context 'cloud' means SaaS (software as a service), where the vendor manages the infrastructure, which you access via a browser across the Internet. There is a

great deal of hype and religion permeating this discussion, but ultimately the decision is driven by deployment and operational needs. At the end of the day, whether you select an endpoint security management service (cloud) or a product involves two considerations:

1. **Scale:** You will hear a lot from cloud providers about infinite scale and the limitations of customer premise offerings. It is true that scalability is the vendor's problem in a cloud/SaaS scenario. That offers some advantages, but local solutions can scale with a suitable deployment architecture.

2. **Technology updates and change:** The other big message from cloud advocates is that cloud platforms handle software updates more quickly and transparently than onsite gear. Again, there is truth here, but every endpoint security management vendor has been sending new rules and tests to its devices for years, so it's not like they haven't figured out software distribution and updating.

The 'decision' about cloud versus customer-premises isn't really a decision at all – what is and isn't 'cloud' nowadays is largely a question of semantics. Let's get back to your requirements. You need to be able to assess your environment, change and remediate, and install agents as needed. As long as you can meet your requirements, *where* the management console runs isn't really significant.

Regardless of where the console resides, there will be an on-site component – which might be a dedicated appliance, a virtual machine, a dissolvable agent, or some combination. Don't get caught up in hype. Focus on problems you need to solve and the best way to solve them.

Vendor Selection Considerations

Inevitably, after doing your research to figure out which platforms can meet your requirements, you will need to define a short list and ultimately choose something. One of the inevitable decision points involves large versus small vendors. Given the pace of mergers and acquisitions in the security space even small vendors may not remain independent and small forever. Ultimately working with a larger vendor is all about leverage. One type is pricing leverage achieved by buying multiple products and services from the vendor and negotiating a nice discount on all their products. But smaller vendors can get aggressive on pricing as well, and sometimes have even more flexibility to sell cheaper.

Another type of leverage is platform leverage, which involves using multiple products managed via a single platform. The larger endpoint security management vendors also have more active defense products (such as anti-malware and network security) you might use, and an integrated console can make your life easier.

Given the importance of *intelligence* to understanding patches, configurations, and file integrity, it is important to consider the size and breadth of the vendor's research team and customer base.

Given the importance of *intelligence* to understanding patches, configurations, and file integrity, it is important to consider the size and breadth of the vendor's research team and customer base. Keeping policies current and issuing effective updates (for configuration or file hashes, for example) requires access to a huge dataset and a serious analysis capability to figure out what needs to be done. You will probably hear a lot about *Big Data,* but that's just the buzzword *du jour*. It's more about the vendor making the investment to keep their platform current, given the dynamic nature of the security business.

You will want to ensure the vendor has the ability to support your environment, wherever that is. Local support is best when dealing with endpoints, since you may not have a capable staffer in a remote office to troubleshoot the issue. But as time goes on we will see better collaboration and remote management/troubleshooting tools, making centralized support increasingly viable for a global customer base. Depending on the sophistication of your local IT teams, you may want to ensure you can get deployment assistance from the vendor as well – especially in remote locations, where it could be very expensive to send your own folks to deploy the system.

Purchasing Cycle

In terms of the purchasing cycle, there is no need to reinvent the wheel. Some organizations are formal and issue RFI/RFP (requests for information/proposals) to gather information. Others work with resellers or rely on personal contacts to learn about alternatives and negotiate deals. However you buy products and services, you are likely to go through the same basic process:

1. **Define requirements:** Don't minimize the need to do internal fact finding and requirement gathering before engaging with vendors. Know what you're buying and why. Understand what's working in your environment and what's not. Then you'll know what you need the endpoint security management platform to do.

2. **Establish short list:** This may be a formal or informal process. Ultimately you need to find a handful of vendors who can meet your requirements. Talk to them and dig deeper into their products and services to figure out which vendors can really solve your problems.

3. **Test products:** You will want to set up a testbed and let the tools do their thing. Depending on which controls you are looking to implement, you can run all sorts of tests during a proof of concept. Figuring out the device overhead of any agentry is key, as well as the user experience of setting policies, managing alerts, and remediating issues.

4. **Try support:** Make sure you put a number of calls into the vendor's support group. Both during typical business hours and off-hours to understand how they'll support you when it counts.

5. **Negotiate:** We could write a book about vendor negotiation, but for now suffice it to say leverage is good. Try to negotiate with at least two vendors to get them competing for your business. And don't believe the vendors when they say end of quarter discounts don't happen. Unless the sales rep is way ahead of their quota, they'll deal at the end of the quarter.

We could go much deeper into purchasing – it's a discipline like any other aspect of a security professional's job. But the high level process outlined above should serve you well.

Summary: 10 Questions to Ask Your Endpoint Security Management Vendor

Normally we wrap up each paper with a nice summary that goes through the high points of our research and summarizes what you need to know. But this is a Buyer's Guide, so we figured it would be more useful to summarize with 10 questions. With apologies to Alex Trebek, here are the 10 key questions we would ask if we were buying an endpoint security management product or service.

1. What specific controls do you offer for endpoint management? Can the policies for all controls be managed via your console?
2. Does your organization have an in-house research team? How does their work make your endpoint security management product better?
3. What products, devices, and applications are supported by your endpoint security management offerings?
4. What standards and/or benchmarks are offered out of the box for your configuration management offering?
5. What kind of agentry is required for your products? Is the agent persistent or dissolvable? How are updates distributed to managed devices? How do you ensure agents are not tampered with?
6. How do you handle remote and disconnected devices?
7. What is your plan to extend your offering to mobile devices and/or virtual desktops (VDI)?
8. Where does your management console run? Do we need a dedicated appliance? What kind of hierarchical management does your environment support? How customizable is the management interface?
9. What kind of reports are available out of the box? What is involved in customizing specific reports?

10. What have you done to ensure the security of your endpoint security management platform? Is strong authentication supported? Have you done an application pen test on your console? Does your engineering team use any kind of secure software development process?

Of course we could have written another 10 questions. But these hit the highlights of device and application coverage, research/intelligence, platform consistency/integration, and management console capabilities. This list cannot replace a more comprehensive RFI/RFP, but can give you a quick idea of whether a vendor's product family can meet your requirements.

The one aspect of buying endpoint security management that we haven't really discussed appears in question 5 (agents) and question 10 – the security of the management capability itself. Attacking the management plane is like attacking a bank rather than individual account holders. If the attacker can gain control of the endpoint security management system, then they can apply malicious patches, change configurations, drop or block file integrity monitoring alerts, and allow bulk file transfers to thumb drives. But that's just the beginning of the risks if your management environment is compromised.

Attacking the management plane is like attacking a bank rather than individual account holders.

We focused on the management aspects of endpoint security in this paper, but remember that we are talking about endpoint **security**, which means making sure the environment remains secure – both at the management console and agent levels.

As we have described, endpoint security management is mature technology – so look less at specific feature/capability differentiation and more at policy integration, console leverage, and user experience. That approach will usually yield the most effective solution for your environment.

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 or ask via the Securosis Nexus (http://nexus.securosis.com/).

About the Analyst

Mike Rothman, Analyst/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 elQnetworks 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/> 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:

- **The Securosis Nexus**: The Securosis Nexus is an online environment to help you get your job done better and faster. It provides pragmatic research on security topics that tells you exactly what you need to know, backed with industry-leading expert advice to answer your questions. The Nexus was designed to be fast and easy to use, and to get you the information you need as quickly as possible. Access it at <https://nexus.securosis.com/>.

- **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 licensed 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/or 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/>.