Threat Detection Evolution: What Practitioners Need to Know

Version 1.5
Released: July 28, 2015

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.

This report is licensed by AlienVault whose support allows us to release it for free.

All content was developed independently.

www.alienvault.com

AlienVault's mission is to enable organizations with limited resources to accelerate and simplify their ability to detect and respond to the growing landscape of cyber threats. Our Unified Security Management (USM) platform provides all of the essential security controls required for complete security visibility, and is designed to enable any IT or security practitioner to benefit from results on day one. Powered by threat intelligence from AlienVault Labs and the AlienVault Open Threat Exchange -- the world's largest crowd-sourced threat intelligence network -- AlienVault USM delivers a unified, simple and affordable solution for threat detection, incident response and compliance management.

Copyright

Threat Detection Evolution

Table of Contents

Why Evolve Threat Detection?

Prevention isn't enough, whether you deploy it on the network or endpoints or both. It's not clear that it ever was, but there is additional proof every day that adversaries cannot be reliably stopped. We see the beginning of the long-awaited shift of focus and funding, from prevention to detection and investigation. But security practitioners have been trying to make sense of security data for years to shorten the window between compromise and detection… largely unsuccessfully.

Not to worry — we haven't become the latest security Chicken Little, warning everyone that the sky is falling. Mostly because it fell a long time ago, and we have been picking up the pieces ever since. It can be exhausting to chase alert after alert, never really knowing which are false positives and which indicate real active adversaries in your environment. Something has to change. We need to advance the practice of detection, to provide better and more actionable alerts. This requires thinking more broadly about detection, and starting to integrate the various different security monitoring systems in use today.

This paper will link our recent research on detection and threat intelligence, in the context of Threat Detection Evolution.

> It's not that all this data is useless. But it's not really useful either. There is value in having the data, but you cannot unlock it without integration, normalization, and analytics.

(Mostly) Useless Data

We have no lack of security data. All devices stream data all the time. Network devices, security devices, servers, and endpoints all generate tons of log data. Then you collect vulnerability data, configuration data, and possibly network flows or even network packets. You look for specific attacks with tools like intrusion detection devices and SIEM, which generate lots of alerts.

You probably have all this security data in a variety of places, with separate alerting policies implemented within each monitoring tool. It's hard enough to stay on top of a handful of consoles generating alerts, but past a dozen or so it's no longer really feasible to get a consistent view of your environment.

It's not that all this data is useless. But it's not really useful either. There is value in having the data, but you cannot unlock it without integration, normalization, and analytics. We have heard it said that finding attackers is like finding a needle in a stack of needles. It's not a question of *whether* there is a needle — you need to figure out *which* needle is poking you.

The traffic and activity on a typical network throw off so much data that it is trivial for adversaries to hide in plain sight, obfuscating their malicious behavior in a morass of legitimate activity. It's very difficult to figure out what's important until it's too late. And it's not getting easier — cloud computing and mobility are disrupting the traditional order of how technology is delivered and information is consumed by employees, customers, and business partners, so there will be more data and more activity to complicate threat detection.

Minding the Store…

In most discussions with practitioners, sooner or later we get around to the challenge of finding skilled staff to implement the security program. It's not a funding thing — companies are willing to invest, given the high profile of threats. The challenge is resource availability, and unfortunately there is no easy fix. The security industry is facing a large enough skills gap that there is no obvious answer.

It's hard for a n00b to come in and be productive their first couple years. Even those with formal (read: academic) training in security disciplines need a couple years of operational experience to become productive.

Why can't security practitioners be identified? What are the constraints on training more people to do security? It is actually counterintuitive, because security isn't a typical job. It's hard for a n00b to come in and be productive their first couple years. Even those with formal (read: academic) training in security disciplines need a couple years of operational experience to become productive. And a particular mindset is required to handle a job where true success is a myth. It's not a matter of *whether* an organization will be breached — it's *when*, and that is hard for most people to deal with day after day.

Additionally, if your organization is not a Global 1000 company or major consulting firm, finding qualified staff is even harder because you have many of the same problems as a large enterprise, but far less budget and skills available to solve them.

Clearly what we have been doing is insufficient to address the current problem. We need to look at the problem differently. It's not a challenge that can be fixed by throwing people at it, because there aren't enough people. It's not a challenge that can be fixed by throwing products at it either — organizations both large and small have been trying for years, and failing. Our industry needs to evolve its tactics to focus on doing the most important things more efficiently.

Efficiency and Integration

When you don't have enough staff you need to make existing staff more efficient. That typically involves two different tactics:

1. **Minimize False Positives and Negatives:** The thing that burns up more time than anything else is chasing alerts into ratholes and then finding out that they are false positives. So making sure alerts represent real risk is the best efficiency increase you can manage. You also want to minimize false negatives, when you miss the attack, have devices compromised and then spend a ton of time cleaning it up. You need to focus on minimizing errors to get better utilization out of your limited staff.

2. **Automate:** The other aspect of increasing efficiency is automation of non-strategic functions where possible. There isn't much value in making individual IPS rule changes based on reliable threat intel or vulnerability data. Once you can trust your automation you can have folks spend time on tasks that aren't suited to automation, like triaging possible attacks.

The other way to make better use of your staff is integration. The security business has grown incrementally to address specific problems. For example when first connecting to the Internet you needed a firewall to provide access control for inbound connections. Soon enough your network was being attacked, so you deployed an IPS to address those attacks. Then you wanted to control employee web use so you installed a web filter. Then you needed to see which devices were vulnerable and bought a vulnerability scanner, and so on and so forth.

This security sprawl continues, with new advanced security technologies to be deployed on the network, on endpoints, within your data center, and in the cloud. Of course you can't turn off the old controls, so a smaller organization may need to manage 7-10 different security products and services — all with separate consoles and policies. Larger organizations can have dozens. An integrated solution offers leverage by combining all those policies, offering a streamlined user experience for faster response.

The Goal: Risk-based Prioritization

To delve a bit into the land of motherhood and apple pie, organizations have been trying for a long time to allocate scarce resources based on potential impact to the organization. The mythical unicorn of security: prioritized alerts with context on what is actually at risk within your organization. There is no generic answer. What presents risk to one organization might not to another. Your threat detection approach needs to reflect these differences.

An evolved view of threat detection isn't just about finding attacks. It's about finding the attacks that present the greatest risk to your organization, and enabling an efficient and effective response.

An evolved view of threat detection isn't just about finding attacks. It's about finding the attacks that present the greatest risk to your organization, and enabling an efficient and effective response. This involves integrating a bunch of existing security data sources (both internal and external) and monitors, then performing contextual analysis on that data to prioritize them based on importance to your organization.

What Data to Collect?

Now that we have explained why detection must evolve, let's work through the mechanics of what that actually means. It comes down to two functions: security data collection and analytics of collected data. First we'll work through which data is helpful and where it comes from.

Threat detection requires two main types of security data. The first is internal data: security data collected from devices and other assets within your control. It's the stuff the PCI-DSS has been telling you to collect for years. Second is external data, more commonly known as threat intelligence. Here's the rub: there is no useful intelligence in external threat data without *context* for how it relates to your organization. But let's not put the cart before the horse. We need to understand what security data we have already, before worrying about external data.

Internal Data

You have likely heard a lot about *continuous monitoring* because the term is so shiny and attractive to security types. The problem we described in [Vulnerability Management Evolution]() is that 'continuous' can have a variety of different definitions, depending on who you are talking to. We have a rather constructionist view (meaning look at the dictionary) and figure the term means "without cessation." But often continuous monitoring doesn't actually add much value over regular and reliable daily monitoring.

> 'Continuous' can have a variety of different definitions, depending on who you are talking to. We have a rather constructionist view (meaning look at the dictionary) and figure the term means "without cessation."

We prefer *consistent* monitoring of internal resources. That may be continuous, for high-profile assets at great risk of compromise. Or possibly weekly for devices and servers that don't change much or access high-value data. The key is to be consistent about when you collect data from resources, and to ensure the data is reliable.

There are many data sources you might collect from for detection, including:

- **Logs:** The good news is that pretty much all your technology assets generate logs in some way, shape, or form. Whether it's a security or network device, a server, an endpoint, or even mobile. Odds are you can't manage to *collect* data from everything, so you'll need to choose which devices to monitor, but pretty much everything generates log data.

- **Vulnerability Data:** When trying to detect or figure out a potential issue, knowing which devices are vulnerable to what can be important for narrowing down your search. If you know a certain attack targets a certain vulnerability, and you only have a handful of devices that haven't been patched to address that vulnerability, you know where to look.

- **Configuration Data:** Like vulnerability data, configuration data offers valuable context for understanding whether a device could be exploited by a specific attack.

- **File Integrity:** File integrity monitoring provides important information for figuring out which key files have changed. A system file tampered with outside any authorized change may indicate nefarious activity, and should be checked out.

- **Network Flows:** Network flow data can identify patterns of typical network activity — which enables you to look for patterns which *aren't* quite normal that could represent reconnaissance, lateral movement, or even exfiltration.

Once you decide what data to collect, you have to figure out from where and how much. This involves selecting logical collection points and deciding where to aggregate the data. These choices depend on the architecture of your technology stack. Many organizations opt for a centralized aggregation point to facilitate end-to-end analysis, but it depends on the size of the organization. Large enterprises may not be able to collect everything in one place, and should consider some kind of hierarchical collection/aggregation strategy where data is stored and analyzed locally, and a subset sent upstream for central analysis.

> Once you decide what data to collect, you have figure out from where and how much. This involves selecting logical collection points and deciding where to aggregate the data.

Finally, we need to mention the role of the cloud in collection and aggregation, because almost everything is being offered either in the cloud or *as a Service* nowadays. Cloud-based aggregation and analysis depend on a few things. The first is the amount of data. Moving logs or flow records is not a big deal because they are pretty small and highly compressible. Moving network packets is a much more serious endeavor, particularly to a cloud service. The other key determinant is data sensitivity: some organizations are not comfortable with their key security data outside their control in someone else's data center/service. That's an organizational and cultural issue, but we've seen a much greater level of comfort with cloud-based log aggregation over the past year, and expect it to become far more commonplace inside the 2-year planning horizon.

The other key aspect of internal data collection is integration and normalization. Different data sources have different data formats, which requires us to normalize the data before datasets can be compared. That involves compromise in terms of granularity of common data formats, and can favor an integrated approach where all data sources are already integrated into a common security data store. Then you as the practitioner don't really need to worry about making all those compromises — instead you can bet on your vendor or service provider having already done the work.

Also consider the availability of resources for dealing with these disparate data sources. The key issue, mentioned above, remains the skills shortage — so starting a data aggregation/collection effort that depends on skilled resources to manage normalization and integration of data might not be the best idea. This doesn't really have much to do with the size of the organization — it's more about the sophistication of staff. Security data integration is an advanced function that can be beyond even large organizations with less mature security efforts.

Ultimately your goal is visibility into your entire technology infrastructure. An end-to-end view of what's happening in your environment, wherever your data is, gives you a basis for evolving your detection capabilities.

An end-to-end view of what's happening in your environment, wherever your data is, gives you a basis for evolving your detection capabilities.

External Data

We have published a lot of research on threat intel to date, most recently a series on Applied Threat Intelligence, which summarized the three main use cases we see for external data.

There are plenty of sources of external data nowadays. The main types are:

- **Commercial integrated:** Every security vendor seems to have a research group providing some type of intelligence. This data is usually very tightly integrated into their product or service. There may be a separate charge for intelligence, beyond the base cost of the product or service.
- **Commercial standalone:** Standalone threat intel is an emerging security market. These vendors typically offer an aggregation platform to collect external data and integrate it into your controls and monitoring systems. Some also gather industry-specific data because attacks tend to cluster in specific industries.
- **ISAC:** Information Sharing and Analysis Centers are industry-specific organizations that aggregate data across an industry to share it among their members. The best known ISAC is for the financial industry, although many other industry associations are spinning up their own ISACs as well.
- **OSINT:** Finally there is open source intel, publicly available sources for things like malware samples and IP reputation, which can be queried and/or have their intel integrated directly into user systems.

How does this external data play into an evolved threat detection capability? As we mentioned above, external data without context isn't very helpful. You don't know which of the alerts or notifications apply to your environment, so you create a lot of extra work figuring it out. And the idea is *not* to create more work.

Use external threat data to look for specific instances of an attack. These indicators from other attacks can be used pinpoint that activity in your network, even if you have never seen the attack before.

How can you provide that context? Use external threat data to look for specific instances of an attack. As we described in Threat Intelligence and Security Monitoring, you can use indicators from other attacks to pinpoint that activity in your network, even if you have never seen the attack before. Historically you were restricted to alerting on conditions/correlations you knew about firsthand, so this is a big deal.

For example, let's think back to retrospection. Let's say you didn't know about a heretofore unknown attack like Duqu 2.0, and received a set of indicators from a threat intelligence provider. You could look for those indicators within your network. Even if you don't find that specific attack, you could set your monitoring system (typically a SIEM or an IPS) to watch for those indicators. Threat intelligence enables you to almost jump through time, looking for attacks that haven't happened to you — yet…

Default to the Process

Once again it all comes back to process. We mapped that out in TI+SM and Threat Intelligence and Incident Response. You need a process to procure, collect, and utilize threat intelligence, from whatever sources it comes.

Identifying Malicious Activity

Once you have the data aggregated, you need to analyze it for indications of compromised devices and/or malicious activity within your organization.

Know Your Assets

You know the old business adage: you can't manage what you can't see. In security monitoring you need to discover new assets — and changes to existing ones — to monitor them, and ultimately to figure out when a device has been compromised. A key aspect of threat detection remains discovery. All devices, especially those pesky rogue wireless access points and mobile devices, provide attack surface for adversaries.

How can you make sure you are continuously discovering devices? You scan your address space. Of course there is active scanning, but that runs periodically. To fill in between active scans, passive scanning watches network traffic streaming by to identify devices you haven't seen or which have changed recently. Once a device is identified passively, you can launch an active scan to figure out what it's doing, and whether it is legitimate. Don't forget to discover your *entire* address space — both IPv4 and IPv6.

> A key aspect of threat detection remains discovery. Surprise is the enemy of a security professional, so it is essential to always be aware of network topology and devices on the network.

Most discovery efforts focus on PCs and servers on the internal network. But that may not be enough any more — it is typically endpoints that end up compromised, so you should discover both full computers and mobile devices. Finally, you will need to figure out how to discover assets in cloud computing environments. This requires integration with cloud consoles to ensure you know about new cloud-based resources and can monitor them appropriately.

After you have a handle on the devices within your environment, the next step is to classify them. We recommend a simple classification, involving roughly 4 groupings. The most important bucket contains critical devices with access to private information and/or valuable intellectual property. Next look for devices behaving maliciously. These devices may not have sensitive information, but adversaries can move laterally from them to critical devices. Then you have dormant devices, which may have connected to a command and control infrastructure but aren't *currently* doing anything malicious. Finally, there are all the other devices which haven't done anything suspicious — which you likely don't have time to worry about.

Finally, we continue to harp on the criticality of a consistent threat detection process, including discovery and classification. As with data collection, your technology environment is dynamic, so what you saw 10 minutes ago will have changed within another 20 minutes — if not sooner. You need a strong process to ensure you always understand what is happening in your environment.

The C Word

Correlation has always been a challenge for security folks. It’s not because the math doesn’t work. Math works just fine. Event correlation has been a challenge because you needed to know what to look for at a very granular level. Given the kinds of attacks and advanced adversaries many organizations face, you cannot afford to count on knowing what’s coming, so it’s hard to find new and innovative attacks via traditional correlation. This has led to the generally poor perceptions of SIEMs and IDS/IPS.

But that doesn’t mean correlation is useless for security. Looking for common attributes, and linking events together into meaningful models of possible attacks, provides a useful way to investigate security events. And you don’t want to succumb to the same attacks over and over again, so it is still important to look for indicators of attacks that have been used against you. Even better if you can detect indicators reported by other organizations via threat intelligence, and avoid those new attacks entirely.

> Correlation has always been a challenge for security folks. It’s not because the math doesn’t work. Math works just fine.

Additionally you can (and should) stage out a number of reasonable attack patterns via threat modeling to look for common attacks. Your vendor or service provider’s research team has likely built in some of these common patterns to kickstart your efforts at building out correlation rules. These research teams also keep their own correlation rules current based on what they see in the wild. Of course you can never know all possible attacks. So you also need to apply behavioral and other advanced analytical techniques to catch attacks you don’t yet know about.

Looking for Outliers

Technology systems have typical activity patterns. Whether network traffic, log events, transactions, or any other data source, you can establish an activity profile for how systems normally behave. Once the profile is established you look for anomalous activity, or outliers, that may represent malicious activity. These outliers could be anything, from any data source you collect.

With a massive trove of data, you can take advantage of advanced “Big Data” analytics (we don’t like the overly vague term, but the techniques are handy). New technologies can reduce a huge amount of data to something you can reasonably scan for abnormal activity patterns. You need an iterative process to refine thresholds and baseline over time. Yes, that means ongoing care and feeding of your security analytics. Activity evolves over time, so today’s normal might be anomalous next month.

Setting up these profiles and maintaining the analytics typically requires advanced skills. The new term for these professionals is data scientists. This skill set is expensive, but a key aspect of detecting threats is looking for outliers, which requires data scientists, so pay up. Just ensure you have sufficient resources to investigate alerts coming from your analytics engine, because if you aren't staffed to triage and validate alerts, you will see no benefit from earlier threat detection.

Alternatively, organizations without these sophisticated internal resources should consider allowing a vendor or service provider to update and tune their correlation rules and analytics for detection. This is especially helpful as organizations embrace more advanced analytics without internal data scientists to run the math.

> Just ensure you have sufficient resources to investigate alerts coming from your analytics engine, because if you aren't staffed to triage and validate alerts, you will see no benefit from earlier threat detection.

Visualization and Drill down

Given the challenges of finding skilled resources for triage and validation, you'll need to supplement internal skills with technology-accelerated functions, including better visualization and a built-in workflow to validate and triage alerts. You want a straightforward graphical metaphor to help categorize and prioritize alerts, and a way to dig into alerts to identify the root cause.

The only way to get a feel for whether a particular visual metaphor will work for you is to actually use it. That's why a proof of concept (PoC) is so important when looking at detection technologies and services. You pump some of your data into the tool, generate alerts, and validate them as you would in a production deployment. Even better, you have skilled resources from the vendor or channel partner to help stand up the system, perform initial configuration, and work through some alerts. Take advantage of these resources to kickstart your efforts.

Integration

Standalone analytics can work, especially for very specialized use cases such as large financial institutions addressing the insider threat. But a more generic detection platform may have greater impact in resource-constrained environments, where the expertise for general security monitoring is lacking. Not having to perform manual triage and validation of issues can save a ton of time and supplement your internal skill sets, especially if you leverage a vendor's security research and/or threat intelligence services.

So another key criteria for evolving threat detection is flexible integration with additional security data sources, emerging analytic techniques, advanced visualization engines, and operational workflow tools. Over time we expect this evolved threat detection capability to morph into the core security monitoring platform (either being subsumed by existing players or disrupting the status quo) — collecting internal security data, absorbing threat intelligence from a number of external sources, providing analytics to detect attacks, and ultimately sending information on to operational systems and controls to secure the environment.

Putting the Plan into Action for Better Detection

Let's work through a quick scenario to see how these concepts come together to enhance your ability to detect attacks. We will assume you work for a mid-sized super-regional retailer with 75 stores, 6 distribution centers, and headquarters. Your situation may be a bit different, especially if you work in a massive enterprise, but the concepts are the same.

Each of your locations is connected via an Internet-based VPN that works well. You've been gradually upgrading the perimeter network at HQ and within the distribution centers by implementing NGFW technology and turning on IPS on the devices. Each store has a low-end security gateway that provides separate networks for internal systems (requiring domain authentication) and customer Internet access. There are minimal IT staff and capabilities outside HQ. A technology lead is identified for each location, but they can barely tell you which lights are blinking on the boxes, so the entire environment is built to be remotely managed.

Given the sheer number of breaches reported by retailer after retailer, you know that the fact you haven't suffered a successful compromise (that you know about) is mostly good luck. You've been doing this too long to assume you are secure.

In terms of other controls, the big project over the past year has been deploying whitelisting on all fixed function devices in distribution centers and stores, including PoS systems and warehouse computers. This was a major undertaking to tune the environment so whitelisting did not break systems, but after a period of bumpiness the technology is working well. The high-profile retail attacks of 2014 freed up budget for the whitelisting project, but aside from that your security program is right out of the PCI-DSS playbook: simple logging, vulnerability scanning, IPS, and AV deployed to pass PCI assessment; but not much more.

Given the sheer number of breaches reported by retailer after retailer, you know that the fact you haven't suffered a successful compromise (that you know about) is mostly good luck. Getting ahead of PoS attacks with whitelisting has helped, but you've been doing this too long to assume you are secure. You know the simple logging and vulnerability scanning you are doing can easily be evaded, so you decide it is time to think more broadly about threat detection. But with so many different technologies and options, how do you get started? What do you do first?

Getting Started

The first step is always to leverage what you already have. The good news is that you've been logging and vulnerability scanning for years. The data isn't particularly actionable, but it's there. So you can start by aggregating it into a common place. Fortunately you don't need to spend a ton of money to aggregate your security data. Maybe it's a SIEM or possibly an offering that aggregates your security data in the cloud. Either way you'll start by putting all your security data in one place, getting rid of duplicate data, and normalizing your data sources, so you can start doing some analysis on a common dataset.

Once you have your data in one place, you can start setting up alerts to detect common attack patterns in your data. The good news is that all the aggregation technologies (SIEM and cloud-based monitoring) offer options. Some capabilities are more sophisticated than others, but you'll be able to get started with out-of-the-box capabilities. Even open source tools offer alerting rules to get you started. Additionally, security monitoring vendors invest significantly in research to define and optimize the rules that ship with their products.

> Even open source tools offer alerting rules to get you started. Additionally, security monitoring vendors invest significantly in research to define and optimize the rules that ship with their products.

One of the most straightforward attack patterns to look for involves privilege escalation after obvious reconnaissance. Yes, this is simple detection, but it illustrates the concept. Now that you have your server and IPS logs in one place, you can look for increased network port scans (usually indicating reconnaissance) and then privilege escalation on a server on one of the networks being searched. This is a typical rule/policy that ships with a SIEM or security monitoring service. But you could just as easily build this into your system to get started. Odds are that once you start looking for these patterns you'll find something. Let's assume you don't because you've done a good job thus far on the security fundamentals.

After starting by going through your first group of alerts, next you can look for assets in your environment which you don't know about. That entails either active or passive discovery of devices on the network. Start by scanning your entire address space to see what's there. You probably shouldn't do that during business hours, but a habit of checking consistently — perhaps weekly or monthly — is helpful. In between active scans you can also passively listen for network devices sending traffic passively, by either looking at network flow records or deploying a passive scanning capability specifically to look for new devices.

Let's say you discover your development shop has been testing out private cloud-based technologies to make better use of hardware in the data center. The only reason you noticed was passive discovery of a new set of devices communicating with back-end data stores. Armed with this information, you can meet with that business leader to make sure they take proper precautions to securely deploy their systems.

Between alerts generated from new rules and dealing with the new technology initiative you didn't know about, you feel pretty good about your new threat detection capability. But you're still looking for the stuff you *already* know you should look for. What really scares you is what you *don't* know to look for.

More Advanced Detection

To look for activity you don't know about you need to first define normal in your environment. Traffic that is not 'normal' provides a good indicator of potential attack. Activity outliers are a good place to start because network traffic and transaction flows tend to be reasonably stable in most environments. So you start with anomaly detection by spending a week or so training your detection system, setting baselines for network traffic and system activity.

Once you start getting alerts based on anomalies, you will spend a bit of time refining thresholds and decreasing the noise you see from alerts. This tuning time may be irritating but it's a necessary evil to optimize the system and ensure your alerts identify activity you need to investigate. And it turns out to be a good thing you set up the baselines, because you were able to detect emerging adversary activity in a distribution center. The attackers got in by targeting a warehouse manager with a phishing message, and they were in the midst of burrowing deeper into your environment when you saw strange traffic from that distribution center, targeting the Finance group to access payment information.

As expected, you found malicious activity in your environment. You just didn't have the optics to see it until you deployed your new detection capability. With the new detection system and some time wading through the initial alerts, you got a quick and substantial win from your investment.

> As expected, you found malicious activity in your environment. You just didn't have the optics to see it until you deployed your new detection capability.

Threat Intelligence

On the back of your high-profile win detecting attackers, you now want to start taking advantage of attacks you haven't seen. That means integrating threat intelligence to *benefit from the misfortune of others*. You first need to figure out what external data sources make sense for your environment. Your detection/monitoring vendor offers an open source threat intelligence service, so that first decision was pretty easy. At least for initial experimenting, lower-cost options are better.

Over time, as you refine your use of threat intel it may make sense to integrate other commercially available data — especially relating to trading communities because adversaries often target companies in the same industry. But for now your initial vendor feed will do the trick. So you turn on the feed and start working through the alerts. Again, this requires an investment of time to tune the alerts, but can yield specific results. Let's say you are able to detect a traffic pattern typical of an emerging malware attack kit based on alerts from your IPS. Without those specific indicators, you wouldn't have known that traffic was malicious.

Once you get comfortable with your vendor-supplied threat intel and have your system sufficiently tuned, you can start thinking about other sources. Given your presence in the retail space, and the fact that you already sold senior management on the need to participate in the Retail Information Sharing and Analysis Center (ISAC), using their indicators is a logical next step.

Keep in mind that the objective for leveraging this external data is to start looking for attacks you don't know exist because you haven't seen them. Nothing is perfect, so you'll also want to keep using out-of-the-box alerts and baselines on the monitoring systems. But if you can get ahead of the game a bit by looking for emerging attacks, you can shorten the window between attack and detection.

Taking Detection to the Next Level

The good news is that your new detection capability has shown value almost immediately. But as we discussed, it required significant tuning and demands considerable care and feeding over time. And you still face significant resource constraints, both at headquarters and in distribution centers and stores. So it makes sense to look for places where you can automate remediation.

Automation based on your evolved detection capability is about containing damage. So you want to get potentially compromised devices out of harm's way as quickly as possible. You can quarantine devices as soon as they behave suspiciously. You can directly integrate your monitoring system with either your network switches or some type of Network Access Control for this level of automation. Further, you could also integrate with egress firewalls to block traffic to destinations with poor IP reputations or and packets that look like command and control activity.

Automation based on your evolved detection capability is about containing damage. So you want to get potentially compromised devices out of harm's way as quickly as possible.

The key to any automation is *trust*. You need to trust automation before you can let it block traffic or quarantine devices. Obviously the downside to blocking legitimate traffic can be severe, so you first need to be comfortable with the validity of your alerts, and then with your integration, before you are ready to actually block traffic or quarantine devices programmatically.

We suggest a slow road to automation, recognizing the need to both tune and refine your detection system, and to integrate it with active network controls. Of course automation's potential is awesome. Imagine being able to see a device acting outside normal parameters, take it off the network, start an investigation, and block any other traffic to destinations the suspect device was communicating to — all automatically. Yes, it takes time and sophistication to get there. But it is possible today, and the technologies are maturing rapidly.

Summary

It’s clear that we need threat detection to evolve to catch modern adversaries. We made the case for more advanced data collection and analytics, and integration of external threat intelligence, to evolve detection.

The success criterion for any detection process is that effective action results from alerts. That requires ensuring alerts are legitimate, provide some context to allow prioritization of remediation given the risk to your organization, and visualization of threats to drive prioritization decisions. This context comes from both internal data (similar behaviors on multiple devices can indicate an outbreak) and external data (new indicators can retrospectively identify ongoing adversary campaigns within your environment).

Detection is an ongoing process, requiring consistent tuning and effort to optimize your efforts and results. But the investment can dramatically shorten the window between attack and detection, and that’s about the best you can do in today’s environment of advanced attackers and limited defenders (in terms of both skills and resources).

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

About the Analyst

Mike Rothman, Analyst and President

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

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

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

About Securosis

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

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