Building an Early Warning System

Version 1.4

Released: January 22, 2013

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 Lookingglass Cyber Solutions

Lookingglass Cyber Solutions is the world leader in Over-The-Horizon Cyber Threat Suppression. The Lookingglass product suite provides revolutionary solutions which deliver deep visibility into the Internet threat landscape. With its ScoutVisionTM and CloudScoutTM solutions, Lookingglass offers alert and warning capabilities that empower customers to continuously monitor for over the horizon threats from outside a client's network and control, such as the presence of botnets, hosts associated with cyber criminal networks, unexpected route changes and the loss of network resiliency. Lookingglass' outside-in approach accounts for a client's entire enterprise cyber ecosystem including the extended enterprise, and other networks beyond their control. This oversight ensures business partners or service providers are not the proximate cause of security breaches or data loss. For more information, visit www.LGScout.com.

Copyright

Table of Contents

Introduction

Getting ahead of the attackers is the holy grail to security folks. A few years back some vendors sold their customers a bill of goods, claiming they could "get ahead of the threat." That didn't work out very well, and most of the world appreciates that security is inherently reactive. The realistic objective is to reduce the time it takes to react under attack, in order to contain the eventual damage. We call this Reacting Faster and Better. Under this philosophy, the most important thing is to build an effective incident response process. But that's not the end of the game. You can shrink the window of exploitation by leveraging cutting-edge research to help focus your efforts more effectively, by looking in the places attackers are most likely to strike. You need an *Early Warning System* (EWS) for perspective on what is coming at you.

Pragmatic Intelligence

Back in 2007 when *The Pragmatic CSO* was written, prioritization was a key part of the operational methodology it espoused. Over the past 5 years we have kept focus on the importance of prioritizing your limited funding, resources, and expertise on the highest-value activities. To get a feel for how this concept works, let's excerpt a small section from *The Pragmatic CSO*:

> *A key operational discipline is figuring out the most likely exposure and working to eliminate it. This is particularly hard because many CSOs run from emergency to emergency without ever getting a chance to manage their security environment or even spend 10 minutes thinking about what is next. Unfortunately, what's next has already happened. Clearly this situation must be addressed.*
>
> *"A good hockey player plays where the puck is. A great hockey player plays where the puck is going to be." – Wayne Gretzky*
>
> *The great ones, in whatever pursuit, figure out how to anticipate what is most likely to happen, so they are ready if it does. Some think it's luck, others figure it's a talent bestowed by a higher power. Actually, in most cases, it is the result of a tremendous amount of hard work. The ability to anticipate is especially critical in security because of the unlimited number of possible attacks across an infinite attack surface. You cannot cover all the bases, so you need to be focused and choose correctly. What is the best way to choose correctly? You need an "inside man" working on your behalf to figure out what the bad guys are working on.*
>
> *Thus, security research plays a critical role in the life of a Pragmatic CSO. It's hard to believe, but Pragmatic CSOs read a lot. They are plugged into the underground networks of researchers that spend time penetrating the hacker networks and tracking down the bot masters to figure out what they are working on. If you know what the bad guys are focused on, you can get a real good idea about what they are planning to strike next. Even though you don't have to spend money to get connected with the research folks, a number of*

> *services focus on reporting new exploits and figuring out what is most likely to be attacked on any given day.*

Of course context is everything, so although third party research may give you a clue to what the next exploit or botnet looks like, it cannot tell you how the attack will be used *against your defenses*. You need to provide that context, which requires looking at the situation from two different perspectives:

1. **In Here:** This is the internal perspective gleaned from what's really happening on your network. Whether the platform to aggregate and analyze the data is a SIEM or a Vulnerability Management platform or any other technology, the point is the same. The foundation for context is a clear understanding of what's going on *within* your environment. Only then you can move on to analyzing what's exposed and determining what needs to be fixed *right now*.
2. **Out There:** The reverse perspective looks at the macro environment, starting with an understanding of attacker tactics and exploits, and figuring out how they will affect you. If you know about attacks you can preemptively implement protection. Obviously you need to walk before you run, so getting a handle on your internal security data is a necessary first step. But once you are there, factoring in the external view can really help narrow down your attack surface.

None of this is new. Law enforcement has been doing this, well, forever. The goal is to penetrate the adversary, learn their methods, and take action *before* an attack. Even in security there is a lot of precedent for this kind of approach. Back at TruSecure (now part of Verizon Business) over a decade ago, the security program was based on performing external threat research and using it to prioritize the controls to be implemented to address imminent attacks. Amazingly enough it worked. But this approach fell out of favor over the past 5-7 years as the entire industry got weighed down by the compliance albatross.

None of this is new. Law enforcement has been doing this, well, forever. The goal is to penetrate the adversary, learn their methods, and take action *before* an attack.

Now that the pendulum is swinging back toward actually securing stuff, we see a resurgence of threat intelligence as a way to make our defenses more effective and efficient. Let's quickly run through the history of security research, now typically called *threat intelligence*.

The Evolution of Threat Intelligence

Back in the day, "security research" really meant anti-virus research. The AV companies looked at viruses, built signatures, and moved on to the next. It was a fairly collegial environment and AV companies shared the malware they discovered, making sure everyone was protected within a couple hours. The next wave of research resulted from the avalanche of spam, which required security companies to build global networks of honeypots to capture bad email directly, create signatures to identify it, and distribute the signatures to their gateways.

Of course that lasted only until the spammers became more effective at evading signatures, which drove heavier reliance on behavioral indicators to infer which files were malware and which messages were spam. This required security vendors to spend time evaluating behavior and tuning their detection cocktails to maintain efficiency. At about this time IP and file reputation started becoming more widely used. An IP address that sends out spam is likely to continue sending spam and to launch other attacks, so give it a bad reputation score and block future messages from that IP.

Then we reached another inflection point, where attackers started using networks of compromised devices (botnets) to defeat reputation and evade detection. They made increasing use of polymorphic malware, rendering traditional signature-based detection largely useless. Attackers also got very sophisticated about masking communications between bots and their controllers.

So security research evolved as well, investigating all these techniques and data sources, aggregating and analyzing attackers, to get a better handle on the threat landscape. The problem is that this research is proprietary and unavailable to most people. It is usually under the umbrella of a security product or service, and increasingly used by vendors as a competitive differentiator.

These days proprietary security research is table stakes for any security vendor, and the industry has gotten much better at publicizing its findings via researcher blogs and other media. Much more information is available than ever before, but what does this mean for you? How can you leverage threat intelligence to provide that elusive *Early Warning System*?

Much more information is available than ever before, but what does this mean for you? How can you leverage threat intelligence to provide that elusive *Early Warning System*?

That's what this paper is all about. We will define a process for integrating threat intelligence into your security program, and then dig into each aspect of the process. This includes baselining internal data sources, leveraging external threat feeds, performing the analysis to put all this information into the context of your business, and finally building a scenario so you can see how the Early Warning system works in practice.

The Early Warning Process

We started by focusing on the increasing importance of threat intelligence for combating advanced attackers by understanding the tactics they are using *right now* against your defenses. With this intelligence, combined with information about what's happening in your environment, you can effectively prioritize your efforts and make better, more efficient use of your limited security resources.

Security folks can learn plenty from how law enforcement takes preventative action based on its best assessment of the risk – within the constraints of funding, resources, and expertise. Of course they cannot cover everything – nobody can. An attacker using fairly simple reconnaissance tactics can understand the visible controls and build an attack plan to evade them. Sound familiar? Law enforcement needs something to help focus its efforts. Something to give them a clue about where to look for the next attack.

That's where *threat intelligence* comes in. Law enforcement has folks who monitor threat sources (radical groups, weapon purchases, etc.) in order to identify patterns that could represent a tangible threat to a specific target. When law enforcement finds something it may act directly to neutralize the threat or work with targets to ensure they improve controls and watch for that particular class of threat.

> Investment and commitment are required to look beyond your walls in order to shorten the elapsed time between attack, detection and remediation.

Of course commercial entities operate under different constraints – they cannot rely on an external party such as government to provide sufficient threat intelligence. So commercial enterprises need to both deploy controls and build their own analysis capabilities to gather and process threat intelligence, which are neither easy nor quick. Investment and commitment are required to look beyond your walls in order to shorten the elapsed time between attack, detection and remediation. This involves a structured process to focus not only on what you see internally, but also on what's happening externally.

That's what the Early Warning System is all about.

The Early Warning Process

An Early Warning System (EWS) implements a systematic process to collect data internally, gather intelligence from third parties, and then analyze this information for particular dangers to your environment. We will list the components here and dig into them through this paper.

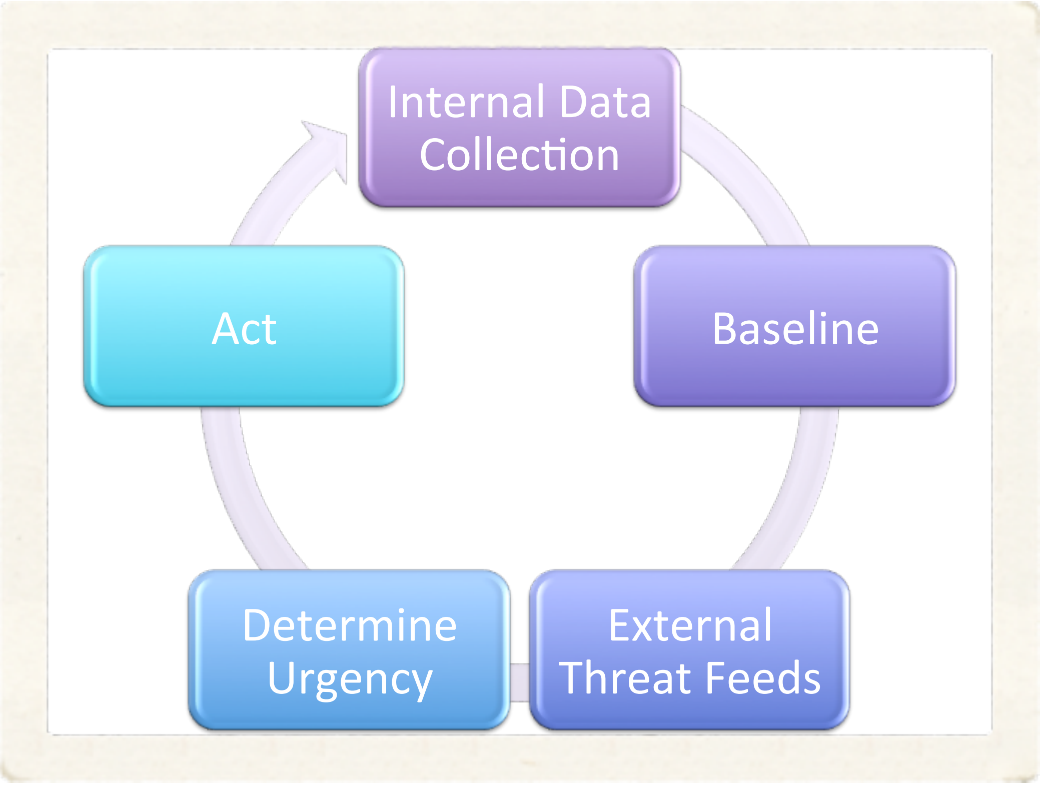

Internal Data Collection/Baselining

The first two steps look inward. We have done a ton of research over the past few years into technologies that enable you to analyze your security data. Whether your aggregation platform works from the perspective of SIEM / Log Management, Vulnerability Management, or even network forensics doesn't really matter. But you need a tool to aggregate your security data, with the ability to include external information.

Once the data is aggregated, the next step involves establishing a number of baselines to identify 'normal' for your environment. Of course *normal* is not really the right term – no environment is really normal. But you need to start somewhere. Equipped with baselines describing what typically happens in your environment, you can start looking for anomalies – deviations from those baselines – which might indicate something funky and bear further investigation.

As a technology foundation for EWS, your existing tools could be entirely suitable. But most data aggregation tools are built specifically to meet compliance requirements. They are designed to put a *lot* of security data in one place and generate reports, but are generally not so good at analysis or pattern detection – especially if you don't know the patterns to look for in advance. An EWS needs to be built on strong analytical tools – with a threat-centric view of the world rather than a focus on compliance reporting.

> An EWS needs to be built on strong analytical tools – with a threat-centric view of the world rather than a focus on compliance reporting.

External Threat Feeds

Once your internal data is in order, you can cross-reference it against attacks seen in the wild via external threat research. You will quickly see there is no lack of threat data. Every vendor has a security research team, and they are constantly generating content – both for external consumption (mostly public relations) and for integration into their own products and services. There are also providers who generate stand-alone threat data feeds for integration into existing tools. So you need to identify the threat intelligence sources to use, and to manage the overlap between different sources. Finally, you need to integrate your feeds into the tool(s) you use to aggregate security data to enable the kind of analysis needed for the next step in the EW process.

Determine Urgency

The next part of the process is to provide context, for both the internal and external data you gathered, and apply it systematically to your specific situation to determine whether you need to take action. This requires more than just security data and an analytics engine to evaluate the data and find patterns. The key to benefitting from an Early Warning is to *interpret* the data, and determine the degree of risk each threat poses specifically *to your environment*. This is neither cheap, nor easy to automate. We know of several organizations which employ dozens of folks specifically for processing threat intelligence and assessing risk. This is not common (and these folks have very mature security programs), but it's simply not realistic to buy a toaster (or any simple security appliance) and expect it to pump out actionable threat intelligence.

Act

An EWS needs to be built on strong analytical tools – with a threat-centric view of the world rather than a focus on compliance reporting. Finally you have information which might indicate an attack is coming. Of course you won't have precision on when or where. So what do you do? Do you take action? Do you just monitor for indicators of the attack commencing? Do you do nothing and wait until systems are down before mopping up the mess? At one time or another you will likely find yourself doing each of these things. Equipped with contextual threat intelligence from your EWS, you can implement workarounds and remediations to address the issues you found. Or not — depending on your culture and ability to survive bets that might be wrong.

That is a key concept to accept. Building an Early Warning System doesn't mean you will always be right. Just like any other system designed to predict the future, it *will* be wrong, and you *will* see false positives. So one success criterion for implementing an EWS is realism, about what it can do and what it can't. Optimize the resources at your disposal by focusing on the most likely (successful) attacks.

Equipped with contextual threat intelligence from your EWS, you can implement workarounds and remediations to address the issues you found. Or not — depending on your culture and ability to survive bets that might be wrong.

Internal Data Collection and Baselining

Third-party threat intelligence, as we will discuss later, will tell you what kinds of attacks you are *more likely* to see based on what is happening out in the world. But monitoring your own environment and looking for variations from normal activity tell you whether those attacks actually **are** hitting you or **have already** compromised your systems. This information provides the context to determine the urgency of the attack.

> Monitoring your own environment and looking for variations from normal activity tell you whether those attacks actually **are** hitting you or **have already** compromised your systems.

Internal Data Collection

An EWS needs to be built on strong analytical tools – with a threat-centric view of the world rather than a focus on compliance reporting. The process starts with collecting data from your internal sources for analysis. Most of you already have data aggregated in a log management system, because compliance has been mandating log management for years. More advanced organizations may have a Security Operations Center (SOC) leveraging a SIEM platform to do more security-oriented correlation and forensics to pinpoint and investigate attacks. Either way, you are likely collecting data which will provide the basis for the internal side of your EWS.

Let's take a quick look at the kinds of data you are likely already collecting and their relevance to the EWS:

1. **Network Security Devices:** Your firewalls and IPS devices generate huge logs of what's blocked, what's not, and which rules are effective. The EWS will match attack patterns and traffic against what is known about other attacks; so recognizing port/protocol/destination combinations, or application identifiers for next-generation firewalls, will be helpful.

2. **Identity:** Similarly information about logins, authentication failures, and other identity-related data is useful for matching against attack profiles received from third-party threat intelligence providers.

3. **Application/Database Logs:** Application specific logs are generally less relevant, unless they come from standard applications or components likely to be specifically targeted by attackers. Database transaction logs are generally more useful for identifying unusual database transactions – which might represent bulk data removal, injection attempts, or efforts to bring applications down. Database Activity Monitoring (DAM) logs are useful for determining the patterns of database requests, particularly when monitoring traffic within the database (or on the database server) consumes too many resources.

4. **NetFlow:** Another data type commonly used in SIEM environments is NetFlow – which provides information on protocols, sources, and destinations for network traffic as it traverses devices. NetFlow records are similar to firewall logs but far smaller, making them more useful for high-speed networks. Network flows can identify lateral movement by attackers, as well as large file transfers.

5. **Vulnerability Scans:** These offer an idea of which devices are vulnerable to specific attacks, which is critical for the EWS to help pinpoint which devices are potential targets for which attacks. You don't need to to worry about Windows exploits against Linux servers, for instance; this information enables you to focus monitoring, investigation, and workarounds on the devices more likely to be successfully attacked.

6. **Configuration Data:** The last major security data category is configuration data, which provides information on the configurations and settings of monitored devices. This data helps from an EWS perspective because when you see unexpected configuration changes, that could indicate a malware infection. The intelligence services can provide specific indicators of compromise (changes to devices made by the malware) and allow you to search through your configuration database enables to determine if the attack has been successful against any devices in your environment.

After figuring out which data you will collect, you need to decide where to put it. That means selecting a platform for your Early Warning System. You already have a SIEM/Log Management offering, so that's one possibility. You also likely have a vulnerability management platform, so that's another choice. We are not religious about which technology gets the nod, but a few capabilities are essential for an EWS. Let's not put the cart before the horse, though – we don't yet have enough context on other aspects of the process to understand which platform(s) might make sense. So we will defer the platform decision until later in this paper.

Baseline

Before your internal data is useful to the EWS, you need to define *normal*. As we mentioned earlier, ‘normal’ does not necessarily mean *secure*. If you are anything like almost every other enterprise, you likely have malware and compromised devices on your network already. Sorry to burst your bubble. You need to be able to identify indications of something *different*. Something that could represent an attack, an outbreak, or an exfiltration attempt. To be clear, different doesn’t always mean bad. The alert might be a false positive, or it could represent a new normal to accept, but either way you want to consciously decide if action is required and continuously adapt and evolve the baseline.

> You need to be able to identify indications of something *different*. Something that could represent an attack, an outbreak, or an exfiltration attempt.

Here is a simple process for building a baseline:

1. **Pick data source(s):** Start by picking a single data source and collect some data. Then determine the ranges you see within the data set. As an example think about firewall logs. These log records typically include the type of traffic (ports, protocols, applications, etc.), the destination IP address, the time of day, and whether the packet was blocked. You can pick numerous data sources and do sophisticated data mining, but we will keep it simple.

2. **Monitor the patterns:** Then collect traffic for a while, typically a few days to a week, and start analyzing it. Get your inner statistician on and start calculating averages, means, medians, and frequencies for your data set. In our example you might determine that 15% of your inbound web traffic during lunchtime is SSL destined for your shopping cart application.

3. **Define the initial thresholds:** From the initial patterns you can set thresholds, with variations beyond them indicating a potential problem. Maybe you set the initial thresholds 2 standard deviations above the mean for a traffic type. You look at the medians and means to figure out which initial threshold makes sense. You don’t need to be precise with the initial thresholds – you don’t yet have enough data or knowledge to know what represents an attack – but they are a place to start. Getting back to our firewall example, a spike in outbound SSL traffic to 30% might indicate an exfiltration. Or it could indicate a bunch of people shopping on their lunch hour. But it’s not normal so you will check it out.

4. **Refine and iterate:** As with any modeling effort, you will be wrong when you start. Over time you will refine the thresholds based on the traffic you see, the attacks you identify, and the threat intelligence you receive. The goal is to tighten the thresholds, reduce the false positives, and ensure the system makes you more efficient.

We need to reiterate the importance of incremental iterative progress for the Early Warning System. As with any project,we would like a quick win to show value soon after deployment (or even during proof of concept), but building a robust Early Warning capability for imminent attacks requires a scientific process. You need to constantly tune your systems, based both on internal factors (how traffic patterns change), and external intelligence that characterizes your adversaries' tactics.

Of course that doesn't mean you won't look for a quick win to show immediate value from the integration of threat intelligence into your security program. But the Early Warning System is a process that requires ongoing effort. Like everything else in security management, this is not *set and forget* technology.

External Threat Feeds

> Of course these threat feeds aren’t a fancy crystal ball that can tell you about an attack before it happens. The attack has already happened, but **hopefully not yet to you.**

This is a great beginning but it still puts you in a *reactive* mode. Even if you can detect an anomaly in your environment, it has already happened and you may be too late to prevent data loss. The next step for Early Warning is to look outside your own environment to figure out what’s happening externally. External threat intelligence provides a sense of current attacks and a set of indicators to look for in your internal data feeds. Of course these threat feeds aren’t a fancy crystal ball that can tell you about an attack before it happens. The attack has already happened, but **hopefully not yet to you.** We have never believed you can get ahead of an attack without a time machine. But you can become aware of an attack in the wild before it’s aimed at you, and take measures to ensure you are protected against it.

Types of Threat Intelligence

There are many different types of threat intelligence, and we are likely to see more emerge as the hype machine ramps up. Let’s quickly review the kinds of intel at your disposal and how they can help with the Early Warning process.

Threats and Malware

Malware analysis is maturing rapidly, and it is becoming commonplace to quickly and thoroughly understand exactly what a malicious code sample does and how to identify behavioral indicators. We described this process in detail in Malware Analysis Quant. For now suffice it to say you aren’t looking for a specific file, but for indicators that a file did something to a device. Fortunately a number of parties offer information services that provide data on specific pieces of malware. You can get an analysis based on a hash of the malware file or upload a file if it hasn’t been seen before. Then the service runs the malware through a sandbox to figure out what it does, profile it, and deliver that data back to you.

What do you do with indicators of compromise? Search your environment for evidence that the malware has executed. Obviously that requires a significant and intrusive search of the configuration files, executables, and registry settings on each device, which typically means some kind of endpoint forensics agent on each one. If that kind of access is available, then malware intelligence can provide a smoking gun for identification of compromised devices.

Vulnerabilities

Most folks never see the feed of new vulnerabilities that show up on a weekly or daily basis. Each scanner vendor updates their products behind the scenes and uses the most current updates to figure out which devices are vulnerable to each new attack. But the ability to detect a new attack is directly related to how often devices get scanned. A slightly different approach involves cross-referencing threat data — which attacks are in use — against vulnerability data to identify devices at risk. For example, if weaponized malware emerges to target a specific vulnerability, it would be extremely useful to have an integrated way to dump out a list of devices vulnerable to that attack. Of course you can do this manually by reading threat intelligence reports and then searching vulnerability scanner output to manually create lists of impacted devices, but will you? Anything that requires additional effort is all too likely to not get done. That's why the Early Warning System needs to be driven by a platform integrating all this intelligence, correlating it, and providing actionable information.

Reputation

Since its emergence as a key data source in the battle against spam, reputation data has rapidly become a component of seemingly every security control. For example, seeing that an IP address in one of your partner networks is compromised should set off alarms, especially if that partner has a direct connection to you. Basically anything can (and should) have a reputation. Devices, IP addressees, URLs, and domains for starters. If you have traffic going to a known bad site, that's a problem. If one of your devices gets a bad reputation – perhaps as a spam relay or DoS attacker – you need to know ASAP.

One specialization of reputation emerging as a separate intelligence feed is *botnet intelligence*. These feeds track command and control traffic globally and use that information to pinpoint malware originators, botnet controllers, and other IP addresses and sites your devices should avoid. Integrating this kind of feed with a firewall or web filter could prevent exfiltration or communication with a controller, and identify an active bot. Including this kind of data in the Early Warning System enables you to use evidence of bad behavior to prioritize remediation activities.

One specialization of reputation emerging as a separate intelligence feed is *botnet intelligence*. These feeds track command and control traffic globally and use that information to pinpoint malware originators, botnet controllers, and other IP addresses and sites your devices should avoid.

Brand Usage

It would be good to get a heads up if a hacktivist group targets your organization, or a band of pirates is stealing your copyrights, so a number of services have emerged to track mentions of companies on the Internet and infer whether they are positive or negative. Copyright violations, brand squatters, and all sorts of other shenanigans can trigger alerts — hopefully *before* extensive damage is done. How does this help with Early Warning? If your organization is a target, you are likely to see several different attack vectors. Think of these services as providing the information to go from DEFCON 5 to DEFCON 3, which might entail tightening the thresholds on other intelligence feeds and monitoring sources in preparation for imminent attack.

Managing the Overlap

With all these disparate data sources it becomes a significant challenge to make sure you don't get the same alerts multiple times. Unless your organization has a money tree in the courtyard, you likely had to rob Peter to pay Paul to just get the first intelligence service in the first place. There isn't any point in paying for the same stuff twice. The first step in determining overlap is to understand how the intelligence vendor gets their data. Do they use honeypots? Do they mine DNS traffic and track new domain registrations? Have they built a cloud-based malware analysis/sandboxing capability? You can categorize vendors by their tactics, to pick the best for your requirements.

Don't fall for the marketing hyperbole about proprietary algorithms, Big Data analysis, and staff linguists penetrating hacker dens and other stories straight out of a spy novel.

In order to pick the best vendor need to compare their services for comprehensiveness, timeliness, and accuracy. Yes, we're talking about a bake-off. Sign up for trials for a number of services and monitor their feeds for a week or so. Does one provider consistently identify new threats earlier? Is their information correct? Do they provide more detailed and actionable analysis?

Don't fall for the marketing hyperbole about proprietary algorithms, Big Data analysis, and staff linguists penetrating hacker dens and other stories straight out of a spy novel. It's really about the data and how useful it is to your Early Warning efforts. There are also more discreet ways to test these offerings. Send them a real malware sample you found. Do they identify it correctly? Quickly? And you can find a blacklisted IP (or a million) to test the reputation services. Buyer beware, and make sure you put each intelligence provider through its paces before you commit.

The good news is that — given the analysis and resource requirements to manage these intelligence feeds — aggregators will emerge sooner than later to take many of these different sources, Q/A, deduplicate, normalize, and validate — hoping to provide a cleaner and more valuable feed. Test these aggregated services just like anything else. Put each through its paces and evaluate on the same criteria.

The last point here is the importance of *short agreements*, especially up front. You cannot know how these services will work for you until you actually start using them. Many of these intelligence companies are startups, and may not be around in 3-4 years. Once you identify a set of core intelligence feeds longer deals can be cut, but we recommend not doing that until your Early Warning process matures and your intelligence vendor establishes a track record.

Caveats

As interesting and promising as threat intelligence is, we need to reiterate its limitations. First, increasingly targeted attacks mean you may be facing something that has been custom-built for your environment. A third-party feed cannot help with that. We are referring to tactics such as custom phishing emails targeting your CEO or someone in your finance department. If they fall for it and click the link, it's game over. No threat intelligence service can tell you to monitor "Bob in Finance," or the CEO's son.

Similarly, if your organization is targeted by a persistent attacker (you know who you are) you may be facing novel zero-day attacks. No intelligence service will see these before they hit you – they should send you a thank-you note for each brand-new malware sample.

But nobody's perfect, and sometimes intelligence feed are wrong. They could create false positives that spin you around in circles and waste precious time. If you are only buying data your downside risk is contained. Threat intelligence errors are more problematic if you have an active control that starts blocking traffic based on a false positive. We have already seen one vendor identify another vendor's honeypot as a botnet controller and raise all sorts of alarms. Or where an endpoint protection product flags a legitimate application update (Skype) as malware. This is an emerging inexact science. Remember the old adage: *trust, but verify*.

> Threat intelligence is great, but it rarely provides a true smoking gun. You need to systematically interpret the intelligence in the context of your environment.

You will hear the term *threat intelligence* in a variety of contexts. Endpoint protection vendors talk about how their threat intelligence helps them defend against advanced malware; firewall vendors claim their reputation networks help block attackers at the perimeter. Once you stop laughing at the claims, the point is that unless you get a **feed** of threat information you can integrate into other tools and check against your internal data, the intelligence is useless for Early Warnings.

Threat intelligence is great, but it rarely provides a true smoking gun. You need to systematically interpret the intelligence in the context of your environment and situation to determine the appropriate course of action. No machine does that, so you need HUMINT (human intelligence) — which brings us to determining the urgency of alerts generated by the Early Warning System.

Determining Urgency

Collecting internal security data and looking at external threat intelligence are all well and good, but you still have lots of *data* without usable *information*. Now we will focus on the analysis aspect of the Early Warning System. You might think this is just rehashing a lot of the work done through our SIEM, Incident Response, and Network Forensics research because all those functions also leverage data in order to identify attacks. The biggest difference is that in an Early Warning context you don't know what you're looking for. Years ago, US Defense Secretary Donald Rumsfeld described this as looking for "unknown unknowns".

> Early Warning turns traditional security analysis on its head.

Early Warning turns traditional security analysis on its head. Using traditional tools and tactics, including those mentioned above, you look for patterns in the data. If it's already in your security data, that means it has already happened to you. The traditional approaches require you to know what you are looking for – by modeling threats, baselining your environment, and then looking for things out of the ordinary. Early Warning brings a new facet to your analysis: Looking for attacks that may not have happened to you yet.

As a security professional your BS detector is probably howling right now. Most of you gave up on proactively fighting threats long ago. Is this Early Warning thingy magical? Does a unicorn deliver it to your loading dock? Of course not. But EWS analysis helps you focus, and enables you to more effectively mine internal security data. It offers some context for the reams of data you have collected. By using threat intelligence to more specifically search your environment for the most likely areas of attack, you can make informed guesses as to what will come next. This helps you figure out the relevance and urgency of implementing controls to block these emerging attacks.

So you aren't *really* looking for "unknown unknowns". You're looking for signs of emerging attacks, using indicators found by *others*. Which at least beats waiting until your data is exfiltrated to figure out a that new Trojan is circulating. Much better to learn for the misfortunes of others and head off attackers before they use the latest attack on *you*. It comes down to looking at both external and internal data, and deciding to how urgently you need to take action. We call this **Early Warning Urgency**, described by a very simple formula.

Relevance * Likelihood * Proximity = Early Warning Urgency

Relevance

The first order of business is to determine the relevance of any threat intelligence to your organization. This should be based on the specifics of the threat, and whether it can be used against you. Like the attack path analysis described in Vulnerability Management Evolution, real vulnerabilities which do not exist in your environment pose a very low risk. For example, you probably don't need to worry about being attacked by StuxNet if you don't have any control systems. That doesn't mean you shouldn't pay any attention to the attack – it uses a number of interesting Windows exploits, and may evolve in the future – but if you don't have any control systems its relevance is low. There are two aspects to relevance:

1. **Attack surface:** Are you vulnerable to the specific attack vector? Weaponized Windows 2000 exploits aren't relevant if you don't have any Windows 2000 systems in your environment. Once you have patched all instances of a specific vulnerability on your devices, you get a respite from worrying about that exploit. This is how the asset base and vulnerability information within your internally collected data provide context to determine Early Warning urgency.

2. **Intelligence Reliability:** You need to keep re-evaluating each threat intelligence feed to determine its usefulness. If a particular feed triggers many false positives it is less relevant. On the other hand, if a feed usually nails a certain type of attack, you should take warnings of that type particularly seriously.

Note that attack surface isn't necessarily restricted to your own assets and environment. Service providers, business partners, and even customers represent indirect risks to your environment – if one of them is compromised, the attack might have a direct path to your assets. We discuss that scenario under Proximity, below.

Likelihood

When trying to assess the likelihood of an actual attack based on an Early Warning, you need to consider the *attacker*. This is where adversary analysis comes into play, as we discussed in Defending Against Denial of Service. Threat intelligence includes speculation regarding the adversary; this helps you determine the likelihood of a successful attack, based on the competence and motive of the attacker. State-sponsored attackers, for instance, generally demand greater diligence than pranksters. You can also weigh the type of information targeted by an attack to determine your risk. You probably don't need to pay much attention to credit card stealing trojans if you don't process or store credit cards.

It's also a good idea to monitor your industry and competitors, since if another organization with similar characteristics to you in terms of customers, intellectual property, expertise, etc. is being targeted, you are more likely to be targeted as well. This kind of industry data can also be used to monitor trends. For instance, retailers may be more likely to be attacked during their peak seasons, which would increase the likelihood number during those periods. As we've mentioned, information sharing networks and services will be critical to this effort, though be wary of finding sources with an acceptable signal to noise ratio.

Likelihood is a *squishy* concept, and most risk analysis folks use all sorts of statistical models and analysis techniques to solidify their assessments. We certainly like the idea of quantifying attack likelihood with fine granularity, but we try to be realistic about the amount of data you will have to analyze. The *likelihood* variable tends to be more art than science; but over time, as threat intelligence services aggregate more data over a longer period, they should provide better and more quantified analysis.

Proximity

How early do you want the warning to be? An Early Warning System can track not only direct attacks on your environment, but also indirect attacks on organizations and individuals you connect with. We call this *proximity.* Direct attacks have a higher proximity factor and greater urgency. An attack on *you* is more serious (to you, at least) than one on your neighbor. The attack isn't material (or real) until it is launched directly against you, but you will want to factor other parties into your Early Warning System.

Let's start with business partners. If a business partner is compromised, the attacker may be able to jump from their network to yours, using a direct network connection or stolen credentials. Many business partners have *trusted* access to, or credentials on, your systems — which makes a bad day if they get pwned. So you should start proximity analysis by categorizing each different type of business partner and grouping them by their access to your critical data. How can you determine if a business partner has been compromised? Emerging proprietary intelligence services and other industry information sharing groups monitor organizations and industries, assess their vulnerabilities, and determine the risks they pose to partners. But this is third party information and may be flawed, so intelligence reliability must also be factored in.

> If a business partner is compromised, the attacker may be able to jump from their network to yours, using a direct network connection or stolen credentials.

Likewise, you can and should monitor service providers as potential indirect attack vectors. As more and more critical data moves to SaaS environments and infrastructure migrates to cloud computing providers, the risk from providers should be factored into your Early Warning analysis. Again, services exist to monitor service providers (as they do with other business partners) which can alert you to risks. Also ensure your service provider has some kind of notification SLA in place to give you a heads-up when they detect an attack on their infrastructure.

Customers tend to be one step removed, compared to business partners and service providers. Customers typically don’t have the same kind of direct access as partners or providers, but can still pose a risk, such as login credentials being stolen and used by an attacker. You can track customer issues by monitoring news feeds, SEC filings for breach disclosures, and data breach reporting sites such as DataLossDB. The proprietary threat intelligence services can also monitor specific high-risk customers.

> Like an “asset criticality” rating in a risk scoring calculation, you can tune the numbers a bit to reflect your professional judgement of the risk presented by each group.

The relationship between proximity and urgency is different for each company, and must be evaluated in terms of available resources. Clearly direct attacks are urgent, but the importance of various other constituencies depends on their access to critical information and their track records. Like an “asset criticality” rating in a risk scoring calculation, you can tune the numbers a bit to reflect your professional judgement of the risk presented by each group. A certain business partner might be a train wreck from a security standpoint, and warrant an increased proximity factor, due to the threat posed by any attack on them. On the other hand a mature and advanced business partner can likely handle most attacks, leaving you free to focus on different partners and customers.

Deploying the EWS

Now that we have covered the concepts behind the Early Warning System it's time to put them into practice. We start by integrating a number of disparate technology and information sources as the basis of the system – building the technology platform. We need the EWS to aggregate third-party intelligence feeds and scan for those indicators within your environment to highlight attack conditions. When we consider important capabilities of the EWS, a few major capabilities stand out:

1. **Open:** The job of the EWS is to aggregate information, which means it needs to be easy to get information *in*. Intelligence feeds are typically just data (often XML), which makes integration relatively simple. But consider how to extract information from other security sources such as SIEM, vulnerability management, identity, endpoint protection, and network security; and to get it all into the system. The point is not to build yet another aggregation point — it is to take whatever is important from each of those other sources and leverage it to determine Early Warning Urgency.

2. **Scalable:** You will use a lot of data for broad Early Warning analysis, so scalability is an important consideration. But *computational scalability* is likely to be more important – you will be searching and mining the aggregated data intensively, so you need robust indexing.

3. **Search:** Early Warning doesn't lend itself to absolute answers. Using threat intelligence you evaluate the urgency of an issue and look for its indicators in your environment. So the technology needs to make it easy for you to search all your data sources and then identify at-risk assets based on the indicators you found.

4. **Urgency Scoring:** Early Warning is all about making *bets* on which attackers and attacks and assets are most important to worry about, so you need a flexible scoring mechanism. As we mentioned earlier, we are fans of quantification and statistical analysis; but for an EWS you need a way to weigh assets, intelligence sources, and attacks so you can calculate an urgency score. Which might be as simple as red/yellow/green.

Some other capabilities can be useful in the Early Warning process – including traditional security capabilities such as alerting and thresholding. Again, you don't know quite what you are looking for initially, but once you determine that a specific attack requires active monitoring you will want to set up appropriate alerts within the system. Alternatively, you could take an attack pattern and load it into an existing SIEM or other security analytics solution. Similarly, reporting is important, as you look to evaluate your intelligence feeds and your accuracy in pinpointing urgent attacks. As with more traditional tools customization of alerts, dashboards, and reports enables you to configure the tool to your requirements.

That brings us to the question of whether you can repurpose existing technology as an Early Warning System. Let's first look at the most obvious candidate: your SIEM/Log Management platform. Go back to the key requirements above and you will see that integration may be the most important criterion. The good news is that most SIEMs are built to accept data from a variety of different sources. The most significant impediment right now is the relative immaturity of threat intelligence integration. Go into the process with your eyes open, and understand that you will need to handle much of the integration yourself.

That brings us to the question of whether you can repurpose existing technology as an Early Warning System.

The other logical candidate is the vulnerability management platform — especially in light of its evolution toward becoming a more functional asset repository, with granular detail on attack paths and configurations. VM platforms aren't there yet — alerting and searching tend to be weaker due to their technological heritage. But over time we will see both SIEM and VM systems mature as legitimate security management platforms. In the meantime your VM system will feed the EWS, so make sure you are comfortable getting data out of it.

Big Data vs. "A Lot of Data"

Big Data means analysis via technologies like Hadoop, MapReduce, NoSQL, etc. These technologies are great, and they show tremendous promise for helping to more effectively identify security attacks. But they may not be the best choices for an Early Warning System.

While we are talking about the EWS platform, we need to address the elephant in the room: *Big Data*. We see that term used to market *everything* relating to security management and analytics. Any broad security analysis requires digesting, indexing, and analyzing a lot of security data. In our vernacular, Big Data means analysis via technologies like Hadoop, MapReduce, NoSQL, etc. These technologies are great, and they show tremendous promise for helping to more effectively identify security attacks.

But they may not be the best choices for an Early Warning System. Remember that SIEM technology evolved as vendors moved to purpose-built datastores and analysis engines because relational databases ran out of steam. The same concepts apply to the EWS, and the main requirement for the technology back-end is its ability to handle all the security data scalably. The underlying technology doesn’t matter much, so long as it can digest and process the data. We know there will be a mountain of data, from all sorts of places, in all sorts of formats. So focus on openness, scalability, and customization.

Turning Urgency into Action

Once you get an Early Warning alert you need to figure out whether it requires action, and if so what kind to take. Validation and remediation are beyond our scope here — we have already covered them in Malware Analysis Quant, Evolving Endpoint Malware Detection, Implementing and Managing Patch and Configuration Management, and other papers which examined the different aspects of active defense and remediation. So we will just touch on the high-level concepts here.

1. **Validate Urgency:** The first order of business is to validate the intelligence and determine the actual risk. The Early Warning alert was triggered by a particular situation, such as a weaponized exploit found in the wild or on vulnerable devices. Perhaps a partner network was compromised by a specific attack. In this step you validate the risk and take it from concept to reality by finding exposed devices, or perhaps evidence of attack or successful compromise. In a perfect world you would select an attack scenario and your EWS platform would mine the data and spit out a list of at-risk devices. And at some point the technology will reach this level of automation, but in the meantime we need to do a bunch of manual searching and mining to identify devices at risk.

2. **Remediate:** Once you determine an attack is legitimate and demands action, you will decide what action to take. If you haven’t seen the attack yet you can implement a workaround such as additional firewall/IPS rules or a more restrictive endpoint configuration for at-risk devices. Or perhaps you will just add a new SIEM rule to scan for the scenario identified in the warning, to alert immediately if the attack materializes. You need to balance being proactive against extra work defending against attacks that never happen. It’s obviously much better to be safe than pwned, but resource realities make figuring out what *not* to do as important as what to do.

Publicizing Quick Wins

Over time, as with any security discipline, you will refine the verification/validation/investigation process. Focus on what worked and what didn’t and tune the process accordingly. It can be a bit bumpy when you first deploy an EWS, as you receive a bunch of alerts which lead you down a bunch of dead ends — similar to your SIEM tuning process (loads of fun, as you remember). But security is widely regarded as overhead, so you need to focus on getting a Quick Win with any new security technology.

Over time, as with any security discipline, you will refine the verification/validation/investigation process. Focus on what worked and what didn’t and tune the process accordingly.

An EWS *will* find stuff you didn’t know about and help you get ahead of attacks. But that success story won’t tell itself, so when the process succeeds – likely early on – you will need to publicize it and tell success stories early and often. There are two key areas to focus on. The first is finding proof of an attack in progress, which you successfully remediate thanks to threat intelligence and the EWS. This illustrates that you *will* be compromised, so success is a matter of containing the damage and preventing data loss. The value of the EWS is in shortening the window between exploit and detection.

The second scenario which tells a great story is taking preemptive steps after a business partner is compromised. We have all sat in meetings to discuss, “Will that happen to us?” after learning of a breach at a business partner or competitor. If you can definitively say that you get threat intel on emerging attacks (particularly on competitors or partners) and then evaluate whether action needs to be taken, that allays the fear of senior management that Security has no idea what’s happening until it is already over — too late. Even better if you can discuss how preemptive workarounds and remediations, implemented in response to threat intelligence, blocked an actual attack.

You can also explain how threat intelligence helped you evolve security tactics based on what’s happening to other organizations. For instance, if you see what looks like a denial of service (DoS) attack on a set of web servers, but already know from your intelligence efforts that DoS is a frequent decoy to obscure exfiltration activities, you have better context to be more sensitive to exfiltration attempts. Finally, to whatever degree you quantify the time you spend remediating issues and cleaning up compromises, you can show how much you saved by using external feeds to refine your efforts and prioritize your activities.

Summary

Our adversaries have too many weapons at their disposal for anyone to expect to effectively secure *all* your information from *all* the attacks you face. So you need to be much smarter about what you do, and much more diligent about reacting quickly to attacks in progress. The last few years have seen a wave of security information management (SIEM) projects designed to help mine internal security data, watch for attack patterns, and identify attacks before attackers make off with the goodies. Many organizations have substantially improved their security postures with these investments.

The next wave of protection involves looking outside the walls of your own environment to leverage what's happening in the broader world, in order to better prioritize your efforts. The critical limitations of SIEM are the need to know what to look for, and only being able to react *after* it happens in your environment. Early Warning changes this with external threat intelligence. With a mushrooming variety of threat intelligence sources ready to detail attacks, malware, and tactics seen in the wild; organizations can not only look for attacks before they hit, but implement preemptive controls to guard against them.

An Early Warning System requires a mature security program. In order for these concepts to be useful, you need to already have a strong security data aggregation and analysis capability. You also must be able to respond quickly to attacks, with a refined operational process to implement protection without getting caught up in extensive change management shenanigans. Organizations which want to push their security programs to the next level can rapidly integrate information about emerging attacks, and gauge the vulnerability of partner networks, to continue closing the window between attack and detection. Ultimately that is the measure of success for security practitioners.

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 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:

- **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 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/>.