



Leveraging Threat Intelligence in Security Monitoring

Version 1.5

Released: February 21, 2014

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 Norse Corporation.



norse-corp.com

Norse is the leading innovator in the live threat intelligence security market. With the goal of transforming the traditionally reactive IT security industry, Norse offers proactive, intelligence-based security solutions that enable organizations to identify and defend against the advanced cyberthreats of today and tomorrow. Norse's synchronous, global platform is a patent-pending infrastructure-based technology that continuously collects and analyzes real-time, high-risk Internet traffic to identify the sources of cyber attacks and fraud. Norse is the only provider of live, actionable, cyberthreat intelligence that enables organizations to prevent financial fraud and proactively defend against today's most advanced cyber threats including zero day and advanced persistent threats.

Copyright

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



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

Leveraging Threat Intelligence in Security Monitoring

Table of Contents

Benefiting from the Misfortune of Others	4
Revisiting Security Monitoring	8
The New TI + SM Process	12
Quick Wins	16
Summary	20
About the Analyst	21
About Securosis	22

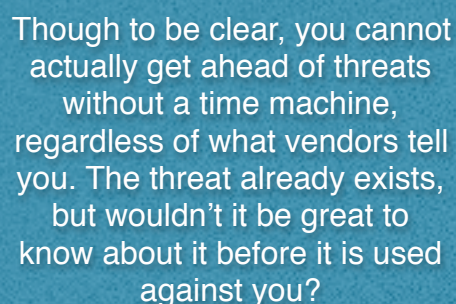
Benefiting from the Misfortune of Others

Threat intelligence (TI) is currently a hot topic because it promises to close the gap between attackers and defenders. Accordingly we have done considerable research into TI over the past year. We started by talking about the [Early Warning System](https://securosis.com/research/publication/building-an-early-warning-system)¹, a monitoring concept that leverages threat intelligence feeds to look for emerging attacks before they hit you. Then we dove into [the kinds of TI you can extract from network traffic](https://securosis.com/research/publication/network-based-threat-intelligence-searching-for-the-smoking-gun)², the ability to identify malicious IPs and senders by [gathering TI through email](https://securosis.com/research/publication/email-based-threat-intelligence-to-catch-a-phish)³, and finally a view of the external world through [EcoSystem TI](https://securosis.com/research/threat-intelligence-for-ecosystem-risk-management)⁴.

There are many different types of threat intelligence feeds and many ways to apply the technology — both to increase the effectiveness of alerting and to implement preemptive workarounds based on likely attacks observed on other networks. That's why we say threat intelligence enables you to *benefit from the misfortune of others*. By understanding attack patterns and other nuggets of information gleaned from attacks on other organizations, you can be better prepared when they come for you.

Though to be clear, you cannot actually get ahead of threats without a time machine, regardless of what vendors tell you. The threat already exists, but wouldn't it be great to know about it before it is used against you?

As the TV networks promote their summer reruns, "If you haven't seen it, it's new to you!"



Though to be clear, you cannot actually get ahead of threats without a time machine, regardless of what vendors tell you. The threat already exists, but wouldn't it be great to know about it before it is used against you?

Shortening the Window

One of the most compelling uses for threat intelligence is helping to detect attacks earlier. By looking for attack patterns identified via threat intelligence in your security monitoring and analytics processes, you can shorten the window between compromise and detection.

¹ <https://securosis.com/research/publication/building-an-early-warning-system>

² <https://securosis.com/research/publication/network-based-threat-intelligence-searching-for-the-smoking-gun>

³ <https://securosis.com/research/publication/email-based-threat-intelligence-to-catch-a-phish>

⁴ <https://securosis.com/research/threat-intelligence-for-ecosystem-risk-management>

This paper will go into depth on how to update your security monitoring process to integrate malware analysis and threat intelligence. We will be using parts of our [Network Security Operations Quant](#)⁵ and [Malware Analysis Quant](#)⁶ process maps to present an updated Threat Intelligence + Security Monitoring Process Model which brings the two ideas together.

Before we can dive into process maps we need to set the stage, by revisiting the kinds of threat intelligence highlighted in our research.

Threat Intelligence Sources

You can get effective threat intelligence from a number of different sources. We divide security monitoring feeds into four high-level categories:

1. Compromised Devices
2. Malware Indicators
3. Reputation
4. Command and Control Networks

Compromised Devices

The first category of TI provides external notification that a device is acting compromised by communicating with known bad sites or participating in botnet-like activities. Services are emerging to mine large volumes of Internet traffic to identify such devices. These services are no-touch: you don't need to install anything on your own network to get a verdict on devices within it.

How does it work? Intelligence providers penetrate botnets and monitor traffic on C&C networks. With this information they build lists of (compromised) devices that appear to be participating.

Of course these services might detect your own internal honeypots or other malware analysis activities, so make sure you have some means of determining which devices *should* show up and which shouldn't.

Malware Indicators

Malware analysis continues to mature rapidly, getting better and better at understanding exactly what malicious code does to devices. This enables you to define both technical and behavioral indicators to seek within your environment, as Malware Analysis Quant described in gory detail. This is essential because the central strategy of classical AV — file blacklisting — is no longer effective. We need new indicators to detect malware by what it *does* rather than what it looks like.

A number of companies offer information on specific malware samples. You can upload a hash of a malware file: if the recipient has seen it already they will recognize the hash and return the analysis on file; otherwise you upload the whole file for analysis. These services run malware samples through

⁵ <https://securosis.com/Research/Publication/network-security-operations-quant-report>

⁶ <https://securosis.com/research/publication/malware-analysis-quant-report>

proprietary sandbox environments and other analysis engines to figure out what they do, build detailed profiles, and provide comprehensive reports which include specific behaviors and indicators that can be integrated into monitoring platforms and security controls. These profiles enable you to look for the *behavior* of malware rather than depending on matching file hashes.

You can also draw conclusions from the kinds of indicators you find. Have the tactics represented by these indicators been tied to specific adversaries? Do these kinds of activities occur during reconnaissance, exploitation, or exfiltration? Internal analysis can enrich these indicators with additional context for better decisions about your best next step.

Reputation

Since its emergence years ago as a primary data source in the battle against spam, reputation data seems to have been integrated into every security control. The most common reputation data is based on IP addresses and provides a dynamic list of known bad and/or suspicious addresses. This has a variety of uses — learning that a partner's IP address has been compromised, for instance, should set off alarms, especially if that partner has a direct connection to your network. Traffic to known malware distribution sites, phishing sites, command and control nodes, spam relays, and other sites with bad reputations should be investigated because those kinds of activity are caused by compromised devices.

IP reputation has evolved since its introduction, now featuring “scores” to assess the relative maliciousness of an address, as well as being able to factor in additional context like Tor nodes/anonymous proxies, geo-location and device ID to further refine the reputation. Additionally, the providers of reputation data continually assess the addresses to either reduce or increase the risk of the address depending on recent activity, as opposed to permanently blacklisting an address.

Pretty much everything in your environment can (and should) have a reputation. Devices, URLs, domains, and files, for starters. If you have traffic going to a known bad site, weird traffic coming from a vulnerable contractor-owned device, or even a recognized bad file showing up when a salesperson connects to the corporate network, you have something to investigate. If something in your environment develops a bad reputation — perhaps as a spam relay or DoS attacker — you need to know ASAP, hopefully *before* your entire network gets blacklisted.

These feeds track global C&C traffic and pinpoint malware originators, botnet controllers, and other IP addresses and sites you should be looking for as you monitor your environment.

C&C Traffic Patterns

One specialized type of reputation which is often packaged as a separate feed is intelligence on command and control (C&C) networks. These feeds track global C&C traffic and pinpoint malware originators, botnet controllers, and other IP addresses and sites you should be looking for as you monitor your environment. As mentioned above, these services can also help identify suspicious devices within your network

that are communicating with malware controllers. Integrating this kind of network-based threat

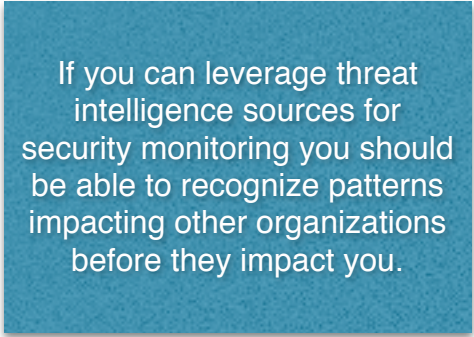
intelligence with an egress firewall or web filter might enable you to prevent exfiltration, or begin more aggressive monitoring (perhaps capturing network packets) to identify exactly what attackers are doing.

Of course advanced attackers do not make analyzing C&C traffic easy. They work hard to obscure their communications, using compromised devices with 'good' reputations as C&C nodes to circumvent reputation filters, and changing locations frequently using a variety of sophisticated Domain Generating Algorithms (DGA). Accurately identifying C&C traffic is currently a kind of black magic, but it is a critical aspect of intelligence.

Challenges of Using TI for Security Monitoring

That all sounds cool, right? If you can leverage threat intelligence sources for security monitoring you should be able to recognize patterns impacting other organizations before they impact you. That's the concept, anyway — as always there are challenges to making it work.

1. **Integration of the data:** TI isn't much good until you get it into your security environment. The first step is to make sure any threat feeds can be integrated easily.
2. **Updating rules/alerts/reports:** Once the data is in you need to actually look for the specified patterns and indicators. That requires a bunch of work on an ongoing basis to update the security monitoring platform. Realistically, you need to automate the process — attacks appear and change too rapidly for manual updates to keep up.
3. **Triage and validation:** Finally, with TI and updated rules, you will start seeing alerts based on emerging patterns and indicators. Someone still needs to validate attacks and take action. Given the severe security resource and skill constraints in many organizations, unless this process is adequately resourced much of this effort won't make much difference to your ability to detect attackers earlier.



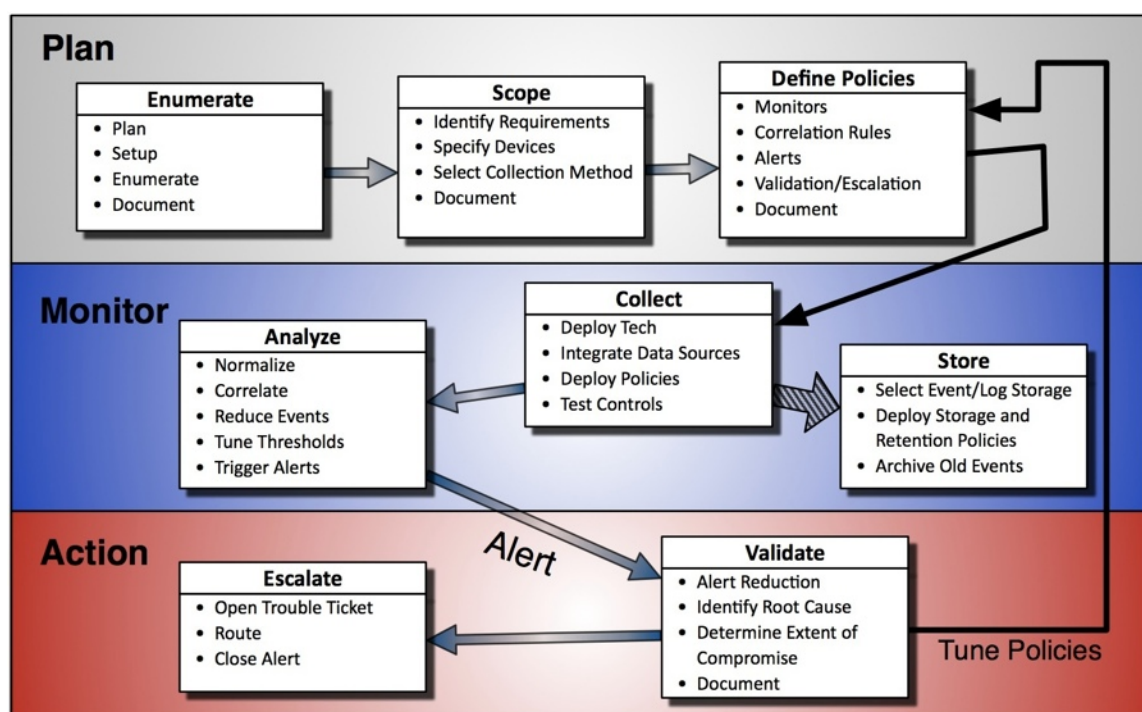
If you can leverage threat intelligence sources for security monitoring you should be able to recognize patterns impacting other organizations before they impact you.

But none of these challenges is insurmountable. It is a question of being aware of the demands at the beginning of the process, and factoring the realities in to avoid issues down the line.

Revisiting Security Monitoring

Now let's revisit our definition of network security monitoring, for context on how monitoring processes need to adapt to leverage threat intelligence. In our Network Security Operations Quant research we detailed all the gory tasks involved in monitoring. That research was about firewalls and IPS devices, but it is easy enough to expand those concepts to other key devices.

The Network Security Monitoring Process



We can break the key activities for monitoring network security devices into three distinct phases:

Phase 1: Plan

In this phase we define the depth and breadth of our monitoring activities. These are not one-time tasks, but processes to revisit quarterly and after incidents that trigger a policy review.

1. **Enumerate:** Find all the security, network, and server devices that are relevant to the security of the environment.

2. **Scope:** Decide which devices are in scope for monitoring. This involves identifying the asset owner; profiling the device to understand data, compliance, and policy requirements; and assessing the feasibility of collecting data from it.
3. **Develop policies:** Determine the depth and breadth of the monitoring process. This consists of two parts: organizational policies (which devices will be monitored and why) and device & alerting policies (which data will be collected from devices and how often). These policies govern the activities of any network, security, computing, application, or data capture/forensics device.

More on Policies

For device types in scope of the monitoring process, develop device and alerting policies to detect potential incidents which require investigation and validation. These policies require a QA process to

A tuning step must be built into process of managing alerting policy definitions — over time alert policies need to evolve to track changes in targets to defend and adversaries' tactics.

test and refine the effectiveness of alerts. A tuning step must be built into process of managing alerting policy definitions — over time alert policies need to evolve to track changes in targets to defend and adversaries' tactics.

Finally, monitoring is part of a larger security operations process, so policies are required for workflow and incident response. They define how monitoring information is leveraged by other operational teams and how potential incidents are identified, validated, and investigated.

Phase 2: Monitor

This phase puts monitoring policies to use, gathering data and analyzing it to identify areas for validation and investigation. All collected data is stored for compliance, trending, and reporting.

1. **Collect:** Collect alerts and log records based on the policies defined in the Plan phase. This can be performed within a single-element log/event manager or abstracted into a broader Security Information and Event Management (SIEM) system for multiple devices and device types.
2. **Store:** Collected data must be stored for future access, for both compliance and forensics.
3. **Analyze:** The collected data is analyzed to identify potential incidents based on alerting policies defined in Phase 1. This may involve numerous analysis techniques, including simple rule matching (resource availability, activity, network traffic to/from resource, time-based rules, etc.) and/or a correlated combination of many rules to detect likely attack patterns.

Phase 3: Action

When an alert fires you need to investigate and determine whether further action is necessary.

1. **Validate/Investigate:** Investigate and validate alerts. Is it a false positive? Is it a real issue that requires further action? If the latter escalate into the incident response process. If this was not a 'good' alert do policies need tuning?
2. **Action/Escalate:** Take action to remediate the issue. This may involve a hand-off or escalation to Operations.

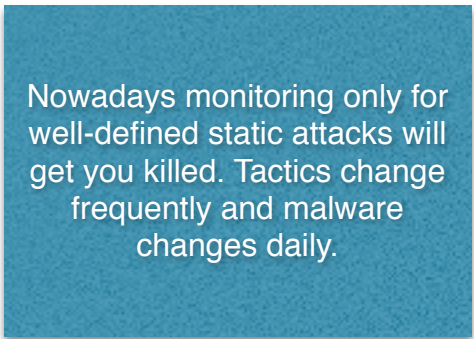
After validating a few alerts you can determine whether policies must be changed or tuned. Tuning policies must be a recurring feedback loop rather than a one-time activity — networks and attacks are dynamic and require ongoing diligence to ensure monitoring and alerting policies remain relevant and sufficient.

What Has Changed

Security monitoring has undergone significant change over the past few years. We detailed many of the changes in our [Security Management 2.5](#) research. The first big change is the need to analyze much more data from many more sources — we will go into detail below.

Additionally, the kind of analysis performed on the collected data is different. Setting up rules for a security monitoring environment was traditionally a static process — you would build a threat model and then define rules to look for that kind of attack. This approach requires you to know what to look for. For reasonably static attacks this approach can work.

But nowadays monitoring only for well-defined static attacks will get you killed. Tactics change frequently and malware changes daily. Sure, there are always activity patterns that indicate a likely attack, but attackers have gotten proficient at evading traditional SIEMs. Security practitioners need to adapt detection techniques accordingly, and threat intelligence provides a means to stay on top of emerging attacks that may not fit into a common pattern.



Nowadays monitoring only for well-defined static attacks will get you killed. Tactics change frequently and malware changes daily.

To detect these new attacks you need to rely much more on looking for variation from normalcy in collected data to trigger alerts and investigation. But how can you perform that kind of analysis on what might be dozens of disparate data sources? Big data, of course. Kidding aside, big data is actually the answer, and it is no overstatement to say big data technologies will fundamentally change security monitoring over time.

Broadening Data Sources

In [Security Management 2.5: Platform Evolution](https://securosis.com/blog/security-management-2.5-platform-evolution)⁷, we explained that to keep pace with advanced attackers security monitoring platforms must *do more with more data*. More data to analyze opens up very interesting possibilities. You can integrate data from identity stores to trace behavior back to users. You can pull information from applications to look for application misuse, or visitors gaming legitimate application functionality such as search and shopping carts. You can pull telemetry from server and endpoint devices to search for specific indicators of compromise — which might represent a smoking gun and flag a successful attack.

We have always advocated collecting more data, even though the monitoring platforms had trouble scaling analysis to higher volumes. As we mentioned, security monitoring platforms increasingly leverage advanced data stores and support much different (and more advanced) analytics to detect patterns among many different data sources. This makes all that data much more useful for detecting attacks.

That doesn't mean the tools will magically do instantaneous big data analytics on a zillion data sources overnight — monitoring platforms are just now integrating these technologies. But the future is promising, and this kind of unstructured analysis is critical to detecting nimble attacks from innovative attackers.

⁷ <https://securosis.com/blog/security-management-2.5-platform-evolution>

The New TI + SM Process

Now we need to put these concepts together, for a better feel for how threat intelligence (TI) fits into the mix by opening up a wealth of additional analysis possibilities. As you integrate threat intelligence into your security monitoring (SM) process, you can generate more accurate alerts from your security monitoring platform, lowering the signal to noise ratio because the alerts are based on what is actually happening in the wild.

As you integrate threat intelligence into your security monitoring (SM) process, you can generate more accurate alerts from your security monitoring platform, lowering the signal to noise ratio because the alerts are based on what is actually happening in the wild.

Developing Threat Intelligence

Before you can leverage TI in SM you need to gather and aggregate the intelligence in a way that can be cleanly integrated into the SM platform. We have already mentioned four different TI sources; so let's go through them for a better feel for how to gather information.

1. **Compromised devices:** When we talk about actionable information, a clear indication of a compromised device is the most valuable intelligence: the proverbial smoking gun. There are many ways you might conclude a device is compromised. The first is by monitoring network traffic for clear indicators of command and control traffic originating from the device, such as DNS requests whose frequency and content indicate a domain generating algorithm (DGA) to find botnet controllers. Monitoring traffic from the device can also show files or other sensitive data being sent, indicating exfiltration or (via network traffic analysis) a remote access trojan. As described above, you could also penetrate major bot networks to monitor their traffic, which can enable you to identify bots without on-premise monitoring.
2. **Malware indicators:** As described in [Malware Analysis Quant](https://securosis.com/research/publication/malware-analysis-quant-report)⁸, you can build a lab and do both static and dynamic analysis of malware samples to identify specific indicators of how it compromises devices. This is obviously not for the faint of heart — thorough and useful analysis requires significant investment, resources, and expertise.
3. **Reputation:** IP reputation data (usually provided as a list of bad IP addresses with relative risk scores) can trigger alerts, and may even be used to block outbound traffic headed for bad

⁸ <https://securosis.com/research/publication/malware-analysis-quant-report>

networks when integrated with network security devices. You can also alert and monitor on the reputations of other resources — including URLs, files, domains, and even specific devices. Of course reputation scoring requires analyzing a large amount of network traffic — a significant

chunk of the Internet — and tens of millions of files to observe useful patterns among emerging attacks.

Given the demands of gathering sufficient information to analyze, and the challenge of detecting and codifying appropriate patterns, most organizations look for a commercial provider to develop and provide this threat intelligence as a feed that can be directly integrated into security monitoring platforms.

Given the demands of gathering sufficient information to analyze, and the challenge of detecting and codifying appropriate patterns, most organizations look for a commercial provider to develop and provide this threat intelligence as a feed that can be directly integrated into security monitoring platforms. This enables internal security folks to spend their time figuring out how the TI applies to their environment to make alerts and reports more relevant. Internal security folks also need to validate and assess TI sources on an ongoing basis because accuracy and timeliness can vary amongst providers over time. TI must be kept fresh to be valuable.

Evolving the Monitoring Process

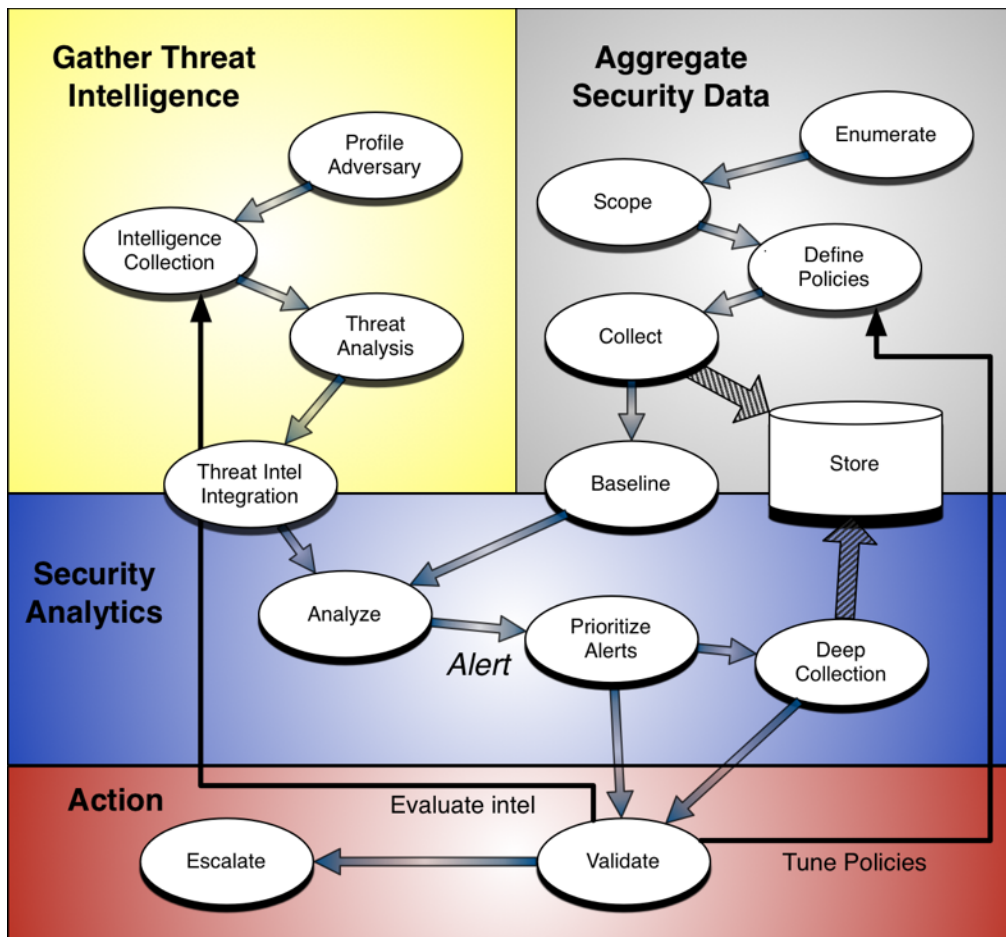
Armed with a variety of threat intelligence sources, you need to take a critical look at your security monitoring process to figure out how it needs to change to accommodate them. First let's turn back the clock to revisit the early days of SIEM. Historically SIEM products were driven by a defined ruleset to trigger alerts, but that requires you to know what to look for *before the attack hits*. You cannot profile advanced attacks before you observe them, so you cannot count on knowing what to look for. We need to think differently about monitoring.

We continue to recommend identifying normal patterns on your network with a baseline, and then building policies to detect deviation. Supplement baselines with emerging indicators identified by TI.

But don't minimize the amount of work required to keep all these rules current in the security monitoring platform — both types of rules are increasingly dynamic. Baselines are constantly changing, so your definition of 'normal' needs ongoing scrutiny and revision. Threat intelligence is inherently dynamic. You need to look for new indicators and network traffic patterns in near real time for any hope of keeping up with hourly changes of C&C nodes and malware distribution sites. Significant automation is required if you want your organization's monitoring environment to keep pace with attackers.

The New Security Monitoring Process Model

Against this backdrop we need to evolve the security monitoring process to both gather threat intelligence and integrate TI. With the new process below you can more effectively handle your organization's rapidly changing attack surface and improve your monitoring results.



Gather Threat Intelligence

The new addition to the process is gathering threat intelligence. As described above, there are a number of different sources you should integrate into your security monitoring environment. Here are brief descriptions of the steps of a broad threat intelligence program:

1. **Profile adversary:** As we covered in the [CISO's Guide to Advanced Attackers](https://securosis.com/research/publication/the-cisos-guide-to-advanced-attackers)⁹, it is critical to understand who is most likely to attack you so you can profile their TTPs (Tactics, Techniques, and Procedures).
2. **Gather samples:** The next step is to gather a large amount of data, which can be analyzed to define specific indicators from the TI feed (IP addresses, malware indicators, device changes, executables, etc.).
3. **Analyze data and distill threat intelligence:** Once the data is aggregated you can mine the repository of samples to identify suspicious activity and distill that down into patterns and indicators of attacks being seen in the wild. Given the rapidly changing nature of attacks, you'll also need ongoing validation and testing of the TI indicators to ensure they remain accurate and timely.

⁹ <https://securosis.com/research/publication/the-cisos-guide-to-advanced-attackers>

Aggregate Security Data

The steps involved in aggregating security data are largely unchanged. You still need to enumerate which devices to monitor in your environment, scope the kinds of data you will get from them, and define collection policies and correlation rules. Then you can move on to actively collecting data and storing it in a repository to allow flexible, fast, and efficient analysis and searching.

Security Analytics

The security monitoring process now has two distinct sources to analyze, correlate, and alert on — external threat intelligence and internal security data — so this phase needs some changes.

1. **Automate TI integration:** Given the volume of TI information and its rate of change, the only way to effectively leverage external TI is to automate ingestion of data into the security monitoring platform; you also need to automatically update alerts, reports, and dashboards.
2. **Baseline environment:** You don't really know what kinds of attacks you are looking for yet, so you will want to gather a baseline of 'normal' activity within your environment and then look for anomalies, which might indicate compromise and warrant further investigation.
3. **Analyze security data:** The analysis process still involves normalizing, correlating, reducing, and tuning the data and rules to generate useful and accurate alerts.
4. **Alert:** When a device shows one or more indicators that it has been compromised, an alert triggers.
5. **Prioritize alerts:** Prioritize alerts based on the number, frequency, and types of indicators which triggered them; use these priorities to decide which devices to further inspect and in what order. Integrated threat intelligence can help by providing additional context allowing responders to prioritize the threats so that analysts can investigate the highest risk cases first.
6. **Deep collection:** Depending on the priority of the alert, you might want to collect more detailed telemetry from the device, and perhaps start capturing network packet data to and from it to facilitate validation and identification of possible compromise — as well as to facilitate forensic investigation if it comes to that.

Action

Once you have an alert and have gathered data about the device and attack, you need to determine whether the device was actually compromised or the alert was a false positive. If a device has been compromised you need to escalate, either to an operations team for remediation/clean-up, or to an investigation team for more detailed incident response and analysis. To ensure both processes improve constantly, you should learn from each validation step: critically evaluate the intelligence, as well as the policy and/or rule that triggered the alert.

Quick Wins

So far we have explained the value of threat intelligence (TI) and discussed how combining it with security monitoring (SM) can help detect attacks faster — based on not only what *you* are seeing, but also on what is happening elsewhere. But that is all still theoretical. How can *you* apply an integrated process to shorten the window between compromise and detection? How can you get a quick win for the integration of TI and SM, to build momentum for your efforts and get your peers in Operations involved? Finally, how can you reliably turn a quick win into sustainable leverage, producing increased accuracy and better prioritization of alerts from a SM platform?

Let's say you work for a big retailer with thousands of stores. You do hundreds of millions of credit card transactions a month, and have credit card data for tens of millions of customers. Your organization is a high-profile target, and one 'benefit' being a large Tier 1 merchant is that the assessors show up pretty much every quarter to make sure your ROC (Report On Compliance) is still accurate. You can play the compensating control fandango to a point (and you do), but senior management understands the need to avoid becoming the latest object lesson on data breaches. So senior management gives you a bunch of resources and money to spend, with the clear expectation that you will ensure private data remains private.

But this is the real world and your organization is a big company. The nature of your business involves technology assets all over the place and employees who come and go, especially around the holidays. They all have access to the corporate network, and no matter how much time you spend educating them folks make mistakes. This long preamble is just to illustrate the reality of your situation. Your odds of keeping all attackers out range from nil to less than nil. So security monitoring is a key aspect of your plan to detect attackers before they can do too much damage. The good news is that you already aggregate a bunch of log data, mostly because you need to (thanks, PCI). You can build on this foundation and integrate TI to start looking for attack patterns and other suspicious activity others have seen. This can give you early warning of imminent attacks.

Your odds of keeping all attackers out range from nil and less than nil. So security monitoring is a key aspect of your plan to detect attackers before they can do too much damage.

Low Hanging Fruit

With any new technology project you should show value quickly and then parlay it into sustainable advantage. Let's focus on obvious stuff that integrating TI into the monitoring process can yield to

provide your quick wins. There are a couple areas to consider but the path of least resistance tends to be finding devices which are already compromised and remediating them quickly. A couple fairly reliable TI sources can yield this kind of information.

Once you identify the suspicious device you need to collect detailed data. Optimally you will get deep endpoint (or server) telemetry including all file activity, registry and other configuration values, and a forensic capture of the device. Though a new generation of devices (think Internet of Things) may not be able to provide this kind of detailed endpoint telemetry. Thus, the reputation data, traffic anomalies, and context from TI information may be the only way to assess these devices. For a broader view of what's going on you will also want to capture network traffic to and from it. Armed with that information you can search for specific malware indicators and other clear manifestations of attack, to understand the compromise and remediate it quickly.

Adding Value to Baselines

At this point you have likely found some devices with issues, and acted decisively to remediate the issues and contain the damage. Once the compromised devices you detected are dealt with, you can get a bit more strategic about what to look for. You have been collecting data for a while (thanks again, PCI), so you can now develop a reasonable baseline of normal activity. Of course you will exclude compromised devices, and you will then be able to set alerts on activity that is not normal. That is Security Monitoring 201, and not particularly novel.

But if you integrate TI into the process you can get a more accurate picture of why a device isn't acting normally. Most current malware isn't easy to find if you haven't seen the indicators before, but you can leverage TI to look for emerging attacks seen by other organizations.

Let's make this a bit more tangible by going back to our example of the large retailer. As with most big companies, you have a bunch of externally accessible devices, which serve up a variety of things to customers. Not all of them have access to mission-critical data (unless you screwed up your network segmentation), so they may not get much scrutiny or monitoring focus. But you can still track traffic in and out to see if or when they violate your established baselines for network traffic flow.

Once you find a suspicious device, integrated TI enables you to search for indicators and other behavior patterns you wouldn't otherwise have known to look for, which address a key limitation in traditional monitoring processes — only looking for things you know about.

If you see an externally accessible web server start sending traffic to a bunch of other devices within its network segment that is probably suspicious. Normally such web servers only send internal traffic to the application server farm running their applications. Communicating with other internal hosts is likely abnormal, so you start pulling additional telemetry from devices and capturing traffic.

Once you find a suspicious device, integrated TI enables you to search for indicators and other behavior patterns

you wouldn't otherwise have known to look for, which address a key limitation in traditional monitoring processes — only looking for things you know about. Integrated TI changes those dynamics, allowing you to identify traffic heading to an emerging botnet. Or perhaps detecting new files associated with a little-known malware kit. You might not have seen these attacks before, but your TI provider probably has.

Without TI, when you identify a suspicious device, you are basically shooting in the dark. You have a device acting strangely but don't know why. You can reset the device to your gold standard, but that doesn't help with any other devices compromised by the same attack, and you are much less likely to learn how to avoid that compromise next time. In this scenario TI helps you confirm the attack faster and eradicate it from your environment because you know what to look for.

Adding Adversary Analysis

As we continue to add sophistication to the TI you integrate into the SM process, you can next look to integrate adversary-specific data. In our retailer example you know credit card data is the most valuable stuff you have, and it draws a certain class of adversaries to target your organization. Like any other organization, these attackers have traits and patterns that can be profiled. They likely use certain types of malware and a handful of botnets, which enables you to look for them specifically.

Your TI provider can provide specific information about these actors. More sophisticated TI offerings add customized scoring metrics (based on types of devices and/or threat actors) that provide another attribute for prioritizing alerts, as described in the new security monitoring process. If you know you are likely to be targeted by organized crime faction X (CFX), and you get an alert of a common attack from CFX, you can make that a top priority. This makes TI more valuable operationally and helps you stay focused on the most significant risk.

Even without specific intel 'scores', being able to look at information about specific actors lets you turn the tables, if only a bit. You can proactively look for the attacks commonly associated with those actors and identify issues before compromised devices starts behaving egregiously.

Sustainable Wins

But a manual approach is not good enough. To achieve the true potential of integrated TI and SM you need to remove humans from the process — at least the front end.

Regardless of how many threat intelligence sources are integrated into your security monitoring process, you can get a quick win by having humans mine security data for TI indicators manually. It is not particularly efficient but it is quite possible, and the way most early TI offerings provided value. But a manual approach is not good enough. To achieve the true potential of integrated TI and SM you need to remove humans from the process — at least the front end. This automation has 3 facets:

1. **Integration of machine data:** You cannot use the cool automated capabilities of a security monitoring platform if you cannot get the data into it. So the first step in automating the process is to feed data into the system using either the monitoring platform's API or a common data format such as [STIX](http://stix.mitre.org/)¹⁰. It is too early to know the ultimate TI data format winner so be sure your TI process/provider can pump data into the monitoring platform you use without extensive manual effort.
2. **TI-based alerts:** Manually configuring rules and alerts based on TI doesn't scale. Fully taking advantage of the integration requires you to start looking for indicators immediately once new indicators from the TI service are fed into the monitoring platform. Of course fully automated rule updating can consume significant compute resources and may throw a bunch of alerts — at least until you tune the system — so you may want a different type or priority of TI-generated alerts to reduce the time you waste on false positives. Just keep in mind that getting a bunch of alerts don't mean they are *all* bogus — which is why prioritization is key, as discussed above.
3. **TI efficiency reports:** Finally you will want to be able to generate reports about the number of alerts generated from each TI source, without hurting your brain building them. These reports show value from the TI investment — especially if you can show how TI identified an attack earlier than you would have detected it otherwise as a quick win. Additionally, if you use multiple TI vendors, these reports enable you to compare them based on actual results. Optimizing TI spending can save real money.

TI can make a huge difference by increasing the accuracy of alerts and prioritizing them more effectively — to help your staff verify, investigate, and remediate the most dangerous attacks.

Finally, we need to acknowledge the difficulty of finding sufficient carbon-based resources to keep pace with the number and sophistication of attacks. Fortunately TI can make a huge difference by increasing the accuracy of alerts and prioritizing them more effectively — to help your staff verify, investigate, and remediate the most dangerous attacks.

¹⁰ <http://stix.mitre.org/>

Summary

Our summary from the [Early Warning System](#)¹¹ paper written in early 2013 is still very relevant today:

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 now look for attacks before they hit, as well as implement preemptive controls to guard against them.

With a mushrooming variety of threat intelligence sources ready to detail attacks, malware, and tactics seen in the wild; organizations can now look for attacks before they hit, as well as implement preemptive controls to guard against them.

This comes down to making your security monitoring smarter by leveraging visibility outside your little corner of the world. The increasing availability and maturity of threat intelligence sources, and their ability to quickly integrate into your rule set, will make monitoring efforts more efficient and effective.

The good news is that with better automation, the benefits of integrating threat intelligence and security monitoring become accessible to companies with less mature security programs, which should be exciting to practitioners who cannot keep pace with attackers in today's very dynamic environment.

If you have any questions on this topic, or want to discuss your situation specifically, feel free to send us a note at info@securosis.com or ask via the Securosis Nexus (<http://nexus.securosis.com/>).

¹¹ <https://securosis.com/research/publication/building-an-early-warning-system>

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:

- **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/>.