



Applied Threat Intelligence

Version 1.5

Released: March 27, 2015

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 Intel Security, whose support
allows us to release it for free.
All content was developed independently.**



www.intelsecurity.com

McAfee is now part of Intel Security. With its Security Connected strategy, innovative approach to hardware-enhanced security, and unique McAfee Global Threat Intelligence, Intel Security is intensively focused on developing proactive, proven security solutions and services that protect systems, networks, and mobile devices for business and personal use around the world. Intel Security is combining the experience and expertise of McAfee with the innovation and proven performance of Intel to make security an essential ingredient in every architecture and on every computing platform. The mission of Intel Security is to give everyone the confidence to live and work safely and securely in the digital world.

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/>



Applied Threat Intelligence

Table of Contents

Defining Threat Intelligence	4
Use Case #1: Preventative Controls	9
Use Case #2: Security Monitoring	14
Use Case #3: Incident Response/Management	19
Building a TI Program	24
Summary	29
About the Analyst	30
About Securosis	31

Defining Threat Intelligence

Security practitioners have basically failed to keep pace with adversaries for the past decade. It is a sad story but true. So threat intelligence (TI) has garnered considerable interest as a counter to adversaries launching new attacks using new techniques, and an alternative to the broken negative security model. Looking specifically for attacks you have seen is an excellent way to remain defenseless against all the new attacks you *haven't* seen yet. If your organization hasn't seen the new attacks and updated your controls and monitors to look for the new patterns, you are out of luck without some way of accounting for them.

Let's daydream for a minute. What if you could see attacks without actually being attacked yourself? What if you could benefit from the experience of higher-profile targets, learn what adversaries are trying against those organizations, and then watch for those patterns in your own environment? That would improve your odds of detecting and preventing attacks. It doesn't put defenders on an even footing with attackers, but it certainly helps.

Knowing what attacks may be coming at you doesn't help if your security operations functions cannot detect the patterns, block the attacks, or use the data to investigate possible compromise.

So what's the catch? It is easy to buy data but hard to take full advantage of it. Knowing what attacks may be coming at you doesn't help if your security operations functions cannot detect the patterns, block the attacks, or use the data to investigate possible compromise. Without those capabilities it's all just more useless data, and you already have plenty of that. As we discussed in detail in both [Leveraging Threat Intelligence in Security Monitoring](#) and [Leveraging Threat Intelligence in Incident Response/Management](#), TI can only help if your security program evolves to take advantage of intelligence data.

As we wrote in our TI + SM paper:

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.

But TI is not just useful for security monitoring and analytics. You can leverage it in almost every aspect of your security program. This paper will briefly revisit how processes need to change (as discussed in the papers mentioned above) and then focus on how to *use* Threat Intelligence to improve your ability to detect, prevent, and investigate attacks. Evolving your processes is great. Impacting your security posture is better. *A lot better.*

Defining Threat Intelligence

We cannot write about TI without acknowledging that, with a broad enough definition, pretty much any security data qualifies as threat intelligence. New technologies like anti-virus and intrusion detection (yes, that's sarcasm, folks) have been driven by security research data since they emerged 10-15 years ago. Those `.dat` files you (still) send all over your network? Yup, those are TI. The IPS rules and vulnerability scanner updates your products download? That's all TI too.

Over the past couple years we have seen new kinds of TI sources emerge — including IP reputation, Indicators of Compromise (IoC), command and control patterns, etc. There is a lot of data out there. And that's great, because without this raw material you have nothing to work with but what you see in your own environment.

Let's throw some stuff against the wall to see what sticks. Here is a starter definition of threat intelligence:

Threat Intelligence is security data that provides the ability to prepare to detect, prevent, or investigate emerging attacks before your organization is attacked.

That definition is intentionally quite broad because we don't want to exclude any interesting security data. Notice this definition doesn't limit TI to external data either, although in most cases TI is externally sourced. Organizations with very advanced security programs can perform proactive research on potential adversaries, and develop proprietary intelligence to identify likely attack vectors and techniques, but even those advanced organizations rely on third-party data sources to make internal tools and processes more effective.

Adversary Analysis

A good place to start your threat intelligence process is by figuring out who your adversaries are, because the attacks you see vary greatly based on your attackers' missions, and their assessment of the easiest and most effective ways to compromise your environment. The analysis is comprised of a few steps:

- **Evaluate the mission:** Start by learning what's important in your environment so you can identify interesting targets. They usually break down into a few discrete categories — including intellectual property, protected customer data, and business operations information.
- **Profile the adversary:** To defend yourself you need to know not only what adversaries are likely to look for, but also what tactics different types of attackers typically use. Figure out which categories of attacker you are likely to face. Types include unsophisticated (using widely available tools), organized crime, competitors, and state-sponsored. Each class has a different range of capabilities.
- **Identify likely attack scenarios:** Based on the adversary's probable mission and typical tactics, put on your attacker hat to figure out which path you would most likely take to achieve it. Regardless of whether the attack is in progress or has already taken place, you are trying to assess and contain the damage. Hopefully considering attack scenarios and various paths to achieve their mission will prove or disprove your hypothesis.

Figure out which categories of attacker you are likely to face. Types include unsophisticated (using widely available tools), organized crime, competitors, and state-sponsored.

Remember that you do not need to be exactly right about the scenario. You need to make assumptions about what the attacker will do, and you cannot predict their actions perfectly. Your objective is to get a head start on response, narrowing down investigation by focusing on specific devices and attacks. Nor do you need a 200-page dossier on each adversary — focus on information needed to understand an attacker and what they are likely to do. That kind of information will improve your security posture.

Collecting Data

Next start gathering data to help identify and detect the activity of these potential adversaries in your environment. You can get effective threat intelligence from a number of different sources. We divide security monitoring feeds into five high-level categories:

- **Compromised devices:** This feed provides external notification that a device is suspiciously 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.
- **Malware indicators:** Malware analysis continues to mature rapidly, getting better and better at understanding exactly what malicious code does to devices. This data enables you to define both technical and behavioral indicators to search out within your environment, as described in gory detail in [Malware Analysis Quant](#).
- **IP reputation:** The most common reputation data is based on IP addresses, and provides a dynamic list of known bad and/or suspicious addresses. IP reputation has evolved since

its introduction, now featuring scores to reflect the relative maliciousness of each address. Reputation services may also factor in additional context such as Tor nodes & anonymous proxies, command and control indicators, geolocation, and device ID, to further refine reputation.

- **Command and control networks:** One specialized type of reputation assessment which is often packaged as a separate feed is intelligence on Command and Control (C&C) networks. These feeds track global C&C traffic to pinpoint malware originators, botnet controllers, and other IP addresses and sites you should watch for as you monitor your environment.
- **Phishing messages:** Current advanced attacks tend to start with a simple email. Given the ubiquity of email and the ease of adding links to messages, attackers typically find email the path of least resistance to a foothold in your environment. Isolating and analyzing phishing email can yield valuable information about attackers and their tactics.

What Has Changed

We have had many of these data sources for years. So why are we talking about TI as a separate function? Because attackers' increasing sophistication requires us to leverage more and better data to keep pace.

We started seeing advanced security organizations staffing up their own threat intelligence groups a couple years ago. Those teams are tasked with understanding the organization's attack surface, figuring out what's at risk, and deciding what is most likely to be attacked.

We started seeing advanced security organizations staffing up their own threat intelligence groups a couple years ago. Those teams are tasked with understanding the organization's attack surface, figuring out what's at risk, and deciding what is most likely to be attacked. They provide context for which of the countless threats out there actually need to be dealt with; and what needs to be done to prevent, detect, and/or investigate potential attacks. These internal TI organizations consume external data to supplement internal collection and research efforts, and have been willing to pay for it, which created a new market for security data.

Another key change in the threat intelligence landscape has been the emergence of standards, specifically STIX and TAXII, which enabled quicker and better integration of TI into security processes. STIX provides a common data format for interchange of intelligence, and TAXII adds a mechanism and protocols to send it from originators to consumers. Before these standards organizations needed to perform custom integrations with all their active controls and security monitors, which were ponderous and didn't scale.

Addressing the Challenges

You just hit the EZ Button, gather some threat intelligence, and find the attackers in a hot minute, leaving plenty of time for golf. Awesome, right? Unfortunately it doesn't actually work like that. Threat intelligence is an emerging capability within security programs. So as an industry we need to overcome a few challenges to operationalize this approach:

- **Aggregate the data:** Where do you collect intelligence? You already have systems that can and should automatically integrate intelligence, and use it within rules or an analytics engine. The more automation the better, so personnel can focus on preventing attacks and figuring out what happened.
- **Analyze the data:** How can you know what's important within the massive quantity of data at your disposal? You need to tune your intelligence feeds and refine rules in your controls and monitors over time. As you leverage intelligence in your security program you get a feel for what works and what isn't so useful.
- **Actionable data:** This takes TI to the next level, with tools automatically updating controls and searching your environment based on threat intelligence feeds. Hopefully you can block attacks and identify attack indicators before an attacker exfiltrates your data. Existing tools such as firewalls, endpoint security, and SIEM can and should leverage threat intelligence. You will also want your forensics tools to play along, with the ability to leverage external intelligence.
- **False positives/false flags:** Unfortunately threat intelligence is still more art than science. See if your provider can prioritize or rank alerts. Then you can use the most urgent intelligence earlier and more extensively. Another aspect of threat intelligence to be careful of is disinformation. Many adversaries shift tactics, borrowing from other adversaries to confuse you. That is another reason not to simply profile an adversary, but to cross-reference against other information to make sure that adversary makes sense in context.

You now have a decent idea of what we mean by threat intelligence, so next we will focus on how TI can be used effectively in common use cases in the threat management lifecycle. Our lifecycle starts with preventing the attack (more realistically, trying to) with active security controls on both the network and endpoints. If prevention fails, you work to detect attacks in progress via security monitoring/alerting. And finally, if the compromise has already happened, you will perform some kind of incident response/management. All these use cases can benefit greatly by applying threat intelligence to their processes.

Use Case #1 : Preventative Controls

Ideally you will stop adversaries before they have a chance to compromise devices and steal data. So our first use case will focus on how threat intelligence can be used in preventative controls. By 'preventative' we mean any control that is in the flow and can prevent attacks. These include:

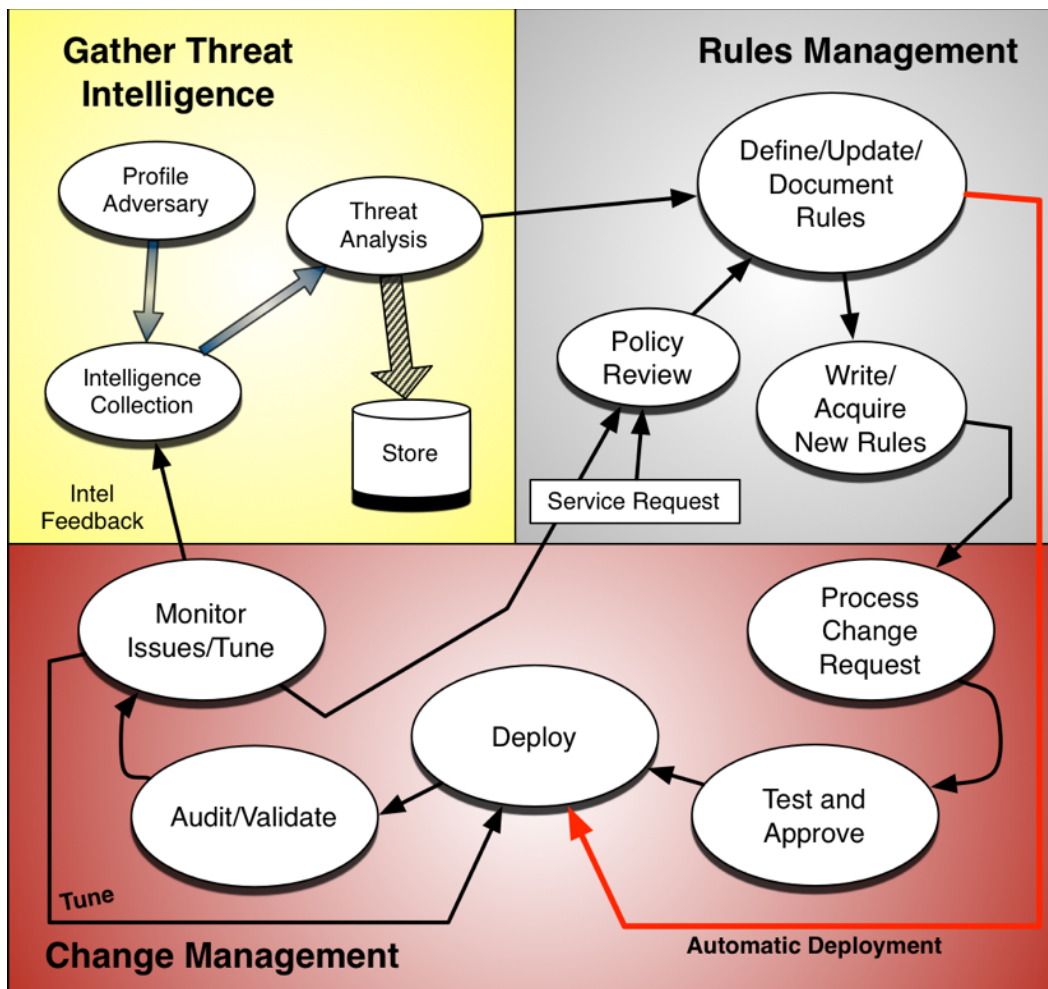
- **Network security devices:** This category encompasses firewalls (including next-generation models) and intrusion prevention systems. But you may also include devices such as web application firewalls which operate at different levels in the stack but are inline and can block attacks.
- **Content security devices/services:** Web and email filters can also function as preventative controls because they inspect traffic as it passes through, and can enforce policies to block attacks.
- **Endpoint security technologies:** Protecting an endpoint is a broad category, and can include traditional endpoint protection (anti-malware) and newfangled advanced endpoint protection technologies such as isolation and advanced heuristics. We examined the state of endpoint security in our [Advanced Endpoint Protection](#) paper, so check that out for detail on the technologies.

TI + Preventative Controls

The best way to understand how to apply TI is through a process map. So we dusted off our complicated Network Security Operations (NSO) process map from [NSO Quant](#), simplified a bit, and added threat intelligence. See the result below.

Rule Management

The process starts with managing the rules that underlie preventative controls. This includes attack signatures and the policies & rules controlling attack response. The process trigger is typically a service request (such as open this port for that customer), signature update, policy update, or threat intelligence alert (drop traffic from this set of botnet IPs). We will talk more about threat intelligence sources later.



1. **Policy Review:** Given the infinite variety of potential monitoring and blocking policies available through preventative controls, keeping rules current is critical. As you decide what policies to deploy, consider the severe performance hit (and false positive implications) of too many policies.
2. **Define/Update/Document Rules:** This next step involves defining the depth and breadth of security policies, including actions (block, alert, log, etc.) to take if an attack is detected — whether via rule violation, signature trigger, threat intelligence, or another method. Initial policy deployment should include a Q/A process to ensure no rules impair critical applications' ability to communicate, either internally or externally.
3. **Write/Acquire New Rules:** Locate the signature, acquire it, and validate the integrity of the signature file(s). These days most signatures are downloaded, so ensure the download completed successfully. Perform an initial evaluation of each signature to determine whether it applies within your organization, what type of attack it detects, and whether it is relevant in your environment. This initial prioritization phase determines the nature of each new/

updated signature, its relevance and general priority for your organization, and any possible workarounds.

Change Management

In this phase rule additions, changes, updates, and deletions are handled in the devices implementing the controls.

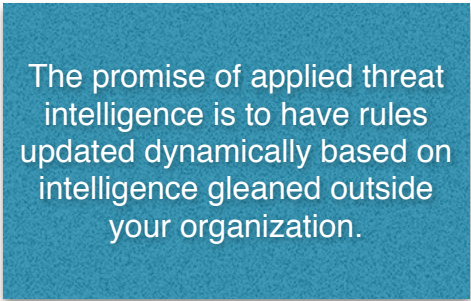
1. **Process Change Request:** Based on the trigger within the Rules Management process, a change to the preventative control(s) is requested. The change's priority is based on the nature of the rule update and risk of the relevant attack. Then build out a deployment schedule based on priority, scheduled maintenance windows, and other factors. This usually involves the participation of multiple stakeholders — ranging from application, network, and system owners, to business unit representatives if downtime or changes to application use models are anticipated.
2. **Test and Approve:** This step includes development of test criteria, performance of any required testing, analysis of results, and release approval of the signature/rule change once it meets your requirements. This is critical if you are looking to automate rules based on threat intelligence, as we will discuss later. Changes may be implemented in log-only mode to observe their impact before committing to blocking mode in production. This is critical for rules implemented as a result of threat intelligence because you will need to develop comfort with the accuracy and timeliness of your TI before you let it drive changes to your controls. With an understanding of the impact of the change(s), the request is either approved or denied.
3. **Deploy:** Prepare the target devices for deployment, deliver the change, and return them to normal operation. Verify that changes were properly deployed, including successful installation and operation. This might include use of vulnerability assessment tools or application test scripts to ensure there is no disruption to production systems.
4. **Audit/Validate:** Part of the full process of making the change is not only having the Operations team confirm it during the Deploy step, but also having another entity (internal or external, but not part of Ops) audit it to provide separation of duties. This includes validating the change to ensure policies were properly updated and matching it against a specific request. This closes the loop to ensure there is a documentation trail for every change. Depending on how automated you want this process to be, this step may not apply.
5. **Monitor Issues/Tune:** The final step of change management is a burn-in period when each rule change is scrutinized for unintended consequences such as unacceptable performance impact, false positives, security exposures, or undesirable application impact. For threat intelligence-based dynamic rules, false positives are the issue of most concern. The testing process in the Test and Approve step is intended to minimize these issues, but there are inevitable variances between test environments and production networks, so we recommend a probationary period for each new or updated rule... just in case.

Automatic Deployment

The promise of applied threat intelligence is to have rules updated dynamically based on intelligence gleaned outside your organization. It allows you to arbitrage time, providing opportunity to prepare for attacks that may eventually hit you.

We joke in conference talks about how security folks hate the idea of Skynet tweaking their defenses. There is still substantial resistance to updating firewall access control rules or IPS blocking actions without human intervention. But we expect this resistance to ebb as cloud computing continues to gain traction, especially in enterprise environments. *It is only possible to manage an environment at cloud speed and scale with automation.*

So automation is the reality in pretty much every virtualized environment, and it is making inroads in non-virtual security as well.



The promise of applied threat intelligence is to have rules updated dynamically based on intelligence gleaned outside your organization.

What can you do to get comfortable with automation? Automate things! No, we aren't being cheeky. You need to start simple — perhaps by implementing blocking rules based on very bad IP reputation scores. Or maybe taking an endpoint off the network if it's deemed to be acting like a bot. Monitor your environment closely to ensure minimal false positives. Tune your rules if necessary, and then move on to another use case.

Automated response makes sense in other situations — not only when acting on threat intelligence. In case of a data breach, lockdown, or zero-day attack (either imminent or in progress), you might want to implement (temporary) blocks or workarounds automatically based on predefined policies. If you detect a device or cloud instance acting strangely you could automatically move it to a quarantine network (or security zone) for investigation or just remove it from the network to contain the damage. This doesn't require human intervention, so long as you are comfortable with your trigger criteria.

Useful TI

Now let's consider collecting external data useful for preventing attacks. This includes the following types of threat intelligence:

- **File reputation:** File reputation can be thought of as a fancy name for traditional AV, but whatever you call it, malware proliferates via file transmission. Polymorphic malware makes signature matching much harder, but not impossible. The ability to block known-bad files close to the edge of the network is valuable — the closer the better.

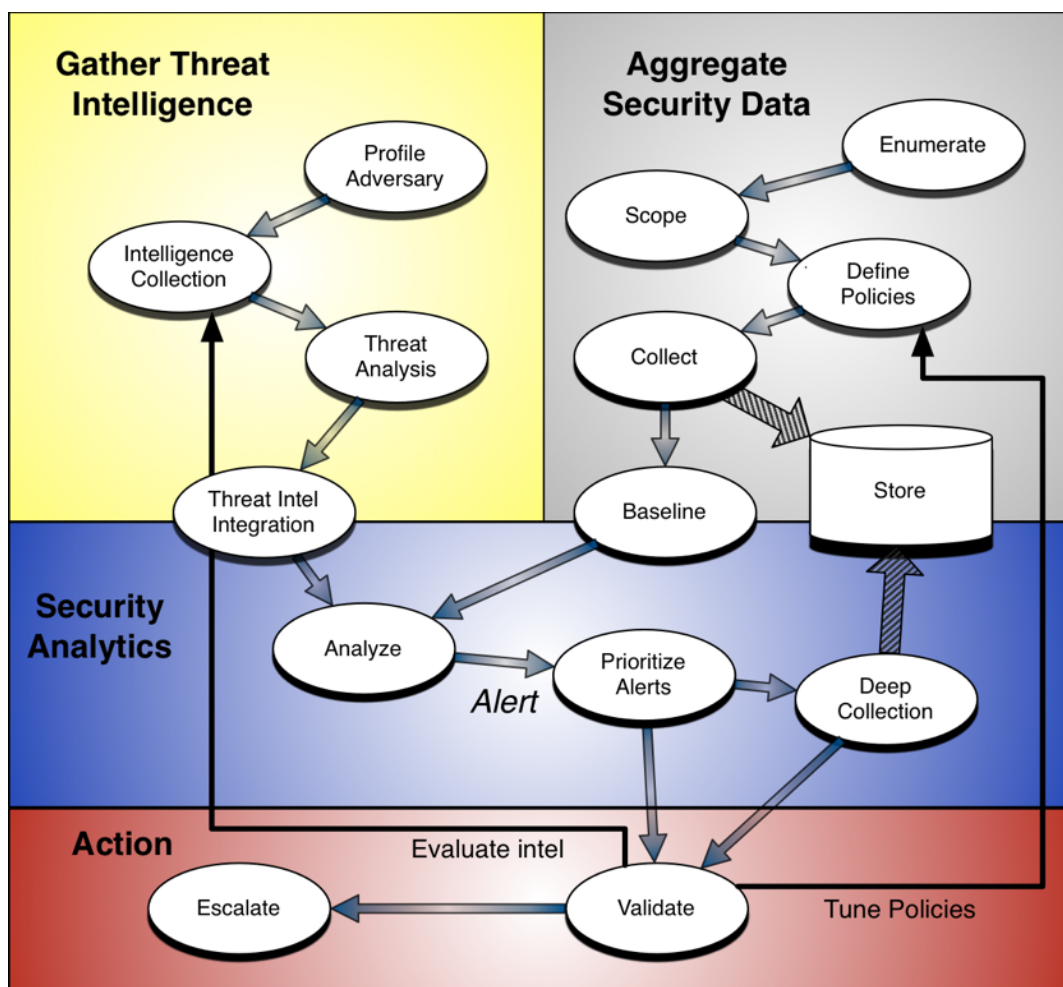
- **Adversary networks:** Some networks are just no good. They are run by hosting companies who provide safe haven for spammers, bot masters, and other online crime factions. There is no reason your networks should communicate with those networks. You can use a dynamic list of known bad and/or suspicious addresses to block ingress and egress traffic. As with any intelligence feed you should monitor effectiveness — bad networks fluctuate by the second, so keeping current is critical.
- **Malware indicators:** Malware analysis continues to mature rapidly, getting better and better at understanding exactly what malicious code does to devices, especially on endpoints. The shiny new term for an attack signature is Indicator of Compromise (IoC). But whatever you call it, an IoC is a handy machine-readable way to identify registry, configuration, or system file changes that indicate what malicious code does to devices. Looking for attacks in the detailed telemetry gathered from endpoints and networks enables you to prevent attacks on the way in, and benefit from others' misfortune.
- **Command and control patterns:** One specialized type of adversary network detection is intelligence on Command and Control (C&C) networks. These feeds track global C&C traffic to pinpoint malware originators, botnet controllers, and other IP addresses and sites to watch for as you monitor communications.
- **Phishing sites:** Current advanced attacks tend to start with a simple email. Given the ubiquity of email and the ease of adding links to messages, attackers typically find email the path of least resistance to a foothold in your environment. Isolating and analyzing phishing email can yield valuable information about attackers and their tactics, and give you something to block on your web filters and email security services.

To ensure you don't add unnecessary downtime or block critical traffic, iterate and tune your processes aggressively. Add new data sources and use cases, but not too fast. Make sure you don't automate a bad process, which would cause false positives and system downtime. Slow and steady wins this race.

Use Case #2: Security Monitoring

Threat intelligence can help detect attacks earlier by benefiting from the misfortune of others and looking for attack patterns which have been used against higher profile targets. Security monitoring is necessary because you simply cannot prevent everything. So you need to get better and faster at responding. The first step is improving detection, to shorten the window between compromise and discovery.

Before we jump into how let's see a security monitoring process with threat intelligence.



TI + SM

We will put the cart a bit before the horse: we will assume you already collect threat intelligence, as described earlier. But of course you cannot just wake up, pray a bit, and have compelling TI fall into your lap. You need to build a process and ecosystem to find, assess, and integrate TI, but we haven't described that process in any detail yet. We will defer that discussion a little, until you understand the context of the problem — then the techniques for systematically gathering TI will make more sense.

But of course you cannot just wake up, pray a bit, and have compelling TI fall into your lap. You need to build a process and ecosystem to find, assess, and integrate TI.

Aggregate Security Data

The steps involved in aggregating security data are fairly straightforward. You need to enumerate devices to monitor in your environment, scope out the kinds of data you will get from them, and define collection policies and correlation rules — all described in gory detail in [Network Security Operations Quant](#). Then you can move on to actively collect data and store it in a repository for flexible, fast, and efficient search and analysis.

Security Analytics

The security monitoring process now has two distinct categories of data to analyze, correlate, and alert on: external threat intelligence and internal security data.

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 data ingestion 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 may indicate compromise and warrant further investigation.
3. **Analyze security data:** The analysis process still involves normalizing, correlating, reducing, and tuning data and rules to generate useful and accurate alerts.
4. **Alert:** When a device shows one or more indicators of compromise, 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 investigate further, and in what order. Integrated threat intelligence can help by providing additional context, enabling responders to prioritize threats so analysts can investigate the highest risks 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. This can facilitate validation and identification of compromise, as well as forensic investigation if it comes to that.

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.

Action

Once you have an alert, and have gathered data about the device and attack, you need to determine whether it was actually compromised or the alert was a false positive. If a device has been compromised you may need to escalate — either to an operations team for remediation/clean-up or to an investigation team for more thorough 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.

Check out our [Leverage TI in Security Monitoring](#) paper for a much deeper discussion.

Useful TI

We are trying to detect attacks faster in this use case (rather than prevent or investigate them), so the most useful types of TI are strong indicators of problems. Let's review some threat intelligence sources, and how they fit this use case:

- **Compromised devices:** The most useful kind of TI is a service telling you there is a cesspool of malware on your network. This “smoking gun” can be identified by a number of different indicators, detailed below. But if you can get a product/service to identify those devices from analysis of TI data you can save yourself considerable effort analyzing and identifying suspicious devices.

Of course it may not be overly obvious which devices are compromised, so analyzing specific TI data types are helpful for detecting attacks:

- **File reputation:** Folks pooh-poo file reputation, but the fact is that a lot of malware still travels around through the tried and true avenue of file transmission. Polymorphic malware does make signature matching much harder, but not impossible, so tracking files can be helpful for detecting attacks and pinpointing the extent of an outbreak — as we will discuss in detail later in this paper.
- **Indicators of Compromise:** The shiny new term for an attack signature is Indicator of Compromise. But whatever you call it, an IoC is a handy machine-readable way to identify registry, configuration, and system file changes that indicate what malicious code does to devices. This kind of detailed telemetry from endpoints and networks enables you to detect attacks as they happen.

- **IP reputation:** At this point, given the popularity of spoofing addresses, we cannot recommend making a firm malware-vs.-clean judgement based solely on IP reputation, but if a device is communicating with known bad addresses and showing other indicators (which can be identified through the wonders of correlation — perhaps in a SIEM) you have strong evidence of compromise.
- **C&C patterns:** The last TI data source for this use case is a behavioral analog of IP reputation. You don't necessarily need to worry about where the device is communicating *to* — instead you can focus on *how* it is communicating. There are ways to analyze DNS traffic to find botnet controllers, and those findings can be confirmed by monitoring network traffic (either on the device or at the egress point).

Of course any security monitoring tool must understand how to parse these specific patterns and indicators, which is a good segue into integration with your enterprise monitoring tools/platform.

SIEM-tegration

The discussion so far all begs one question: How and where should you use TI data? We suggest you start with the existing tool, likely a SIEM of some sort (or its shiny cousin, a security monitoring/analytics platform). To understand the leverage points we need to revisit what SIEM does in the first place.

SIEM looks for patterns in security data via correlation and fancy math. But its Achilles' heel is that you need to know what to look for to generate an alert, and that involves building complex rules to look for specific attack scenarios and threat models. Many organizations spend considerable time and money figuring out these complex rules, and then even more tuning to produce some semblance of actionable alerts. And if you ask SOC (Security Operations Center) staff, they will be happy to explain, colorfully, how many of those alerts — even after tuning — turn out to be false positives.

Availability of the data types above changes what you should look for — at least initially. If you know there are a handful (or a couple handfuls) of attacks prevalent in the wild at any given moment, you can look for those. That is far more efficient and effective than scanning for every possible attack.

But that only works if you can keep your rules up to date, with fresh threat intelligence on the latest attacks. Unless you have some kind of *savant* who can parse a threat intelligence feed and build SIEM rules instantly, you will need to automate loading and rule building in the monitoring tool. We have already discussed the emergence of standards like STIX and TAXXI to facilitate the integration of threat feeds into your security monitoring platform. But even without standards, many TI providers built custom integrations to feed data into leading SIEM and security analytics products/services to extract value from their data.

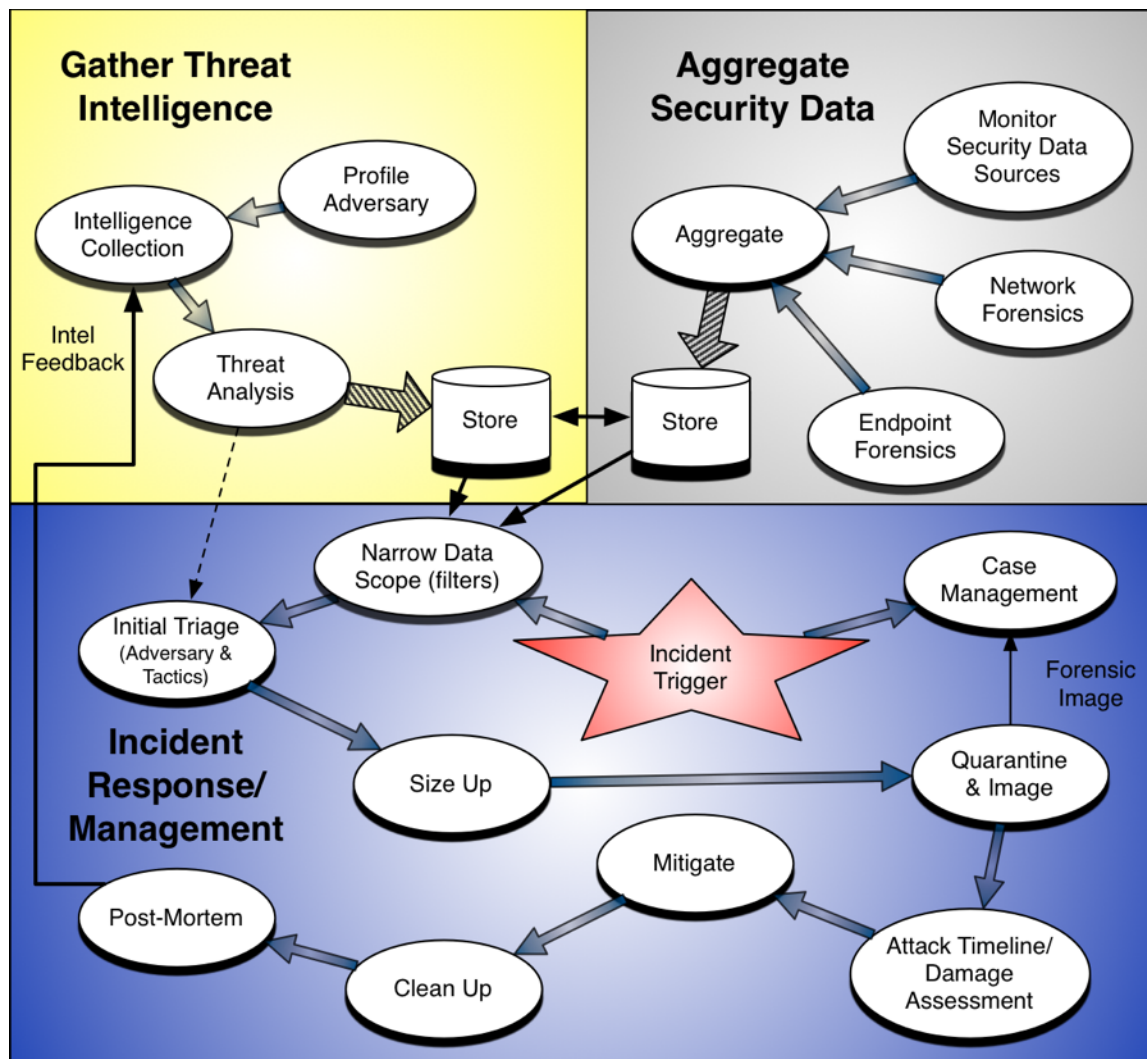
Actionable TI

Why is this approach better than just looking for patterns like privilege escalation or reconnaissance, as we learned in SIEM school? Because TI data represents attacks happening *right now* on other networks. Attacks you otherwise wouldn't see or know to look for until too late. In a security monitoring context leveraging TI enables you to focus your validation/triage efforts, shorten the window between compromise and detection, and ultimately make better use of scarce resources which need to be directed at the most important current risk. That is what we call *actionable intelligence*.

Why is this approach better than just looking for patterns like privilege escalation or reconnaissance, as we learned in SIEM school? Because TI data represents attacks happening right now on other networks.

Use Case #3: Incident Response/Management

Similar to the way threat intelligence helps with security monitoring, you can use TI to focus investigations on the devices most likely to be impacted, help identify adversaries, and lay out their tactics to streamline your response.



TI + IR/M

Let's start by revisiting the incident response and management process, and then list which types of TI data can be most useful and where.

You can get full descriptions of all the steps in our [Leveraging TI in Incident Response/Management](#) paper.

Trigger and Escalate

The incident management process starts with a trigger kicking off a response, and the basic information you need to figure out what's going on depends on what triggered the alert. You may get alerts from all over the place, including monitoring systems and the help desk. But not all alerts require a full incident response — much of what you deal with on a day-to-day basis is handled by existing security processes.

Where do you draw the line between a full response and a cursory look? That depends entirely on your organization. Regardless of the criteria you choose, all parties (including management, ops, security, etc.) must be clear on which situations require a full investigation and which do not before you can decide whether to pull the trigger. Once you escalate an appropriate resource is assigned and triage begins.

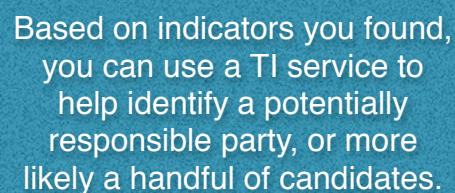
Triage

Before you do anything you need to define accountabilities within the team. That means specifying the incident handler and lining up resources based on the expertise needed. Perhaps you need some Windows gurus to isolate a specific vulnerability in XP. Or a Linux jockey to understand how system configurations were changed. Every response varies a bit, and you want to make sure you have the right team in place.

As you narrow down the scope of data needing analysis, you might filter on the segments attacked or logs of the application in question. You might collect forensics from all endpoints at a certain office if you believe the incident was contained. Data reduction is necessary to keep the data to investigate manageable.

Analyze

You might have an initial idea of who is attacking you, how they are doing it, and their mission based on the alert that triggered the response, but now you need to prove your hypothesis. This is where threat intelligence plays a huge role in accelerating response. Based on indicators you found, you can use a TI service to help identify a potentially responsible party, or more likely a handful of candidates. You probably don't need legal attribution, but this information can help you understand an attacker and their tactics.



Based on indicators you found, you can use a TI service to help identify a potentially responsible party, or more likely a handful of candidates.

Then you need to size up and scope out the damage. The goal here is to take the initial information provided and supplement it quickly to determine the extent and scope of the incident. To determine scope dig into collected data to establish the systems, networks, and data involved. Don't worry about pinpointing every affected device at this point — your goal is to size the incident and generate ideas for how best to mitigate it. Finally, based on your initial assessment, use your predefined criteria to decide whether a formal investigation is in order. If yes, start thinking about chain of custody and some kind of case management system to track the evidence.

Quarantine and Image

Once you have a handle (however tenuous) on the situation, you need to figure out how to contain the damage. This usually involves taking a device offline and starting the investigation. You could move it onto a separate network without access to anything real, or disconnect it from the network altogether. You could turn the device off. Regardless of what you decide, do not act rashly — you need to make sure things do not get worse, and avoid destroying or contaminating evidence. Many malware kits (and attackers) wipe devices which are powered down or disconnected from the network, so be careful.

Next you take a forensic image of the affected devices. You need to make sure your responders understand how the law works in case of prosecution, especially what provides a basis for reasonable doubt in court.

Investigate

All this work is a precursor to the full investigation, when you dig deep into the attack to understand exactly what happened. We like timelines to structure your investigation because they help illustrate what happened and when. Start with the initial attack vector and follow the adversary as they systematically moved to achieve their mission. To ensure a complete cleanup the investigation must pinpoint exactly which devices were affected and review either very detailed metadata derived from network traffic, or the actual exfiltrated data via full packet capture from perimeter networks.

Investigation is more art than science, and you can never actually know everything, so focus on what you *do* know. At some point a device was compromised or a system/network was taken down by an attack. You might have had data exfiltrated. Systematically fill in gaps to understand what the attacker did and how. Focus on completeness of the investigation — a missed compromised device is sure to mean reinfection somewhere down the line. Then perform a damage assessment to determine, as closely as possible, what was lost.

Mitigation

There are many ways to ensure an attack doesn't happen again. Temporary measures include shutting down access to certain devices via specific protocols, and locking down traffic in and out of critical servers. Or possibly blocking outbound communication to certain regions based on adversary intelligence. Also consider more 'permanent' mitigations, such as a service or product to block denial of service attacks.

Once you have a list of mitigation activities, marshal operational resources to work through them. We favor remediating affected devices in one fell swoop (a “big bang”), rather than incremental cleaning/reimaging. We have found it more effective to eradicate the adversary from your environment as quickly as possible, because a slow cleanup provides opportunity for them to dig deeper.

The mitigation is complete once you have halted the damage and regained the ability to continue operations. Your environment may not be pretty when you finish mitigation, with a bunch of temporary workarounds to protect information and make sure devices are no longer affected.

Clean up

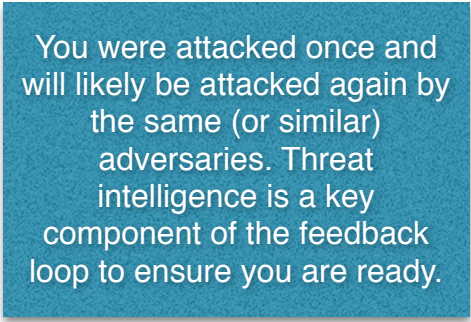
Now take a step back and clean up any disruptions to normal business operations, making sure you are confident that particular attack will never work again. Incident managers focus on completing the investigation and cleaning out temporary controls, while Operations handles updating software and restoring normal function. This could include updating patches on all systems, checking for and cleaning up malware, restoring systems from backups and bringing them back up to date, etc.

Postmortem

Your last step is to analyze the response process itself. What can you identify as opportunities for improvement? Should you change the team or your response technology (tools)? Don't make the same mistakes again, and be honest with yourself about what needs to improve.

You cannot completely prevent attacks, so you need to optimize your response process to detect and manage problems as quickly and efficiently as possible — which brings us full circle back to threat intelligence. You also need to learn about your adversary during this process.

You were attacked once and will likely be attacked again by the same (or similar) adversaries. Threat intelligence is a key component of the feedback loop to ensure you are ready.



You were attacked once and will likely be attacked again by the same (or similar) adversaries. Threat intelligence is a key component of the feedback loop to ensure you are ready.

Useful TI

Now let's delve into collecting external data that will be useful to streamline investigation. This involves gathering threat intelligence, including the following types:

- **Compromised devices:** The most actionable intelligence you can get is a clear indication of compromised devices. This provides an excellent place to begin investigation and manage your response. There are many ways you might conclude a device is compromised. The first is clear indicators of command and control traffic coming from the device, such as DNS requests whose frequency and content indicate a Domain Generating Algorithm (DGA) to locate botnet controllers. Monitoring network traffic from the device can also catch files or other sensitive data being transmitted, indicating exfiltration or a remote access trojan.

- **Malware indicators:** You can build a lab and perform both static and dynamic analysis of malware samples to identify specific indications of how malware compromises devices. This is a major commitment (detailed in Malware Analysis Quant). Thorough and useful analysis requires significant investment, resources, and expertise. The good news is that numerous commercial services now offer indicators in formats you can use to easily search through collected security data.
- **Adversary networks:** IP reputation data can help you determine the extent of compromise, especially if it is broken up into groups of adversaries. If during your initial investigation you find malware typically associated with Adversary A, you can look for traffic going to networks associated with that adversary. Effective and efficient response requires focus, and knowing which devices may have been compromised in a single attack helps isolate and dig deeper into that attack.

Given the scarcity of security team resources, many organizations select a commercial provider to develop and provide this external threat intelligence, or leverage threat intel data provided as part of a product or service. Stand-alone threat intelligence is typically packaged as a feed for direct integration into incident response/monitoring platforms. Wrapping it all together produces the process map above. This map encompasses profiling the adversary, collecting intelligence, analyzing threats, and then integrating threat intelligence into incident response.

Action Requires Automation

The key to making this entire process run is automation. We talk about automation a lot these days, with good reason. Technology infrastructure changes too quickly to keep up with much of anything manually, especially in the heat of an investigation. You need to pull threat intelligence in a machine-readable format, and pump it into an analysis platform, without human intervention.

Technology infrastructure changes too quickly to keep up with much of anything manually, especially in the heat of an investigation. You need to pull threat intelligence in a machine-readable format, and pump it into an analysis platform, without human intervention.

Building a TI Program

The last piece of this puzzle is building a repeatable process to collect, aggregate, and analyze threat intelligence. This should include a number of different information sources, as well as various internal and external data analyses to provide context and clarify what the intelligence means to you.

As with pretty much everything in security, handling TI is not “set and forget”. You need to build repeatable process to select data providers and continually reassess the value of those investments. You will need to focus on integration; as we described, data isn’t helpful if you cannot use it in practice. And your degree of comfort in automating processes based on threat intelligence will impact day-to-day operational responsibilities.

As with pretty much everything in security, handling TI is not “set and forget”. You need to build repeatable process to select data providers and continually reassess the value of those investments.

First you need to decide where the threat intelligence function will fit organizationally. Larger organizations tend to formalize an intelligence group, while smaller entities need to add intelligence gathering and analysis to the task lists of existing staff. Out of all the things that could land on a security professional, an intelligence research responsibility isn’t bad. It provides exposure to cutting-edge attacks and makes a difference in your defenses, so that’s how you should sell it to overworked staffers who don’t want yet another thing on their to-do lists.

But every long journey begins with the first step, so let’s turn to collecting intelligence.

Gather Intelligence

Early in the intelligence gathering process you focused efforts with analysis of your adversaries. Who they are, what they are most likely to try, and what kinds of tactics they use to achieve their missions — you need to tackle all the questions above. With those answers you can focus on intelligence sources that address your probable adversaries. Then identify the kinds of data you need, based on the use case you are addressing. Depending on which use cases you are trying to address, you will know whether to focus on malware indicators, compromised devices, IP reputation, command and control indicators, or something else.

Then start shopping. Some folks love to shop, others not so much. But it’s a necessary evil; fortunately recent growth in the threat intelligence market provides plenty of options. Let’s break down a few categories of intelligence providers, along with the particular value of each:

- **Commercial:** These providers employ research teams to conduct proprietary research, and tend to achieve high visibility by merchandising findings with fancy exploit names and logos, spy thriller stories of how adversary groups compromise organizations and steal data, and shiny maps of global attacks. They tend to offer particular strength against specific adversary classes. Look for solid references from your industry peers.
- **OSINT:** Open Source Intelligence (OSINT) providers are commercial entities that specialize in mining the huge numbers of information security sources available on the Internet. Their approach is all about categorization and leverage because plenty of information is available free. These folks know where to find it and how to categorize it. They normalize data and provide it through a feed or portal to make it useful for your organization. As with commercial sources, the question is how valuable any particular source is to you. You already have too much data — you need providers who can help you wade through it.
- **ISAC:** There are many Information Sharing and Analysis Centers (ISAC), mostly for specific industries, to communicate current attacks and threat data among peers. As with OSINT, quality can be an issue, but this data tends to be industry specific so its relevance is pretty well assured. Participating in an ISAC obligates you to contribute data back to the collective, which we think is awesome. The system works much better when organizations both contribute and consume intelligence, but we understand the cultural constraints. So you will need to make sure senior management is okay with it before committing to an ISAC.

Another aspect of choosing intelligence providers is figuring out whether you are looking for generic or company-specific information. OSINT providers are more generic, while commercial offerings may go deeper. Additionally, various ‘Cadillac’ offerings dedicate analysts to your organization — proactively searching grey markets, carder forums, botnets, and other places for intelligence relevant to you.

Managing Overlap

With disparate data sources it is a challenge to ensure you don’t waste time on multiple instances of the same alert. One key to detecting overlap is understanding 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 help pick the best fit for your requirements.

To choose services you need to compare comprehensiveness, timeliness, and accuracy. Sign up for trials of a number of services and monitor their feeds for a week or so.

To choose services you need to compare comprehensiveness, timeliness, and accuracy. Sign up for trials of 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? How easy will it be to integrate their data into your environment and your use cases?

Don't fall for marketing hyperbole about proprietary algorithms, big data analysis, staff linguists penetrating hacker dens, or other stories straight out of a spy novel. It all comes down to data and how useful it is to your security program. Buyer beware, and make sure you put each intelligence provider through its paces before you commit.

Our last point 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 might not be around in 3-4 years. Once you identify core intelligence feeds longer deals can be cut, but we recommend against doing so before your TI process matures and your intelligence vendor establishes a track record addressing your needs.

To Platform or Not to Platform

Now that you have chosen intelligence feeds, what will you do with the data? Do you need a stand-alone platform to aggregate it all? Will you need to stand up yet another system in your environment, or can you leverage something in the cloud? Will you actually use your intelligence providers' shiny portals? Or do you expect alerts to show up in the existing monitoring platforms, or be sent via email or SMS?

There are many questions to answer as part of operationalizing TI process. First you need to figure out whether existing technology can fit the bill. Existing security providers (specifically SIEM and network security vendors) now offer threat intelligence 'supermarkets' which enable you to easily buy and integrate data into monitoring and control environments. Even if your platform vendors don't offer a way to easily buy TI, many support standards such as STIX and TAXII to facilitate integration.

If you have a dedicated team to evaluate and leverage TI, have multiple monitoring and/or enforcement points, or want more flexibility in how broadly you use TI, you should probably consider a separate intelligence platform or 'clearinghouse' to manage TI feeds.

We are focusing on Applied Threat Intelligence, so your decision hinges on how you will use it. If you have a dedicated team to evaluate and leverage TI, have multiple monitoring and/or enforcement points, or want more flexibility in how broadly you use TI, you should probably consider a separate intelligence platform or 'clearinghouse' to manage TI feeds.

Selecting the Platform

When you evaluate stand-alone threat intelligence platforms, there are a few key selection criteria.

1. **Open:** The TI platform's task is to aggregate information so it must be easy to get information into it. Intelligence feeds are typically just data (often XML), and increasingly distributed in industry-standard formats such as STIX, which make integration relatively straightforward. But make sure any platform you select will support the data feeds you

need. Make sure you can use the data that's important to you, and not be restricted by your platform.

2. **Scalable:** You will use a lot of data in your threat intelligence process, so scalability is essential. But computational scalability is likely more important than storage scalability — you will be intensively searching and mining aggregated data so you need robust indexing. Unfortunately scalability is hard to test in a lab, so ensure your proof of concept testbed is a close match to your production environment, and that you can extrapolate how the platform will scale in your production environment.
3. **Search:** Threat intelligence, like the rest of security, doesn't lend itself to absolute answers. So make TI the start of your process of figuring out what happened in your environment, and leverage the data for your key use cases, as we described earlier. One clear requirement for all use cases is search. So make sure your platform makes it easy to search all your TI data sources.
4. **Urgency scoring:** Applied Threat Intelligence is all about betting on which attackers, attacks, and assets are most important to worry about, so you will find considerable value in a flexible scoring mechanism. Scoring factors should include assets, intelligence sources, and attacks, so you can calculate a useful urgency score. It might be as simple as red/yellow/green, depending on the sophistication of your security program.

Determining Relevance in the Heat of Battle

How can you actually use the threat intelligence you painstakingly collected and aggregated?

Relevance to your organization depends on the specifics of the vulnerability and whether it can be used against you. Focus on real potential exploits — a vulnerability which does not exist in your environment is not your concern. For example you probably don't need to worry about financial malware if you don't hold or have access to credit card data. That doesn't mean you shouldn't pay any attention to these attacks — many exploits leverage a variety of interesting tactics, which might become part of a relevant attack in the future. Relevance encompasses two aspects:

1. **Attack surface:** Are you vulnerable to this specific attack vector? Weaponized Windows 2000 exploits aren't relevant if you don't have any Windows 2000 systems. Once you have patched all instances of a specific vulnerability on your devices, you get a respite from worrying about the exploit. Your asset base and internally collected vulnerability information provide this essential context.
2. **Intelligence reliability:** You need to continually reevaluate each threat intelligence feed to confirm its usefulness. A feed which triggers many false positives is less relevant. On the other hand, if a feed usually nails a certain type of attack, you should take those warnings particularly seriously. Note that attack surface may not be restricted to your own assets and environment. Service providers, business partners, and even customers represent indirect risks — if one of them is compromised, an attacker might have a direct path to you.

Constantly Evaluating Intelligence

How can you determine the reliability of a TI source? Threat data ages very quickly and TI sources such as IP reputation can change hourly. Any system you use to aggregate threat intelligence should be able to report on the number of alerts generated from each TI source, without hurting your brain building reports. These reports show value from your TI investment — it is a quick win if you can show how TI identified an attack earlier than you would have detected it otherwise. Additionally, if you use multiple TI vendors, these reports enable you to compare them based on actual results.

Marketing Success Internally

Over time, as with any security discipline, you will refine your verification/validation/investigation process. Focus on what worked and didn't and tune your process accordingly. It can be bumpy when you start really using TI sources — typically you start by receiving a large number of alerts, and following them down a bunch of dead ends. It might remind you, not so fondly, of the SIEM tuning process. But security is widely regarded as overhead, so you need a Quick Win with any new security technology.

TI will find stuff you didn't know about and help you get ahead of attacks you haven't seen yet. But that success story won't tell itself, so when the process succeeds — likely early on — you will need to publicize it early and often. A good place to start is discovery of an attack in progress. You can show how you successfully detected and remediated the attack using threat intelligence. This shows that you *will* be compromised (which must be constantly reinforced to senior management), so success is a matter of containing damage and preventing data loss. The value of TI in this context is in shortening the window between exploit and detection.

TI will find stuff you didn't know about and help you get ahead of attacks you haven't seen yet. But that success story won't tell itself, so when the process succeeds — likely early on — you will need to publicize it early and often.

You can also explain how threat intelligence helped you evolve security tactics based on what is 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 used by that particular adversary is often a decoy to obscure exfiltration, you have sufficient 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 using threat intelligence to refine efforts and prioritize activities.

Summary

As we have discussed through this paper, threat intelligence can even up the battle between attackers and defenders, to a degree. But to accomplish this you must be able to gather relevant TI and leverage it in key threat management processes. This starts with using TI in preventative controls at the front end of your security process to disrupt attacks, either within the network or on endpoint devices. In the event your preventative controls fail to prevent an attack, the sooner and faster you can respond, the more likely you can prevent a catastrophic breach. This involves supplementing your security monitoring processes with external data.

Finally, should an adversary successfully compromise devices, and possibly even exfiltrate data, you will race against time to identify the root cause of the attack and ensure the attack is contained via an effective and efficient incident response process.

None of this is *easy*, but nothing worthwhile in security ever is. It is about continuously improving your processes to favorably impact your security posture. Applying threat intelligence within these processes can provide leverage to make your controls and people work better.

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

About the Analyst

Mike Rothman, Analyst and President

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

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

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

About Securosis

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

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

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors.

Additionally, Securosis partners with security testing labs to provide unique product evaluations that combine in-depth technical analysis with high-level product, architecture, and market analysis. For more information about Securosis, visit our website: <<http://securosis.com/>>.