How to Use Threat Intelligence for Security Monitoring and Incident Response

Published: 24 February 2020 **ID:** G00463498

Analyst(s): Michael Clark, Augusto Barros

Threat Intelligence is becoming a ubiquitous capability in many security tools. It is a key aspect of security architecture that helps security and risk management technical professionals detect, triage and investigate threats. This research provides guidance on how to use TI capabilities.

Key Findings

- Threat intelligence (TI) improves an organization's detection and response capability by increasing alert quality, reducing investigation time, and adding coverage for the latest attacks and adversaries.
- Modern security tools can ingest and leverage threat intelligence. However, they often don't include guidance on the best way to utilize it.
- Using threat intelligence improperly will result in more noise and false positives. Proper upfront planning for TI usage is critical.

Recommendations

As a security and risk management technical professional focused on security operations, you should:

- Collect TI requirements based on the threats faced and technology use cases. Tactical use cases deliver TI to your security controls, while strategic use cases leverage TI to educate and inform stakeholders.
- Curate threat intelligence before delivering it to stakeholders and security controls, by applying scores, expirations and enrichments.
- Deliver tactical threat intelligence to your existing security controls by using API- or TAXII-based integrations. Deliver strategic threat intelligence to stakeholders within your organization by creating regular reporting.

- Assess the effectiveness of the threat intelligence by tracking metrics and describing the impact of TI. Use information about a threat in combination with observables attributed to that threat to demonstrate losses prevented.

Table of Contents

List of Tables

List of Figures

Problem Statement

Threat intelligence is often included as a standard feature of security products. In this case, "standard feature" may mean any one of the following:

- The product leverages threat intelligence provided by the vendor automatically.
- The product can import threat intelligence from third-party sources.
- The product supports some combination of both capabilities.

Unfortunately, such security products usually provide very little information given about the TI and how to use it effectively. In the worst-case scenario, this lack of guidance leads to more alert noise and false positives.

Like with other security controls, TI needs to be configured properly, and the results need to be understood and reacted to appropriately. This research provides guidance on properly using threat intelligence, including:

- How to determine what type of the threat intelligence the organization needs
- How to manage that threat intelligence
- How to gauge the performance of the threat intelligence

The Gartner Approach

Getting started with threat intelligence can be a daunting task. It involves understanding what your organization needs to know, what your security controls can handle, and how to leverage the intelligence in a timely and effective matter. If threat intelligence is not handled correctly, it can exacerbate alert fatigue and waste time.

This guidance framework leads organizations through all the steps required to effectively use threat intelligence in security monitoring and incident response.

The Guidance Framework

Using threat intelligence in security monitoring is a five-step process: collect, curate, deliver, improve and assess (see Figure 1). The Prework section lays the base for incorporating threat intelligence into your security monitoring and incident response programs. Then, the next two sections map the five-step process to the two main types of threat intelligence: strategic and tactical.

Figure 1. Guidance Framework for Using Threat Intelligence

Prework

If you examine most security products, you will likely see threat intelligence listed somewhere in their feature set. Its inclusion has almost become standard, as TI is a very useful tool in the defender's toolbox. However, threat intelligence should not be treated as a checkbox or some background process that will just always work. If used improperly, threat intelligence can cause more harm, such as increased false positives, than good to your security program. To best leverage threat intelligence to accomplish your goals, you first need to understand what threat intelligence is.

Definition of Threat Intelligence

Gartner defines threat intelligence as "evidence-based knowledge (e.g., context, mechanisms, indicators, implications and action-oriented advice) about existing or emerging menaces or hazards to assets." (See "How Gartner Defines Threat Intelligence.") TI can be very specific, such as indicators of compromise (IOCs). It can also be more high-level, such as a report on an adversary.

Examples of threat intelligence include:

- IP address of a command and control (C&C) server
- MD5 hash of a malicious executable with context
- Report on a threat actor who is known to target the financial sector
- Network intrusion prevention system (NIPS) rule that will detect malware communications delivered as part of a report

This broad definition encompasses several types of threat intelligence with distinct uses and value for an organization. Threat intelligence data is not equivalent to NIPS signatures, antivirus (AV) updates, YARA signatures or other ready-made security updates. These do describe patterns of malicious activities, such as exploits or malware, but they do not include the additional context, which can range from malicious infrastructure locations to targeting priorities of foreign intelligence agencies.

> Threat intelligence is knowledge about who or what is on the other side, as well as how they operate. Trying to defend against every possible threat with limited resources is a losing proposition. Use TI to protect your organization from all the relevant adversaries. Further, use TI to decide which adversaries are in fact relevant.

Threat intelligence can be organized into two broad classes:

- **Strategic TI** includes reports and other human-readable products on threat actors and their intentions, affiliations, interests, goals, capabilities, plans and campaigns. Strategic TI is typically produced by human analysts and is likely consumed by humans as well. Strategic TI is often associated with decisions that are longer-term in nature, such as what new programs to implement, what processes to change or what new infrastructures to invest in.
- **Tactical TI** often consists of IOCs, such as IP addresses, domains, URL or hash lists, and other system-level or network-level artifacts. These artifacts can be matched to what is observed on information systems. Tactical TI is most often consumed by security controls, but it is also manually looked at or used during investigations and incident response. Tactics, techniques and procedures (TTPs) are another example of tactical TI. Tactical TI is typically associated with decisions that are shorter-term in nature, such as urgent alerts to personnel, invocation of an existing escalation process or configuration changes on existing infrastructure.

Indicators of compromise are becoming less reliable as time goes on due to their static nature. Attackers can easily change IOCs such as hashes, IP addresses and domain names. IOCs are a view of a past situation. They are also a view of specific incidents. Sophisticated attackers will often use different infrastructure for different targets. For example, a command-and-control server may be used only for a single target. The popular Pyramid of Pain diagram illustrates how easy or difficult it is to change an IOC, depending on its type. At the base of the pyramid are file hashes, which are trivial for an adversary to alter. At the top are TTPs, which require much more effort to alter.

This is not to say that IOCs don't have value. Many attacks still use static IOCs or reuse infrastructure. IOCs are also very useful for identifying signs of compromise that may have occurred in the past. Much like AV signatures, IOCs can be unreliable if they are the only detection method being relied upon.

Difference in TI Usage by Type

Table 1 clarifies the differences between these broad types of threat intelligence.

Table 1. Difference Between TI Types

	Strategic	Tactical
Created By	Humans using technical and nontechnical sources	Machines or humans
Consumed By	Humans	Machines and humans
Delivery Time Frame	Days to years	Seconds to hours
Useful Life Span	Long	Usually short
Resistance to Change	Durable	Fragile
Focus	Planning and high-level decisions	Detection, triage and response
Examples	Information targeted, organization affiliation of the threat actor, intentions, preferred tools and threat actor profiles	IP, domain, URL, MD5, hostname and filename

Source: Gartner (February 2020)

Identify the primary assets of the organization that should be protected early in the requirements-gathering process. Going into this initiative with a strategy of protecting anything and everything will lead to poor results. Resources and, in most cases, budgets are limited. It is important to focus on the critical types of assets for the business, such as intellectual property or financial information that lives in a database.

In addition to critical asset types, other requirements could be based on geographical location or industry vertical. An adversary may use either of those criteria to target an attack, rather than target the attack against your organization specifically.

Many organizations face the challenge of figuring out how to tell good threat intelligence products from bad ones. After all, many vendors that offer similarly sounding data feeds of TI make the same claims, but feature wildly different prices — from free to hundreds of thousands of dollars per year.

When discussing TI, you need to keep some general requirements in mind:

- **Correctness:** Although this requirement is sometimes reduced to “low false-positive ratio,” there is more to correctness than this. Correctness means that data is validated, cross-checked, refined and removed when it is no longer appropriate.
- **Breadth:** This requirement is critical for organizations operating in many regions and industries; their TI should cover the relevant locations, industries, IT asset types under risk and threat actor types. Internet threats are global, and TI should be as well.

- **Timeliness:** Given the long delays in identifying incidents, it is useful to receive an indicator today that helps you discover a compromise that happened last month. However, it is much more useful to discover compromises as, or immediately after, they happen.
- **Context:** Although a list of known malicious entities has value, clarifying the type, history, nature and intention of the threat actors behind those entities enhances the value of TI.
- **Relevance:** A long list of threat indicators with all the right context data — coming from a trusted source and on a timely basis — is still of limited use if it is not connected to the threats operating in your environment.
- **Durability:** This parameter of intelligence is more difficult to understand because it involves how long the TI remains valid and how difficult it is for the attacker to change away from that TI. An IP address for an exfiltration site can be changed easily, but preference for a tool tends to be stickier.
- **Structured and linked format:** Receiving an email that warns you about an imminent attack is helpful, but only structured data formats can make threat intelligence scalable enough to realize benefits, relate to other intelligence and ultimately create a better defense.
- **Cost:** The price of the intelligence may not be a hard requirement. Some organizations report that advance attack warnings and IOCs that enable them to unearth advanced threats in their environments are always worth the cost.

Threat Intelligence Platforms

As an organization's use of tactical threat intelligence expands and matures, additional tools may be required to perform the work in an efficient way. Threat intelligence platforms (TIPs) have emerged as repositories for both strategic and tactical TI. They provide capabilities to make searching for, relating, creating and improving TI easier tasks.

As shown in Figure 2, a TIP can act as a central store for threat intelligence. Once an organization starts to consume multiple sources of threat intelligence with multiple security controls, management can become complicated. A TIP can aggregate all that information and provide a single place for security controls or stakeholders to obtain the information. TIP vendors such as Anomali, ThreatConnect and ThreatQuotient offer different components and capabilities. Without a TIP, each control will connect out to the TI sources on its own, and there will be limited control over the information.

Figure 2. Threat Intelligence Architectures

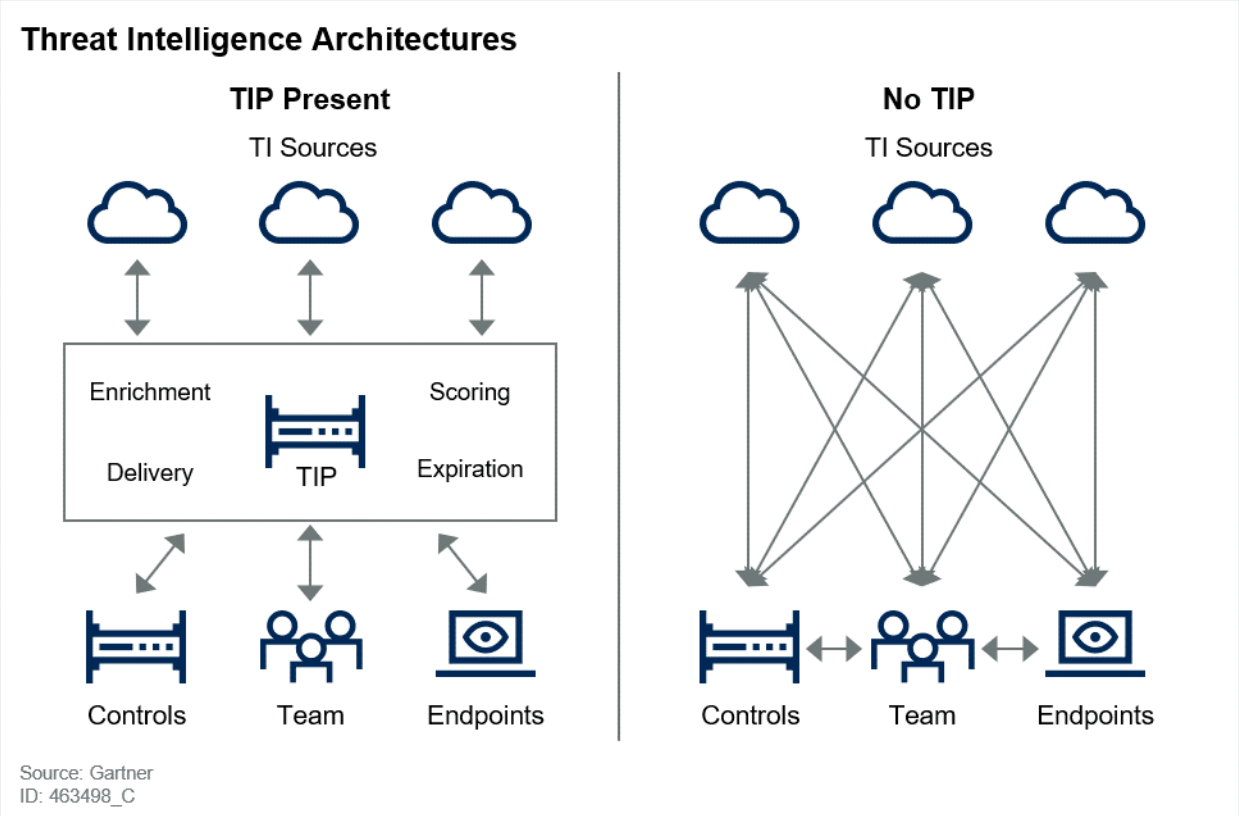

Table 2 lists essential TIP capabilities. Given that relatively few organizations have deployed such platforms, the capability matrix is expected to significantly evolve in the future.

Table 2. TIP Capabilities

Platform Capability	Details

Collect TI Data	■ Collect:
	  ■ Free-form data (e.g., emails and PDFs) — ideally with entity recognition (such as IP address and email address), tactical intelligence, and strategic intelligence
	  ■ Standard data and transport formats (e.g., Trusted Automated eXchange of Indicator Information [TAXII]/STIX and OpenIOC)
	  ■ Other structured data (e.g., XML and CSV)
	■ Collect from browser via copy-paste-parse
	■ Manually add from other tools (export/import)
	■ Collect incident metadata or full incident files, local reversing results, and local TI data
	■ Retain data for some time based on policy

Refine TI Data (TI Fusion)	■ Fuse structured and unstructured information from various sources ■ Relate tactical intelligence to strategic intelligence ■ Analyze, enrich and associate threat indicators over time ■ Correlate across data feeds, data types and sources ■ Support workflows, and automate enrichment processes where possible ■ Index for faster search ■ Score and prioritize IOCs
Enable Analysts to Search, Report on and Analyze Data	■ Search keywords and parameters ■ Review reports and trends, identify changing patterns, and so on
Enable Analyst Collaboration and Group Workflows	■ Support shared comments and workflows ■ Support manual ranking of data sources/feeds by reliability, detection quality and so on ■ Support manual data enrichment with local details, investigation results and other sources
Disseminate TI to Various People Inside the Organization	■ Support portal access for reports and search, as well as data manipulation ■ Email reports based on conditions ■ Support API access for internal tools ■ Export data
Distribute TI to Tools	■ Integrate actionable intelligence into existing network defenses and sensors

Platform Capability	Details
	■ Convert and export data into signatures, checks and so on ■ Feed data into security information and event management (SIEM) for context and matching ■ Match in real time (elsewhere) ■ Match historically (elsewhere)
Enable Measurement of the Organization's TI Efforts	■ Collect metrics on the usage of TI data to rank feeds and sources ■ Report high-level metrics that are useful for overall risk management
Enable Cross-Organization Intelligence Sharing	■ Tag data for sharing with particular trust circles/communities ■ Export/make available select data (filtered, aggregated and sanitized) to other organizations, industry groups and so on ■ Transfer data over supported protocols
Investigate	■ Provide an interface for using TI as a basis for investigation
Classify	■ Support data classification using the Traffic Light Protocol (TLP) or custom schemes

Source: Gartner (February 2020)

TIPs often contain automation and orchestration capabilities, with some of them evolving into full security orchestration, automation and response (SOAR) tools. They can be very valuable for informing decision-making steps in an automation playbook. Vendors, such as ThreatConnect, incorporate robust automated playbook capabilities.

In addition to automation platforms, TIPs can be used as investigation platforms. Some vendors include a graph-based platform that allows analysts to visually see the relationships between TI and data from security control integrations.

Type 1: Strategic Threat Intelligence

Strategic threat intelligence comes in the form of a human-readable report or brief. It often contains detailed information on adversaries, malware, vulnerabilities and campaigns. Besides general situational awareness, strategic TI is useful for understanding whether the organization's security posture is aligned to how threats operate.

Step 1: Collect

Gartner clients report using the dimensions listed in Table 3 to compare TI providers, but most large clients end up engaging more than one provider, due to differences in regional and actor coverage. For example, one organization may choose to procure subscriptions from multiple strategic intelligence providers to achieve better coverage of cybercrime actors. By contrast, another may require more human-sourced intelligence (HUMINT) in its country of origin.

Table 3. Strategic TI Provider Comparison

Dimension	Details
Content Coverage	■ Regions and countries ■ Types of actors (crime, espionage/state, hacktivism, etc.)
Resources	■ Analysts in countries ■ Analyst language coverage ■ Analysts available for custom projects ■ Types of custom intelligence projects delivered
Depth	■ Relation to TTPs, campaigns, tactical intelligence feeds, etc. ■ Availability of detection-ready content
Capabilities	■ APIs, portals, integrations with TIPs and document-sharing platforms ■ Actionability of reports and other content (how often reports lead to action)
Timeliness	■ Frequency of content base updates
Customization	■ Made-to-order intelligence products (such as reports) ■ Access to analysts
Formats	■ Available data formats (e.g., PDF, XML and HTML) and other TI presentation and delivery methods

Source: Gartner (February 2020)

These criteria can form a basis for comparing the providers and tweaking the product to specific requirements.

Step 2: Curate

When you receive strategic TI, perform a fusion process before delivering that intelligence to stakeholders. The fusion process is like that of enrichment, in that it appends context from other sources. For example, if the strategic TI is about a new threat actor and that actor's techniques, then information about how your organization is exposed to those techniques would be valuable to add for readers. Once the strategic TI has been enriched and packaged in a way that is useful for the readers, it should then be delivered. A document directly from an intelligence vendor may lack the important context that will make it impactful for the stakeholders who will receive it.

> The converse is also true: If strategic TI offers no value to the stakeholders, it does not need to be passed on. Do not overwhelm recipients with too much TI, as it will lose its impact and eventually be ignored when it matters.

Step 3: Deliver

Strategic TI is delivered using more traditional methods. Many TI providers provide a portal where the report can be viewed securely. Email and public-viewable webpages are also common. TIPs can also receive the TI document and automatically store it, parse it for tactical TI, and relate all the TI received together.

You can disseminate strategic TI to your organization in the same way that the TI providers do. If the report contains sensitive data that you wouldn't want publicly available, you should use a secure portal. If a TIP is present, it could act as that secure repository, where only those that need access have it.

As an alternative, some organizations use regularly scheduled conference calls to discuss the TI. This approach leads to a much more interactive experience where stakeholders can ask questions. However, the availability of key stakeholders can prove to be challenging.

Steps 4 and 5: Improve and Assess

Strategic TI's effects are not directly linked to detection and response occurrences. Instead, strategic TI can influence an organization's security posture, leading to a more qualitative impact on the organization's overall security position. Strategic TI metrics reported by clients include:

- Security decisions, such as patching prioritization, made based on TI
- Security architecture changes made based on TI
- Number of intelligence reports communicated to senior management and other stakeholders

When you are reporting on the impact of TI, you must tailor the presentation to the audience. Quantitative metrics become less relevant the higher up the reporting chain you go. Instead, a better strategy is to form a story around the metrics using the context provided from the sources of threat intelligence. This story should include the adversaries involved, their motives, and the risk that was averted due to the actions the threat intelligence enabled you to take. For example, TI can enable a vulnerability to be reprioritized as it is seen being exploited in the wild. If there is a known adversary who targets your vertical or geographic area for financial gain, the potential cost of an attack that was averted will resonate with many stakeholders.

Type 2: Tactical Threat Intelligence

Tactical TI often comes in a machine-readable format that allows it to be easily leveraged by security controls. IOCs are the most common forms of tactical TI, as they describe discrete pieces of information that can be detected to generate incident alerts.

Step 1: Collect

Many organizations struggle when selecting threat intelligence data feeds for tactical TI, because it is not clear to them which ones are better than others — or what "better" even means. How do you compare tactical TI feeds? The feeds may contain IP addresses, DNS names, URLs, MD5s and other artifacts, with as much relevant context data as possible.

Using the defined requirements, Tables 4, 5 and 6 showcase the attributes used to compare TI feeds.

Table 4. Simple Comparison Criteria for Tactical TI Feeds

Attribute	Usefulness
Number of Entries	Usefulness depends on the quality of the TI.
Geographic and Industry Coverage	This can be useful if part of the collection requirements.
Classification (e.g., Malware, Exfiltration and C&C)	Classification is the most basic context that should be expected from a feed.
Additional Context Data, or Extended Schema Fields	The more context, the better.
Update Timing	This is useful if tactical TI can be utilized at the same speed. Some consider frequently updated feeds to be better, because the vendor may be doing more work to keep the data fresh.
Customization and Filtering Available	These capabilities reduce the amount of data the organization needs to filter and consume to achieve the same result.
Access to Data via APIs, Query Interfaces, etc.; Prebuilt Integrations With Tool Vendors	APIs are essential for faster, automated collection, enrichment and utilization of TI data. Common ones are REST and TAXII.
Support for Standard Formats (e.g., STIX, OpenIOC, YARA)	Standards make it easy to integrate with many security tools.

Source: Gartner (February 2020)

Table 4 lists the TI feed attributes that are easy to evaluate. For example, it is easy to compare the number of indicators in the providers' databases or the formats in which you can obtain the information. This information is usually readily available. Understanding the context that is delivered

in a TI report should be a big part of the evaluation criteria for a TI feed. Without enough context, analysts will not be able to make informed decisions.

Not all TI providers cover the same kinds of intelligence, which can make some attribute comparisons difficult. For example, adversary information is not included by all TI providers. Additionally, providers that do have adversary information may or may not cover the adversaries that are threats to your organization. The same is true with attributes like confidence. They require special processing, which some providers may not do. Table 5 lists the TI feed criteria that can be challenging to compare.

Table 5. Challenging Comparison Criteria for Tactical TI Feeds

Attribute	Usefulness
Threat Type and Actor Coverage	Tailoring the TI feed to the threats and adversaries your organization faces will make the feed more relevant.
Confidence	Confidence can help with decision making, but the measure provided is not tailored to your organization and may not be relevant.
Overlap With Already-Available TI Data	The TI feed may not add substantially to the stable of already-utilized feeds. Multiple sources reporting the same TI may increase trust in the data.
Preprocessing by the Provider	Enriching with more context can save time and money.

Source: Gartner (February 2020)

Some of the attributes that show value can only be seen once the TI feed is operational in your environment. These attributes are listed in Table 6. Providers also generally don't allow potential customers to view their database of intelligence, which makes understanding how frequently IOCs will be seen in your environment difficult to predict. Once you have access to multiple sources of TI, determining attributes such as exclusivity becomes possible. TI providers don't usually provide that measure, as they would need access to competitors' intelligence. Some can provide comparisons against open-source intelligence (OSINT) feeds.

Table 6. Hard Comparison Criteria for Tactical TI Feeds

Attribute	Usefulness for the TI User
Frequency of Matches With Your IT Environment	This is what makes the feed relevant and actionable.
Frequency of Matches in Your Environment Not Connected to an Ongoing Investigation	This is actionable in the best possible way, but impossible to know in advance.
Frequency of False Positives in Your Environment	This metric is useful in combination with the above attributes. Feeds used for live blocking or real-time alerting must have a low false-positive ratio.
Durability (i.e., Continued Use for Detection Over Time)	Durability makes the feed reliably actionable.
Exclusivity	Combining exclusivity with relevance metrics ensures that the provider's TI feed has unique threats that are relevant.

Source: Gartner (February 2020)

Step 2: Curate

Once your organization starts receiving threat intelligence, whether strategic or tactical, it needs to be managed, or "curated" (see Figure 3). If the threat intelligence is automatically being ingested by a security control from the vendor, there won't be any opportunity to curate it. While this makes dealing with threat intelligence simpler, it also removes the opportunity to enhance or modify the data before it starts being used.

Figure 3. Curation Life Cycle

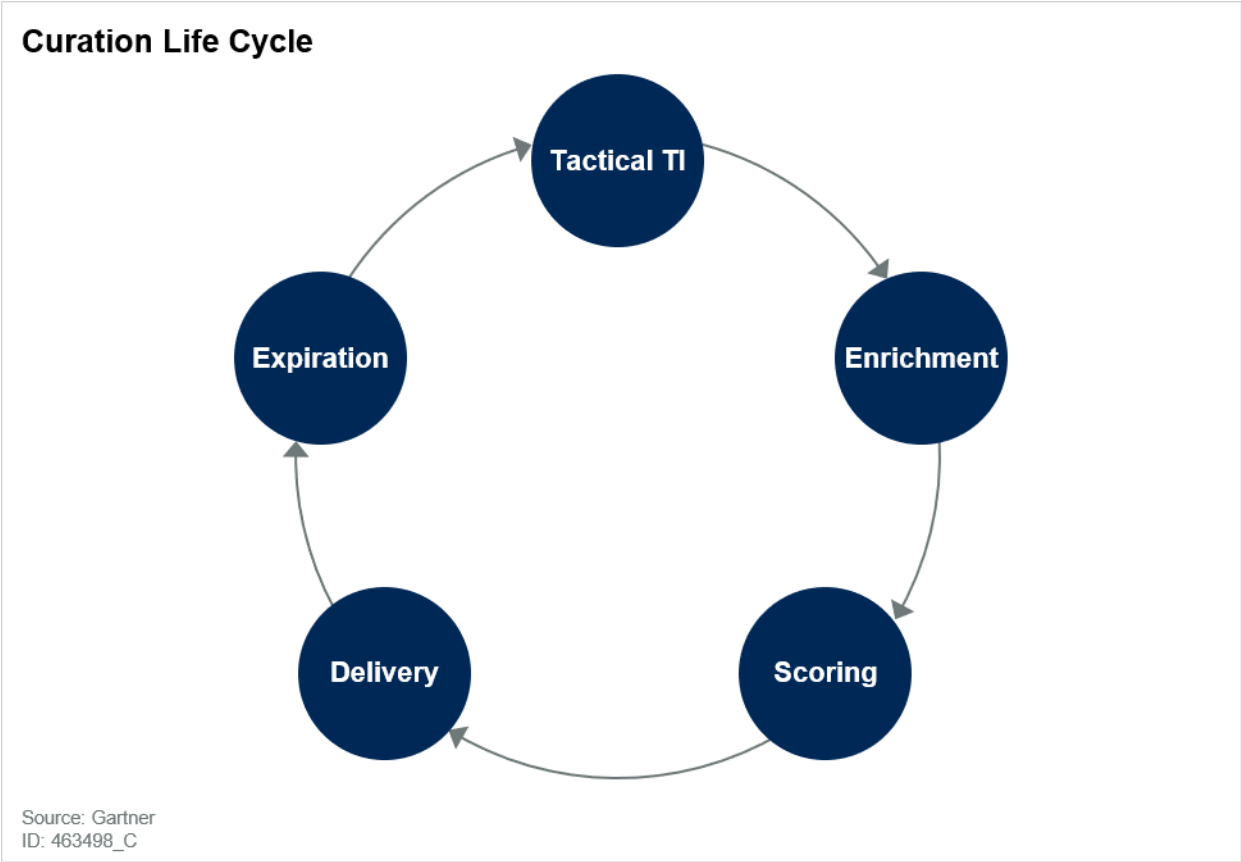

Curation of tactical TI should leverage a repeatable process. As TI is received, it should be enriched with context that allows an analyst or tool to make better decisions based on it. With the additional context, the TI can then be scored or prioritized. The scoring step will help determine which TI should be delivered to security controls throughout the environment. At some point, TI may no longer be valid, so it should be removed from security controls and marked as expired. If the TI is seen from an intelligence provider again, the curation process begins again.

Enrichment

For the sake of simplicity, we will assume tactical threat intelligence consists of IOCs. An IOC, such as an IP address, is not very helpful without context. Some tactical TI feeds may include context, while others may not. If the IOC lacks the necessary context, it can be enriched with more data.

For example, the following could be an enrichment process for IP address IOCs: “For each external IP, get WHOIS data, geolocation data and files associated with the IP.” Such tactical TI data enrichment provides a way to validate weak TI signals. Providers such as DomainTools, Farsight Security, RiskIQ and VirusTotal are useful for enrichment purposes.

The following example shows a generalized version of a typical enrichment process:

1. Collect TI data such as IP address lists, DNS domain reputations, file and process data, and other indicators from public, community and commercial sources.
2. Store data in a way that takes care of duplicate data and links related TI together. Retain counts of how many times each entity appears across sources. Any and all deduplication must be 100% lossless.
3. Enrich newly loaded data with required external context, such as DNS, WHOIS and hashes from public sources like VirusTotal.
4. Relate enriched TI data to enable future insights and analysis. Example relations include IP to domain to URL, domain to registrant, IP to ISP and geographical area, IP to campaign, and file hash to malware family. Relate across TI domains, such as malware data to phishing to intrusion data. (Ideally, relating to strategic TI, such as adversary interests and goals, may happen here.)
5. Relate newly acquired threat intelligence to existing threat intelligence (e.g., a malicious URL to other malicious activities stemming from an IP address where that URL is hosted). This correlation helps to uncover additional traces that can be used to detect and discover threats in the environment.
6. Validate and contextualize the newly loaded data with other historical records, such as logs and flows, as well as with old incident records and existing proven TI datasets.
7. Update any existing adversary profiles based on newly loaded, enriched and related TI data.
8. If desired, adjust confidence and priority ratings for each entity. For example, a new bad IP from an ISP with a long history of hosting malware may be elevated in priority, or an entity that shows up on many lists may be upvoted in confidence.
9. Convert TI data into tool-specific formats based on the entry type. For example, convert malware IPs into network intrusion detection system (NIDS) signatures, other bad IPs into SIEM watchlists, email subjects into data loss prevention (DLP) rules, and file hashes into endpoint detection and response (EDR) solutions.
10. If part of the process, share the intelligence with a trusted circle of other organizations to enable them to detect threats that your organization encountered.

The above process is one of the main reasons TIPs were created. If done manually, this process would be very time-consuming. TIPs can take the steps above and automate most of the process while storing and maintaining relationships between the data.

Scoring

The resources for security monitoring tools, analysts and responders to process TI are finite. In the case of technology, there may be hard limits on how many IOCs a control can hold. For analysts and responders, the resulting alert volume from matched IOCs may exhaust their time and patience. For these reasons, being able to prioritize or score indicators becomes beneficial. Once a score is assigned, decisions can be made on whether to deploy an IOC to security controls.

An example of simple scoring is sending IOCs to a SIEM for matching against ingested logs. However, the SIEM technology has a limit of 10,000 items that can be held for matching. When you are looking at the IOCs provided by the TI feed, the source provides a confidence attribute. The IOC's score can be adjusted based on what that value is. If possible, several attributes provided in the context can be used to make this determination.

An example of complex scoring is sending IOCs to a SIEM, where the TI feed being used provides thousands of IOCs a week. Analysts are reporting a high number of false positives on IP addresses belonging to large web hosts. Using enrichment, a TIP can automatically add WHOIS data to an IOC as an attribute. Once the web hosts that are causing false positives are identified, indicators can start to be scored lower based on autonomous system number (ASN) or other identifying information from the WHOIS data.

When you are dealing with multiple TI feeds, the knowledge of what TI feed is most reliable, trusted and useful comes only after operational usage. Therefore, scoring intelligence sources based on some TI goals and confidence is another part of this process. Prioritizing intelligence is a step that cannot be missed!

Expiration

Tactical TI, specifically an IOC, is generally not valid forever. Attackers change their infrastructure and tools. Sites that are compromised may eventually be fixed, making their IP addresses or domain names legitimate once again. If matching continues against these IOCs, it will generate false positives. These defunct IOCs will also take up space in the lists on your security controls. Thus, they could prevent new IOCs from being detected or hamper performance.

It is important to periodically remove old IOCs through an expiration process. Some TI feeds may handle this expiration process by pruning IOCs at some regular interval or after confirming that they are no longer malicious. Usually with a vendor TI feed, this pruning process will all happen automatically. When using a TIP or TI integration, you will need to determine whether the deletion of IOCs is automated similar to the addition of IOCs.

After an indicator is expired, it still has the potential to appear again in the future. This reappearance starts the whole process over again.

Step 3: Deliver

For TI to reach its potential, it must be delivered to where it is needed in a timely manner. In the case of tactical TI, that would be your security infrastructure. For strategic TI, it would be the stakeholders and other people who need to know. Strategic TI is most often delivered by traditional means, such as email. Websites, TIPs and other TI portals also often host the TI, but require the user to retrieve it themselves.

Tactical TI can be delivered in several ways, most of which will depend on the security control that needs to ingest the data. The most common way to transport the TI is via REST API calls. These are basically just web/HTTP requests. Some tools will pull the TI in, while others will need it pushed to them. This transport method will usually operate over Port 443 with Transport Layer Security/Secure

Sockets Layer (TLS/SSL). Depending on where the TI is located, the security controls may need to reach out to the internet (see Figure 2). The format of the data can be anything from XML to CSV.

There is an effort to standardize how TI should be communicated called Structured Threat Information Expression (STIX). STIX is a schema that allows everyone who uses it to speak the same language. For example, if there is an IP address indicator of compromise, it will be named and represented the same way in all contexts. This makes the entire process of sharing TI easier for everyone involved, as there should be no extra parsing or normalization required.

In STIX 1.x, the data was delivered in an XML format. STIX 2 uses JSON instead. STIX TI can also be delivered as discrete files and over network protocols.

One of the challenges that comes with using STIX is versioning support. Some vendors are quicker than others to update their version of STIX, so there is the possibility of incompatibilities. It is worth mentioning that format conversion is a popular use case for TIPs.

Another standard commonly used with STIX is Trusted Automated eXchange of Indicator Information (TAXII). TAXII is a transport mechanism for STIX data. It basically provides a protocol for security controls and TIPs to exchange STIX information. The communication can be bidirectional, which allows for TI information to be pulled and pushed in an automated fashion.

Detection

Using tactical TI for threat detection is the most common use case. Blocking based on TI is less common, due to the possibility of disruption to the environment. The SIEM is the most common recipient of TI for detection purposes, as it was purpose-built to perform rule matching on log data from many sources. Figure 4 depicts a real-world setup.

Figure 4. Using TI With a SIEM

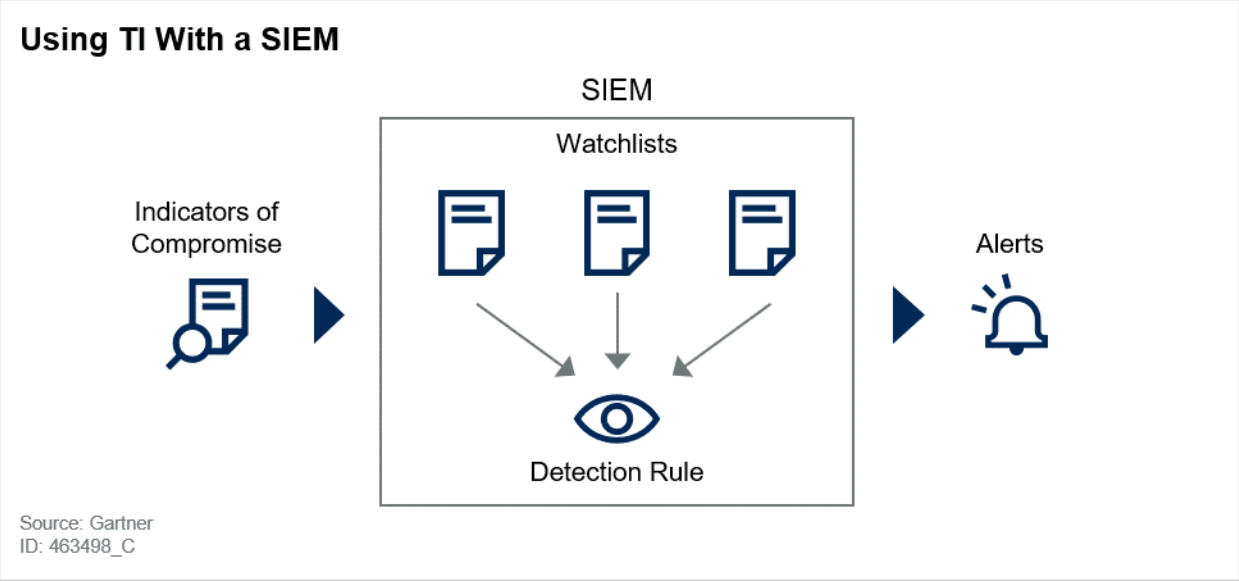

The following is an example of how TI can be used with a SIEM for threat detection:

1. Your TI provider has delivered a list of IP addresses known to be used by the PlugX malware family. The context included with the TI shows the IP addresses are command-and-control servers.
2. The goal is to check whether there have been any connections to these IP addresses through the firewall.
3. By using an API integration or making a manual effort, you add these IP addresses to a watchlist in the SIEM.
4. On the SIEM, IOCs are stored in watchlists. Depending on your SIEM, each watchlist may be able to store only a single type of IOC.
5. Next, the SIEM will need a rule or use case to begin matching the IOC list against events. It should match any IP address fields against data in the watchlist.
6. Optionally, a rear-view search can be run in order to detect past instances of the IOCs.
7. If any of the IOCs in the watchlist are seen in the logs, an alert will be raised.

While the above example applies to SIEMs, it is also very common to use tactical TI such as MD5 hashes to scan endpoints for malicious files. In addition, NIPS signatures can often be built in an automated fashion using tactical TI such as fully qualified domain names (FQDNs) and URLs.

Using tactical TI for prevention can be a risky proposition. It is not uncommon for a legitimate FQDN, IP address or hash to be included in a threat intelligence feed. This can cause serious issues if those incorrect IOCs are sent out to security controls and blocked. Proper curation can help discover these situations before they cause any harm.

A common prevention scenario among some Gartner clients involves blocking IP addresses on their perimeter firewall. One benefit of this approach is that it can cut down on the noise in the firewall and application logs by filtering out the traffic. However, those clients have also reported legitimate IP addresses being included in their threat feeds and causing network issues due to being blocked.

Due to performance limitations, firewalls generally are not the ideal place to store large lists of IOCs to block. There are specialty products, such as Bandura Cyber's Threat Intelligence Gateway, for this type of prevention if required. These products are purpose-built to match large amounts of IOCs and block them if necessary.

Some TIPs provide a matching capability, allowing analysts to remain in the same interface to perform their investigation. This method leverages SIEM APIs and allows work to be offloaded from the SIEM. For example, on an indicator entry, there may be a button to look for any logs that contain the indicator value. This matching function can be performed in bulk on groups of indicators, such as campaigns or adversaries. Being able to check your environment for any evidence of an adversary with a single click is very powerful.

Threat Hunting

Threat intelligence is an important part of threat hunting. Using TI for threat hunting should not be limited to IOCs, as they are easily handled by detection tools. If no detection tool is available for a certain type of IOC, then threat hunters will still need to use that information. What's more important to threat hunters is a type of TI called "tools, techniques and procedures" (TTPs). TTPs can be defined as "patterns of activities or methods associated with a specific threat actor or group of threat actors."[1]

During the compromise of a system, it is often necessary for an adversary to execute custom code for the purposes of a backdoor. There are many ways to accomplish this goal, but adversaries usually standardize their operation around one or two techniques. For example, APT32 is an adversary believed to operate out of Vietnam. APT32 is known to use a technique called "DLL sideloading," which takes advantage of how Windows looks for DLLs it is asked to load.

DLL sideloading involves placing a malicious DLL in a spot where a legitimate executable will load it before seeing the real DLL. Once the malicious DLL is loaded, the adversary's code is run in what looks to be a normal process. This technique is not something that can be detected by signature-based technology. Instead, it is a pattern of behavior that is discovered through forensics and log analysis.

A recent development that has made threat hunting more effective is the ability to codify many of these TTPs using the MITRE ATT&CK framework. Reports on new adversaries and their typical attack methods can be represented in ATT&CK, which translates those attributes into specific artifacts or behaviors a threat hunter can search for. IOCs can be changed (some more easily than others), but TTPs are usually more difficult for an attacker to alter. For example, a new IP address or FQDN can easily be obtained to change where a command-and-control server is located. Alternatively, changing the malware or communication protocols requires much more effort on the attacker's side.

In today's threat landscape, attackers have turned to TTPs. TTP attacks involve using tools that are already on the system, also known as "living off the land." This paradigm, in turn, has led to TI that describes which tools are being used and what kind of behaviors are suspicious (e.g., PowerShell being executed with certain command line parameters). In order to leverage this type of TI, threat hunters require more visibility, such as EDR tools or enhanced endpoint audit logging. It is important to keep these requirements in mind if threat hunting will be one of the uses of TI.

Alert Prioritization

Threat intelligence can also be used to help prioritize alerts instead of just triggering them. For example, an alert shows an outgoing connection to an uncommon country. While this connection could be legitimate, a TI source has labeled the FQDN as a command-and-control server. With that added information, this alert should be given priority. Alert fatigue is a continuing problem with many Gartner clients. If TI is used improperly, the result can be an increase in the number of low-quality alerts.

A common way to accomplish alert prioritization is with a SOAR platform, such as Demisto, Splunk Phantom or Swimlane. These automation-focused platforms will trigger alerts before the analyst is involved in the process. From there, the SOAR platform can use TI to enrich the alert and make decisions.

For example, in Figure 5, a SIEM has a use case to detect HTTP traffic egressing through the firewall via Port 80. The organization uses a secure web gateway, which makes this type of occurrence suspicious. It could be a simple misconfiguration or a malicious application that is not proxy-aware. When the SOAR platform receives the alert from the SIEM, it can check TI to determine whether the destination of the connection is known to be used for malware command and control. If the destination is malicious, the SOAR platform can raise the severity of the alert or issue a ticket.

Figure 5. Sample SOAR Workflow With Threat Intelligence

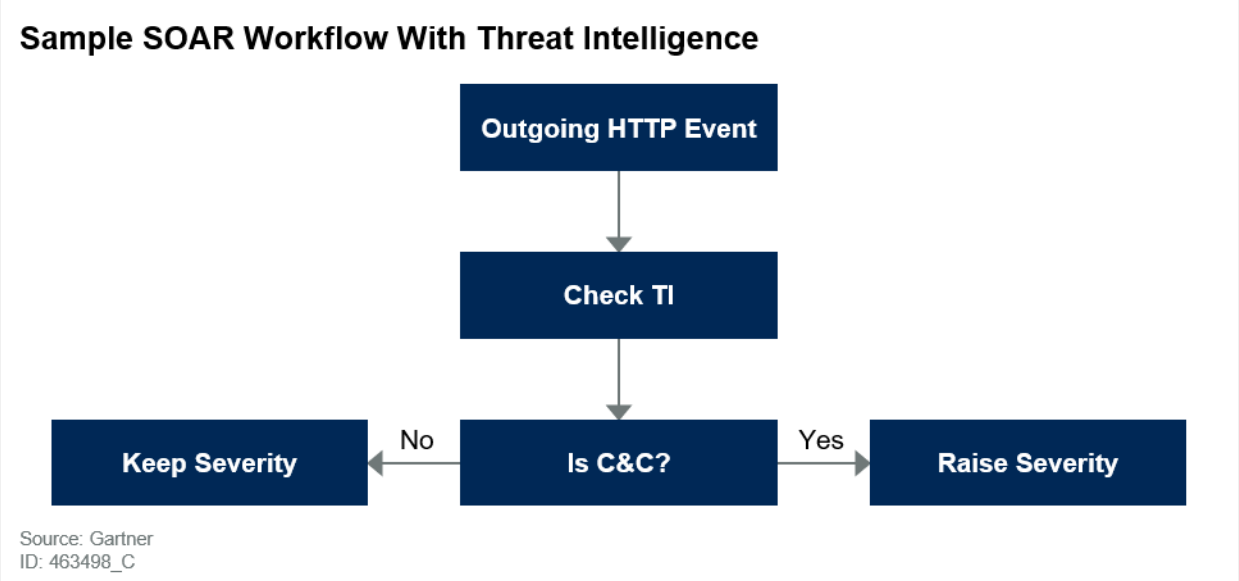

By using TI in conjunction with either built-in automation capabilities or a separate automation platform, organizations can help lessen alert fatigue. As mentioned above, alerts can be automatically closed or escalated. The process does not have to be binary though. Instead, the automation can create tickets with differing severity based on the TI it was able to gather about the incident. This allows analysts to view all the alerts, but still get to the most severe ones first. Using TI to judge the severity allows for more granularity than most built-in rule-based systems, which classify an alert based on a static definition.

Incident Response

When an incident response (IR) is needed, time is of the essence. Strategic and tactical TI can save an IR team a significant amount of time if they can attribute an attack to a known campaign or adversary. TI will provide detailed information on common IOCs and TTPs. This information will give

the IR team a map of the immediate artifacts to look for. Once the IR team understands the threat it is dealing with and has verified that the TI is accurate, containment of the incident can happen quickly.

TI is not a perfect solution, especially with IR. The IOCs and TTPs obtained from TI all happened in the past. It is quite possible that that adversaries have modified their toolkits and TTPs in some way since the TI was gathered. This scenario leads to situations where the TI doesn't quite match what an IR team is seeing during an investigation. Very rarely do the results of an investigation match up exactly with a TI report. Also, many adversary groups share some subset of TTPs. These reasons make attribution very difficult. However, for most IR teams, the goal is incident containment and recovery, not attribution.

Another use of strategic TI is to help prepare the organization for a potential incident. TI can provide valuable insight into which types of threats may target your organization and how they operate. This information can be used for tabletop exercises that provide realistic scenarios for teams to work through. For more information about tabletop exercises, see "How to Implement a Computer Security Incident Response Program."

The result of an IR should be a report that includes a detailed recap of the incident and a list of IOCs found. These IOCs should be added to the existing threat data that the organization is monitoring for. IR teams are often one of the major creators of IOC/TTP information.

Steps 4 and 5: Improve and Assess

To realize the benefits that TI brings, you need to assess how the TI is performing and then make modifications based on the findings. Just as your organization evolves over time, the TI you're using and the way you're using it must evolve too. Replace TI feeds that are not bringing value with others that are more relevant. Over time, adjust your enrichment/fusion processes to make the information more useful to your analysts.

Tactical TI is easily handled with more traditional metrics, such as number of alerts resulting from TI. The list below includes common examples of metrics used:

- **False positives generated from TI feeds:** False positives can waste a lot of an analyst's time and can cause confidence in the accuracy of TI-based alerts to erode. If a TI feed is found to be generating too many false positives, then improving curation or replacing the feed should be considered.
- **True positive alerts from TI:** Alerts that are generated from TI and deemed to be accurate should be measured to show that the feed is relevant, even if no escalation occurs. For example, a download is blocked based on an MD5 hash provided by a TI feed. While there may not be a need to escalate this issue to an incident, the TI proved useful.
- **Incidents discovered primarily due to TI:** TI that results in an incident being created, from either detection or threat hunting, is one of the best ways to show TI's value to a security program. Without the TI, the incident may have gone undiscovered.

- **Improvements in time to containment due to use of TI:** During an incident response effort, reducing the time to containment is critical. TI can shorten this time by providing the responders with information upfront.
- **Incident loss reduction due to early discovery:** Measuring risk avoidance in terms of financial loss can be a powerful metric. For example, a preventive block is put into place for an FQDN received from a TI feed. The domain was seen hosting ransomware. A day later, a phishing campaign is launched against your organization, and the SIEM reports attempts to access that domain. Without that TI, the attack may have been successful, resulting in significant damage.
- **Malicious activities blocked using TI per security device, per intelligence source, etc.:** By tracking which devices and TI feeds are blocking, you can determine where the most value comes from. For example, the metrics may show Feed A results in much more preventive actions than Feed B. Is there value in continuing to receive Feed B?
- **Threat hunts launched and the number of findings:** TI will often provide information about new TTPs in addition to IOCs. When relevant information comes in, hunts should be launched. Tracking the frequency of these occurrences and their results will show the value of the TI. Even if there are no findings, TI provides a level of assurance against that particular threat.

Risks and Pitfalls

Threat intelligence is a key aspect of security operations that will help you detect, triage and investigate threats. However, when embarking on a TI initiative, you must be vigilant to avoid the following risks and pitfalls:

- **Selecting the wrong type of threat intelligence for your organization:** Select TI that is relevant to your organizational needs with respect to adversaries, asset types, industry and geographical locations. For example, if the TI provider does not cover adversaries that target your industry, a great deal of the benefits provided by TI will be lost. If the TI is sourced from unreliable sources, such as generic “honeypots,” it will result in lower-quality alerts that will contribute to increased alert fatigue.
- **Lacking the expertise to properly understand and curate the threat intelligence:** If an organization wants to move beyond importing feeds and into curating intelligence, it will require additional expertise. The management and exploitation of intelligence is a skill and requires a significant amount of time. According to Gartner clients, analysts rarely have the time to devote to such tasks.
- **Deploying too many indicators of compromise or irrelevant ones:** Some security controls have hard limits on the amount of user-defined information, such as IOCs, that they can match against. Firewalls are a typical example of technologies where such limits may exist. Deploying too many indicators may run up against these hard limits or hurt the performance of the security control. Also, the more IOCs that are pushed out, the more potential false-positive alerts that could be generated. Higher quality is better than higher quantity.

- **Lacking a plan to measure the effectiveness of your threat intelligence:** Without metrics and qualitative reports resulting from threat intelligence, it will be difficult to judge the worth of the investments being made. Also, without a plan, it will be hard to tell which TI sources are beneficial and which ones are just causing more wasted time for the analysts. Before using TI, plan which metrics to track.

Gartner Recommended Reading

Some documents may not be available as part of your current Gartner subscription.

"How Gartner Defines Threat Intelligence"

"How to Implement a Computer Security Incident Response Program"

"Information Sharing as an Industry Imperative to Improve Security"

"Market Guide for Security Threat Intelligence Products and Services"

"Innovation Insight for Machine-Readable Threat Intelligence"

"How to Hunt for Security Threats"

"Managing or Setting Up Threat Intelligence and Cyber Hunting Efforts"

Evidence

[1] "Definitive Guide to Cyber Threat Intelligence," CyberEdge Group.

GARTNER HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
USA
+1 203 964 0096

Regional Headquarters
AUSTRALIA
BRAZIL
JAPAN
UNITED KINGDOM

For a complete list of worldwide locations,
visit http://www.gartner.com/technology/about.jsp