Buyer's Guide to Threat Intelligence Platforms

THREATQUOTIENT™

TABLE OF CONTENTS

INTRODUCTION

Many organizations are establishing their own threat intelligence operations, building Security Operations Centers (SOCs), incident response (IR) capabilities and threat intelligence teams. In the process they acquire multiple data feeds, from commercial sources, open source, industry and their existing security vendors - each in a different format. They soon realize they lack the manpower and technology to programmatically sift through mountains of disparate global data and to actually use it. Without the proper resources, the data they've invested in fades into the background and becomes more noise, potentially generating significant false positives.

To use threat intelligence more productively, organizations are investing in a threat intelligence platform (TIP). Selecting a TIP is important as it will serve as the foundation for your entire threat operations program, allowing you to understand and act upon the highest priority threats you face, while enabling you to get more from your existing resources - technology and people.

This guide outlines the essential capabilities you need in a threat intelligence platform and core questions to ask vendors so that you make the best decision for your organization.

Defining a Threat Intelligence Platform

A threat intelligence platform empowers SOCs, threat intelligence analysts, incident response, risk management and vulnerability teams to not only respond to events and alerts, but to also anticipate threats and become more proactive. The key to enabling this is that the TIP serves as the central repository for all threat data from both external and internal sources - fostering collaboration, better decision making, proactive measures and accelerated detection and response. The platform must be able to help you understand, prioritize, and act upon the most relevant threats facing your organization. As the threat landscape and your internal environment changes, the TIP can also help you learn about and anticipate potential threats through continuous threat assessment, so you can proactively strengthen defenses.

Streamlining threat operations and management, a TIP should provide the ability to rapidly bring together internal threat information including security event data, malware analysis and adversary analysis and overlay it with external insights for critical context of the who, what, when, how and why of a threat. A TIP should also help with prioritization, so you can automatically filter out noise and understand what matters to your organization based on parameters you set. Through automation and synchronization of threat intelligence, a TIP should allow you to strengthen the configuration and policies of your security infrastructure proactively and accelerate detection and response efforts. Regular updates with pre-processed, contextual and prioritized data, along with the ability to capture feedback and learnings, empowers teams to re-prioritize and anticipate threats to reduce risk now and in the future.

The TIP and Security Orchestration, Automation and Response (SOAR) Technologies

Automation, playbook, orchestration, operation - a range of terms exist to describe an ability to streamline an existing workflow or potentially chain together a group of workflows based on decision-tree logic. There is a fair amount of confusion around this topic, but Gartner has made strides to shed some light by grouping these capabilities under a category defined as security orchestration, automation and response (SOAR) technologies. TIPs are in this category along with security incident response platforms and security automation and orchestration platforms. While there is overlap across these technologies to varying degrees based on vendor, a key difference is that TIPs are focused on aggregating and enriching intelligence from external and internal data sources for prioritization and export, whereas incident response and orchestration platforms offer capabilities that perform more robust steps within a workflow.

For instance, a TIP may take an indicator with a high threat score and in an attempt to gather as much peripheral data about it, query it against several third-party resources (e.g., DomainTools, OpenDNS, PassiveTotal, VirusTotal, etc.). More advanced intelligence platforms can apply any new learned information and either use it (exporting it to the existing sensor grid automatically) or pivot off it for further analysis without human intervention.

Complementing TIPs are the pure-play orchestration and incident response platforms that automate workflows based on a set of predefined actions to perform tasks with or without human intervention. These platforms rely on threat intelligence to improve the efficacy of security technologies and processes as they coordinate response and facilitate collaboration among team members.

Keep in mind, SOAR technologies are in their infancy so make sure you have a clear understanding of what capabilities you need to improve security operations efficiency, quality and efficacy as you conduct your evaluation of TIPs.

Why You Need a TIP

From the boardroom to the SOC, executives and analysts alike can benefit from a TIP as the foundation to their security operations.

- Chief information security officers can reduce risk, improve defenses and execute on strategic and tactical enterprise goals while staying on budget. They can arm their SOCs, IR teams and threat intelligence analysts with a platform to efficiently structure, organize and utilize threat intelligence across the enterprise.
- Security analysts can improve situational understanding, accelerate detection and response, maximize existing security investments and collaborate more effectively as a team.
- Incident response teams can automate prioritization of threats and security incidents, accelerate investigations and push intelligence automatically to detection and response tools.
- Threat intelligence analysts can efficiently structure and organize threat intelligence with context and prioritization to build adversary dossiers, make better decisions and take action.

Here are the primary use cases to consider as you conduct your evaluation:

- **Curated Intelligence:** Turn threat data into threat intelligence through context and automatically prioritize based on user-defined scoring and relevance.
- **Attack Trends:** Investigate attacks and track over time using the data to improve your defensive posture.
- **Intelligence Pivoting:** Utilize campaign, malware and indicator knowledge to identify related attacks and adversaries that may affect your operations.
- **Breach Investigation:** Support scoping and remediation by correlating artifacts of an investigation with a threat library of related indicators and context.
- **Threat Hunting:** Empower your teams to proactively search for malicious activity that has not yet been identified by your sensor grid.
- **Improve Incident Response:** Gain global visibility to adversary tactics, techniques and procedures to improve remediation quality, coverage and speed.
- **Strengthen Sensor Grid:** Make firewall, IDS, IPS, SIEM and other devices smarter with the most accurate and relevant threat data.
- **Operational ROI:** Retrospectively evaluate your intelligence sources' value, versus the relevance of their information to incidents you experience.

EVALUATION CRITERIA

This guide separates evaluation criteria into two areas: technology considerations and business considerations.

Technology

Ingest Data

The evaluation journey begins at the ability to import various data lakes of intelligence; whether it's external feeds, internal technologies, or analysts' analysis, this is by far the utmost important capability of any TIP. The baseline functionality covers the ability to allow end users to parse and index both structured and unstructured data - both of which continue to be critical for analysts to paint that bigger picture - to help coordinate defenses. For unstructured data (e.g., blogs, whitepapers, Twitter posts) the platform absolutely needs to be able to parse and extract de-fanged or neutered data (e.g., www[dot]badguy[dot]com), which is no small feat because there is no industry standard to rely on. Customers also need the ability to set-up customized STIX/TAXII feeds - the industry's latest push to standardize intelligence terminology and feed syntax. The final element in the baseline functionality is the simple ability to add an indicator and its respective associations using an intuitive user interface (UI) to increase analyst efficiency.

However, as the platform technology has advanced, several additional pivotal capabilities have surfaced to better manage risk factors, including the ability to correlate external intelligence and internal information ranging from malware analysis results, incident response tickets, suspicious events from a SIEM or sensor grid, and even vulnerability assessment results. The platform should allow the customer to define additional custom objects in order to expand the types of intelligence managed. For example, can I add a new "indicator type" to ingest, manage and export suspected malicious BitCoin addresses?

Questions to ask:

1. How many out-of-the-box commercial feeds and open source feeds do you have?
2. Do customers have the ability to enable/disable individual feeds?
3. Do customers have the ability to enable/disable components of a feed? (i.e., I only want to import intelligence associated with industries x, y, and z.)

Context

Context is king! Indicators are purely a means to an end and are only really used for "detection." But the context (or attributes) wrapped around the indicator provide additional supporting information to prioritize and inform how an analyst should react to the alert. Due to the importance of the supporting context, it is important to determine if the TIP vendor imports all of the data or if they modify any of the data, or both. Modification can be helpful as a layer of normalization is critical to de-duplication efforts. For instance, if Feed X publishes https://www.badguy.com, Feed Y publishes http://www.badguy.com and Feed Z publishes www.badguy.com, all three should be reconciled into a single indicator of compromise (IOC) entry. Admittedly those are all technically different indicators, however the goal is to efficiently maximize detection strategies with minimal duplication. Data feed normalization helps to consolidate any analyst's comments, better organize associated intelligence and effectively export one IOC in lieu of three IOCs. Given the volume of domains, URLs and IPs hosting malicious websites published in various intelligence feeds, the normalization of indicators can save a significant amount of resources and reduce analyst confusion.

Customer-defined attributes are the most valuable to a team because they are specific to your organization. The TIP must allow customers to add, modify and delete supporting attribute tags to help mold the product around a team and their organization.

Questions to ask:

1. Are customer-initiated or customer-defined IOC context shared across the vendor's other customers?
2. If an indicator was seen more than once, does the subsequent sighting or context of the indicator override the prior sightings?
3. Am I able to define custom attributes to fit the needs of our organization (i.e., Bitcoin addresses, Twitter handles, SWIFT code)?

Scoring

The volume of indicators being published today is exponentially greater than number of indicators most defensive technologies can actually monitor, making it mandatory that TIP technologies allow customers to score and prioritize intelligence. Prioritization is critical to help drive better decision making across security operations, including orchestration and IR platforms as well as TIPs. All intelligence is not created equal and customers need a mechanism to prioritize which indicators they should block, detect to investigate or even disregard because it does not pose a threat. The indicator score must be specific to that organization.

Intelligence scores based solely on a vendor's own research, the industry's opinion or the community's opinion may not necessarily translate to your team, your tools or your mission. The score itself is typically based on the source of the information, but more progressive TIPs will allow the customer to set their own scoring algorithm based on any piece of information within the system. Prioritization based on parameters you set makes the data within the system more beneficial to your team and more accurate for threat management.

Two other important components are how often the score is re-calculated and the scoring range itself. Scores must be in real-time to ensure the actions taken to block, detect or ignore an indicator, and the decisions made during an investigation are based on the latest data in the system. Some vendors will recalculate an indicator's score hourly, daily or even weekly, which could hinder the effectiveness of the customer's actions. Furthermore, scores should also be re-calculated and automatically adjusted whenever the team resolves a confirmed incident. The platform should do this based on a bidirectional integration with a ticketing system versus having an analyst make the changes manually.

Additionally, a score range of 1-5 is not granular enough and 1-1000 is too complex to conceptualize. The standard scoring range of 1-10 or 1-100 offers the most ideal balance. Platforms that allow negative scores offer an additional dimension to a team's scoring strategy to ensure the important intelligence surfaces to the top, above everything else with a "middle of the road" threat score.

Questions to ask:

1. Can customers customize scoring based on their own organization, team, resources and capability without those customizations being broadcasted to your other customers?
2. Is the vendor scoring transparent?
3. Do you support negative scores?
4. Can I control the score associated to feeds?

Expiration

Most indicators have a limited shelf life, meaning that over time they become less and less of a threat. A core function of the platform is the scoring and ranking of intelligence. However, an independent byproduct of that feature is the ability to expire the intelligence. This is not meant as a "hard delete" from the system because the respective context may be paramount if the indicator's threat resurfaces. Rather, the system needs to have an automatic mechanism to determine when not to export the indicator to the various sensor grid detection tools. The expiration methodology should start with the source, but then factor in indicator type and other elements based on a customer's environment. Customer-defined expiration is critical because it should be based on your resources - team and sensor grid technologies. All sensor grid technologies (firewalls, web-proxy, endpoint, IDS/IPS, etc.) have limitations on how much intelligence they can monitor, so the expiration methodology cannot be dictated by vendors, industry or anybody else outside of the customer's environment. Analysts must also be allowed to manually override an expiration date. The final capability required for expiration (as discussed above for scoring) is that the platform can automatically pull in an investigation and automatically adjust the expiration date for the intelligence within the ticket versus having an analyst make the changes manually for every confirmed incident the team resolves.

Questions to ask:

1. What is the vendor's approach to expiring intelligence?
2. Can I adapt the expiration methodology to align with my customized scoring and capabilities of my sensor grid technologies?
3. Can the TIP automatically adjust expiration dates based on parameters I set?

Correlate Internal and External Data

The most important and valuable feature of a platform is its ability to correlate external and internal intelligence and overlay it with internal network activity with as little manual intervention as possible. The more control and automation a customer's team can leverage at the intersection with their SIEM, malware sandbox, ticketing system, vulnerability management system, asset management system, etc., the more value gained from the platform. This is where customers need to put a lot of focus on evaluating on-premise versus cloud platforms (especially multi-tenant platforms) because integrating across tools adds significant network overhead to deploying, managing and optimizing cloud-based systems, including requiring additional open ports in the firewall. Relying on the TIP provider's own disaster recovery and network resilience can also limit control and introduce risk.

Automatically correlating and deploying threat intelligence to your sensor grid is only half the battle. Re-ingesting the post-mortem results from investigations and alerts will help the platform self-tune through scoring and prioritization. If the intelligence was accurate, then the threat score increases and is automatically re-adjusted to help defend the customer against future attacks. However, if the intelligence was inaccurate (i.e., false positive), the threat score can decrease at the customer's discretion.

Questions to ask:

1. If bidirectional data is enabled, does your company have sole ownership rights to my company's data within the system?
2. To address the integrations I need, must I open additional ports on the firewall?
3. Do we need to pay more for application programming interface (API) use for integrating internal and external data? Is the additional API cost a flat fee or is it a "pay-by-the-drink" model?

Integrations

As mentioned previously, a huge value proposition for any platform is connectivity to the organization's ecosystem of tools - SIEM, malware sandbox, ticketing system, IDS/IPS sensors, firewalls, DNS, web-proxies, endpoint solutions, vulnerability management solutions, data-leak prevention (DLP) technologies, etc. The more technologies that can exchange data, the less manual work required by the analyst and the higher the efficacy of an operations team. Integrations have two primary factors to consider - the direction and degree of integration. The direction of integration includes unidirectional and bidirectional.

Unidirectional is the most basic integration, for instance, from the intelligence platform into a firewall, IDS/IPS, or endpoint solution. This is a purely defensive strategy and the most common integration, moving the automatically scored highest threats from your intelligence platform into the trenches of your sensor grid for detection and blocking. A common misconception is that you can bypass pushing data to a sensor grid and only send the unidirectional feed into your SIEM. This loses efficiency because most organizations don't have the budget or infrastructure to funnel 100% (or even 70%) of their logs and network traffic through a SIEM, which means your highest intelligence threats are only being correlated against a subset of your data and likely only the SIEM's escalated alerts.

Another common and critical unidirectional integration is from your malware sandbox into your TIP. By definition, sandbox technologies monitor and capture attacks. Pulling that data into an intelligence platform is a huge benefit for correlating internal curated threat intelligence and pivoting to malware hashes, command and control channels, import hashes, compile timestamps, mutexes, packers, attributed malware family, and other associated tags. Admittedly, depending on your sandbox's capability this could be a pretty large feed by itself (some sandboxes cannot aggregate specimens across operating system and application detonations) so ideally the integration should be able to be configured to ingest malware results deemed to pose a threat.

Bidirectional integrations (i.e., push and pull data to the tool, or getting information back into the TIP) are the wave of the future because they offer a full circle automated capability which empowers analysts to make significantly faster and better decisions using the data already at their fingertips. There are three distinct bidirectional uses cases that are becoming a team's core modus operandi for improving cyber defenses including SIEM or log repository, ticketing system and vulnerability management solution.

SIEM or Log Repository:

Bidirectional interconnectivity between the intelligence platform and the customer's SIEM or log repository offers the biggest time savings and is commonly referred to as the "rear-view mirror" search. The workflow is simple yet extremely powerful. As intelligence is ingested into the platform and scored, the threats with the higher scores are then queried against the customer's data archive (since a majority of the time intelligence is shared shortly after attacks are launched). This provides the biggest benefit because without intelligence platforms this workflow is often completely skipped due to the painstaking effort and the amount of time it takes to gather all unscored information and perform the search. By automating the workflow, analysts can now focus on higher priorities and when there is a rear-view mirror event, all the information is instantly at the analyst's fingertips without having to log into several applications or wait for the data to display.

Ticketing System:

The bidirectional interconnectivity between the ticketing system and an intelligence platform is unique because there are core workflows for starting at both the ticketing system or the platform. In the case where the integration flows as follows: *ticketing system -> intelligence platform -> ticketing system*, the process provides an enrichment benefit to help jumpstart an investigation. As teams have collected and expanded intelligence for nearly a decade, they have amassed a significant amount of data. Unfortunately, all the supporting data cannot be exported to the sensor grid. Instead, only a subset of the supporting intelligence (usually an executive summary or the latest information) is exported. However, by using a TIP, when a ticket is created and populated with indicators, the ticketing system can be configured to automatically query the intelligence platform for any and all the supporting information. This drives efficiencies, quality and efficacy by accelerating the investigation with deeper intelligence than just the executive summary or latest information from the indicator.

As mentioned, the inverse workflow also is critical to hone the best intelligence possible. In the case where the integration flows as follows: *intelligence platform -> ticketing system -> intelligence platform*, the ability for the TIP to tune itself becomes center-stage. Ticketing systems hold the final outcome of each and every investigation - true incident, false positive, or benign traffic - so pulling that data back into the platform and validating or re-scoring the indicator fine-tunes the system. If the incident is deemed a false positive, then the re-ingestion of that information provides a negative impact on the intelligence and a lower score is re-calculated. If the incident is categorized as benign traffic and re-ingested, then the score may remain the same. If the incident was a true infection, the re-ingestion can hold steady at a high threat score or even increase the threat level of that indicator in order to extend the indicator's expiration date. This workflow is also one of several allowing companies to dictate which source of intelligence is quantitatively their most valuable.

Vulnerability Management Solution:

The next phase of bidirectional interconnectivity is overlaying the attacker's attempts, the internal alerts and, more importantly, the internal vulnerabilities to discover possible attack routes and jump ahead of the adversary. It is critical that a TIP has the ability to ingest vulnerability data, match that against an attacker's tactics, techniques and procedures (TTPs) and then automatically query a customer's environment to determine which endpoints as well as whose endpoints are most likely to face the attack. Risk management and patching solutions are poised to patch the most critical infrastructure first to better protect the "crown jewels." However, until now that prioritization has been done without a core component - threat intelligence. With a TIP, companies can re-adjust their patching prioritization based on the adversary's historical attacks and previous lateral movement attempts. That combination is valuable because it allows defenders to stay vigilant as adversaries will mimic previous successful movements.

Questions to ask:

1. Do I need to pay extra for API use for integrations?
2. Do you have bidirectional integration with all the SIEMs, ticketing systems and vulnerability management solutions?
3. What other tools do you support with bidirectional integration?
4. Do I need to engage professional services to handle integrations?
5. Does SIEM information become shared information?

Notifications

A platform must be able to streamline many of the repetitive efforts, including the ability for the system to notify an analyst when a certain adversary attack is discovered, if certain adversary infrastructure is active again or even if a certain keyword is found in an intelligence report saved in the system. Analysts should be able to raise notifications or alerts on any object within the system - indicator, incident, adversary, etc. - and even objects they create including Bitcoin addresses, honeypot account codes, or vulnerabilities in key applications within their own environment. The alert notification can range from user interface notification on a dashboard to an email notification.

Questions to ask:

1. Can an analyst create an alert list within your dashboard on any object in the system?
2. Is the alert notification within the UI, email or another third-party client (e.g., Slack, HipChat, RSS feed, etc.)?

Export

The true value of a platform is not only to aggregate and prioritize intelligence, but also to export or transport the data for other systems or analysts to consume. Exporting data seems like an easy feature, but there are several hurdles including format, sequence of export (because most tools cannot handle the volume pushed to them), which supplemental tags are needed to support the IOCs and what output file format to use. To deliver the most value, exporting must be done in a way that facilitates the use of all features supported by a particular tool. Analysts also require the ability to create new exports based on their own needs, their role, their investigation or purely for their exploratory research. For instance, an export may require all the intelligence revolving around:

- a certain adversary, malware family or exploit kit
- within the past 8 months
- targeting “my” and adjacent industries
- with a threat score higher than 7 (out of 10)
- export the data in a JSON or even STIX format

These are five of the most common export elements, but a platform should allow you to have nearly limitless capabilities to manipulate and craft the intelligence in a manner to empower detection and blocking but also investigations and hunting expeditions. Finally, the export should support baseline users as well as seasoned analysts who require advanced scripting capabilities.

Questions to ask:

1. Does the export include the most common out-of-the-box file formats (e.g., CSV, JSON, CIF, etc.)?
2. Can an analyst export any object and any supporting context?
3. Does the export support a scripting language to allow comprehensive control over the type and format of information being exported (i.e., can an analyst define what intelligence is being exported and output it in a technology specific format within a single UI)?
4. Can an analyst configure multiple export feeds (i.e., by sensor technology, per geographic location or to support daily exploratory research)?

Sharing and Collaboration

To this point we've addressed sharing in terms of ingesting data from external feeds and an organization's ecosystem of tools, and exporting threat intelligence to other systems or analysts to consume. The ability to normalize structured and unstructured data as well as support bidirectional integration are essential to reduce data fragmentation and gaps in defenses. But there's another aspect to sharing and collaboration - the human element.

As a central repository, a TIP should enable teams to work together more efficiently, continuously augment and enrich threat intelligence and share learnings from any location, at any time in order to accelerate threat detection and response. The ability to collaborate provides teams with utmost control over the who, what, when and how of the threat - the context that allows them to prioritize and focus on mitigating the greatest risks to their organization. To facilitate the use of the TIP for sharing and collaboration, the SOC, IR team, threat intel analysts and network team must all be able to use and update the TIP as part of their existing workflow. Commentary and data can be stored for longer periods of time than with other tools, such as SIEMs, and instantaneously accessed by all team members to share information for better decisions. This also reduces the challenge of "brain drain" that occurs when team members leave the organization; knowledge is captured and retained despite any personnel turnover. Integrating into existing systems - including, but not limited to SIEM, log repositories, ticketing systems, incident response platforms, orchestration and automation tools - will allow disparate teams to use the tools and interfaces they already know and trust, and still benefit from and act on that intelligence.

Another aspect to sharing and collaboration is sharing your enriched threat data externally or with communities such as Information Sharing and Analysis Centers (ISACs). Technology vendors use the threat data you share to enhance their products, like threat intelligence feeds, for other customers. Organized by industry, ISACs also share data across member organizations in your specific sector. As you evaluate membership in threat intelligence communities, be sure to understand the level of control you have over what, when and how much data is shared.

Questions to ask:

1. Can the TIP serve as a shared workbench for all members of the broader security team (i.e., IR, threat intel, hunters, management, etc.)?
2. Are we able to integrate the collaborative functionality into our existing workflows? If so, how and is there additional cost involved with this integration?
3. Can we opt-in and opt-out of sharing data with a vendor or community?
4. Is the shared data anonymized and how?
5. How is the shared data used?
6. Is the TIP vendor assuming ownership rights to any data shared within its platform?

Business

Deployment Options

Deployment options and definitions can vary based on the vendor so be careful to assess each one individually and not make any assumptions. General guidelines for evaluation, along with pros and cons are outlined below.

Off-Premise Architectures:

These solutions are the easiest to turn on, hit the ground running and can offer ease of management and financial savings. But they also pose challenges when integrating with your on-premise sensors and systems. To do so requires either opening more ports in a firewall which can be extremely challenging to justify to the IT team, or standing up a small on-premise integration server to serve as the local liaison to organize feeds. If your team is going through the trouble of implementing an on-premise integration server, it is definitely worth considering reaping the full benefits of the on-premise platform instead.

Utilizing a cloud platform, whether multi-tenant or single tenant, also means relying on having a 100% uptime line-of-sight to the platform. It is rare, but in several recent cases large cloud provider networks have gone offline, and as a result, several platform vendors' capabilities have been offline as well. Consider also the extreme case where a customer is dealing with a massive breach and needs to disconnect from the Internet to control it. Losing access to their intelligence platform is counterintuitive to its purpose.

The other potential shortcoming when leveraging platforms that are cloud instances is the fact that, in turn, customers are losing ownership rights to the intelligence put into the vendor's cloud. Typically, a cloud vendor's terms and conditions state that customers are surrendering ownership of any customer data put into the vendor's cloud solution and the vendor has the rights to do whatever they would like with the data - even repackage it in their own data feed. It is critical to note that most teams collect intelligence from within circles of trust where the information shared cannot be shared outside the team's organization (or even the individual in certain cases). Yet, putting that intelligence into a vendor's cloud platform will absolutely violate the terms of the agreement. In most cases this is done inadvertently, but can result in a negative outcome if that vendor uses that data in a report, data feed or even during their own research efforts. Finally, cloud providers introduce an additional attack vector to monitor and protect as adversaries are becoming extremely proficient at stepping through large cloud providers undetected.

Questions to ask:

1. Am I the sole owner of my data?
2. Is my data being co-mingled with other customers in a multi-tenant environment?
3. What capabilities do I lose if I do not deploy an on-premise integration server? (If the vendor says "None!", turn and run.)
4. If a cloud provider's cloud infrastructure goes offline, what exact functionality is lost?
5. Can you share copies of your quarterly third-party performed penetration tests?

On-Premise Platforms:

These solutions have their own set of pros and cons. For on-premise deployments which leverage physical hardware, customers can expect an initial hurdle getting the system provisioned and installed. Most organizations need to work with their datacenter team which can take two hours or two weeks. Utilizing an on-premise virtual machine is usually a swifter process to stand - up, but often customers must befriend their virtual administrators in order to allocate dedicated virtual memory which every box in the cluster is fighting for. Frequently, in virtual load balancing, core services (such as payroll or the financial database) are prioritized for memory over security tools. On-premise deployments also typically have to sort through the complex web of local proxy configurations-especially in global or government WAN setups. However, once deployed, the benefits of a true on-premise approach are significant including the ease of integration across a customer's toolset as well as always having access to your data and maintaining ownership.

Pricing Models

As you establish your threat intelligence program, you need to understand who, how and where you will use the TIP so that you can accurately evaluate pricing models across vendors.

This first part - the who - is covered by the annual subscription fee and number of user licenses. The subscription fee for the platform is usually based on the capacity of the platform and often includes any maintenance and management fees which should be minimal. User license packages typically start at five to ten users and may step up all the way to an unlimited option. Obviously, the more user licenses, the lower the price per user. An annual subscription fee and user licenses should provide a level of predictability. To plan and budget appropriately, take into consideration the tactical users (security analysts, intelligence analysts, etc.) but also as the platform harbors collaboration (which is your goal), forecast the access need for the risk management, vulnerability assessment or fraud teams.

The second component of the pricing evaluation - the how - requires understanding your data import and export requirements. To get the most use out of your threat intelligence you must be able to easily and affordably integrate the TIP with your existing defenses. As discussed throughout this guide, integrations allow you to increase security posture and the value you get from your existing security investments. Some vendors charge a fee, as much as several thousand dollars per integration, which can easily double the total cost of the TIP when you consider integrating to your SIEM, firewalls, anti-virus, endpoint detection and response, intrusion detection/prevention, web application firewalls, IR ticketing systems, vulnerability management solution, case management systems, proxies, etc. Likewise, consider the threat feeds and any custom data you plan to integrate into the TIP and understand if there are any costs associated with these integrations. Two very important notes to take into account: first, many organizations grow organically through mergers and acquisitions and rather than unify security technologies to conform with the parent organization they maintain steady-state which can double the number of integrations you need to purchase from the vendor (if they price integrations "pay-by-the-drink"). Second, whether through merger and acquisition or just a cognizant business decision to maintain federated independent business units throughout the organization, at some point all analysts should be using the platform. Your five-person security analyst team in Scottsdale, AZ could easily morph to a 25-person security analyst team to include the personnel from four of your other business units.

NOTE: If you are a managed security service provider (MSSP), understand if there is a specific price list or if you must purchase off the commercial list. TIP providers interested in working collaboratively with MSSPs will have a different engagement process based on shared success.

The third component is where you choose to deploy the platform. As with any enterprise application deployment, if you deploy the technology in the cloud you need to consider the cost of hosting. If you host the platform on-premises you should factor in your own data center costs and rack space. However, there is a twist with TIPs. If you are evaluating a cloud-based service but know you will need to deploy a private cloud instance for compliance or privacy requirements, be sure to understand if there are any additional costs and tradeoffs in functionality. A TIP designed to run in the cloud often cannot offer full functionality on premises.

Questions to ask:

1. Is there a cost per integration or API with each defense system? If so, what is that cost?
2. Is there a cost associated with integrating custom data or IOCs?
3. What is the total cost of ownership given my business requirements?
4. Are there additional costs associated with a private instance of a cloud-based deployment?
5. Can we adjust the number of user licenses without penalty?
6. Is there an unlimited users license option?
7. Is there special pricing for managed security service providers (MSSPs)?

Support

Receiving assistance with a TIP when needed should be easy. You should have the ability to quickly contact technical support and receive assistance not only with the baseline functionality of the platform, but also to request new features from the vendor or to report bugs or other challenges with using the platform and have these issues resolved efficiently. A good support organization should be able to provide you with documentation detailing their process for handling customer-reported issues and answer any questions you have about how your issues will be treated.

Your TIP provider should also be able to easily provide you with a user manual and any supplemental documentation regarding interaction with the software. This could include:

- an API or software development kit (SDK) guide
- a product knowledge base
- release note archives
- online video webinars

The technical support department should be easily accessible during hours that align with yours, but at a minimum, during normal business hours (Monday-Friday, 8am-5pm). Support should be reachable via email and phone, but providers may also offer more immediate assistance via chat or private Slack channels. Slack is an increasingly popular and effective way to communicate among technical teams as many already use these channels as part of their existing processes and workflows.

Questions to ask:

1. How do I contact support?
2. What SLAs are offered in regard to support tickets?
3. How do I update support tickets?
4. How can I escalate an issue?
5. How are feature requests handled and how quickly are they addressed?
6. How are bug reports handled and how quickly are bugs typically resolved?
7. What is your return merchandise authorization (RMA) policy? (if applicable)

CONCLUSION

As you look to establish your own threat intelligence operations and select a threat intelligence platform solution, there are several criteria to consider. The evaluation process can be overwhelming. But armed with a guide that outlines the core components, technical and business considerations, key questions to ask and potential hidden risks, you can navigate the process successfully and find the right platform to meet your requirements.

Evaluation Quick Reference: TIP Buyer’s Guide

Ingest Data	
How many out-of-the-box commercial feeds and open source feeds do you have?	
Do customers have the ability to enable/disable individual feeds?	
Do customers have the ability to enable/disable components of a feed? (i.e., I only want to import intelligence associated with industries x, y, and z.)	
Context	
Are customer-initiated or customer-defined IOC context shared across the vendor’s other customers?	
If an indicator was seen more than once, does the subsequent sighting of the indicator override the prior sightings?	
Am I able to define custom attributes to fit the needs of our organization (i.e., Bitcoin addresses, Twitter handles, SWIFT code)?	
Scoring	
Can customers customize scoring based on their own organization, team, resources and capability without those customizations being broadcasted to your other customers?	
Can I design the scoring algorithm to determine what information it is based on?	
Do you support negative scores?	
Can I control the score associated to feeds?	
Expiration	
What is the vendor’s approach to expiring intelligence?	
Can I adapt the expiration methodology to align with my customized scoring and capabilities of my sensor grid technologies?	
Can the TIP automatically adjust expiration dates based on parameters I set?	
Correlate Internal and External Data	
Do we need to pay more for API use for integrating internal and external data?	
If bidirectional data is enabled, does your company have ownership rights to my company’s data within the system?	
To address the integrations I need, must I open additional ports on the firewall?	

Integrations	
Do I need to pay extra for API use for integrations?	
Do you have bidirectional integration with all the SIEMs, ticketing systems and vulnerability management solutions?	
What other tools do you support with bidirectional integration?	
Do I need to engage professional services to handle integrations?	
Does SIEM information become shared information?	
Notifications	
Can an analyst create an alert list within your dashboard on any object in the system?	
Is the alert notification within the UI, email, or another third-party client (e.g., Slack, HipChat, RSS feed, etc.)?	
Export	
Does the export include the most common out-of-the-box file formats (e.g., CSV, JSON, CIF, etc.)?	
Can an analyst export any object and any supporting context?	
Does the export support a scripting language to allow comprehensive control over the type and format of information being exported (i.e., can an analyst define what intelligence is being exported and output it in a technology specific format within a single UI)?	
Can an analyst configure multiple export feeds (i.e., by sensor technology, per geographic location or to support daily exploratory research)?	
Sharing and Collaboration	
Can the TIP serve as a shared workbench for all members of the security team?	
Are we able to integrate the collaborative functionality into our existing workflow? If so, how and is there additional cost involved with this integration?	
Can we opt-in and opt-out of sharing data with a vendor or community?	
Is the shared data anonymized and how?	
If data is shared, how is it used by the vendor?	
Is the TIP vendor assuming ownership rights to any data shared within its platform?	

Deployment Options	
Am I the sole owner of my data?	
Is my data being co-mingled with other customers in a multi-tenant environment?	
What capabilities do I lose if I do not deploy an on-premise integration server?	
If a cloud provider's infrastructure goes offline, what functionality is lost?	
Can you share copies of your quarterly third-party performed penetration tests?	
Pricing Models	
Is there a cost per integration or API with each defense system? If so, what is that cost?	
Is there a cost associated with integrating custom data or IOCs?	
What is the total cost of ownership given my business requirements?	
Are there additional costs associated with a private instance of cloud-based deployment?	
Can we adjust the number of user licenses without penalty?	
Is there an unlimited users license option?	
Is there special pricing for managed security services (MSSPs)?	
Support	
Are there various methods to contact support?	
What SLAs are offered in regard to support tickets?	
How do I update support tickets?	
How can I escalate an issue?	
How are features requests handled and how quickly are they addressed?	
How are bug reports handled and how quickly are bugs typically resolved?	
What is the RMA policy? (If applicable)	