THREATQUOTIENT

Buyer's Guide to Threat Intelligence Platforms

www.threatquotient.com

Contents

INTRODUCTION

Many organizations are establishing their own threat intelligence operations, building Security Operations Centers (SOCs), incident response capabilities and threat intelligence teams. In the process they acquire multiple data feeds, some from commercial sources, some open source, some industry and some from 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 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.

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

Additionally, when thinking of threat intelligence many organizations fail to include internal data — the telemetry, content and data created by each layer in their security architecture, on-premises and in the cloud. In addition to the SIEM, this includes data from modern security tools and technologies, like Endpoint Detection and Response (EDR), Network Detection and Response (NDR) and Cloud Detection and Response (CDR). Not only is this data high fidelity, it's also free!

To use all their threat intelligence and data more productively, many organizations are investing in a threat intelligence platform (TIP). Selecting a TIP is important as it will serve as the foundation for your entire security 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 TIP 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 also to anticipate threats and become more proactive. The key to enabling this is that the TIP serves as the central repository for all threat data and intelligence from both external and internal sources. This creates a single source of truth enriched with context to understand 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 focus on 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.

With these capabilities, a TIP supports multiple users and use cases, beyond threat intelligence management to include threat hunting, incident response, spear phishing, alert triage and vulnerability management.

© ThreatQuotient, Inc.

The TIP as an enabler of SOAR and XDR

Gartner defines security orchestration, automation, and response (SOAR) as being grounded in the convergence of three technology solutions — security orchestration and automation, threat intelligence platforms and security incident response platforms. TIP capabilities are important because a data-driven approach to SOAR provides high confidence in the intelligence being used, the decisions that are made and the incident response workflows that are executed. Whereas a process-driven approach to SOAR, focusing on automating playbooks, is inherently complex because playbooks need to be individually configured and maintained. When playbooks are data-driven the intelligence resides in the platform and not in individual playbooks which provides for simpler configuration and maintenance and more efficient and effective automation. As you conduct your evaluation of TIPs, consider use cases and the gaps needed to fill to improve security operations efficiency, quality and efficacy. This perspective will prove useful in determining if a TIP can also meet your SOAR needs. See the discussion of some of the top use cases a TIP can help you address.

With respect to Extended Detection and Response (XDR) solutions, Gartner defines this emerging category as solutions that "automatically collect and correlate data from multiple security products to improve threat detection and provide an incident response capability." The data management aspect, consuming and correlating the necessary data for effective detection and response across the enterprise, is a core capability of a TIP. Organizations looking to evolve their security roadmap to an XDR architecture, should evaluate a TIP with that end goal in mind and its ability to support and enable a successful evolution.

As you conduct your evaluation of TIPs, consider use cases and the gaps needed to fill to improve security operations efficiency, quality and efficacy.

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.

- **CISOs** can reduce risk, improve defenses and execute on strategic and tactical enterprise goals while staying on budget. They can arm their SOCs, Incident Response 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.

© ThreatQuotient, Inc.

HERE ARE THE PRIMARY USE CASES TO CONSIDER AS YOU CONDUCT YOUR EVALUATION:

THREAT INTELLIGENCE MANAGEMENT: Automatically aggregate, normalize, de-duplicate, analyze and turn threat data into threat intelligence through context and automatic prioritization 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 teams to proactively search for malicious activity that has not yet been identified by your sensor grid.

INCIDENT RESPONSE: Gain global visibility to adversary tactics, techniques and procedures (TTPs) to improve remediation quality, coverage and speed.

VULNERABILITY MANAGEMENT: Prioritize and respond to software, hardware and network vulnerabilities based on active and relevant adversary TTPs.

SPEAR PHISHING: Track and understand methods used by attackers to target organizations in order to identify users that may have succumbed to a malicious email and mitigate risk.

ALERT TRIAGE: Contextualize and automatically prioritize alerts so that teams can focus on the most relevant alerts and collaborate to address alerts in the "gray zone".

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.

© ThreatQuotient, Inc.

EVALUATION CRITERIA

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

TECHNOLOGY

Consume Structured and Unstructured Data

The evaluation journey begins at the ability to import various data lakes of intelligence; including from internal technologies, external feeds and analysts' analysis. This starts by aggregating internal data from across the entire ecosystem — the telemetry, content and data created by each layer in your security architecture, on-premises and in the cloud. In addition to the SIEM, this includes data from modern security tools and technologies, like Endpoint Detection and Response (EDR), Network Detection and Response (NDR) and Cloud Detection and Response (CDR).

End users must also be able to parse and index structured and unstructured data from external feeds — 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 . 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 structure/syntax. And when new threats emerge around events, for example COVID-19 or the SolarWinds Orion security breach, what's needed are custom connectors to any type of threat intelligence feed that can be written and deployed within hours so organizations can begin ingesting threat data from new sources quickly. 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 or to fit a specific use case. For example, can I create a set of objects to fit a vulnerability management use case or to support an intelligence requirement to comply with a specific workflow?

© ThreatQuotient, Inc.

TO ENSURE A TIP SUPPORTS THIS BASELINE FUNCTIONALITY, THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU WILL NEED TO ASK:

1. How many "out-of-the-box" commercial feeds and/or 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.)
4. Do customers have the ability to write their own feeds in a documented language?
5. Is it possible for customers to adapt the data model to specific use cases or risks associated with their unique environment?

Context / Transparency

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 provides additional supporting information to prioritize and inform how an analyst/team should react to the alert. Due to the importance of the supporting context, it is important to determine if the TIP vendor imports all the data and/or if they modify any of the data. Modification can be helpful as a layer of normalization is critical to de-duplication efforts. However, normalization and unification of data must be done while preserving context. 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 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. It is also important to have the ability to translate data back into the necessary formats for use with the existing tools teams are using.

Context also makes it possible to map source information into a specific model. When a feed provides information about associated objects like malware, campaigns or attack patterns, it is important that the TIP has the ability to retrieve this information and map it to the customer's object model. This can be done by either creating relationships with existing objects or by creating these objects and the relationships between them to reflect the full context of what has been provided.

Along the same line, it is important for a TIP to allow you to use your preferred analytic framework. For example, allowing customers to use the Lockheed Martin Cyber Kill Chain model or the MITRE ATT&CK framework and providing the relevant object and/or attribute capabilities to do so.

© ThreatQuotient, Inc.

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/delete supporting attribute tags to help mold the product around a team and organization.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU WILL NEED TO ASK.

1. Are customer-defined IOC tags/context/attributes 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. Can I control and adjust the score/confidence associated with feeds?
4. Can I ingest all levels of threat intelligence from vendors (from strategic to tactical)?
5. Can I map information to an analyst framework (e.g., MITRE ATT&CK)?

Scoring / Prioritization

The volume of indicators being published today is exponentially greater than the 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 "real-time" to ensure the actions taken to block/detect/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. Further, scores should also be re-calculated and automatically adjusted whenever the team resolves a confirmed incident. The platform should do this based on a bi-directional integration with a ticketing system versus having an analyst make the changes manually.

© ThreatQuotient, Inc.

Additionally, a score range of 1–5 is not granular enough and 1–1000 is too complex to conceptualize. The standard scoring range of –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.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU WILL NEED 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 ingest all levels of threat intelligence from vendors (from strategic to tactical)?
5. Can I control the score/confidence associated with 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 or "bump" it back. The final capability required for expiration (as discussed above for scoring) is that the platform can automatically pull in an investigation/alert 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.

© ThreatQuotient, Inc.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU WILL NEED 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 + External Data

The most important and valuable feature of a platform is its ability to correlate internal and external 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, SOAR solution, 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 vs. 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.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU WILL NEED TO ASK.

1. If bi-directional data is enabled, does your company have sole ownership rights to my company's data within the system?
2. Do we need to pay more for 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?
3. To address the integrations I need, must I open additional ports on the firewall?
4. Are post-mortem results incorporated back into the platform for learning and improvement?

© ThreatQuotient, Inc.

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, SOAR solutions, XDR solutions, 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 efficiency of an operations team. Integrations have two primary factors to consider — the direction and degree of integration. The direction of integration includes uni-directional and bidirectional.

Uni-direction is the most basic integration and encompasses a single direction 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/or 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 and/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 uni-directional 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 families, and other associated tags. Admittedly, depending on your sandbox's capability this could be a pretty large feed by itself (i.e., 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.

Bi-directional 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. The TIP vendor should offer a software development kit (SDK) and open APIs so that it is easy for customers to build their own integrations or customize integrations, both uni- and bi-directional.

There are four distinct bi-directional use cases that are becoming a team's core 'modus operandi' for improving cyber defenses including SIEM or log repository, ticketing system, vulnerability management solution, and SOAR solution.

SIEM or Log Repository

Bi-directional 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

© ThreatQuotient, Inc.

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, now analysts can 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 bi-directional 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. So, instead, only a subset of the supporting intelligence (i.e., usually an executive summary and/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 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: i*ntelligence 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.

This is also a way of capitalizing on lessons learned from incident response. Not only is previously known information enriched as described above, but additional information like IOCs can be found and automatically added to the platform. So by creating this feedback loop, not only do customers gain better visibility and understanding of already known threats, they are also building organizational memory which makes it possible to understand and respond to newly discovered threats more quickly.

Vulnerability Management Solution

The next phase of bi-directional 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 have 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

© ThreatQuotient, Inc.

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.

SOAR Solution

Bi-directional integration is also important when working with SOAR solutions. First, the intelligence platform should contextualize data before it is ingested by the SOAR solution for further action. This approach minimizes the number of playbooks the SOAR solution executes by as much as 80% and ensures the output is relevant and high priority thereby saving valuable analyst time. The threat intelligence platform should also capture data after playbook execution by the SOAR solution, storing it for further analysis. This closed-loop approach uses data derived from each completed playbook execution to improve the quality of security operations.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU WILL NEED TO ASK.

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

© ThreatQuotient, Inc.

Notifications / Alerts

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, investigation/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 (UI) notification on a dashboard to an email notification.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU WILL NEED TO ASK.

1. Can an analyst create an alert list within your dashboard on any object/node 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:

1. a certain adversary, malware family, or exploit kit
2. within the past 8 months
3. targeting "my" and adjacent industries
4. with a threat score higher than 7 (out of 10)
5. 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. And finally, the export should support baseline users as well as seasoned analysts who require advanced scripting capabilities.

© ThreatQuotient, Inc.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU WILL NEED 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 and/or to support daily/weekly 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 bi-directional integration are essential to reduce data fragmentation and gaps in defenses.

However, a TIP should also be able to help you improve data utilization by enabling teams and organizations that make up your entire enterprise to share that information. Think about the following scenarios:

- Government entities with distinct threat intelligence teams and missions that are federated and need to collaborate and share relevant intelligence.
- Commercial organizations with locations worldwide or segmented business units that have different risk profiles based on geographic-, partner- and sector-specific nuances.
- Managed Security Services Providers (MSSPs) that provide multi-sector or geographic coverage to their customers.

A subset of data needs to be sent to each team or location for consistent detection around the globe and to ensure global security risk is covered. The data that is transferred should be curated for local consumption, based on parameters set by the entities who will be receiving the data. To achieve this goal, the TIP should be able to exchange information with other TIP technologies and of course other instances of the TIP. It should be able to offer granular selection of the information that needs to be shared and the ability to anonymize the data.

There's also another aspect to sharing and collaboration — the human element.

As a central repository, a TIP should enable teams to work together more efficiently, and 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

© ThreatQuotient, Inc.

organization. To facilitate the use of the TIP for sharing and collaboration, the SOC, incident response 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, SOAR 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 and/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.

The ultimate form of collaboration involves tasking and coordination to conduct investigations efficiently and effectively. Most security operations or investigations are rife with chaos as teams act independently and inefficiently with limited visibility into the tasks other teams or team members are performing. However, with a single collaborative environment that fuses together threat data, evidence and users, all team members involved in the investigation process can collaborate. Rather than working in parallel, they can automatically see how the work of others impacts and further benefits their own work. Managers of all the security teams can see the analysis unfolding, which allows them to act when and how they need to, coordinating tasks between teams and monitoring timelines and results. Embedding collaboration into the investigation process ensures that teams work together efficiently to take the right actions faster to more effectively mitigate risk.

TO ENSURE A TIP SUPPORTS THIS BASELINE FUNCTIONALITY, THERE ARE SEVERAL CORE QUESTIONS YOU WILL NEED 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 the threat data shared with other parts of the organization be curated for local consumption?
4. Can we opt-in and opt-out of sharing data with a vendor or community?
5. Is the shared data anonymized and how?
6. How is the shared data used?
7. Is the TIP vendor assuming ownership rights to any data shared within its platform?

© ThreatQuotient, Inc.

Data-Driven Automation and Investigations

Security teams require the ability to automate "Tier 1" repetitive, low-risk, time-consuming tasks, and tools that aid in investigation when an analyst is working on high impact, time-sensitive "Tier 2 / Tier 3" incidents. The best approach provides a balance between automation and manual investigation ensuring that teams always have the best tool for the job, and follows a data-driven approach to improve the speed and thoroughness of the work.

Automation capabilities should be native to the threat intelligence platform and also available via integration to other solutions. In either case, a data-driven approach is necessary. Data-driven automation enables teams to define triggers for a playbook execution whenever specific conditions are met. For example, "automatically enrich an event when an IOC exceeds a specified score and the threat targets a specific industry." This data contextualization reduces unnecessary automation by ensuring action is only taken against relevant and high priority alerts. It also allows decision logic to be isolated from playbooks which simplifies ongoing maintenance and changes.

Threat intelligence platforms with a data-driven approach to automation also capture the results of each completed automation and store it for further analysis. This closed-loop model enables the security team to continually improve operations.

Security teams have access to dozens of technologies, feeds, and third-party data sources which can present a challenge in bringing this wealth of data together into a common work surface to investigate and stop attacks.

To support teams working "Tier 2 / Tier 3" level events, TIPs should provide broad situational awareness with visualization tools that also enable documentation and collaboration. Best-in-class TIPs help analysts fuse together threat data, evidence and users, ultimately accelerating the analysis of active threats and their remediation. With the ability to build data-driven incident, adversary, and campaign timelines, analysts can quickly understand how an incident unfolded up to the eventual response. As with automation, it is important for manual investigations to take a data-driven approach, whether responding to alerts or proactively hunting threats.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU WILL NEED TO ASK.

1. Does the TIP support data-driven automation natively and through API integration with SOAR platforms?
2. Can the TIP reduce unnecessary playbook executions by providing contextualized data to filter out irrelevant, low priority events?
3. Does the TIP provide a feedback loop to continually improve the way automation is implemented in the SOC?
4. Does the TIP provide a common work surface to investigate and document incidents?

© ThreatQuotient, Inc.

BUSINESS

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 size/capacity of the platform and often includes any maintenance and management fees which should be minimal. User license packages typically start at five to `10 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 fosters collaboration (which is your goal) forecast the risk management team or vulnerability assessment team or even fraud team to need and want access.

The second component of the pricing evaluation — the how — requires you to understand 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, EDR, IDS/IPS, 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 consider. 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"). Secondly, whether through merger and acquisition or just a business decision to maintain federated independent business units throughout the organization, at some point all analysts should be using the platform. So your 5-person security analyst team in Scottsdale AZ could easily morph to a 25-person security analyst team to include the personnel from 4 of your other business units.

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/features. A TIP designed to run in the cloud often cannot offer full functionality on premises.

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.

© ThreatQuotient, Inc.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU SHOULD ASK.

1. Is there a cost per integration/API with each defense system? If so, what is that cost?
2. Is there a cost associated with integrating custom data/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 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 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.

© ThreatQuotient, Inc.

THERE ARE SEVERAL CORE VENDOR QUESTIONS YOU SHOULD 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 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.

© ThreatQuotient, Inc.

EVALUATION QUICK REFERENCE: TIP BUYER'S GUIDE

CONSUME STRUCTURED AND UNSTRUCTURED DATA	
How many "out-of-the-box" commercial feeds and/or 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.)	
Do customers have the ability to write their own feeds in a documented language?	
Is it possible for customers to adapt the data model to specific use cases or risks associated with their unique environment?	
CONTEXT / TRANSPARENCY	
Are customer-defined IOC tags/context/attributes shared across the vendor's other customers?	
If an indicator was seen more than once, does the subsequent sighting or context of the indicator override the prior sightings?	
Can I control and adjust the score/confidence associated with feeds?	
Can I ingest all levels of threat intelligence from vendors (from strategic to tactical)?	
Can I map information to an analyst framework (e.g., MITRE ATT&CK)?	
SCORING / PRIORITIZATION	
Can customers customize scoring based on their own organization, team, resources and capability without those customizations being broadcasted to your other customers?	
Is the vendor scoring transparent?	
Do you support "negative" scores?	
Can I ingest all levels of threat intelligence from vendors (from strategic to tactical)?	
Can I control the score/confidence associated with 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 + EXTERNAL DATA	
If bi-directional data is enabled, does your company have sole ownership rights to my company's data within the system?	
Do we need to pay more for 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?	
To address the integrations I need, must I open additional ports on the firewall?	
Are post-mortem results incorporated back into the platform for learning and improvement?	

© ThreatQuotient, Inc.

INTEGRATIONS	
Do I need to pay extra for API use for integrations?	
Do you have bi-directional integration with all the SIEMs, ticketing systems, vulnerability management solutions and SOAR solutions?	
What other tools do you support with bi-directional integration?	
Do I need to engage professional services to handle integrations?	
Does SIEM information become "Shared Information"?	
NOTIFICATIONS / ALERTS	
Can an analyst create an alert list within your dashboard on any object/node 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 and/or to support daily/weekly exploratory research)?	
SHARING AND COLLABORATION	
Can the TIP serve as a shared workbench for all members of the broader security team (i.e., IR, threat intel, hunters, management, etc.)?	
Are we able to integrate the collaborative functionality into our existing workflows? If so, how and is there additional cost involved with this integration?	
Can the threat data shared with other parts of the organization be curated for local consumption?	
Can we opt-in and opt-out of sharing data with a vendor or community?	
Is the shared data anonymized and how?	
How is the shared data used?	
Is the TIP vendor assuming ownership rights to any data shared within its platform?	
DATA-DRIVEN AUTOMATION AND INVESTIGATIONS	
Does the TIP support data-driven automation natively and through API integration with SOAR platforms?	
Can the TIP reduce unnecessary playbook executions by providing contextualized data to filter out irrelevant, low priority events?	
Does the TIP provide a feedback loop to continually improve the way automation is implemented in the SOC?	
Does the TIP provide a common work surface to investigate and document incidents?	

© ThreatQuotient, Inc.

PRICING MODELS	
Is there a cost per integration/API with each defense system? If so, what is that cost?	
Is there a cost associated with integrating custom data/IOCs?	
What is the total cost of ownership given my business requirements?	
Are there additional costs associated with a private instance of a 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 service providers (MSSPs)?	
SUPPORT	
How do I 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 feature requests handled and how quickly are they addressed?	
How are bug reports handled and how quickly are bugs typically resolved?	
What is your RMA policy? (if applicable)	

© ThreatQuotient, Inc.

ThreatQuotient improves security operations by fusing together disparate data sources, tools and teams to accelerate threat detection and response. ThreatQuotient's data-driven security operations platform helps teams prioritize, automate and collaborate on security incidents; enables more focused decision making; and maximizes limited resources by integrating existing processes and technologies into a unified workspace. The result is reduced noise, clear priority threats, and the ability to automate processes with high fidelity data. ThreatQuotient's industry leading data management, orchestration and automation capabilities support multiple use cases including incident response, threat hunting, spear phishing, alert triage, vulnerability prioritization and threat intelligence management. ThreatQuotient is headquartered in Northern Virginia with international operations based out of Europe and APAC. For more information, visit **www.threatquotient.com**.

© ThreatQuotient, Inc.

TQ-RPT01-1121-01