SCALING SECURITY OPERATIONS THROUGH THREATQ'S™ ROI

by Ryan W. Trost, ThreatQuotient

THREATQUOTIENT™

WHITEPAPER

SCALING SECURITY OPERATIONS THROUGH THREATQ'S ROI

INTRODUCTION

Security operations teams are expanding their role within the enterprise enclave and wrangling cyber intelligence as a top initiative for survival. Teams are finding their stride, maximizing the benefits of intelligence as ThreatQ™ plays a cornerstone role to help augment their success. ThreatQ enables teams to make a significant leap forward in their current capability to manage intelligence, automate processes and identify and triage investigations exponentially faster. This paper will discuss how ThreatQ delivers return on investment (ROI), focusing specifically on these three capabilities:

- Intelligence Scoring and Prioritization
- Automation and Sensor Grid Hardening
- Time-to-Detect/Time-to-Respond Improvement

By deploying ThreatQ to provide structure to cyber intelligence and perform core workflows, organizations can regain the analytical productivity of between 6 and 9 full-time enterprise Security Operations Center (SOC) analysts.

SCORING: THE GROUND TRUTH OF SCORING INTELLIGENCE

The sheer volume of indicators being published on a daily basis continues to climb at a staggering rate and unfortunately, the industry is starting to asphyxiate itself trying to manually validate the intelligence.[1] In 2016, Carnegie Mellon released a whitepaper[2] which analyzed the open source Blacklist Ecosystem consisting of over 180M indicators of compromise (IOCs) — limited to IPs and FQDN. IOCs trigger alerts, which, in turn, initiate analyst investigations. However, alerts are generally not created equal and manpower is wasted chasing ghost alerts [false positives]. A customer-defined scoring methodology allows the team to dictate their own risk posture based on their resources, detection tools and other team priorities.

ThreatQ's scoring and prioritization capability automates what currently takes an entire team of analysts to do, including managing the intelligence lifecycle of gathering, sanitizing, researching and deploying threat intelligence. In addition, since scoring is customer-specific and customer-defined, it empowers analyst teams to determine their own risk levels and apply that configuration to all incoming intelligence. The result is process optimization with an ROI of 2-3 full-time employees (FTEs).

One of the primary purposes of scoring intelligence is to have the utmost control and accuracy over the whole dataset — not to rely on an external rating, feed or community to solely dictate what is more evil versus less evil in your environment.

THE POWER OF THREATQ'S SCORING METHODOLOGY IS TO:

1) Allow customers to control their own destiny by accurately reflecting their risk temperature
2) Allow customers to customize the scoring algorithm — *because nobody knows your environment better than you*
3) Provide a scoring range that is digestible: 1-5 isn't drastic enough and 1-1000 is too difficult to conceptualize and deviates significantly from industry best practices
4) Find middle ground — too few scoring elements does not accurately reflect the risk, whereas too many elements overwhelms teams and is ignored
5) Offer score transparency AT ALL TIMES, displaying how the score was calculated!
6) Ensure that scores reflect local environmental variables (i.e., sandbox results, ticket results, observed dates, etc.)
7) Most importantly, update the indicator score every time a new piece of information is appended to it

1 Often, teams will "trust but verify" external intelligence to minimize the operational impact a bad piece of intelligence might incur. Some adversaries are known for hiding malicious files in legitimate hosting services to increase the probability of unfiltered access.

2 resources.sei.cmu.edu/asset_files/WhitePaper/2016_019_001_466029.pdf

THE OPERATIONAL SCORING ALGORITHM IN MOTION

Ingesting IOCs across several feeds over a four-month span produced nearly 1M IOCs. When a customer applies the ThreatQ customized scoring framework based on their risk levels, it filters the intelligence into a manageable subset. ThreatQ offers a spectrum of scoring approaches that evolve with a team's experience level. Here's the progression:

Entry Approach: Newer teams will leverage all the indicators equally. A quick view shows a total number of ingested indicators in one dataset is **910,321**.

Maturing Approach: More experienced teams will factor in the commercial alert's severity. The breakdown of indicators by the provider's confidence score (low/medium/high) is below, which still shows a hefty volume of indicators.

Malicious Confidence = HIGH = **319,754**

Malicious Confidence = MEDIUM = 497,969

Malicious Confidence = LOW = 80

Seasoned Approach: The most powerful capability is applying a custom score that aligns with your resources, threats and capabilities. In this use case, let's maintain the vantage point of an SOC Manager within the defense industry and apply higher custom scores to the attributes that are more meaningful to "my" organization. In Figure 1, we focus on "increasing" the pertinent risk scores versus strategically allocating negative values for attributes on the opposite end of the spectrum which pose little to no risk to the organization.

By applying the custom scoring algorithm, the ~1M is broken down into risk categories as follows:

Risk Category	Number of Indicators	Percentage
Very High	27,358	~3%
High	45,651	~5%
Medium	312,623	~34%
Low	248,211	~27%
Very Low	276,478	~30%

Figure 1. The breakdown of intelligence scores after applying ThreatQ's user-defined custom scoring algorithm.

ThreatQ's ability to automatically re-score the ~1M indicators to less than 10% [~72,500 indicators] without requiring constant analyst intervention eliminates the need for a team of analysts performing the typical tasks of the intelligence lifecycle. This also provides a level of threat uniformity, as each analyst's threshold of risk is slightly different, leading to an unbalanced threat-scoring standard.

With ThreatQ, customers quickly become more strategic about WHAT intelligence warrants immediate deployment versus requires additional research. Figure 2 demonstrates how simplistic, granular and yet powerful the user-defined scoring configuration is. Intelligence with more confidence can be escalated to block technologies based on the risk it poses to the organization. For example, threat intelligence with higher threat scores, and thus more reliable, will be deployed to blocking technologies (i.e., firewalls, IPS, web-proxy, endpoint, etc.), whereas intelligence that poses less of

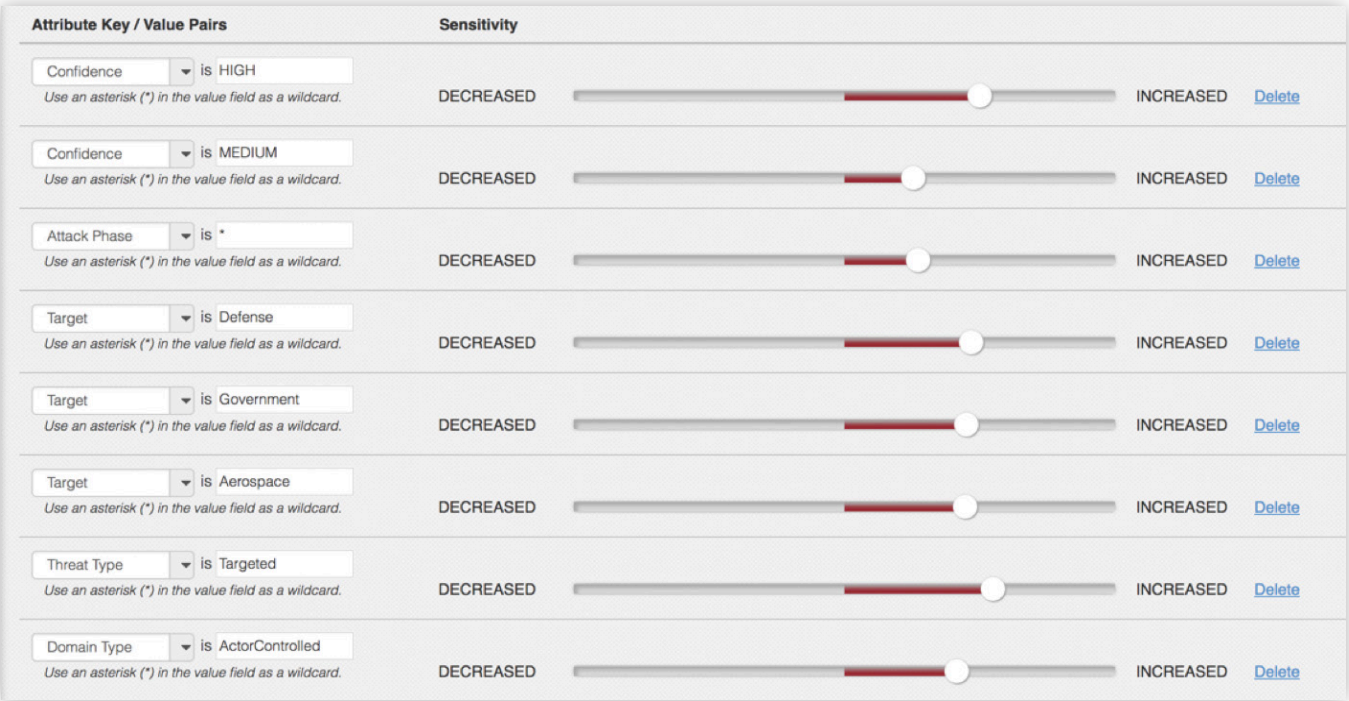

Figure 2. An example of how a customer can define their scoring temperature.

a threat [read: less reliable] will be distributed to detection technologies (i.e., IDS, NetFlow, etc.) to minimize any operational impact due to false positives. This is a critical ROI for companies with limited infrastructure tools already pushed to sensor limits.

HARDEN THE SENSOR GRID: ACCURATELY THROTTLE THE FLOODGATES

ThreatQ's ability to immediately and automatically harden your sensor grid based on IOC score, custom export and bi-directional rear-view mirror search alleviates much of the current, manual and fragmented effort required. ThreatQ automates the process to update the sensor stack within potentially minutes of being imported, providing an ROI of 1-2 FTEs. In a previous life, my team would collect industry intelligence from various sources. However, since we were not network administrators, my team lacked the authority to deploy intelligence. We had to rely 100 percent on the network and system administrators to push intelligence to the sensor stack (i.e. firewalls, web-proxy, SMTP defenses, etc.). Unfortunately, their priority was not my team's priority, so it would often take 24-48 hours (or more) before our gathered intelligence was deployed. This operational hurdle caused my team significant heartburn, as we didn't exactly know "when" the latest intelligence would be deployed.

To ensure maximum threat defenses, teams deploy intelligence hourly to align with most industry intelligence providers'. The FTE ROI revolving around the sensor grid is not dependent on the volume of intelligence, but rather, the hourly disruption of the engineer's workflow. This includes each hour network engineer(s) need to stop what they are doing, log into EACH sensor technology (i.e., firewall, router, email, web-proxy, DNS, endpoint, etc.), upload and QA the latest intelligence, and, finally, return to where he/she originally left off. Automating the application of intelligence to the sensor stack strengthens defenses by orders of magnitude and also unburdens the network/infrastructure team, freeing them up to stay focused on their priorities.

As previously mentioned, the ability for ThreatQ to apply your customer-defined scoring algorithm plays a pivotal role in what intelligence is exported. ThreatQ's powerful export query language, called SMARTY, provides customers with a nearly endless capability to export exactly what they need. This allows users to not only identify WHAT information they want to export to the sensor grid, but also provides the ability to dictate the output format of choice. Below are two sample exports to demonstrate the possible complexity your team is empowered to handle.

Example #1 – Export IOCs associated with *ExploitKit* OR *MalwareFamily* OR *AdversaryName* AND with a ThreatScore >= 65 to ALL BLOCKING TECHNOLOGIES

This is a standard customer export that will deploy high-risk intelligence associated with a certain exploit kit, malware family or adversary set. The export will run hourly to ensure the latest internal and external intelligence is distributed to "blocking" security infrastructure.

Example #2 – Export IOCs associated with CVE-2017-7477 AND where vulnerable_host_systems > 1 to ALL BLOCKING TECHNOLOGIES.

This demonstrates ThreatQ's ability to overlay intelligence with a company's internal vulnerability data to help determine threat levels. This export will gather only the indicators associated with CVE-2017-7477 if, and only if, the organization has a vulnerable host. This marriage between intelligence and vulnerability management, as demonstrated in Figure 3, is the wave of the future to help organizations standardize defenses across departments.

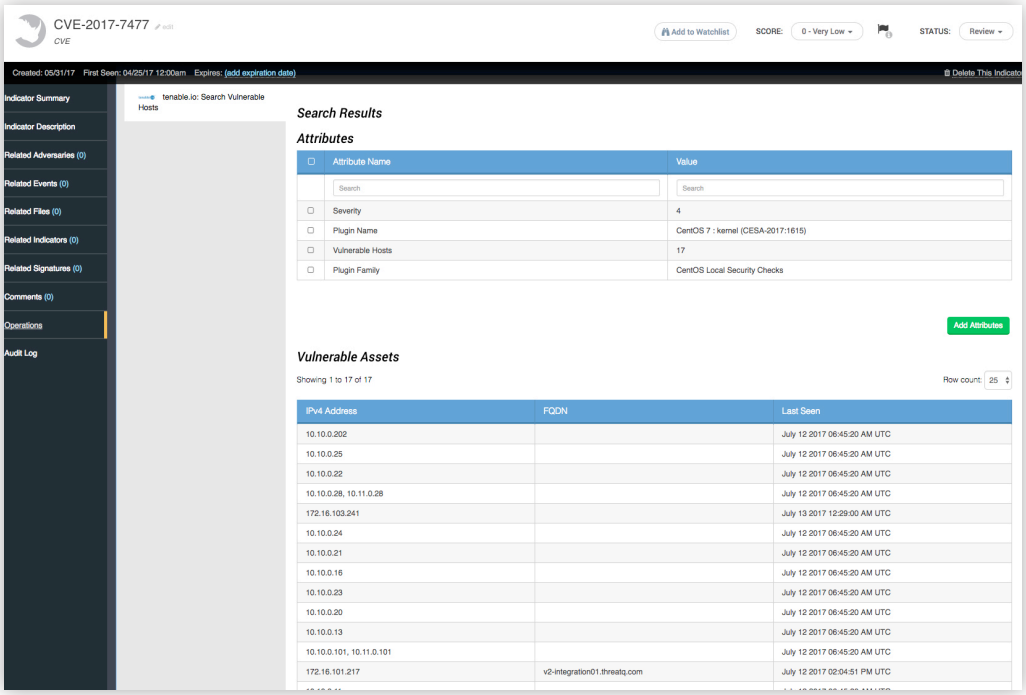

Figure 3. A screen capture of CVE-2017-7477 with related attributes and list of vulnerable assets.

TIME-TO-DETECT (TTD) OR TIME-TO-RESPOND (TTR): KNOW WHAT YOU DON'T ALREADY KNOW... FASTER

ThreatQ's ability to customize scoring and prioritization, self-tune its Threat Library™ and deploy intelligence automatically to an organization's defenses significantly helps the analyst re-gain time. However, ThreatQ's ability to minimize adversary dwell time provides another ROI of 3 to 4 FTEs. The industry has come to the realization that intrusions will happen, so the goal is to minimize time-to-respond. SANS recently published a report by Matt Bromiley, "The 2016 SANS Incident Response Survey,"[3] which states 21 percent of the survey participants reported adversary dwell times 2-7 days before detection and another 2-7 days before remediation efforts. That's nearly half-a-month of lateral movement, privilege escalation, data and credential exfiltration, preparing data ransom, or simply deploying sleeper backdoors for future use.

The TTD/TTR ROI stems largely from the time it takes analysts to search log repositories for high threat/risk intelligence results from the scoring algorithm. Log repositories are vast data lakes of information and often, rear-view mirror searches beyond near-term tactical correlation capabilities can take a considerable amount of time — search times increasing exponentially the further back the search requires. Industry best practices dictate that a precautionary due diligence rear-view mirror search is a minimum of 30-45 days in order to identify any previous attack attempts — successful or mitigated. The ROI estimate is based on performing thorough search queries on 50 alerts — a 10% hit rate on the 500 high risk indicators per day.

The TTD/TTR capabilities that allow ThreatQ to deliver significant ROI also provide the following benefits:

- Provide as much situational understanding as possible
- Maximize bi-directional integration with core tools
- Leverage the ThreatQ operations to further extend automation

SITUATIONAL UNDERSTANDING

Situational awareness is all the context surrounding an indicator or piece of intelligence. Indicators are largely arbitrary pieces of information, but the context is what helps the analyst or incident responder determine the appropriate course of action. The more context around the piece of intelligence, the faster they can make a decision. Historically, when the alert was triggered in the SIEM, the analyst immediately opened several web browsers and proceeded to research the indicator. This took a considerable amount of time, as the analyst needed to click between screens to intuitively summarize the threat data. ThreatQ eliminates all of that by centralizing the context and packaging that up in the SIEM to provide analysts with situational understanding. When the alert is triggered, all the additional contextual information or even specific intelligence reports are already at the analyst's fingertips — saving 5-10 minutes of research per indicator/event.

BI-DIRECTIONAL INTEGRATION

Most intelligence spreads "after the fact," meaning teams are implementing blocks after an intrusion has occurred. To mitigate risk at this point, you need to be able to look back over your log data and apply the intelligence to it as well. Why? Because once adversaries have a foothold into the organization, they will pivot to "new" C2 infrastructure or move laterally to minimize the probability of detection. This makes it even more difficult to trace patient zero. Several commercial intelligence feeds realize that "lengthy" rear-view mirror searches can have a negative operational impact on customer technologies, so they include a "Last Seen" attribute respective to an indicator. This is a huge benefit to organizations gleaning large amounts of internal and external intelligence because they can now stay laser-focused with their log repository search. In this example, we are going to leverage a feed's "Last Seen" date and create an operation to take that date and search "+/- 5 days." The + and – five days is arbitrary and can be defined by the user, but is meant to offer some wiggle room to identify the attack campaign while lessening the burden on already overwhelmed technologies. This search operation becomes more paramount if the "Last Seen" date extends beyond what the SIEM/log repository keeps in memory.

```
Function search(indicator)
    if indicator.last_seen:
        last_seen = indicator.last_seen
        last_seen_range = last seen - 5 + " - " + last_seen + 5
    else:
        last_seen_range = now() - "30d"
    results = search_siem(indicator,last_seen_range)
    return results
```

3 https://www.sans.org/reading-room/whitepapers/incident/incident-response-capabilities-2016-2016-incident-response-survey-37047

THREATQ OPERATIONS

Leveraging the ThreatQ Operations module through the SDK/API capability helps to deliver significant ROI benefits. Best practice dictates that organizations treat intelligence from a sandbox and/or ticketing system differently than intelligence gleaned from external entities. This is primarily because intelligence related to an active or direct attack warrants a longer rear-view mirror search in order to identify previous successful infections or attempts. Why can't you treat all intelligence like this? Based on the volume of intelligence and the duration of the search, such an approach will likely bring your SIEM or log repository to a crawl or experience packet-loss. By automating several historically manual, time-consuming and repeatable tasks through ThreatQ's RESTful API access, customers can save a significant amount of time and further augment their ROI.

THE BOTTOM LINE

Throughout this whitepaper, we have explored several use cases catering to both mature and less mature intelligence efforts, but putting everything together offers the most compelling ROI case.

There are two critical components to highlight in the examples:

1) ThreatQ is performing four of the five steps automatically without analyst intervention, saving substantial time and effort, and
2) these steps are cyclical and repeated, regardless of a team's skillset, resources, budget or capability, which demonstrates the additional value of ThreatQ to augment operational routines by automatically adjusting the self-tuning library. If, for instance, the analyst had deemed the triaged event a false positive, ThreatQ can apply a negative score to the intelligence so that in the future, that piece of intelligence is automatically categorized as less of a risk against the organization.

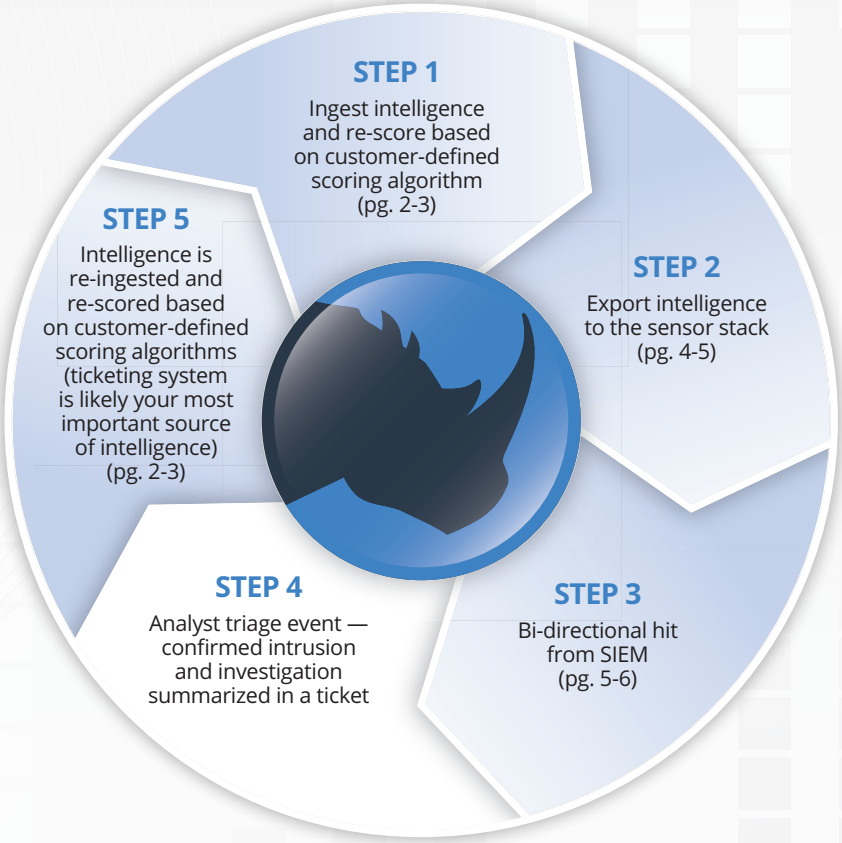

CONCLUSION

Large and small threat intelligence teams benefit by using ThreatQ to apply customer-defined scoring of intelligence, quickly deploy threat data to the existing sensor grid, and cornerstone workflows to focus on TTD/TTR. By deploying ThreatQ to provide structure to cyber intelligence and perform core workflows, organizations can regain the analytical productivity of 6-9 full-time enterprise SOC analysts. The ability to automate the indicator lifecycle, export the highest threat risks to blocking and detecting technologies, and significantly accelerate incident response will improve risk posture, mature team workflows and provide measurable key performance indicators management teams will support.

If you are looking to optimize current processes and get out from under the deluge of threat data, noise and false positives, or trying to build a team but do not have the necessary resources, ThreatQ can empower your team to scale. Through prioritization and automation, the team members you do have will be much more efficient and effective. Let ThreatQ bring immediate ROI to your organization!

ABOUT THREATQUOTIENT™

ThreatQuotient understands that the foundation of intelligence-driven security is people. The company's open and extensible threat intelligence platform, ThreatQ™, empowers defenders to ensure the right threat intelligence is utilized within the right tools, at the right time. Leading global companies are using ThreatQ as the cornerstone of their threat intelligence operations and management system, increasing security effectiveness and efficiency.

For additional information, please visit threatq.com.

TQ_ThreatQ-ROI-Whitepaper_Rev1