THREATQUOTIENT™

Security Operations Platforms: An Assessment of the Economic Benefits of Six Common Use Cases

Threat actors continue to work faster and show greater sophistication in their tactics, techniques, and procedures. The ease with which breaches can be monetized puts companies of all sizes at risk while the attack surface continues to grow as a result of cloud, remote workers, and an increasingly digital supply chain.

The battle against these threats and others continues to wage on where staffing shortages plus siloed organizations and disparate technologies limit security teams' ability to defend against attacks. This ultimately slants the advantage towards threat actors even more.

In this constantly changing landscape, many teams are turning to Security Operations Platforms as a way to combat the challenges they face when protecting their organization from cyber attacks. Investing in a Security Operations Platform is a highly strategic decision. Choosing the right platform for a Security Operation Center (SOC) is arguably more important than choosing any point security product. The Security Operations Platform will become a central part of the security infrastructure, effectively acting as the operating system and data translation layer for all security investments.

Security Operations Platforms support a wide range of use cases and produce a number of economic benefits while helping SOC teams to work more efficiently. This paper aims to quantify those benefits by framing an estimated Return on Investment (ROI) for six common Security Operations Platform use cases including: Spear Phishing, Threat Hunting, Alert Triage, Incident Response,Vulnerability Prioritization, and Threat Intelligence Management.

The following table summarizes the highlights for each use case:

	Spear Phishing	Threat Hunting	Alert Triage	Incident Response	Vulnerability Management	Threat Intelligence Management
Annual Savings Realized with ThreatQ	$279,552	$150,758	$186,318	$228,096	$186,624	$142,128

The calculations are based on industry research, plus experience working with multiple clients. The results speak for themselves; each use case provided enough savings to produce a positive ROI and short payback period after factoring in the cost of a ThreatQ license.

SPEAR PHISHING

Spear phishing is the practice of sending fraudulent emails that targets specific individuals or organizations for the purpose of gaining unauthorized access to confidential information.

Spear phishing emails contain a wealth of hidden evidence that can be used to track and understand the methods used by attackers to target the organization. By extracting that information, analysts can better understand what to look for to identify other users that may have succumbed to the trick.

Armed with this evidence, analysts can discover associations between multiple spear phishing messages to understand a wider campaign that may be underway. Identifying malware samples across campaigns, and associating them with adversary profiles (and therefore intentions) notably improves the ability to respond.

Conducting this level of analysis can be difficult and laborious. Typically, analysts must discover these associations by manually sifting through messages and correlating the information they discover about the campaign with external data on adversaries and their methods.

ThreatQ simplifies and automates the process of parsing and analyzing spear phish emails for prevention and response. With a centralized Threat Library that aggregates all the external threat data organizations subscribe to along with internal threat and event data for context and relevance, analysts are in a position to begin to analyze and determine which emails to focus on.

Recipients of suspicious emails forward the email to an inbox that ThreatQ monitors continuously. Comparing indicators from the email against the data in the Threat Library, ThreatQ determines high risk emails versus low risk, allowing prioritization and noise reduction.

On high-priority items, ThreatQ automatically performs rear-view mirror searches on email logs using SMTP-specific IOCs – email subject, email sender, email filename/attachments. Analysts are able to identify spear phish attacks that might have fallen through the cracks because they were not identified as malicious at the time.

Going a step further, analysts can query to identify all the spear phish recipients and then overlap those findings with vulnerability scan results to determine the scope and help accelerate response and containment.

ThreatQ drives a number of benefits for teams working to protect their organization against spear phishing including:

- Triage spear phish faster and more effectively based on analyst familiarity of adversary TTPs.
- Improved spear phishing attribution.
- Increased understanding of the environment and susceptibility to spear phish attacks.
- Proactive protection against spear phishing attacks.

As a result of these benefits, ThreatQuotient clients have reported considerable efficiency gains related to detection and response when using ThreatQ for spear phishing protection which can be used to estimate a financial return from ThreatQ use:

FTEs Assigned to Phishing Analysis	2
Portion of Functional Tasks Suited for Automation	70%
Expected Efficiency Gains from Automation	80%
FTE Fully Burdened Hourly Rate	$120
Expected Annual Savings	$279,552

FTE = Full Time Equivalent

One ThreatQ client reported time savings of more than 2,300 hours per year related to tasks including parsing and analyzing spear phish emails for prevention and response. Using the fully loaded hourly rate for an analyst responsible for spear phishing protection and the efficiency gains from using ThreatQ, the client was able to realize annual savings of nearly $280,000 from the spear phishing use case alone.

THREAT HUNTING

Threat hunting is the practice of proactively and iteratively searching for abnormal activity within networks and systems for signs of compromise.

Analysts use threat hunting to identify nefarious activity that has not triggered a sensor grid alert as well as potential hopping points an attacker might leverage in the future. While great in theory, there are several challenges to threat hunting. Many security teams don't know where to begin because they lack the ability to prioritize threats for relevance to their environment. Threat hunting also requires specific knowledge and expertise which limits the practice to a few highly skilled analysts. It is also difficult to see the big picture of what is happening across the environment when security teams and tools operate in silos.

When analysts do gain access to what they need, they must quickly find indicators that might reveal adversaries that are staying below the radar – either leveraging Remote Function Call (RFC) protocols or organizational policy thresholds without raising alerts. They also must be skilled at connecting historical attacks with other open source resources to understand an attacker's tactics, techniques and procedures (TTPs) and how they might move laterally when inside the environment. It is extremely time consuming to sift through logs manually to determine which are relevant and to correlate logs with massive volumes of external threat intelligence and other internal data to identify malicious activity. Organizations can end up with a few high-value resources spending inordinate amounts of time potentially chasing ghosts.

The goal of threat hunting is to mitigate the risk once an adversary infiltrates the network. To be effective, threat hunting must start with the threat. The ThreatQ DataLinq Engine includes the ability to centralize and prioritize vast amounts of threat data from external and internal sources so that analysts can automatically determine the highly important items to hunt for within the environment.

ThreatQ TDR Orchestrator enables analysts to automate the prioritization of threat data to support hypothesis development, and the correlation and scoring of internal and external data during the hunt. ThreatQ TDR Orchestrator can also help to automate the response once the team is ready to take action after the hunt.

ThreatQ Investigations allows analysts to conduct investigations collaboratively to search for and compare indicators across infrastructure and find matches between high-risk IOCs and internal log data that indicate possible connections.

Once a match is discovered, analysts can slowly cast the net wider and identify second-tier indicators and attributes (i.e., malware associations, adversary relationships, similar event indicators, etc.). These capabilities enable analysts to engage in threat hunting and follow the prescribed lifecycle, similar to that of any scientific experiment.

ThreatQ drives a number of benefits for teams engaged in threat hunting including:

- Proactively block similar attacks in the future by developing a signature, or identifying new IOCs to detect and block depending on confidence rating.
- Adjust corporate policy to align with new defense rules/signatures.
- Achieve true fusion analysis, leveraging the intelligence and understanding of teams and tools across the organization.

- Develop better intelligence collection methodologies.
- Develop better intelligence practices.
- Find and stop evil before the attack.
- Mitigate risk when an adversary infiltrates infrastructure.
- Orchestrate and synchronize threat intelligence management across all teams and tools so they can work in concert and increase effectiveness, efficiency and productivity.
- Automate the response actions or create an investigation for the team to collaborate.

As a result of these benefits, ThreatQuotient clients have reported considerable efficiency gains related to correlating hypotheses, hunting, and taking action when using ThreatQ for threat hunting which can be used to estimate a financial return from ThreatQ use:

Threat Hunting (based on industry research, TQ client estimates)	**Time Spent Monthly**	
	Without TQ	**With TQ**
Hypothesizing including: Using internal/external data to prioritize threats Visualizing & collaborating on hypotheses	43	17
Hunting including: Correlate logs with internal/external data Import data from frameworks like MITRE ATT&CK to map adversaries to TTP’s	75	15
Acting including: Reporting to detection engineering for signature development Reporting to IR team for remediation	38	19
Totals	156	51
Monthly Savings in Hours		105
Annual Savings in Hours		1,256
Hourly Rate / Analyst		$120
Total Annual Savings		$150,758

Based on industry research and savings reported by ThreatQ clients, an organization with a single headcount focused on Threat Hunting might expect to realize time savings of more than 1,200 hours per year related to tasks including prioritizing threat data, visualizing and collaborating on hypothesis development, correlating data, and reporting findings to other stakeholder groups.. Using the fully loaded hourly rate for an analyst responsible for threat hunting and the efficiency gains from using ThreatQ, yields annual savings of nearly $151,000 from the threat hunting use case alone.

ALERT TRIAGE

Alert triage is the process of efficiently and accurately going through alerts and investigating them to determine the severity of the threat and whether or not the alert should be escalated to incident response.

Analysts are inundated by the number of alerts that require human attention, generated by noisy SIEM rules and default defense infrastructure.

In an attempt to reduce the volume and velocity of security alerts they must tackle on a daily basis, analysts apply external threat data and threat intelligence feeds directly to the SIEM, but challenges continue for two main reasons. First, the amount of external threat data is staggering. Sending all of this data directly to the SIEM for correlation results in an excessive amount of non-contextual alerts, each of which requires significant work by an analyst to research. Second, there is a lack of decision support capabilities in current tools to provide additional context and understanding to determine relevance, before applying threat intelligence feeds directly to the SIEM. Prioritization is imperative to focus and determine the appropriate next actions to take during the alert triage process.

Users can stop the useless alerts before they happen by only feeding threat intelligence that is relevant to the organization into the SIEM for correlation. Machine-to-machine communication allows the SOC analyst to work within their chosen tool, and still have an impact on the continuous tuning of the company's intelligence.

For alerts in the 'gray zone' of importance (med-high, but not high/very high priority), simplify triage with a tool that enables visualization and collaboration. For high priority alerts, include the ability to import alerts into an investigation for visual associations and understanding, and to perform analysis and engage with various collaborators as needed. Additionally, teams can handle routine alerts with automation from enrichment to the response.

In an adjacent workflow within the SIEM where the analyst lives, gain the ability to import highly relevant investigation alert data. Most SIEMs allow three to five additional pieces of context to accompany an indicator of compromise (IOC). Including threat score, IOC source(s), existing ticket numbers + outcome, adversary attribution, etc. will allow an analyst to make very quick and accurate triage decisions.

Learn from and reduce false positives automatically and improve the quality of alerts. If a false positive does slip through, simple feedback can allow for automated tuning of the threat repository. Likewise, the ability to build more accurate SIEM rules based on threat intelligence directly improves the quality of future alerts.

More effectively manage alert triage spikes when a wave of new STIX intelligence hits the sensors. By automatically ingesting the STIX package and transferring that text into the ThreatQ Investigations visualization, an analyst can digest that shared intelligence and take immediate actions.

ThreatQ drives a number of benefits for teams triaging alerts including:

- Improved alert triage process.
- Greater focus – get to the alerts that matter faster, by eliminating the ones that do not.
- Faster investigation response time.
- Better decision making.
- Accelerated resolution – close security alerts more quickly and accurately.
- Instant understanding with visualization of alerts and context (weaponized delivery, malware dropper, CVE, C2, false flags, adversary overlap, shared infrastructure, etc.).

As a result of these benefits, ThreatQuotient clients have reported considerable efficiency gains when using threat intelligence to reduce the raw alert load and improve operations with visualization, collaboration and the automation of context. These improvements to the alert triage function can be used to estimate a financial return from ThreatQ use:

Alert Triage (based on industry research, TQ client estimates)	
Number of alerts per analyst / per month	1,100
Percentage of alerts deemed irrelevant after Filtering with Threat Intel	45%
Alerts Removed from Queue	495
Avg time spent triaging an alert (min)	10
Time savings per analyst / per month (hrs)	83
Remaining Alerts requiring Triage	605
Avg time spent triaging an alert (min)	10
Total time spent triaging (hrs)	101
Percentage of time spent gathering / sharing context	50%
Time spent gathering / sharing context (hrs)	50
Efficiency gains from visualization, collaboration, automation of context gathering/ sharing	93%
Time savings per analyst / per month (hrs)	47
Total Time savings per analyst / per mth (hrs)	129
Annual Savings in Hours	1,553
Hourly Rate / Analyst	$120
Total Annual Savings	$186,318

Based on industry research and savings reported by ThreatQ clients an organization could realize time savings of more than 1500 hours per year related to tasks including filtering alerts for relevancy, as well as improving the visualization, collaboration and automation of context gathering/sharing. Using the fully loaded hourly rate for an analyst responsible for alert triage and the efficiency gains from using ThreatQ, yields annual savings of more than $180,000 from the alert triage use case alone.

INCIDENT RESPONSE

Incident response is an organized approach to the process by which an organization handles the aftermath of a cyberattack or data breach with the goal of limiting damage and reducing recovery time and cost.

Once an event/alert is escalated to an incident, the investigation gains resources and visibility. Additional efforts need to be applied as quickly as possible to understand the scope, impact and the actions required to mitigate damage and recover. Gathering all the required information is a difficult and often manual process, and it comes in a great variety of formats from many different teams and tools.

If an incident under investigation can be related to a known campaign or adversary, the analysis and response time can be drastically reduced, since key tactics, techniques and procedures (TTPs) are already documented providing the proverbial breadcrumbs that lead to hiding places to look. Maintaining adversary profiles and historical incident response reports provides a jumpstart to any incident response investigation. But there is typically no central repository to store, share and update key learnings across teams, and no easy way to work collaboratively to accelerate investigation and response.

ThreatQ is designed to support the fact that incident response is a team sport. Start by importing an event/investigation along with any peripheral intelligence into a shared investigation environment. This instantly allows an incident responder to quickly assess what other research has been performed and by whom, what tasks need to be assigned, and how all the data relates. The ability to include the necessary resources from outside the immediate security department (i.e., database administrators, application specialists, etc.) ensures complete situational understanding and engages the full set of capabilities of the organization. As the necessary responders from around the organization complete tasks and publish them to the larger incident canvas, the team progresses towards identifying patient-0 and re-arming the organization against the next wave of attacks.

If a team knows their attackers' tactics, techniques and procedures (TTPs) , then as that intelligence comes in, they can be scored appropriately and even be added to a "watchlist" for visibility. This is a subtle and proactive way to keep a finger on the pulse of malicious activity. When adversary profiles are frequently updated and maintained with the latest attributes, new analysts can learn about the adversary exponentially faster.

IR teams tend to work within specialized IR platforms. Bi-directional integrations with a security operations platform ensures that the user can focus on their processes and procedures without the need to switch back and forth between multiple interfaces and platforms.

Documenting investigations that can be correlated to future cases, results in organizational memory and ability to correlate investigations that may have seemed to be separate, but are in fact part of a single campaign.

ThreatQ drives a number of benefits for teams handing incident response including:

- Better analysis is performed.
- Faster response time and time to resolution.
- More incidents can be completed.
- Current incident resolution is faster by applying past learnings.
- Better team collaboration and productivity.
- Increased new hire 'time-to-value' (TTV).
- Faster and more complete understanding of how to orchestrate a coordinated response.

As a result of these benefits, ThreatQuotient clients have reported considerable efficiency gains related to identifying, scoping and mitigating intrusions when using ThreatQ for incident response which can be used to estimate a financial return from ThreatQ use:

IR Process (based on TQ Client Estimates)	Without TQ	With TQ	
Identify the intrusion (hours per month)	132	112	
Scoping the intrusion (hours per month)	88	22	
Mitigating the intrusion (hours per month)	132	59	
Totals	352	194	
Monthly Savings in Hours		158	
Annual Savings in Hours		1,901	
Hourly Rate / Analyst		$120	
Total Annual Savings			$228,096

The ThreatQ client reported time savings of nearly 2,000 hours per year related to tasks including creating defensive procedures, maintaining awareness of TTPs, conducting malware analysis, updating signatures, linking activity to known adversaries, stopping intruder access, and monitoring for changes in TTPs. Using the fully loaded hourly rate for an analyst responsible for incident response and the efficiency gains from using ThreatQ, the client was able to realize annual savings of nearly $230,0000 across the team responsible for the incident response use case alone.

VULNERABILITY PRIORITIZATION

Vulnerability management is the practice of continuously discovering, classifying, prioritizing and responding to software, hardware and network vulnerabilities.

Even for mature organizations it is simply impossible to patch and mitigate every vulnerability present in an enterprise network, leading teams to prioritize mitigation based on limited and inward-facing data such as:

- Server versus workstation
- Employee role
- Asset criticality
- Vulnerability score
- Patch availability

Despite this level of prioritization, patching remains one of the most time-consuming vulnerability management tasks. Patching also has limited effectiveness because it does not take into account knowledge of how that vulnerability is actively being exploited in the wild, and the risks associated by those adversaries leveraging it to a company's specific environment.

Since vulnerability is only as bad as the threat exploiting it and the impact on the organization. Security teams must take a data-driven approach to prioritizing vulnerabilities with knowledge about how vulnerabilities are being exploited.

ThreatQ allows security teams to focus their vulnerability management resources where the risk is greatest through the following three steps:

1. Understand the threats and which vulnerabilities threat actors are leveraging to determine relevance to the organization's environment and prioritize which vulnerabilities to address first. For example, a vulnerability related to a specific adversary campaign and IOCs that have been seen in an organization's SIEM and/or ticketing system should be addressed immediately. A vulnerability that has related threats and IOCs but they have not been known to target the organization's specific industry should be watched but is a lower priority. A vulnerability with no known adversaries using it or associated IOCs may indicate it is not being exploited in the real world yet, and can be deprioritized for now.
2. Overlap adversaries that target the company with CVEs the adversaries use, historical victimology targets and vulnerability scan results for those targets to create a superior risk profile.
3. Reassess and re-prioritize on a continuous and ongoing basis as adversaries change tactics, techniques and procedures (TTPs), systems and applications evolve, and their usage within the organization's environment does as well.

ThreatQ drives a number of benefits for teams engaged in vulnerability management including:

- Better situational awareness of attackers, their motivations and one's own environment.
- Clear priorities on what actions to take first to address which vulnerabilities.
- Automate the prioritization of vulnerabilities based in internal/external data and other factors.
- Ability to focus on the vulnerabilities that are the most relevant based on the organization's risk profile.
- A superior risk profile based on deeper insights into adversaries, their tactics, techniques and procedures (TTPs) and relevance to the organization.
- Better investment and resource decisions.

As a result of these benefits, ThreatQuotient clients have reported efficiency gains as high as 90% when using ThreatQ for vulnerability prioritization which can be used to estimate a financial return from ThreatQ use:

Hours/month spent on patching activities	144
Hourly Rate / Analyst	$120
Total Annual Cost	$207,360
TQ Efficiency Gains	90%
Total Annual Savings	$186,624

A ThreatQ client reported the team spending up to 144 hours per month on patching activities related to vulnerability management. Using the fully loaded hourly rate for an analyst responsible for patching and the efficiency gains from using ThreatQ, the client was able to realize annual savings of more than $186,000 from the vulnerability prioritization use case alone.

THREAT INTELLIGENCE MANAGEMENT

Threat intelligence management is the practice of aggregating, analyzing, enriching and de-duplicating internal and external threat data in order to understand threats to your environment.

Analysts are bombarded with millions of threat data points every day from multiple sources in multiple formats. This includes external data from commercial sources, open source, industry and existing security vendors as well as data from internal sources. Each point product within their internal layers of defense, SIEM and other systems within their security infrastructure generates a massive amount of log and event data and alerts. The noise level is deafening.

Analysts need a way to automatically ingest, consolidate, normalize and de-duplicate threat intelligence data in one manageable location. While this external cyber threat data is commonly well-defined and understood, additional context from within the organization can vary wildly between industry verticals and companies. It's vital that the threat intelligence management solution be able to consume and store these different data types as well as provide the capability to tailor data models to fit security teams' needs.

The next step is to prioritize the vast amounts of threat data aggregated in this central repository. However, what is a priority to one company may not be relevant to another. What is needed is the ability for analysts to control how scoring, prioritization and expiration should be done – tell the system what is more important and less important once, and let the system automatically score and re-score when new data and context is learned. As more data comes in, the threat intelligence management system will automatically tune itself, creating a threat library that provides consistent information tailored specifically for the company.

The repository serves as a centralized memory to facilitate future investigations. Security teams can operate from a single source of truth, passively collaborating through the instantaneous sharing of knowledge and using their tools of choice to improve security posture and reduce the window of exposure and breach.

Integration with an ecosystem of data sources is streamlined and cost effective using open APIs at no additional cost, and can be further tailored with an SDK. For broad visibility, the system must be designed to be integrated with all systems that provide or leverage threat data within the organization.

ThreatQ drives a number of benefits for teams engaged in threat intelligence management including:

- Contextualized, relevant intelligence in a database that is customized for the organization's environment and risk profile.
- Focus, noise reduction and decision support during investigations and triage.
- Greater shared understanding of relationships across objects and object types to better support investigations and threat intelligence management.
- The freedom to spend more time performing analysis versus manual tasks.
- Orchestrated and synchronized threat intelligence management across all teams and tools so they can work in concert and increase effectiveness, efficiency and productivity.

As a result of these benefits, ThreatQuotient clients have reported considerable efficiency gains when using ThreatQ for threat intelligence management response which can be used to estimate a financial return from ThreatQ use:

	Without TQ	With TQ	
Hours/month spent on report generation/trending	210	111	
Monthly Savings in Hours		99	
Annual Savings in Hours		1,184	
Hourly Rate / Analyst		$120	
Total Annual Savings			$142,128

A ThreatQ client reported time savings of nearly 1,200 hours per year related to report generation and trending alone. Using the fully loaded hourly rate for an analyst responsible for threat intelligence management and the efficiency gains from using ThreatQ, the client was able to realize annual savings of more than $140,000 from the use case alone.

CONCLUSION

All data is security data because data that provides the context needed to make the best decisions and take the right actions isn't limited to a few tools and feeds, it's everywhere. And harnessing all that data is problematic. No one understands this better than SOC teams battling to work smarter and faster all while facing internal challenges including staffing shortages, siloed organizations and disparate technologies, plus the ever-advancing threat.

Data is a common thread that runs through the six use cases presented in this paper, and most others that security teams face. It's for this reason that a data-driven approach is crucial when selecting a Security Operations Platform.

ThreatQ DataLinq Engine takes a unique approach to make sense of data in order to accelerate detection, investigation and response. The DataLinq Engine starts by enabling data in different formats and languages from different vendors and systems to work together. From there, it focuses on getting the right data to the right systems and teams at the right time to make security operations more data driven, efficient and effective.

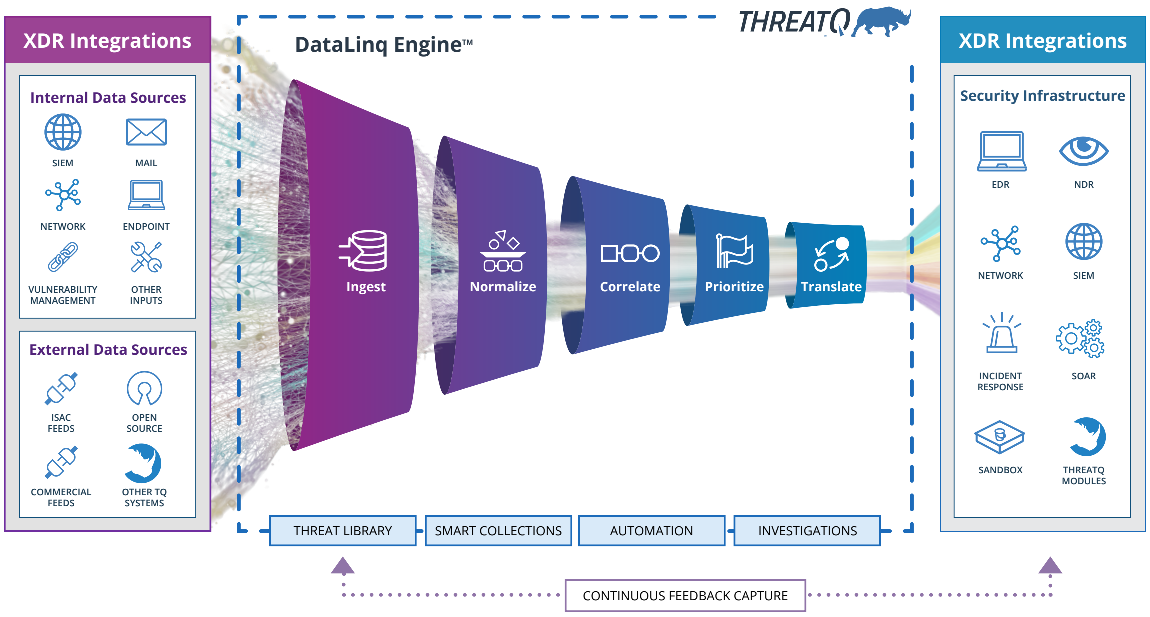

The DataLinq is a foundational technology in the ThreatQuotient product family. To learn more about how to take a strategic approach to using data to accelerate detection, investigation and response, contact us at info@threatq.com for a demo.