

3 Ways to Apply a Risk-Based Approach to Threat Detection, Investigation and Response

Published 16 November 2022 - ID G00719034 - 26 min read

By Analyst(s): Jonathan Nunez, Pete Shoard, Al Price

Initiatives: [Security Operations](#)

Effective threat detection, investigation and response is not simply about buying and implementing security technology. This research will help SRM leaders understand how risk applies to their critical assets and how the impact of a cyber event might affect the organization.

Overview

Key Findings

- The evolution of cyberthreats outpace security teams' ability to implement effective and measurable defenses and security visibility tools. Even with modern security products, it's still incumbent on the organization to determine its risk position as they battle to prioritize resources and security incident remediation.
- Incident investigations often lead to dead ends, as asset information is often missing, including asset ownership and overall relevancy. In combination, these issues make it challenging for responders to determine what to do when they validate a compromise.
- Organizations still widely manage cyber risk in isolation from the rest of the business, which commonly results in a lack of resourcing. The absence of effective risk reporting and awareness can have a direct impact on security budgets and overall security risk management.

Recommendations

Security and risk management (SRM) leaders responsible for implementing threat detection, investigation and response (TDIR) capabilities as a part of security operations (SecOps) should:

- Break through the silos and open dialogue by establishing a quorum of business leaders to openly discuss cybersecurity and its requirements. Allow the business to be part of the conversation and therefore champions of the capability, elevating the security program to a business function rather than an I&O underpinning.
- Reduce unnecessary delays in investigation by ensuring that threat detection use cases are fully enriched with internal business context at the point which alerts are generated. Utilizing this information as a key factor in security alerting can greatly decrease false positives, leading to increased actionability, less time spent on investigations and better utilization of existing staff.
- Enable incident responders to make effective prioritization and response decisions by centrally recording asset-based and business-level risk information. Through the elevation of risk context, responders are able to resolve alerts with a higher probability of business disruption first.

Introduction

While the technology landscape continues to evolve in all directions, security teams continue to battle with resourcing and achieving their desired objectives. There are many factors organizations face which exacerbate threat detection and response accuracy and latency, including but not limited to lack of resources, siloed operations and poor tool implementation. Combined, these issues can slow your ability to detect and respond to threats, potentially missing an opportunity to thwart disrupting impacts.

Through the application of a risk-based approach, organizations can expect to create measurable efficiency gains in threat detection and increase their ability to respond to threats in a timely manner. This method aims to redirect efforts to where they matter most, by focusing on probable points of entry while prioritizing critical assets. Such a process can be defined in four high-level steps (Figure 1):

- Collect and aggregate risk information with business context, adding to your asset inventory.
- Integrate this information into threat detection and investigation technologies and processes.
- Make this information available to security incident responders.
- Utilize the outcomes to deliver accurate feedback to the program.

Figure 1. Risk-Based Threat Detection, Investigation & Response Workflow

Risk-Based Threat Detection, Investigation & Response Workflow



Source: Gartner
719034_C

Gartner.

Analysis

What is risk? At its highest level, risk is the potential for exposure to danger. For most organizations, this can come in many forms, including financial, health-and-safety-based and reputational, among others (see [Standardize Risk Language From The Bottom Up](#)). However, technology has only recently become the binding factor for these risks, as organizations have become increasingly reliant on the internet and automation to produce output.

Business-dependent technologies are a focal point for criminals moving into cyberspace, as anonymity is now a commodity, making the dash for profits an exceedingly easy gain.

Therefore, SecOps must consider and understand business risk and the impact cyber elements have on these risks. However, the question remains: How do these inundated security technologists reduce the noise and achieve their objectives in an environment where time is a limiting factor?

Many academic institutions, security vendors and even government agencies have increased awareness of cyber-risk formulas that contextualize and prioritize decision making, specific to the risks inherent in cyberspace. While there are many versions of these equations, they all seem to share a similar theme: Risk is the sum of your assets, active threats, resident vulnerabilities and potential organizational impact.

Identifying where this data exists within your environment and weaving it into your security program can help it become more agile to the threats and the weakest links that those threats may exploit. Today, most SecOps programs operate as a fixed object, where technologies and processes are created, deployed and infrequently revisited for substantive changes. This is much like a manufacturing line, which is limited in its ability to shift to the demands of the business and the evolution of threat activity. By shifting the focus and leveraging risk information, security programs will naturally mature to an agile posture. Instead of operating like a manufacturing product line, they'll behave more like a lighthouse, directing resources and technologies to the weakest links in the corporate chain (the likely point of entry for an attacker). This, in turn, aims to minimize the amount of damage a threat actor can do. So instead of undertaking a massive task, SRM leaders are relying on their defense-in-depth models to play their part, while allowing SecOps to triage the blind or weak spots first.

Here are Gartner's top three recommendations for enabling risk-based threat detection, investigation and response:

Break Through the Silos and Open the Dialogue

Security resource gaps, including lack of adequate budget and staffing, are often a direct result of ineffective performance and efficacy reporting. Executive decision makers may also not be privy to the inner workings of a security operations center (SOC) As such, reporting remains a key tool for them to make informed decisions, which most often leverage risk as a basis to allocate resources. According to the [SANS 2022 SOC Survey](#), only 39% of respondents attested to a close working relationship with executive leaders for SOC funding allocation, stating a lack of executive urgency to act as a primary factor. ¹ Security programs are often misaligned from the rest of the business, as their function is commonly seen as an overlay to IT operations rather than an integrated, business-supporting capability. To help executives make the most informed decisions, SRM leaders should cultivate relationships with key stakeholders and report effective risk-based metrics, promoting a business-integrated security capability (see: [Tool: Board Briefing – How to Communicate the Cyber-Risk Posture of Your Organization](#)).

Relationships with Key Stakeholders

Get to know key players and decision makers within the overall organization, not just those who are already plugged into the security ecosystem. Many organizations struggle with this principle for a variety of reasons, from work overload to a general unawareness (due to lack of documentation and attrition). These same organizations also find that different parts of the business are actually managing some aspect of security in isolation, potentially duplicating aspects of security spending and producing operational gaps in knowledge and performance. Fostering relationships with a wider cadre of leaders within an organization (see Table 1) can help promote the critical services SecOps provides, creating champions for continued investment and also help source new requirements (see: [Quick Answer: What Are CISOs Using Cyber-Risk Quantification For?](#)). Through building these key relationships and gathering comprehensive requirements, SRM leaders are able to leverage the cyber-risk information elements (CRIEs) to extract all the relevant data points for actionable executive reporting.

Table 1: (Example) Security Operations Requirements Committee

(Enlarged table in Appendix)

<i>(Example) Security Operations Requirements Committee</i>	
<i>Example Business Stakeholders</i> ↓	<i>Basis/Potential Requirement</i> ↓
Chief Executive Officer (or delegate)	General business risks; highest-level view of risk pertaining to the overall business.
CIO	All I&O functions and the protective monitoring thereof.
Chief Financial Officer	Financial loss due to cybercrime such as ransomware and extortion, as well as general business disruption from cyber attacks, resulting in financial loss due to lack of availability or regulatory fines.
Chief Legal Counsel	Regulatory, legal or compliance violations due to cyberattacks (data breach/data leakage). Potential contractual violations due to data breach.
Human Resources Officer	Insider threats.
Head of Marketing	Brand protection.
<i>The stakeholders listed in the above graphic should be considered as an example; organizations should use the context of their business to establish a listing of relevant stakeholders.</i>	

Source: Gartner (November 2022)

Reporting Effective Risk-Based Metrics

Metrics are a mainstay within security programs, from measuring the efficacy of vulnerability management programs to monitoring incident resolution times (see [Outcome-Driven Metrics for Cybersecurity in the Digital Era](#)). Security metrics are essentially measurements collected and reported to identify whether a service-level agreement or objective is being met or violated. These baseline thresholds are typically established by executive management in accordance with their tolerance for risk.

Establishing the correct metrics starts by sourcing the correct requirements and then capturing the relevant telemetry to enable empirical measurements. SecOps service reporting should report risk that ties back to the business requirements, which is therefore business risk, utilizing metrics to quantify the message they’re trying to deliver. However, many organizations are still not capturing the relevant metrics to entice adequate support from business leaders, such as tying monetary value to security interventions, or the inverse. While capturing and measuring effective metrics is a multifaceted security governance function, leveraging cyber-risk information to source relevant telemetry is a healthy way to start improving security reporting (see [Effective Metrics Practices for Cybersecurity Leaders](#)).

When these executive reports are sourced from the requirements of the business stakeholders themselves, SRM leaders are able to articulate their security program outcomes in the language that will resonate with the reader, instead of presuming the reader will understand SecOps outcomes.

A good way to start putting this logic into practice is by establishing a quorum to make decisions and give authoritative requirements from a business leader's standpoint. This quorum or committee should have the authority to opine and make decisions on behalf of the business unit it is representing. It should also have the visibility needed to participate in a roundtable with a selection of peers from around the business; the table above (Table 1) or a delegation thereof may be a good place to start. To ensure viability of this group's success, SRM leaders should approach it bidirectionally, establishing the administrative elements of the committee, as well as the committee's expectations and responsibilities (see Figure 2).

Achieving executive buy-in is critically important when establishing a quorum. This can prevent potential pushback from other leaders that may initially see this as just another meeting to discuss something that is outside of their purview or area of influence. Secondly, identifying the appropriate stakeholders as noted in the table above is almost equally important, as they are the ones that you will solicit the requirements from, and the ones who will evaluate the effectiveness of your program. They are also typically the leaders in direct communication with the CEO, and therefore are in a position of greater influence.

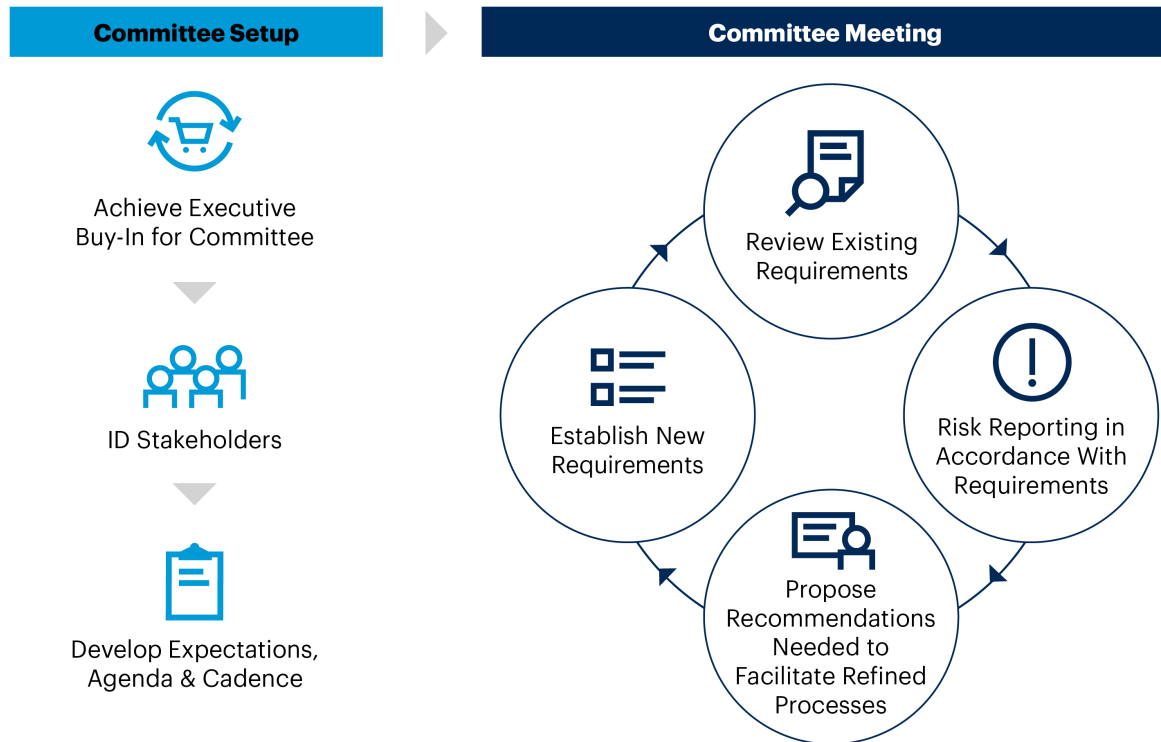
Lastly, establishing expectations and an agenda is of high importance, so the committee understands what's to be discussed and comes prepared to make decisions so topics and issues don't go unaddressed for a lengthy period of time. Issues discussed in this setting are risk-oriented, so unresolved issues would likely have direct organizational impact. Some organizations may choose to codify this collection of administrative activities in a security charter, endorsed by the CEO or a delegate. The charter will help prevent confusion on any of the aforementioned items, and will also demonstrate irrefutable support from executive management.

Once the administrative undertakings have been completed for roundtable discussions, it's incumbent on the SRM leader, or committee chair, to facilitate the meetings in a manner that places the requirements, and reporting against those requirements, front and center. As shown in Figure 2 below, there are four key actions that should not be missed in each gathering:

- **Review existing requirements:** Utilize the roundtable group to go over the list of existing requirements and allow the team to validate/revalidate at each engagement. This is essential, as SRM leaders will actively use the requirements for all applicable SecOps activities under this methodology. If the scope of the business function changes, therefore invalidating the existing requirements, this is an opportunity to flag the issue proactively.
- **Report against requirements:** SRM leaders should use this time to report against these requirements to the stakeholders directly and allow an opportunity for direct feedback. All reporting should have supporting metrics to quantify and validate the progress, or deficiency statements made during the report-out. This is also an opportunity for stakeholders to directly contribute to how security is reporting information back to the business. Adjustments may be needed and this is the group to authoritatively define those changes in collaboration with security.
- **Propose recommendations:** During the course of review and reporting, there will be opportunities for stakeholders and SRM leaders to provide recommendations on how to better position resources to reduce risk. This may be through requirement adjustments and/or additional investments needed to adequately defend against new or expanding threats. This is the forum to have those discussions, as the relevant information is being presented in a manner that is consumable by the decision makers.
- **Gather new requirements:** Gather any new requirements that arise from this committee; this step provides SRM leaders with an opportunity to capture additional or new requirements proactively as the business needs change. As business leaders work on their own expanding priorities, they are also able to leverage this group to discuss how the changing priorities and expanding attack surface can be addressed by SecOps. In turn, this conversation should evolve into capturing new requirements, and the group discussion around resourcing to absorb those new requirements.

Figure 2. Security Operations Requirements Committee (Illustrative)

Security Operations Requirements Committee
Illustrative



Source: Gartner
719034_C

Enrich Risk Information Into Threat Detection Processes

Threat detection is a capability primarily operated by SecOps teams, which typically includes sending IT information (including security appliance information) to a centralized log management repository and analyzing that data for anomalies. Today, it's commonly accomplished through the use of security information and event management (SIEM) systems, although in the most useful sense, it's a combination of log sources that aggregate IT visibility in the SIEM for inspection. One challenge SecOps teams have is a lack of context when receiving security alerts in the SIEM. There are two primary ways to analyze or process this telemetry to produce alerts: signature-based or heuristically, using machine learning. In the former, we think about SIEMs requiring analyst-defined signatures to detect threats; in this sense, typically the only threats detected are what an analyst has prescribed or predetermined as "bad." Conversely, the latter automatically detects deviations to 'normal' network traffic and alerts the analyst when standard deviation or anomaly thresholds trigger a warning (see [Emerging Technologies: Patterns in How Providers Position AI for Security Attack Detection](#)).

However, even with sophisticated detection tools, SecOps teams continue to struggle with knowing what to triage first as they wade through significantly high volumes of data.

To alleviate some of these challenges, SRM leaders should advocate for the integration of CRIEs into threat detection processes.

Cyber risk varies in its measurement; to be effective, organizations must define at least four core areas to measure and collect data: sums of assets, resident vulnerabilities, active threats and organizational impact. Using this understanding, you can start to map out what data sources or information elements are available in your organization for collection, aggregation, correlation and enrichment (see Table 2).

Asset Information

A well-populated and maintained configuration management database (CMDB), as an example, would likely provide an effective asset inventory with associated asset criticality, which can then be used to determine impact. In this sense, one data source can provide data for two variables in the equation. However, many organizations have challenges maintaining such an inventory. In this case, consider using a CMDB in conjunction with other platforms that discover and classify assets. This includes but is not limited to vulnerability assessment (VA) tools, external attack surface management (EASM) platforms or cyber asset attack surface management (CAASM) technologies.

Exposure Information

Resident vulnerability data should likely be sourced from your VA platform or provider. However, for a more comprehensive vulnerability dataset, it's worth considering additional sources (where available) to increase the qualitative value of this category, including but not limited to attack surface management, breach attack simulation, penetration testing and enterprise risk registers.

Active Exploits in the Wild

Lastly, to complete the cyber-risk equation, leverage an understanding of how active exploits in the wild can *impact* your organization. This information can be obtained from open-source, commercial or information-sharing cyberthreat intelligence (CTI) sources. CTI sources can illuminate what threat actors are actively targeting and exploiting; however, when correlated with your organizational vulnerabilities and associated impact, the resulting risk data becomes quantifiable enough to prioritize decision making. To implement effective threat intelligence, Gartner recommends adopting the six-step threat intelligence life cycle, ² which starts with establishing your organizational priority intelligence requirements (PIRs). PIRs allow you to narrow down CTI collection, analysis and production to the threats the business cares about. Skipping this step can have the inverse effect, flooding your security operations team with potential false positives and increased overhead in technological and analytical processing.

Table 2: Examples of Cyber Risk Information Elements

(Enlarged table in Appendix)

Exemplar Data Sources ↓	Example Uses ↓
Information Element: Asset Inventory	
CMDB	Asset information; asset record of authority
CAASM	Asset information; addresses asset gaps in CMDB
EASM/VM	New, dynamically discovered assets
Information Element: Threat Exposure	
EASM/VM	Vulnerability enumeration, increased risk on external attack surface vulnerabilities, supply chain risk indicators
Security Control Validation Tools/Services	Attack path mapping, impact simulation, security control validation
Enterprise Risk Register	Known enterprise IT risks/vulnerabilities
Information Element: Active Exploits in the Wild	
Cyberthreat Intelligence	Indicators of compromise, threat actor tactics/techniques/procedures (TTPs), dark web chatter, exploitation of weaponization, leaked credentials, targeting information, etc.

Source: Gartner (November 2022)

Once the CRIEs have been sourced, it should then be incumbent on the SecOps content creators, commonly referred to as the development, security and operations team (see [12 Things to Get Right for Successful DevSecOps](#)), to synthesize the information for enrichment and correlation. Essentially, they will add all of the known information about an asset into a machine-readable format. This amalgamation of contextualized asset data can then be leveraged to curate customized, risk-based threat detection logic for the SIEM and for threat hunting purposes. Ultimately, the idea is to gather as much information about your digital assets as possible, including threat information, and use that collective intelligence to inform threat detection operations. This information should enable your entire security program to agilely respond to the threat landscape, actively providing tangible, measurable risk reduction.

Using the [cybersecurity mesh architecture](#) (CSMA) as a high-level reference, obtain the CRIEs as an extract of the schematic's security intelligence layer and use them to enrich threat detection processes. Protective monitoring practices have long been seen as fixed objects, much like a manufacturing production line: data goes in and "fix" recommendations come out. However, that system is no longer able to keep pace with the evolving threatscape. A modern risk-based methodology can help SOC's transform their responses to these threats by shifting their focus from what's happening externally to specific protection of their enterprises in the areas of greatest need. In essence, SOC's will redirect resources to the weakest and most important links, ensuring business continuity. While transformational goals should be aligned toward the CSMA, this document can serve as a guide for how to start reshaping how you detect and respond to threats.

According to [Mandiant's 2022 M-Trends Special Report](#), it took organizations an average of 21 days to detect a successful cyberattack. To underscore the security impact of this statistic, CrowdStrike's Falcon OverWatch team stated that it took threat actors an average of 92 minutes to move laterally across a compromised network.³ So how do the CRIEs translate to better threat detection? Let's take a look at a typical SOC analyst workflow. In a traditional SOC construct, there are three tiers of analysts: Level 1, Level 2 and Level 3.

Level 1 Analysts

These analysts, known as L1s, are typically responsible for making an initial false-positive determination; they are the first line of analysts to review the alerts and make an impression on the findings. Their key decisions boil down to a couple of factors, (1) how much enrichment is included in the finding (i.e., threat intelligence confidence, log source correlations, use case rule logic, etc.) and (2) how much they understand about the “victim” asset. In this instance, these analysts spend much of their time assessing whether they can trust the alert, and trying to determine what the affected asset is. In many cases, the asset information is totally missing from the threat alert process, leading to an inordinate amount of time spent qualifying the other factors in the absence of this information while still trying to gain any information on the asset. Analysts may even use manual methods, such as performing reverse network lookups, emailing or escalating tickets to the I&O teams, or in the worst cases, skipping over the asset information entirely.

Cyber risk information elements help SOC analysts triage the most important findings first, saving time and potentially minimizing organizational impact.

Level 2 Analysts

These analysts, or L2s, are charged with investigating the findings escalated by L1s to further validate the initial assessment. The typical questions they try to answer are:

- Are there any other data sources to help further validate the event as a threat?
- Is this activity normal or expected behavior?

First, it is difficult to comprehensively identify all available log sources for an asset with no associated context. Yes, analysts can and often do refer to the standard security appliances they know they have, to “check their traps” (web proxy, firewall, endpoint detection and response, intrusion prevention and detection systems, etc.). However, without understanding what the asset is and how important it is to the business, they may miss some of the most critical sources for validation.

Consider a SIEM alert or notification of a known bad IP address establishing a TCP connection with an enterprise IT asset. Say that the asset is actually an application server and the attack vector is on the application. Without immediately having the asset information available, the analyst would not know to reach out to the application team for the corresponding application logs where the attack would have been evident. Additionally, CRIE enrichment would have made the application's criticality evident, garnering immediate attention versus sitting idle while the analyst works in chronological order.

The other question that Level 2 analysts try to answer is whether or not the activity noted in the finding is normal or expected behavior. It can be exceedingly difficult for SOC analysts to understand what is normal or expected for enterprise traffic, especially when considering an individual asset or cluster of assets in isolation. Yes, they can look at historical patterns in network traffic, but that is a manual, time-consuming effort and still doesn't account for unexpected traffic deviations (such as the introduction of new assets, maintenance windows, new business processes, etc.).

The evolution of machine learning in threat detection technologies has helped in this regard. But even so, machine learning requires time to learn behavior patterns, and lack of context can incorrectly train the algorithms, incorrectly influencing the data science and resulting in false positives that require systematic tuning. Comprehensive CRIE enrichments can alleviate some of this burden. Analysts reviewing the alerts can quickly determine what may be normal or expected traffic for an asset, or even help to determine whether the exploit was, in fact, correctly targeting the asset.

Consider the application from the previous example. If the CRIE data included vulnerability information for the asset, the analyst may have been able to quickly determine if there was a likelihood of compromise, even before sourcing the application logs from the I&O team.

When time is a limiting factor, information can be an accelerator, even when used as a confidence builder for faster decision making.

Level 3 Analysts

These analysts, also known as L3s, are responsible for reviewing the investigation findings and qualifying the true-positive nature of the incident, validating the event categorization, confirming the incident severity level and escalating to the incident response team for containment and resolution. They require high-fidelity information to refer the event to the incident response (IR) team. Without asset context, the alert from the previous example would have just been a TCP connection from a known bad IP address on an enterprise asset, with no additional information. The SOC would have likely recommended, through due diligence, an escalation because of the established TCP connection. However, none of the additional information would have likely been collected, at least not without significant delays to the L3 for qualification.

The L3s also have the responsibility of establishing and validating incident severity, which helps prioritize the IR team's workload, as well as establish the appropriate metrics for executive reporting and, in some cases, compliance requirements. Not having CRIE enrichment can, and often does, adversely affect SOC analysts and ultimately, the organization. If analysts do not know that the asset is hosting a vulnerable, business-critical application, this will affect how they escalate the event and assess its severity. Collectively, this can cause a snowball effect, as the IR team will likely not prioritize the incident in accordance with its significance, as per the L3's recommendation.

Through the investigative process, SOC analysts can (at times) source the relevant information about an asset and its risk. However, the time it takes to get the information is typically longer than the business would tolerate for their most critical assets, and that is the underlying issue.

Enriching threat detection from the outset is the best use of SOC resources as effective decisions can be made faster, protecting the enterprise from significant outages and related costs.

There is also an unintended consequence to not enriching with CRIE: improper metric measurements. Some metrics have become commonplace within modern SOCs, such as mean time to detect (MTTD), mean time to contain (MTTC) and mean time to resolve (MTTR). Service-level objectives and/or agreements (SLO/SLAs) are measured against these metrics for the identification of adherence or violations to policy.

Without asset context readily available, the probability of SLO/SLA violations can increase exponentially. As an example, an organization with an MTTR of 24 hours for a critical asset may very well violate their SLAs if it isn't known that the asset is critical. In this case, the IR team may not receive the incident until Day 2, 3 or longer. The same can be said for MTTD and MTTC; essentially, if it's not immediately evident, the alert will sit in the queue waiting for its turn to be worked, and the performance measurements will follow suit.

Many organizations today leverage a variation of hybrid SOC deployment models, where the provider is responsible for a portion or all of the detection, investigation or response. However, they should still be held accountable for reporting these metrics and ensuring adherence to the organizationally defined SLAs; the impetus remains the same regarding the utilization of CRIs. (For more information, see [Quick Answer: What Key Questions Should I Ask When Selecting an MDR Provider?](#)).

Use Risk-Based Prioritization for Faster Incident Response

Once the incident responders receive the escalation from the SOC (L3s), they're typically charged with establishing or validating infection boundaries, identifying the root cause of the infection and offering containment and remediation actions. The absence of CRIE information creates challenges in the evidence escalated with the incident, effectively leaving the "heavy lifting" to the IR team, which can cause delays in resolution. In the previous example, the IR team may have to reach out to various parts of the business to:

- Identify the asset/application owners to source the appropriate logs
- Obtain network diagrams to see what the asset is connected to
- Understand the business logic of the application and the underlying data flows to investigate the probabilities of cross-proliferation/lateral movement

While some of these activities are reserved for the IR investigative function, the question around time delay remains, especially considering many organizations leverage third parties for hybrid or cloud-first hosting strategies. Consider that the investigator has to manually identify the asset, asset owner and custodian, only to realize that the application is actually hosted by a cloud service provider (CSP). Now also consider that the IR team has to open a ticket with the CSP, who has a 48-hour SLA to provide the requested logs to the investigator. That 24-hour MTTR SLA for the critical application is long gone, and the IR team hasn't completed or, worse, started their investigation.

According to [IBM's Cost of a Data Breach Report 2022](#), it took organizations an average of 70 days to contain a breach. ⁴ While CRIE enrichments cannot, in isolation, alleviate the burden of reducing this MTTR entirely, it can significantly reduce some of the lag time in containing and resolving incidents. With this information, the IR team will have received appropriately categorized escalations and severity information to work from, as well as a good starting point in terms of asset context. It is beneficial to understand what vulnerabilities or risks are associated with the asset in question. It helps to determine not only the investigative approach, but also the containment and remediation actions that will best harden the affected assets and prevent reinfection. In the example above, if the CSPs patching SLAs are not fast enough to remediate this critical asset, the IR team may recommend virtual patching from the web application firewall in front of the affected application as a mitigating control.

Understanding which active threats are targeting the vulnerabilities of the affected asset can also greatly assist in expediting the investigation. By understanding the threat actor's TTPs for the exploit targeting the vulnerability, responders can gear their investigation toward those specific actions. This can potentially identify the malicious activity faster than regressively searching across the entire enterprise network for any artifacts of anomalous activity. CRIEs should not replace IR due diligence. In this context, it's simply meant to expedite the identification of the infection boundary by narrowing down the TTPs used by the threat actors. To do this effectively, the responders need to understand what the asset is, what vulnerabilities reside on the asset, and sufficient and current intelligence on who's exploiting these vulnerabilities and how.

Incident responders should have as much information at their disposal as needed to be effective in finding a needle in a haystack.

IR teams are working through multiple incidents, most often on a daily basis. They can often be found working long hours into the night and on weekends, trying to make containment decisions based on what they know. CRIE information, while only one piece of the puzzle, can help in answering some of the more difficult questions in a response, leading to faster containment and remediation times.

Evidence

¹ [2022 SANS SOC Survey](#), Sans Institute.

² [How Gartner Defines Threat Intelligence.](#)

³ [2021 CrowdStrike Global Security Attitude Survey, CrowdStrike.](#)

⁴ [Cost of a Data Breach Report 2022, IBM.](#)

[M-Trends 2022: Mandiant Special Report, Mandiant.](#)

Recommended by the Authors

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

[IT Score for Security and Risk Management](#)

[How to Use MITRE ATT&CK to Improve Threat Detection Capabilities](#)

[Use Adversary-Generated Threat Intelligence to Improve Threat Detection and Response](#)

[Security Program Management 101 – How to Select Your Security Frameworks, Controls and Processes](#)

© 2022 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)."

Table 1: (Example) Security Operations Requirements Committee

<i>(Example) Security Operations Requirements Committee</i>	
<i>Example Business Stakeholders</i> ↓	<i>Basis/Potential Requirement</i> ↓
Chief Executive Officer (or delegate)	General business risks; highest-level view of risk pertaining to the overall business.
CIO	All I&O functions and the protective monitoring thereof.
Chief Financial Officer	Financial loss due to cybercrime such as ransomware and extortion, as well as general business disruption from cyber attacks, resulting in financial loss due to lack of availability or regulatory fines.
Chief Legal Counsel	Regulatory, legal or compliance violations due to cyberattacks (data breach/data leakage). Potential contractual violations due to data breach.
Human Resources Officer	Insider threats.
Head of Marketing	Brand protection.
<i>The stakeholders listed in the above graphic should be considered as an example; organizations should use the context of their business to establish a listing of relevant stakeholders.</i>	

Source: Gartner (November 2022)

Table 2: Examples of Cyber Risk Information Elements

<i>Exemplar Data Sources</i> ↓	<i>Example Uses</i> ↓
Information Element: Asset Inventory	
CMDB	Asset information; asset record of authority
CAASM	Asset information; addresses asset gaps in CMDB
EASM/VM	New, dynamically discovered assets
Information Element: Threat Exposure	
EASM/VM	Vulnerability enumeration, increased risk on external attack surface vulnerabilities, supply chain risk indicators
Security Control Validation Tools/Services	Attack path mapping, impact simulation, security control validation
Enterprise Risk Register	Known enterprise IT risks/vulnerabilities
Information Element: Active Exploits in the Wild	
Cyberthreat Intelligence	Indicators of compromise, threat actor tactics/techniques/procedures (TTPs), dark web chatter, exploitation of weaponization, leaked credentials, targeting information, etc.

Source: Gartner (November 2022)