MTTR White Paper

When Every Second Counts

Measuring Meaningful Mean Time To Reduce Costs of a Data Breach

Guest Author:
Brad LaPorte, Lionfish Tech Advisors

ULTRARED

Table Of Contents

Introduction

“Mean Time to Remediate” (MTTR) is a metric cybersecurity teams use to track the amount of time it takes an organization to remediate a vulnerability after it is discovered. The longer a Security team takes to remediate a vulnerability, the more leverage they give to threat actors to discover and exploit that vulnerability. Given the rising global incidence of cyberattacks, Security teams are paying closer attention to their MTTR.

On average, the MTTR for most organizations is between 2–5 months, which is significantly worrying given that most organizations have a backlog of over 100k vulnerabilities. According to IBM and Ponemon, the increase is staggering, with data breach costs increasing 15% over the past 3 years, with 2023 being the highest average total cost in the history of their Cost of a Data Breach Report. The report also stated that costs were significantly lower for some organizations with a more mature cybersecurity posture and higher for organizations that lagged in areas such as security AI, automation, and cloud security. Those savings averaged $1.76 million or 40%. 277 days is the average number of days to identify and contain a data breach. The longer it took to identify and contain, the more costly the breach. Ransomware breaches were the hardest to detect, taking 49 days longer and costing 10% more or $4.54 million. Ransomware attacks caused 8% of analyzed breaches in the report. These costs included escalation, notification, lost business, and response costs but did not include the cost of the ransom payment or the cost of cyber insurance.

Data breach costs accrue over several years. The report found that, on average, 52% of data breach costs were incurred in the first year, 29% in the second year, and 19% more than 2 years after the event. Organizations in highly regulated industries, such as healthcare organizations and financial services, suffered the worst long-tail costs, with the cost of a breach rising in the second and third years compared to low-regulated industries. High data protection regulatory environments incurred 45% of breach costs in the first year, 31% in the second year, and 24% more than 2 years after a breach.

Supply chain breaches took 26 days longer to detect, with the significant impact of threat detection time on data breach damage costs highlighting the importance of efficient risk remediation planning. Data breaches in the United States continue to be

vastly more expensive than in other countries, with an average total of $9.44 million (more than double the global average). Although relatively rare at 6% of occurrences, attacks initiated by malicious insiders were the costliest, at an average of USD 4.90 million, which is 9.6% higher than the global average cost of USD 4.45 million per data breach. It took nearly 11 months (328 days) to identify and contain data breaches resulting from stolen or compromised credentials, on average, and about 10 months (308 days) to resolve breaches initiated by a malicious insider.

In summary, an organization's ability to identify, detect, and remediate issues quickly can mean the difference in millions of incurred damages. The problems of prolonged MTTR time and vulnerability backlog stem from various root causes: siloed cybersecurity tools, compounding cybersecurity data sprawl, fractured cross-functional collaboration, poor or lacking threat/vulnerability intelligence, insufficient cybersecurity budget, and workforce constraints. Until these root causes are resolved, cybersecurity posture will continuously weaken and potentially fall prey to threat actors.

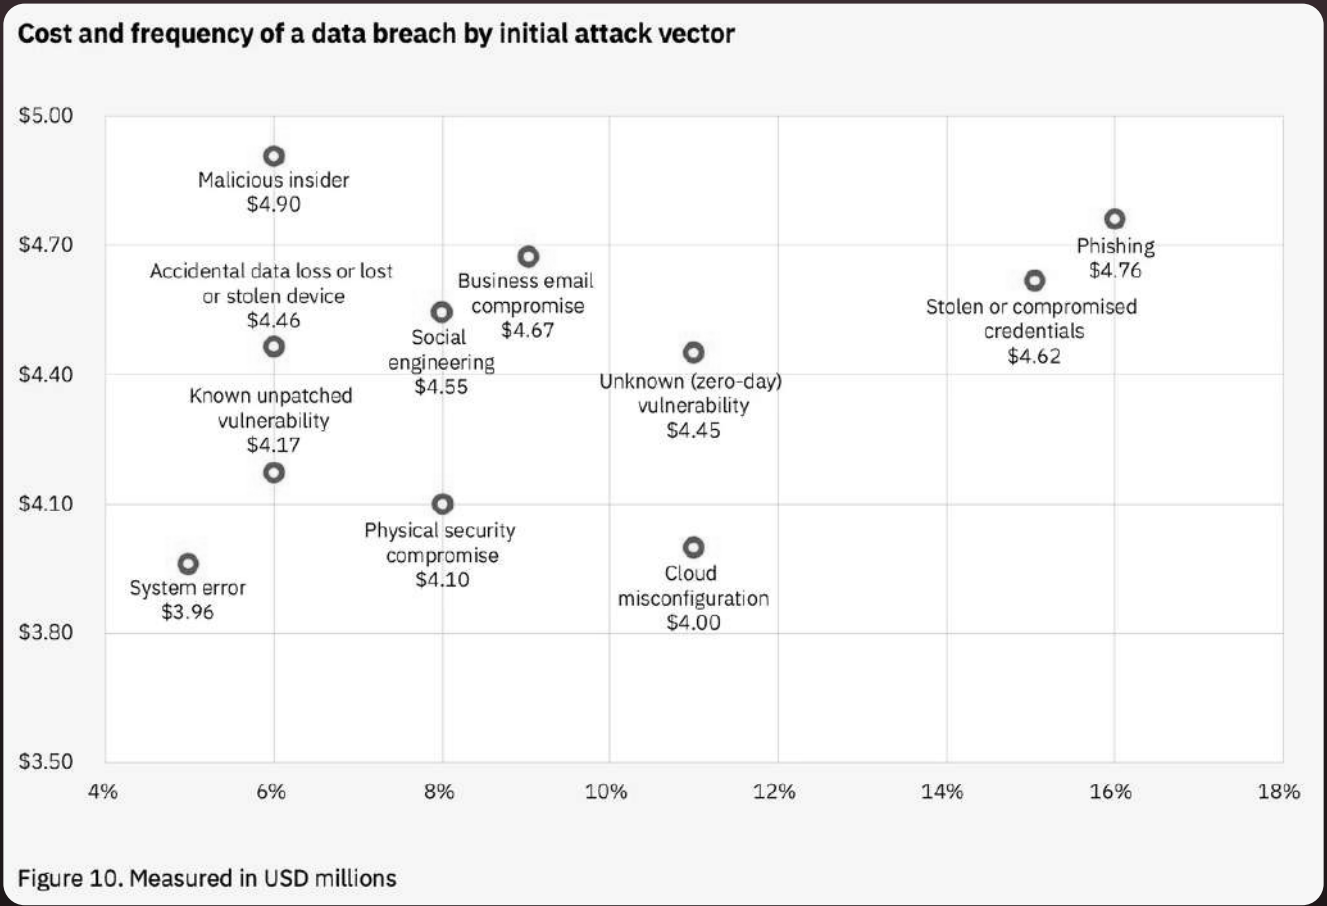

Figure 10. Measured in USD millions

Source: IBM and Ponemon

Understanding the Concept of Mean Time to Recover (MTTR)

MTTR is a critical metric in cybersecurity, encompassing Mean Time to Detect (MTTD), Mean Time to Acknowledge (MTTA), Mean Time to Repair, and Mean Time to Resolve. A reduced MTTR indicates more efficient incident response and recovery, which can have profound impacts on SOCs, particularly regarding their structure and level of effectiveness.

Incident management consists of four stages:

1. Identification (Mean Time To Detect—MTTD and Mean Time To Acknowledge—MTTA)
2. Containment (Mean Time To Respond, Mean Time To Repair)
3. Remediation (Mean Time To Remediation)
4. Continuous Improvement (the interval between incidents—Mean Time Between Incidents (MTBI) and Mean Time To Remediate (MTTR)).

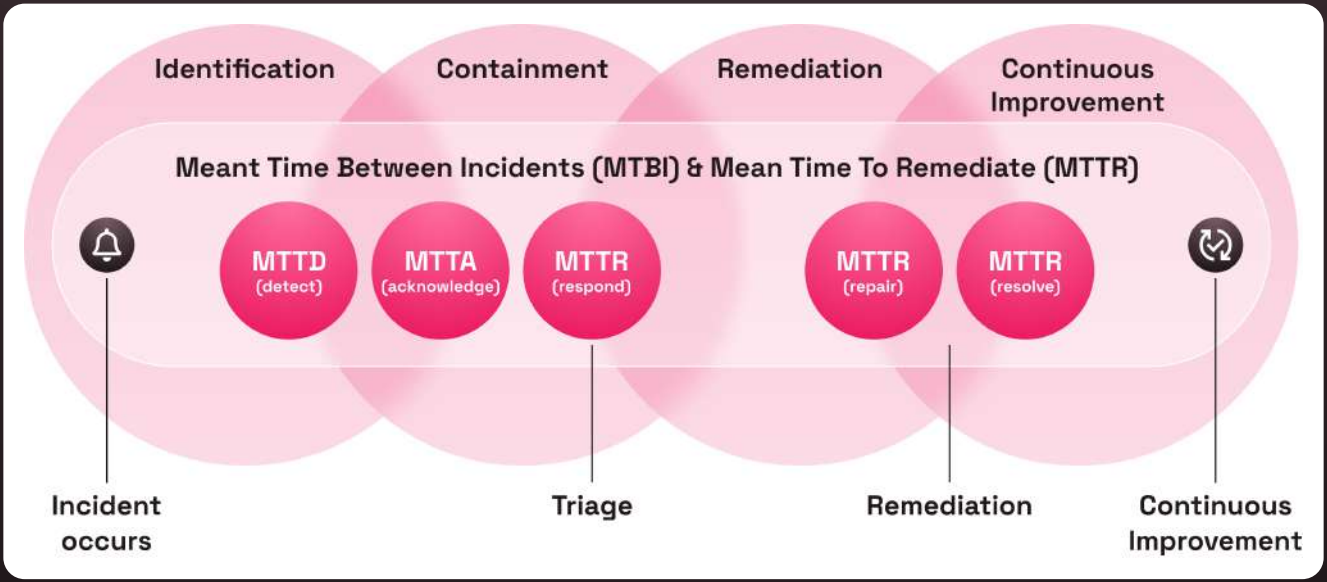

Evolving From RBVM to CTEM

To help guide efficient vulnerability management and reduce MTTR, most organizations follow a "risk-based vulnerability management" approach. According to "risk-based vulnerability management" (RBVM), organizations are encouraged to understand the risks posed by vulnerabilities inherent across their asset estate and to prioritize remediation according to those risks. To better understand where RBVM is today in most organizations, we must first define risk.

Risk = (Threat x Vulnerability x Probability of Occurrence x Impact)/Controls in place[1]

Most vulnerability management tools and surrounding processes baseline the first half of the risk score (threat x vulnerability x probability of occurrence x impact) on vulnerabilities logged in CVE databases and ranked by CVSS scores. Further, they prioritize vulnerabilities on these scores. CVE databases and CVSS scores alone; however, do not fully encompass the definition of risk above. What is largely missing is adequate threat scoring, the probability of occurrence, and control testing to understand the resulting potential of impact. Rarely do vulnerability management practitioners (mostly embedded in Blue Teams) track threats and test the integrity of cybersecurity controls to holistically calculate risk because these tasks are often assigned to internal pen-testers and offensively-focused cybersecurity practitioners (Red Teams). Although it might seem like a simple fix to integrate the two teams and surrounding technology to mobilize change; it has remained challenging for many.

Fortunately, a new program was recently proposed by Gartner, Inc. to help resolve these issues, and it's called "Continuous Threat Exposure Management" (CTEM). CTEM will, as a result of the several proposed changes, drastically reduce MTTR.

"CTEM is a set of processes and capabilities that allow enterprises to continually and consistently evaluate the accessibility, exposure, and exploitability of an enterprise's digital and physical assets" (Gartner). The program follows 5 steps: Scope, Discover, Prioritize, Validate, Mobilize.

[1] NIST 800-30 2002

The CTEM framework clarifies that to reduce MTTR, organizations must first understand their asset landscape, attack surface, and how exploitable their organization is to the most current threat actors. Only then can they properly prioritize vulnerabilities for threat and risk reduction. This understanding is often supported by technology. However, before technology can be brought into question, several programmatic hurdles exist that should be resolved before significant technology investments are made. The programmatic hurdles that exist are delayed vulnerability remediation due to fractured relations between the cybersecurity team, IT team, and business asset owners. IT incurs difficulty in agreeing on the necessary remediation actions as assessed by cybersecurity, justifying the downtime that remediation may cause, and adjusting Service Level Agreement (SLA) times for timely remediation. If there were sufficient evidence to prove that threats are imminent and remediations are necessary, this would help IT teams significantly. Technology would then be an added benefit to streamline cross-functional collaboration and risk reduction.

The range of technology involved in CTEM includes asset discovery, vulnerability scanning, attack surface management, vulnerability prioritization, breach and attack simulation, threat and vulnerability intelligence aggregators, IT Service Management, GRC/IRM, and content collaboration tools.

Although all of these tools exist, it is difficult to find a good number of these tools with added feature benefits existing either natively or, by the majority, fully integrated into one platform.

Source: Gartner

If organizations adopt a CTEM program and use technology that befits the "platform approach"— all in equal parts and working harmoniously together, then organizations will definitely see a reduction in MTTR.

The Current State of SOCs Has Room to Change

Traditionally, SOCs operate on a three-tier model and are covered extensively by NIST's Cybersecurity Framework. However, this model often results in siloed operations and delayed response times. Re-leveling SOCs to a single level promises a more streamlined, faster, and more effective approach to handling security incidents.

The roles of SOC personnel typically break into tiers according to their involvement in an incident's timeline and severity.

The common roles and responsibilities for a SOC team are:

- Security Analyst (Tier One) – Responsible for vulnerabilities monitoring, triaging identified incidents, and escalating those that warrant it.
- Security Analyst (Tier Two) – In charge of investigating and responding to incidents, then executing response and recovery processes to remediate incidents' impact.
- Threat Hunters (Tier Three) – Responsible for assessing IT security infrastructure according to the latest threat intelligence to determine unexpected or stealthy means of network entry.
- Manager (Tier Four) – Responsible for overseeing the entire team and reporting findings, action plans, and threat notifications to the organization's CISO.
- Engineer/Architect – Works alongside other members on SOC teams, designing, developing, and maintaining security infrastructure.

The Journey Towards a One-Level SOC with ULTRA RED

ULTRA RED envisions a future where SOCs are streamlined into a single level, enhancing their ability to respond to and manage security threats. By focusing on reducing MTTR through measures like EASM, data enrichment, and Embedded Security Validation (ESV), ULTRA RED aids in the transformation towards this one-level SOC. This shift promises greater efficiency, faster response times, and enhanced security.

The Role of External Attack Surface Management in Reducing MTTR

External Attack Surface Management (EASM) involves identifying, managing, and securing all external digital assets accessible from the open internet. By employing EASM, red teams can more quickly detect threats, thereby improving MTTD. This early detection can significantly reduce MTTR, enhancing overall cybersecurity efforts.

The Importance of Data Enrichment in Acknowledging Threats

Data enrichment, the process of enhancing, refining, and improving raw data, plays a crucial role in improving MTTA. By providing more detailed, accurate, and actionable information, data enrichment enables red teams to more rapidly acknowledge threats, contributing to a reduced MTTR.

Figure 47. ASM helped accelerate total time to identify and contain a data breach by nearly 12 weeks.
Without an ASM solution, organizations took 260 days to identify a data breach and another 77 days to contain it, for a total of 337 days or about 11 months. Organizations with an ASM solution identified the breach in 193 days and contained it in 61 days. The 254-day total time to identify and contain a breach represented an acceleration of 83 days or about 12 weeks so the data breaches were identified and contained in 75% of the time taken by data breaches at organizations without ASM solutions.

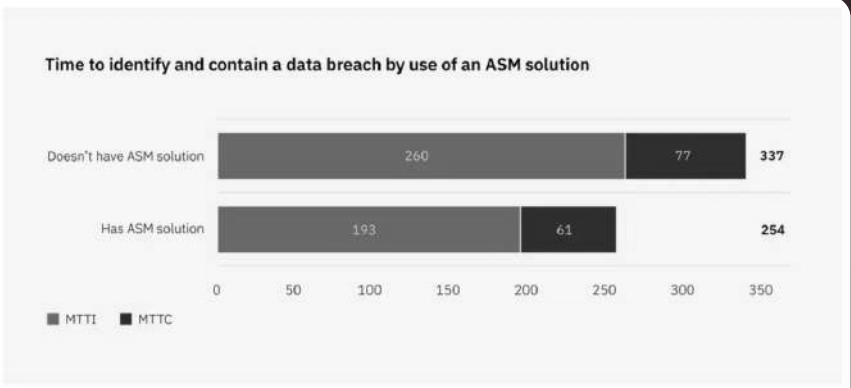

Figure 47. Measured in days

Source: IBM and Ponemon

ASM is a set of processes that aids in the discovery, analysis, remediation, and monitoring of an organization's potential attack surfaces or vulnerabilities. Organizations that deployed an ASM solution were able to identify and contain data breaches 75% of the time of those without an ASM solution.

ASM helped accelerate the total time to identify and contain a data breach by nearly 12 weeks. Without an ASM solution, organizations took 260 days to identify a data breach and another 77 days to contain it, for a total of 337 days or about 11 months. Organizations with an ASM solution identified the breach in 193 days and contained it in 61 days. The 254–day total time to identify and contain a breach represented an acceleration of 83 days or about 12 weeks, so the data breaches were identified and contained in 75% of the time taken by data breaches at organizations without ASM solutions.

Embedding Security Validation for Faster Remediations and Resolutions

Embedded Security Validation (ESV) is a proactive cybersecurity approach that incorporates security checks and balances throughout the system to help identify and address, and mitigate security issues quickly and with greater accuracy. When a security incident occurs, embedded validation tools can swiftly assess the threat, validate its legitimacy, and provide actionable insights for remediation. This streamlined approach reduces the time required to identify and repair security vulnerabilities or breaches, thus lowering MTTR. Embedded security validation acts as a proactive and reactive safeguard, enhancing an organization's cyber resilience while significantly improving its incident response capabilities.

Introducing ULTRA RED

The ULTRA RED Threat Exposure Management Platform was purpose-built to address key aspects and weak points within vulnerability and threat exposure management to, among a multitude of other benefits, reduce MTTR and vulnerability backlog. Following the principles of CTEM, we have created an all-in-one, fully integrated platform from which all stakeholders in vulnerability and threat exposure management can convene, collaborate, and remediate their most pressing risks and threat exposure. Below, we pair our platform capabilities with the phases of CTEM to paint the full platform story.

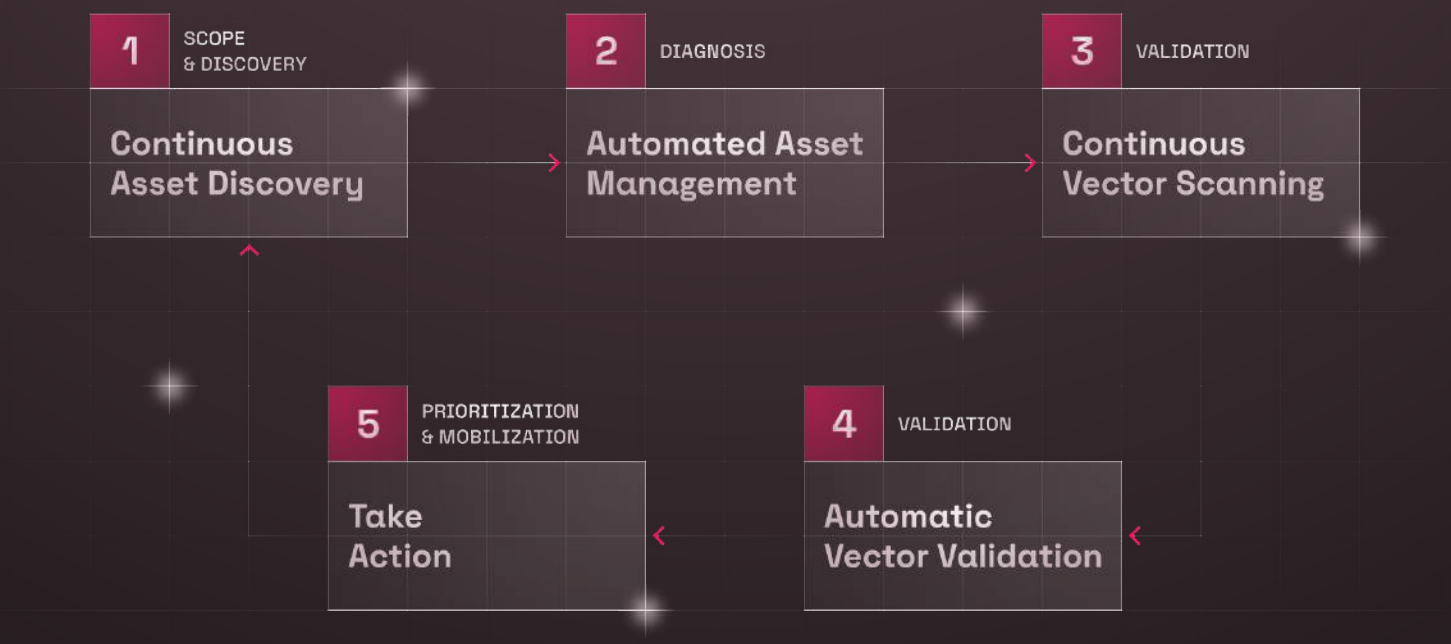

SCOPING & Attack Surface Management (ASM)

Attack surface defense starts with visibility. After all, you can't protect what you don't know about. Manually maintaining an asset register of internet-facing assets and services is impractical, if not impossible, for most organizations. Asset audits are resource-intensive and time-consuming to conduct, and the result is a point-in-time snapshot that is both incomplete and obsolete as soon as it is created.

Blind spots resulting from shadow IT deployments, cloud adoption, and mergers and acquisitions can result in a large pool of organization-owned assets that are unknown to security teams. These unknown assets quickly become the soft underbelly of the attack surface, harboring vulnerabilities that would normally be identified and remediated. Better knowledge to this end can help Red Teams in their adversary-emulation exercises, Blue Teams in fortifying defenses, IT Management Teams in knowing what asset landscape they should keep a keen eye on, and GRC teams in communicating more accurate threats and risks to Security Leadership and Board.

In as little as **3 clicks**, ULTRA RED automatically discovers hosts, domains, and IP addresses that are within the organization's scope, including assets created by shadow IT activities and cloud deployments. This is an “action once and automate indefinitely” functionality which assures that the attack surface is always mapped for changes and threats. Additionally, ULTRA RED utilizes recursive techniques and proprietary technology to check and re-check validity which yields an asset register with an extremely low false positive rate. ULTRA RED's asset discovery is a continual process, uncovering new assets and services as they are introduced to ensure you always have an up-to-date view of your attack surface.

SCOPING/DISCOVERING & Digital Risk Protection Services (DRPS)

Is your company primed to be compromised? DRPS answers this question by bringing to light all exposed identities and credentials in the open web, social media, dark web, and deep web sources. Intelligence also includes potential threat actors, their tactics, and processes used to initiate compromise. DRPS will immediately show you all your exposed digital assets up for grab and potentially already in play by cyber attackers.

VALIDATING & Cyber Threat Intelligence (CTI)

With enhanced CTI integrations in the Ultra Red Threat Exposure Management platform, Red Teams are armed with enough to automate potential cyber attacker hacks. Integrations and enrichments include intelligence like IOCs, potentially compromised accounts, leaked credentials, and enhanced threat actor intelligence. ULTRA RED combines CTI with vulnerability intelligence, cross references assets against a wide range of proprietary scanning findings and documented Dark Net references, and detects weaknesses without impacting the scanned system and its security

infrastructure and controls, ensuring production system continuity. All of ULTRA RED's intelligence is actionable, given in-platform, actionable, and prioritized remediation lists.

CTEM & Vulnerability Management (VM)

The average external attack surface contains a range of assets and services with vulnerabilities of varying severity, and it is impractical to address them all. While zero-day vulnerabilities are a threat to be considered, most breaches are the result of exploiting known vulnerabilities that organizations could have remediated. In fact, malicious actors focus their efforts on recently announced vulnerabilities, scanning the Internet for vulnerable systems to exploit.

Manual efforts, either in-house or through bug-bounty programs, are resource-intensive activities that often serve up a disproportionate number of low and mid-severity vulnerabilities. The real challenge is to identify and focus on those vulnerabilities that represent the greatest risk to the organization. This can best be done using automated methods against a complete asset register. Better knowledge to this end can help Red Teams in their adversary-emulation exercises, Blue Teams in fortifying defenses, IT Management Teams in knowing what to patch, where, and why, and GRC teams in communicating more accurate threats and risks to Security Leadership and Board. Relying on CVSS scores and other external indicators to prioritize vulnerabilities has its limitations because every attack surface is different. Some high-risk vulnerabilities will have mitigating controls in place and others may not have the required prerequisites for successful exploitation.

ULTRA RED continuously scans assets in the Asset Inventory to detect new vulnerabilities and weaknesses. It not only identifies known public vulnerabilities but also cross-references assets against a wide range of proprietary scanning findings and documented Dark Net references. All detected weaknesses are validated without impacting the scanned system and its security infrastructure and controls, ensuring production system continuity. ULTRA RED's vector scoring is based on a combination of the CVSS severity and exploitability. For example, a high-risk CVE may not be considered a vector as there are mitigating controls that prevent it from being

exploited. A vector with a high CVSS score may have a lower Ultra Red score if the weakness exists but cannot be exploited due to missing prerequisites. In both cases, continual monitoring will update their status and priority if the situation changes. Asset risks are categorized into logical groupings such as VPNs, development environments, admin & sensitive information. This helps managers and teams understand the potential impact of affected assets to further assist with prioritization.

ULTRA RED's continuously expanded vulnerabilities are tracked over time, giving them essentially a "timeframe" of how long they are found for. Due to a continuous scanning approach - vectors are tracked over time. Once a vector is detected, it is given a "finding date." Because of the continuous scanning, every vector that was detected and is re-found in the next scan will continue to "live." When a vector is not identified in the next scan, it is identified as an anomaly and re-tested in different time intervals to accommodate for potential network disruptions. When the vector is no longer identified - it is considered "remediated." By measuring the lifetime of a vector (vulnerability), we can estimate the MTTR of an organization and significantly reduce it.

Conclusion

Organizations can greatly enhance their speed, accuracy, and efficiency by incorporating security AI and automation into their operations. Extensive use of security AI and automation has shown significant benefits, resulting in approximately USD 1.8 million in data breach cost savings and over 100 days reduction in the time to identify and contain a breach compared to organizations that do not utilize these technologies.

Embedding security AI and automation throughout the toolsets of security teams is highly advantageous. By leveraging these technologies across threat detection and response tools, analysts can detect new threats with greater accuracy and effectively contextualize and prioritize security alerts. Automation can streamline parts of the threat investigation process and provide recommendations for faster response. Furthermore, AI-driven data security and identity solutions contribute to a proactive security posture by identifying high-risk transactions, ensuring minimal user friction, and effectively correlating suspicious behaviors.

When implementing AI in security operations, it is essential to select technologies with trusted and mature use cases that demonstrate accuracy, effectiveness, and transparency. This helps eliminate potential biases, blind spots, or drift. Organizations should also establish an operational model for AI adoption that supports continuous learning which will enable them to adapt to evolving threats and technological capabilities. By embracing security AI and automation, organizations can significantly enhance their overall cybersecurity posture and mitigate the impact of cyber attacks.

Security teams can benefit from having security AI and automation embedded throughout their toolsets. Using security AI and automation across threat detection and response tools can help analysts detect new threats more accurately and contextualize and triage security alerts more effectively. These technologies can also automate portions of the threat investigation process or recommend actions to speed response. Additionally, AI-driven data security and identity solutions can help drive a proactive security posture by identifying high-risk transactions, protecting them with minimal user friction, and stitching together suspicious behaviors more effectively. When applying AI within your security operations, look for technologies that offer trusted and mature use

cases with demonstrated accuracy, effectiveness, and transparency to eliminate potential bias, blind spots, or drift. Organizations should plan an operational model for AI adoption that supports continuous learning as threats and technology capabilities evolve.

In conclusion, strengthening resiliency against cyber attacks requires organizations to understand their attack surface and implement effective incident response (IR) measures. By identifying their exposure to industry-specific threats and vulnerabilities, organizations can prioritize their security strategies. Tools like ASM and techniques such as adversary simulation provide valuable insights into risk profiles and exploitable vulnerabilities. Having a skilled IR team in place, equipped with the right protocols and tools, significantly reduces costs and response time in the event of a breach. Investing in IR planning, testing, and training has proven to be a top cost mitigator, resulting in lower data breach costs and faster incident resolution. Implementing network segmentation practices is also crucial to limit the spread of attacks and minimize damage, enhancing overall cyber resiliency and reducing the overall blast radius of a cybersecurity incident, or worse, a data breach.