



State of Threat Management

a Filigran Report

Exploring the shift towards intelligence-led Continuous Threat Exposure Management (CTEM), market readiness and what it means for the future of proactive cyber defense.

Based on a third-party survey of over 550 global cybersecurity leaders and practitioners.

Table of Contents

03 Executive Summary

Too much noise. Not enough signal.

07 The Exposure Gap

More tools, more exposure, less clarity.

11 The Prioritization Bottleneck

Visibility without prioritization is just more noise.

17 Validate to Assess your Readiness

Stop guessing. Start validating.

21 The Future of Risk Exposure Management

The shift to intelligence-led, continuous exposure management.

30 Methodology & Contributions

01

Executive Summary



Cybersecurity professionals are inundated with signals: threat intelligence feeds, vulnerabilities, misconfigurations and attack surface data, alongside an ever-expanding set of solutions designed to make sense of it all. By all accounts this should represent progress; offering more data, more tools and more insight into potential exposure.

At the same time, AI and automation have become increasingly embedded across the security stack, yet this has not necessarily translated into stronger security outcomes. The reality is more complex: **while cybersecurity teams have never had more visibility, many are no clearer about their actual state of risk.**

In this report, we examine how organizations are currently managing cyber exposure, and identify the gap between widespread awareness of risk and the ability to act on it effectively. Based on an independent survey and insights from 550 global security decision makers and practitioners, the research explores how threat intelligence, exposure management, validation, and automation are used in practice, and where they fall short. It provides security leaders, executives, and stakeholders with a clear view of the operational challenges limiting risk reduction today, while highlighting the strategic priorities required to move from fragmented, manual approaches toward continuous, intelligence driven threat exposure management.

Modern security teams at a glance:

84% say attacks they face exploit risks that are known, but not prioritized

97% report difficulties determining whether exposures are actually exploitable

38% only utilize threat intelligence within a continuous, fully automated validation process

41% only have a fully consolidated view of cyber risk exposure

88% agree that without greater automation, it is difficult for them to keep up with the volume of risks they need to assess

Despite increased visibility, only 41% of organizations have a fully consolidated view of cyber risk exposure, highlighting the extent to which fragmentation still limits understanding. Even when risks are identified, prioritization and validation remain significant challenges. For example, 84% say the attacks they face often exploit risks that are already known but not prioritised, while 97% report difficulties determining whether exposures are actually exploitable.

At the same time, manual processes continue to constrain progress. Many organizations still rely on manual approaches for vulnerability assessment (85%) and threat analysis (86%), creating a bottleneck between detection and action. This is further compounded by the fact that only 38% currently use threat intelligence within a continuous, fully automated validation process, reinforcing the gap between point-in-time insight and continuous risk management.

As a result, the challenge is no longer identifying risk, it is operationalizing it. Bridging this gap requires a shift from fragmented insight to continuous, evidence-driven action.




Figure 1: To what extent does your organization rely on manual processes across the following areas of exposure management?

Over the next 12-24 months, organizations expect to invest in:

76% Cyber risk quantification and GRC tools

75% Exposure assessment tools

This report explores why exposure management has become a strategic priority — and why Continuous Threat Exposure Management (CTEM) is gaining traction as a framework for addressing it.

Having a unified view of threat exposure, along with an integrated capability to move seamlessly from threat intelligence to prioritization, security validation, and remediation is essential to reducing real, exploitable risk. In this context, AI and automation are not just being adopted, they are becoming critical enablers of a more continuous, scalable approach to managing exposure.

02

The Exposure Gap



Understanding cyber risk exposure has become significantly more complex in recent years. Security teams are surrounded by data, from vulnerability scanning tools and cloud posture management to external attack surface monitoring and threat intelligence platforms. They get lots of alerts but are they any closer to understanding what threats actually puts them at risk and take any risk reduction decisions?

As environments expand across cloud, on-premise, and third-party ecosystems, this growing mix of tools creates a fragmented picture.

The challenge is no longer collecting alerts, but connecting internal telemetry (what you own and how it's configured) with external telemetry (how attackers behave and what they're targeting) and unifying the signals. Only then can organizations answer a more important question: what actually matters right now?

“Correlating data across multiple tools is time-consuming, visibility gaps make it difficult to assess real business risk quickly.”

Senior IT security decision maker,
Business and professional services, US

That lack of coherence becomes clearer when looking at how organizations manage visibility. The vast majority (93%) experience challenges in maintaining an accurate, up-to-date view of their attack surface, particularly when it comes to seeing across assets and environments (41%), consolidating attack surface data (38%) and adding meaningful risk context (21%). As a result, only 41% have a fully consolidated view of their cyber risk exposure.

In other words, visibility exists — but it is incomplete, inconsistent, and difficult to act on.

Tools currently used to assess cyber risk exposure across organizations




Figure 1: Which of the following tools, if any, does your organization currently use to assess cyber risk exposure?

Cyber risk exposure assessment tool adoption based on organizations' CTEM maturity stage

Organizations without a CTEM framework in place, but with plans to implement one, rely on a markedly different tool mix than those running an established program. Among CTEM-established organizations, GRC adoption is 22 points higher and pen-test reliance 27 points lower — a shift that signals a move away from point-in-time testing toward continuous risk visibility. Custom tool usage, notably, holds steady across both groups.

The pattern is clear: CTEM maturity correlates with a more mature approach to cyber risk assessment tooling.




Figure 2: Which of the following tools, if any, does your organization currently use to assess cyber risk exposure? Results organized by CTEM maturity.

Note: "Planning CTEM" represent organizations without a formal CTEM framework currently in place, but with plans to implement one. "Established CTEM" represents organizations operating a fully established CTEM program.

Threat intelligence: Widely used but under-delivering

While almost all organizations (99%) use it within their SOC, fewer than half (45%) have fully integrated and operationalized threat-intelligence. Organizations are pulling in an average of 14 different threat intelligence feeds, 9 of which are open-source. While this increases the breadth of available insight, it's often not operationalized, requiring teams to manually structure, contextualize, and prioritize what matters. On paper, that should create a comprehensive view; in practice, it often reinforces silos, from exposure data and remediation workflows.

This is where Continuous Threat Exposure Management (CTEM) begins. Its initial stages, scoping and discovery, focus on defining what matters most to the business and aligning disparate signals into a coherent view of exposure. By defining what matters most to the business and building a contextualized view of exposure around it, organizations can start to connect these signals in a meaningful way. The goal is not more data, but better alignment, turning fragmented signals into a shared understanding of risk.

But even with that understanding in place, a new challenge quickly emerges: when everything looks like a risk, how do you decide what actually matters most?

REGIONAL SPOTLIGHT:

Cyber risk exposure visibility across regions

Levels of visibility vary across regions. While 52% of organizations in North America report a fully consolidated view of cyber risk exposure, this compares to just 37% in EMEA and 31% in APAC.

A similar pattern can also be seen in their approach to validation, with 51% of organizations in North America using threat intel within a continuous, fully automated validation process, compared to 35% in EMEA and 27% in APAC.

These differences highlight how organizations are at different stages in their exposure management journey, particularly in how they connect intelligence and validate risk. As exposure continues to evolve rapidly, advancing these capabilities will be critical to enabling faster, more confident decision-making and reducing cyber risk in a consistent and timely way.



03

The Prioritization Bottleneck



Even with better visibility, most organizations face a new problem: too much to act on, and not enough clarity on what requires immediate attention. As the volume and sophistication of threats increases, the challenge shifts from identifying risk to deciding what actually needs to be addressed.

84%

agree that as the volume and sophistication of threats increases, it is becoming more difficult for their security tools to clearly highlight which issues need immediate attention

This is where many security teams begin to lose momentum. Without a structured, intelligence-led approach that highlights the threats that really matter and the attack paths targeting vulnerable assets, prioritization becomes difficult to sustain. And without clearly defined Priority Intelligence Requirements (PIRs), prioritization remains largely manual, limiting its scalability. Nearly half of organizations either completely or mostly rely on manual processes for vulnerability assessment (48%) and threat analysis (45%), creating a bottleneck between detection and action. As a result, teams struggle to scale their efforts and keep pace with the volume of risk.

This issue becomes even more pronounced when validating risk. Almost all organizations (97%) report difficulties determining whether exposures are actually exploitable. In practice, this process is still heavily constrained by manual effort, limited context, and a lack of integration across tools and workflows, making it both time-consuming and difficult to scale.

Challenges faced when validating whether security risks are actually exploitable

Validating whether a risk is actually exploitable should be a routine security function. Instead, for most organizations, it remains a friction-heavy, under-resourced process that competes for attention it rarely wins.

The barriers aren't random. They cluster around two compounding problems: teams lack the confidence to test aggressively (because testing itself carries perceived risk) and when they do, the process is slow, manual, and poorly integrated into existing workflows. Each gap reinforces the next: without the right skills and tooling, teams fall back on effort-intensive methods that burn time without scaling.

The result is a function that exists on paper but underperforms in practice, leaving organizations making priority decisions based on unconfirmed exposure rather than validated threat.




Figure 3: What are the main challenges your organization faces when trying to validate whether security risks are actually exploitable? Combination of responses ranked first, second and third.

THE SENIOR CONFIDENCE GAP:

Highlighting the difference in perception between decision-makers and practitioners

IT security practitioners are more likely to report that validation is not well integrated into existing security processes than senior IT decision-makers (46% vs 36%), reflecting their hands-on role in working across tools and workflows. Practitioners are also 10 percentage points more likely to claim a fully consolidated view of cyber risk exposure, yet 37% cite a lack of leadership buy-in as a barrier to improving exposure management, compared to just 27% of decision-makers.

This points to a confidence gap between strategy and execution. Practitioners feel they can see the risks, but lack the strategic support to act on them — while senior decision-makers are less likely to recognize these operational frictions. This raises questions about how well current approaches support operational needs in practice.

The real cost of poor prioritization: wasted time

As a result, many organizations are not struggling to identify risk but to act on it. In our survey, 84% agreed that the cyber-attacks they face often exploit vulnerabilities that are already known, but not prioritized. Similarly, 82% of respondents say manual processes make it harder to determine which risks require immediate attention. When prioritization relies on stitching together data across disconnected tools, it becomes slow and inconsistent. The impact of this is tangible: most take more than a day to detect (72%), respond (71%), and remediate (83%) incidents, extending the window of exposure.

The average analyst spends about 42% of their working week on dead-end investigations.

17 hours per analyst, per week, lost to chasing risks that turn out not to matter.

2.1 full working days lost per week




Figure 4: Approximately what percentage of your security team's time is spent investigating potential risks that later prove to be low priority or not exploitable?

*Based on a standard 40-hour working week.

The need for structured prioritization

89%

agree that reducing alert noise would help them identify which alerts represent genuine business risk, so critical threats are less likely to be missed or delayed

82%

agree manual processes make it harder to determine which risks require immediate attention

This is driving demand for a more structured approach to prioritization. Rather than reacting to everything, organizations need a clear way to define which threats are relevant to their environment. This is reflected in the fact that 95% say they would benefit from establishing Priority Intelligence Requirements (PIRs) to better focus on the threats that matter most.

Within the CTEM framework, prioritization is a necessary stage – where PIRs help connect threat intelligence, business context, and exposure data to guide decision-making and ensure effort is focused where it will have the greatest impact.

But even with clearer priorities, the next step is proving what actually matters. Which risks are truly exploitable — and can your security controls stop them?

04

Validate to Assess Your Readiness



As systems and environments change more quickly, periodic assessments alone are no longer enough to manage risk effectively. For many, validation remains a point-in-time activity in an environment that is anything but static.

88%

agree as systems and environments change more quickly, periodic assessments alone are no longer enough to manage risk effectively

Current approaches reflect this tension. While breach and attack simulation (44%) and penetration testing (41%) are widely used, they provide only snapshots of exposure – and in fast-changing environments, those snapshots quickly become outdated. Only 38% report using threat intelligence within a continuous, fully automated validation process, highlighting the gap between point-in-time testing and continuous exposure management.

This gap is not simply a matter of having the right tools, it reflects deeper operational challenges in how those tools are integrated and used. Nearly half (47%) report barriers with integrating existing tools or processes when improving exposure management, while 42% cite a lack of visibility into real-world risk. In practice, validation remains fragmented, often disconnected from the broader security workflow, limiting its ability to inform decision-making in real time.

“Determining exploitability versus theoretical risk remains difficult, especially when asset context and exposure data are incomplete.”


Senior IT security decisionmaker, Business and professional services, US

INDUSTRY SPOTLIGHT: CTEM Across Industries

While most agree that periodic assessments are no longer enough to manage risk effectively, the strength of that view varies by sector.

Those in business and professional services (58%), financial services (43%), and energy (45%) sectors are more likely to strongly agree - likely reflecting environments where sensitive data, critical systems, or high-value assets heighten the need to keep pace with rapidly evolving threats.

At a more strategic level, industry variance is even striking. For instance, 73% of energy sector respondents report a fully established CTEM program, compared to just 33% in retail. Financial services, despite its heavy regulatory burden, sits at 40%. This 39 point difference highlights how sector-specific risk pressures are shaping the pace of adoption, with sectors operating critical infrastructure or managing high-value assets moving faster toward full CTEM adoption.



Tools sprawl is real – but direction is clear

When evaluating new exposure management capabilities, the focus is on tools that integrate with existing security platforms, deliver accurate and reliable risk insights, and introduce automation to reduce manual effort. This reflects a shift away from isolated testing toward validation that is continuous, connected, and operational.

Factors that matter the most when evaluating new exposure management tools or capabilities

Top 4 Factors



Other considered factors



Figure 5: When evaluating new exposure management tools or capabilities, which factors matter most to your organization?

TOOL SPRAWL:

A Catalyst for Change

This trend is also being driven by the practical challenges of managing increasingly complex security stacks. For example, 31% of organizations cite having too many tools to manage effectively as a barrier to maintaining accurate and up-to-date visibility into their attack surface.

At the same time, 42% say that investment in exposure management capabilities is being driven by the need to consolidate or rationalise existing security tools, highlighting how tool sprawl is both an active pain point and a catalyst for change.

31%

say too many tools prevents them from maintaining accurate and up-to-date visibility into their attack surface.

42%

say investment in exposure management capabilities is driven by the need to consolidate existing security tools.

“The most frustrating and time consuming part is consolidating and analyzing data from multiple security tools and systems to get a clear, unified view of our actual risk exposure.”

Senior IT security decision maker, IT and technology, Canada

This should drive a move toward continuous, intelligence-led validation, increasingly enabled by Adversarial Exposure Validation (AEV) tools. By combining threat intelligence with real-world attack emulation, organizations can move beyond theoretical risk and understand what can actually be exploited in their environment, providing the evidence needed to assess exposure and control effectiveness with confidence.

Once that clarity is established, the focus shifts from understanding risk to acting on it — quickly, consistently, and at scale.

05

The Future of Risk Exposure Management



Exposure management is no longer a static function — it is an evolving discipline, shaped by shifting market dynamics, rapid technological change, and an increasingly complex threat landscape. Two key trends are emerging that point to the future of exposure management: the accelerating adoption of AI and automation, and the growing demand for structured, risk-reduction frameworks.

AI and automation - the fastest moving curve in the dataset

Once risk is understood and validated, the challenge becomes execution. Acting on exposure at speed is no longer simply an efficiency issue but a structural constraint. As environments grow more complex and interconnected, the volume of decisions required to manage exposure has outpaced what manual processes can realistically support.

AI and automation are becoming an increasingly important part of how organizations manage exposure. Adoption is already underway, with 37% of exposure management processes currently AI-driven, rising to 59% over the next two years, on average. This reflects a shift toward more continuous, operational decision-making, where insight must translate directly into action.

88%

agree that without greater automation, it is difficult for security teams to keep up with the volume of risks they need to assess

AI and automation are most valuable in the areas where execution has historically broken down: detecting exposures (59%), validating whether they are exploitable (54%), and prioritizing remediation based on real business risk (53%). These are not isolated use cases, but the core stages of exposure management where manual approaches have struggled to keep pace.

Areas of exposure management most likely to benefit from AI and automation

These findings align closely with the five stages of the CTEM cycle: Scoping, Discovery, Prioritization, Validation, and Mobilization. Organizations in the planning phase gravitate toward AI at the front of the cycle, particularly for discovery and detection, while mature programs look to AI further downstream, at the validation and prioritization stages. This pattern has direct implications for how AI capabilities should be positioned based on buyer maturity.

Put simply, the AI bottleneck shifts down the workflow as CTEM programs mature. Planning organizations want AI to surface threats; mature organizations want AI to validate and rank them. Among planning-stage organizations, detection scores 68% versus just 38% for prioritization — a 30-point gap. Among mature organizations, that gap nearly disappears, sitting at just one percentage point.




Figure 6: Which areas of your organisation's exposure management workflow do you believe would benefit the most from AI and automation? Results organized by CTEM maturity.

Note: "Planning CTEM" represent organizations without a formal CTEM framework currently in place, but with plans to implement one. "Established CTEM" represents organizations operating a fully established CTEM program.

Detection, validation, and prioritization are no longer separate activities, but part of a unified process where insight must translate directly into action. AI enables this by improving the speed and consistency of how threat intelligence is applied to exposure data, while automation ensures those insights are acted on in a timely and coordinated way.

“Identifying and prioritizing cyber risks is largely manual, a strong push for automation is needed.”

IT security manager/practitioner, Public sector, Japan

The need for structured cyber risk reduction

As organizations look to move towards continuous exposure management and proactive security, investment is increasingly focused on both the technologies that identify risk and the strategic initiatives that help manage it. Over the next 12–24 months, around three quarters of organizations plan to invest in cyber risk quantification and GRC capabilities, as well as exposure assessment technologies. This reflects a growing recognition that visibility or validation alone is not enough: outcomes need to be evidence-based and communicated in a way that drive actions and informs investment decisions to continuously reduce cyber risk.

Over the next 12-24 months, organizations plan to increase investment in:

76% Cyber risk quantification and GRC tools

75% Exposure assessments (e.g. vulnerability management, attack surface management)

At its core, this is about connecting technical exposure to business impact. Exposure management provides the insight required by identifying vulnerabilities, attack paths, and areas of weakness, while GRC capabilities translate that insight into measurable, decision-ready risk. Together, they enable a more consistent and actionable understanding of cyber risk.

“I think it’s about taking into account the probabilities of the impact of risks. Evaluating the risks, the potential impact, and the possibility of avoidance.”

Senior IT security decision maker, Retail, France

REGIONAL SPOTLIGHT:
From Exposure to Risk Quantification

While investment in cyber risk quantification and GRC capabilities is increasing globally, the pace varies by region. Organizations in North America are most likely to report plans to increase investment (86%), compared to 77% in APAC and 65% in EMEA.

This builds on the earlier differences seen in visibility and validation maturity, where organizations in North America were more likely to report a consolidated view of risk and more advanced validation approaches. As a result, higher levels of investment likely reflect a continuation of this progression — moving from understanding exposure to embedding it into measurable, decision-ready risk.



Frameworks such as CTEM bring these elements together into a continuous process, linking visibility, prioritization, validation, and remediation with risk-based decision-making. This ensures that what is identified and assessed is also acted upon in a structured and repeatable way.

The factors driving this investment reinforce the challenges seen earlier in this survey report. The increasing frequency and sophistication of cyber attacks is a clear external pressure, but equally important are internal challenges — namely the need to better prioritize risks and to base decisions on validated, real-world insight. These are the same issues that have limited organizations’ ability to move from visibility, to prioritization, to effective action.

Factors driving investment in cyber exposure management tools

Investment in exposure management is driven by roughly equal internal and external forces, but the individual scores tell a more nuanced story.

Attack frequency and sophistication leads all factors at 54%, confirming that the external threat environment remains the primary catalyst for security investment. On the internal side, the need to better prioritize risk (48%) and growing exposure volumes (47%) signal that operational pressure is building from within, independent of what's happening outside. Tool consolidation at 42% adds another layer: organizations aren't just investing more, they're rationalizing what they already have. These findings suggest that investment decisions are no longer reactive to threats alone and are increasingly shaped by the limits of existing programs.

External Pressure

from outside the organization

45%

Average Total

Attack frequency

54%

Increase in the frequency or sophistication of cyber

Regulation & compliance

43%

Regulatory or compliance requirements

Executive/board scrutiny

39%

Increased pressure or scrutiny from executives or the

Operational Pain

internal capacity issues

46%

Average Total

Need to prioritize risk

48%

Need to better prioritize cyber risks and exposures

Volume of exposures

47%

Growing volume of exposures becoming difficult to manage

Tool consolidation

42%

Need to consolidate or rationalize existing security

Figure 7: Which factors, if any, are driving your organization's investment in cyber exposure management tools?

Taken together, this points to a more mature approach to managing cyber risk. Rather than reacting to individual issues, organizations are building the capability to continuously assess, contextualize, and reduce exposure over time. In this model, exposure management and GRC are no longer separate functions, but part of a connected discipline. Using this approach, cyber risk can be understood, communicated, and reduced in line with real-world threat.

The biggest barrier is organizational, not technical

Security teams know exactly what they need but fragmented tools, manual processes and unvalidated risk create limitations. Visibility has improved, but without context, validation, and coordination, it is not translating into meaningful risk reduction. The result is a growing gap between what is known and what is addressed. Closing this exposure gap will require tighter integration between intelligence and exposure, greater automation, and a shift from periodic assessment to continuous, intelligence driven risk validation.

By aligning technical insight with real-world threat and business impact, CTEM enables organizations to move beyond identifying potential issues toward consistently reducing those that matter most. It transforms exposure management from a collection of activities into a discipline, one that is measurable, repeatable, and outcome-driven.

The direction is clear. Investment is increasing, automation is accelerating, and validation is becoming central to decision-making. But progress will depend on more than technology alone.

Organizations that succeed will be those that break down silos, align teams around shared priorities, and embed continuous decision-making into their operations. They will treat exposure not as a static list of issues, but as a dynamic, evolving risk that must be constantly understood and reduced.

Because in today's threat landscape, knowing where you are exposed is only the starting point.

What matters is how effectively, and how quickly, you can do something about it.

From framework to execution

Adopting CTEM requires more than defining a process, it demands collaboration across different security teams and the ability to connect threat intelligence, exposure, and validation in a way that supports confident, real-time decision-making.

This is where many organizations continue to face challenges, particularly in aligning intelligence with real-world exposure and proving what is truly exploitable.

At Filigran, we see this as the critical step in making exposure management operational. Operationalizing threat intelligence, pairing it with continuous, adversarial exposure validation and looping it with cyber risk reduction enables organizations to move from assumption to evidence, and from insight to action.

Solutions such as Filigran's eXtended Threat Management (XTM) Platform combine threat intelligence, exposure validation, and risk quantification capabilities, which are designed not as standalone tools, but as part of a continuous, intelligence-led approach – underpinned by an agentic foundation – focused on one outcome: reducing real-world, exploitable risk.

eXtended Threat Management (XTM) Platform



Survey Methodology

We conducted a global quantitative survey between February and March 2026, gathering insights from 550 security decision-makers and professionals.

Respondents were required to have at least an understanding of their organization's exposure management processes and to work in organizations with a minimum of 1,000 employees.

Participants spanned a range of sectors, with a particular focus on financial services, technology, and government/public sector organizations.

The study covered multiple countries, including the United States (150), Canada (50), France (50), Germany (50), United Kingdom (50), Australia (50), Singapore (50), Japan (50), and the UAE (50).



About Filigran

Filigran, founded in France in 2022, stands out in the cybersecurity landscape with its unique open-source, threat-informed approach to Continuous Threat Exposure Management (CTEM). Filigran's eXtended Threat Management (XTM) platform delivers proactive security by combining threat intelligence, exposure validation, and cyber risk quantification — underpinned by an agentic foundation.

About Vanson Bourne:

Vanson Bourne is an independent specialist in market research for the technology sector. Their reputation for robust and credible research-based analysis is founded upon rigorous research principles and their ability to seek the opinions of senior decision makers across technical and business functions, in all business sectors and all major markets. For more information, visit www.vansonbourne.com.



To find out more, visit

filigran.io