

SANS 2026 Exposure Management Survey

Every Tool,
No Assembly:
How Exposure
Management
Stalled Between
Adoption and
Operation



Written by Jonathan Risto

October 2026

Table of Contents

2026 Exposure Management Survey Key Findings ... 3

Introduction..... 4

The Adoption-Effectiveness Gap 6

The Visibility Gap That Surfaces Everywhere 9

A Data Analysis Problem..... 11

Remediation Is a Coordination Gap 13

What the Field Gets Right..... 16

Conclusion 18

Survey Demographics..... 19

Sponsor.....20

About the Author



Jonathan Risto

SANS Principal Instructor
Technical Director for the Cyber
Posture Management Program,
Government of Canada

Jonathan brings 25 years of experience across network design, IP telephony, service development, security, and project management. He is currently the Technical Director for the Cyber Posture Management Program at the Government of Canada, where he conducts cybersecurity research focused on vulnerability management and automated remediation. He is also the Founder and Principal Consultant of ZenzizenSec Inc., advising organizations on scalable security programs, and the author and instructor for **LDR516: Strategic Vulnerability and Threat Management**. The course focuses on helping leaders build exposure-driven programs that align cyber risk with business priorities.

 [READ MORE HERE](#)

2026 Exposure Management Survey Key Findings

Exposure management (EM) has hit near-universal adoption, but almost nobody has it running well. Nearly every organization has a program in place, yet only a handful have seen real risk reduction over the past year, and fewer still would call their approach continuous. That gap comes down to one chain of failure. Asset visibility is unreliable, so the data feeding into EM is fragmented, and that fragmentation stalls out in organizational turf wars over who owns the fix. None of this gets solved with another tool. It takes consolidation, automation, and integration to turn a pile of disconnected parts into one program that actually holds together.

— EXPOSURE DEBT IS GROWING

38%

saw minimal or no measurable reduction in exposure risk over the past year, evidence that exposure debt is accumulating faster than programs can retire it.

— HOW REMEDIATION GETS STARTED

32% | 24%

32% of organizations still route remediation through a security analyst manually creating tickets, while only **24%** have moved to automated workflows.

! *"This is a coordination gap, not a capability gap."*

— TOP EXPOSURE VISIBILITY GAPS



— EXPOSURE DATA IS SCATTERED

44% | 16%

44% pull exposure findings from six or more disconnected systems, while only **16%** have fully centralized their data in a single platform.

— TOP REMEDIATION BARRIERS



! *The barriers are organizational, not technical.*

— ADOPTION OUTPACES MATURITY

91% | 7%

91% have implemented some form of exposure management, but only **7%** run it continuously, combining integration, automation, and ongoing measurement into one cycle.

! *"Continuous is the whole point. Everything short of it is a snapshot of yesterday's risk in a landscape that moved overnight."*

— PRIORITIZATION INSTINCTS ARE SOUND

44% | 43%

44% rank asset criticality and known exploit availability as their top prioritization factors, effectively tied with CVSS severity at **43%**.

! *Risk-based is the most common self-described approach (36%), ahead of severity-based (20%), and 91% use at least one external threat signal.*

— KEY TAKEAWAYS

01 Fix ownership first

Clear, named accountability for remediation timelines closes the coordination gap faster than any new tool.

02 Consolidate before automating

Centralizing fragmented exposure data is the precondition for automation, not a parallel project.

03 Measure outcomes, not adoption

91% have adopted EM, but only **15%** show real risk reduction and only **15%** are confident in their prioritization. Adoption isn't the win, those two numbers are.

Introduction

EM is the process of continuously discovering, prioritizing, validating, and remediating anything across an organization's environment that could lead to compromise.¹ That scope is deliberately broader than classic vulnerability management. A misconfigured cloud identity, an overprivileged service account, or an unmanaged SaaS integration all qualify as exposures in the same way an unpatched Common Vulnerabilities and Exposures (CVE) does. The distinction matters, because the attack surface has expanded well beyond what a scanner alone can map, and a program scoped only to known vulnerabilities will, by definition, miss everything else.

This survey set out to measure how far organizations have come in building their EM programs and there appears to be a central gap between adoption and effectiveness. Ninety-one percent of respondents say their organization implements exposure management, yet among those with a program, only 7% describe their approach as truly continuous, meaning integrated, automated, and measured on an ongoing basis.² Just 15% of that same group report significant risk reduction over the past year, a sign that adoption has become routine while effectiveness lags behind it. The sections that follow trace that gap to its structural roots, starting with unreliable asset visibility that feeds fragmented data, which then stalls in organizational coordination problems no single tool can close. The last section turns to what the field already gets right, because the direction of travel matters as much as the distance still to cover.

Key data points to note:

1. 38% of organizations report minimal or no measurable reduction in exposure risk over the past year, a sign that exposure debt is accumulating faster than programs can retire it.³ The missing pieces are decision capacity and remediation ability. The gap between finding an exposure and fixing it closes only when clear ownership names who is accountable for the fix, rather than depending on one team persuading another to act.
2. 91% of organizations have implemented some form of EM, but only 7% run it continuously, 15% report a significant reduction in overall exposure risk, and 15% are very confident they are prioritizing the right exposures first. Adoption is nearly universal, but maturity and outcomes are not.
3. 48% of organizations cite unmanaged or unknown assets as their top visibility gap, and 46% point to unreliable or incomplete asset inventory as their top prioritization blocker, with lack of asset context and insufficient automation clustered close behind at 42% each. When the same underlying weakness surfaces across separate questions, it points to a root cause rather than a coincidence.

¹ This survey instrument covers discovery, prioritization, and remediation; validation practices were not measured but are planned for 2027.

² The 91% adoption figure reflects all 410 respondents, while the 7% and 15% figures reflect the 371 respondents who went on to detail their approach to exposure management and risk reduction.

³ Exposure debt describes unresolved exposure obligations accumulating when discovery and scoping outpace an organization's capacity to prioritize, validate, and remediate. See <https://zenizensec.com/tag/exposure-debt>

4. 44% of organizations pull exposure findings from six or more disconnected systems, another 43% sit at three to five, and only 16% have fully centralized their data in a single platform. That fragmentation resurfaces as a blocker for both prioritization and remediation, because findings that live in separate consoles cannot be correlated, deduplicated, or acted on as a single picture of risk.
5. 44% of organizations point to responsibility spread across multiple teams as a top remediation barrier, with risk of operational disruption and competing business priorities each cited by 39% and unclear system ownership by 37%. The barriers are organizational, not technical. Remediation still typically starts with a security analyst manually creating tickets (32%), and only 24% of organizations have moved to automated workflows. This is a coordination gap, not a capability gap.
6. 44% of organizations now rank asset criticality and known exploit availability as their top prioritization factors, effectively matching the CVSS severity score at 43%. Risk-based is the most common self-described approach (36%), ahead of severity-based (20%), and 91% of organizations use at least one external threat signal to enrich their decisions. The field has learned to ask the right questions about what to fix first, which makes the gap between that judgment and actual remediation outcomes all the more telling.

Put these six together and you get one chain, not six separate problems. Bad visibility feeds fragmented data, fragmented data stalls out on unclear ownership, and unclear ownership is what turns a manageable backlog into exposure debt that never stops growing. That's the order the rest of this paper follows: visibility, then data, then remediation, then a look at what's already working, before circling back to what it actually takes to close the loop.

The Adoption-Effectiveness Gap

EM has crossed the adoption threshold, with more than 91% confirming their organization runs a program. That 91% adoption rate, however, masks a wide range of maturity. Only 7% of respondents describe their approach as truly continuous, combining integration, automation, and ongoing measurement into a single cycle (see Figure 1). That sounds like progress, and to some extent it is. But it coexists with a sizable share of organizations still making decisions based on a severity score or reacting to vulnerabilities only as they surface. All combined, nearly two in five organizations are running their programs on approaches the field has been actively moving away from.

Q: Which statement best describes your organization's current approach to exposure management?




Figure 1. Current Approach to Exposure Management

Expert Corner

“Only 7% of organizations run exposure management continuously, and continuous is the whole point. Everything short of it is a snapshot of yesterday's risk in a landscape that moved overnight.”



Chris Dale

SANS Principal Instructor and
Chief Hacking Officer at River Security

COURSE TAUGHT

SEC504:
Hacker Tools, Techniques,
and Incident Handling

[VIEW PROFILE](#)

Significant risk reduction over the past year remains uncommon, reported by just 15% of respondents, while the rest of the sample clusters in more modest or unmeasured territory (see Figure 2). Confidence in prioritization tells a similar story. Only 15% are highly confident they are fixing the right things first (see Figure 3) and the majority land at somewhat confident, a qualified answer that amounts to “We think so, but we cannot prove it.” The gap reflects where most EM programs actually sit today—still maturing rather than fully operational—and where addressing the issue is less about adding a new capability and more about connecting the capabilities organizations already have.

Q: Over the past 12 months, how much has your organization reduced its overall exposure risk?




Figure 2. Risk Reduction Over the Past 12 Months⁴

Q: How confident are you that your organization is prioritizing the right exposures first?




Figure 3. Confidence in Prioritization

⁴ Note that this question asked respondents to assess their own risk reduction without defining risk, so responses likely reflect different underlying measures

Underneath the averages sits a long tail of organizations operating with minimal infrastructure with 26% reporting minimal or no automation in their EM programs (see Figure 4). Nine percent use no external threat signals at all, meaning they prioritize without reference to what is being actively exploited in the wild. Remediation still starts, more often than not, with a security analyst manually creating tickets (32%), while only 24% have moved to automated workflows and 9% have no consistent remediation process at all. These represent a substantial portion of organizations that would, on the top-line question, count as “implementing exposure management.”

Q: Where does your organization currently use automation in exposure management?
Select all that apply.




Figure 4. Where Organizations Use Automation in Exposure Management

Last year’s SANS Attack Surface Management (ASM) survey puts that disconnect in sharper relief.⁵ In that survey, 78% of respondents said they expected their ASM solution to deliver continuous monitoring or adaptive scanning, and nearly 59% preferred a daily scanning cadence. The two surveys cover different populations and different years, so the comparison is directional rather than longitudinal, but the direction is telling. The demand for continuous, always-on visibility at the discovery layer is nearly universal, yet only 7% of organizations in the EM survey have translated that expectation into a continuous end-to-end EM program.

Most organizations have moved well beyond doing nothing, and even a moderate reduction in exposure risk reflects real work by real teams. But adoption alone has not delivered what the discipline promises and the sections that follow outline why. The chain starts with assets organizations cannot reliably see, runs through data they cannot consolidate, and stalls at coordination gaps between the teams that find problems and the teams responsible for fixing them. The fix at every link is the same, better governance, process, and clear ownership.

⁵ SANS 2025 Attack Surface & Vulnerability Management Survey: Hackers Don’t Wait—Why Should We?. October 2025.
www.sans.org/webcasts/sans-2025-attack-surface-vulnerability-management-survey-hackers-dont-wait-why-should-we

The Visibility Gap That Surfaces Everywhere

When organizations are asked where their largest exposure visibility gaps lie, unmanaged or unknown assets top the responses at 48% (see Figure 5). And when asked what most prevents them from prioritizing exposures effectively, unreliable or incomplete asset inventory leads again at 46%, with lack of asset context and insufficient automation clustered just behind it in a near-tie (see Figure 6). The same underlying weakness, an incomplete grip on what the organization actually owns, surfaces as the top visibility problem and the top prioritization blocker simultaneously. When one issue leads two separate questions, it may not be a finding, but a root cause.

Q: Where are your largest exposure visibility gaps?
Select your top 3.




Figure 5. Largest Exposure Visibility Gaps

Q: What most prevents your organization from prioritizing exposures effectively?
Select your top 3.



Figure 6. Barriers to Effective Prioritization

The discovery tooling landscape explains part of that pattern. Vulnerability scanners lead by a wide margin at 92%, which makes sense since they are the oldest and most operationally embedded tool in the stack (see Figure 7). Coverage thins out after that, across threat intelligence feeds, asset inventories, external attack surface management, and identity and privilege analysis, with manual discovery bringing up the rear. Identity and privilege analysis stands out because it lags behind, even though compromised credentials and overprivileged accounts now drive a substantial share of breaches. Organizations are best equipped to find the exposures they have always looked for and least equipped to find the ones attackers are increasingly targeting.

Q: What sources does your organization use to discover exposures?
Select all that apply.




Figure 7. Sources Used to Discover Exposures

In the open-text responses to the visibility question, several respondents described inherited environments from mergers and acquisitions that remain invisible to enterprise security tools until they are fully migrated into the organization’s identity stack. Others pointed to cloud assets spun up by automation pipelines without being linked to a business service owner, and plant floor assets visible to operations but not to security tooling. Each example names a specific blind spot sitting outside the reach of the scanners organizations rely on most. And, in each case, the root issue is a missing process rather than a missing tool. Nobody wrote security into the migration playbook or the provisioning pipeline does not assign an owner or the integration path between OT operations and enterprise security does not exist.

An organization that does not know what it owns cannot prioritize what it finds, because every prioritization decision depends on asset context: what the asset does, who owns it, and what would break if it were compromised. That context is exactly what the prioritization blocker data says is missing. The visibility gap does not just make discovery harder. It undermines everything downstream.

A Data Analysis Problem

With scanners everywhere and organizations finding exposures at scale, discovery doesn't seem to be the problem. Instead, the challenge is what happens once findings land across systems that do not talk to each other, with no shared schema and no single place to connect them. The effort required to make sense of what the tools produce is where programs stall.

Forty-three percent of organizations draw exposure findings from three to five systems, with 25% pulling from six to 10 (see Figure 8). In addition, what those systems produce is rarely consolidated as only 16% of respondents report fully centralized exposure data in a single platform, and the largest group describes their data as partially centralized (47%) with some still siloed (see Figure 9). It's clear that the missing component is a defined and shared schema with integrated playbooks that connect discovery to remediation.

Q: How many systems contribute exposure findings in your organization?




Figure 8. Number of Systems Contributing Exposure Findings

Q: How centralized is your exposure data?




Figure 9. Centralization of Exposure Data

That fragmentation does not stay contained in the tooling layer but instead resurfaces as a direct blocker to two separate stages of the process. Among the top barriers to effective prioritization, exposure data fragmented across tools ranks at 41%, sitting in the same leading cluster as unreliable asset inventory and insufficient automation. Tool fragmentation shows up at 32% as a barrier to getting exposures fixed. That means the same root cause (data scattered across disconnected tools) is doing damage twice: once when teams try to decide what to fix first and again when teams try to get it fixed. It's not two separate problems. It's one structural problem draining effectiveness at two different points in the pipeline. This is not a data problem, but a data analysis problem. Organizations are not short on findings, instead, they are drowning in them, spread across disconnected dashboards (41%) with no common schema and no shared asset key. The analyst moves, at times, between more than 10 systems (19%) with no single place to weigh one finding against another. In practice, what most organizations have built is a data swamp where the raw material is there, but extracting value from it requires consolidation, automation, and integration that the program has not delivered.

We have lots of data available in organizations, but a program's maturity is being held back due to the inability to access and leverage scattered pieces of data quickly to assist with the decision-making process.

The capability most directly aimed at that analysis bottleneck is also the one organization are least deliberately deploying. Forty-seven percent of respondents report that AI is not currently used in their EM program. Among those who do use it, the applications cluster around the analytical tasks where fragmented data creates the most friction, with correlating exposures across data sources, exposure prioritization, and remediation guidance in the lead (see Figure 10).

Q: How is AI currently used in your exposure management program?

Select all that apply.




Figure 10. How AI Is Currently Used in Exposure Management

That stated absence of AI, however, likely overstates the reality. Nearly half of the respondents say they do not use AI at all, but open-text responses reveal a pattern with several respondents indicating that their vendors embed AI in their tooling but the organization does not classify that as "AI use" because no one made a deliberate decision to deploy it. AI is arriving in EM the way it is arriving in much of enterprise software, embedded in vendor platforms rather than adopted as an explicit organizational capability. Whether organizations are not using AI at all or are using it without recognizing it, neither position reflects the deliberate, strategic deployment that the analysis bottleneck demands.

The structural fix is not a single tool or a single purchase but rather the work of consolidating fragmented data into a shared view and automating the correlation that analysts currently perform by hand. The systems that discover exposures need to connect to the systems that track their resolution. Most organizations have not done that work yet, and they are fighting the same data swamp with manual effort. The sections that follow show where that effort breaks down.

Remediation Is a Coordination Gap

The biggest barriers to remediating organizational exposures are organizational issues, with the top barrier being that the responsibility is spread across multiple teams (see Figure 11). The barriers you might expect to dominate—tool fragmentation, staffing limitations, and lack of executive support—rank consistently lower. Governance, decision capacity, and remediation coordination all come before the technical concerns. The teams tasked with remediation have the tools and, broadly, the mandate, but they lack a clear line from “this needs to be fixed” to a named person who will fix it by a specific date.

Q: What are the biggest barriers to remediating exposures in your organization?
Select your top 3.




Figure 11. Biggest Barriers to Remediating Exposures

The structural reason is clear: The EM function owns the outcome (i.e., visibility into risk, prioritized findings, and a timeline for remediation) but it does not control the teams doing the actual work. Remediation falls to IT operations, application owners, cloud platform teams, and third-party vendors. For all of them, patching a vulnerability is a secondary duty competing with release schedules, uptime commitments, and their own backlog. The result is that the security team can set the priority and create the ticket, but it cannot make the work happen faster than those teams are willing or able to move.

The process for how remediation activities are initiated is primarily from security analysts manually creating and assigning tickets or SecOps teams coordinating on a case-by-case basis (see Figure 12). This ultimately means that for the majority, remediation begins with a human writing a ticket and hoping the right team picks it up.

When organizations want to improve exposure management, they often reach for a new tool because a new tool improves discovery. It rarely improves governance, decision capacity, or remediation throughput, so the program ends up carrying more exposure debt without addressing what created it.

Q: How are remediation activities typically initiated in your organization?




Figure 12. How Remediation Activities Are Initiated

The data shows that remediation is primarily a coordination gap rather than a capability gap. Organizations are not failing because they lack the technical ability to apply a patch or reconfigure a system. They are failing because the organizational machinery between runs on manual handoffs, unclear ownership, and competing priorities that the EM function has no authority to override. The best method for addressing this issue lies in governance with a clear ownership of assets, clear accountability for remediation timelines, and a process that does not depend on one team persuading another to care.

Expert Corner

“Vulnerability and exposure management programs are going through one of the largest transformations in recent memory. While programs work at maturing this functional area, the backlog of older unmitigated risks keeps growing. One constant remains, though: The biggest barriers to remediating organizational exposures are organization issues, with the top barrier being that the responsibility is spread across multiple teams.

No matter where your organization is in its maturity of vulnerability and exposure management, have clear responsible and accountable parties across all tasks. This removes any ambiguity across the organization and sets your maturing vulnerability and exposure management programs up for continued success in the future.”



Kevin Garvey

Certified SANS Instructor and
Senior Cybersecurity Practice Manager
at World Wide Technology

COURSES TAUGHT

LDR512:

Security Leadership Essentials for Managers

LDR516:

Strategic Vulnerability and Threat Management

 [VIEW PROFILE](#)

What the Field Gets Right

When it comes to prioritizing exposures, asset criticality and business impact, along with known exploit availability, rise to the top with raw CVSS severity score close behind (see Figure 13). For years, the industry defaulted to CVSS as the primary triage mechanism, a score that measures technical severity in isolation without accounting for whether a vulnerability is actually being exploited or whether the affected asset matters to the business. The fact that business impact and exploitability now match it represents a genuine shift in how the field thinks about risk. That shift is reinforced at the program level, where risk-based is the most common self-described approach to EM at 36%, ahead of severity-based at 20%.

Q: What factors most influence how your organization prioritizes exposures?
Select your top 3.




Figure 13. Factors That Most Influence Prioritization

The external signals organizations feed into that prioritization process tell a similar story. The CISA Known Exploited Vulnerabilities catalog leads, followed by vendor exploit advisories and commercial or open source threat intelligence feeds (see Figure 14). Enriching prioritization with real-world exploit context has moved from an advanced practice to a baseline expectation. Last year’s SANS Attack Surface Management survey pointed in the same direction, with 75% of respondents integrating threat intelligence feeds into their ASM platforms, indicating a consistent pattern where threat intelligence integration is expected and no longer a differentiator.

Q: Which external threat signals does your organization use to prioritize exposures?
Select all that apply.




Figure 14. External Threat Signals Used to Prioritize Exposures

Regulation is moving the same way. CISA’s BOD 26-04 supersedes BOD 19-02 and, in doing so, removes the requirement that federal civilian agencies use CVSS to drive prioritization. NIS2 obliges essential and important entities to document risk management measures. The audiences differ, but the direction is consistent: Severity alone is no longer accepted as sufficient basis for deciding what to fix, which is what Figure 14 shows practitioners already doing.

The environments covered by EM programs lead with on-premises infrastructure, public cloud infrastructure, and external internet-facing assets (see Figure 15). That level of consistency across three architecturally different environments is meaningful since discovery in each typically depends on different tooling and different data sources, so clustering above 74% in all three suggests programs have extended coverage across the traditional perimeter as a whole rather than excelling in whatever area they started with. However, we are still missing visibility and action across numerous segments.

OT and IoT is the clear outlier, with only 36% of programs extending coverage there. That does not necessarily mean OT goes unmanaged. It often sits with a dedicated OT team rather than the EM function, so the exposure may be getting assessed, but not as part of the organization’s holistic view of risk, which is itself another manifestation of the fragmentation this survey keeps surfacing. Other areas that hover near 55% coverage are developer environments (54%) and third-party systems (57%).

Q: Which environments does your exposure management program cover?
Select all that apply.




Figure 15. Environments Covered by Exposure Management Programs

Conclusion

Organizations have all the right terminology and have purchased all the right tools, but the survey data shows the same outcome regardless: The individual pieces (discovery, prioritization, remediation) aren't functioning as one connected system. It means they have the parts, but they don't have a working machine.

That failure is not evenly distributed, which is what makes it solvable. Prioritization has matured, threat intelligence is now a baselining practice, and coverage across the traditional environment is broad. Organizations are asking the right questions about what to fix first. What they have not done yet is connect what they see to what they fix, and that missing connection is where the outcomes leak away.

The 7% describing their programs as continuous selected an option defined as combining integration, automation, and ongoing measurement. So, by their own account, they have the three things the rest of the field reports lacking. The survey does not isolate how they got there. What it does show is where everyone else is stuck: 16% have fully centralized exposure data, 24% have automated remediation workflows, and 44% name responsibility spread across teams as a top barrier to getting fixes done. This structural work is harder to purchase than tooling, which is exactly why so few have it.

The remaining work will come from governance and accountability that crosses team boundaries. Someone senior enough needs to tie asset ownership to remediation timelines and make it stick. These organizational choices—not new tools or purchases—are what determines whether EM programs improve.

The scarce resources in exposure management are decision authority and remediation throughput. Neither arrives with a purchase order.

Survey Demographics

The 2026 SANS Exposure Management Survey received responses from IT and security professionals globally. All respondents confirmed their organization is currently implementing EM functions and capabilities. See Figure 16 for details.



Figure 16. Demographics of Survey Respondents

Sponsor

SANS would like to thank this survey's sponsor:





About the SANS Research Program

The SANS Research Program is a key initiative by the SANS Institute and a premier global provider of cybersecurity research and information. SANS Research Program is designed to provide cybersecurity practitioners and leaders with data-driven insights, thought leadership, and solutions that help them better understand and respond to evolving security challenges. All content is authored by SANS instructor experts from around the world who apply their years of experience from hands-on practitioner work in the field, advisory roles, and the classroom to provide education, guidance, and actionable insights that help make the cyber world a safer place.

To learn about sponsorship opportunities for research, content, and in-person or virtual events, email us at Sponsorships@sans.org or go to www.sans.org/sponsorship.