

SANS 2026 Threat Hunting Survey

The Evolution of
Threat Hunting

Written by Josh Lemon

September 2026



Table of Contents

2026 Threat Hunting Survey Key Highlights 3

Introduction..... 4

How Long Have Organizations Been Hunting? 5

Hunt Cadence, Duration, and Maturity 6

Where Organizations Hunt: Priority and Pain..... 8

What Hunters Find:
Attacker Profiles and Techniques 9

Knowing What and Where to Hunt:
Threat Intelligence and Landscape Tracking 11

Methodology: Defined, Ad Hoc, or Absent?..... 13

Outsourcing Threat Hunting 15

Measuring What Matters: Hunt Effectiveness
and Outcomes 16

Barriers to Success 18

The Investment Horizon 19

Conclusion 21

Respondent Demographics..... 22

Sponsor..... 23

About the Author

Josh Lemon

Principal Instructor



Josh Lemon is the Founder and Chief of Digital Forensics and Incident Response at SoteriaSec, a boutique cybersecurity firm operating across Australia and Singapore. He has more than two decades of experience investigating large-scale security incidents for multinational organizations, government agencies, law enforcement, and law firms, and is regularly engaged as an expert witness in legal proceedings involving data breaches and cyberattacks. Before establishing SoteriaSec, Josh held senior incident response and security leadership roles at Uptycs, Ankura, Salesforce, the Commonwealth Bank of Australia, and BAE Systems Applied Intelligence, building and running the operational teams that respond when things go wrong. He is a Principal Instructor for the SANS Institute’s **FOR508: Advanced Incident Response, Threat Hunting, and Digital Forensics**; **FOR572: Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response**; and **FOR509: Enterprise Cloud Forensics and Incident Response** courses.

 [READ MORE ABOUT JOSH](#)

2026 Threat Hunting Survey Key Highlights

This year's data points to a threat hunting community that is more experienced than ever, yet still held back by the same gaps that have shown up in every prior edition of this survey. The highlights below cover where programs are gaining ground, where adversaries continue to outpace detection, and where the friction has grown rather than eased.

— EXPERIENCE RUNS DEEP

82%

Have been threat hunting for at least two years—an accumulated, operational knowledge base, not early experimentation.

— CLOUD IS THE HARDEST HUNT

32%

Identify cloud infrastructure as the hardest environment to hunt in, more than any other area, as adoption keeps outpacing capability.


— AI INTEREST HOLDS STEADY

47% → 39%

39% ranked incorporating AI/ML into hunting tools as a top planned improvement, down from 47% last year—a shift from aspiration to grounded evaluation.

| Early responses describe agentic hunting frameworks that automate repetitive work while keeping human judgment at the center.

— WHAT HUNTERS ACTUALLY FIND



— DATA IS THE NEW BOTTLENECK

34% → 50%

50% rank data quality or quantity as their primary barrier, rising steadily since 2023 and up 16% from last year—for the first time, data has overtaken skilled staff.


— FORMAL MEASUREMENT SLIPS

64% → 40%

40% of programs formally measure hunt outcomes, down from 64% in 2024. Without it, demonstrating value to leadership and justifying headcount both get harder.

| “You tend to find less when you are not looking for it.”

— LIVING OFF THE LAND LEADS



| Ransomware is the exception: Targeted exfiltration narrowly edges out living off the land, 65% to 63%.

— TOP BARRIERS TO SUCCESS



— PROGRAM HEALTH: OUTSOURCING, METHODOLOGY, AND OUTCOMES

58% in-house

Held steady for a second year; outsourcing dropped to 18% from 30% in 2025.

37%

Have a formally defined methodology, down from 46% in 2025 as ad hoc approaches (39%) creep back.

70% / 11%

Report improved security posture, but only 11% report a large (50%+) gain—down from 47% in 2022.

— KEY TAKEAWAYS

01 Fix the data problem first

Prioritize normalization and telemetry coverage—especially cloud and identity logs—over adding more hunts.

02 Re-formalize methodology

Staffing pressure is dragging programs back to ad hoc. Anchor hunts in a repeatable framework like PEAK or TaHiT.

03 Measure what you hunt

Only 40% formally track effectiveness. Without it, the investment case for staff and tooling gets harder to make.

Introduction

The 2026 SANS Threat Hunting Survey is part of a multiyear longitudinal study that has tracked how organizations approach proactive threat detection since 2021. Each year, we survey practitioners and security leaders across industries and organization sizes to build an independent picture of how the threat-hunting community is evolving. This year's respondent base of 500 participants spans technology, banking and finance, government, healthcare, and a dozen other sectors, with operational reach across North America, Europe, Asia, and Latin America.¹ This diversity is important because the benchmarks and frustrations captured in this report reflect real programs at real organizations, from teams still building their first playbook to those running automated, intelligence-driven hunts across multiple cloud environments. More than 80% of respondents have been hunting for at least two years, which means this data is shaped by accumulated experience rather than early experimentation.

Four themes anchored our expectations going into this year's survey:

1. The speed at which adversaries are moving
2. The continued dominance of malware-free tactics
3. The maturity of cloud-hunting capabilities
4. The role of AI-assisted hunting in modern programs

The data delivers on all four, though not always in the direction anticipated. Living-off-the-land techniques dominate across every threat actor category, leading outright for nation-states and organized crime, and running a close second to targeted exfiltration for ransomware groups, confirming that the shift away from custom malware and toward blending into legitimate system activity is now the default adversary posture rather than a niche capability. Cloud remains the hardest environment to hunt in by a significant margin, improving only incrementally year over year, while enterprise reliance on cloud infrastructure continues to grow. And AI and ML, which had been climbing steadily as a planned improvement in prior surveys, appear to be plateauing. Intent has declined two years running, suggesting the community is moving from aspiration toward a more grounded assessment of what these tools can deliver today.

What this year's findings make clear is that the foundational challenges have not changed as much as the industry might hope. Data quality has overtaken skilled staff as the top barrier for the first time in this survey's history, formal measurement of hunt effectiveness continues to decline, and ad hoc methodologies are creeping back after several years of progress toward more defined approaches. At the same time, the shift toward internal capability over outsourcing is accelerating, investment in both staff and tooling is becoming more balanced, and early signals in this year's responses suggest that some programs are beginning to experiment with agentic hunting frameworks that keep human judgment at the center while letting automation handle the repetitive work. The gap between where threat hunting is and where it needs to be is real, but so is the momentum of a community that has been doing this long enough to know what good looks like.

Interest in broad AI adoption is declining for the second-year running, even as early signals point to programs beginning to experiment with agentic hunting frameworks that automate repetitive work while keeping human judgment at the center.

¹ Survey demographics are located at the end of this report.

How Long Have Organizations Been Hunting?

This year, we changed the way we ask respondents about their threat hunting history.

Rather than a simple yes-or-no question about whether an organization actively performs threat hunting (which is how the question was framed in prior surveys) we asked how long respondents have been doing it. The shift gives us a much richer picture of who is in the room. As it turns out, the room is fairly experienced with 33% of respondents having threat hunted two to three years and 27% for three to five years (see Figure 1). That means more than 80% of our respondents have been at this for at least two years, which is genuinely encouraging. This tells us that the findings are shaped by accumulated, hard-won operational experience across a global respondent base and not from a community that just discovered threat hunting. This report offers readers a collective knowledge base that has been building for years, one that speaks to what works, what does not, and what the rest of the industry is doing.

Q: How long have you been performing threat hunting?




Figure 1. Time Threat Hunting

Hunt Cadence, Duration, and Maturity

Two new questions this year track operational tempo, hunt frequency, and duration, giving us a baseline to follow in future editions. On frequency, the largest group of respondents runs three to five hunts per month at 31%, followed by 25% running one to two, 17% running six to 10, and 12% running more than 10. That last number is worth pausing on. Depending on team size, data availability, and tooling maturity, hunting more than five times per month starts to look less like a well-resourced program and more like a team running itself into the ground—particularly when you factor in reporting, retooling, and follow-up actions that come after each hunt concludes.

On duration, 31% of respondents report that a typical hunt takes nine to 24 hours, roughly two to three working days from scoping through to conclusion. Another 24% complete hunts within a full day, and 13% finish in under four hours. When you combine those durations with the higher-frequency threat hunts, the math quickly becomes uncomfortable for smaller teams. A program running six or more hunts per month, each lasting two to three days, is asking a lot of people who almost certainly have other responsibilities alongside their hunting work.

On the maturity side, the proportion of respondents rating themselves as very mature dropped to 9% from 19% in 2025, while the mature and maturing categories both grew (see Figure 2). Whether that reflects a genuine recalibration of self-assessment standards or a shift in respondent composition is worth watching as more years of cadence data accumulate. Either way, that tempo pressure matters in context, since adversaries operating with living-off-the-land techniques and minimal dwell time are not waiting for monthly hunt cycles to find them.

Q: What do you consider your threat-hunting maturity level?




Figure 2. Threat Hunting Maturity Level, 2024–2026

Three dominant themes emerge from our free-text responses that explain the maturity self-assessments fairly well. Data quality and telemetry coverage appear throughout responses across every region and industry sector. Organizations with workable playbooks routinely describe their ceiling as a data problem rather than a process one, with gaps in cloud logging and identity telemetry cited most often.

The second theme is process consistency: Many respondents describe hunting that works in pockets but is not yet repeatable across teams or geographies, with phrases like “depends on individual analysts” and “not consistently run” recurring throughout. The third is the gap between the existence of a playbook and its use. A number of respondents describe documented processes that are not strictly followed, or hypothesis-based approaches that only some analysts are comfortable executing. The community knows what mature threat hunting looks like. The friction is in the data coming in, the consistency of execution, and the distance between what is written down and what actually happens during a hunt.

Organizations with workable hunting playbooks consistently describe their ceiling as a data problem rather than a process one, with gaps in cloud logging and identity telemetry cited most often across every region and industry sector surveyed.

Where Organizations Hunt: Priority and Pain

Networks, workstations, and servers continue to rank as the most important environments for effective threat hunting, capturing 28%, 24%, and 18% of top importance responses, respectively. That is consistent with where threat actors operate once inside an environment—endpoints and lateral movement paths remain the core hunting ground for most programs. Cloud infrastructure ranks 18% in importance, a slight drop from 21% in 2025, which is honestly a little surprising given how much of the enterprise technology stack now lives there, including AI platforms and SaaS tooling that organizations have been adopting at pace. However, cloud infrastructure is also the hardest environment to hunt in by a significant margin, at 32%, down only slightly from 39% in 2025. Year over year, cloud is getting marginally less painful, but the pace of improvement is not keeping up with the pace of cloud adoption, and that gap is worth watching.

The more notable shift this year is in portable devices. Importance jumped from 1% in 2025 to 9% in 2026, and the difficulty rating sits at 23%, second only to cloud. That combination of rising importance and great difficulty, with a gap between the two that rivals what we see with cloud, is a signal worth paying attention to. As workforces become more mobile and device fleets more fragmented, portable devices are likely to become an increasingly contested hunting terrain. Organizations planning their tooling roadmap would do well to add this to the list now rather than wait until the problem becomes more acute.

Cloud is getting marginally less painful, but the pace of improvement is not keeping up with the pace of cloud adoption, and that gap is worth watching.

What Hunters Find: Attacker Profiles and Techniques

Ransomware remains the most commonly uncovered threat actor type when hunting at 55% this year, representing a rebound from 46% in 2025, following a sharper drop from 64% in 2024 (see Figure 3). Business email compromise (BEC) tells a less certain story. At 43%, it is down notably from 64% last year, but that drop likely has more to do with the survey question itself than with BEC activity declining in the field. This year’s question introduced insider threat and supply chain attacks as standalone categories for the first time, both of which had almost certainly been absorbed into “Other” in prior years, thereby inflating categories like BEC by default. Nation-states followed a similar pattern, down to 26% from 41% in 2025. The restructuring matters because it gives a more accurate picture of what hunters are actually finding, rather than one skewed by a catchall category that was quietly masking real threat types. Insider threat now sits at 26% and supply chain at 23%, both large enough that treating them as “Other” was clearly understating their prevalence.

Q: Which attackers do you usually uncover when threat hunting?
Select all that apply.




Figure 3. Threat Actors

When it comes to what those threat actors are actually doing inside an environment, living off the land dominates across all three groups (see Figure 4). The logic tracks: Nation-states and ransomware groups both have strong incentives to steal data, whether for intelligence value or double-extortion leverage. Organized crime groups, whose primary objective is typically to turn a compromise into cash quickly, score slightly lower across most technique categories for the same reason: They value speed over sophistication.

Q: What techniques do you see used for nation-state attacks?
Select all that apply.




Figure 4. Nation-State Threat Actor Techniques

Custom malware shows up at 53% for both nation-states and organized crime, which tracks with those groups’ resources and objectives. Ransomware sits lower at 44%, although that figure is likely conservative rather than a genuine gap. The ransomware payload itself is custom by most definitions, at least at the point a victim encounters it. The more revealing finding is anti-forensic behavior: 47% of nation-state hunts uncover evidence of deleted traces, compared to 40% for ransomware and 31% for organized crime. Nation-states covering their tracks is expected; extended dwell time is part of the playbook. Ransomware and organized crime doing it at these rates is more surprising, since both groups typically reach a point where victim awareness is unavoidable. It suggests these groups are investing in antifoensic capability earlier in the intrusion, buying time before the ransom demand or monetization event actually lands.

The thread running through all three threat actor profiles is the same: living-off-the-land techniques, targeted exfiltration, and increased operational security. This matters directly for how you build a hunt hypothesis. If you are still writing hunts around known bad hashes or IP addresses, you are hunting for threat actors who stopped behaving that way years ago. The combination of data here points squarely toward behavior-based detection over indicator-based detection. For example, hunts that ask what legitimate tools should not be doing, where data should not be going, and what system activity is statistically unusual, rather than hunts that ask whether something matches a known signature. That shift to hypothesis construction is the practical takeaway for anyone planning their next threat hunt.

Knowing What and Where to Hunt: Threat Intelligence and Landscape Tracking

Vendor and independent blogs and research papers remain the most relied-upon sources for staying current with attacker techniques (see Figure 5). The drop in commercial intelligence providers (47%, down from 62% in 2025) and OSINT sources (41%, down from 63%) is notable. Both categories have fallen sharply from last year, which may reflect cost pressures, tool consolidation, or simply a shift toward building internal research capability, possibly aided by AI tools that make rapid synthesis of threat data more accessible to smaller teams.

Threat hunters draw on multiple external sources to stay current with attacker techniques, and usage of nearly all of them declined this year:

- Vendor blogs and papers dropped to 57% from 76% in 2025
- Independent blogs and papers fell to 56% from 69%
- OSINT providers dropped to 40% from 63%
- Industry peer networks fell to 36% from 56%
- Governmental organizations dropped to 41% from 64% in 2025

The one source that held its ground was respondents’ own research, at 53% this year vs. 57% in 2025; a comparatively small drop in a field where most other categories fell by 15 to 20 points. This matters because where hunters source their intelligence directly shapes what a hunt hypothesis looks like, and a broad pullback from external sources, paired with an organization’s own research holding steady, suggests hunters are leaning more heavily on their own analysis to fill the gap rather than distributing that trust across vendors, peers, and OSINT providers the way they did last year.

Q: How do you stay up to date with the newest attacker techniques? Select all that apply.



Figure 5. Sources of Attacker Techniques

Sourcing intelligence is only half the picture, though. What actually happens once a hunter has it is the more operationally relevant question. On that front, 65% of respondents say threat landscape changes are integrated on an ad hoc basis without a formal process, and 23% report fully automated derivation of new hunts from threat intelligence. Only 6% say the threat landscape has no effect on their hunting at all. On tooling preference, EDR/XDR ranks first, which is exactly what you would expect given that it is where threat actors ultimately land and operate. SIEM sits second, NDR and native cloud logs are tied for third, and memory analysis sits last. Memory's position at the bottom of the list does not reflect its value—hunting in memory remains one of the more reliable ways to find threat actors who have avoided leaving artifacts on disk. However, it does reflect the practical difficulty of accessing and analyzing memory at scale, compounded by operating system vendors progressively locking down direct memory access. Native cloud logs climbing to joint third is the more encouraging signal, suggesting that reliance on cloud provider-native tooling is slowly reducing the friction that has historically made cloud hunting so difficult, and that difficulty, as noted earlier, is significant.

Expert Corner

“ Nearly half of the hunts run to find nation-state activity found traces of anti-forensic measures. That finding aligns with what we see in our day-to-day work, and this should heavily influence how we hunt and the evidence we use. Defenders need to consider evidence retention windows, log integrity risks, and the challenge of data collection at scale. This does force us to think hard about what our hunt hypothesis is and how we can test it. It is also really interesting to see memory analysis ranking last on the tooling list—possibly a sign of how difficult it is becoming to collect RAM at scale, even though it is one of the most valuable evidence sources we can have.



Tarot (Taz) Wake

SANS Certified Instructor,
Managing Director (Information Security)
at Halkyn Consulting, Ltd.

Author of the FOR577: Linux Incident Response and Threat Hunting course
Teaches: FOR508: Advanced Incident Response, Threat Hunting, and Digital Forensics course

[VIEW PROFILE](#)

Methodology: Defined, Ad Hoc, or Absent?

Formally defined threat hunting methodologies have declined to 37% this year, down from 46% in 2025 and 51% in 2024, while ad hoc approaches have crept back up to 39%—a split we last saw closer to the 2023 numbers (see Figure 6). That regression matters because formalized methodology is what makes a hunting program repeatable and defensible: It is the difference between a hunt someone can hand off to a new analyst and one that only works because a specific person happens to know where to look. The gains made between 2023 and 2024 in formalizing methodology appear to be unwinding, and the likely reason sits in the staffing data: 38% of respondents say available human resources drive methodology selection, and a further 39% say it is a combination of both staffing and methodology. Add those together, and the influence of staffing on what gets hunted and how it gets hunted is substantial.

When teams are under-resourced, the methodology tends to bend around what people can actually execute rather than what would be most effective, which is exactly how ad hoc approaches take hold. It is worth noting that scalable, adaptable methodologies exist, and organizations that have not yet landed on one would be well served by exploring frameworks such as PEAK² or TaHiT³ rather than continuing to build something from scratch with whoever happens to be available.

Q: Does your organization use one or more clearly defined methodologies to threat hunting?




Figure 6. Threat Hunting Methodologies

² "Introducing the PEAK Threat Hunting Framework," April 2023, www.splunk.com/en_us/blog/security/peak-threat-hunting-framework.html
³ www.nvb.nl/themas/veilig-bankieren/fi-isac/tahiti

On the question of who shapes and executes methodology, the threat hunting team unsurprisingly leads both contribution and execution. More interesting is that CISOs are the most active nonhunter contributors at 42%, with a further 23% involved in both contributing and performing hunts, the latter being a genuinely surprising figure unless the CISO is effectively the security team, which in smaller organizations is not uncommon. That level of CISO involvement is worth flagging because effective threat hunting depends on skills that stay sharp through regular use: understanding what current threat actors are doing, knowing where the soft spots in a network actually are, and recognizing what evidence a given technique leaves behind. Those are hard skills to maintain if hunting is not a regular part of your role. That raises the question of whether a CISO is always the right person to be performing hunts, even where they are well placed to contribute to shaping them. Incident response (IR) teams are active in both roles at 33% contributing and 40% doing both, reflecting the natural overlap between hunting and IR in many programs. External entities contributing at 28% is worth flagging here because it sets up the outsourcing discussion later in this report; nearly a third of methodologies have some third-party input even before you get to organizations that formally outsource their hunting altogether. On review cadence, 32% update their methodology whenever needed, and 29% review it quarterly. The survey does not tell us what is actually driving those “whenever needed” updates, so it is worth treating that finding as a data point to revisit in future editions rather than a settled trend.

The free-text responses we sought reveal a fairly consistent operational pattern across respondents, regardless of region or industry. The most common approach starts with a written hypothesis derived from threat intelligence, recent incidents, or detection gaps, followed by data source validation, query execution across SIEM and EDR platforms, evidence documentation, and a post-hunt step in which confirmed findings are either converted into detections or assigned to platform owners for remediation. Weekly stand-ups or SOC queue coordination are the most frequently cited mechanisms for prioritizing which hunts run next, and evidence is most commonly stored in case management systems, ticketing platforms, or shared wikis, with reproducibility for another analyst cited as an explicit goal by a notable number of respondents.

On analyst autonomy, the responses land in a consistent place: Deviation from the documented methodology is generally permitted, but only when the reason is recorded. That is a practical and sensible standard because it preserves investigative flexibility while maintaining enough of a paper trail to learn from what was found and how. A smaller number of respondents described more structured approaches, including formal hunt life cycle phases mapped to MITRE ATT&CK® Matrix, multilayer telemetry correlation requirements, and AI-assisted companion prompts functioning as living reference documents for platform-specific query syntax. One respondent described a detailed implementation of the PEAK framework with a formal planning phase conducted a quarter in advance, structured execution, and an explicit keep phase capturing working queries and lessons learned. The same respondent noted their team is already exploring agentic threat hunting frameworks to guide AI-driven hunts within a structured methodology rather than letting automation run unconstrained. That last point is worth noting as something likely to appear with more frequency in next year’s responses.

Formalized methodology is what makes a hunting program repeatable and defensible: It is the difference between a hunt someone can hand off to a new analyst and one that only works because a specific person happens to know where to look.

Outsourcing Threat Hunting

The multiyear trend on outsourcing is now fairly clear: Organizations performing threat hunting in house have held steady at 58% for the second consecutive year, while those outsourcing have dropped to 18% from 30% in 2025 and 37% in 2024, and the linear trendlines in Figure 7 illustrate that direction of travel. The rise in “not applicable” responses to 19% also reflects the growing number of consultancies in the respondent base, which matters for reading the outsourcing numbers correctly. Those respondents are not choosing to keep hunting in house, they are on the supply side of the outsourcing market rather than the demand side, so they should not be read as evidence of organizations rejecting outsourced hunting.

This shift toward internal capability makes intuitive sense, no third party will ever know an organization’s environment, its soft spots, or the creative corners a threat actor might exploit better than the people who work in it every day. For those that do still outsource, the governance picture has improved: 53% now define hunting scope and outcomes collaboratively with their third party, up from 50% in 2025, while sole reliance on the third party to define the hunting ground has dropped to 18% from 30% last year. Organizations that outsource but retain a strong voice in scoping those engagements are getting considerably more value from the arrangement than those handing the keys over entirely.

Q: Does your organization outsource its threat hunting?




Figure 7. Outsourcing Threat Hunting

Measuring What Matters: Hunt Effectiveness and Outcomes

Seventy percent of respondents report that threat hunting has improved their organization’s overall security posture in the past 12 months (see Figure 8). Although the data is encouraging in 2026, the year-over-year picture is more nuanced. The proportion reporting larger improvements of 50% or more has declined steadily from 47% in 2022 to around 11% this year, while the share in the 10–25% range has grown. Whether that reflects genuinely diminishing returns, more conservative self-assessment, or simply a measurement problem is difficult to say cleanly, but the measurement data points toward the latter.

Q: Has threat hunting provided improvement in the overall security of your organization?
If so, estimate the percent improvement during the past 12 months?




Figure 8. Threat Hunting Percent Improvement During Past 12 Months

Only 40% of respondents formally measure the success of their threat hunting, down from 51% in 2025 and 64% in 2024, while those not measuring at all has climbed to 49% (see Figure 9). That is a trend heading in the wrong direction. It is hard to demonstrate the value of threat hunting to leadership, justify headcount, or identify what is actually working when just under half of all programs have no structured way to assess their own outcomes. There is also a reasonable possibility that the apparent decline in reported improvements is at least partly an artifact of fewer people measuring in the first place—you tend to find less when you are not looking for it.

Q: Do you formally measure the success/effectiveness of your threat hunting?




Figure 9. Formally Measure Threat Hunting Success/Effectiveness

For the 40% of respondents who do measure their threat hunting's effectiveness, metrics and SLA-based approaches dominate at 76%, with staff feedback coming in at 39%. Both are reasonable proxies for operational impact and are worth combining, since metrics tell you what changed, and staff feedback often tells you why.

On the specific areas showing improvement, attack surface hardening leads the field at 77% combined “some” or “significant” improvement, narrowly ahead of detection accuracy and fewer false positives at 76%. Breakout time (70%), remediation resource savings (69%), and exfiltration detection (68%) follow closely behind, with all five categories clustered within roughly nine points of each other. That tight spread is arguably more interesting than any single leader; it suggests threat hunting is delivering fairly even improvement across the board rather than excelling in one area and lagging in others. Attack surface hardening taking the top spot makes operational sense: A hunt that surfaces an exposed or misconfigured asset produces an immediate, visible fix, which is a more concrete outcome to point to than a probabilistic reduction in false positives. The challenge is converting that felt improvement into something that shows up in a board report, and that conversion requires measurement infrastructure that, based on this year's data, fewer than half of programs currently have in place.

You tend to find less when you are not looking for it.

Barriers to Success

Data quality or quantity has climbed to the top barrier this year at 50%, a notable shift from prior years, where it ranked lower: 41% in 2025, 44% in 2024, and 34% in 2023. The direction of that trend suggests that the volume of data flowing into hunting programs is now creating as many problems as it solves, with normalization and standards lagging behind collection, a theme that surfaced repeatedly in the methodology free-text responses discussed earlier.

Skilled staff remains a persistent second at 45%, though it has eased from a peak of 73% in 2023 and 61% in 2025, which signals progress, but not resolution. That resolution matters because a skilled-staff shortfall does not just mean fewer hunts get run, it means the hunts that do get run lean more heavily on whoever happens to be available rather than whoever is best placed to construct the hypothesis, which is exactly the ad hoc drift the methodology section flagged earlier. Budget constraints (42%), lack of data standards (39%), tool limitations (37%), and lack of defined processes (36%) round out the remaining barriers, and their stability across all four years of comparable data suggests these are structural costs of running a hunting program rather than problems any one organization is likely to solve outright.


On planned improvements, more internal investigative staff tops the list at 44%, which tracks directly with the staffing barrier and continues a multiyear upward trend from 39% in 2024. Improved contextual awareness from data and tools ranks second at 43%, while incorporating AI and ML ranks third at 39%—down from 48% in 2025 and 47% in 2024. That two-year decline is worth noting directly: Going into this survey, the expectation was that interest in AI-assisted hunting would continue its upward climb. The data says otherwise. Intent appears to be plateauing, which could mean organizations are moving from planning to actual implementation or that a more realistic assessment of what these tools deliver day-to-day is starting to take hold. Either way, threat hunting remains stubbornly dependent on human judgment for now and cannot be wholly AI-ed and ML-ed away by the bots.

Data quality or quantity has climbed to the top barrier this year at 50%, suggesting the volume of data flowing into hunting programs is now creating as many problems as it solves, with normalization and standards lagging behind collection.

The Investment Horizon

Investment intentions for the next 24 months show a more balanced picture between staffing and tooling than prior years, where staffing tended to dominate the conversation. On the staffing side, 26% plan a 10% increase and 19% plan a 25% increase, with 29% reporting no planned change—down from 40% in 2025. This is an encouraging shift because it suggests more programs are moving off the fence (see Figure 10). Planned decreases remain a small minority across both categories.

Q: Does your organization plan to change its investment in the staffing for threat hunting in the next 24 months?
Estimate how much the change in investment might be.



Tooling investment follows a similar pattern, with 33% planning a 10% increase (see Figure 11). Tool investment trends are consistent in comparison to 2025 numbers. The closer alignment between staff and tool investment intentions this year is a positive sign. Historically, organizations have tended to prioritize one over the other, often acquiring tooling without the skilled people to drive it, or building headcount without the platforms to support them effectively. A program that grows both in parallel is generally better positioned. Experienced hunters with capable tooling tend to produce outcomes more meaningful than either element alone.

Q: Does your organization plan to change its investment in the tools for threat hunting in the next 24 months?
Estimate how much the change in investment might be.




Figure 11. Planned Tool Investment Changes

Conclusion

The 2026 data reflects a community that is consolidating rather than racing forward. Programs are more experienced than ever with over 80% of respondents reporting they have been hunting for at least two years. However, the friction points that have appeared in every edition of this survey remain stubbornly in place. Data quality has overtaken skilled staff as the top barrier for the first time, which is a meaningful shift because it suggests organizations are no longer struggling primarily to find people who can hunt, but to give them data worth hunting for. Cloud remains the hardest terrain by a clear margin, improving only incrementally while adoption outpaces hunting capability. And nearly half of all programs still have no formal way to measure whether any of this is working. On the adversary side, living-off-the-land techniques dominate across all threat actor categories, and the data confirms this. Hunting for these actors with signature-based or indicator-driven approaches is optimistic at best.

Looking ahead, the AI and ML story is more measured than prior years suggested it would be. Intent to incorporate these tools has declined two years running, which likely reflects organizations moving from aspiration to the harder work of actual implementation. The teams describing agentic hunting frameworks in their free-text responses are the early signals of where this goes next. For now, the practical gap most worth closing is the measurement one: If your program cannot demonstrate what it finds and what that prevents, the investment case for everything else in this report becomes much harder to make.

Data quality has overtaken skilled staff as the top barrier for the first time, a shift that suggests organizations are no longer struggling to find people who can hunt, but to give them data worth hunting for.

Respondent Demographics

The 2026 SANS Threat Hunting Survey received 500 qualified responses from security professionals. Figure 12 provides the demographic details.




Figure 12. 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.