



Relentless Defense

A global report about cybersecurity in the AI era

Contents

- 1 Executive summary
- 2 Why now?
- 3 How well organizations actually defend
- 4 What's getting in the way
- 5 Seeing what you need to defend
- 6 How fast organizations actually move
- 7 Where AI is already working
- 8 What Relentless Defenders do differently
- 9 What needs to change

Executive summary

We asked 8,000 security leaders from across the globe how they're coping with a threat landscape being reshaped by AI, including whether their processes can keep pace with AI's ability to surface thousands of vulnerabilities at once, and whether they're confident staying ahead of the volume of new threats being discovered.

Less than one in ten are confident they can. And they told us that the challenge isn't coming from a lack of technology. Instead, friction — how well organizations are minimizing internal barriers so security teams can handle change — is now the culprit.

They also told us the problem is only becoming more urgent. Nearly all (91%) of the security leaders we surveyed experienced at least one material disruptive cyber incident in the past year. More than a third (35%) involved an AI-related attack. At the same time, more than a quarter (26%) of respondents state that AI is already significantly reducing operational burden in their security organizations.

As AI becomes critical to defense and a powerful weapon for attackers, success will need more than the best tools and infrastructure. Success will need organizations that can collaborate to prioritize a unified defense.

1

2

3

4

5

6

7

8

9



“

Democratizing AI — enabling every person and organization to harness its potential — demands that we build trust into its foundation. That trust is only realized when security and resilience become inseparable. You cannot govern what you cannot see, and you cannot trust what you cannot govern. This unified approach to AI trust is the only way for organizations and consumers to fully embrace the promise of this technology.



Sunil Potti, SVP and General Manager, Security, Observability, and Data Platform, Cisco

1

2

3

4

5

6

7

8

9



Clearer communication and defined roles would have made the biggest difference. During the incident, there was confusion about who was responsible for key decisions and delays in sharing critical information.

Chief Information Security Officer, USA,
survey respondent

The best performing security teams act as one

Our research shows a small group of organizations significantly outperform the rest by breaking down silos and working as a single, seamless unit with shared visibility. They can act faster, are more than twice as likely to rapidly turn approval into action, and their coverage is more comprehensive. For them, security isn't a cost center that protects business value; it's an integral part of the organization and meaningfully contributes to building that value.

They're the Relentless Defenders — but only 8% make the cut

This group moved beyond established security practices and are continuously optimizing their systems to defend against an ever-changing threat landscape.

They aren't necessarily the biggest spenders; they're working smarter, and prioritizing budgets differently.

And that's something any organization can do.

1

2

3

4

5

6

7

8

9

Why now?

The cost of getting it wrong is climbing. Splunk's 2026 report, [The Hidden Costs of Downtime](#), found that the Global 2000's aggregate downtime costs have reached \$600 billion annually, up 50% in just two years — with regulatory fines and ransomware payouts among the fastest-growing line items, and nearly a third of all downtime now cybersecurity-related. Finance leaders feel the shock as clearly as security teams feel the breach.

Security has become a shared concern across infrastructure, finance, and the C-suite; no longer a line item confined to the Security Operations Center (SOC). Meanwhile, the ground is shifting: frontier AI models can now discover software vulnerabilities at a scale and speed no human team alone can match. The tools available to attackers and defenders alike are changing faster than most organizations can absorb.

Everyone is worried about security, but security alone can't fix it

Every layer of the enterprise is dominated by concerns about security. The boardroom is no longer discussing whether AI will work. It's asking whether the organization is secure enough to let it. So, while most CEOs are more optimistic about AI than they were a year ago and worried about underinvesting, they also ranked infrastructure and agentic AI security among their top priorities for 2026. And most expect security risks to keep rising as AI adoption expands, believing threats are evolving faster than existing controls can adapt ([How CEOs see AI in 2026, Cisco](#)).

If concern about security runs through the entire organization, the support needed to actually defend the organization should be a combined effort. But, for some reason, it isn't.

①

②

③

④

⑤

⑥

⑦

⑧

⑨



Boards have stopped asking whether AI will deliver. They are asking whether their organization is secure enough to use it. Regulation accelerated that shift, but it will not finish it. The organizations treating security as a compliance exercise will struggle. Those building it into how they operate are already ahead.



Chintan Patel, Cisco Chief Technology Officer, EMEA

①

②

③

④

⑤

⑥

⑦

⑧

⑨

①

②

③

④

⑤

⑥

⑦

⑧

⑨

We asked 8,000 security professionals to tell us in their own words what would have made the biggest difference in how their team responded during a recent security incident or near-miss. Their replies show us that the friction slowing security teams down rarely comes from their technical capabilities alone, but from within their organization – procurement delays, infrastructure decisions IT owns, and priorities the C-suite sets elsewhere.

Another big recurring factor is cross-department coordination: practitioners are asking for “centralized communication protocols”, “clear, pre-defined communication channels”, “better cross-departmental communication”, and “clear ownership and faster communication”. Security teams, alone, can’t solve this.

Could it be that the functions voicing the loudest concern about security are the ones standing in the way?

How well organizations actually defend

The average organization scored 64 out of 100

We scored organizations on three essential pillars for maintaining continuous defense against machine-speed attacks, and weighted the scores on each based on the problems practitioners told us about:

● Coverage ● Speed ● Friction

Are there any gaps in what you're protecting?

25%

How fast can you deploy new security controls? How fast can you respond when something goes wrong?

25%

How well are internal barriers minimized to enable your teams to handle change?

50%

Friction was double-weighted because, when asked to describe in their own words what would have made the biggest difference in how their team responded during a recent incident or near-miss, six in ten practitioners flagged organizational rather than technical blockers. When asked what they would change if they could change one thing about how their organization approaches cybersecurity today, roughly the same amount focused on organizational changes.

①

②

③

④

⑤

⑥

⑦

⑧

⑨

From there, we calculated a total defense score for each organization, and used that score to cluster them into four groups.

Only 8% qualified as **Relentless Defenders** — the top segment — achieving the coverage, response speed, and organizational alignment needed to react at the pace agentic AI is setting. Everyone else clusters in the middle, with challenges in one or more pillars. Throughout the rest of this report, we will focus on the differences between the Relentless Defenders and everyone else, categorizing Established, Reactive, and Exposed Defenders as the “Other Defenders”.

Relentless Defenders

High coverage, high speed, and low friction are all in place. These organizations continuously optimize their defenses rather than settling for “good enough.”

8%

Established Defenders

Strong foundations and the ability to recover — but coordination and approval slow them down at the exact moment speed matters most.

39%

Reactive Defenders

Capability exists, but only in parts of the organization rather than across it. Results depend on which team, or which system, happens to be involved.

41%

Exposed Defenders

Still building the basics. Coverage has gaps, changes take too long to land, and the organization typically moves after an incident, rather than ahead of one.

12%

①

②

③

④

⑤

⑥

⑦

⑧

⑨

What's getting in the way

Friction gets in the way even when teams have the tools and the ability to act. The blocker is organizational.

We asked what caused the most significant problems in the hours immediately after an incident.

Well over half don't think their teams collaborate seamlessly

Security practitioners already have a lot of stakeholders, and a need to act quickly. This challenge only grows after an incident, when response coordination across teams and partners breaks down, and decisions from senior leaders don't come fast enough.





The biggest difference would have been a clearer escalation path. During the incident, there was confusion about who had the final authority to shut down the affected systems.

Chief Security Officer, India,
survey respondent

You can't influence without understanding

The security team can't fix ownership alone. Other teams must understand the problem. But the metrics security teams use, if there are any at all, are hard to translate to the rest of the organization.

More than four in ten (43%) surveyed practitioners say what they are tracking is too technical to influence executive decision-makers, while over a third (36%) struggle to translate that data into clear business impact.

For over a quarter (27%), even that problem would be nice to have, but they simply don't have a consistent way to measure performance accurately at all. If it can't be measured, how can you explain it to someone? Without this, securing investment becomes harder and the gap doesn't close.

①

②

③

④

⑤

⑥

⑦

⑧

⑨



Fast, clear communication combined with predefined response procedures enabled the team to make timely decisions and reduce the risk of recurrence.

Director of Security, Thailand,
survey respondent

Wrangling data instead of chasing threats

Underneath both problems is the same missing ingredient: data the team can actually access. Four in ten (40%) security teams say they spend more time manually collecting and correlating data from different systems than they spend dealing with and chasing the threat. Nearly four in ten (39%) say critical insights are getting missed because the data lives somewhere they can't reach.

Which raises the question: if the data is obscured, how much of the environment can security teams really see?

①

②

③

④

⑤

⑥

⑦

⑧

⑨

Seeing what you need to defend

Many teams are defending an environment they can't fully see. For Relentless Defenders, data management is generally more unified. Almost all (98%) state that they have a single, real-time view of their security data, compared with 58% globally. This also translates into how much of their organization they are actively protecting. More than half (51%) of Relentless Defenders report having active security protection deployed across at least 95% of endpoints (laptops, mobile devices, servers and IoT devices), while 20% of the global average meet that same threshold.

AI agents risk becoming a bigger blind spot

The gap extends to non-human identities: 86% of Relentless Defenders have comprehensive controls consistently applied to AI agents, service accounts and automated systems, compared with 41% of all global respondents. And with virtually every organization we surveyed (98%) deploying or planning to deploy AI agents in their security operations, resolving this visibility gap for non-human identities becomes urgent.

①

②

③

④

⑤

⑥

⑦

⑧

⑨

How fast organizations actually move

Only about one in five organizations (21%), globally, say they can fully stand up a new control within six months, even after budget and approval have already cleared. Relentless Defenders are more than twice as likely to have new controls fully operational within the same time. They are also more than twice as likely to run continuous automated threat detection as their primary approach.

Overall, 14% told us they can update controls in real time when a new threat is detected; another quarter (26%) within 24 hours. This leaves around 60% who take more than a day when it matters most. In comparison, 45% of Relentless Defenders can update in real time, and 79% can update within 24 hours.




During a recent phishing near-miss, our team's delayed detection stemmed from fragmented alerting and unclear escalation paths. A centralized real-time alert system paired with pre-defined incident response playbooks would have made the biggest difference.

Chief Security Officer, Sweden,
survey respondent

- ①
- ②
- ③
- ④
- ⑤
- ⑥
- ⑦
- ⑧
- ⑨

How long it takes to implement new security systems



- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9

Where AI is already working

The advantage is moving to the defenders

Defenders are already moving faster with AI, improving how quickly they detect, respond and remediate threats.

But speed isn't the only return. Nearly nine in ten also say AI has reduced the operational burden on their security teams, which, for teams that have been running at capacity for years, is worth as much as the minutes saved.

The advantage is moving to the defenders

AI has improved threat detection speed

By at least 25%



87%

By at least 50%



50%

AI has improved response and remediation speed

By at least 25%



85%

By at least 50%



50%

①

②

③

④


⑤

⑥

⑦

⑧

⑨



AI related security concerns hold back deployments

Fewer than half of organizations have extensively deployed AI-powered capabilities in any single domain. The highest is in network security, the lowest is security operations (investigation), meaning the domain closest to the customer is the one AI is trusted the least with. Just a third have extensively deployed AI agents specifically within security operations, while almost half have them in key areas only, and the rest are still in pilot or planning stages.



AI adoption in the network was inevitable because human reaction time couldn't keep pace with packet-level threats years ago. The network became machine-speed by necessity. Security operations is now hitting that exact same inflection point. The organizations pulling ahead aren't treating AI as a separate chatbot in the SOC; they are treating the network fabric as the sensor and enforcement engine for the entire security stack.



Carl Solder, Cisco Chief Technology Officer, Australia and New Zealand

①

②

③

④

⑤

⑥

⑦

⑧

⑨

Nearly all cite at least one significant AI-related security concern

Across the board, many are holding back for real reasons. Nearly all organizations cite at least one significant AI-related security concern. The top three most common cited fears being sensitive data exposure through AI usage (41%), the worry that the AI systems become compromised (40%) and security and data handling by third party providers (37%).

Lack of governance is in the way

Just over half (54%) are very confident in their ability to detect and respond to a compromised AI agent or autonomous system — meaning almost half are only somewhat confident and would struggle to catch it if their own AI turned into the threat.

There's real concern about delegating security-sensitive actions to AI agents without enough oversight, and most believe new policies and controls are needed to govern what AI agents can access and do. 80% believe agentic AI will enable attackers to run more persistent, adaptive campaigns with less effort.

As one retail security executive put it in [No time to wait: The accelerating impact of AI on campus and branch networks](#), the difficulty isn't the concept, it's the sheer variety: "it's hard to create the guardrails for every possible AI tool that your organization must use."

A VP of IT and digital infrastructure in UK education summed up the mood plainly: "We're just playing catch-up at the moment... it's a worrying time, and I think it'll stay like this for another 18 months or two years."



1

2

3

4

5

6

7

8

9



Response-process readiness, automated alerting and real-time data helped us detect and contain the incident, while close coordination and strong specialist skills helped us make the right decisions.

Director of Information Security, Vietnam,
survey respondent

You can't scale what you can't control

While only 54% of the total sample feel confident they could detect and respond to a compromised AI agent, the Relentless Defenders have built a different reality. They didn't scale AI because they were less worried about the risks; they scaled it because they were more prepared for them.

Raising the ceiling

Relentless Defenders are four times more likely to report threat-detection speed improvements of over 100%. This massive leap isn't because they have "better AI". It is because they have done the hard work of governing identities and unifying data.

AI is already closing the coverage and speed gap and keeping the advantage with the defenders. The constraint isn't the capability of the machine; it is how prepared the organization is to let the machine work. For those who build the guardrails, the ceiling for what AI can achieve is still rising.

①

②

③

④

⑤


⑥

⑦

⑧


⑨

Relentless Defenders did the governance work first, then put AI to work



What Relentless Defenders do differently

The gap isn't closed with bigger budgets. Only two in five (41%) organizations that increased security spending saw fewer incidents. Almost half (45%) saw incidents rise anyway or remained unchanged despite their increased investments. Meanwhile, more than seven in ten (71%) Relentless Defenders reported decreased incidents stemming from their increased budgets, suggesting it isn't the size of their investment that's paying off, but what they are investing in.



Technology is the effect, not the cause

Every capability that follows, from AI governance and identity controls to visibility and vulnerability management, depends on the same thing: a team that's aligned enough to deploy it consistently, and reducing friction enough to keep it that way. A unified team doesn't just work better together; it's also the reason the technology gets configured correctly, rolled out everywhere it's needed, and kept up to date. Friction doesn't just slow people down. It's the reason coverage has gaps, controls fall out of date, and tools never reach their full potential. Remove the friction, and the technology finally gets the chance to do its job.

①

②

③

④

⑤

⑥

⑦

⑧

⑨

Relentless Defenders pull ahead of their peers

● Relentless Defenders
 ● Everyone Else

1. AI Governance



A coordinated team is what makes AI governance possible in the first place: someone has to own the policy, monitor it, and keep it current across every agent and service account. Relentless Defenders apply comprehensive, consistent controls for AI agents and non-human identities across the organization, compared with far fewer Other Defenders.

2. Security controls across different environments



Consistency requires the same standard to be applied everywhere, by teams that are talking to each other. Relentless Defenders are nearly twice as likely as Other Defenders to have comprehensive security controls consistently applied across each environment, including data security & privacy, network security, and application security.

- ①
- ②
- ③
- ④
- ⑤
- ⑥
- ⑦
- ⑧
- ⑨

● Relentless Defenders
 ● Everyone Else

3. Core identity security controls across



Identity gaps usually aren't a technology failure. They're a coordination failure between who owns the tool and who manages the user. Relentless Defenders apply strong, consistent safeguards across groups such as contractors and third-party users, IT administrators, and developers with production access, at nearly double the rate of everyone else.

4. Visibility of anomalies and risks



Visibility depends on data actually reaching the people who need it. Relentless Defenders have better oversight and response across key areas. This suggests that top performers don't just have better tools, but have removed the barriers that stop those tools from being seen and used.

- ①
- ②
- ③
- ④
- ⑤
- ⑥
- ⑦
- ⑧
- ⑨

● Relentless Defenders ● Everyone Else

5. Critical and high severity vulnerability management



Finding a vulnerability and fixing it in time are two different challenges. And it's the second one that's almost entirely organizational: who owns it, who prioritizes it, and how quickly a decision gets made. Relentless Defenders report strong confidence across core capabilities: prioritizing vulnerabilities, remediating them fast enough, and tracking ownership.

1

2

3

4

5

6

7

8

9

● Relentless Defenders
 ● Everyone Else

6. Operational level

The technology gains above don't happen in isolation; they happen because these organizations operate differently:



operate as a single coordinated security team with shared visibility and tools across security, network and IT.



are investing heavily in upskilling existing security team members.

Taken together, these figures paint a consistent picture: Relentless Defenders aren't excelling in one area, but building strength across identity, visibility, and vulnerability management, all made possible by the same underlying advantage: a team working as one, with nothing standing in its way.



AI is forcing technical concepts across boundaries that didn't used to need crossing: technology, security, HR, legal, ethics. This report makes it crystal clear that the organizations getting ahead are the ones breaking down organizational silos and removing barriers, so the right people can make the right decisions together in the first place.



Lauren Wilson, Field CTO & Strategic Advisor, Splunk, a Cisco company

①

②

③

④

⑤

⑥

⑦

⑧

⑨

What needs to change

What would fix this? Five moves, tied directly to what is holding back progress.

- 1. Fix ownership before fixing tools.**
Pre-delegate the decision rights before the incident happens, not during it. Ownership ambiguity is the single biggest reason things go wrong when it matters most.
- 2. Build one picture, not multiple.**
Unify, so the team is looking at one version of the truth. Fragmented visibility is the second-largest failure mode, nearly as common as not knowing who's in charge.
- 3. Practice, don't just document.**
A playbook that's never been run under pressure is just a document.
- 4. Let AI take the first ten minutes.**
Use constrained automation to handle the earliest, most mechanical steps before a human needs to get involved. This goal is cited by nearly a quarter and directly tied to the speed gaps in section:
[How fast organizations actually move.](#)
- 5. Design for the easiest path, not just the compliant one.**
People default to whatever's easiest, so make that the secure option too. If security creates friction for the people using it, they'll route around it. A quarter of organizations already cite this as a priority.

①

②

③

④

⑤

⑥

⑦

⑧

⑨



Our methodology

The findings draw on a survey of **8,000 security professionals** across 30 markets, 27% of them CISOs and 60% with six or more years at their current organization.

Performance is scored out of 100 across three pillars — coverage, speed and friction — each built from a set of indicators, with respondents grouped into four performance categories.

Friction carries 50% of the score, coverage and speed 25% each. That weighting comes from practitioners themselves: asked what they would change about their organization's approach and what would have made the biggest difference in their most recent incident, approximately six in ten named an organizational blocker rather than a technical one.

This report was created in partnership with independent third parties. The survey was conducted through a double-blind survey in April and May 2026 by Satori Experience. The analysis and report was completed by Sandpiper.

Americas Headquarters

Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters

Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters

Cisco Systems International BV Amsterdam
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners.
The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)