TEAM8™

AI, RISK, AND THE ROAD AHEAD

Key Findings from Team8's 2025 **CISO Village Survey**

July 2025

AUTHORS

Amir Zilberstein
Managing Partner, Team8

Liran Grinberg
Co-Founder, Managing Partner, Team8

Ori Barzilay
Partner, Team8

Noa Hen
Director of Strategy, Team8

The Team8 CISO Village is a community of CISOs from the world's leading enterprises. The primary focus of the Village is to facilitate collaboration among the world's most prominent companies with the goal of sharing information and ideas, conducting intimate discussions on industry and technology trends and needs, and generating value and business opportunities for all parties.

By helping Team8 to identify real pain points and understand the requirements of large organizations,members of the Village are first in line to leverage solutions that are purpose-built by Team8's portfolio companies to support their needs.

To contact the Team8 CISO Village, please email cisovillage@team8.vc

Table of Contents

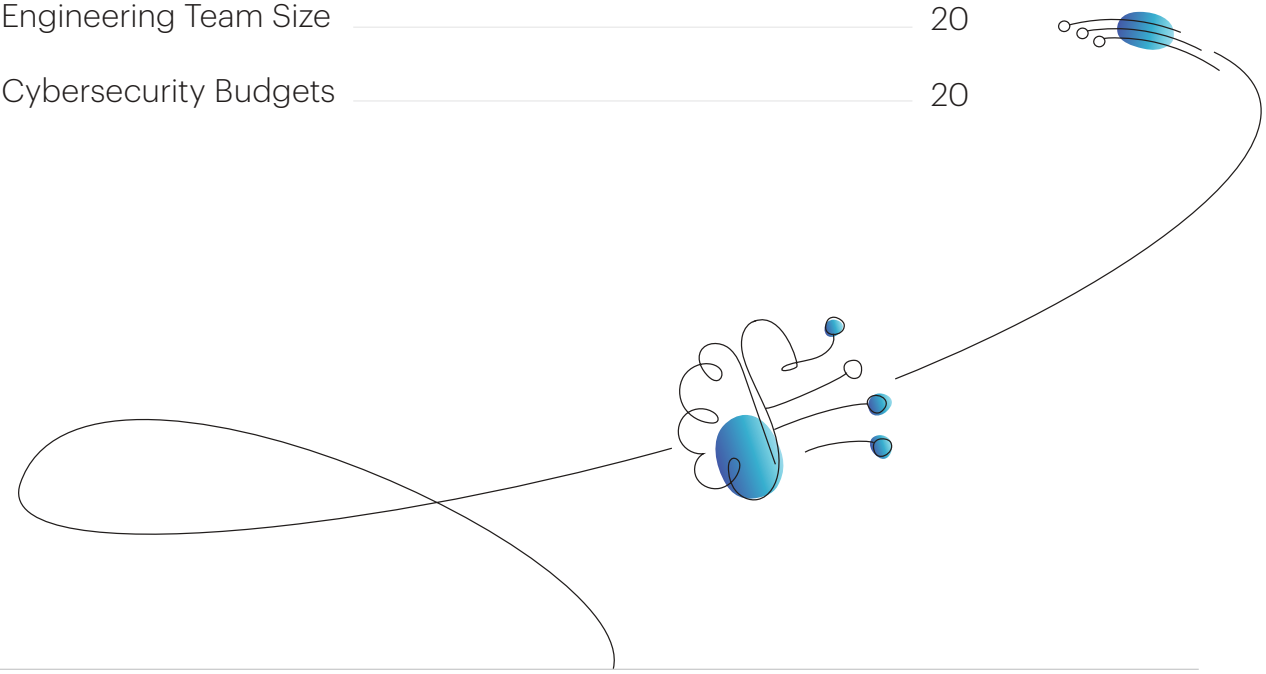

About This Survey and The **Team8 CISO Village**

Team8's CISO Village is a private, exclusive community of security executives from the world's most influential companies. Each year, we host a five-day summit where CISOs engage in off-the-record conversations, share best practices, and shape the future of cybersecurity innovation.

The insights in this report are drawn from our annual CISO Village Survey, completed by over 110 security leaders during the 2025 Summit. With respondents from GuideWire, Respol, Elastic, Centene, and other industry leaders, this survey offers a rare window into how the world's top security teams are thinking, reacting, and building.

Budgets & Buying Behavior

Cybersecurity budgets remain on a strong upward trajectory. This year, **52% reported a budget increase** – a signal of continued investment following an exceptionally high-growth year in 2024, where 70% saw gains.

That spike, in our view, represented a post-COVID correction; What we're seeing in 2025 is not a pullback, but rather a return to steady, strategic growth driven by long-term planning, operational pressure, and the rising complexity of the threat landscape.

Two forces are shaping this reality. On one side, **geopolitical conflict, AI-accelerated attacks, and enterprise-wide adoption of new technologies** are keeping cybersecurity high on the strategic agenda. On the other, **macroeconomic volatility and heightened AI-driven productivity expectations** are pushing CISOs to justify each new hire and tool. In this context, budget growth continues, but with a stronger focus on impact, coverage, and enablement.

A Majority of CISOs Report a Budget Increase in 2025

Decreased
11%
vs. 15% last year

Did not change
37%
vs. 15% last year

Increased
52%
vs. 70% last year

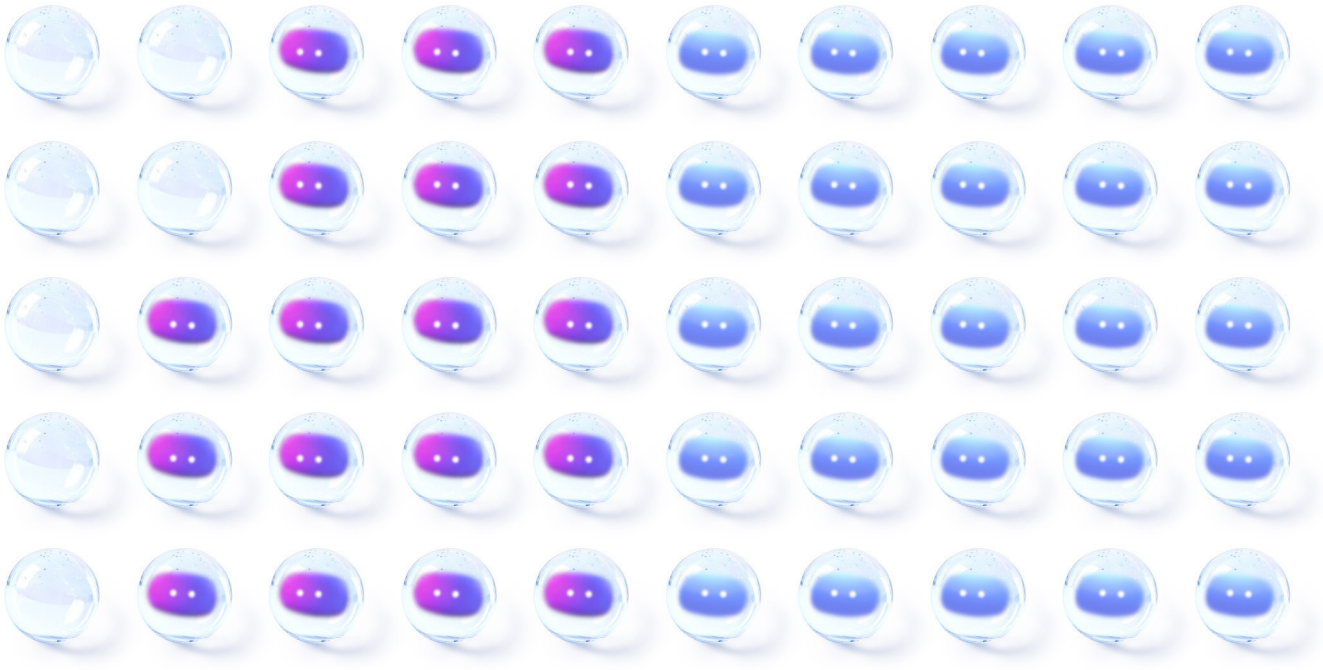

Best-of-Breed Makes a Comeback

potentially marking a turning point in the "platformization" wave that has shaped cybersecurity buying behavior in recent years. While platforms once promised simplicity and scale, many CISOs are finding that breadth often comes at the expense of depth. As IT environments grow more complex, with hybrid infrastructure, tool proliferation and AI implementation, **a checkbox capability on every front isn't enough**. Each attack surface demands a high-quality solution, and platforms frequently fall short when compared to focused, innovation-driven startups.

That doesn't mean CISOs are rejecting consolidation outright—**They're looking to spend smarter**. In today's ROI-driven climate, security leaders aren't just buying technology, they're buying outcomes. New products must deliver measurable value, reduce operational overhead, and align tightly with business needs. Tools that add complexity without tangible benefit are being questioned and often cut.

This mindset opens the door for a new generation of startups: **companies built not just to add security, but to optimize it**. Whether through eliminating redundancy, quantifying risk, or reshaping the stack around what actually works, the next wave of cybersecurity innovation may belong to those who make security simpler, leaner, and more effective.

60%

Prefer best-of-breed tools over best-of-suite platforms

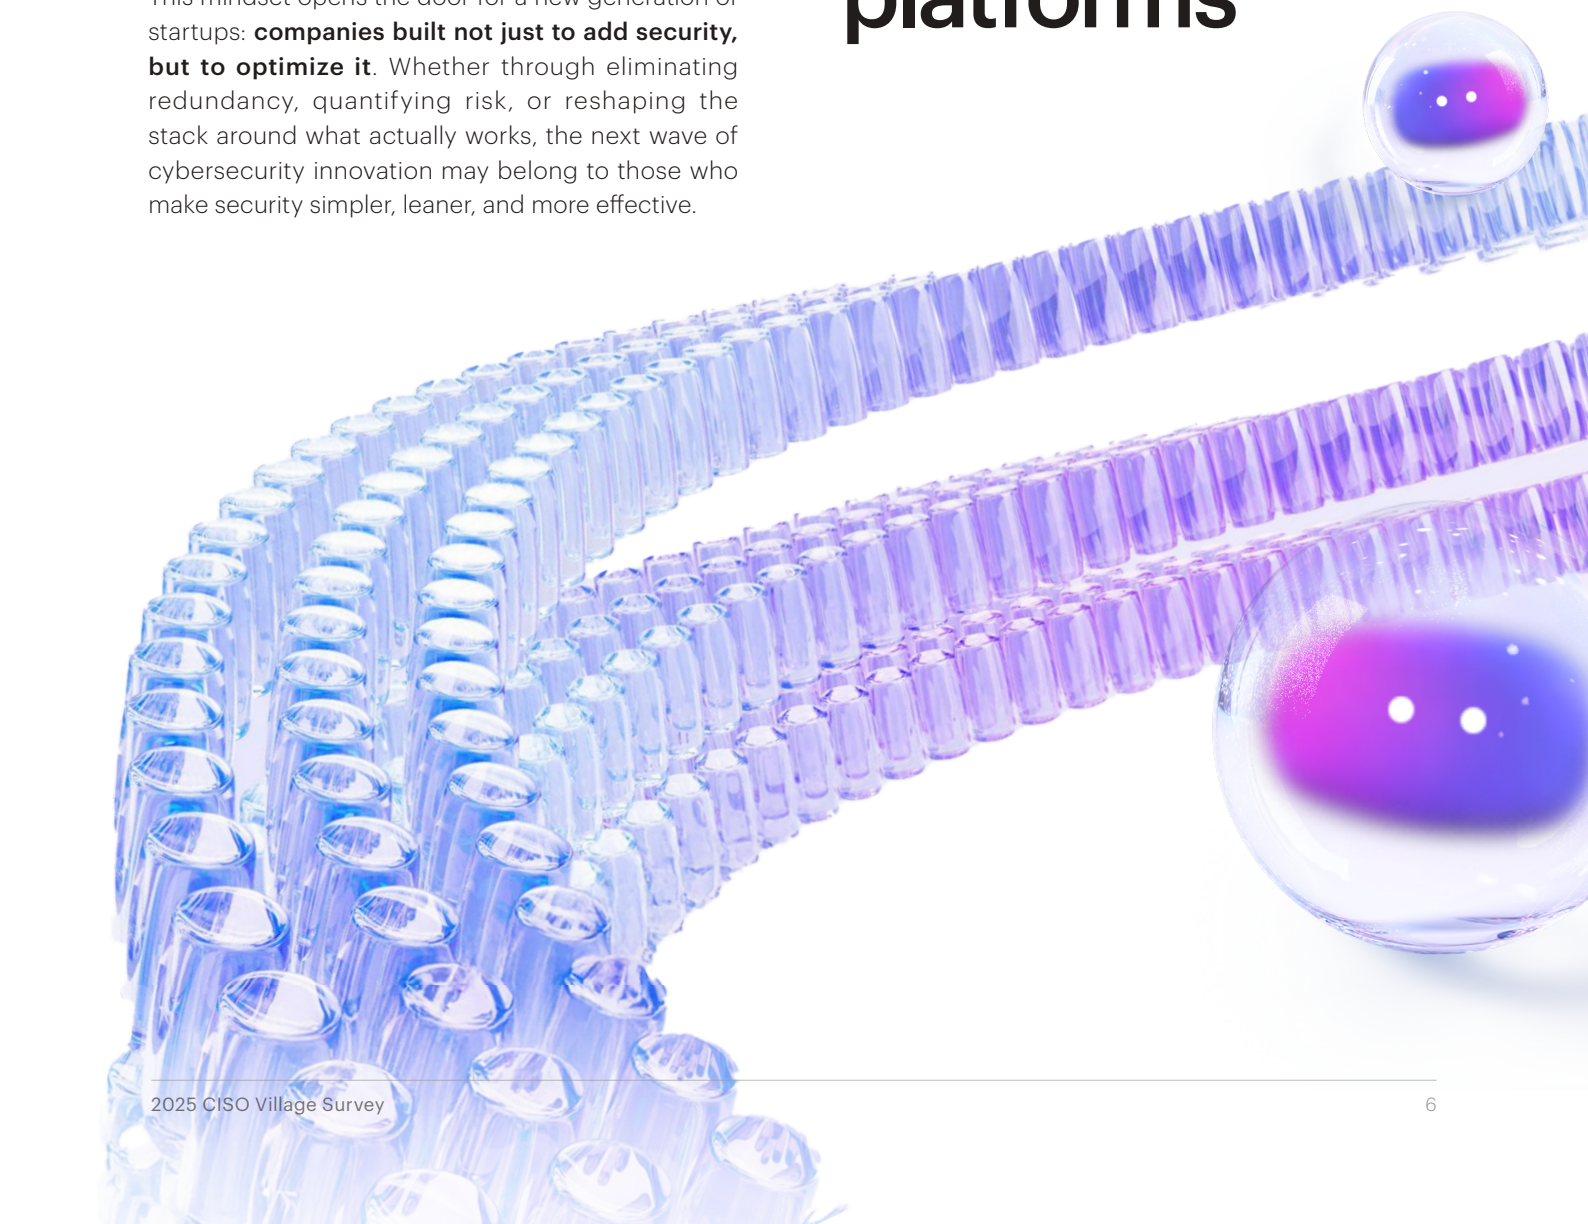

The AI Arms' Race

AI is no longer on the horizon, it's in the kill chain. We are witnessing a true arms race between attackers and defenders. For attackers, AI unlocks novel weapons like deepfakes and voice clones, while also accelerating traditional vectors through automation and scale. While combatting these new threats, defenders are also challenged with defending AI as a new attack surface, introducing new risks. At the same time, AI has become essential for surviving the velocity and scale of modern threats. It offers not only faster detection and response, but a chance to automate manual, resource-intensive processes in a field plagued by persistent talent shortages.

25%

One in four CISOs reported experiencing an AI-generated attack within the past 12 months

2.1 AI-Powered Attacks: The First Wave Has Already Hit

In our survey, **1 in 4 CISOs reported experiencing an AI-generated attack within the past year**. While significant, this number likely underrepresents the true threat landscape. Many AI-driven attacks are difficult to differentiate from traditional, human-led campaigns using existing detection and measurement tools. The most visible examples today remain **AI-generated social engineering attacks**: deepfakes, voice cloning, and real-time impersonation tactics designed to exploit human trust.

One notable example is the recent crackdown on a **North Korean scheme wherein operatives posed as remote IT workers**, infiltrating nearly 100 U.S. companies, including a defense contractor.

But it doesn't stop with deepfakes. Offensively, AI can also be used as a powerful automation tool that augments and accelerates attacks to unprecedented rates. This was initially enabled by **malicious AI copilots** like WormGPT and EvilGPT. These are fine-tuned, open-source LLMs stripped of guardrails that are sold on the darknet. These tools allow even low-skill actors to create sophisticated phishing campaigns and develop obfuscated malware with just a prompt.

Today, these "evil twins" of ChatGPT operate as malicious helpers, but soon they may evolve into **autonomous weapons**. Much like their enterprise counterparts, attackers are now exploring **agentic AI systems**—AI that can plan, decide, and act without human intervention. When this shift happens at scale, we'll move from tool-assisted attacks to autonomous threat execution.

And the future becomes even more complex when AI starts finding the vulnerabilities to exploit on its own. Both **Google and academic researchers have recently demonstrated that large models can detect zero-day vulnerabilities in production code at scale**. For attackers, this is a goldmine: a way to discover and weaponize flaws faster than defenders can identify or patch them. It could compress the zero-day window from months to days and leave enterprises on the back foot.

We are only at the beginning. The AI arms race is accelerating, and the only thing more dangerous than falling behind is standing still.

> *"At Elastic, we've already seen the impact of AI-driven attacks—especially sophisticated phishing and deepfake campaigns that are incredibly convincing and hard to catch. These aren't just email scams anymore. They show up across business communication channels, are tailored to specific individuals, and are designed to fool even well-trained employees. It's clear we're in the early stages of an AI arms race, and right now, the attackers moved first and have the edge. One of our biggest priorities moving into 2026 is rethinking how we defend against this new wave—hardening our detection capabilities and doubling down on making the human layer more resilient. I suspect what we're seeing now is just the beginning. The speed and scale at which AI can be used to automate and amplify attacks is unlike anything we've faced before."*
>
> **Mandy Andress,** CISO
>
> elastic

2.2. Securing Enterprise AI Adoption is The CISO's Top Priority for 2025

In 2025, AI risk has become the defining security challenge for CISOs, outpacing long-standing concerns like vulnerability management, data loss prevention (DLP), and third-party risk. Within this new category, two specific issues rise to the top: securing AI agents (37%) and governing employee use of AI tools (36%).

This shift is hardly surprising. AI introduces an entirely new attack surface, driven by natural language inputs, non-deterministic reasoning, and unprecedented risks introduced by agents. Boards are pushing aggressively for enterprise-wide adoption, and security leaders are expected to enable, not block, this transition. That puts CISOs in the hot seat: charged with mitigating risk in a technology domain that's still poorly understood, moving fast, and lacking mature controls.

While widespread AI-native breaches have yet to materialize in the wild, Team8's CISO Village agrees that they're coming. In the meantime, secure enablement, visibility, and monitoring are emerging as top priorities in every CISO's AI strategy.

2.2.1. AI Agents - Securing The Age of Autonomy

AI agents are not chatbots. They are autonomous software entities that can perceive, reason, and act across enterprise systems, often with read and write access to sensitive environments. Unlike copilots, agents do not just assist; they execute, making decisions and orchestrating tasks without step-by-step instructions.

This shift introduces a new kind of risk. **Adversarial attacks** like prompt injection or tool hijacking can trick agents into harmful behavior. **Misalignment and reasoning flaws** mean agents might take the wrong actions on their own, causing widespread damage to enterprise environments even without an attacker. In both cases, the danger is not what the agent says – it is what it does. And in complex environments, that can mean business disruption, privilege escalation, or even lateral movement.

AI agents are being adopted at an unprecedented pace. **67% of enterprises are deploying agents in 2025**, with another **23% planning to follow in 2026**. Only 9% of CISOs say their organization has no plans to introduce agents. This signals a major architectural shift underway across the enterprise stack.

One of the early signals of this shift was the proliferation of tools embedding agentic capabilities, including Salesforce Agentforce, Zapier Agents, and Zendesk. However, our survey reveals that this is not just a vendor-driven trend. According to our Villagers,**67% of enterprises deploying agents are building them in-house**. This likely reflects both the **need for deep customization and the productivity gains offered by developer-facing agent tools** like Winsurf and Cursor. Rather than waiting for fully packaged solutions, enterprises are assembling tailored agent stacks designed to fit their environments and objectives. At the same time, **59% of enterprises are adopting pre-packaged SaaS agents**, indicating that while the build-versus-buy debate is shifting, most organizations are hedging their bets with a hybrid approach.

But while adoption is widespread, deployment maturity remains early. Many agents today operate in sandboxed environments or support only narrow tasks. They often have limited write permissions and rely on **human-in-the-loop guardrails**. The real shift will come when agents are trusted with more autonomy, evolving from task executors into virtual teammates: systems that can optimize for business outcomes, collaborate across channels, and prioritize their own work queues. According to several leading AI labs, this level of autonomous agent capability may arrive as early as 2027.

67% of enterprises are deploying agents in 2025

23% planning to follow in 2026

67%
Deploying Agents in 2025

23%
Plan to introduce agents in 2026

9%
No plan to introduce AI agents

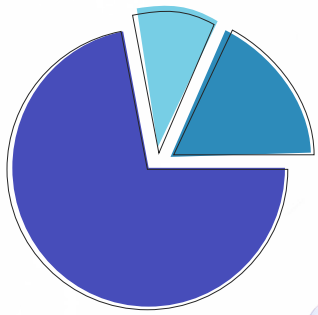

Enterprises Are Developing, Buying and Productizing Agents

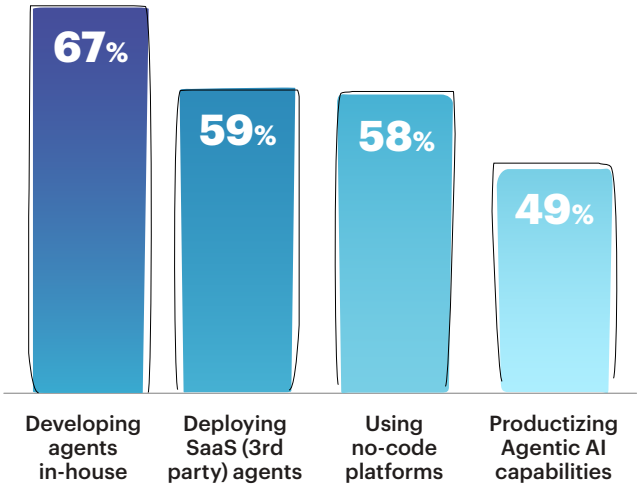

2.2.2. Securing Employee AI Usage

While AI agents are emerging as a new attack surface, **the risk from employee use of AI tools is not fading, it is expanding rapidly and is here to stay**. With the explosion of LLM-powered SaaS products, most employees now interact with AI in their day-to-day work—whether through writing assistants, customer service bots, auto-generated meeting summaries, or code suggestions. This sprawl introduces a persistent and compounding challenge: **shadow AI**.

Shadow AI refers to tools that enter the organization through individual teams or employees without going through IT or security review. These tools often come with unclear policies, opaque model behaviors, and limited safeguards. The risks are significant: employees may unknowingly leak sensitive data into external models, or receive outputs based on internal data that violate access boundaries. In more extreme cases, models can produce harmful or policy-violating responses, exposing the company to reputational and compliance risks.

Unsurprisingly, **securing workforce AI usage was the second most prioritized concern among CISOs**, cited by 36% of respondents.

And that's because **very few organizations are getting it right**. CISOs today face a lose-lose tradeoff: **either restrict access and stifle innovation, or allow usage without controls and accept unmanaged risk.**

As shown below, **48% of enterprises take a restrictive approach**, limiting usage to an allow-list of approved tools. These approaches reduce risk, but at the cost of agility, productivity, and business buy-in.

On the other side, **over 30% allow AI usage with little or no monitoring**, creating a wide and largely invisible attack surface. For most enterprises, the tooling to enable safe-by-default AI usage simply doesn't exist yet. Until it does, **CISOs are left choosing between innovation and control**, when they urgently need both.

The real need and opportunity is for **solutions that can provide meaningful controls without breaking the user experience**. CISOs want to move toward allow-by-default models that are safe, observable, and policy-compliant. Today, most are still waiting for the tools that will let them get there.

How would you describe your organization's current policy on employee use of generative AI tools?

43.4%
Restricted to approved tools

3.5%
All external AI tools are strictly blocked

22.1%
Use allowed and monitored

25.7%
Use allowed with usage policies, but not monitored

5.3%
Use allowed, no formal policy yet and no monitoring

2.2.3. AI for Security

Agentic AI introduces a new possibility: for the first time, machines can begin to **mimic human security work**, not just assist it. That shift couldn't come at a more urgent moment. With **30% of North America's security roles unfilled in 2023** (according to the World Economic Forum) and a persistent global skills gap, CISOs are under immense pressure to scale security operations without growing their teams.

The use cases where CISOs expect AI to replace human labor fall into two broad patterns. The first involves **high-volume, process-heavy tasks** that are essential but burdensome, and often bottlenecked by scarce human time. At the top of this list is the SOC. **77% of CISOs believe SOC analyst responsibilities will be the first to be transformed by AI**, a clear reflection of how overloaded these teams are with alert triage, enrichment, and case management. Just behind the SOC are **third-party risk management (27%)** and **provisioning and access reviews (26%).** Both domains are riddled with form-filling, checklist-based evaluations, and recurring approval flows. These tasks follow structured logic and consistent policies, making them ideal candidates for agentic systems that can execute autonomously at scale.

But CISOs are also eyeing AI for roles that have historically been constrained not by volume, but by expertise. **Pentesting (27%) and threat modeling (22%)** are examples of high-skill functions that require security judgment, contextual analysis, and deep technical understanding. Such resources were previously extremely difficult to hire and scale, but agentic AI could unlock expert-level capabilities across a broader surface area, more frequently and more consistently than ever before.

What are the first use-cases where you expect agentic AI to replace human labor?

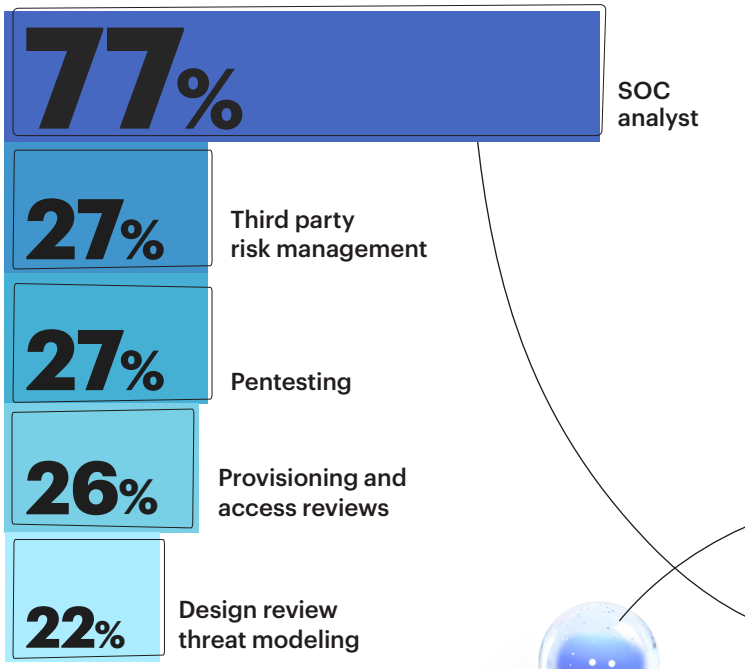

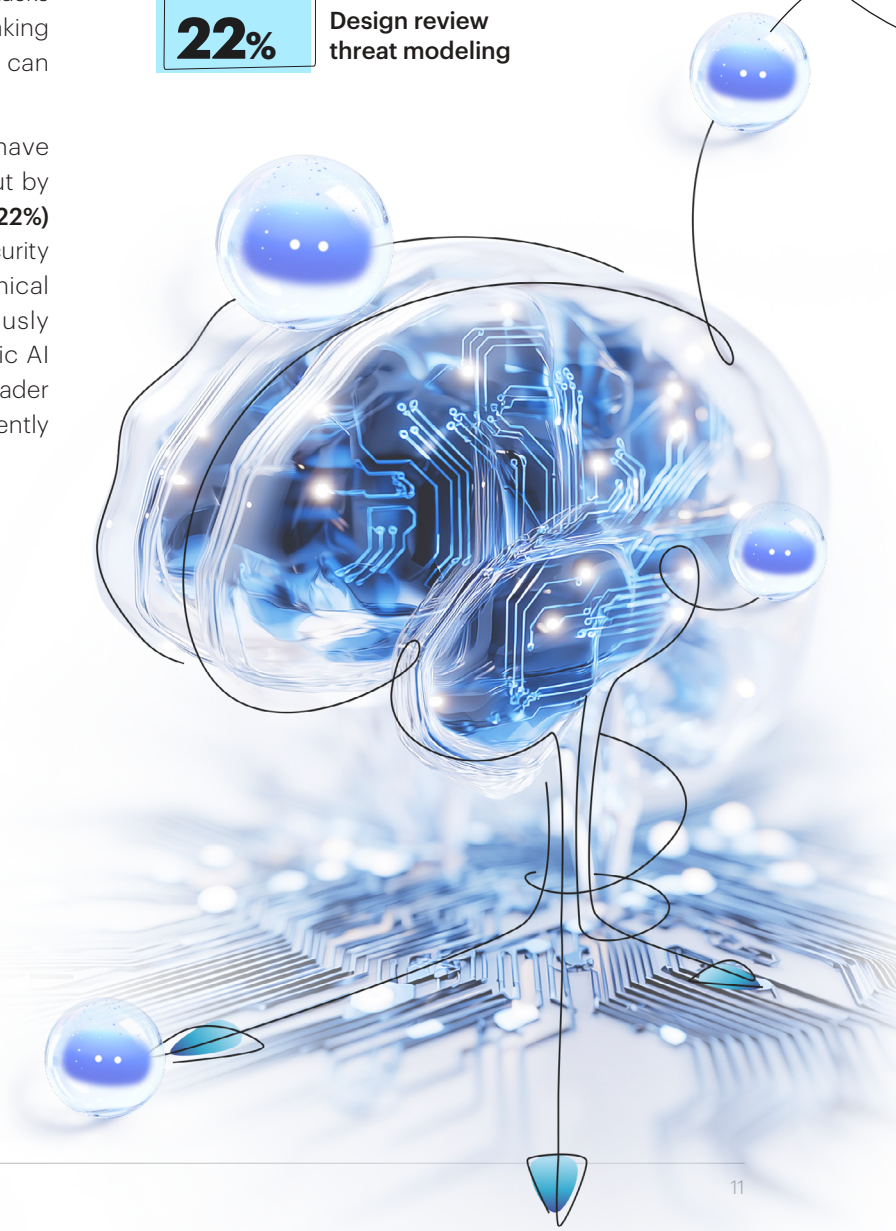

2025 Top CISO Pain Points

Each year, we ask CISOs to name the cybersecurity challenges that remain unsolved, where existing tools fall short and where they're actively seeking new, disruptive solutions. These responses reflect the areas where security leaders believe the market is still open, fragmented, or failing to deliver.

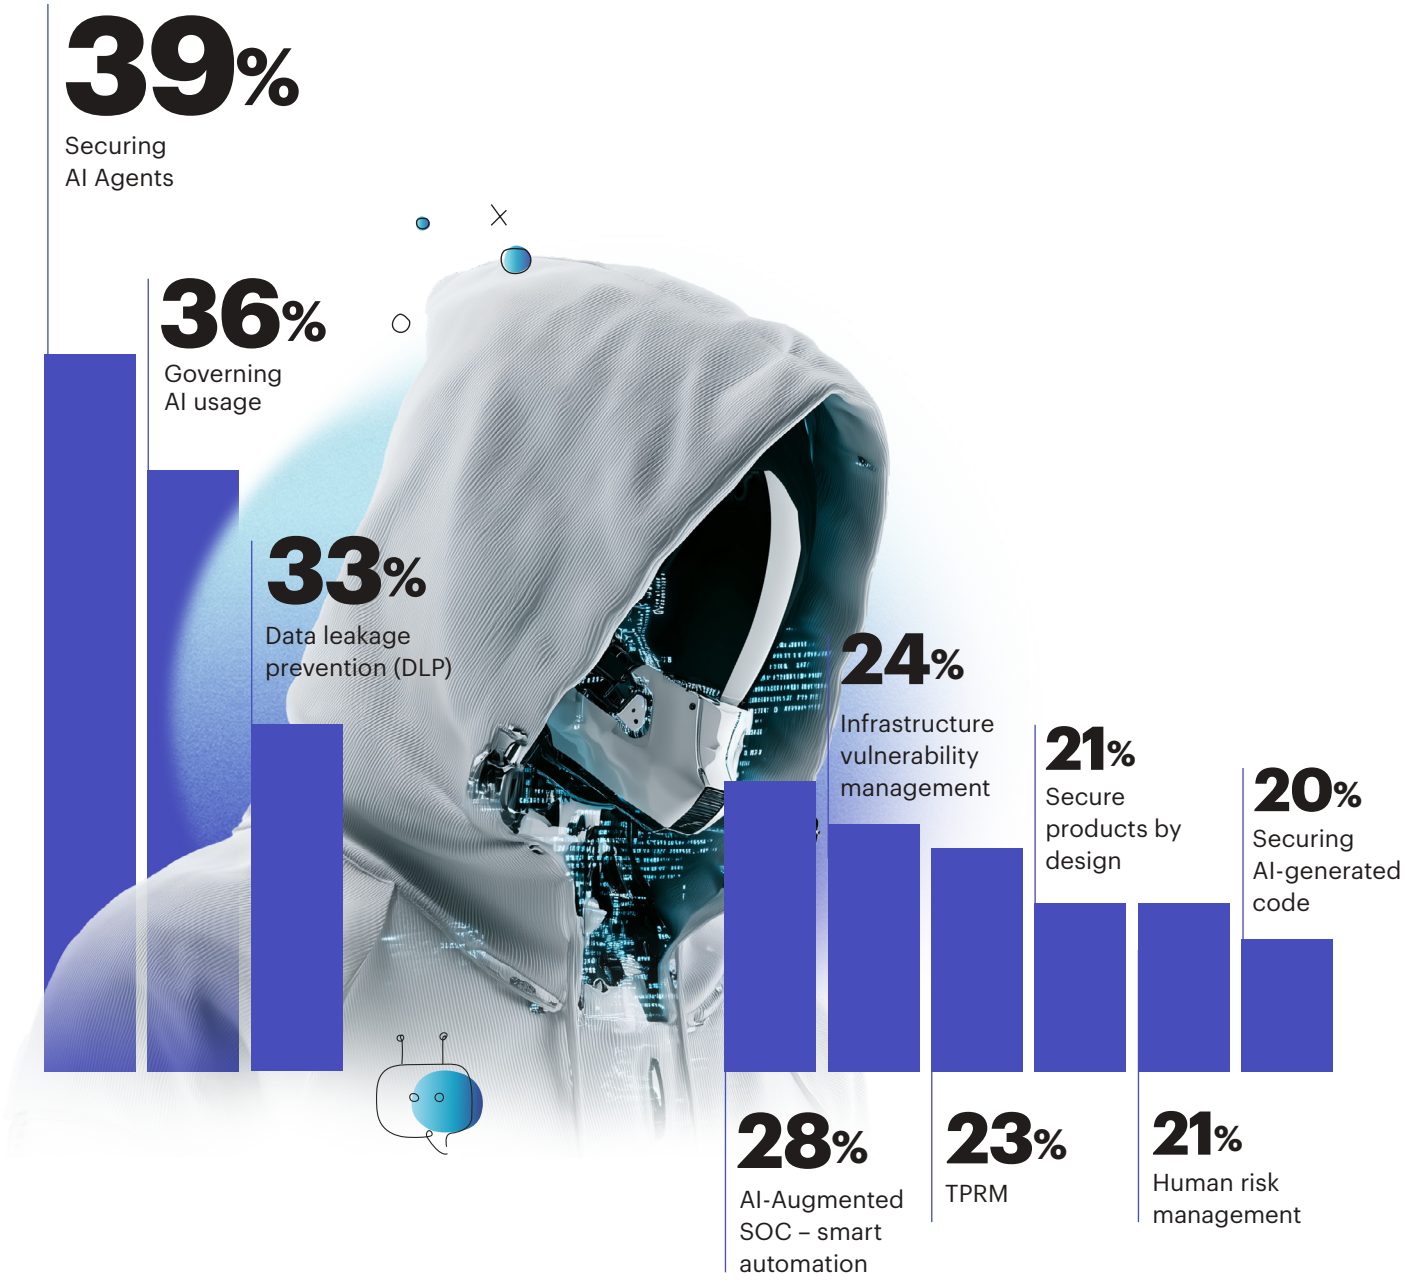

In 2025, the message is clear: **AI has introduced a new wave of risk, and it's now the top innovation priority**. Securing AI agents (39%) and governing employee AI usage (36%) took the first and second spots, surpassing even foundational issues like data loss prevention (33%), SOC automation (28%), and third-party risk management (23%). That marks a significant departure from previous years, where DLP, TPRM, and Identity consistently ranked highest. It signals a shift in CISO focus: from building resilience around known attack surfaces to securing an AI-native enterprise, and leveraging AI to augment cybersecurity defenses.

In the rest of this chapter, we'll dive into the top unsolved challenges, excluding AI-related risks, which were covered under the "AI Arms Race".

Data Leakage Prevention

Data Loss Prevention has been a cybersecurity staple for decades, but ask any CISO and the answer is clear: **it still doesn't work**. The same pain points persist: rigid pattern matching, poor visibility across surfaces, and brittle, outdated policies. Legacy DLP tools rely on regex and manual tagging, struggle with unstructured data, and flood teams with false positives, leading to constant triage and alert fatigue.

The problem has only worsened as data moves fluidly across SaaS apps, endpoints, and collaboration tools. DSPMs have helped identify where sensitive data lives, but not how it moves or leaks in real time. The question now is: **how do we prevent sensitive data from reaching the wrong place without slowing down the business and overwhelming security teams with noise?**

AI offers a promising path forward. For the first time, we can combine **deep content understanding with contextual awareness**, understanding what a document contains, but also why it's being sent, and as part of which broader business process or chain of events. AI models can classify far beyond standard PII, capturing nuanced, sensitive content like acquisition term sheets, investor updates, or internal strategy decks. Paired with behavioral signals and modern sensors, the next generation of DLP can be contextual, adaptive, and invisible to the user.

AI-Augmented SOC

The Security Operations Center remains one of the most labor-intensive functions in cybersecurity, and analysts suffer from repetitive workflows, high false positive volumes, and escalating alert fatigue. At the same time, the industry faces a chronic shortage of skilled talent and burnout. AI presents a clear opportunity to augment analysts, reducing the burden of manual triage and enabling faster, more consistent responses.

When asked what would drive AI adoption in the SOC, CISOs pointed to tangible, operational incentives:

Top Success Metrics for AI Soc Adoption

60%
Proven improvement in threat containment speed

47%
Cost savings vs. additional headcount hiring

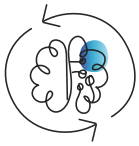

45%
Reduction in false positive escalation

37%
Measurable decrease in analyst overtime/burnout

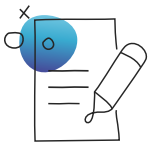

But the path to automation is layered. Many enterprises are not managing the SOC entirely in-house. **45% of CISOs reported a hybrid model**, where detection and response is split between internal teams and external vendors. Among those, **incident response retainers are the most common (60%)**, followed by **MDR services (34%)** and **MSSPs (31%)**.

Still, **augmenting analysts is only part of the puzzle**. The infrastructure that supports SOC operations is itself under strain. **Only 39% of CISOs reported high confidence in their SOC stack**. The rest cite brittle pipelines, ballooning SIEM storage costs, and complex integrations that create blind spots and friction.

In this environment, **AI isn't just a speed boost, it's a structural fix**. The opportunity is to shift from reactive escalation to intelligent prioritization, from alert fatigue to smart filtering, and from brittle tooling to scalable, adaptable SOC infrastructure.

External SecOps Services

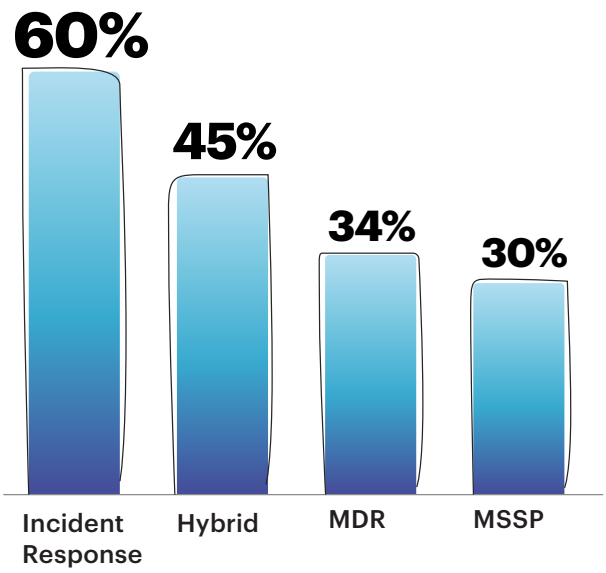

Confidence in SOC Infrastructure

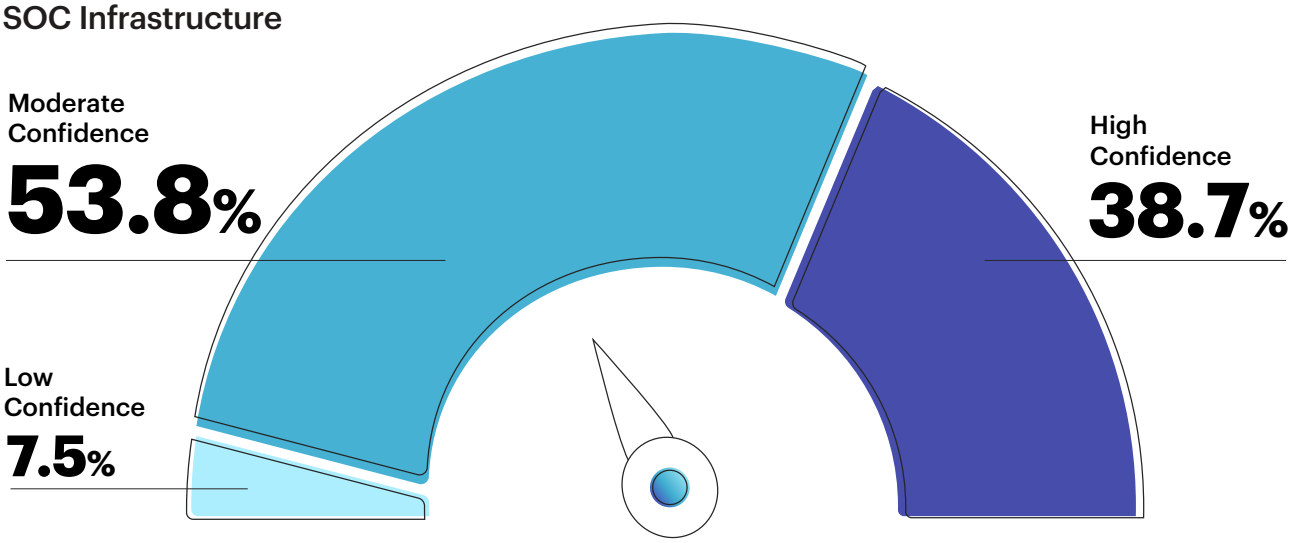

> *Integrating AI agents into our SOC is central to our strategy for evolving cybersecurity operations. This next step aims to enhance our defense mechanisms to match the sophistication of our adversaries, transitioning to a more proactive, adaptive, and precise model. By combining agentic-AI with the expertise of SOC analysts, we will effectively manage a significantly higher volume and complexity of events and threats, which we anticipate will be our next major challenge in the near future".*

Javier Garcia Quintela, Global CISO

repsol

Infrastructure Vulnerability Management

Vulnerability management is one of cybersecurity's wicked problems — **it doesn't stand still, it only gets worse**. As enterprise IT environments become more distributed and dynamic, the volume of vulnerabilities detected by scanners often reaches into the millions, far outpacing what security teams can realistically assess or remediate.

The result is a widening gap between what's discovered and what's fixed. 37% of CISOs report that a significant amount (more than 40%) of critical vulnerabilities remain unpatched beyond SLA, leaving organizations exposed both to regulators and real-world attackers.

Limited resources mandate **aggressive prioritization**. Most enterprises still rely on generic scoring systems like CVSS to triage vulnerabilities. These "one-size-fits-all" metrics lack the nuance of real enterprise context, failing to consider reachability, exploitability, or the blast radius if exploited. This type of reasoning is typically left to human analysts, who must stitch together technical data, business impact, and threat intel on a case-by-case basis.

Agentic AI could change that. With the ability to mimic analyst-level thinking, AI agents could dynamically assess which vulnerabilities pose real risk in a given environment—factoring in internal topology, exposure paths, compensating controls, and business context. That shift could finally bring precision to prioritization.

Looking further ahead, **agentic patching may fundamentally rewrite the problem**. If agents can test, schedule, and deploy patches autonomously, the marginal cost of remediation drops dramatically. In that future, **prioritization might become obsolete**—not because we fixed it, but because we outpaced it.

The biggest barriers that keep vulnerabilities open past SLA

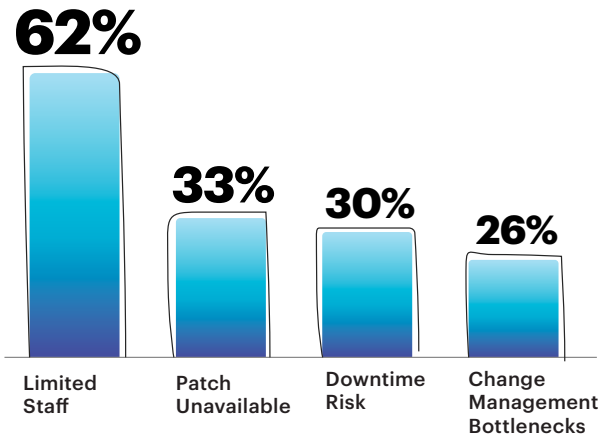

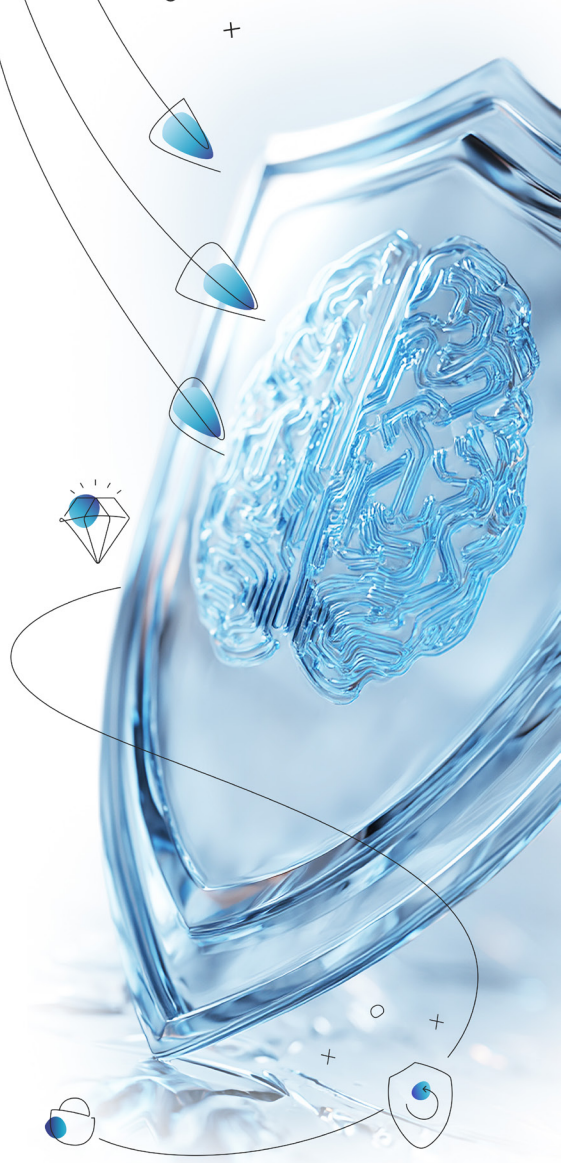

Almost 40% of CISOs report that more than 40% of critical vulnerabilities are not patched within SLA

TPRM

For years, third-party risk management has been one of the most persistent pain points in cybersecurity, **consistently ranking among the top three innovation priorities for CISOs**. In 2025, it remains a major concern, even as some urgency has shifted toward AI-related risks. The core problem hasn't changed: **TPRM is still a slow, manual, check-the-box process that demands massive effort with little meaningful assurance in return.**

Security teams are forced to chase vendors for questionnaire responses, review dense policy documents, and assess security controls based on self-attestation rather than real behavior. Even when done thoroughly, these assessments reflect a moment in time, not an ongoing picture of risk. As third-party ecosystems grow in size and complexity, this model is proving unsustainable.

A new generation of startups is rethinking the category, embedding AI to both eliminate manual overhead and unlock deeper visibility. Instead of having humans read through hundreds of PDFs, AI agents can parse vendor documents, match evidence to control frameworks, and enrich them with threat intelligence and relevant public information automatically. But more importantly, they're beginning to go beyond static assurance: integrating directly into enterprise environments to track how vendors are actually interacting with the enterprise over time.

These systems can detect when access drifts from the original scope, when a vendor's privileges escalate unexpectedly, or when an integration behaves differently than expected. In doing so, **they turn TPRM from a point-in-time audit into a real-time feedback loop**, and from a compliance cost center into a genuine security control.

Secure Products By Design

For years, product security lived under the umbrella of application security, focused on identifying bugs in code through scanning and testing. While effective, this model did not sufficiently address critical security failures like flaws in architecture and business logic. These issues are difficult to detect with traditional tools and even harder to fix once baked into the product.

Secure by design offers a more effective path forward. It means embedding security from the earliest stages of product planning, treating it as a core design and engineering concern, not a downstream QA task. The goal is to reduce risk across the entire lifecycle by anticipating how products will behave, how they'll be misused, and where implicit trust might break down. It's a shift **from reactive scanning to proactive, systemic resilience**.

In 2025, **CISOs are embracing a broader product security paradigm** that embeds security across the entire software development lifecycle, from design to deployment.

Current approach to product security

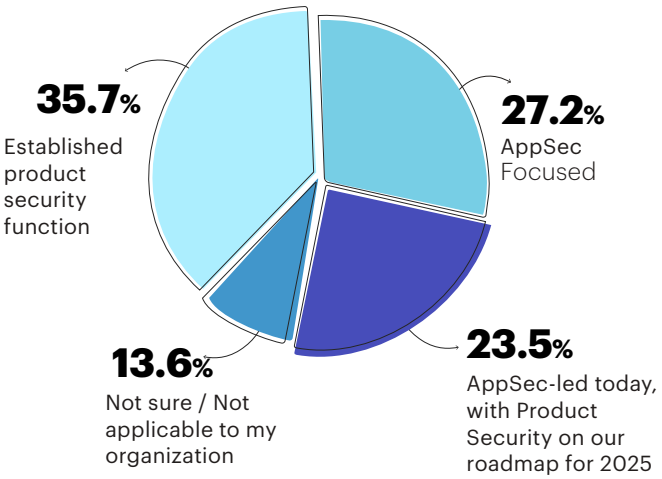

According to the survey, 36% of CISOs have already established a dedicated product security function, covering threat modeling, privacy reviews, architecture decisions, and secure development practices. Another 23% plan to adopt a product security model by 2026, signaling a significant shift toward deeper, more integrated security programs.

What's driving this evolution is the growing realization that the hardest threats to detect are no longer technical bugs, but design flaws and human logic issues.

Hardest issues to detect and address

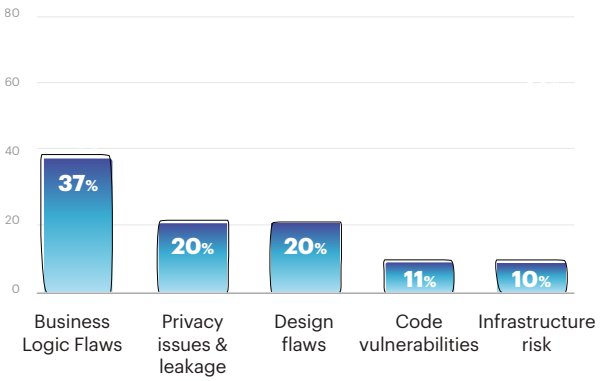

When asked what kinds of flaws are most difficult for their teams or tools to catch, **37% of CISOs cited business logic flaws**, followed by **privacy and data leakage risks (20%)** and **design flaws (20%)**. Traditional AppSec tools simply weren't built for these kinds of problems. They require cross-functional visibility, contextual understanding, and collaboration between security, product, and engineering.

In short: The age of treating security as a QA task is over. Today's products need security decisions made at the whiteboard, not just in the CI/CD pipeline.

> *"The journey from AppSec to ProdSec is a shift from checking code to correctness to designing for secure outcomes. Secure-by-design is a good start, but secure-by-default is better yet. Users shouldn't have to take action for software and services to be secure. That means security decisions need to involve product managers, UX designers, and business teams, not just engineers. It also means product security has to scale to a broader audience through clearer communication, smarter automation, and better collaboration across functions."*
>
> **James Dolph,** CISO
>
> GUIDEWIRE

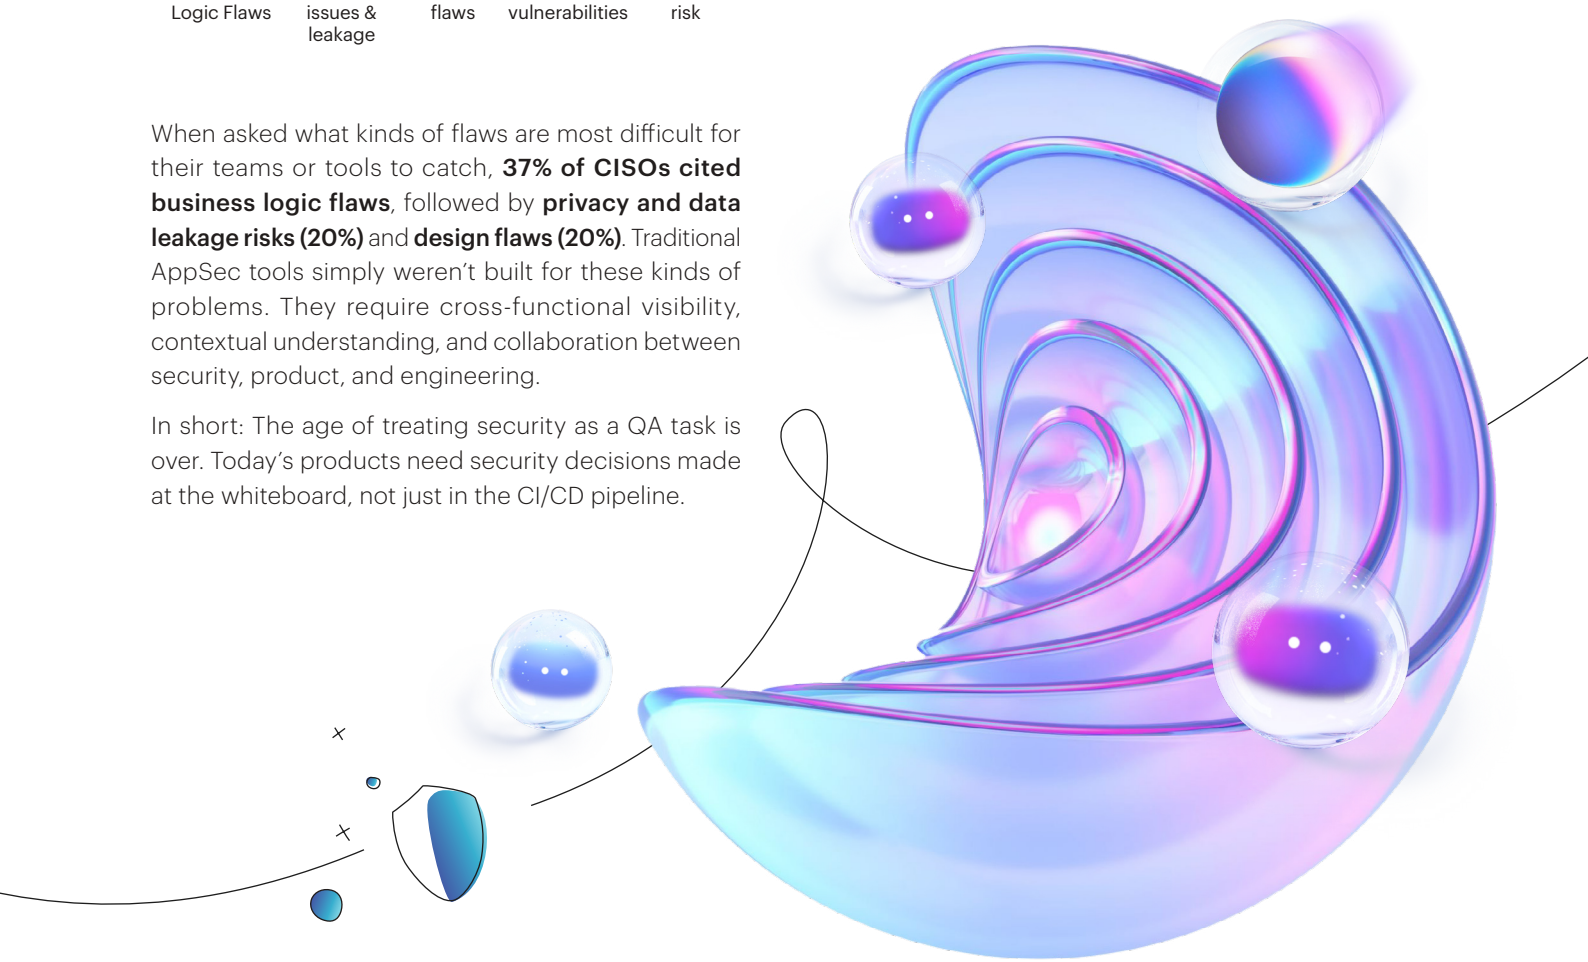

Human Security

The human element has long been security's soft spot, and **AI is turning it into a high-speed attack vector**. The human element has long been security's soft spot—and today's attackers are exploiting it with greater speed and sophistication. Phishing messages are now well-written and highly targeted. Deepfakes can convincingly mimic a CEO's voice or face in real time. Even voicemails and video calls are being used to deceive employees with remarkable realism.

68% of breaches are still caused by human actions, yet our approach to "fixing the human" remains stuck in outdated training modules and blanket awareness campaigns. That means looking at users the way we look at machines: **What's their exposure? What privileges do they have? What behaviors make them vulnerable?**

This category is ready to be reimagined. With AI agents, we can move beyond static training and toward continuous, adaptive and dynamic human risk management. That means personalized security guidance, policies that adjust to behavior in real time, and just-in-time training delivered at the moment of risk to the user who needs it most.

By combining behavioral telemetry with intelligent automation, we can generate user-level risk scores, spot patterns before they escalate, and turn training from a checkbox into a meaningful behavior-changing intervention. Done right, this shift could finally help reduce the human attack surface.

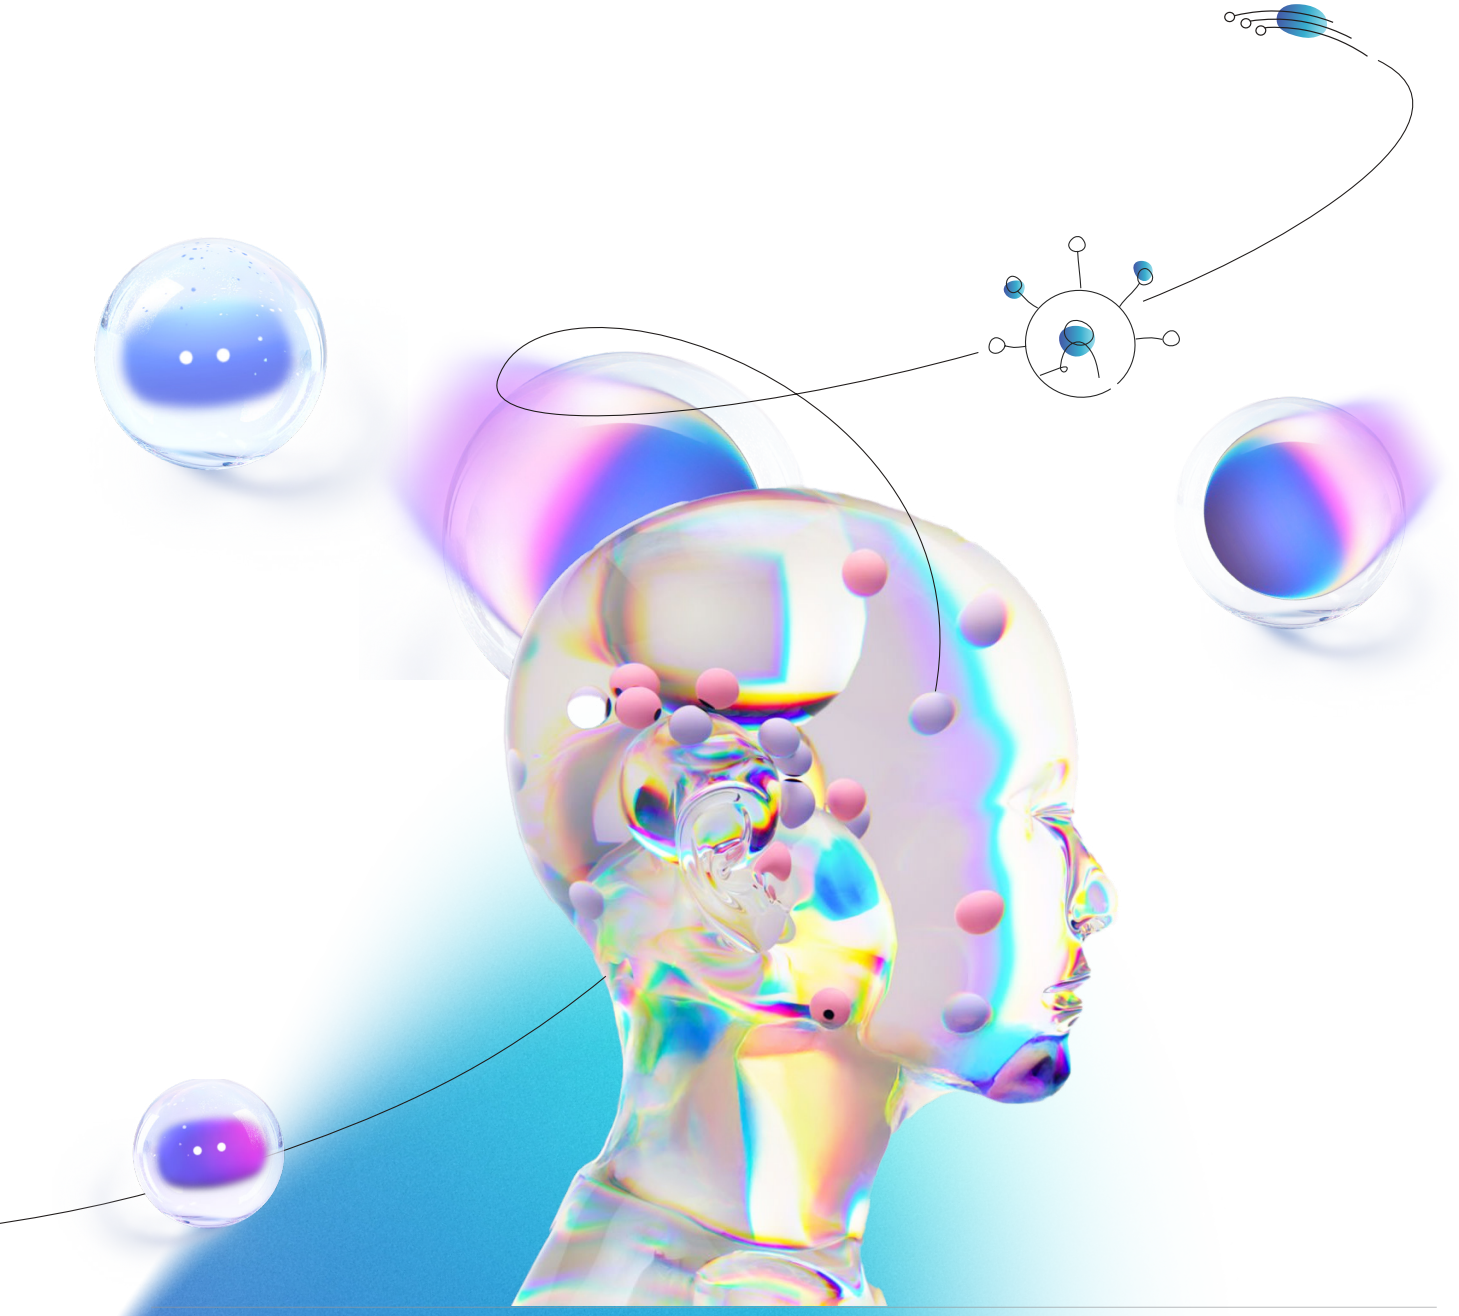

Methodology and Demographics

Each year, Team8's CISO Village Survey offers a rare glimpse into the minds of security executives leading some of the world's most prominent enterprises. The 2025 edition was launched at the annual CISO Village Summit, hosted across Miami and the Florida Keys. This year's theme, **"AI as your force multiplier"**, brought together over 110 CISOs from Team8's 600+ member community, including many from the Fortune 500.

The Summit featured voices from the cutting edge of security and AI, including Anthropic CISO Jason Clinton and SemiAnalysis founder Dylan Patel, as well as security leaders like Rich Baich (CISO, AT&T), Branden Newman (CTO and former CISO, MGM), and Phil Venables (former CISO, Google Cloud). In keynotes, workshops, and candid peer sessions, our Villagers shared lessons from the frontlines, including real breach stories, and debated how AI is reshaping the enterprise security model.

This report presents the survey results gathered from over 110 CISOs attending the Summit. Our goal is to enrich the broader cybersecurity community with the challenges, priorities, and opportunities shaping enterprise security in the AI era.

4.1. Industry Breakdown

Respondents represent large enterprises across various verticals. The largest representation came from the **technology sector (34.5%)**, followed by **financial services (27.4%)**. Industrial and manufacturing companies made up **15%** of participants, with the remainder spanning health and **pharma (8%)**, **retail and e-commerce (3.5%)**, and **others (6.2%)**.

This report presents the survey results gathered from over 110 CISOs attending the Summit. Our goal is to enrich the broader cybersecurity community with the challenges, priorities, and opportunities shaping enterprise security in the AI era.

What is your Industry

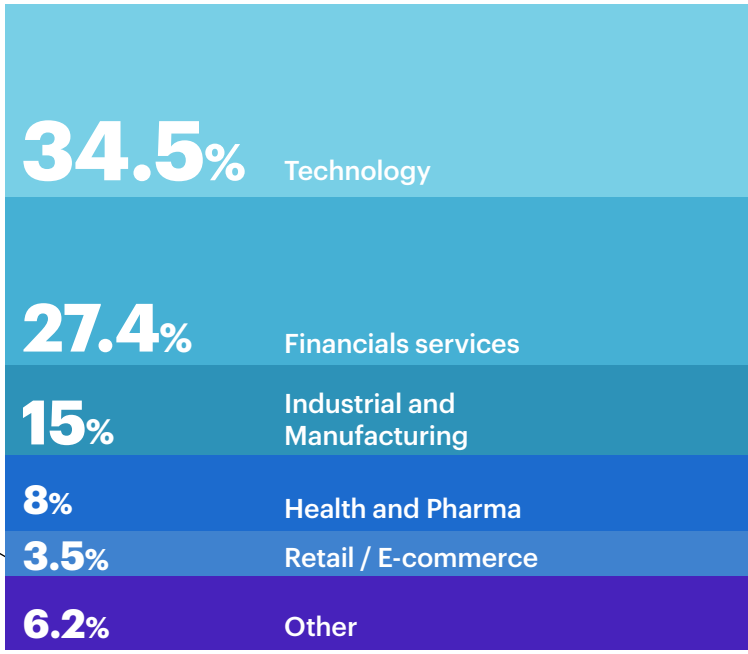

4.2. Security Team Size

This year's data reflects a broad range of cybersecurity team sizes, with a noticeable lean toward leaner teams: **39%** of respondents reported having **20 or fewer security professionals**, while another **27.4% r**eported teams between 21–100 people. **18.6%** had security teams exceeding 500 professionals.

What is the current size of your cybersecurity organization?

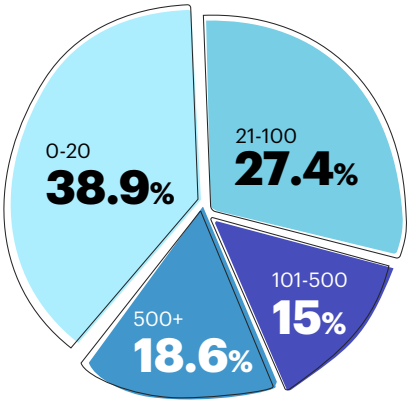

4.3. Engineering Team Size

CISO strategies are often shaped by the size of the engineering organization they support. Over half of the companies surveyed reported engineering teams larger than 500 people. Specifically, 24.8% have over 1,000 engineers, and 17.7% have teams of 501–1,000. 11.5% reported teams of 50 or fewer.

What is the current size of your engineering / CTO organization?

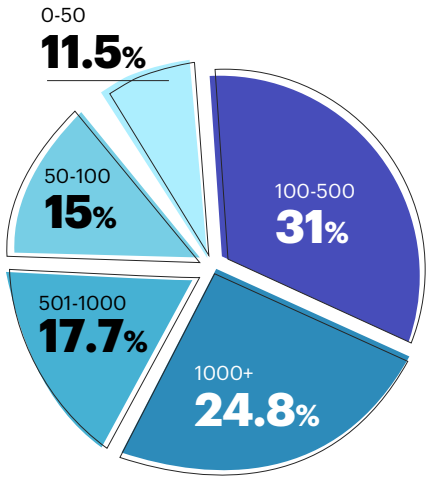

4.4. Cybersecurity Budgets

The companies surveyed manage meaningful cybersecurity budgets. The median budget sits in the **$3M–$10M range**, with **38.1%** of CISOs reporting within that bracket. Nearly **18.6%** reported budgets below $3M, while at the top end, **7.1%** reported budgets above $100M.

What is your organization's cybersecurity budget?

$3M-$5M
38.1%

<$3M
18.6%

$10M-$20M
17.7%

$20M-$50M
11.5%

$100M+
7.1%

$50M-$100M
7.1%

This demographic spread reflects a diverse but security-forward sample of large enterprises navigating rapid technological change. Their responses in this report surface early signals of where the enterprise security market is heading next.

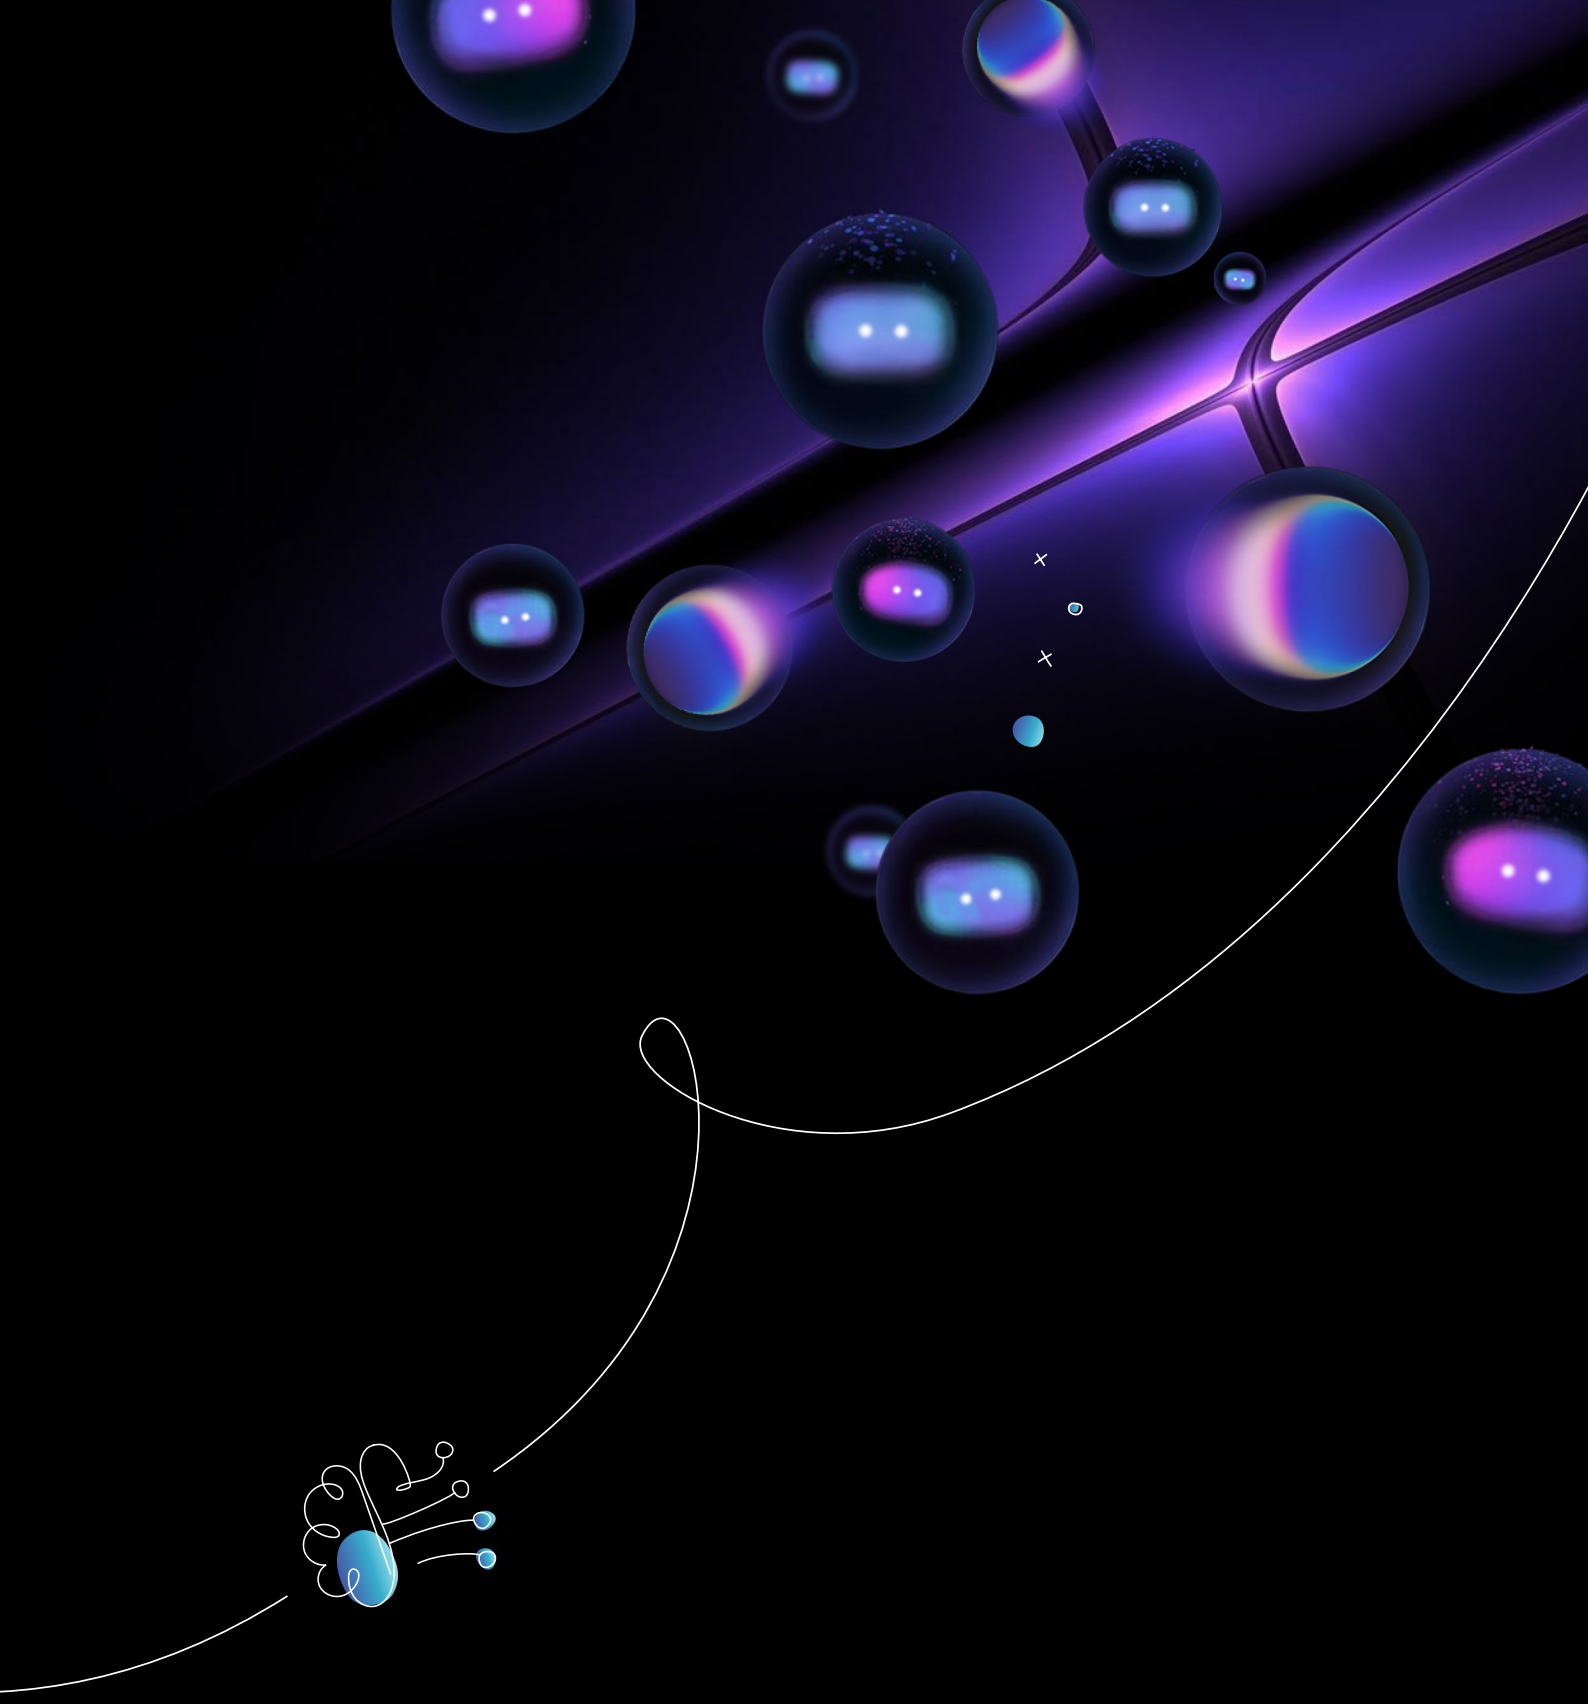

For more information

Contact us at: cisovillage@team8.vc | www.team8.vc

TEAM8™