WITNESS AI

2026 Enterprise AI Risk Survey

The Hidden Cost of Enterprise AI

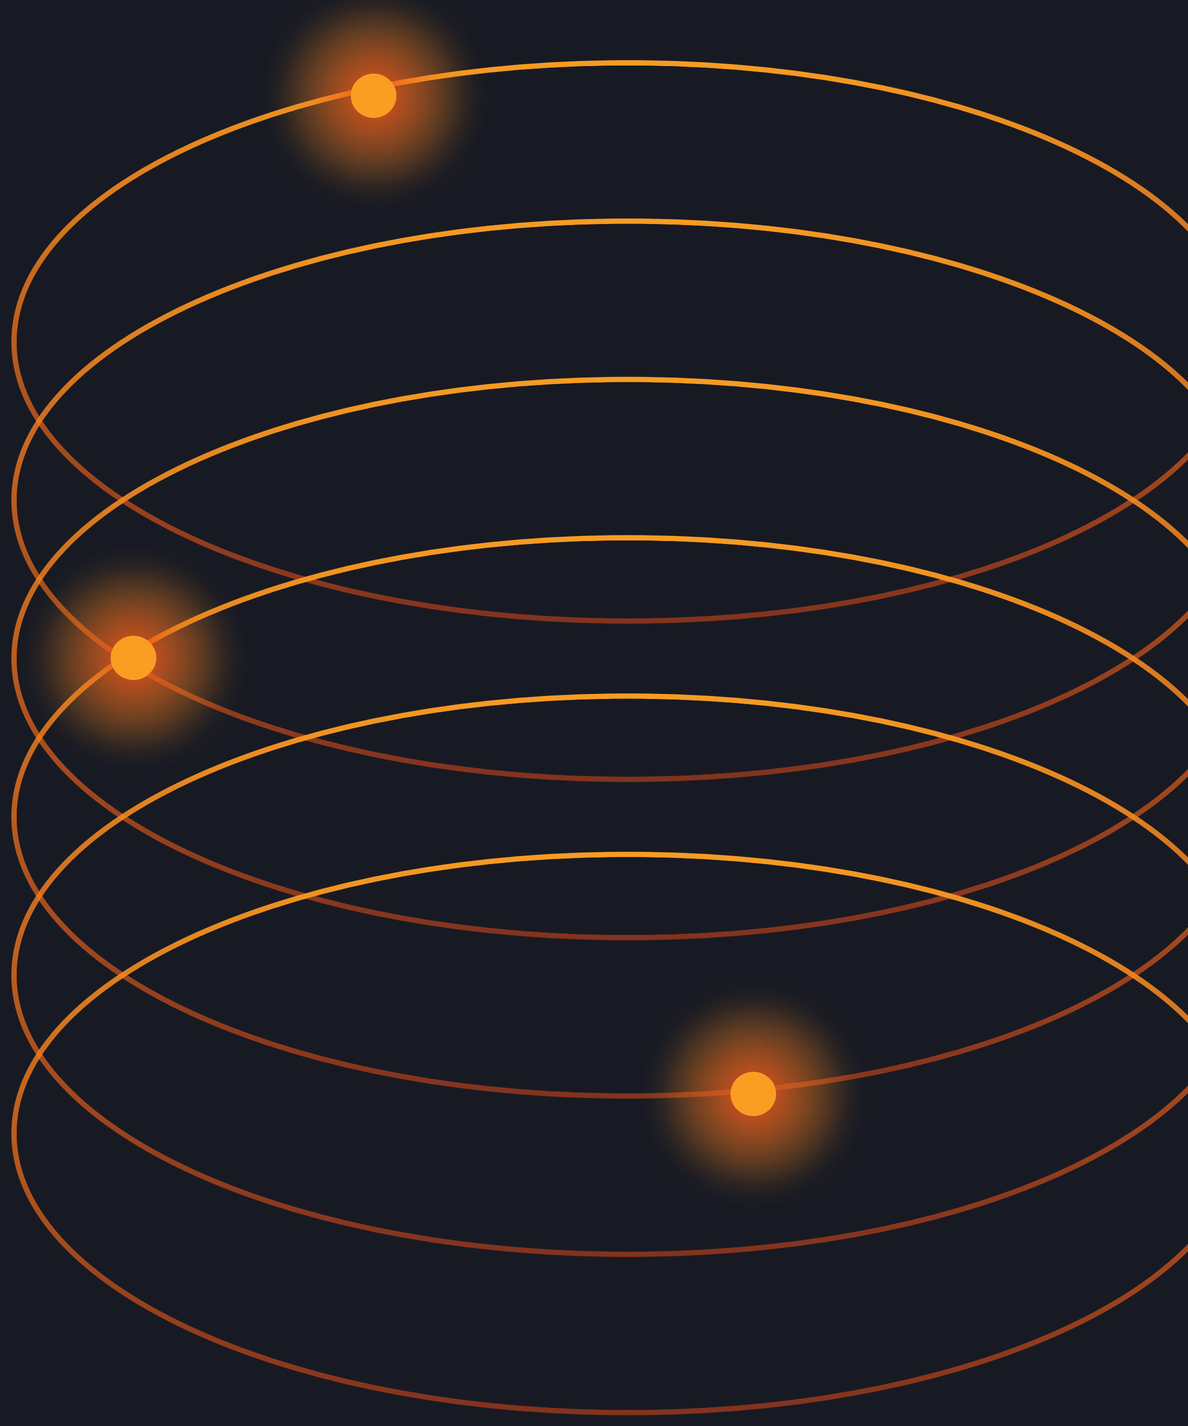

Enterprises are deploying AI agents faster than they can measure, predict, or contain the financial fallout. New research from WitnessAI finds that most organizations have already experienced AI-related incidents, in some cases significant cost overruns, and rising regulatory exposure — even as confidence in AI's payoff remains far from settled.

Seventy percent of respondents said they are already using or piloting AI agents. Yet among those organizations, fewer than one in five (18%) said all agents are formally inventoried and approved by security, and just 49% have continuous monitoring in place.

Earlier this year, the public conversation transitioned from risk to productivity. Instead of "is this safe?" executives started asking, "are we getting what we're paying for?" That shift happened before most organizations had fully answered the first question. Where most enterprise AI surveys measure concern, this report measures cost: the incident, compliance, and revenue-exposure figures business leaders can act on.

To better understand how organizations are managing AI risk, WitnessAI commissioned The Harris Poll to conduct a survey of 300 senior enterprise decision-makers in spring 2026, including 100 C-suite leaders and 200 vice presidents at various levels responsible for technology, security, risk, compliance, finance, and AI strategy. The findings reveal a gap within financial risk management: organizations are deploying AI agents faster than they are building the financial visibility, ownership structures, and risk management practices needed to manage them responsibly.

Key Findings

AI agents are moving forward despite widespread concern.
Ninety-one percent of respondents said they are concerned that AI agents increase financial risk exposure, yet 64% believe their value outweighs those risks.

AI ROI remains difficult to prove.
Only 9% said that more than three-quarters of AI initiatives have delivered measurable financial return, while 68% reported that AI projects exceeded budget at least sometime during the past year.

AI incidents are already creating significant financial consequences.
21% classified their most significant AI incident costs at $1 million or more. The cost of all AI-related incidents in the past 12 months was $2 million or higher for 43% of those surveyed.

Key Findings

Agent deployment is outpacing oversight. Seventy percent are using or piloting autonomous AI agents, but only 18% have all agents formally inventoried and just 49% have continuous monitoring in place.

Risk management investment is rising, but governance bottlenecks remain the leading contributor to AI ROI underperformance — a sign that risk management practices haven't caught up with how fast AI is being deployed.

AI risk has no consensus owner. Responsibility defaults to the CIO or IT, but no single role commands a majority; just 46% of organizations said their CFO actively models AI-specific risk and ROI, leaving finance under-integrated.

Senior leaders and the teams running deployments don't see the same risk. 68% of C-suite respondents reported full confidence in their visibility into AI tools and agents versus 46% of vice presidents, even as senior leaders reported experiencing larger incidents and higher costs than the VPs.

Section 1

The Promise of AI Outweighs the Risk for Enterprise Leaders

Organizations are making a deliberate bet that the benefits of AI outweigh the risks of adopting it before risk management and measurement practices have fully matured.

The data makes clear that enterprise leaders understand that AI agents introduce new forms of financial and operational risk. Yet the findings also show that those concerns have done little to slow deployment as competitive pressure and expectations for productivity gains continue to ramp up.

Nearly every respondent (91%) expressed concern that AI agents increase financial risk exposure, including 58% who describe themselves as either extremely or somewhat concerned (Table 1). At the same time, 64% said they believe the value AI agents can deliver outweighs those risks.

Together, those findings illustrate the central tension shaping enterprise AI strategy today. Organizations increasingly view AI as an essential business capability even as many acknowledge they have not fully developed the visibility, measurement frameworks, risk management practices, and accountability structures needed to support it.

Rather than delaying adoption until every policy and control is in place, many organizations are deploying AI while building the processes needed to measure and manage its impact at the same time.

TABLE 1

How concerned are you that AI agents increase your organization's financial risk exposure?*

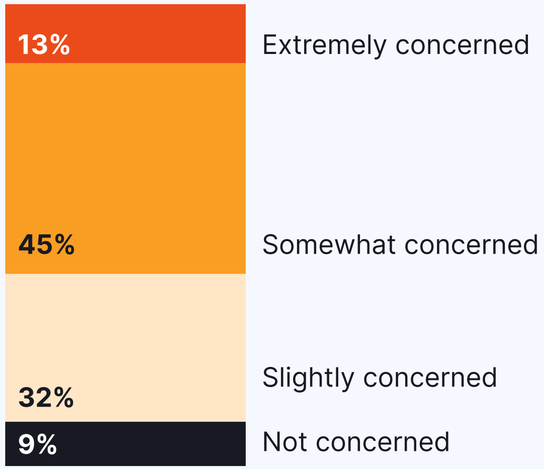

*Percentages may not sum to 100% due to rounding.

> “It’s competitive pressure, board pressure, a CEO who's decreed that every department will deploy agents. **If we don't do it, our competitors will.** Both the risk and the value were assumptions, not estimates. The bet was that risk could be managed and value would appear."

Rick Caccia
WitnessAI CEO
and Co-Founder

Section 2

The Business Case for AI Remains Difficult to Measure

Despite widespread investment in AI, most organizations struggle to quantify what that investment is actually returning, and the early signals are mixed. AI initiatives are exceeding budgets, underperforming against ROI expectations, and running into governance bottlenecks that slow deployment without reducing risk.

The data suggests the problem is both operational and structural. Organizations lack the measurement frameworks to know which bets are paying off and which are compounding cost.

Investing in AI has been the easy part; proving it pays off has not.

Only 9% of respondents said that more than three-quarters of their AI initiatives have delivered measurable financial return, and 43% put that figure at one-quarter to one-half.

Part of the difficulty in proving that return is that AI's impact rarely shows up in a single metric — it surfaces instead as fragments scattered across budgets, productivity reports, and individual business units, making a single ROI figure hard to pin down.

TABLE 2

How often have your AI initiatives exceeded their original budget?

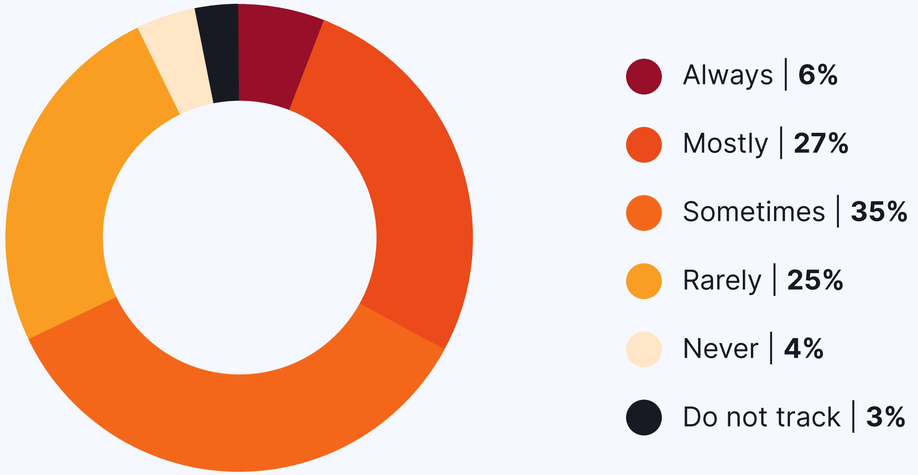

Budget discipline compounds the problem: 68% said AI initiatives have exceeded budget at least sometimes in the past year, 33% said budgets are mostly or always exceeded, and just 4% said initiatives never exceed budget (Table 2). Without a unified view of cost and return, executives may find themselves increasing AI investment without a clear understanding of whether it's paying off.

Section 3

Risk Management Investment is Rising, but Outcomes Are Lagging

The survey also reveals that organizations recognize the need for stronger oversight as AI adoption expands and, as a result, they are dedicating substantial resources to risk management, security, and governance activities.

Governance effort isn't translating to results.

More than half (54%) of respondents said they allocate between 21% and 45% of their AI budgets to risk management and governance. In other words, risk management now claims a real line in the AI budget, not a footnote.

Even so, organizations are struggling to balance oversight with execution. Seventeen percent identified governance bottlenecks as the leading reason AI initiatives underperform against ROI expectations, followed by inability to scale AI initiatives (15%) and tool cost overruns (13%) (Table 3). Separately, when asked whether governance itself constrains AI adoption, 18% described their governance posture as a significant deployment bottleneck.

The downstream effects are reflected in the data: 30% reported that unmanaged or poorly governed AI usage led to cost overruns, and 27% said it resulted in delayed or canceled AI initiatives.

This helps explain why traditional governance — built around committees, written policy, and periodic review — struggles to keep pace. "By the time you spend six months defining a policy and getting it approved by a committee, 53 new things have come out and your project is a waste of time," said WitnessAI Chief Product Officer Daniel Graves. The survey data backs him up. Governance effort is creating friction without reducing exposure, slowing down approved projects while unapproved ones move forward regardless.

TABLE 3

What is the single biggest reason your AI initiatives have underperformed against ROI expectations?

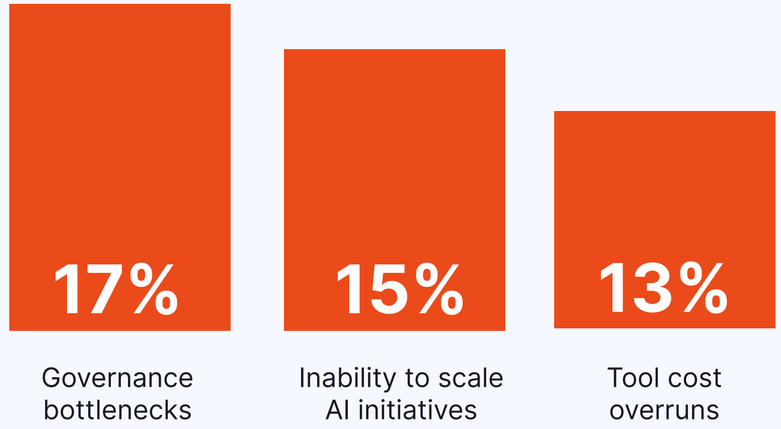

Section 4

The Financial Consequences of AI Risk Are Increasingly Visible

AI risk has traditionally been measured in technical terms: hallucinations, bias, cybersecurity vulnerabilities. The survey suggests organizations are increasingly evaluating AI through a financial lens as incidents, compliance obligations, and regulatory exposure become easier to quantify.

Eighty-six percent of respondents said they investigated one or more AI-related security or operational incidents during the previous 12 months (Table 4).

Among organizations with deployed AI agents, 12% report minimal or no formal oversight, which helps explain the scale of financial consequences the data reveals. A total of 21% of respondents classified their most significant AI incident costs at $1 million or more, including 12% between $1 million and $4.9 million, 7% between $5 million and $14.9 million, and 2% at $15 million or more. These include legal fees, remediation, loss of IP value, and regulatory fines.

When asked separately about the total estimated cost of all AI-related incidents over the past 12 months, not just the single most significant one, 43% of those who reported incidents said they faced financial implications of $2 million or higher. That included 17% who estimated total annual costs between $10 million and $24.9 million and 5% who estimated between $25 million and $50 million (Table 5).

TABLE 4

How many AI-related security or operational incidents has your organization investigated in the past 12 months?

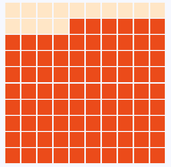

86%

investigated one or more **AI-related security** or **operational incidents** during the previous 12 months

While companies spent the last year building AI governance councils to figure out processes, departments moved fast to roll AI out anyway. The result is a cost burden arriving from two directions at once. The incident costs are detailed to the right. A second, separate wave is building from rising token and usage costs, as AI vendors shift from flat-rate to usage-based pricing.

Internal AI adoption incentive structures are rewarding volume of use over judgment about when AI should be used at all.

Many companies have found leaderboards and compensation-based motivation to be expensive, with reports of AI budgets being exhausted within months. Some developers have reported monthly bills in the thousands.

TABLE 5

What was the total estimated cost of all AI-related incidents in the past 12 months?

Total $2 million or more (NET)

43%

$2 million-$9.9 million

21%

$10 million-$24.9 million

17%

$25 million-$50 million

5%

Organizations are also assigning financial value to potential future exposure.

More than one-third (36%) estimated that between 3% and 5% of annual revenue could be at risk from a rogue AI agent event involving sensitive company or customer information. Nearly one-fifth (18%) say the regulatory impact could be double that (Table 6).

Among C-suite executives, 44% estimated 3%-5% of annual revenue could be at risk, compared with 32% of VPs. This suggests senior leadership increasingly views AI risk as a material business issue rather than solely a technology concern.

TABLE 6

What percentage of your organization's annual revenue could be at risk from regulatory penalties in the event that unmanaged AI (such as a rogue AI agent) exposes sensitive company and customer data?

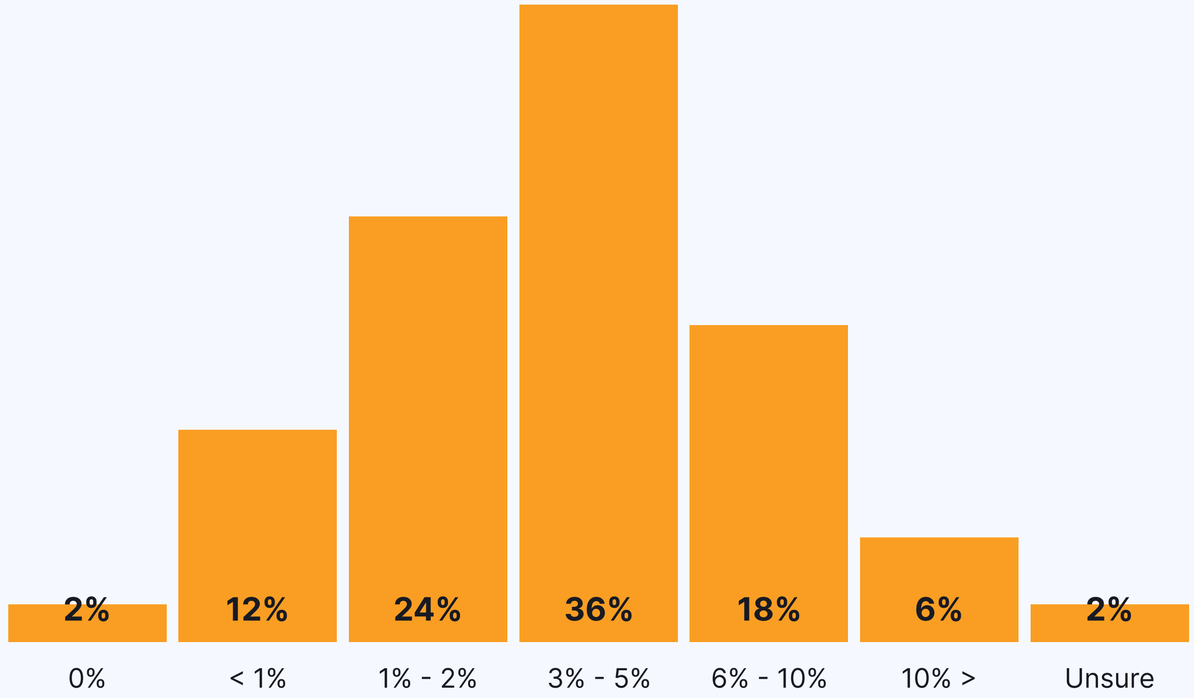

The cost of regulatory compliance is also becoming more visible. Twenty-nine percent estimated spending between $500,000 and $999,999 annually to achieve AI regulatory compliance under current manual processes, while another 29% estimated annual costs between $1 million and $3 million (Table 7).

Add up incident costs, usage costs, compliance spend, and revenue-at-risk estimates, and a pattern emerges. AI risk now shows up on the balance sheet, not just the threat model.

TABLE 7

What is your estimated cost to achieve compliance with upcoming AI regulations under your current manual processes?

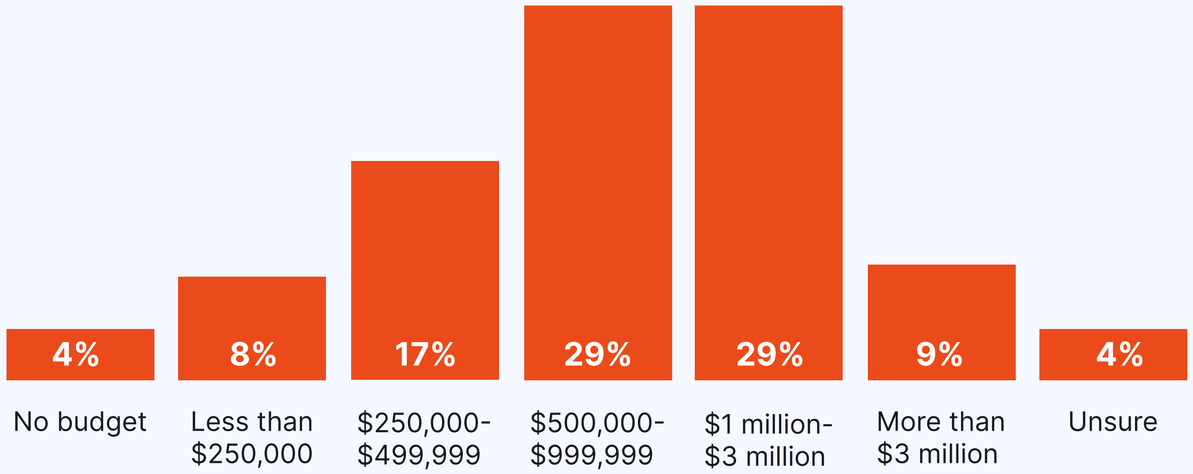

> "The financial exposure organizations are quantifying here matches what we're seeing across our portfolio. **AI-related incidents are already producing material losses**, and most companies haven't built the instrumentation to see them coming."

Barmak Meftah
Co-Founder & General Partner at Ballistic Ventures

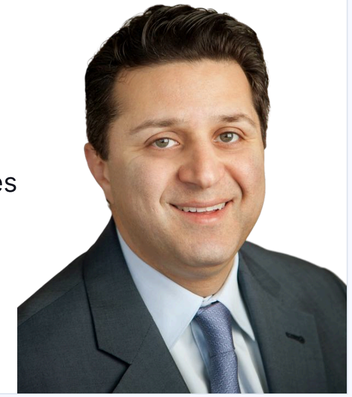

Section 5

Accountability Structures Have Not Kept Pace With AI Adoption

As AI systems become more autonomous, many organizations remain uncertain about who is responsible for managing AI risk and who ultimately owns the consequences when something goes wrong.

When asked who is primarily responsible for managing AI risk, the most common answers were CIO (30%), CISO (15%), and shared responsibility across roles (14%), with the remainder spread across CTOs, Chief AI Officers, legal teams, and dedicated governance teams (Table 8).

Similar ambiguity appears when respondents consider liability. Only 6% identified the CISO as primarily liable when an AI agent causes financial or regulatory harm, while 26% assigned primary liability to the CIO and 18% said they believe responsibility is shared across multiple roles.

Finance also remains unevenly integrated into AI risk management efforts. Only 46% reported that their CFO actively models AI-specific ROI and risk, while 38% said finance reviews AI spending but does not separately model AI-related risk exposure. As organizations assign increasingly large budgets to AI initiatives and quantify potential revenue exposure in percentage points of annual revenue, the absence of consistent financial modeling may represent a significant blind spot.

TABLE 8

Which role is primarily responsible for managing and mitigating AI-related risk within your organization? (Top 3)

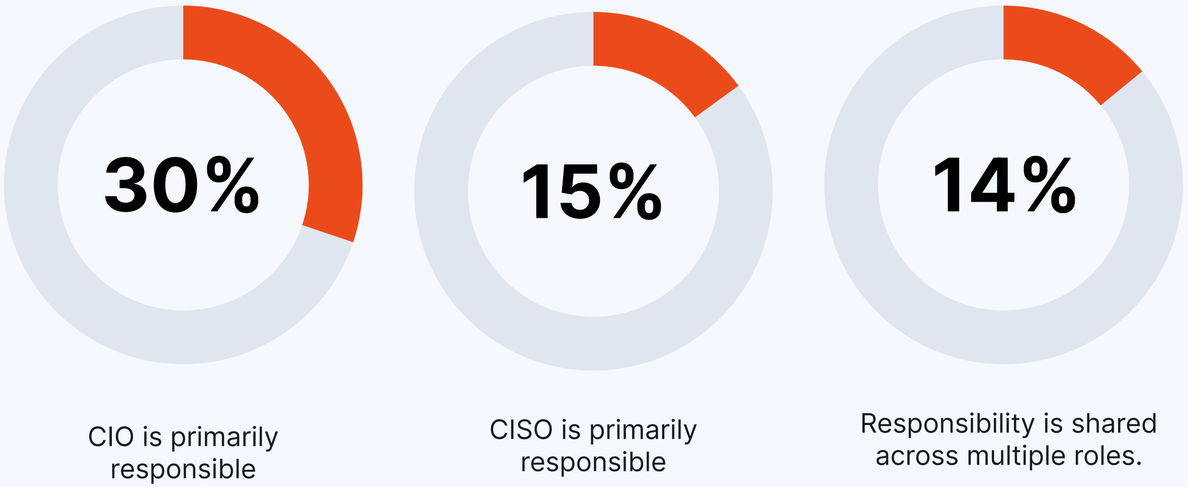

These statistics highlight a systemic inability to reconcile responsibility for risk, a failure driven more by the velocity of adoption than by oversight.

Given that incident expenses, regulatory costs, and revenue exposure have reached the levels documented in this research, the lack of such an approach is creating material financial impact.

> “
>
> This isn't an academic exercise. Organizations are **struggling to keep up with the sheer speed of change**. While some have introduced new AI risk roles, others let the burden fall to the CISO by default, regardless of whether they possess the necessary mandate or technical insight."

Daniel Graves
Chief Product Officer,
WitnessAI

Section 6

Executive Leadership and VP-Level Decision-Makers View AI Risk Differently

The survey reveals a consistent split between how the C-suite and the various vice president levels, who are closer to deployment, describe their organization's handling of AI risk. The top leaders report more visibility, firmer governance, and clearer lines of ownership than the various VP levels.

That perception gap is widest on accountability.

When asked who is primarily responsible for managing AI risk, 20% of VPs said responsibility is shared across roles, compared with just 4% of C-suite respondents. The same pattern holds for liability when an AI agent causes financial or regulatory harm: 21% of VPs called it shared, versus 11% of the C-suite.

However, senior leaders are more likely to name a specific owner. On finance specifically, 12% of the C-suite assigned AI risk ownership to the CFO, versus 2% of VPs. Leadership tends to believe ownership is assigned, while VPs are far more likely to call it diffuse. In practice, shared ownership often functions as no ownership at all.

The split runs through confidence and maturity as well. C-suite respondents consistently report higher levels of visibility and confidence in the following areas (Table 9):

TABLE 9*

Full confidence in visibility into AI tools, models, and agents accessing company data.

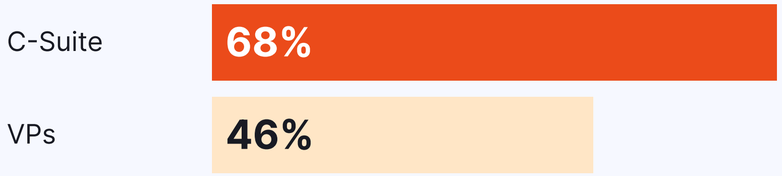

Formal, enforced, and organization-wide AI governance maturity.

Real-time visibility into AI activity across human-to-AI interactions.

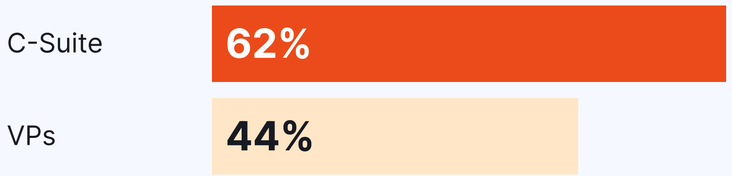

*Each metric represents responses to a separate survey question. Percentages should not be summed or directly compared.

> The board asks about risk, the **CISO points to the products purchased** — an E5 license, Purview, some AI security tooling — and the C-suite concludes that the risks are covered.
>
> It's a long-existing gap between risk and security spend."

Rick Caccia
WitnessAI CEO
and Co-Founder

What makes the disconnect between the roles notable is that senior leaders also see the larger financial stakes. On incident costs, C-suite respondents were more likely to say their most significant AI incident reached $1 million or more (28% vs. 18% of VPs) (Table 10). On regulatory exposure, C-suite leaders were about twice as likely to estimate annual compliance costs between $500,000 and $1 million (44% vs. 21% of VPs) (Table 11).

The result is confidence at the top that isn't always shared by the people running the deployments. When leadership believes coverage is in place and the implementation layer does not, that gap shapes how budgets get allocated and whether an accurate picture of exposure ever reaches the board.

TABLE 10

What was the approximate financial impact of your most significant AI-related incident in the past 12 months?

$1 million+

VPs	18%
C-suite	28%

TABLE 11

What is your estimated cost to achieve compliance with upcoming AI regulations under your current manual processes?

$500,000 - $999,999

VPs	21%
C-suite	44%

Conclusion

The Next Phase of Enterprise AI

The survey findings reveal a consistent theme across organizations: AI adoption continues to accelerate because leaders believe the technology represents a strategic business opportunity. Yet many organizations are still developing the visibility, measurement frameworks, risk management practices, and accountability structures necessary to manage that opportunity effectively.

Enterprise AI maturity starts with visibility.

“Solutions exist today that show you what data is being used within AI tools and where, so the priority should be understanding your actual exposure as soon as possible," said Rick Caccia. From there, the work is connecting that visibility to financial discipline. That means measuring cost against value, assigning clear ownership for risk, and building risk management approaches that enable innovation rather than just constraining it.

One pitfall is getting locked into a single, rigid platform. “This problem is moving too fast for one-size-fits-all checkbox solutions, and consolidation won't happen quickly enough to keep up," Caccia said. "Flexibility matters more than familiarity.”

Organizations that build these capabilities will capture more of AI's value while managing less of its downside. The ones that don't are running an experiment they can't see the results of yet.

Sidebar

Shadow AI Remains a Visibility Challenge

Shadow AI is often assumed to be a problem of unsanctioned use by individual departments, such as engineers or marketing teams adopting tools without IT's knowledge or approval.

The data tells a different story. IT and infrastructure was found to be the largest single source of shadow AI activity, at 47%, ahead of sales and business development (34%) and marketing and communications (33%). According to the findings, the department most responsible for governing AI use within the organization is also the department most likely to be operating outside its own policies (Table 12).

TABLE 12

In which departments do ungoverned or unsanctioned AI usage most commonly occur?*

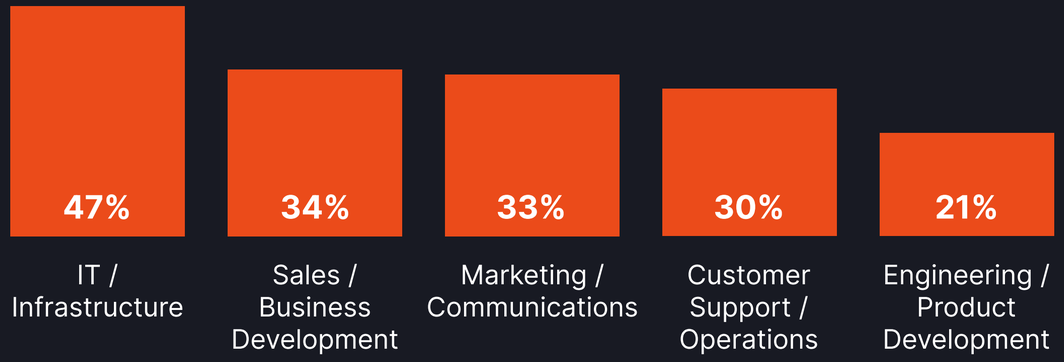

*Respondents could select multiple departments

WITNESS AI

About Witness AI

WitnessAI is the AI security platform enterprises trust to govern and protect all AI activity.

We provide complete visibility into every interaction including employees and autonomous agents, even in native apps where legacy tools are blind. Unlike traditional security that relies on outdated keywords, our AI-native platform understands intent, enabling intelligent policies that stop novel threats like prompt injection while empowering productivity. Our enterprise-first, single-tenant architecture ensures data sovereignty and compliance. WitnessAI transforms security from a bottleneck into the enabler of your AI strategy as the confidence layer for enterprise AI.

METHODOLOGY

This research was conducted online in the United States by The Harris Poll on behalf of WitnessAI from May 11 to June 1, 2026, among 300 senior executives: 100 C-level leaders and 200 vice presidents who have input into AI-related decision making at organizations with at least 1,000 employees. Percentages may not sum to 100% due to rounding. The sampling precision is accurate to within ±5 percentage points.