ENTERPRISE AI GOVERNANCE RESEARCH · 2026

peridot

# Shadow AI Benchmark Report

Produced by Peridot — Enterprise AI Governance Platform

Survey findings from 400 enterprise IT and security leaders at organizations with 5,000+ employees across 8 industries

n = 400 | 5,000+ employee orgs | 8 industries | CISOs · CIOs · VPs IT & Security

## KEY FINDINGS AT A GLANCE

| | |
|---|---|
| **86%** expect AI app volume to increase significantly in 12 months | **75%** cannot fully detect data sent to external AI services |
| **57%** have significant gaps or are unprepared for regulatory scrutiny | **60%** have no clear ownership framework for AI-generated application incidents |
| **44%** report fewer than 25% of AI tools are formally approved by IT | **43%** have not audited for AI-related data exposure incidents |

METHODOLOGY

# Assessing the State of Enterprise AI Governance

To understand how the world's largest organizations are navigating the tension between rapid AI adoption and the risks of Shadow AI, we conducted a comprehensive multi-method research study designed to move beyond surface-level trends — focusing instead on the practical realities of governing AI within complex, highly regulated environments.

Our findings are the result of a rigorous two-pronged approach: 20 in-depth qualitative interviews with senior executive leaders, conducted under Chatham House Rule, followed by a quantitative poll of 400 IT and security professionals at organizations with 5,000+ employees.

## RESEARCH DESIGN & PARTICIPANT SELECTION

**The Enforcers** 45%

CISOs, Deputy CISOs, and VPs/Directors of Security — providing frontline perspective on threat vectors and unauthorized LLM usage.

**The Enablers** 32%

CIOs, VPs of IT, and Enterprise Architects — representing AI's role in broader IT strategy and tech stack modernization.

**The Guardians** 23%

Heads of IT Risk/GRC and Chief Data/Privacy Officers — acknowledging that AI risk is inextricably linked to data sovereignty and privacy compliance.

## INDUSTRY MIX: HIGH-STAKES ENVIRONMENTS

Technology & Software respondents were capped at 14% to prevent the 'tech-echo-chamber' effect common in vendor-led research. We over-indexed on industries with the highest regulatory and operational stakes.

| Industry | Weight | Strategic rationale |
|---|---|---|
| Financial Services | 20% | High regulatory exposure and systemic risk contexts. |
| Healthcare & Life Sciences | 16% | Acute sensitivity around HIPAA and patient data. |
| Manufacturing & Industrial | 14% | Operational technology (OT) and IP protection focus. |
| Public Sector & Energy | 18% | Critical infrastructure and unique risk postures. |
| Retail & Professional Services | 18% | Massive consumer data and confidentiality requirements. |
| Technology & Software | 14% | Capped to prevent tech-echo-chamber skew. |

## ENTERPRISE SCALE & COMPLEXITY

Minimum employee count

**5,000+**

Every participating organization

Exceed 25,000 employees

**50%**

Half the cohort — complex environments

Global Enterprise tier

**20%**

75,000+ employees — multi-jurisdictional

Special emphasis was placed on the Global Enterprise tier (75,000+ employees). This segment represents the most complex governance environments in the world, where Shadow AI is not a department-level issue but a sprawling, multi-jurisdictional challenge involving thousands of decentralized endpoints.

### THE QUALITATIVE CORE: 20 DEEP-DIVE INTERVIEWS

While the quantitative poll provided the 'what,' our 20 executive interviews provided the 'how' and 'why.' Discussions were held under Chatham House Rule to ensure candid disclosure of internal struggles, budget shifts, and cultural friction. Participating organizations include:

**FINANCIAL & CONSUMER POWERHOUSES — JPMORGAN CHASE · AMERICAN EXPRESS**
Perspectives on balancing aggressive cybersecurity posture with consumer-facing AI innovation at global scale.

**HEALTHCARE & INDUSTRIAL LEADERS — MAYO CLINIC · KAISER PERMANENTE · SIEMENS**
Complexities of governing AI in life-critical and highly compliant environments where data exposure carries patient risk.

**SCALE & LOGISTICS TITANS — WALMART · SHELL**
Managing AI governance across massive, global, and decentralized workforces operating across multiple jurisdictions.

**HIGH-SENSITIVITY INNOVATORS — LOCKHEED MARTIN · BOEING · DELOITTE**
The intersection of AI with national security, supply chain integrity, and professional client confidentiality.

### DATA INTEGRITY & VALIDATION

#### 01 Identity Verification

Respondents vetted via professional networking profiles to confirm title and company size accuracy.

#### 02 Logic Scrubbing

Surveys with speed-running patterns or contradictory logic were discarded before analysis.

#### 03 Thematic Synthesis

Qualitative transcripts were coded for recurring themes, then used to weight and contextualize quantitative results.

---

This methodology represents a shift away from 'general pulse' surveys. By focusing exclusively on the largest global enterprises — the organizations with the most to lose and the most to gain — this study provides a blueprint for IT and security leaders tasked with building a secure, governed, and productive AI future.

SECTION 1

# Visibility

Q1 · TOOL INVENTORY

## Most large enterprises are running 200+ AI tools — and many can't reliably estimate the number at all.

| Largest response group | Second response | No reliable estimate |
|---|---|---|
| **41%** | **33%** | **18%** |
| Report 200+ tools | Report 51–200 tools | Cannot quantify |

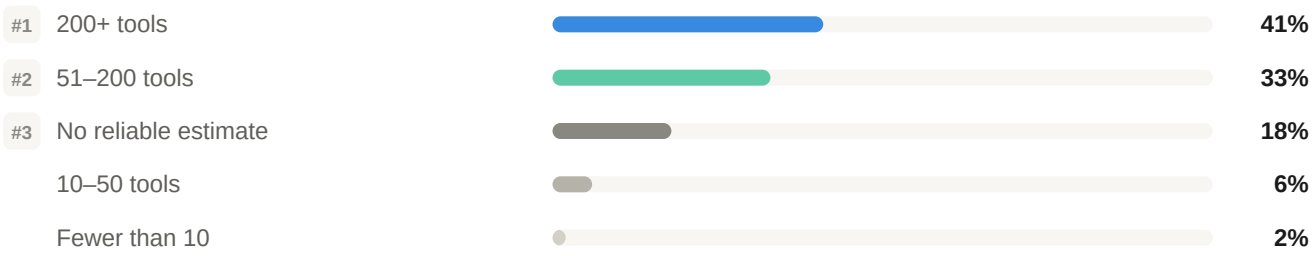


**PEER BENCHMARK**

Organizations reporting 200+ tools are at enterprise-scale AI sprawl — the plurality position, not the outlier. AI arrived embedded in tools already approved: Copilot, Salesforce Einstein, GitHub, Workday. If your count is under 50, your inventory methodology has blind spots.

**RISK SIGNAL · HIGH**

AI tool sprawl creates unmanaged vendor data-sharing agreements, unreviewed training data consent clauses, and shadow exposure across the stack — at a scale most DLP tools were not designed for.

If you're presenting an AI inventory to your board or a regulator, the number should carry a confidence interval, not a clean figure. The CISOs most exposed in the next 24 months are those who reported a tidy count and later discovered it was off by an order of magnitude.

---

Q2 · APPROVAL RATE

## Fewer than one in four AI tools is formally approved at most large enterprises — and nearly 1 in 5 CISOs cannot answer this question at all.

| Largest response | Second response | No visibility |
|---|---|---|
| **44%** | **29%** | **17%** |
| Less than 25% approved | 25–50% approved | Cannot estimate |

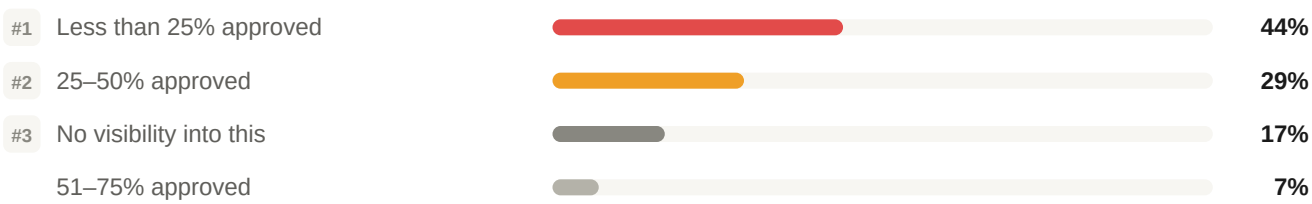

| | |
|---|---|
| More than 75% approved | 3% |

**PEER BENCHMARK**

Less than 25% approved is the plurality answer. Organizations reporting 75%+ represent only 3% of peers — the top maturity tier. CIOs anchor on what was formally procured; CISOs anchor on what they can see on the network. Those are two very different numbers.

**RISK SIGNAL · CRITICAL**

For every sanctioned tool, three are operating without vendor risk assessment, data handling review, or contractual data processing safeguards. This is where regulatory exposure accumulates silently.

Financial services leaders may report higher approval rates — not because shadow usage is lower, but because faster procurement cycles reduce the incentive to go rogue. The underlying shadow usage may be comparable; the visibility into it is what differs.

---

Q3 · AI-GENERATED APPLICATIONS

## Nearly half of large enterprises do not track AI-generated applications — creating an invisible and accumulating technical debt liability.

Don't track this
**48%**
No tracking mechanism

1–10 apps reported
**24%**
Low-count tracking

11–50 apps reported
**16%**
Mid-range visibility

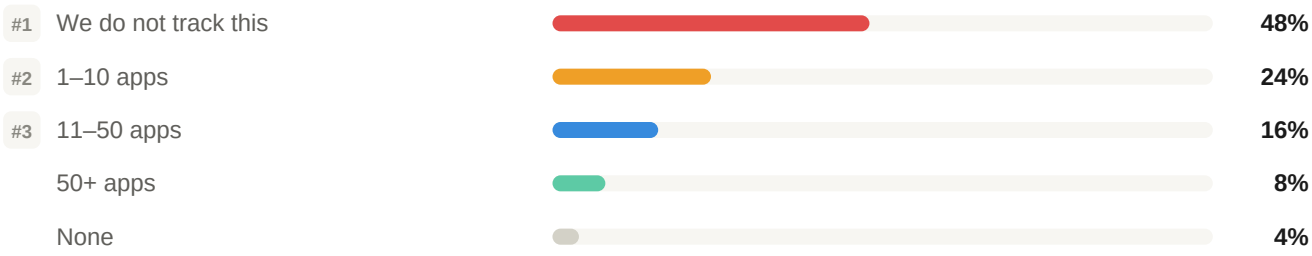


**IF YOU'RE A CISO/CIO**

An AI-generated application leaves no distinguishing metadata in your CMDB, no different flag in SAST tooling, no separate deployment pipeline entry. You are accumulating technical debt you cannot quantify or audit.

**INDUSTRY VARIATION**

Manufacturing and logistics leaders skew most heavily toward 'we do not track this' — the segment with highest risk of AI-generated apps running adjacent to operational technology environments.

The maintenance liability is the underreported risk. Applications built today by non-engineers will require debugging, patching, and decommissioning — by staff who may not understand how they were built, using codebases with no human author to consult.

---

Q4 · INVENTORY CONFIDENCE (1–10 SCALE)

## The average enterprise IT leader rates their AI inventory confidence at 4 out of 10 — meaning most could not produce a reliable inventory under regulatory pressure.

| Most common band | Survey mean score | Score 7+ peers |
|---|---|---|
| **3–4/10** | **4.1/10** | **Top 12%** |
| 38% of respondents | Cross-sector average | Top confidence tier |

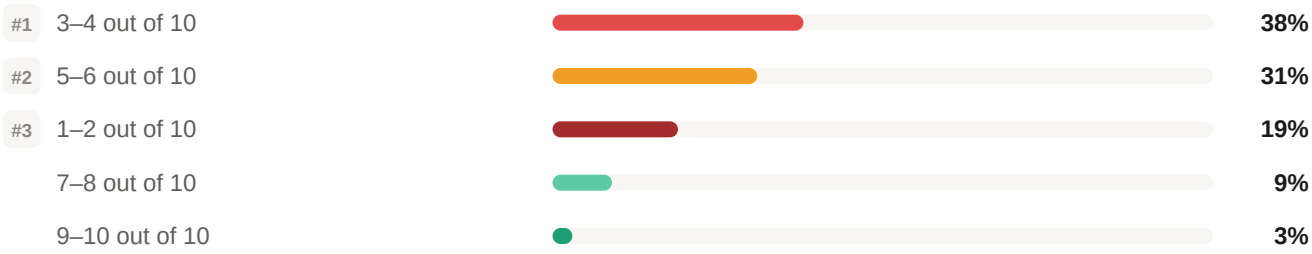


**PEER BENCHMARK**

Scoring 7+ places you in the top 12% for inventory confidence. Scoring 1–2 is the bottom quartile, shared by 19% of respondents — disproportionately government and public sector. CISOs consistently score lower than CIOs: not less capable, but measuring the real environment, not just the sanctioned one.

**RISK SIGNAL · CRITICAL**

A mean score of 4.1 means the average large enterprise CISO could not produce a defensible AI inventory within 24 hours of a regulatory request. This is board-level exposure, not an operational gap.

Any AI inventory presented to your board should explicitly state what it covers and what it excludes. A board that later discovers the inventory missed browser-based AI usage, personally expensed subscriptions, or embedded SaaS AI features will have legitimate questions about everything else you reported.

---

PERIDOT AI GOVERNANCE RESEARCH

SECTION 2

# Shadow AI Reality

Q5 · UNSANCTIONED TOOL DETECTION

## 39% of large enterprises are not actively tracking unsanctioned AI — meaning their incident count reflects what they found, not what exists.

| Not tracking | Found 1–5 instances | Found 6–20 instances |
|---|---|---|
| **39%** | **31%** | **17%** |
| No active tracking | Likely undercounting | Active monitoring signal |

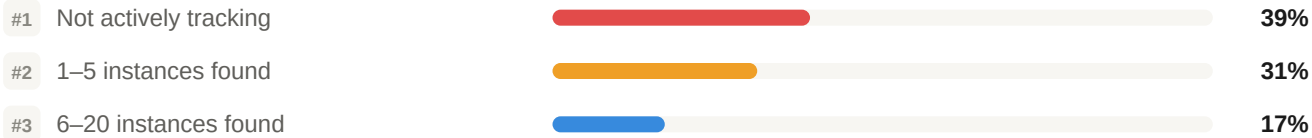

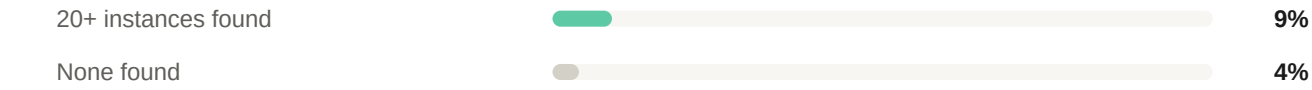


**IF YOU'RE A CISO/CIO**

Organizations reporting 20+ incidents are not the worst environments — they are the best-monitored ones. Higher detection reflects deployed tooling, not worse posture. If your count is zero, the honest question is whether you have looked.

**RISK SIGNAL · CRITICAL**

In a post-incident regulatory proceeding, 'we were not tracking it' is not a defensible posture. The question will be whether you should have known — and at enterprise scale, the answer is almost certainly yes.

The 1–5 count is an artifact of detection capability, not actual prevalence. Organizations with active AI monitoring find 10–20x more instances than those running passive or periodic reviews.

*Peridot observation: Based on patterns observed across enterprise environments working to regain control over AI usage, these findings are consistent with what our teams encounter during initial governance assessments.*

---

Q6 · DEPARTMENT DISTRIBUTION (SELECT ALL THAT APPLY)

## Engineering and marketing lead unsanctioned AI usage — but finance's third-place ranking signals a data exposure risk most governance programs have underweighted.

Engineering / Dev
**71%**
Top department

Marketing / Content
**68%**
Second department

Finance / Accounting
**52%**
Third — underweighted risk


**INDUSTRY VARIATION**

Healthcare skews toward higher HR representation. Financial services shows elevated legal and compliance usage. Technology companies show near-universal engineering department usage. C-Suite figure is systematically underreported — actual prevalence estimated significantly higher.

**RISK SIGNAL · HIGH**

Senior leaders are among the heaviest consumer AI users, operate with least device oversight, and handle the most sensitive strategic information. A combination that creates high-impact exposure that DLP tools are almost certainly not catching.

The finance number deserves more attention than it typically gets. Financial analysts and FP&A; teams are using AI for modeling, variance analysis, and reporting — with data that sits at the intersection of your most sensitive financial information and your most significant regulatory obligations.

Q7 · NON-TECHNICAL AI APP CREATION

## 65% of large enterprises confirm non-technical employees are creating and deploying AI-powered tools — the vibe coding risk is already inside the enterprise.

| Combined 'yes' response | Yes, frequently | Not monitoring this |
|---|---|---|
| **65%** | **27%** | **21%** |
| Frequently or occasionally | Active app creation | No visibility |

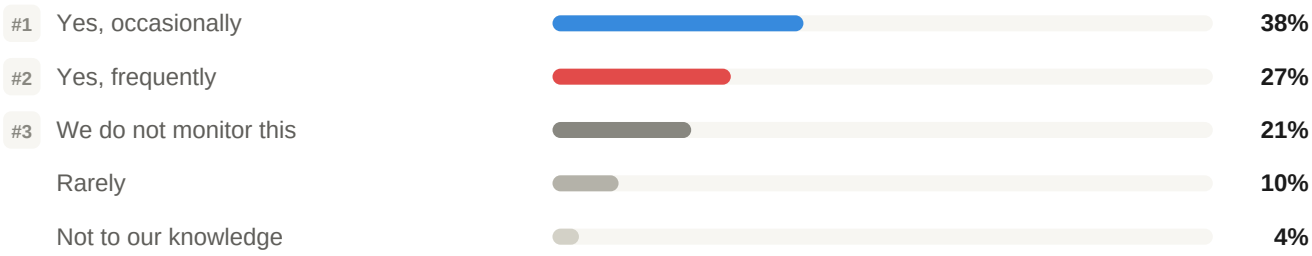


**PEER BENCHMARK**

Technology companies report 'frequently' at the highest rate of any industry. Healthcare and manufacturing are most likely to report 'we do not monitor this' — meaning actual prevalence may exceed the already high cross-sector average.

**IF YOU'RE A CISO/CIO**

The instinct to prohibit non-technical AI development is understandable but strategically counterproductive. The most effective organizations build lightweight review processes that operate at business speed — making the governed path easier than the ungoverned one.

The barrier to building a deployable application has dropped from months of engineering time to an afternoon of natural language prompting. Your marketing manager can build a customer-facing chatbot. None of them are thinking about input validation, access controls, or data residency when they do it.

PERIDOT AI GOVERNANCE RESEARCH

SECTION 3

# Risk & Incidents

Q8 · DATA EXPOSURE INCIDENTS

## 43% of large enterprises haven't audited for AI-related data exposure — operating on assumption rather than evidence that no incident has occurred.

| Haven't audited | Suspected, unconfirmed | Confirmed incident |
|---|---|---|
| **43%** | **28%** | **7%** |
| No audit conducted | Signals without confirmation | Likely undercounted |

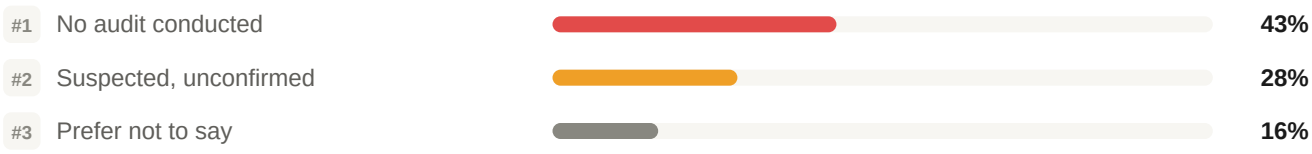


**RISK SIGNAL · CRITICAL**

The 28% 'suspected, unconfirmed' implies a material uninvestigated exposure signal at more than 1 in 4 large enterprises. For publicly traded organizations, this may constitute an undischarged SEC disclosure assessment obligation.

**IF YOU'RE A CISO/CIO**

'We haven't audited for it' is not a defensible posture at this scale. If an incident is traced to an unsanctioned AI tool, regulators and plaintiffs' counsel will ask not whether you knew — but whether you should have known.

The 'prefer not to say' category will be disproportionately represented by publicly traded companies with disclosure sensitivities. If you have a suspected exposure touching material information, your legal team needs to be in the conversation now.

---

Q9 · DETECTION & INVESTIGATION SPEED

## 60% of large enterprises cannot commit to detecting and investigating an AI incident within one week — incompatible with GDPR's 72-hour breach notification requirement.

Cannot commit to 1 week
**60%**
Longer or no process

No defined process
**24%**
Would improvise response

Within hours
**3%**
Top maturity tier

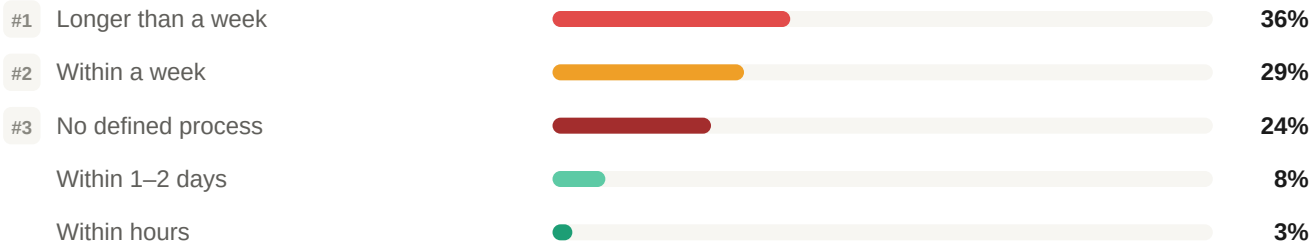


**RISK SIGNAL · CRITICAL**

GDPR requires notification within 72 hours. HIPAA within 60 days of discovery. SEC rules require material incident disclosure within 4 business days of determining materiality. A greater-than-one-week detection timeline structurally prevents compliance in many scenarios.

**IF YOU'RE A CISO/CIO**

The cost of building an AI incident response playbook before an incident is a fraction of the cost of building it during one — under regulatory time pressure, without a defined process, with legal counsel on the call.

Most organizations have not built AI-specific incident response playbooks. They would retrofit existing data breach procedures — adding significant time. Healthcare and financial services cluster around 'within a week' due to existing breach response infrastructure. Manufacturing and retail dominate 'no defined process.'

---

Q10 · DATA EXPOSURE TYPES (SELECT TOP THREE)

## Source code and IP rank second in exposure concern — yet most enterprises have no controls preventing developers from pasting proprietary code into AI assistants daily.

| Customer PII | IP / Source code | Credentials / Auth |
|---|---|---|
| **67%** | **61%** | **27%** |
| Top exposure concern | Second — often uncontrolled | Most underestimated risk |

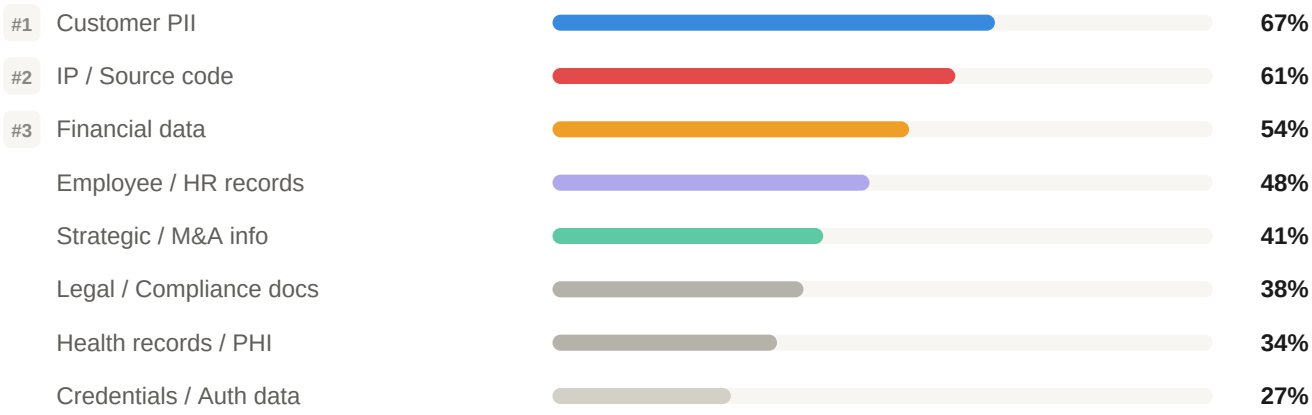


**INDUSTRY VARIATION**

PHI ranks first in healthcare. Source code ranks first among technology companies. Financial data ranks first in financial services. Credentials are systematically underestimated across all sectors.

**RISK SIGNAL · HIGH**

Credential exposure is the most underestimated finding in this section. AI-generated code has a documented tendency to include hardcoded API keys and authentication tokens. If developers are using AI to write infrastructure code without output scanning, you likely have credentials in production you haven't checked.

* Credential figure is systematically underestimated. Most security leaders do not intuitively connect AI tool usage to credential exposure, but AI tools trained on code repositories frequently produce code with embedded secrets.

*Peridot observation: Based on patterns observed across enterprise environments working to regain control over AI usage, these findings are consistent with what our teams encounter during initial governance assessments.*

SECTION 4

# Governance & Accountability

Q11 · POLICY STATUS

## More than half of large enterprises are operating without an effective, enforced AI governance policy — with 34% still in development as the environment expands daily.

No effective policy
**55%**
In development or nonexistent

Policy in development
**34%**
Largest single response

Comprehensive + enforced
**9%**
Top maturity tier

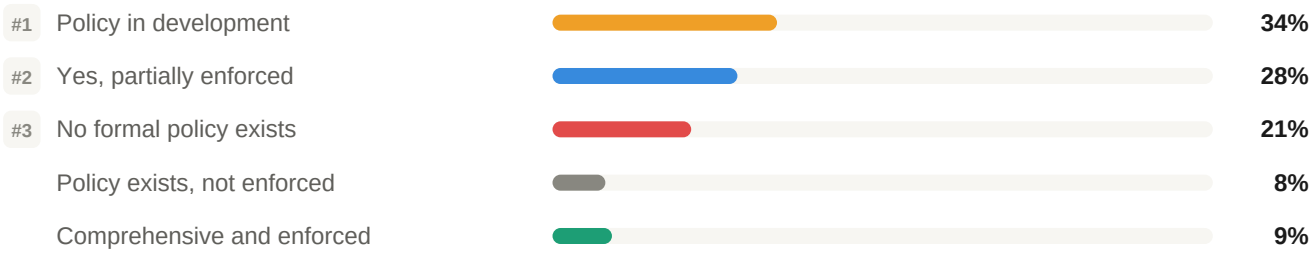


**PEER BENCHMARK**
Only 9% of peers report comprehensive, enforced AI policy. Financial services organizations are disproportionately represented in this tier, driven by OCC, SEC, and FRB guidance. Technology companies show weaker enforcement despite more sophisticated IT functions.

**RISK SIGNAL · CRITICAL**
A documented policy that is demonstrably not enforced can create worse legal exposure than no policy at all — establishing awareness of the risk without demonstrating remediation. In litigation or regulatory proceedings, this is a material distinction.

Policies in development for more than six months in a fast-moving technology landscape are not governance — they are documentation of intention. Meanwhile, the environment they are meant to govern continues to expand daily.

Q12 · INCIDENT OWNERSHIP

## 60% of large enterprises have no clear accountability framework for AI-generated application incidents — a gap that will be resolved reactively for most organizations.

| Ownership unclear/undefined | Partially defined | Clearly defined |
|---|---|---|
| **60%** | **33%** | **7%** |
| No or not defined | Incomplete framework | Documented and tested |

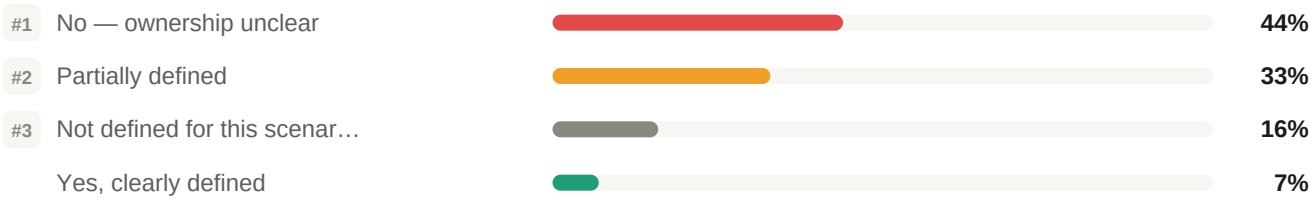


**INDUSTRY NOTE**

Financial services leaders under SR 11-7 model risk management guidance face a specific gap. AI-generated applications performing calculations or automating decisions may meet the functional model definition — without the ownership, validation, and monitoring infrastructure that guidance requires.

**IF YOU'RE A CISO/CIO**

The 7% who report clearly defined ownership almost all share one characteristic: they experienced a prior AI-related incident and were forced to resolve the ownership question in real time. Don't wait for the incident to define the framework.

AI-generated applications blur every category that accountability frameworks rely on. Who is the developer? The employee who wrote the prompts. Who approved deployment? In many cases, no one in IT. Existing RACI matrices and SDLC policies were not designed for this scenario.

---

Q13 · REGULATORY PREPAREDNESS

## 57% of large enterprise IT leaders characterize their AI governance posture as having significant gaps or being outright unprepared for regulatory scrutiny.

| Significant gaps or unprepared | Significant gaps | Fully prepared |
|---|---|---|
| **57%** | **41%** | **6%** |
| Combined response | Largest single response | Top maturity tier |

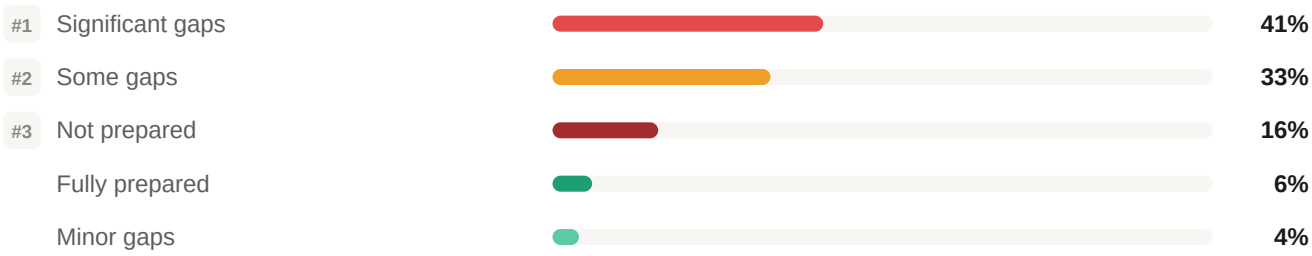

**PEER BENCHMARK**

Financial services organizations are the most prepared — but still cluster around 'some gaps' because AI governance specifically lags behind existing model risk frameworks. Government and public sector report the worst preparedness despite often having the most formal compliance requirements.

**RISK SIGNAL · CRITICAL**

The regulatory environment is accelerating faster than most internal governance programs. The SEC, EU AI Act, OCC/FRB/FDIC joint guidance, and state-level AI legislation are all moving on parallel tracks. The question is not whether regulators will ask — it is whether you will be ready.

This question produces the sharpest industry variation in the survey. Being ahead of manufacturing or retail on AI governance does not mean being ready for the specific questions regulators are now asking about generative AI, third-party AI vendor management, and AI-assisted decision-making in regulated contexts.

---

PERIDOT AI GOVERNANCE RESEARCH

SECTION 5

## Tooling & Control

Q14 · SECURITY TOOL CAPABILITY

### 75% of large enterprises cannot fully detect what data employees are sending to external AI services — making policy enforcement structurally impossible at most organizations.

Partial or no detection
**75%**
Cannot see full data flow

No capability
**29%**
Tools not designed for this

Full capability
**5%**
AI-specific tooling deployed

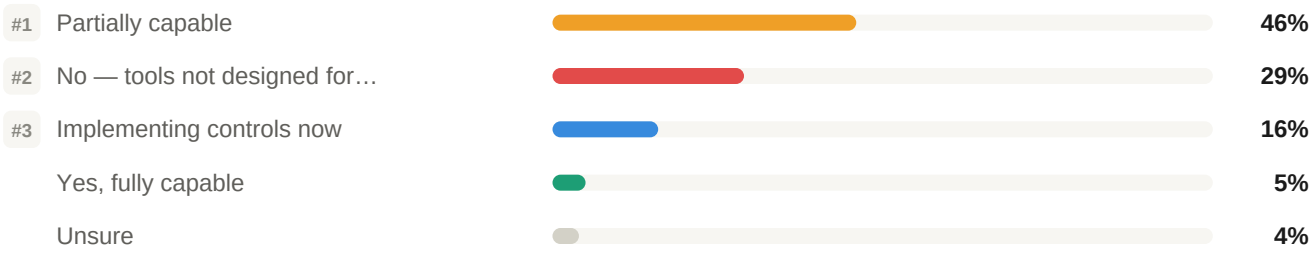


**PEER BENCHMARK**

Most DLP and CASB tools can block known AI domains at the network layer — but cannot inspect encrypted browser sessions, detect personal account usage, or monitor API calls from shadow applications. 'Partial' typically means the lowest-sophistication controls are in place, not that the gap is manageable.

**RISK SIGNAL · CRITICAL**

You cannot enforce a policy you cannot monitor. You cannot investigate an incident you cannot detect. You cannot quantify data exposure risk if you cannot see your data flows. The governance problem is real, but the tooling problem is what makes it intractable.

Evaluate AI security tooling against your specific gap — browser session inspection, data classification in prompts, or shadow application detection — rather than a generic checklist. The vendor landscape has matured significantly in the past 18 months and point solutions now claim enterprise-grade capability. Test that claim against your actual exposure vector.

---

Q15 · PRIMARY BARRIER (SELECT ONE)

## The #1 barrier to Shadow AI control is adoption speed outpacing governance — a structural asymmetry that no policy document alone will resolve.

Speed vs. policy
**34%**
Top barrier selected

Lack of visibility
**27%**
Second barrier

"We have it under control"
**1%**
Near-zero confidence

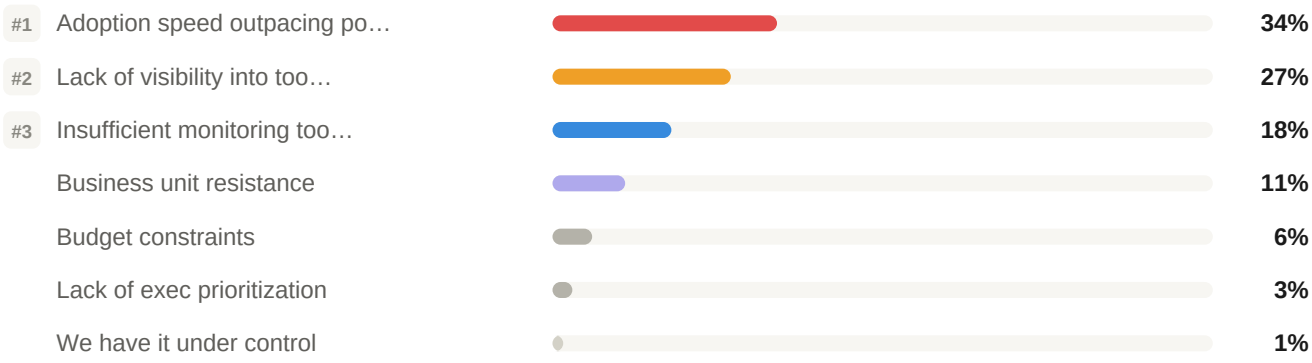


**INDUSTRY VARIATION**
Business unit resistance ranks fourth overall but first in technology companies, where engineering culture creates resistance that policy alone cannot overcome. Budget ranks higher at mid-sized enterprises in the 5,000–10,000 band. 'We have it under control' selected by under 1% of respondents.

**RISK SIGNAL · HIGH**
AI adoption is driven by individual behavior at consumer technology speed. AI governance is driven by institutional process at organizational change management speed. These are not compatible speeds, and the gap is widening, not narrowing.

Organizations managing Shadow AI most effectively are not those with the most comprehensive policies — they are those that built lightweight, fast-cycle governance: AI tool fast-track procurement, developer sandbox environments, usage monitoring with automated alerting. The goal is to make the governed path easier than the ungoverned one.

*Peridot observation: Based on patterns observed across enterprise environments working to regain control over AI usage, these findings are consistent with what our teams encounter during initial governance assessments.*

SECTION 6

# Forward Outlook

Q16 · VOLUME TRAJECTORY

## 86% of large enterprise IT leaders expect AI-generated application volume to increase significantly over 12 months — against a backdrop of controls that are currently inadequate at most organizations.

| Expect increase | Increase significantly | Expect decrease |
|---|---|---|
| **86%** | **52%** | **2%** |
| Significant + somewhat combined | Majority response | Statistically negligible |

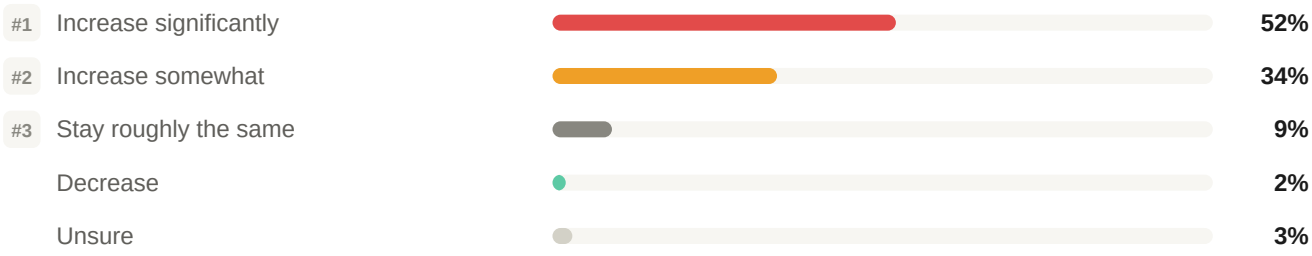


**RISK SIGNAL · CRITICAL**

Read this alongside the governance gaps documented throughout this report: 60% without clear incident ownership, 57% unprepared for regulatory scrutiny, 75% without full AI data flow visibility. A risk surface that 86% of leaders expect to grow significantly, while controls remain materially incomplete, is the definition of an accelerating exposure.

**IF YOU'RE A CISO/CIO**

This is your most actionable board-level data point. The question your board should be asking is not whether AI-generated application volume will increase — it will. The question is whether your governance infrastructure will grow at a comparable rate. The data suggests it will not, at most organizations.

Technology companies approach 95% combined 'increase.' Manufacturing and industrial show the widest spread. The 'decrease' answer is statistically negligible — under 2%. This finding pairs with the governance gaps surfaced earlier: the problem is accelerating while the controls lag.

Q17 · BOARD-LEVEL ATTENTION

## Only 36% of large enterprises discuss AI risk at board level on a regular basis — the majority treat it as an occasional reactive topic rather than a governed risk.

Regular board discussion
**36%**
Quarterly or more

Occasionally
**39%**
Largest response — reactive

Rarely or never
**18%**
Not a standing topic

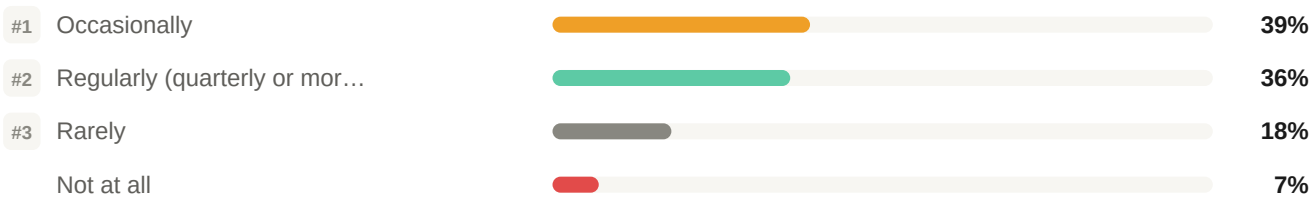


**PEER BENCHMARK**
Financial services boards lead all sectors — driven by regulatory pressure and direct revenue connection to AI decisions. Manufacturing and retail show highest rates of 'rarely.' That framing will change when a significant incident forces the reframe. The gap between 'occasionally' and 'regularly' is a difference in governance models, not degree.

**IF YOU'RE A CISO/CIO**
Boards do not sustain attention on topics they cannot measure. To move from 'occasionally' to 'regularly,' build a reportable metrics framework — sanctioned vs. unsanctioned tool count, policy coverage rate, detection capability score, incident count. A dashboard creates the standing agenda item that a one-time presentation cannot.

Quarterly board-level discussion implies a reporting framework, defined metrics, and executive accountability for progress. Occasional discussion implies the topic surfaces reactively after an industry incident or regulatory announcement — without systematic oversight infrastructure. These are not points on a spectrum; they are different governance models.

ABOUT

## peridot

Enterprise AI Governance Platform

Peridot helps enterprise IT and security teams discover, monitor, and govern AI usage across their organization — including AI tools, data flows, and AI-generated applications.

- Discover all AI tools in use — sanctioned and unsanctioned — across every environment
- Monitor data flow into external AI services with real-time visibility and alerting
- Track AI-generated applications across cloud, on-prem, and hybrid infrastructure
- Enable audit-ready AI governance with policy enforcement and compliance reporting

WHAT THIS MEANS FOR YOU

If your responses align with the majority of organizations in this report:

- You likely have 2-4x more AI tools in use than your current inventory reflects.
- You likely have active data exposure pathways your current security stack cannot detect.
- You are likely not audit-ready for a regulatory, board, or executive-level inquiry.

**If your CEO or board requested a complete AI governance review this week, your current posture would likely not hold up under scrutiny.**

*The gap between AI adoption and governance is widening — not stabilizing. Delayed visibility compounds risk.*

## Close Your AI Governance Gaps

Most organizations in your position uncover critical gaps within the first 30 minutes of an audit — including unknown AI tools, unmonitored data exposure, and policy gaps that would not withstand regulatory review.

**Request Your AI Governance Audit (48-hr turnaround)**

No prep required. No disruption to your team. | peridot.company/audit
Used by security teams preparing for board reviews and regulatory audits.

*Survey conducted among 400 enterprise IT and security leaders at organizations with 5,000+ employees across 8 industries. Respondent mix: 25% CISO/Deputy CISO, 20% CIO/VP IT, 20% VP/Director of Security, 15% Head of IT Risk/GRC, 12% Enterprise Architect/Head of IT Strategy, 8% CDO/CPO. All figures represent cross-sector averages unless otherwise noted.* 