Kiteworks

2026

Data Security and Compliance Risk Forecast Report

AI Adoption Is Accelerating. Governance Is Stalling. The Reckoning Is Coming.

REPORT

Contents

Executive Summary

The Year Everything Gets Real

2026 is the year AI data security moves from "emerging concern" to "operational reality." Every organization we surveyed—every single one—has agentic AI on their roadmap. The question isn't whether AI will touch your sensitive data. It already does.

The uncomfortable truth: most organizations aren't ready. They've started the work. Very few have finished it.

This report identifies 15 predictions for enterprise data security in 2026, based on a survey of 225 security, IT, and risk leaders across 10 industries and 8 regions. What we found is a market in transition: significant gaps in AI-specific capabilities, and a widening divide between organizations with board attention on AI governance and those without.

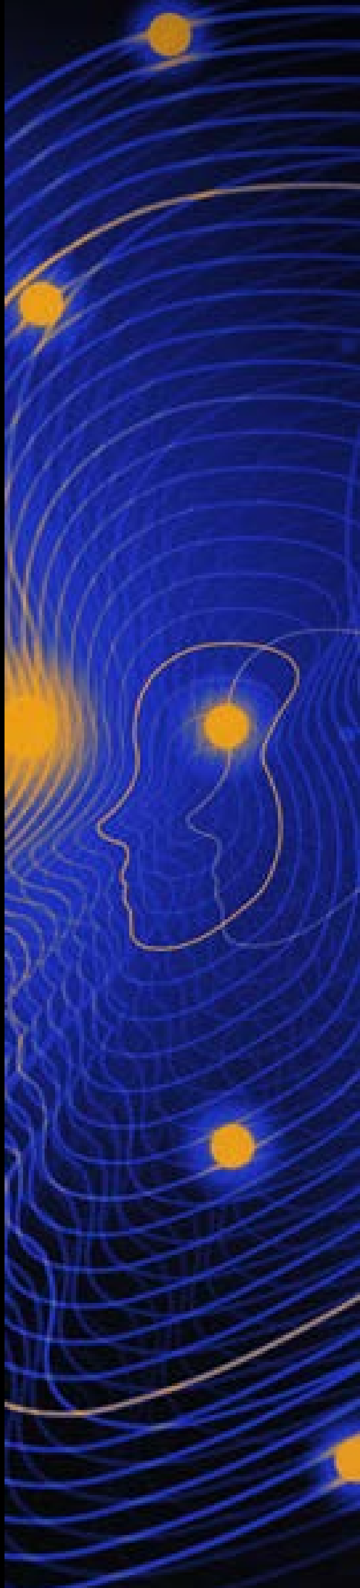

The 15 Predictions at a Glance

	Prediction	Key Metric	Confidence
1	DSPM becomes the default baseline	**61%** can't enforce tagging	High
2	Data governance goes "managed-by-default"	**37%** below Managed maturity	High
3	Centralized AI gateways become the control plane	**57%** non-centralized; **33%** of government has no dedicated AI controls	High
4	Agentic AI goes mainstream	**100%** on roadmap; **37%** to **40%** have containment	High
5	Containment controls become the battleground	**63%** lack purpose binding; **60%** lack kill switch	Medium
6	AI risks dominate the security agenda	**30%** cite third-party AI; only **36%** have visibility	High
7	Supply chain expands to AI attestations	**72%** no SBOM; legacy MFT can't support AI	Medium
8	Third-party risk pivots to visibility	**89%** never practiced IR with partners	High
9	IR becomes AI-infused	**60%** lack AI anomaly detection	Medium
10	Audit trails become the keystone	**33%** lack trails; **61%** fragmented logs	High
11	Training-data controls become regulatory requirements	**78%** can't validate; **53%** can't recover	High
12	AI governance hits every boardroom	**54%** of boards not engaged	High
13	EU AI Act creates a global template	**22-33 point** control gap	High
14	PQC moves to mainstream	**84%** haven't implemented	Medium
15	Data sovereignty becomes AI imperative	**29%** cite cross-border AI exposure	High

Gaps That Matter

Containment Gap: Organizations have invested in watching what AI does—human-in-the-loop (59%), monitoring (58%). They haven't invested in stopping it—kill switch (40%), purpose binding (37%). That's a **15-20-point gap** between observing and acting. **60%+** can't terminate a misbehaving AI agent or enforce purpose limitations.

Keystone Capabilities: Evidence-quality audit trails and AI training-data recovery predict overall maturity better than industry, region, or size. Organizations with audit trails show **+20-32-point advantages** on every AI metric. But **61%** have fragmented logs across systems—not actionable evidence.

Board Effect: **54%** of boards aren't engaged on AI governance. Those organizations are **26-28 points** behind on every AI maturity metric. This is the strongest correlation in the survey.

Data Sovereignty Gap: Organizations have solved sovereignty for storage—not for AI processing. **29%** cite cross-border AI transfers as exposure, but only **36%** have visibility into where data is processed, trained, or inferred.

Critical Outliers

Government is a generation behind: 90% lack purpose binding, 76% lack kill switch, 33% have no dedicated AI controls—while handling citizen data and critical infrastructure.

Australia is the benchmark: +10-20 points on nearly every metric, with the strongest pipelines. Leading on AI adoption AND controls.

15 Predictions

PREDICTION #1:

DSPM Becomes the Default Data-Protection Baseline

By the end of 2026, DSPM will be a baseline expectation for mid-to-large enterprises—but most will still be struggling with enforcement. The topline numbers mask the real problem: 86% have DSPM protocols in place, but only 39% can enforce tagging and classification across channels. Having a DSPM tool is one thing. Making it work requires consistent data classification, policy enforcement, and coverage across every channel where sensitive data moves. Most organizations aren't there.

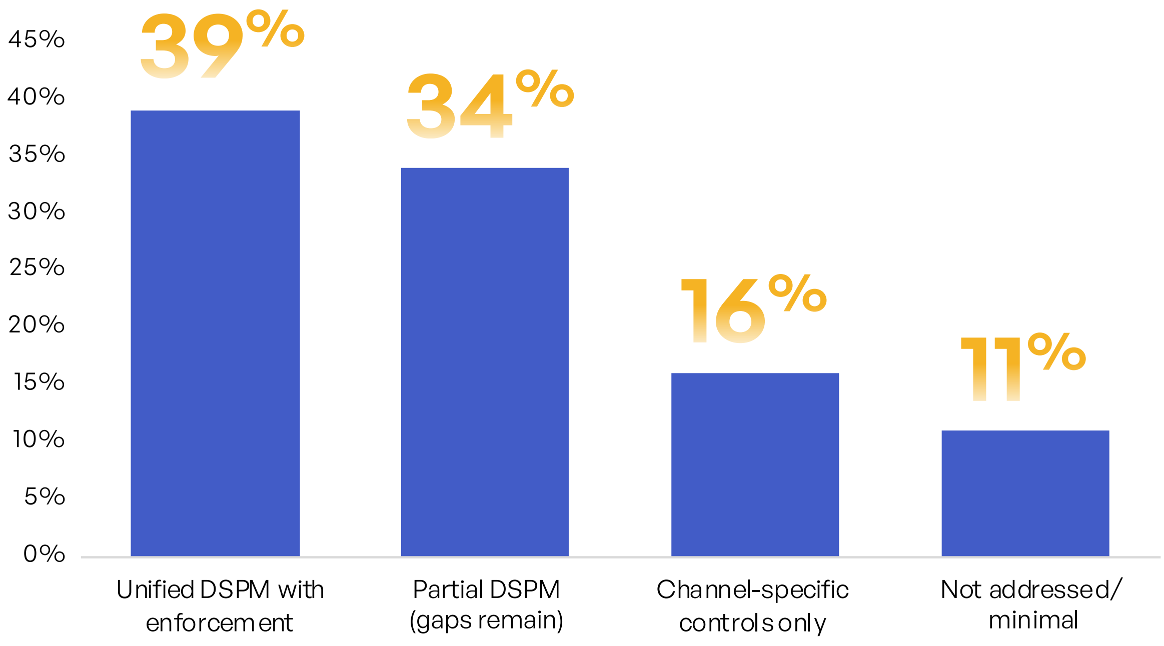

DSPM Effectiveness Level

61% can't enforce tagging consistently. 34% have partial coverage with known gaps—the tools are deployed but classification isn't propagating across systems, or policies aren't triggering when they should. Another 16% have only channel-specific controls: data classified in one system loses its tags when it moves to another. And 11% have nothing meaningful in place.

Segment	Operationalization Gap	Not Yet Using
Government	**34 points**	48%
Global Average	22 points	36%
UAE	19 points	31%
Australia	18 points	22%

Government has the widest gap: 86% have protocols on paper, but only 52% are using them operationally—and "using" doesn't mean "enforcing." Nearly half of government organizations have DSPM policies sitting in documentation while sensitive data flows untagged through production systems. Even Australia, the leader, still has 22% not yet operational.

The uncomfortable truth: DSPM without enforcement is just expensive monitoring. By the end of 2026, most organizations will have DSPM. Far fewer will have closed the gap between detecting sensitive data and controlling where it goes.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that DSPM becomes expected; lower confidence on closing the enforcement gap)

PREDICTION #2:

Data Governance Operating Models Go "Managed-by-Default"

"Managed" governance maturity will be the baseline expectation—but most organizations won't meet it. The aspiration is policy-as-code underpinning DSPM, IR, and compliance. The reality: 37% of organizations are still below "Managed" maturity, running governance models that exist on paper but don't execute consistently.

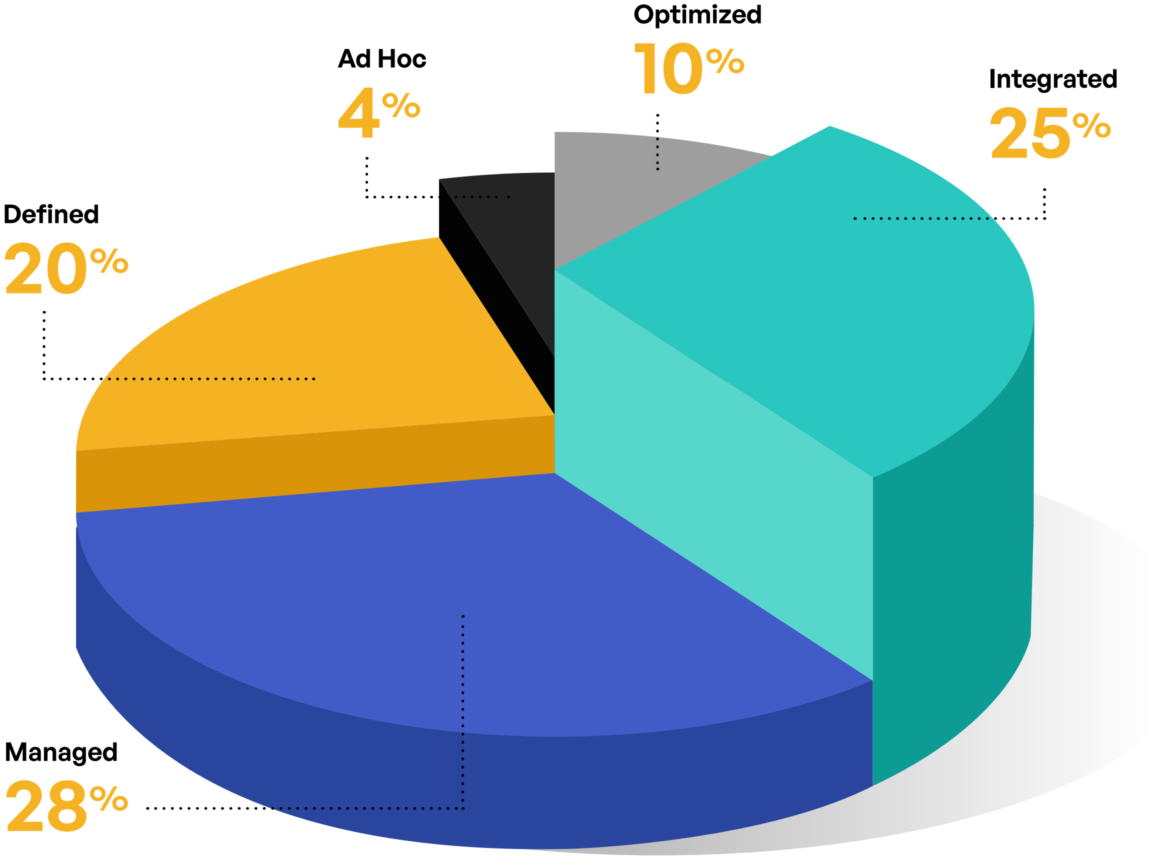

Maturity Levels

Only 28% have reached "Managed"—defined metrics, consistent execution, some automation. Below that, 20% are stuck at "Defined" (policies documented but not reliably followed) and 4% remain ad hoc. That's nearly a quarter of organizations where governance is more aspiration than operation. Even the 25% at "Integrated" often have gaps between what the model says and what happens.

Note: Percentages shown exclude an additional 13% of respondents who selected "No response / Not applicable." These respondents are treated as "below Managed" in the 37% figure.

Industry	Governance-to-Automation Gap	Still Manual/Periodic
Government	**24 points**	38%
Healthcare	**20 points**	32%
Financial Services	6 points	15%
Technology	4 points	12%

Healthcare and Government show the widest gaps. Government has 86% with formal governance models but only 62% using automated compliance—a 24-point chasm. These organizations have the documentation. They don't have the automation to make it real. 38% of government organizations still rely on manual or periodic compliance processes, which means evidence collection happens quarterly or annually rather than continuously.

25% of all organizations still use manual or periodic compliance as their primary approach. In a regulatory environment that increasingly expects continuous evidence, periodic compliance is a liability waiting to surface.

The uncomfortable truth: Most organizations have governance models they can't operationalize at the speed their AI deployments and regulatory requirements demand.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that "Managed" becomes the expectation; lower confidence that most will get there)

PREDICTION #3:

Centralized AI Data Gateways Become the Control Plane for AI

Centralized AI data gateways will be the expected architecture for governing sensitive data flowing through models and agents. Most organizations aren't there. Only 43% have a centralized gateway today. The remaining 57% are fragmented, partial, or flying blind.

26% are partial, ad hoc, or have nothing at all. 19% have cobbled together point solutions without coherent policy—controls that made sense when they had one AI pilot but don't scale to five or 10 use cases running simultaneously. And 7% of enterprises have no dedicated controls whatsoever for how AI systems access sensitive data. These organizations have deployed AI. They just haven't governed it.

Even the 27% with "distributed controls and clear policies" face a scalability problem. Distributed works when you have one copilot. It doesn't work when you're running internal copilots, workflow agents, API integrations, document generation, and decision-making systems across multiple business units—each with its own policy interpretation.

AI Data Governance Approach

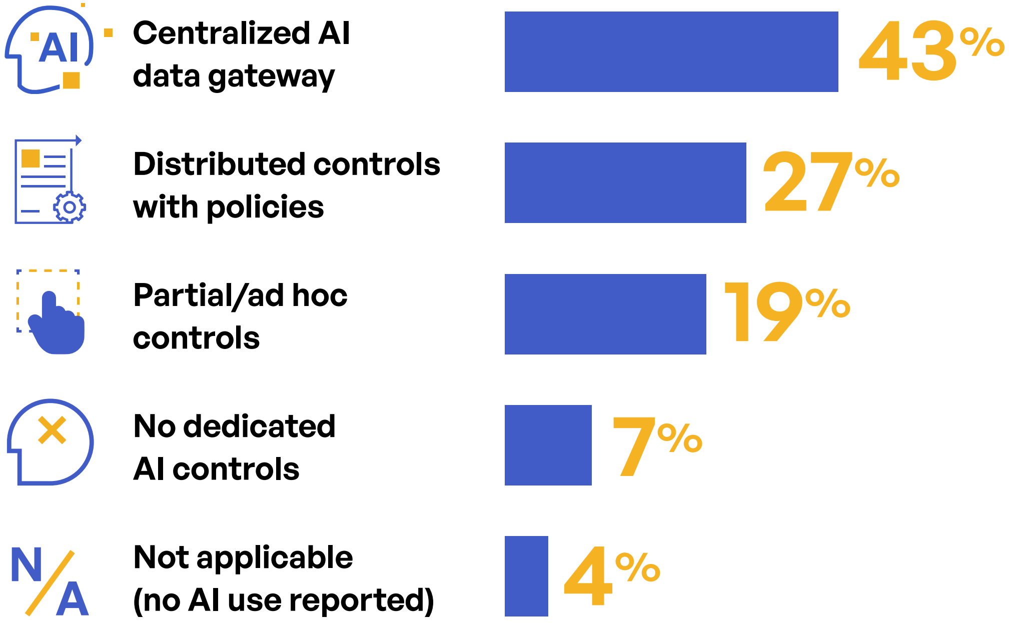

Industry	No Centralized Gateway	No Dedicated Controls
Government	**90%**	**33%**
Healthcare	77%	14%
Financial Services	60%	5%
Technology	44%	3%
Professional Services	33%	0%

Government is the crisis. 90% lack centralized AI governance. One-third have no dedicated AI data controls at all—not partial, not ad hoc, nothing. These are organizations handling citizen data, classified information, and critical infrastructure. AI is already in these environments. Governance isn't.

Healthcare isn't far behind: 77% without centralized gateways and 14% with nothing dedicated to AI. Even Financial Services—heavily regulated, highly targeted—has 60% without centralization and 5% with no dedicated controls.

The gap between AI deployment velocity and AI governance maturity is widening. Most organizations will spend 2026 trying to retrofit centralized controls onto AI systems that were deployed without them.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that centralized gateways become the expected architecture; lower confidence that most will close the gap, particularly in government and healthcare)

PREDICTION #4:

Agentic AI Use Cases Go Mainstream—and Touch Critical Channels

AI and agents will be embedded into core business and security workflows in every industry. Every organization in our survey has agentic AI on their roadmap—zero exceptions. The problem isn't adoption. It's that organizations are deploying AI far faster than they're governing it.

Use Case	Existing or Planned	Controls Typically in Place
Internal copilots	39%	Moderate
File/document generation	34%	Low-Moderate
Data extraction/ enrichment	34%	Low
Email composition	33%	Low
API/integration agents	33%	Low
Autonomous workflow agents	33%	**Very Low**
SFTP/MFT automation	27%	**Very Low**
Decision-making agents	24%	**Very Low**

A third of organizations are planning autonomous workflow agents—systems that take actions without human approval for each step. A quarter are planning decision-making agents. These aren't chatbots. These are systems that will access sensitive data, integrate with critical infrastructure, and execute business logic autonomously. Yet purpose binding sits at 37% and kill switches at 40%. Organizations are deploying agents they can't constrain or terminate.

The MFT channel is a particular concern: 27% are planning AI-driven MFT automation, but MFT security adoption is only 46%. More than half of organizations lack adequate MFT security—and they're about to add autonomous agents to that channel.

Industry	No API Agents Planned	No Decision Agents Planned	Modernization Gap
Government	**95%**	90%	Severe
Healthcare	82%	86%	Significant
Financial Services	65%	73%	Moderate
Technology	53%	69%	Lower

Government faces a different problem: 95% have no API agents planned, and 90% have no decision-making agents on the roadmap. While this might look like prudent caution, it's also a modernization gap that will widen as other sectors automate. When Government does adopt—and it will—organizations will be starting from zero on both deployment and governance.

Healthcare's conservatism (82% without API agents planned) may provide a temporary buffer, but it also means less experience with AI governance when adoption accelerates. The organizations deploying cautiously now aren't necessarily building the governance muscles they'll need later.

The uncomfortable reality: 100% of organizations have AI on the roadmap, but only 37% to 40% have the containment controls to manage it when something goes wrong.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that agentic AI goes mainstream; high confidence that governance will lag deployment through 2026)

Sidebar: AI Agent Swarms Move From Theory to Field Use

In mid-September 2025, Anthropic reported detecting and disrupting a cyber-espionage operation it attributes (with high confidence) to a Chinese state-sponsored group it calls GTG-1002. The actor used Claude Code plus Model Context Protocol (MCP) tools and ran multiple Claude instances in groups as autonomous "orchestrators" to execute major parts of the intrusion life cycle—reconnaissance, vulnerability discovery, exploitation, lateral movement, credential harvesting, and data analysis.[1]

Anthropic says the campaign targeted ~30 entities and that AI executed ~80-90% of tactical work, with humans stepping in only at a few critical decision points (roughly 4-6 per campaign)—for example, approving escalation from recon to exploitation and deciding what to exfiltrate.

One defensive insight: Anthropic observed the AI sometimes overstated findings or fabricated data (e.g., "working" credentials that failed), forcing validation and slowing attackers down.

What to do now: treat agent runtimes + tool connectors as privileged infrastructure—lock down who/what can run tools, enforce allowlists, monitor high-rate automation, and maintain a fast "kill switch" for suspicious agent activity.

[1] Anthropic, "Disrupting the first reported AI-orchestrated cyber espionage campaign" (Nov. 13, 2025) and full report (Nov. 17, 2025).

PREDICTION #5:

Containment Controls Become the AI Security Battleground

63% of organizations can't enforce purpose limitations on AI agents. 60% can't quickly terminate an agent that's misbehaving. 55% can't isolate AI systems from broader network access. These are the containment controls—the ability to stop AI when something goes wrong—and they're the largest gaps in the entire survey.

Control	Not In Place	Pipeline	Projected Still Missing (2026)
Purpose binding	**63%**	39%	~24-36%
Kill switch	**60%**	34%	~26-36%
Network isolation	**55%**	34%	~21-31%
Input validation	**54%**	36%	~18-28%
Data minimization	44%	33%	~11-18%
Continuous monitoring	42%	24%	~18-25%
Human-in-the-loop	41%	23%	~18-24%

The governance-vs.-containment gap is the central problem. Organizations have invested in watching—human-in-the-loop (59%), continuous monitoring (58%), data minimization (56%). They haven't invested in stopping. Purpose binding, kill switches, and network isolation all trail by 15 to 20 points. Most organizations can observe an AI agent doing something unexpected. They can't prevent it from exceeding its authorized scope, quickly shut it down, or isolate it from sensitive systems.

Governance
(monitoring, human-in-loop, minimization)

56-59% **Gap: Moderate**

Containment
(kill switch, purpose binding, isolation)

37-45% **Gap: Severe**

The pipelines are the largest in the survey—39% for purpose binding, 34% for kill switch. Organizations have identified precisely the right gaps. But pipelines don't equal execution. Historically, 60-70% of security roadmaps actually ship. If only 70% of these pipelines execute, purpose binding lands at ~64% (36% still missing) and kill switches at ~64% (36% still missing). Even the optimistic projections leave a quarter of organizations without basic containment controls at the end of 2026.

Difference

15-20 **points**

Segment	Missing Purpose Binding	Missing Kill Switch	Missing Isolation
Government	**90%**	**76%**	**81%**
Just starting AI	**81%**	**79%**	**72%**
Healthcare	68%	59%	55%
Global Average	63%	60%	55%
Technology	49%	46%	44%
Australia	48%	43%	39%

Government is the crisis: 90% lack purpose binding, 76% lack kill switches, 81% lack network isolation. These organizations are deploying AI agents they cannot constrain, cannot terminate, and cannot isolate from sensitive systems. Organizations just starting their AI journey are nearly as exposed—79-81% missing containment controls—and they're about to accelerate deployment.

Australia shows what's possible: 48% missing purpose binding (vs. 63% global) and 43% missing kill switch (vs. 60% global), with the strongest pipelines to close the remaining gaps. They're not just ahead—they're pulling further ahead.

The investment intent is clear. Organizations know what's broken. Whether they fix it before an incident forces them to is the open question. The governance-vs.-containment gap will narrow through 2026—but it won't close.

CONFIDENCE LEVEL: ■ ■ ■

MEDIUM

(that containment controls improve; low confidence they catch up with deployment velocity, particularly in government and organizations just starting AI adoption)

PREDICTION #6:

AI Data Security and Privacy Remain the Fastest-Growing Risk Cluster

AI-related risks will dominate security and privacy agendas through 2026—and most organizations aren't equipped to address them. The top concerns point to exposures that existing controls don't cover, including third-party AI vendor handling (30%), training data poisoning (29%), PII leakage via outputs (27%), insider threats amplified by AI (26%). These aren't traditional threat vectors. Most security programs weren't built for them.

	Top Security Risks	Typical Control Maturity
30%	Third-party AI vendor handling	WEAK only 36% have visibility
29%	Training data poisoning	VERY WEAK 22% have pre-training validation
27%	PII leakage via outputs/ embeddings	WEAK 37% have purpose binding
26%	Insider threats amplified by AI	MODERATE 59% have human-in-the-loop
23%	Shadow AI	VERY WEAK few have discovery tools

The #1 security concern—third-party AI vendor handling—is also one of the least controlled. Only 36% have visibility into how partners handle data in AI systems. Organizations are worried about a risk they can't see. Training data poisoning ranks #2, but only 22% have pre-training validation in place. 78% are training or fine-tuning models without validating input data integrity.

Top Privacy Exposures		Current Mitigation
Personal data in prompts	35%	Mostly policy, rarely technical
Cross-border transfers via AI vendors	29%	Contractual only for most
PII leakage in outputs	26%	37% have purpose binding
Lack of consent for AI processing	24%	Often not tracked

The #1 privacy exposure—personal data in prompts—is the simplest failure mode. Employees paste customer information into AI assistants every day. 35% of organizations cite this as a top exposure, but technical controls to prevent it are rare. Most rely on policy and training. Policy doesn't stop someone from pasting a customer list into ChatGPT at 11 p.m.

Industry	Third-Party AI Handling Concern	Visibility Into Partner AI
Manufacturing	**52%**	Low
Healthcare	**50%**	Low
Financial Services	33%	Moderate
Technology	34%	Moderate
Global Average	30%	**36%**

Manufacturing and Healthcare are most exposed—over 50% cite third-party AI handling as a top concern, but these industries also trail on visibility and AI-specific controls. They see the risk clearly. They lack the tools to manage it.

The uncomfortable pattern: organizations can articulate the AI risks they face. They haven't built the controls to address them. The risk cluster is growing faster than the control portfolio.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that AI risks dominate the agenda; high confidence that control gaps will persist through 2026)

PREDICTION #7:

Software Supply Chain Security Expands to Include AI Model Attestations

72% of organizations can't produce a reliable inventory of their software components. The AI supply chain is even worse: There's no standard for AI model attestations, and almost no one is tracking model provenance. Software supply chain security is maturing—but not fast enough, and not broadly enough to include AI.

Supply Chain Control	Not In Place	Exposure
SBOM management	**72%**	Can't identify component vulnerabilities
Continuous dependency monitoring	**71%**	Vulnerabilities go undetected
Zero-trust deployment	**65%**	Compromised code can execute
Vendor security attestations	**63%**	Trusting without verifying
Secure SDLC	**59%**	Vulnerabilities introduced in development
Code vulnerability scanning	**56%**	Known vulnerabilities missed

When the next Log4j happens, 72% of organizations will scramble to determine exposure because they don't have SBOM. 71% won't catch it through continuous monitoring because they don't have any. The basics aren't in place—and AI makes it worse.

The problem extends beyond application code to the infrastructure organizations use to move sensitive data. Legacy file sharing and managed file transfer (MFT) solutions—many built on decades-old protocols—lack the security capabilities modern threats require: granular access controls, real-time DLP, zero-trust architecture, evidence-quality audit trails, and AI-aware policy enforcement. Organizations are running AI workloads and sensitive data exchanges through infrastructure that predates the threat landscape they now face. Modernizing data exchange technology isn't optional—it's a supply chain security requirement.

Region	No SBOM	No Code Scanning	No Secure SDLC
United States	**71%**	58%	62%
Global Average	**72%**	56%	59%
UAE	65%	38%	35%
Australia	52%	39%	43%

The U.S. trails badly: 71% without SBOM, 58% without code scanning, 62% without secure SDLC. Australia and UAE are significantly ahead—but even there, half or more lack SBOM management.

35% cite AI supply chain risks in their top three concerns—compromised models, poisoned training data, missing AI attestations. They're right to be concerned. There's no standard AI SBOM format. No widely adopted attestation framework for AI model supply chains. Organizations know they need this. The tooling and standards don't exist yet, and organizations aren't building workarounds. Meanwhile, they're exchanging AI models, training data, and inference results through legacy transfer infrastructure that can't enforce the policies or provide the visibility AI governance requires.

CONFIDENCE LEVEL: ■ ■ ■

MEDIUM

(that SBOM and AI attestations grow; dependent on regulatory push and standard development)

PREDICTION #8:

Third-Party Risk Management Pivots to Visibility and AI Handling

The annual vendor questionnaire is dying—but 89% of organizations have nothing to replace it with. Third-party risk programs need to pivot from checkbox assessments to continuous, AI-aware monitoring of partner data handling. Most won't make it.

	Top Third-Party Priorities		Current Capability
1	End-to-end visibility gaps	46%	**Only 11% have practiced IR with partners**
2	Partners' AI data handling	36%	**Only 36% have any visibility**
3	Partner compliance gaps	32%	Questionnaire-dependent
4	Inconsistent policy enforcement	31%	Manual for most
5	Unauthorized onward sharing	31%	Rarely tracked

46% cite visibility gaps as their #1 priority—and they're right to worry. Only 36% have any visibility into how partners handle data in AI systems. The rest are trusting contracts and questionnaires to protect them from risks they can't see.

The resilience gaps are severe. 87% lack joint IR playbooks with partners. 89% have never practiced incident response with their third-party vendors. When a partner gets breached—and partners get breached—nearly nine out of 10 organizations will improvise their response. No playbook. No practice. No coordinated plan.

Third-Party Control	Not In Place	Gap Severity
Joint incident response exercises with partners	**89%**	Critical
Joint IR playbooks	**87%**	Critical
Automated kill switch for partner access	**84%**	Severe
Zero-trust access	63%	Significant
External identity/lifecycle management	60%	Significant
Data classification for partner exchanges	57%	Moderate
Secure private data exchange	52%	Moderate

Segment	Visibility Gap Concern	Unauthorized Sharing Concern
Manufacturing	**67%**	38%
Germany	48%	**60%**
Global Average	46%	31%

Manufacturing sees blind spots everywhere—67% cite visibility gaps, 21 points above average. Complex, multi-tier supply chains with almost no insight into how data moves through them. Germany stands out on unauthorized onward sharing at 60%—nearly double the global average. GDPR enforcement has taught German organizations that they're liable for what their partners do with data. Everyone else will learn the same lesson eventually.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that visibility becomes the priority; low confidence that most organizations will achieve it)

PREDICTION #9:

Incident Response Becomes AI-Infused

60% of organizations lack AI-powered anomaly detection. 51% are running manual IR playbooks. 52% haven't tested their RTO/RPO. The foundational capabilities exist—68% have immutable backups, 67% have audit trails—but the AI-specific detection and response capabilities that modern threats require are missing.

IR Capability	Not In Place	Gap Severity
AI anomaly detection	60%	Critical
Automated IR playbooks	51%	Significant
RTO/RPO testing	52%	Significant
Partner notification protocols	48%	Moderate
Evidence-quality audit trails	33%	Moderate
Immutable backups	32%	Moderate

AI anomaly detection is the largest gap with the largest pipeline—43%, the highest for any IR capability. Organizations know they need it. But going from 40% to the projected 83% requires tool procurement, data pipeline construction, model tuning, alert triage processes, and staff training. That's not a flip-the-switch deployment. Assume 60% to 70% pipeline execution, and AI anomaly detection lands at 65% to 70%—leaving 30% to 35% still blind to AI-specific threats at the end of 2026.

Industry	No AI Anomaly Detection	No Automated Playbooks	No RTO/RPO Testing
Government	**76%**	**76%**	**67%**
Healthcare	**64%**	**68%**	**77%**
Professional Services	47%	40%	47%
Technology	31%	44%	38%

Government and Healthcare are in the worst position. 76% of government organizations lack AI anomaly detection; 76% are running manual playbooks. Healthcare handles PHI with 64% missing AI anomaly detection and 77% not testing RTO/RPO—they don't know how long recovery will take until they're in the middle of an incident.

The IR gap connects directly to training-data recovery (see Prediction #11). 53% can't recover AI training data after an incident—meaning even organizations with strong detection and response can't remediate compromised models. IR that stops at "we contained the breach" without addressing "we cleaned the AI" is incomplete.

CONFIDENCE LEVEL: ■ ■ ■

MEDIUM

(that AI-infused IR improves; execution risk is significant, particularly in government and healthcare)

PREDICTION #10:

Evidence-Quality Audit Trails Become the Keystone of AI Governance

33% of organizations lack evidence-quality audit trails. That one gap predicts nearly everything else. Organizations without audit trails show dramatically lower maturity across every AI dimension—not by a few points, but by 20 to 32 points. Audit trails aren't just a compliance artifact. They're the foundation that makes everything else possible.

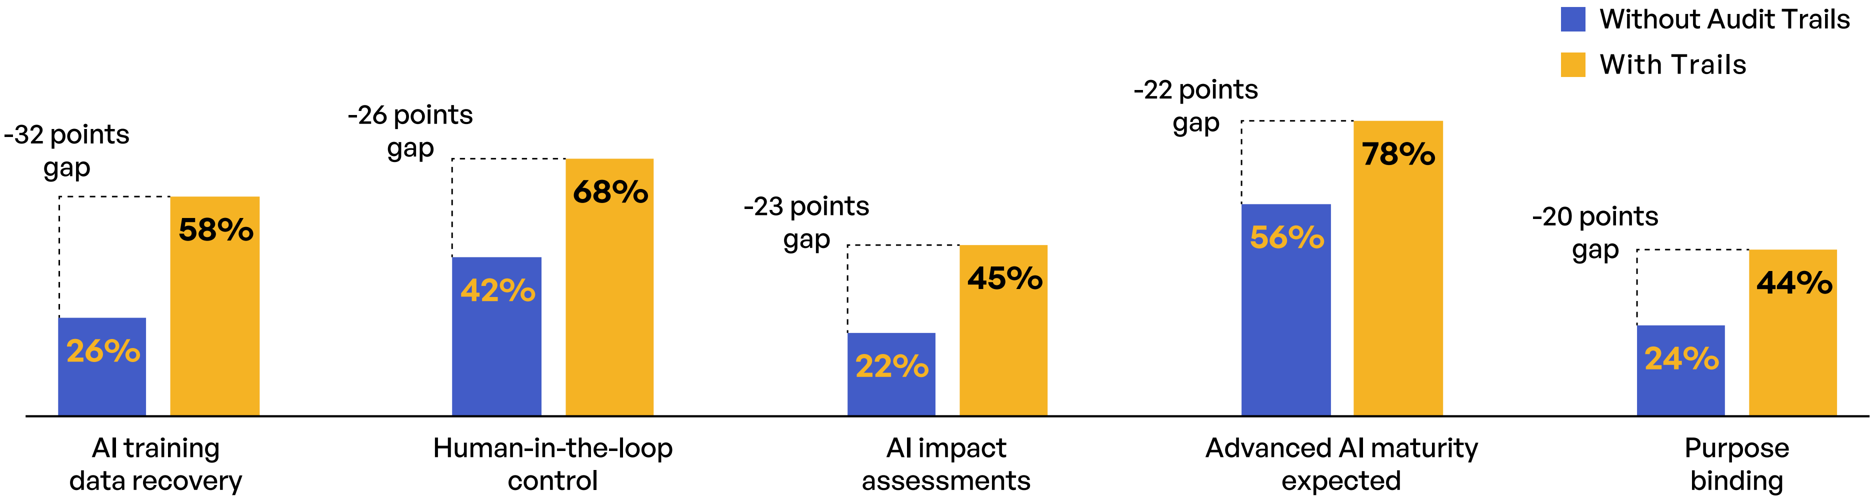

Organizations without audit trails are half as likely to have AI training data recovery (26% vs. 58%). They're 20 points behind on purpose binding, 26 points behind on human-in-the-loop controls. These aren't small differences—they're categorically different maturity tiers.

The problem isn't just missing audit trails—it's fragmented ones. Only 39% of organizations have unified data exchange with enforcement; 61% are running partial, channel-specific, or minimal approaches. That fragmentation shows up in the logs. Disaggregated data exchange—separate systems for email, file sharing, MFT, cloud storage, AI tools—produces logs scattered across platforms, each in its own format with its own retention policy. When an incident occurs or an auditor asks questions, security teams spend hours—sometimes days—manually correlating logs across systems, trying to reconstruct what happened.

Data Exchange Approach	Percentage	Audit Trail Quality
Unified with enforcement	39%	Evidence-quality possible
Partial (gaps remain)	34%	Fragmented logs, manual correlation
Channel-specific only	16%	Siloed logs, major gaps
Minimal/not addressed	11%	Little to no evidence

The logs exist. They just aren't aggregated, normalized, or actionable. 61% of organizations are trying to build evidence-quality audit trails on top of fragmented data exchange infrastructure—a foundation that can't support it. This creates both risk (gaps in visibility, delayed detection, incomplete evidence) and operational inefficiency (manual correlation, inconsistent retention, duplicated effort). Evidence-quality audit trails require a unified view across all channels where sensitive data moves—not a patchwork of system-specific logs that no one has time to stitch together.

The correlation between audit trails and everything else is stronger than industry, region, or organization size. Organizations that take governance seriously start with the ability to prove what happened—and that requires unified data exchange infrastructure, not just logging tools bolted onto fragmented systems. The 33% without evidence-quality trails and the 61% with fragmented data exchange are behind on almost everything else.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that audit trails are recognized as the keystone capability; the correlation is already clear in the data)

PREDICTION #11:

Training-Data Controls and "Unlearning-Ready" Architectures Become Regulatory Requirements

78% of organizations can't validate data before it enters training pipelines. 77% can't trace where their training data came from. 53% can't recover training data after an incident. The "right to be forgotten" is coming for AI. Almost no one is ready.

Not In Place	78%	77%	65%	62%	59%	53%
Training Data Capability	Pre-training validation	Provenance & lineage	Dataset access controls	Encryption for training data	Data minimization for AI	AI training data recovery
Regulatory Exposure	Can't prove data quality to regulators	Can't answer "where did this data come from?"	Can't prove authorized use	Training data exposed in breach	Excess PII in models	Can't "unlearn" or remediate

When a regulator asks, "How do you know there's no PII in your model?"—78% of organizations can't answer. When a data subject exercises deletion rights under GDPR, CCPA/CPRA, or emerging AI regulations—53% have no mechanism to remove their data from trained models. They'll either retrain from scratch (expensive, often impractical) or hope no one asks (increasingly risky).

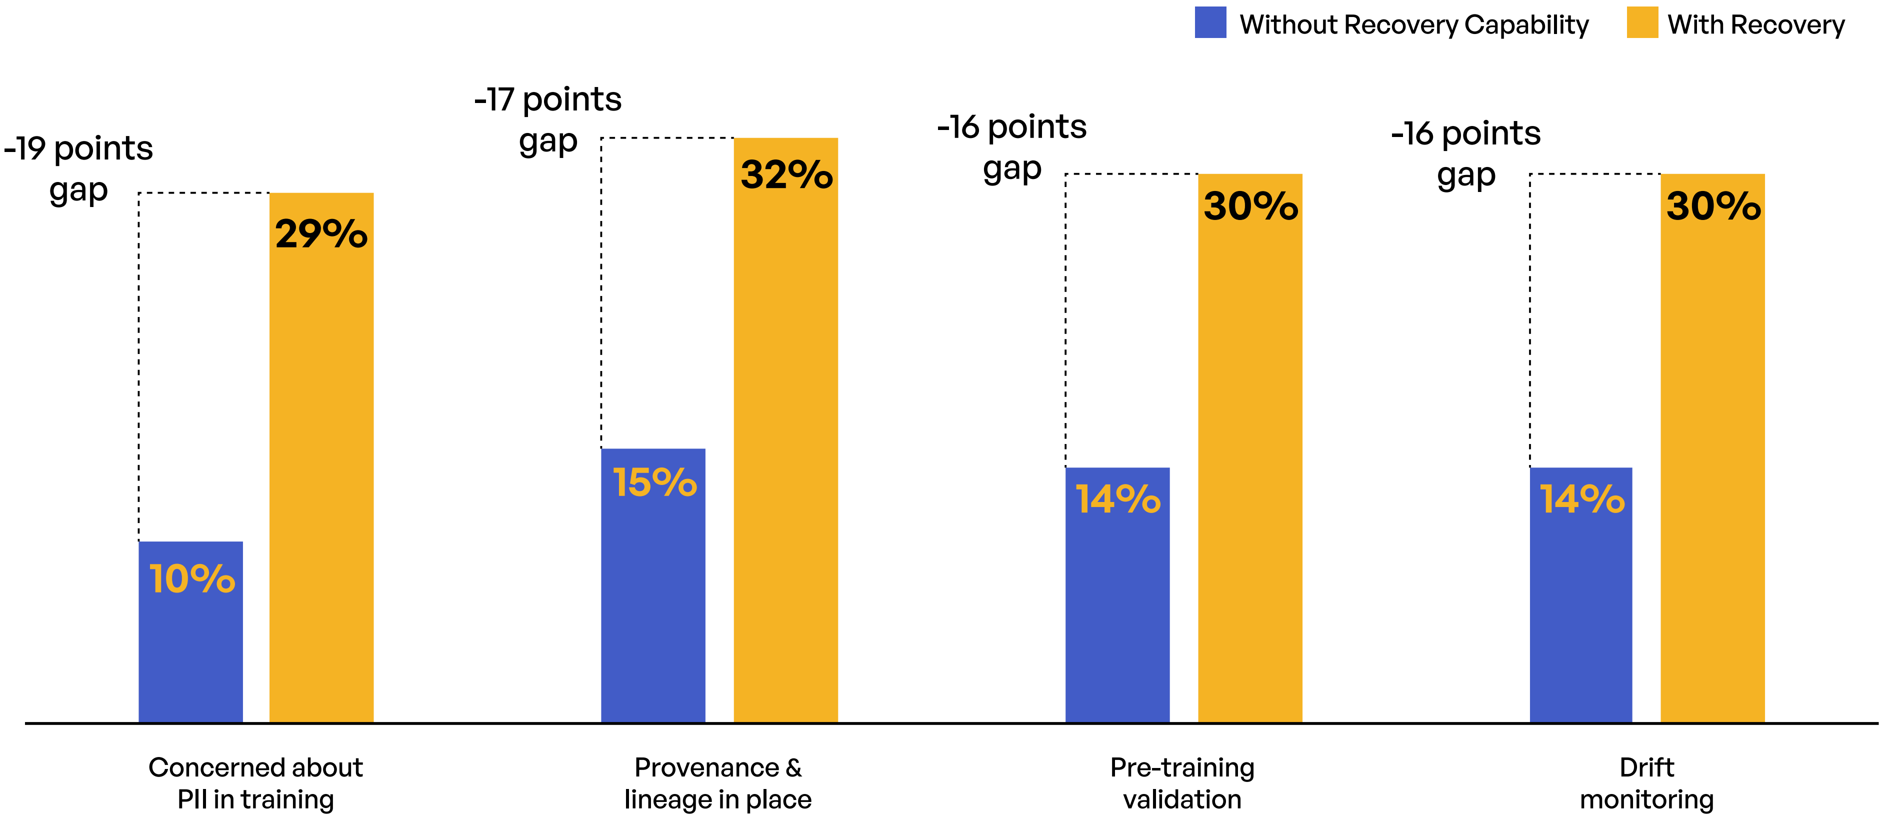

Organizations without training-data recovery are less concerned about PII in training (10% vs. 29%)—not because they have less exposure, but because they're less aware of the risk. The capability gap tracks directly to the awareness gap: Organizations that haven't built recovery don't see the problem they've created.

The regulatory trajectory is unmistakable:

Regulation	Training Data Requirement
GDPR Article 17	Right to erasure extends to derived data
EU AI Act	Training data documentation and governance
CCPA/CPRA	Deletion rights include inferences
Emerging state laws	Following CCPA pattern

This connects directly to IR capabilities (see Prediction #9). Training-data recovery isn't just a compliance capability—it's an incident response capability. When a model is compromised, poisoned, or found to contain unauthorized data, organizations need to remediate. 53% can't. Their incident response stops at containment; they have no path to remediation that doesn't involve starting over.

The organizations that can prove how training data is governed, traced, validated, and "forgotten" will have competitive and compliance advantage. The 77% to 78% that can't trace or validate, and the 53% that can't recover, will face increasingly uncomfortable questions from regulators, auditors, and data subjects.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that training-data controls become differentiators; the regulatory trajectory makes this inevitable)

PREDICTION #12:

AI Governance Becomes a Board-Level Risk Domain Everywhere

54% of boards don't have AI governance in their top five topics. That gap correlates with dramatically lower maturity on every AI metric—26 to 28 points lower on impact assessments, purpose binding, and expected maturity. Where boards aren't paying attention, organizations aren't investing.

Board Topic	Not in Top 5	Consequence
AI governance/ responsible AI	54%	**26-28** point maturity gap
Third-party/ vendor risk	65%	Visibility gaps persist
Regulatory compliance	60%	Reactive posture
Data breach incidents	58%	Incident response gaps
Data privacy & protection	57%	Privacy controls lag

AI governance is already the #2 board topic at 46%—but that means the majority still aren't prioritizing it. The gap matters because board attention is the single strongest predictor of AI maturity in the survey.

Organizations without board engagement are half as likely to conduct AI impact assessments (24% vs. 52%). They're 26 points behind on purpose binding, 24 points behind on human-in-the-loop controls. When boards don't ask about AI governance, organizations don't build it.

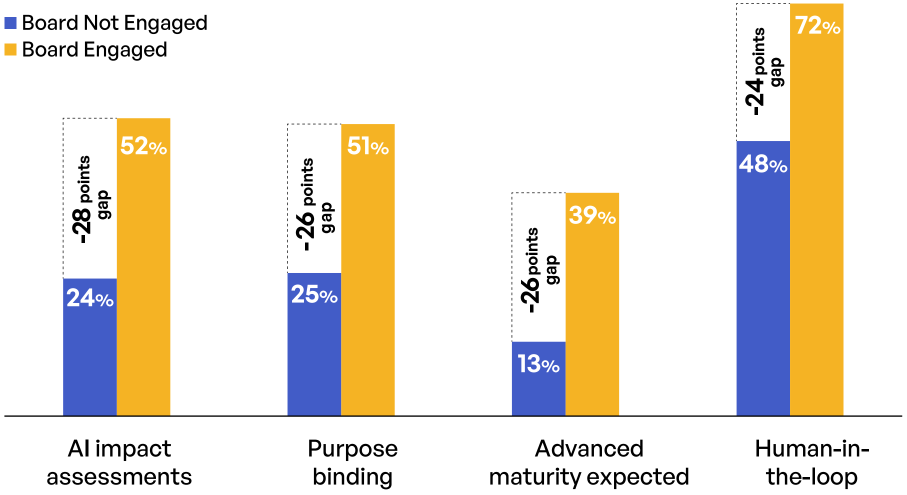

Industry	Board Not Engaged	Gap to Leaders
Government	**71%**	-51 points vs. Pro Services
Healthcare	55%	-35 points
Technology	47%	-27 points
Financial Services	40%	-20 points
Professional Services	20%	Benchmark

Government is the outlier: 71% of boards aren't engaged on AI governance. Professional Services leads at 80% engagement—a 51-point gap. Government handles citizen data, classified information, critical infrastructure—with the least board oversight on AI risk.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that board engagement becomes universal; lower confidence on government closing the gap)

PREDICTION #13:

EU AI Act Compliance Creates a Global Governance Template

Organizations not impacted by the EU AI Act are 22-33 points behind on every major AI control. 74% lack AI impact assessments. 72% lack purpose binding. 84% haven't conducted AI red-teaming. The EU AI Act isn't just a European regulation—it's becoming the definition of what "good AI governance" looks like. Organizations outside its scope are falling behind.

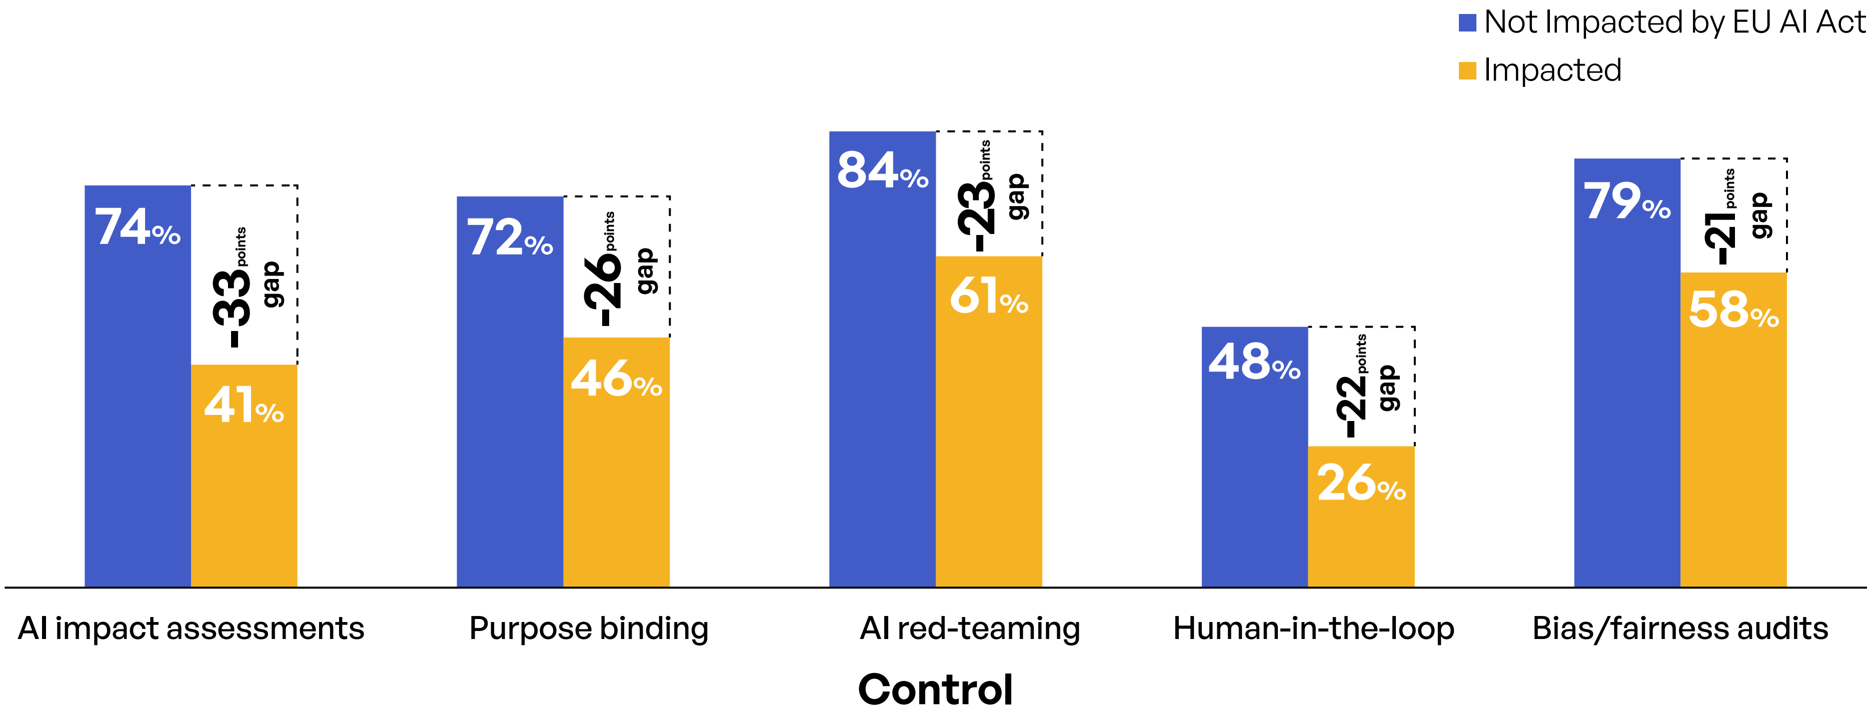

The gaps are categorical. Organizations under EU AI Act pressure are building governance infrastructure. Organizations outside that pressure largely aren't. The Act is creating a two-tier market.

Region	Not Impacted by EU AI Act	Governance Gap Exposure
United States	**82%**	High
Saudi Arabia	**86%**	High
Australia	74%	Moderate-High
United Kingdom	56%	Moderate
Germany	45%	Lower
France	40%	Lower

82% of U.S. organizations aren't feeling EU AI Act pressure—yet. But the regulation spreads through supply chain requirements, multinational operations, and competitive benchmarking. Organizations that dismiss it as "a European problem" will find themselves 22-33 points behind on AI governance as the framework becomes the global baseline.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that EU AI Act becomes global template; the maturity gaps are already visible)

PREDICTION #14:

Post-Quantum Cryptography Moves From Early Adopter to Mainstream

84% of organizations haven't implemented post-quantum cryptography (PQC). 48% aren't using it at all, and we believe this number is even worse (overconfidence on the part of survey respondents). The "harvest now, decrypt later" threat is already active adversaries can capture encrypted data today and wait for quantum computers to break it. For data that needs to stay confidential for decades, the window to act is closing.

PQC Status			
Not using PQC	**Piloting/evaluating**	**Implemented**	**Not applicable/don't know**
48%	**~18%**	**16%**	**18%**
Fully exposed to harvest attacks	Aware but not protected	Protected	Unaware of exposure

Only 16% have implemented PQC. The other 84% are either piloting, ignoring, or unaware. For organizations handling medical records, financial data, classified information, or anything else that needs confidentiality beyond 2030—48% are fully exposed, and another 18% don't know enough to assess their risk.

Driver	Status	Implication
NIST PQC standards	Finalized 2024	No more "waiting for standards" excuse
OMB M-23-02	Active	Federal crypto inventory required
NSA CNSA 2.0	Phased through 2030s	National security systems must migrate
"Harvest now, decrypt later"	Active threat	Long-lived data already at risk
Platform PQC support	Expanding 2025-2026	Implementation getting easier

No strong regional or industry leaders exist yet—everyone is early. That's concerning given that government, defense, and financial services should be leading and aren't. The migration timeline extends to 2028-2030, but organizations that haven't started planning are already behind.

CONFIDENCE LEVEL: 🟨 🟨 🟨

MEDIUM

(that PQC awareness grows; implementation will lag awareness significantly)

PREDICTION #15:

Data Sovereignty Becomes an AI Governance Imperative

29% cite cross-border transfers via AI vendors as a top privacy exposure. 34% cite cross-border data transfer mechanisms as a top regulatory priority. But most organizations have solved sovereignty for storage—not for AI. They know where their data resides. They don't know where it's processed, trained, or inferred.

Data Sovereignty Exposure	Percentage Citing	Current Control Maturity
Cross-border transfers via AI vendors	29%	Contractual only for most
Third-party AI vendor data handling	30%	**Only 36% have visibility**
Lack of consent for AI processing	24%	Often not tracked
Partner AI/LLM tools exposing exchanged data	29%	Rarely governed

Traditional sovereignty controls address data at rest: which data center, which country, which legal jurisdiction. AI breaks that model. A prompt sent to a cloud AI vendor may be processed in a different jurisdiction, used to fine-tune models hosted elsewhere, or generate outputs that traverse multiple borders before returning. 29% recognize cross-border AI transfers as an exposure—but recognition isn't control.

Data Sovereignty Gap	Consequence
Know where data is stored, not where it's processed	Processing jurisdiction determines legal exposure
No visibility into AI vendor data handling	Can't verify contractual sovereignty claims
No control over training data location	Models may train on data across jurisdictions
Inference location unknown	Real-time processing may violate residency requirements

Region	Third-Party AI Handling Concern	Unauthorized Sharing Concern
Middle East (UAE/Saudi)	42-45%	35-40%
Germany	38%	60%
Manufacturing	52%	38%
Global Average	30%	31%

The Middle East leads on sovereignty concerns—42% to 45% cite third-party AI vendor handling as a top risk, driven by explicit data localization requirements. Germany stands out at 60% concerned about unauthorized onward sharing—nearly double the global average—because GDPR enforcement has made data flow liability concrete. These regions see the problem clearly. Most others are still catching up.

The regulatory trajectory is tightening. EU data boundary requirements, Middle East localization mandates, China's data export restrictions, and emerging U.S. state laws all assume organizations can demonstrate where data is processed—not just stored. AI complicates every one of these requirements because processing is distributed, dynamic, and often opaque.

Regulatory Pressure	Sovereignty Requirement
GDPR / EU Data Boundary	Processing location matters, not just storage
Middle East localization	In-country processing for sensitive categories
China PIPL	Cross-border transfer restrictions for AI
U.S. state laws (emerging)	Increasingly following GDPR patterns
EU AI Act	Transparency on where AI systems operate

The gap: Organizations have invested in sovereign storage infrastructure. They haven't extended sovereignty controls to AI processing. 30% cite third-party AI vendor handling as a top security concern, but only 36% have any visibility into how partners handle data in AI systems. The rest are relying on contracts and hoping vendors comply.

For organizations exchanging sensitive data with partners, customers, or AI vendors across borders, the sovereignty question is shifting from "where is the data stored?" to "where is it processed, who can access it, and can you prove it?" Most can't answer the second set of questions. As AI becomes embedded in data exchange workflows, the organizations that can demonstrate processing sovereignty—not just storage sovereignty—will have regulatory and competitive advantage. The majority, still governing storage while ignoring processing, will face increasingly difficult compliance conversations.

CONFIDENCE LEVEL: ■ ■ ■ ■ ■

HIGH

(that sovereignty requirements expand to AI processing; the regulatory trajectory is already clear in EU, Middle East, and emerging elsewhere)

Thematic Deep Dive

The Agentic AI Security Imperative

The Adoption Curve

Agentic AI has crossed the threshold from pilot project to production system. The problem: controls haven't kept pace. Organizations deploy advanced use cases with controls designed for basic ones.

Use Case Maturity	Established	Scaling	Emerging	Early
Examples	Internal copilots	Document gen, email composition	API agents, workflow automation	Decision-making agents
Adoption	39%	33-34%	33%	24%
Control Maturity	Moderate gaps	Significant gaps	Severe gaps	**Almost no controls**

Technology and Professional Services are furthest along on advanced use cases—API agents (47%), MFT automation (53%), code generation (50%)—but even they have containment gaps. Healthcare stays conservative at 59% copilots, which provides temporary buffer but no governance experience for when adoption accelerates. Government barely registers on advanced automation—only 5% on API agents compared to 36% elsewhere—creating a modernization gap that will compound when they inevitably adopt.

The MFT channel highlights the disconnect: 27% are planning AI-driven MFT automation, but MFT security adoption is only 46%. Organizations are adding autonomous agents to channels they haven't secured.

Data Governance vs. Containment Gap

This is the central tension of agentic AI security—and it won't resolve by 2026.

Control Category

	Examples	Not In Place	Gap Severity
Governance	Human-in-the-loop, monitoring, minimization	41-44%	Moderate
Containment	Kill switch, purpose binding, isolation	55-63%	Severe
Difference			15-20points

Organizations have invested in watching. They haven't invested in stopping:

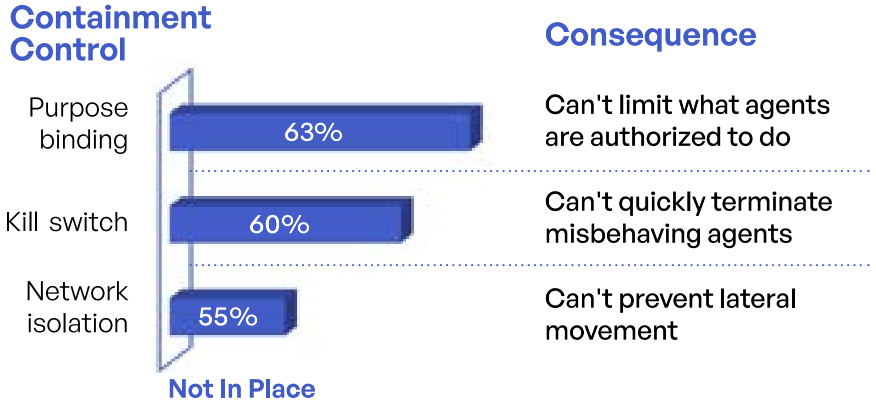

Why the gap persists: Governance is easier to deploy—logging doesn't require architecture changes. Governance satisfies auditors—"we're monitoring" sounds like control. Containment reveals capability gaps organizations would rather not discover.

The pipelines are aimed at the right targets—purpose binding has the highest pipeline in the survey (39%), kill switch has 34%. But even if 80% execute, 24% to 26% of organizations will still lack basic containment at the end of 2026. If only 60% to 70% execute—more realistic historically—36%+ will still be missing them.

AI Intensity Creates Two Worlds

Organizations with extensive AI use look nothing like organizations just starting.

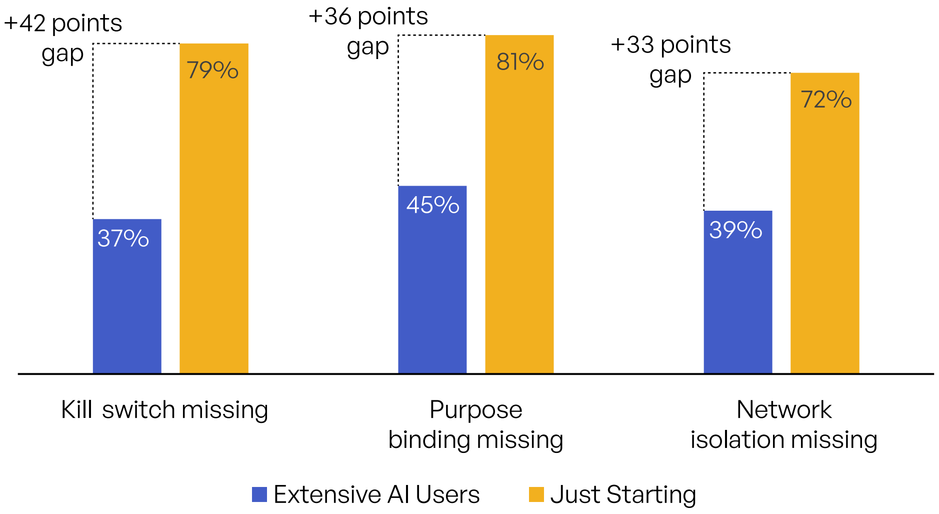

Containment Control

The organizations deploying the most AI are governing it best. The organizations just starting have almost nothing—and they're about to accelerate deployment.

This creates bifurcation: leaders pull further ahead while laggards fall further behind. The next wave of AI incidents will likely come from organizations rushing to deploy without the governance infrastructure that experienced organizations have built through trial and error.

The Data Sovereignty Dimension

Data sovereignty adds another layer of exposure (see Prediction #15). Organizations have solved sovereignty for storage—not for AI processing.

29%

Cite cross-border AI transfers as exposure

Only **36%**

Have visibility into partner AI data handling

Minority

Know where data is processed, not just stored

Data Sovereignty Gap

A prompt sent to a cloud AI vendor may be processed in a different jurisdiction, used to fine-tune models hosted elsewhere, or generate outputs that traverse multiple borders. Traditional data residency controls don't address this. Organizations governing AI storage while ignoring AI processing will face compliance problems as sovereignty requirements expand.

Notable Outliers

Australia is the benchmark:

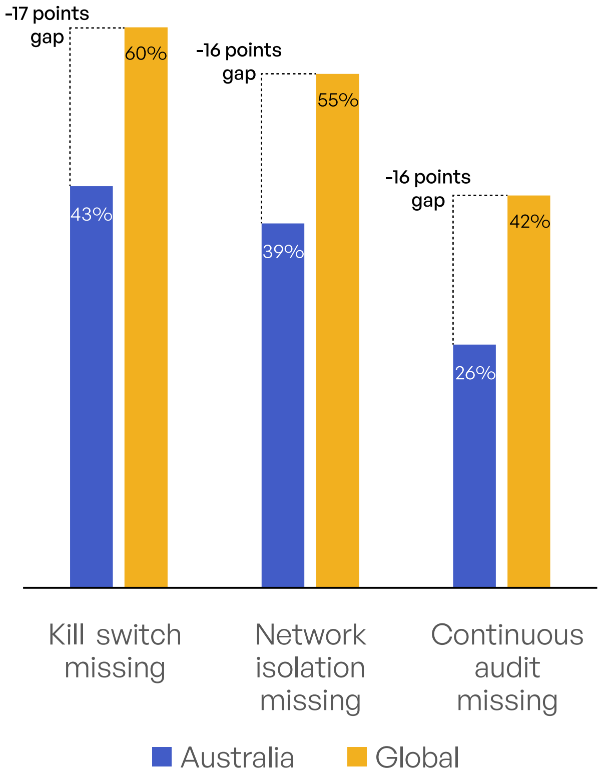

Plus, the strongest pipelines. They have both higher AI adoption and higher controls—compounding advantage, not trading off.

Government is the crisis:

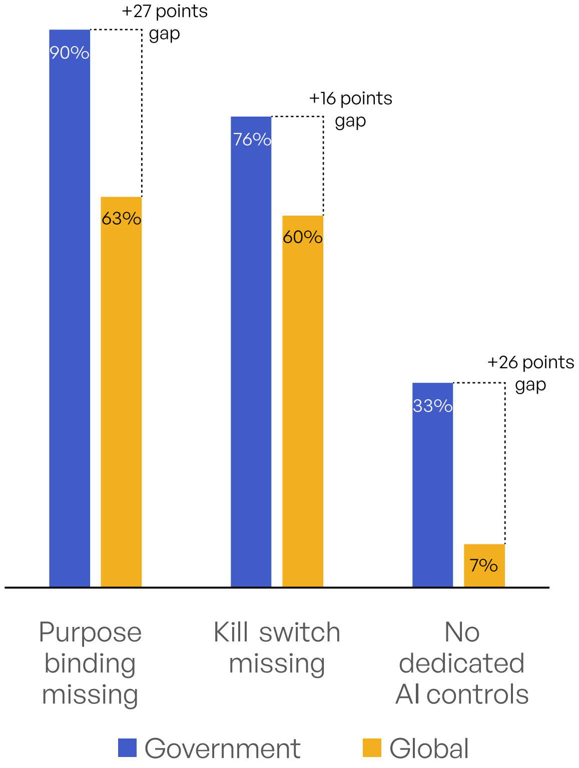

These organizations handle citizen data, classified information, critical infrastructure—with AI controls a generation behind everyone else.

Professional Services is the pressure cooker:

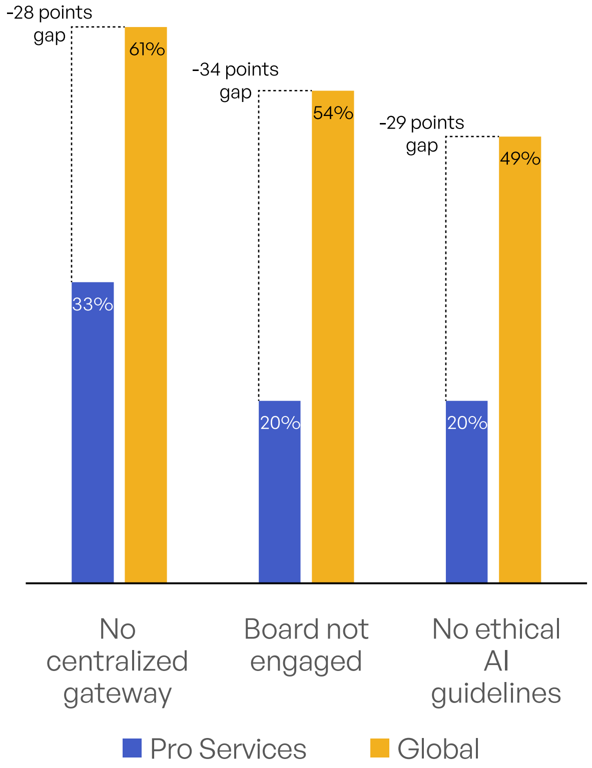

Client data exposure drives aggressive governance. The fear is appropriate; the response is rational. If you want to see what AI governance under pressure looks like, study Professional Services.

Bottom Line

100% of organizations have AI on the roadmap. 63% can't enforce purpose limitations. 60% can't terminate misbehaving agents. 55% can't isolate AI from sensitive systems. Organizations just starting are 33 to 42 points behind on containment—and accelerating deployment anyway.

The governance-vs.-containment gap will narrow through 2026. It won't close. The organizations that close it first will be demonstrably more resilient. The organizations that don't will learn the same lessons the hard way—likely through incident.

Strategic Recommendations

Priority Actions by Timeline

Immediate (Q1-Q2 2026)

Action	Why Now
Close the kill-switch gap	60% can't terminate AI agents quickly; incident will expose this
Implement purpose binding	63% have no limits on agent authorization; largest gap in survey
Audit your audit trails	33% lack them; 61% have fragmented logs that aren't actionable
Inventory agentic AI use cases	Can't govern what you don't know about; shadow AI proliferating
Assess third-party AI exposure	36% have visibility; the rest are trusting contracts blindly
Map AI data sovereignty exposure	29% cite cross-border AI as risk; most don't know where data is processed

Medium-Term (H2 2026)

Action	Why Now
Deploy AI anomaly detection	60% gap; largest IR capability missing
Build training-data governance framework	EU AI Act requires it; deletion requests are coming; 78% can't validate
Require third-party AI attestations	Include in 2026 contract renewals; questionnaires won't cut it
Establish joint IR playbooks with critical vendors	87% lack this; improvisation isn't response
Practice IR with partners	89% have never run joint tabletops; first time shouldn't be live incident
Consolidate data exchange infrastructure	61% fragmented; can't build evidence-quality trails on scattered systems
Implement centralized AI data gateway	61% are fragmented or have nothing; control plane for all AI governance
Modernize legacy file transfer infrastructure	Legacy MFT lacks AI-aware controls, real-time DLP, evidence-quality logging
Build unlearning-ready architecture	53% can't recover training data; regulators will ask
Extend sovereignty controls to AI processing	Storage sovereignty isn't enough; processing location matters

Long-Term (2027+)

Action	Timeline
Complete PQC migration	84% haven't implemented; "harvest now, decrypt later" already active
Cryptographic inventory and prioritization	Identify long-lived data requiring PQC protection first

Key Actions by Role

Role	Top 3 Actions
CISO/CIO	Get AI governance on board agenda (54% not engaged); demand containment controls not just monitoring; fund keystone capabilities (audit trails + training-data recovery)
IT/Infrastructure	Map AI data flows including cross-border processing; consolidate fragmented data exchange (61% scattered); close MFT security gap (46% adoption, 27% planning AI automation)
DevSecOps	Expand SBOM to AI models (72% lack SBOM entirely); integrate AI security into CI/CD; establish training-data validation (78% can't validate)
Line of Business	Know where AI touches your data and where it's processed; demand vendor AI visibility (only 36% have it); participate in use case governance
Board	Make AI governance standing agenda item (46% have it, 54% don't); ask specifically about containment controls; benchmark against industry and region, not size

Industry Callouts

Government is a generation behind—not incrementally behind. 90% lack purpose binding. 76% lack kill switch. 33% have no dedicated AI controls at all. 71% of boards aren't engaged. This requires transformation: adopt EU AI Act framework as baseline even if not legally required; treat this as a multi-year modernization program, not a checklist.

Healthcare shows severe IR gaps despite PHI sensitivity. 77% not testing RTO/RPO. 64% lack AI anomaly detection. 68% running manual playbooks. Prioritize ruthlessly: audit trails first (keystone capability), then detection and response. You can't afford to discover recovery time during an incident.

Manufacturing sees blind spots everywhere—67% cite visibility gaps, 21 points above average. Complex, multi-tier supply chains with almost no insight into how data moves through them. Third-party visibility isn't optional; it's existential.

Technology is leading but moving fast. 31% lack AI anomaly detection (vs. 60% global), but advanced use case adoption is aggressive. Maintain control deployment in lockstep with AI deployment. Your advanced use cases require advanced governance—don't assume current controls scale.

Professional Services has the highest governance posture—80% board attention, 67% centralized gateway, 80% ethical AI guidelines. Client data exposure drives this. Every control should be evaluated through the lens of "what happens if client data leaks?" The fear is appropriate.

Financial Services is heavily regulated and heavily targeted—but AI governance is still fragmented: 60% lack a centralized AI data gateway and 5% have no dedicated AI controls. Even with a relatively small governance-to-automation gap, 15% still rely on manual/periodic compliance, which won’t hold up as evidence expectations shift to continuous.

Path Forward

1. **The governance-vs.-containment gap is the central challenge.** 63% can't enforce purpose limitations. 60% can't terminate misbehaving agents. Organizations can watch but can't stop. Close this gap first.

2. **Audit trails and training-data recovery are keystone capabilities.** They predict everything else—+20-to-32-point advantages across all AI metrics. But 61% have fragmented logs that aren't actionable. Unified data exchange infrastructure comes before evidence-quality trails.

3. **Board attention is the strongest predictor of maturity.** 54% of boards aren't engaged; those organizations are 26 to 28 points behind on everything. If AI governance isn't on your board's agenda, put it there.

4. **Sovereignty has expanded from storage to processing.** Knowing where data resides isn't enough. 29% cite cross-border AI transfers as exposure, but only 36% have visibility into where data is processed, trained, or inferred.

5. **Legacy infrastructure can't support AI governance.** Disaggregated file sharing and decades-old MFT solutions lack the security capabilities modern AI governance requires. You can't build containment controls, evidence-quality audit trails, or sovereignty assurance on fragmented infrastructure.

The predictions say where the market is headed. The gaps say where you're exposed. What happens to your organization depends on what you do next.

Legal Disclaimer

The information provided in this report is for general informational purposes only and should not be construed as professional advice. Kiteworks and Centiment make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained in this report. Any reliance you place on such information is strictly at your own risk. None of the sponsoring or contributing organizations shall be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from loss of data or profits arising out of, or in connection with, the use of this report. Readers should consult with qualified legal counsel and cybersecurity professionals when addressing specific compliance requirements. The data in this report was analyzed using AI and the content was generated with AI assistance. While AI enhances analytical capabilities, it can occasionally produce errors or biased information that should be considered when reviewing these findings.

About Centiment

Centiment is a market research firm specializing in data collection and analysis for the cybersecurity and technology sectors. The company delivers actionable insights through customized survey design, targeted respondent recruitment, and sophisticated analytics. Centiment’s proprietary research platform ensures exceptional data quality through AI-driven verification and expert human oversight. The company serves Fortune 500 enterprises, technology vendors, and government agencies, providing intelligence for strategic decisions in evolving markets. Headquartered in Denver, Centiment conducts research globally to help organizations understand complex technology landscapes and cybersecurity trends.

About the Research

225 security, IT, compliance, and risk leaders across 10 industries and 8 regions. 97% represent organizations with 1,000+ employees. Survey fielded Q4 2025.

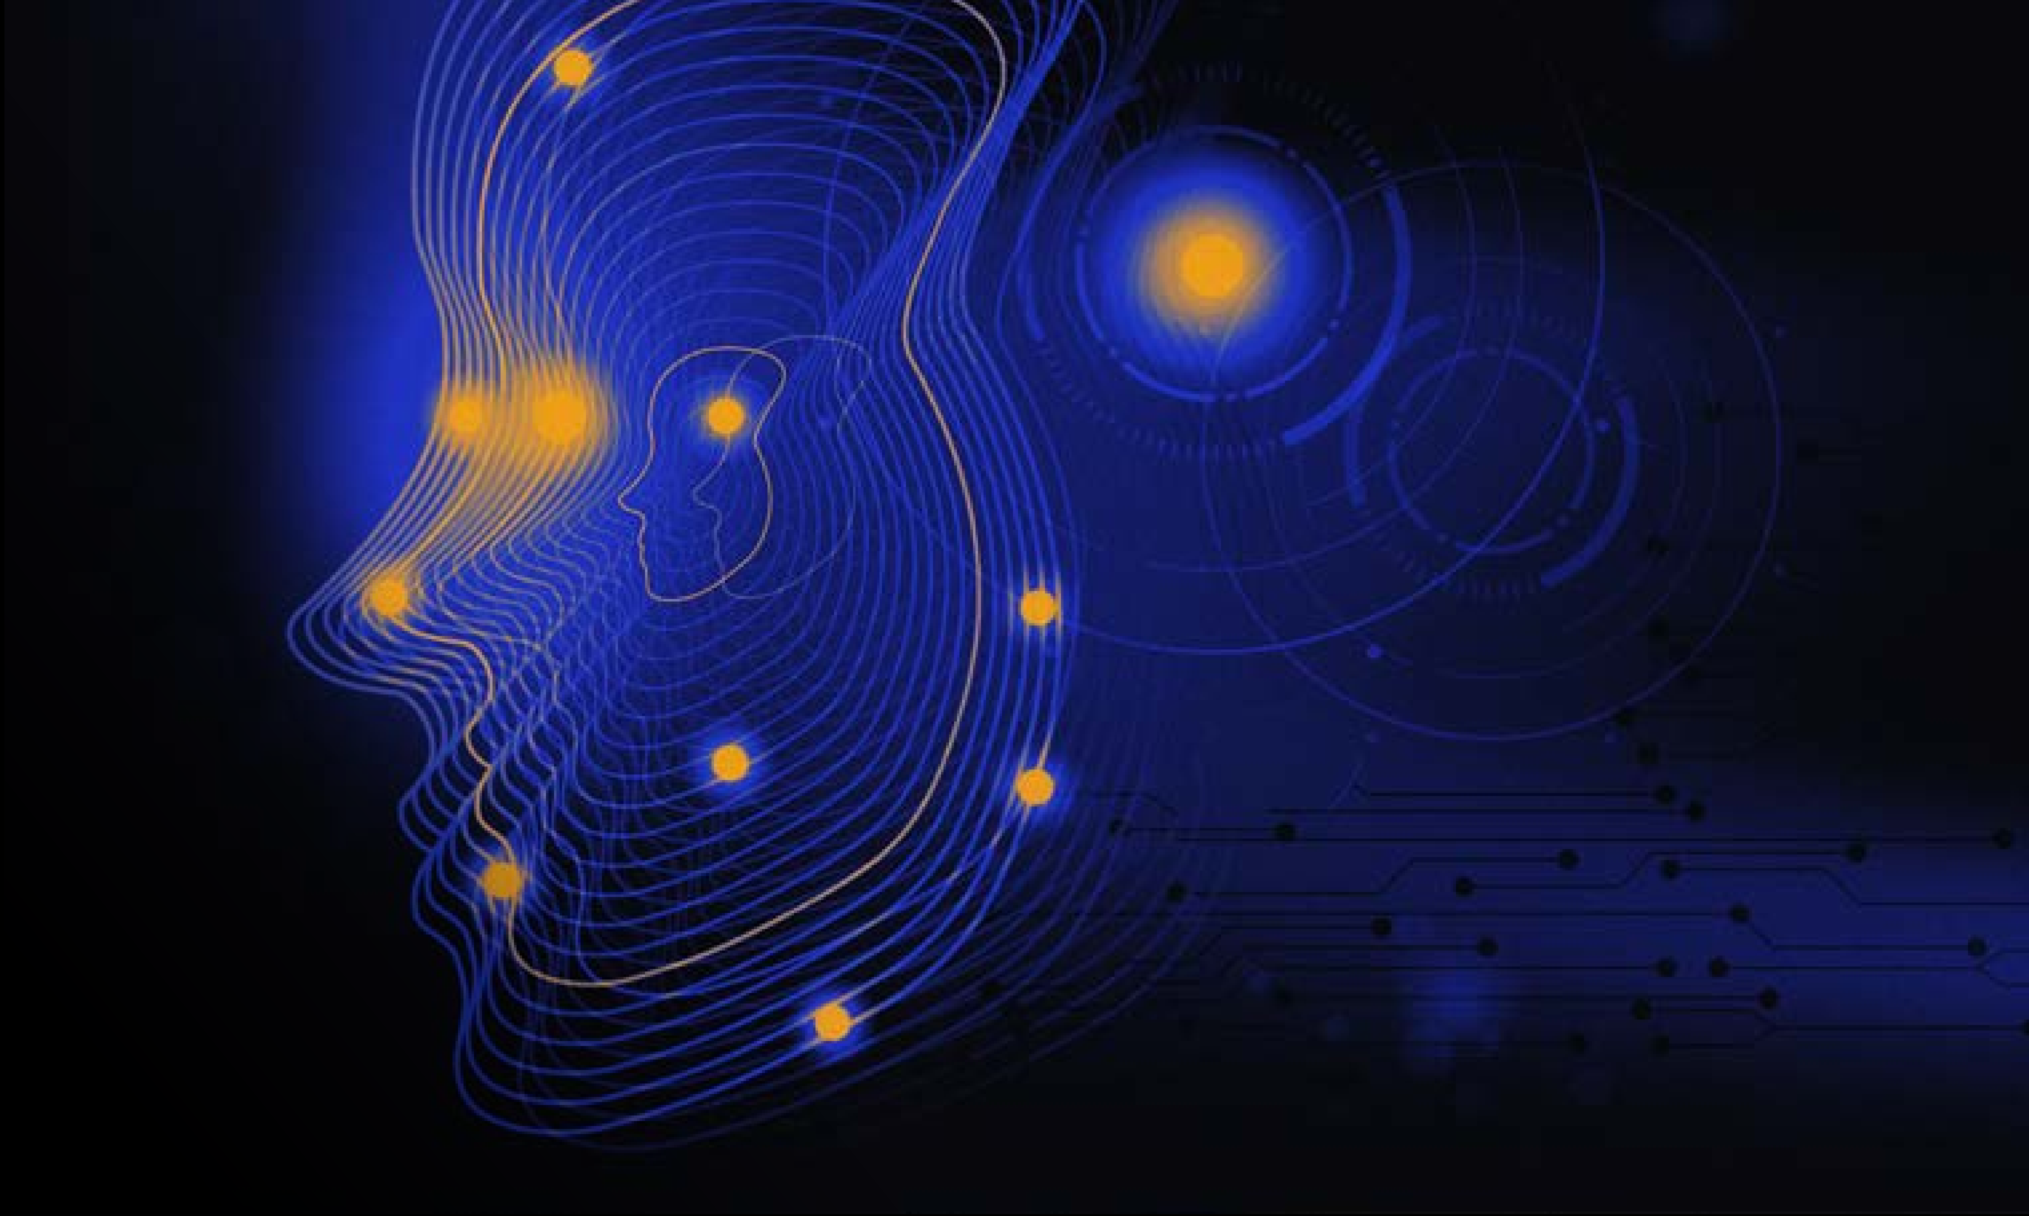

Kiteworks

Kiteworks' mission is to empower organizations to effectively manage risk in every send, share, receive, and use of private data. The Kiteworks platform provides customers with a Private Data Network that delivers data governance, compliance, and protection. The platform unifies, tracks, controls, and secures sensitive data moving within, into, and out of their organization, significantly improving risk management and ensuring regulatory compliance on all private data exchanges. Headquartered in Silicon Valley, Kiteworks protects over 100 million end-users and over 1,500 global enterprises and government agencies.

December 2025

www.kiteworks.com