Kiteworks

EXECUTIVE SUMMARY

2026 Data Security and Compliance Risk Report

THE AI DATA GOVERNANCE GAP: PRIVATE DATA, STALLED COMPLIANCE, AND THE SECURITY RECKONING OF 2026

Overview

Over 450 security and compliance professionals told Kiteworks and Centiment what their organizations are actually doing with AI and sensitive data. The answers are not encouraging. The 2026 Data Security and Compliance Risk survey measures how organizations govern sensitive data in an environment where AI deployment, multi-channel fragmentation, and accelerating regulatory requirements are converging simultaneously.

Three composite scores anchor the analysis, all computed directly from survey responses rather than self-reported estimates:

Data Security Maturity Score (DSMS)

an index from 0 to 100 built from 11 binary security controls (encryption, MFT, SIEM integration, kill switches, and others).

Survey mean: **39/100**

AI Governance Maturity Score (AIGMS)

an index from 0 to 100 built from 19 binary AI data governance capabilities (purpose binding, behavioral monitoring, AI-specific DLP, audit trail generation, and others).

Survey mean: **35/100**

approximately 7 of 19 capabilities deployed

Data Security and Compliance Readiness Index (DSCRI)

DSMS × (AIGMS/100), a multiplicative combination that penalizes single-dimension excellence.

Survey mean: **16.2/100**

The average organization in this survey sits in Tier 2 — Developing — on both the DSMS and AIGMS frameworks.

This report is the accountability document for the 2026 Data Security and Compliance Risk: Forecast Report, published by Kiteworks in December 2025 based on a survey of 225 security and risk leaders. That report made 15 predictions about the state of AI data governance, containment controls, compliance readiness, and third-party risk entering 2026. This survey — 459 respondents, collected in Q2 2026 — tests those predictions against reported organizational outcomes. The verdict, repeated throughout this report: **The Forecast was correct about the direction. It was optimistic about the scale.**

Not a Forecast, a Reported Outcome

74% of organizations experienced at least one general security incident in the past 12 months. Among the 64% that have deployed AI, 64% experienced an AI-specific incident. Across all respondents, 80% experienced at least one incident of either type. The World Economic Forum's Global Cybersecurity Outlook 2026, published in January 2026, finds that 87% of cyber leaders identified AI-related vulnerabilities as the fastest-growing cyber risk of 2025, and CEOs rank data leaks as their highest concern associated with generative AI. 63% experienced a compliance consequence: an audit finding, a required remediation plan, a board escalation, a contractual penalty, or a formal regulatory investigation. 65% discovered employees using unapproved AI tools with organizational data. All of this happened in the 12 months before survey completion. None of it is projected.

The Governance Deficit in One View

Each row below is a governance control that organizations deploying the row above it should already have in place. The cascade is not a progression — most organizations deploying AI have not deployed the controls in parallel. Governance is not catching up to deployment.

Stage	% of Organizations
AI deployed in production environments	64%
AI kill switch formally deployed	30%
AI-specific DLP technically enforced	27%
Purpose binding applied to AI agents	26%
AI access logs forwarded to SIEM	33%
Can produce complete AI audit record within one business day	27%
Can produce complete AI audit record within one hour	17%

No AI containment control measured in this survey is deployed by more than 35% of organizations. 50% cannot produce a complete AI data access audit record within one business day. 73% lack purpose binding controls on sensitive data — meaning AI agents are not restricted to authorized tasks and data scopes, and policy exists without technical enforcement. The gap between what organizations say they govern and what the technical infrastructure actually enforces is the central story this data tells.

Investment intent is aligned with the problem: 43% name AI security and governance as their top investment priority for the next 12 months. But intent and deployment are different things. The 7% of organizations that have reached Tier 4 — with 9 or more of 11 measured security controls deployed — did not get there through stated priorities. They got there through architectural discipline: governance enforcement that is technical rather than behavioral, automated rather than manual, measurable rather than attested.

The Maturity Gap in Numbers

Data Security Maturity Score (DSMS): 39/100 mean index · 70% score in Tier 1 or Tier 2 (index below 45) · 7% have reached Tier 4 — Advanced (index 82–100)

AI Governance Maturity Score (AIGMS): 35/100 mean index · 71% score in Tier 1 or Tier 2 (index below 45) · 9% have reached Tier 4 — Advanced (index 73–100)

Overarching outcomes (past 12 months): 80% experienced at least one security incident · 63% experienced a compliance consequence · 65% discovered shadow AI usage

Data Security and Compliance Readiness Index (DSCRI): 16.2/100 (DSMS × AI governance factor, survey mean)

Forecast Check — 15 Predictions, One Verdict

The 2026 Forecast predicted 100% of organizations would have agentic AI on their roadmap with fewer than 40% having containment controls to manage it. The 2026 Annual Survey found 80% experienced at least one security incident of any type. The Forecast Report five months ago modeled the exposure. This survey measured the consequence.

About the Survey

This report is based on primary research with 459 security, compliance, and technology professionals — cybersecurity, IT, risk, and compliance leaders at organizations with 1,000 or more employees (92% of the sample), spanning 10 industries and 3 primary geographic regions.

Region distribution:
EU/UK 50% (n=231), North America 32% (n=147), MEA 12% (n=54), other 6% (n=27)

Industry distribution:
Technology 36%, Financial Services 16%, Manufacturing 16%, Healthcare 8%, Professional Services 6%, Legal 5%, Education 4%, Energy and Utilities 4%, Federal Government 4%, Defense Contractors 1%

Organization size:
46% have 1,000–4,999 employees; 24% have 5,000–9,999; 14% have more than 25,000

The survey covered seven primary topic domains: AI data governance and containment controls; sensitive data exchange channels; detection and response capabilities; compliance automation and evidence readiness; regulatory framework requirements; human and behavioral dimensions of AI data risk; and planned investment priorities. Responses were collected in Q2 2026 and are self-reported by survey participants.

The Security Maturity Framework

The DSMS is computed from 11 binary security control variables spanning general data security infrastructure (automated classification, DLP enforcement, approved-channel enforcement, SIEM integration, board-level reporting, third-party vendor controls, continuous compliance evidence) and AI-specific data security controls (AI-specific DLP, AI access logs to SIEM, AI kill switch). The average organization has deployed 4.3 of the 11 controls.

DSMS Tier	Index Range	Controls Deployed	% of Orgs
Tier 1 — Nascent	0–18	0–2 of 11	24% (n=111)
Tier 2 — Developing	27–45	3–5 of 11	46% (n=212)
Tier 3 — Established	55–73	6–8 of 11	23% (n=104)
Tier 4 — Advanced	82–100	9–11 of 11	7% (n=32)

The AIGMS is computed from 19 binary AI data governance capabilities across four dimensions: Technical Data Containment (5 capabilities — kill switch, SIEM integration for AI logs, zero-trust access controls, network isolation, single control point for access revocation), Data Audit and Traceability (6 capabilities — complete audit log coverage, output-to-source traceability, one-business-day audit records, formal detection process, rollback capability, board-level reporting), Data Policy Enforcement (4 capabilities — AI-specific DLP, purpose binding, human-in-the-loop review, data minimization), and Organizational Data Governance (4 capabilities — dedicated ownership, documented responsibilities, contractually enforced third-party controls, externally validated compliance).

AIGMS Tier	Index Range	Capabilities Deployed	Governance Gaps Remaining	% of Orgs
Tier 1 — Nascent	0–26	0–5 of 19	14–19	48% (n=220)
Tier 2 — Developing	27–44	6–8 of 19	11–13	23% (n=106)
Tier 3 — Established	45–72	9–13 of 19	6–10	20% (n=93)
Tier 4 — Advanced	73–100	14–19 of 19	0–5	9% (n=40)

How DSMS and AIGMS interact: the two indices define four distinct organizational risk profiles: **False Confidence** (high DSMS, low AIGMS — strong general security, no AI governance; the modal Technology-sector profile; feels protected, is not); **AI-Ready** (high DSMS, high AIGMS — the destination, 11% of organizations); **Dual Exposure** (low DSMS, low AIGMS — Federal Government, Defense Contractors, MEA; incidents are a matter of schedule, not chance); and **Foundation First** (low DSMS, high AIGMS — rare, typically AI-native organizations that skipped foundational infrastructure).

The survey median DSCRI is 11.5 — half of all organizations score below 12. By sector, the readiness gap is narrowest in Defense Contractors (10 points), where both DSMS and AIGMS are simultaneously low, and widest in Financial Services (25 points), where strong security infrastructure has not been matched by equivalent AI governance investment. A small cohort of top performers illustrates what high readiness requires: the highest-scoring organization achieved a DSCRI of 84 — a perfect DSMS (100) combined with an AIGMS of 84. Only 19 respondents (4%) exceeded a DSCRI of 70.

Key Findings

1. AI Data Governance and Containment

No AI containment control measured in this survey is deployed by more than 31% of organizations. The modal enterprise has deployed AI into production and built almost none of the infrastructure that would tell it when something goes wrong.

64% of organizations have already directly integrated AI into internal systems or are actively using external AI platforms; among those with AI deployed, 70% are running 3 or more distinct AI use cases in production simultaneously. AI-specific anomaly detection and behavioral monitoring is deployed by 31% of organizations (69% have none). AI kill switch capability is deployed by 21% (79% have no automated mechanism to terminate a misbehaving AI agent). Human-in-the-loop review for high-risk AI actions is deployed by 30% (70% have no human checkpoint). AI-specific DLP policies are deployed by 28% (72% have no technical control governing data flows into AI systems). Purpose binding for AI agents is deployed by 26% (74% do not restrict AI agents to authorized tasks and data scopes). Board-level AI data governance reporting is present at 46% of organizations with AI systems deployed (30% of all organizations) — the highest-rated control.

23% of organizations with AI in production have never tested their AI agent termination capability — an untested kill switch is an assumption, not a control. 22% of organizations with AI deployed have had to revise, roll back, or restrict at least one AI deployment in the past 12 months due to data security concerns. 54% have no standing AI data governance agenda item at board level. 72% cannot trace AI outputs to the source data that produced them; 74% cannot log which AI model version produced a given output; 75% cannot reconstruct source records for a specific AI output on demand.

Forecast Check — Containment Controls: The Forecast report projected purpose binding gaps at 63% and kill switch gaps at 60%, predicting these would narrow modestly through 2026 pipelines. The Annual Survey found purpose binding absent at 74% and kill switch absent at 79%. The gap widened, not narrowed.

Forecast Check — Board Governance: The Forecast report found that 54% of boards were not engaged on AI governance and rated this the strongest correlation in the Forecast survey. The Annual Survey confirms the 54% figure precisely — unchanged six months later.

At current adversary speeds, the absence of machine-speed AI containment is a structural vulnerability, not a governance gap. CrowdStrike's 2026 Global Threat Report documents AI-enabled lateral movement in as little as 27 seconds; Mandiant's M-Trends 2026 report corroborates this trajectory, with median time from initial access to secondary threat group handoff falling from more than eight hours in 2022 to 22 seconds in 2025. An organization without an automated kill switch has no machine-speed response to a machine-speed threat.

KEY FINDING

80% of organizations experienced at least one AI-related security incident in the past 12 months. No AI containment control is deployed by more than 35% of organizations. 23% of organizations with AI in production have never tested their AI kill switch. 72% to 75% cannot trace AI outputs to source data.

2. Sensitive Data Exchange and Fragmentation

54% of organizations use 4 or more separate platforms for sensitive data exchange — and 73% have no technical enforcement restricting which of those channels employees can use.

The channel inventory spans every major category: cloud-based file sharing (SharePoint, Box, Dropbox, Google Drive) is used by 78%; email (including encrypted email) by 73%; managed file transfer (MFT) by 55%; API-based integrations by 48%; web forms by 43%; physical media by 29%; AI tools for data processing or exchange by 38%. Only 4% operate a single unified platform for sensitive data exchange. Despite this widespread channel adoption, the Thales 2026 Data Threat Report finds that only 47% of sensitive cloud-resident data is actually encrypted — and only one in three organizations has complete knowledge of where their sensitive data resides.

Only 39% of organizations validate their sensitive data inventory on a defined quarterly-or-more schedule. 44% have not included AI systems in their data inventory at all. Only 27% of organizations have purpose binding in place — technically restricting AI agents to authorized tasks and data scopes; the remaining 73% operate without this control. 55% have not achieved automated classification across all major systems, and only 47% have classification that enforces downstream controls rather than merely applying labels.

Forecast Check — Centralized AI Gateways and DSPM: The 2026 Forecast report identified centralized AI data gateways as the critical control plane and found 57% still non-centralized. The Annual Survey confirms that only 43% have achieved centralized AI gateway control. The Forecast report also flagged DSPM enforcement: 61% could not enforce data tagging consistently across channels; the Annual Survey finds 55% have not achieved automated classification across all major systems. Same structural failure, still unresolved.

Incident data confirms that channel governance failures produce real outcomes: 26% of organizations experienced sensitive data exposure through an AI tool in the past 12 months; 25% experienced data exfiltration via email, file sharing, or MFT; 22% experienced data exposure through web forms. Tier 1 organizations average 5+ sensitive data exchange platforms; Tier 4 organizations average approximately 3. Among the 253 MFT-using organizations in this survey, only 26% stream logs to SIEM in real time and 39% have tamper-evident audit trails in place — corroborating the Kiteworks 2025 MFT Survey Report finding that 63% of MFT users have not integrated their MFT infrastructure with SIEM.

KEY FINDING

54% use 4 or more sensitive data exchange platforms. 73% lack purpose binding controls for AI agents. 55% have not achieved automated classification across all systems. 63% of MFT users are not SIEM-integrated (MFT Survey Report); this Annual Survey finds only 26% of MFT users stream logs to SIEM in real time.

3. Detection and Response

74% of organizations experienced at least one general security incident; 64% of AI-deploying organizations experienced an AI-specific incident; 80% experienced at least one incident of either type in the past 12 months. The detection infrastructure most organizations have built cannot catch these incidents early enough to limit the damage.

The detection gap is structural, not incidental. 37% of organizations have real-time alerting for AI data access anomalies. Only 33% forward AI access logs to a SIEM. Only 29% generate any logs of AI system data access at all. 25% have no formal detection process whatsoever — one in four organizations has no systematic capability to identify when an AI system is accessing data it should not be accessing. 57% of organizations do not have consistent AI log coverage across all systems they know about, or no mapping at all.

The most common AI-related incident types reported in the past 12 months: sensitive data inadvertently exposed to AI (employees pasting confidential information) at 20%; shadow AI usage at 19%; AI output data leakage at 16%; AI-generated phishing at 14%; prompt injection at 9%.

The interaction between DSMS and AIGMS produces the most operationally significant finding in this section: low AIGMS predicts AI incident rates regardless of DSMS tier. An organization in DSMS Tier 3 with AIGMS index 20 will experience more AI-related incidents than an organization in DSMS Tier 1 with AIGMS index 80. General security maturity does not protect against AI-specific exposure. The AI incident rate drops meaningfully with security maturity tier: Tier 1 organizations experience AI incidents at 91% rates; Tier 4 (Advanced) at 34%.

Forecast Check — AI Detection and Incident Response: The Forecast report identified 60% lacking AI anomaly detection as the incident response gap to close. The Annual Survey found 69% lacking behavioral monitoring for AI systems — the detection deficit widened rather than closed. Our Forecast report warned that 89% had never practiced incident response with AI vendor partners. The Annual Survey finds third-party AI vendor governance remains one of the weakest-rated controls: 27% have not evaluated or technically verified whether AI vendors use their data for model training. The Black Kite 2026 Third-Party Breach Report documented 136 verified third-party breach events in 2025, with approximately 26,000 additional affected organizations never publicly named and a median disclosure lag of 73 days. Despite this exposure, 19% of organizations have never tested or exercised their third-party access revocation capability.

KEY FINDING

80% experienced at least one security incident (general or AI-specific). Only 28% generate any AI data access logs. Tier 1 (Nascent) organizations experience AI incidents at 91% rates; Tier 4 (Advanced) at 34%. 23% have never tested their AI kill switch.

4. Compliance and Regulatory Readiness

63% of organizations experienced at least one compliance consequence in the past 12 months — an audit finding, a required remediation plan, a board escalation, a contractual penalty, or a formal regulatory investigation. Despite that exposure, 7% of organizations have taken no action whatsoever on AI-specific regulatory requirements.

50% of organizations cannot produce a complete AI data access audit record within one business day — and 83% cannot produce one within one hour. 10% cannot produce one at all — their logging infrastructure is insufficient to reconstruct what any AI system accessed. Only 17% can produce a complete AI audit record within one hour from existing dashboards. This is not a capability that can be built reactively.

AI regulation has emerged as the dominant compliance burden: 61% of respondents rank AI-specific regulatory requirements in their top 3 compliance challenges; 25% rank it their single biggest challenge — more than any other item on the list. Compliance consequences reported in the past 12 months: customer or partner audit finding requiring corrective action at 27%; required remediation plan from a regulator or assessor at 26%; board or executive escalation of compliance risk at 23%; regulatory fine or penalty at 19%.

Forecast Check — Audit Trails and DSPM Enforcement: The Forecast report found 33% lacking complete audit trails with 61% carrying fragmented logs — rated HIGH confidence. The Annual Survey finds 67% without tamper-evident audit trails and 50% unable to produce a complete AI data access record within one business day. The Forecast modeled fragmentation. The Annual Survey found it worse.

Only 40% apply a consistent evidence approach across all required frameworks. Governance ownership is fragmented in a way that predicts poor outcomes: 39% treat AI data governance as an add-on to an existing function (CISO, CIO) rather than a dedicated responsibility; only 24% have a dedicated team. Tamper-evident audit trails — the specific evidence type that investigators and auditors examine to confirm records have not been modified after the fact — are in place at only 33% of organizations.

Regulatory deadlines are close, not distant: ISO/IEC 42001 compliance is required for 40% of respondents today. DORA operational resilience obligations for financial entities already require audit-quality evidence on short notice. The EU AI Act's transparency, logging, and human oversight obligations for high-risk systems (Articles 12–14) become enforceable on August 2, 2026 — weeks away at the time of this survey, not years. Organizations that cannot produce an AI data access record within one business day will not be positioned to meet these frameworks' audit obligations once each deadline lands — regardless of what their policy documents say.

KEY FINDING

63% experienced a compliance consequence in the past 12 months. 50% cannot produce an AI audit record within one business day. 61% rank AI-specific regulatory requirements in their top 3 compliance challenges. Only 33% have tamper-evident audit trails.

5. The Human Element: Shadow AI

65% of organizations discovered employees using unapproved AI tools with organizational data in the past 12 months. Among those that discovered it, 16% are finding it monthly or more frequently.

Shadow AI is not an edge case: 16% of organizations encounter it monthly or more; 28% encounter it quarterly; 21% have discovered it at least once but rarely. The 2026 Cost of Insider Risks Global Report identifies shadow AI as the leading driver of negligent insider incidents, with employees routinely sharing confidential documents, source code, and strategy through unsanctioned AI channels; average annual insider risk cost has reached $19.5 million per organization.

The data flowing through unsanctioned AI tools is not generic. Among organizations using employee-facing AI chatbots, 36% are routing customer and client data through them; 33% are routing IT credentials and access requests; 31% are routing employee personal and HR data; 30% are routing financial data. The governance infrastructure to prevent this — AI-specific DLP, approved-channel enforcement, purpose binding — is deployed in fewer than 30% of organizations. The Cisco 2026 Data and Privacy Benchmark Study documents the governance gap in motion: outright organizational bans on AI tool usage dropped from 28% of organizations in 2025 to 7% in 2026 — a 21-point decline — without a corresponding increase in technical controls replacing those bans.

Organizational bans on AI tool usage dropped from 28% to 7% in one year — a 21-point decline — without a corresponding increase in technical controls to replace them.

Discovery without remediation is documentation of a problem, not resolution of one. Fewer than half of organizations that discovered shadow AI subsequently deployed technical controls to prevent recurrence — the most common response was updated policy guidance, a behavioral intervention applied to an architectural problem. 36% cite workforce training on AI data security as a top investment priority, making training the most commonly planned response to a problem training cannot structurally solve. The organizations that have reduced shadow AI incidence are not the organizations with the most training programs — they are the organizations with AI-specific DLP deployed at the transmission layer.

KEY FINDING

65% discovered shadow AI usage in the past 12 months. Only 28% have AI-specific DLP. Only 27% technically enforce approved-channel restrictions. 27% have not evaluated or verified whether AI vendors use their data for model training.

Cross-Analyses: Variations by Key Factors

The survey mean DSMS Index of 39 and mean AIGMS index of 34.7 represent the aggregate condition. Disaggregation by industry, region, and organization size reveals substantial variation with direct implications for how organizations in each cohort should prioritize remediation investment.

Industry

Financial Services and Energy and Utilities lead the DSMS distribution (mean indices of 43.3 and 43.2), driven by regulatory mandate: DORA and model risk management requirements have operationalized control deployment in ways voluntary frameworks have not. Manufacturing (43.0), Healthcare (42.3), and Technology (40.7) follow. Defense Contractors records the lowest DSMS of any sector (15.2), though its sample size of three respondents limits generalizability. Federal Government is the lowest sector with a meaningful sample size (22.7), more than 16 points below the survey mean of 39.

Forecast Check — Government as Critical Outlier: The Forecast report identified Government as the critical outlier — a generation behind: 90% lacking centralized AI gateways, 33% with no dedicated AI controls, 76% missing kill switches, 90% missing purpose binding. The Annual Survey confirms Federal Government records the lowest DSMS of any sector with a meaningful sample size, and Defense Contractors records the lowest AIGMS of any sector. Six months later, unchanged.

Defense Contractors operate under CMMC, a mandatory compliance framework with contractual enforcement, yet record a DSCRI of just 3.75 (DSMS 15 × AIGMS 25/100), the lowest composite risk index of any sector. CMMC 2.0 Level 2 requires 110 NIST SP 800-171 practices including access control, audit and accountability, incident response, and system and communications protection: the foundational controls the DSMS framework measures. Defense Contractors at DSMS 15 are not describing a future compliance challenge. They are describing a current compliance deficit in a framework where noncompliance produces contract loss, not an audit finding.

Technology presents a risk profile its aggregate scores obscure. DSMS of 40.7 is above the survey mean, and its AIGMS of 35.3 tracks the survey mean almost exactly, numbers that suggest an average posture. But 65% of Technology respondents fall in the Exposed quadrant (DSMS below 50, AIGMS below 50), and 16% sit in the Fortified quadrant: strong security infrastructure paired with weak AI governance. The infrastructure exists. The governance perimeter around AI systems does not.

By DSCRI: Energy and Utilities leads at 19.6, Financial Services second at 18.5, Technology sits just above the survey mean DSCRI of 16.2 (Technology DSCRI 16.8), and Federal Government (7.0) and Defense Contractors (4.8) sit at the bottom of the distribution.

Representation analysis tells a related but distinct story. Healthcare produces leaders (organizations in the Resilient quadrant, DSMS ≥45 and AIGMS ≥45) at 1.57 times the expected rate, the highest of any sector and a little counterintuitive given Healthcare's reputation as a security-cautious industry. Energy and Utilities follows at 1.38 times. Technology, despite its reputation as technically sophisticated, produces leaders at only 1.04 times the expected rate. The starkest findings are in Legal, Education, and Federal Government: Legal organizations appear in the leadership quadrant at 0.42 times the expected rate and in the laggard quadrant at 1.39 times; Federal Government's laggard representation index is 1.70, the highest of any sector with a meaningful sample.

Region

North America (32% of the sample) reports the highest AI incident rate of any region at 84%, and the most frequent shadow AI discovery rate: 24% of North American organizations discover shadow AI monthly or more, versus 11% in EU/UK. Overall shadow AI discovery (any frequency) is highest in the Middle East at 76%, with North America at 73% and Europe at 57%. North America also shows the highest AI tool data exposure incident rates, 36% compared to 20% in EU/UK.

EU/UK organizations (56% of the sample) carry the highest rates of AI regulatory concern. 29% rank AI-specific regulatory requirements as their top compliance challenge, versus 24% in North America, reflecting the concrete enforcement environment created by GDPR, the EU AI Act, and NIS2 notification timelines. Notably, EU/UK organizations report the lowest compliance consequence rates of any region at 57%, versus 76% in the Middle East and 69% in North America.

MEA organizations (12% of the sample) report AI incident rates of 78%, above Europe's 75%, alongside the lowest DSMS scores of any region (mean 35.7 of 100). MEA also carries the highest compliance consequence rate of any region at 76%, and reports the highest third-party access revocation testing rates driven by actual incidents rather than planned exercises: 50% of MEA organizations have tested third-party revocation through actual incident response, versus 35% in North America. That is an organization profile responding to incidents and enforcement actions rather than building preventively.

North America is overrepresented in the leadership quadrant at 1.32 times the expected rate. Europe is underrepresented in the leadership quadrant at 0.86 times despite carrying the strongest regulatory pressure of any region in the survey: the region is producing compliance activity but not, at this stage, systematically better governance architecture. The Middle East is underrepresented in leadership at 0.80 times and overrepresented in the laggard quadrant at 1.15 times, the highest laggard overrepresentation of any region.

Organization Size

Enterprise organizations (10,000–25,000 employees) show the highest shadow AI discovery rates and the strongest leadership quadrant representation. Larger organizations carry higher DSMS scores up to the Enterprise band, after which Global organizations decline. AI governance maturity does not scale proportionally with deployment at any size.

The 1,000–4,999 employee cohort (46% of the sample) shows the most fragmented sensitive data exchange environments. These organizations have the resources to adopt multiple specialized platforms but have not yet built the centralized governance architecture to manage them. Enterprise organizations (10,000–25,000 employees) report the highest shadow AI discovery rate of any size band at 70%. DSMS scores rise from Mid (38.9) through Large (40.2) to Enterprise (41.1), but Global organizations (37.1) drop below even the 1,000–4,999 band: coordination complexity at global scale erodes the deployment advantage that scale otherwise provides.

54% of all respondents allocate 25% or more of their IT budget to cybersecurity. Despite this, the mean DSMS across the full sample is 39, Tier 2, developing maturity. Cybersecurity spending does not automatically produce control deployment: the high-spend cohort (≥25% of IT budget, n=248) has a mean DSMS of 39.7, less than one point above the full-sample mean of 39.2. Tier 4 organizations are not distinguished by budget allocation. They are distinguished by architectural discipline: classification that enforces downstream behavior, SIEM integration covering every sensitive data channel, kill switches tested before they are needed, and audit trails that exist before a regulator requests them.

Enterprise organizations (10,000–25,000 employees, n=63) are overrepresented in the leadership quadrant at 1.25 times the expected rate, the highest of any size band with a meaningful sample. Global organizations (over 25,000 employees, n=65) underperform relative to their size: they appear in the leadership quadrant at only 0.55 times the expected rate, the lowest of any size band with a meaningful sample, despite carrying the highest mean security budgets in the survey. Global organizations are not failing at security investment. They are failing to convert security spending into governance architecture.

Implications and Recommendations

Business and Strategic Implications

The 80% incident rate across both general and AI-specific incidents — and the 63% compliance consequence rate — are not forecasts. They are reported outcomes from the 12 months prior to survey completion. The organizations experiencing these outcomes are not a subset of low-maturity outliers — 70% of the full sample is in Tier 1 or Tier 2. These organizations are operating AI systems that are generating incidents they cannot detect, cannot contain at machine speed, and cannot document for regulators when asked.

The cost structure of data security incidents rewards early detection and rapid containment. IBM Security's 2025 Cost of a Data Breach Report documents that breaches contained within 200 days cost organizations roughly $1.1–$1.9 million less on average than those that take longer to contain. Every detection gap documented in this survey — the 73% without validated AI output traceability, the 80% without tamper-evident audit trails, the 91% without complete AI log coverage — directly extends the containment timeline and the associated cost.

Operational Recommendations

1

Classify and tag sensitive data — and make classification enforce downstream controls, not just apply labels.

Data Security Posture Management (DSPM) provides the continuous discovery layer that keeps the inventory current as data moves across systems, channels, and AI pipelines. A Data Policy Engine translates those classifications into enforceable controls across every channel — email, MFT, file sharing, web forms, APIs, and AI tools — ensuring that a sensitivity label triggers encryption, access restriction, or blocking rather than just adding metadata. 55% of organizations have not achieved automated classification across all major systems.

2

Deploy AI-specific DLP that technically blocks transmission to unapproved channels and AI tools — and extend that enforcement through a Data Policy Engine that applies consistent policy across every sensitive data channel from a single control point.

Only 27% have AI-specific DLP today. Shadow AI that is discovered but not blocked at the transmission layer produces the same compliance consequence as shadow AI that is never discovered.

3

Integrate all sensitive data exchange infrastructure — including MFT and AI systems — with SIEM.

63% of MFT users are not SIEM-integrated. Only 33% forward AI access logs to SIEM. Infrastructure outside the SIEM perimeter is infrastructure outside the detection perimeter.

4

Implement and test an AI kill switch.

23% of organizations with AI in production have never tested their AI agent termination capability. A kill switch that has not been tested is an assumption, not a control. Test it on a planned schedule, not for the first time during an active incident.

5

Build AI audit trails that can produce complete access records within one business day.

50% cannot currently produce an audit record within one business day. DORA, NIS2, and the EU AI Act operate on timelines that 50% of organizations cannot currently meet.

6

Assign dedicated AI data governance ownership — not as an add-on to an existing role.

38% currently treat AI data governance as an add-on to an existing role. Dedicated ownership is the structural condition that makes the other five recommendations actionable and sustainable.

7

Consolidate sensitive data exchange platforms.

Organizations in Tier 3 and Tier 4 most commonly operate 4–6 platforms; Tier 1 and Tier 2 organizations most commonly operate 2–3 platforms. Maturity does not reduce platform count — it reflects investment in integrating more channels into a unified governance and audit framework. Each unintegrated platform is an additional SIEM gap, classification gap, and visibility gap.

The Path From Tier 2 to Tier 4

The survey data is unambiguous: organizations that have progressed to Tier 3 or Tier 4 experience AI incident rates of 34%, versus 91% for organizations that have not progressed past Tier 1 (Nascent). The progression is sequential and specific:

Tier Progression	Controls Required	What It Closes
Tier 1 → Tier 2 Nascent → Developing (index 0–18 to 27–45)	Sensitive data classified with controls enforcing downstream access restrictions (DSPM + Data Policy Engine); approved-channel restrictions technically enforced at the platform level; all sensitive data exchange channels unified under a single governance and audit framework	DSMS controls c1–c3: classification, DLP, channel enforcement
Tier 2 → Tier 3 Developing → Established (index 27–45 to 55–73)	AI-specific DLP blocking sensitive data from reaching unapproved AI tools; AI access logs forwarded to SIEM for all production AI systems; AI kill switch implemented and tested on a defined schedule within the past 12 months; complete AI data access audit record producible within one business day; zero-trust access controls applied to AI agent data access	DSMS controls c4–c7: SIEM integration, AI-specific DLP, audit readiness
Tier 3 → Tier 4 Established → Advanced (index 55–73 to 82–100)	Achieve AIGMS ≥ 42 — deploy at least 8 of 19 AI governance capabilities; third-party AI vendor controls contractually enforced and technically verified; dedicated AI data governance ownership assigned; AI regulation compliance externally validated; compliance evidence generated continuously from operational systems	DSMS controls c8–c11: AI governance, third-party controls, continuous evidence

Investment Priorities for the Next 12 Months

Priority	% Citing as Top Priority
AI Security Controls	43%
Workforce Security Training	36%
Compliance Automation	30%
SIEM/SOC Capabilities	28%
Data Sovereignty/Localization	25%
Data Loss Prevention (DLP)	24%

Conclusion

The 9% of organizations in AIGMS Tier 4 are not outliers. They are the destination. What separates them from the 71% concentrated in AIGMS Tiers 1 and 2 — and from the 70% in DSMS Tiers 1 and 2 — is not budget, industry, or headcount. It is architectural discipline: governance enforcement that is technical rather than behavioral, automated rather than manual, measurable rather than attested.

The survey data is unambiguous about where the next 12 months will produce consequences. Data sovereignty tops the list — 61% of organizations rank it their single biggest compliance challenge, yet only 29% use a technical mechanism to enforce it. The gap between what organizations say their sovereignty posture is and what their architecture actually enforces is where regulatory exposure lives. Ungoverned AI systems accessing sensitive data is the second-highest risk named by respondents. Shadow AI is already affecting 65% of organizations. Compliance consequences are already affecting 63%. These are not risk forecasts. They are the current state.

The path from DSMS Tier 2 to Tier 3 and Tier 4, and from AIGMS Tier 2 to Tier 3 and Tier 4, is specific and sequential. It begins with classification that enforces behavior rather than applying labels. It requires SIEM integration for every channel that carries sensitive data, including AI. It requires kill switches that are tested before they are needed — 79% of organizations do not have one. It requires audit trails that can produce records in hours, not weeks — 50% cannot produce a complete audit record within one business day. None of these are aspirational. They are the specific controls that separate the 22% that avoided all incidents from the 80% that experienced at least one.

The intent is there. 43% of organizations name AI security and governance as their top investment priority for the next 12 months. But intent and deployment are different things. Only 9% of organizations have reached AIGMS Tier 4. Only 7% have reached DSMS Tier 4. The organizations that close that gap will not do it through stated priorities. They will do it through the specific architectural decisions this report documents.

This report does not describe a coming threat. It describes the current state. The organizations that treat these numbers as a baseline — rather than someone else's problem — are the ones that will not be answering the same questions when this survey runs again next year.