Kiteworks

ANNUAL SURVEY REPORT

2026 Data Security and Compliance Risk

The AI Data Governance Gap: Private Data, Stalled Compliance, and the Security Reckoning of 2026

Contents

LETTER FROM THE EDITOR

A New Way to Measure What Actually Matters

Every year we publish this survey, and I find myself explaining to executive audiences why two organizations with similar security budgets can have dramatically different outcomes. This year, we finally have the arithmetic to answer that question directly — three indices that, taken together, tell you precisely where an organization stands.

The first is the Data Security Maturity Score (DSMS). Built from 11 binary security controls — encryption, MFT, SIEM integration, kill switches, and others — the DSMS measures the security infrastructure an organization has deployed for secure data exchange, normalized to a 0–100 index. The survey mean is 39 out of 100. That means the average organization has implemented fewer than half the controls we measured. It is not a pessimistic interpretation. It is arithmetic.

The second is the AI Governance Maturity Score (AIGMS). Across 19 AI data governance capabilities — purpose binding, behavioral monitoring, AI-specific DLP, audit trail generation, and others — AIGMS measures what is present — the AI governance capabilities an organization has deployed. A higher AIGMS means better governance. The survey mean is 35 out of 100. The average organization has deployed approximately 7 of the 19 capabilities we measured. That is the governance infrastructure gap that makes AI adoption a liability rather than an asset.

The third — and the one I want to dwell on — is the Data Security and Compliance Readiness Index (DSCRI). The formula is DSMS × (AIGMS/100). It is multiplicative by design. Security maturity, scaled by the proportion of AI governance capabilities deployed. If your DSMS is 39 and your AIGMS is 35, your DSCRI is approximately 14 out of 100. That number is not a flaw in the methodology. It is the point. You cannot offset a 65% AI data governance gap with incremental security investment. The two dimensions must advance together.

What the DSCRI makes explicit is the leverage asymmetry. At the survey mean DSMS of 39, raising AIGMS from 35 to 60 adds approximately 10 points to your DSCRI. Raising DSMS from 39 to 55 — adding four security controls — while AIGMS stays at 35 adds fewer than 6. AI governance investment, at current organizational baselines, returns more readiness improvement per dollar than additional security spend. That is a resource allocation insight, not a theoretical one.

The findings in this report are not comfortable reading. Seventy-four percent of organizations experienced at least one general security incident. 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. Sixty-three percent faced a compliance consequence. The DSCRI of 18 makes clear why. The average organization has enough security maturity to worry about AI risk, but not enough AI data governance to contain it. That is the gap this report documents — and the gap the recommendations in the section titled "Implications and Recommendations" are designed to close.

The maturity gap is not distributed evenly. Plotting all 459 respondents on a DSMS × AIGMS grid reveals meaningful structural differences by industry, company size, and region that aggregate means obscure. The figures at the end of the report visualize these distributions; the representation analysis in each figure's sidebar makes the comparisons group-size-adjusted.

Patrick Spencer, Ph.D.
SVP, Americas Marketing & Industry Research, Kiteworks

Executive Summary

Over 450 security and compliance professionals told us what their organizations are doing with AI and sensitive data. The answers are not encouraging. The survey measures how organizations govern sensitive data in environments where AI deployment, multi-channel fragmentation, and accelerating regulatory requirements are converging simultaneously. Two composite scores anchor the analysis: the Data Security Maturity Score (DSMS), expressed as an index from 0 to 100, with a survey mean index of 39; and the AI Governance Maturity Score (AIGMS), expressed as an index from 0 to 100, with a survey mean index of 35 (35% of maximum possible maturity — approximately 7 of 19 governance capabilities deployed). The average organization in this survey is in Tier 2 — Developing — on both the DSMS and AIGMS frameworks.

The defining finding is not a projected risk. It is a reported outcome. 74% percent of organizations experienced at least one general security incident. 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 single highest concern associated with generative AI.[1] 63% experienced a compliance consequence — an audit finding, a required remediation plan, a board escalation, a contractual penalty, or a formal regulatory investigation.[2] 65% discovered employees using unapproved AI tools with organizational data. These numbers represent what happened in the 12 months prior to survey completion.

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: The Forecast was correct about the direction. It was optimistic about the scale.

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. Our 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.

The Governance Deficit in One View

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%

Each row is a governance control that the preceding row's organizations should 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. This survey measures the cost of that gap.

The governance gap that produced these outcomes is measurable and specific. 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[3] — 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 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 path from Tier 2 to Tier 4 is documented in this report.

The Maturity Gap in Numbers

Data Security Maturity Score (DSMS)

39 / 100
Mean DSMS index

70%
Score in Tier 1 or Tier 2 (index below 45)

7%
Reached Tier 4 — Advanced (index 82 – 100)

AI Governance Maturity Score (AIGMS)

35 / 100
Mean AIGMS index

71%
Score in Tier 1 or Tier 2 (index below 45)

9%
Reached Tier 4 — Advanced (index 73 – 100)

Overarching Outcomes (past 12 months)

80%
Experienced at least one security incident in the past 12 months

63%
Experienced a compliance consequence

65%
Discovered shadow AI usage

Note: Tier definitions and full distributions: see pages 10–12.

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

About The Survey

This report is based on primary research conducted with 459 security, compliance, and technology professionals. Participants are cybersecurity, IT, risk, and compliance leaders at organizations with 1,000 or more employees (92% of the sample), distributed across 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:

1. AI data governance and containment controls
2. Sensitive data exchange channels
3. Detection and response capabilities
4. Compliance automation and evidence readiness
5. Regulatory framework requirements
6. Human and behavioral dimensions of AI data risk
7. Planned investment priorities

Responses were collected in the second quarter of 2026. All findings are self-reported by survey participants; figures reflect the share of 459 respondents unless otherwise noted.

FIGURE 1

Survey Respondents by Organization Size

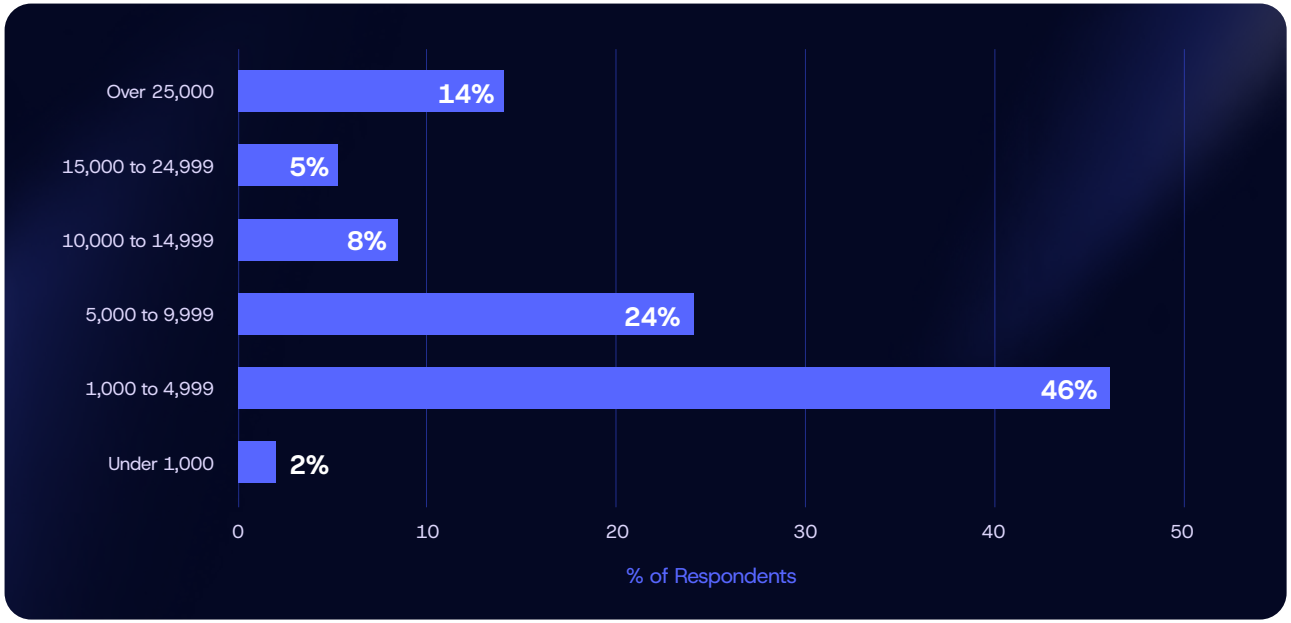

FIGURE 2

Respondent Role/Function

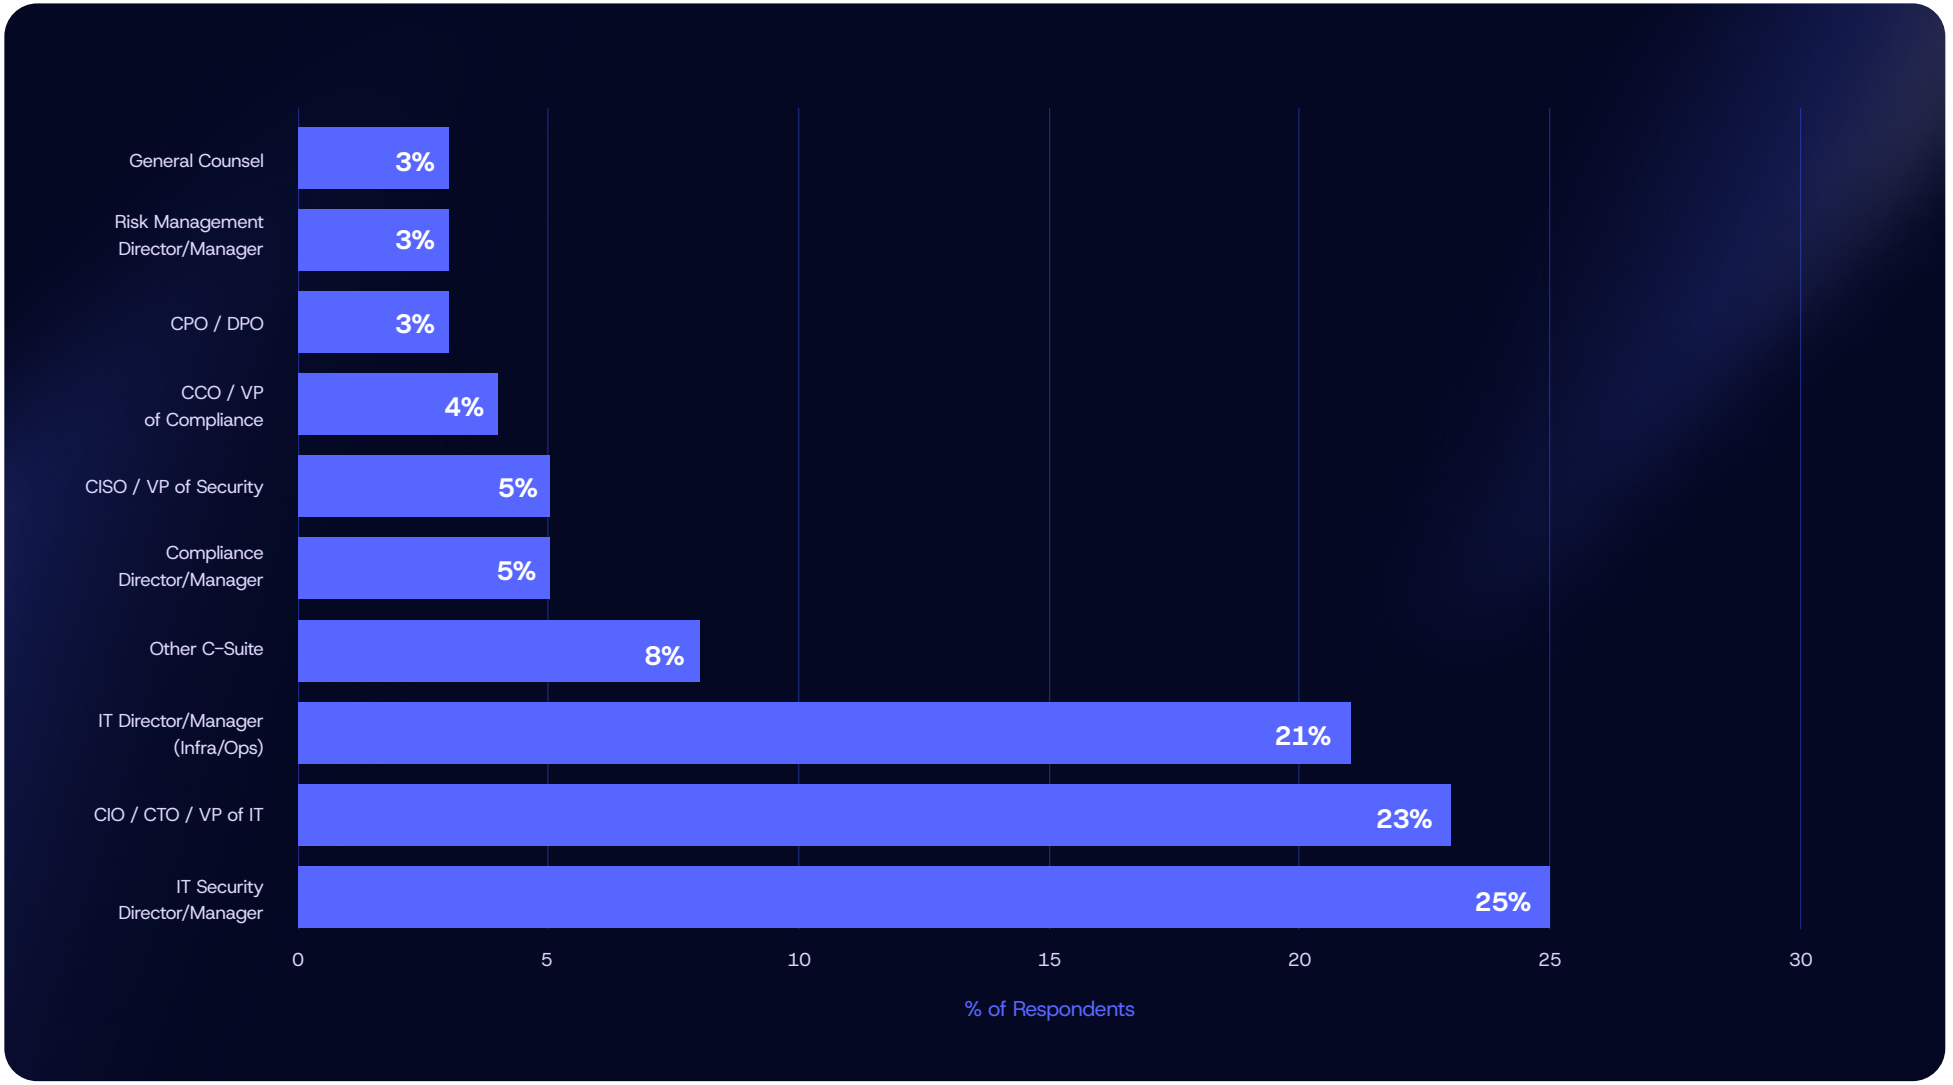

The Security Maturity Framework

Three composite scores are used throughout this report to anchor cross-sectional analysis. All three are computed directly from survey responses — not from self-reported estimates. The Data Security Maturity Score (DSMS) measures the breadth of operational security controls an organization has deployed. The AI Governance Maturity Score (AIGMS) measures the depth of AI data governance capabilities an organization has deployed. The Data Security and Compliance Risk Index (DSCRI) combines both into a single composite that reflects overall organizational readiness. Together, they provide a three-dimensional picture of organizational risk posture — breadth of security controls, depth of AI governance, and combined readiness — that is reproducible, comparable across cohorts, and grounded in specific, verifiable control states.

Data Security Maturity Score (DSMS)

The DSMS is a composite score expressed as an index from 0 to 100, computed from 11 binary security control variables measured in this survey. Each control is scored 1 (deployed) or 0 (not deployed). The DSMS is the sum of all 11 binary scores for each organization.

$$\text{DSMS} = \Sigma\left(C_1 + C_2 + C_3 + C_4 + C_5 + C_6 + C_7 + C_8 + C_9 + C_{10} + C_{11}\right)$$

Where each $c_n \in \{0, 1\}$ and the 11 controls are:

Control	Variable
C_1	Automated data classification deployed across all major systems
C_2	DLP policies technically enforced across all sensitive data channels
C_3	Approved-channel restrictions enforced (not policy-only)
C_4	SIEM integration for all sensitive data exchange infrastructure
C_5	DLP policies deployed
C_6	Board-level AI data governance reporting (standing agenda item)
C_7	Third-party AI vendor data controls contractually enforced and technically verified — not reliant on attestation alone
C_8	AI-specific DLP policies deployed
C_9	AI access logs forwarded to SIEM
C_{10}	AI kill switch / agent termination capability deployed
C_{11}	Compliance evidence generated continuously from operational systems

These 11 controls span general data security infrastructure (C_1–C_7, C_{11}) and AI-specific data security controls (C_8–C_{10}). The split reflects the report's scope: overall data security posture as the foundation, with AI data security controls as the emerging layer where the greatest gaps exist.

Maturity Baseline: Year-Over-Year

The 2026 Forecast Report modeled expected control deployment rates for 2026. The Annual Survey measures what was actually deployed. The gap between forecast and outcome is the most direct measure of whether the industry is closing the governance deficit or widening it.

Metric	2026 Forecast Projection	2026 Annual Survey Measured
Mean DSMS Index	—	39
Mean AIGMS Index	—	35
AI kill switch deployed	60% lacking	70% lacking
Behavioral monitoring in place	60% lacking	69% lacking
Tamper-evident audit trails	33% lacking	67% lacking
Complete audit record within 1 business day	73% unable	73% unable

The detection deficit widened rather than closed. The audit trail deficit held exactly at the forecast level. Neither DSMS nor AIGMS has a 2026 Annual Survey equivalent for direct index comparison; establishing a measured baseline here enables year-over-year tracking beginning with the 2027 Annual Survey.

FIGURE 3

Data Security Maturity Score (DSMS)

11 binary security controls · Index 0–100 · Higher = more controls deployed

Survey mean **39 / 100**

Average controls deployed **4.3 of 11**

TIER	INDEX RANGE	CONTROLS DEPLOYED	% OF ORGS
Tier 1 **Nascent**	0 – 18	0 – 2 of 11 controls	24% n = 111 orgs
Tier 2 **Developing**	27 – 45	3 – 5 of 11 controls	46% n = 212 orgs
Tier 3 **Established**	55 – 73	6 – 8 of 11 controls	23% n = 104 orgs
Tier 4 **Advanced**	82 – 100	9 – 11 of 11 controls	7% n = 32 orgs

Index gaps (19–26, 46–54, 74–81) cannot occur naturally given the 11-control structure and contain no respondents.
DSMS = Σ(c1–c11) / 11 × 100

FIGURE 4

Data Security Maturity Score (DSMS) Distribution

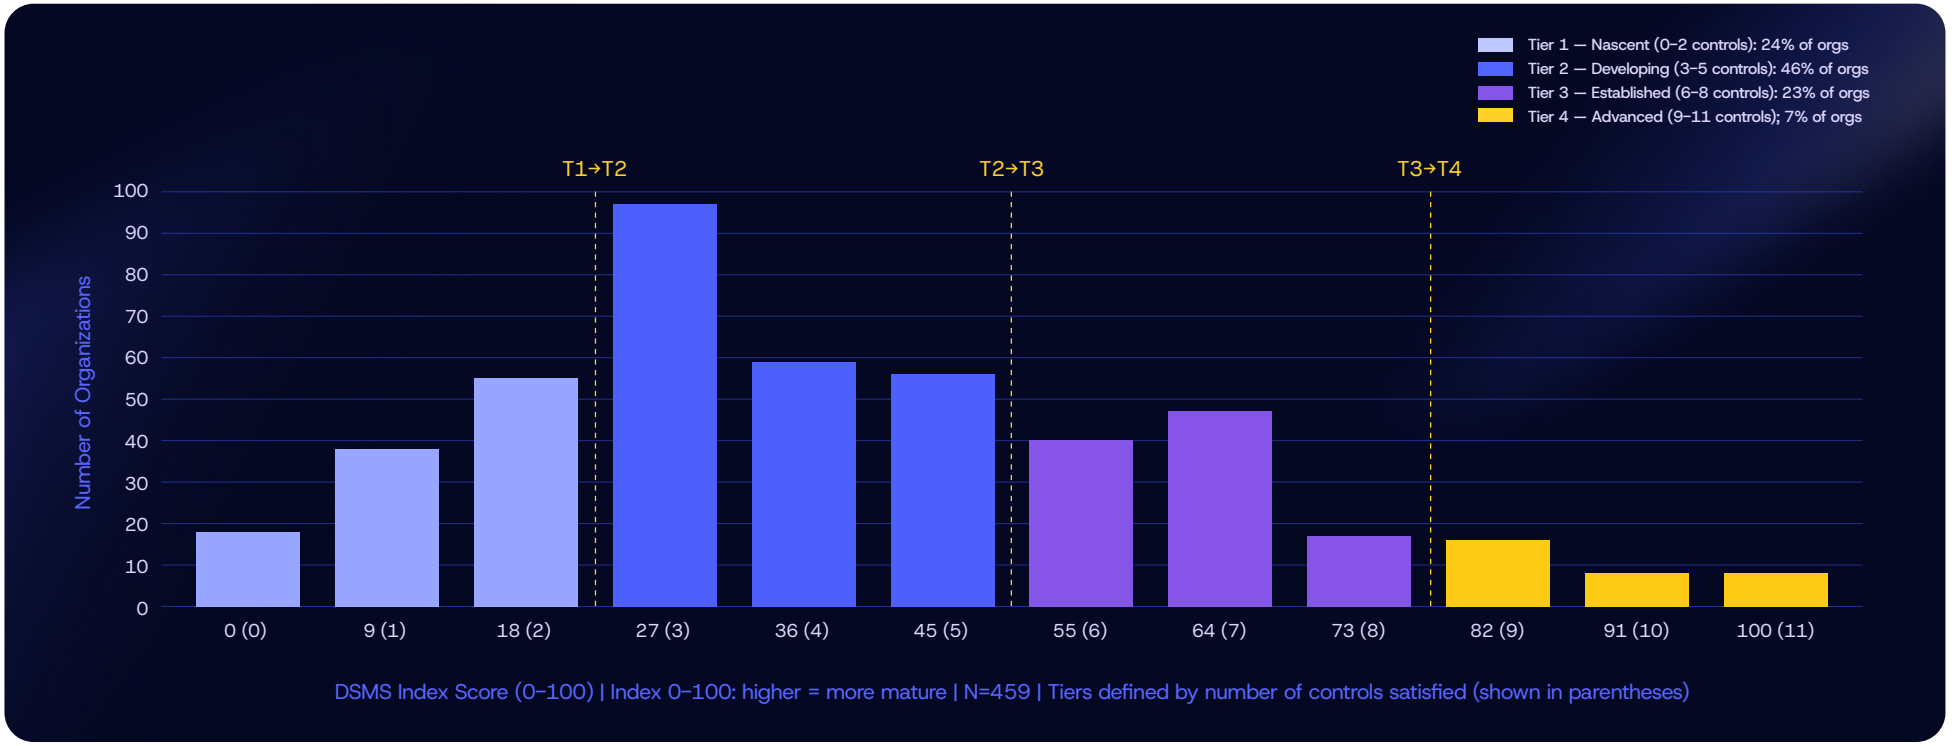

For readability, DSMS and AIGMS values throughout this report are expressed as index scores from 0 to 100, computed by normalizing the raw score against the maximum possible value.

The DSMS is computed from operational controls — not self-reported confidence. That distinction matters. When asked what they know about their own security infrastructure without consulting a colleague, 49% of respondents cannot identify their SIEM or log management platform. 41% cannot identify the audit standard or framework they are audited against.[4] Organizations cannot score well on the DSMS by believing they have controls in place. The controls either exist or they do not. The visibility gaps suggest that self-reported maturity assessments would produce materially higher scores than the control-based DSMS — which is precisely why the DSMS is control-based.

FIGURE 5

AI Governance Maturity Score (AIGMS)

19 binary governance capabilities · Index 0–100 · Higher = more capabilities deployed

Survey mean **35 / 100**

Average capabilities deployed **6.7 of 19**

TIER	INDEX RANGE	CAPABILITIES DEPLOYED	GOVERNANCE GAPS REMAINING	% OF ORGS
Tier 1 **Nascent**	0–26	0 – 5 of 19 capabilities	14 – 19 gaps remaining	48% n = 220 orgs
Tier 2 **Developing**	27–44	6 – 8 of 19 capabilities	11 – 13 gaps remaining	23% n = 106 orgs
Tier 3 **Established**	45–72	9 – 13 of 19 capabilities	6 – 10 gaps remaining	20% n = 93 orgs
Tier 4 **Advanced**	73–100	14 – 19 of 19 capabilities	0 – 5 gaps remaining	9% n = 40 orgs

AIGMS = 100 − (governance gaps / 19) × 100 · Each capability scored 1 (deployed) or 0 (not deployed) across 4 dimensions: Technical Data Containment, Data Audit & Traceability, Data Policy Enforcement, Organizational Data Governance

FIGURE 6

Mean Data Security Maturity Score (DSMS) By Industry

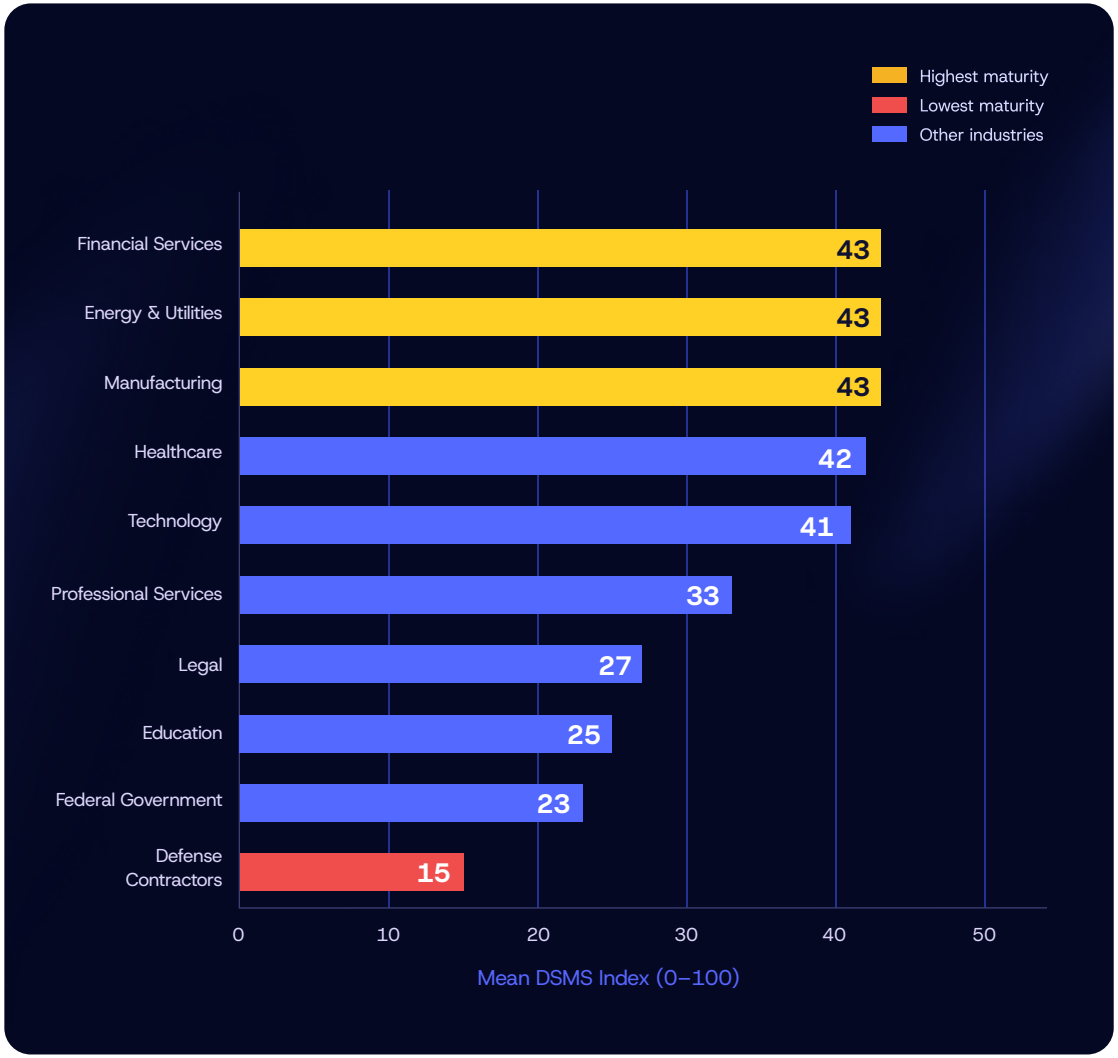

FIGURE 7

Mean Data Security Maturity Score (DSMS) By Region

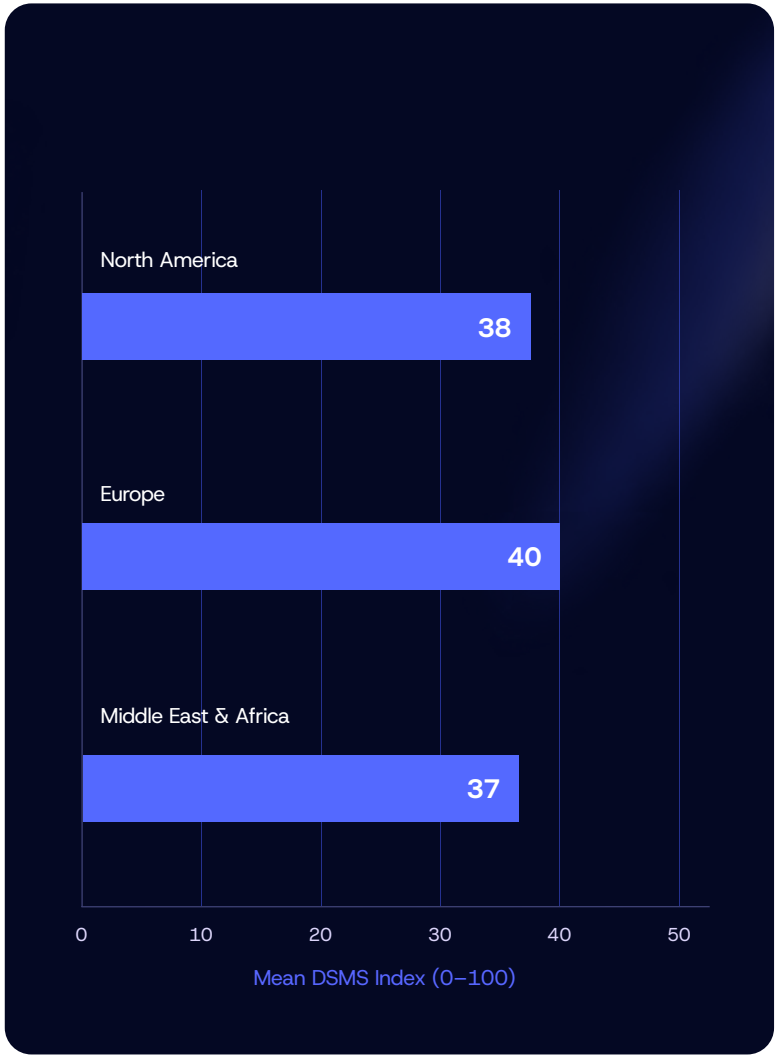

AI Governance Maturity Score (AIGMS)

Each of 19 AI data governance capabilities is scored: $g_n = 1$ if the capability is present and operational, $g_n = 0$ if absent. AIGMS is the sum of all 19 scores — the count of capabilities deployed. A score of 7 means 7 capabilities are operational. The survey mean index of 35 means the average organization has deployed approximately 7 of 19 measured AI data governance capabilities.

$$\text{AIGMS} = \Sigma(g_n) \text{ for } n = 1 \text{ to } 19, \text{ where } g_n \in \{0, 1\}$$

AIGMS is expressed as an index from 0 to 100 (normalized from 19 capability variables): 0 means no capabilities present, 100 means all 19 capabilities present (full measured maturity). *Like DSMS, higher AIGMS is better.* The score reflects what governance is present — the higher the score, the more AI governance capabilities are deployed — not merely the absence of best practices. An organization with AIGMS index 26 carries materially higher probability of sensitive data exposure, compliance failure, and containment breakdown than one with AIGMS index 79. The score is designed to make that distinction visible.[5]

Dimension 1

Technical Data Containment (5 capabilities)

- AI kill switch / agent termination capability
- SIEM integration for AI system data access logs
- Zero-trust access controls applied to AI agent data access
- Network isolation capability to cut off AI agent data access
- Single control point to revoke AI agent access across all systems

Dimension 2

Data Audit and Traceability (6 capabilities)

- Complete AI audit log coverage across all AI systems
- Output-to-source data traceability (AI outputs traceable to source records)
- Audit record producible within one business day
- Formal process to detect AI systems accessing inappropriate data
- AI deployment rollback or restriction process operational
- Board-level AI data governance reporting (standing agenda item)

Dimension 3

Data Policy Enforcement (4 capabilities)

- AI-specific DLP policies blocking sensitive data transmission to unapproved AI tools
- Purpose binding restricting what data AI agents can access
- Human-in-the-loop review before AI systems access high-risk data
- Data minimization enforced — AI agents receive only minimum necessary data

Dimension 4

Organizational Data Governance (4 capabilities)

- Dedicated AI data governance ownership (not assigned as add-on to existing role)
- AI data governance responsibilities documented in writing
- Third-party AI vendor data controls in place and contractually enforced
- AI regulation compliance externally validated (not self-attested only)

DSMS Survey Mean:

39 / 100

AIGMS Survey Mean:

35 / 100

DSMS and AIGMS values are expressed as index scores throughout this report. A DSMS index of 39 means an organization has deployed the equivalent of 39% of the maximum measurable security control coverage. An AIGMS index of 35 means the organization has deployed approximately 7 of 19 governance capabilities — 35% of maximum maturity — and is missing the remaining 12.

The AIGMS reflects what governance is present — the higher the score, the more capabilities the organization has actually deployed. This represents active risk — not merely the absence of best practices. An organization with AIGMS index 26 carries materially higher incident probability, compliance exposure, and containment failure risk than one with AIGMS index 79. The score is designed to make that distinction visible.

FIGURE 8

Mean AI Governance Maturity Score (AIGMS) by Industry

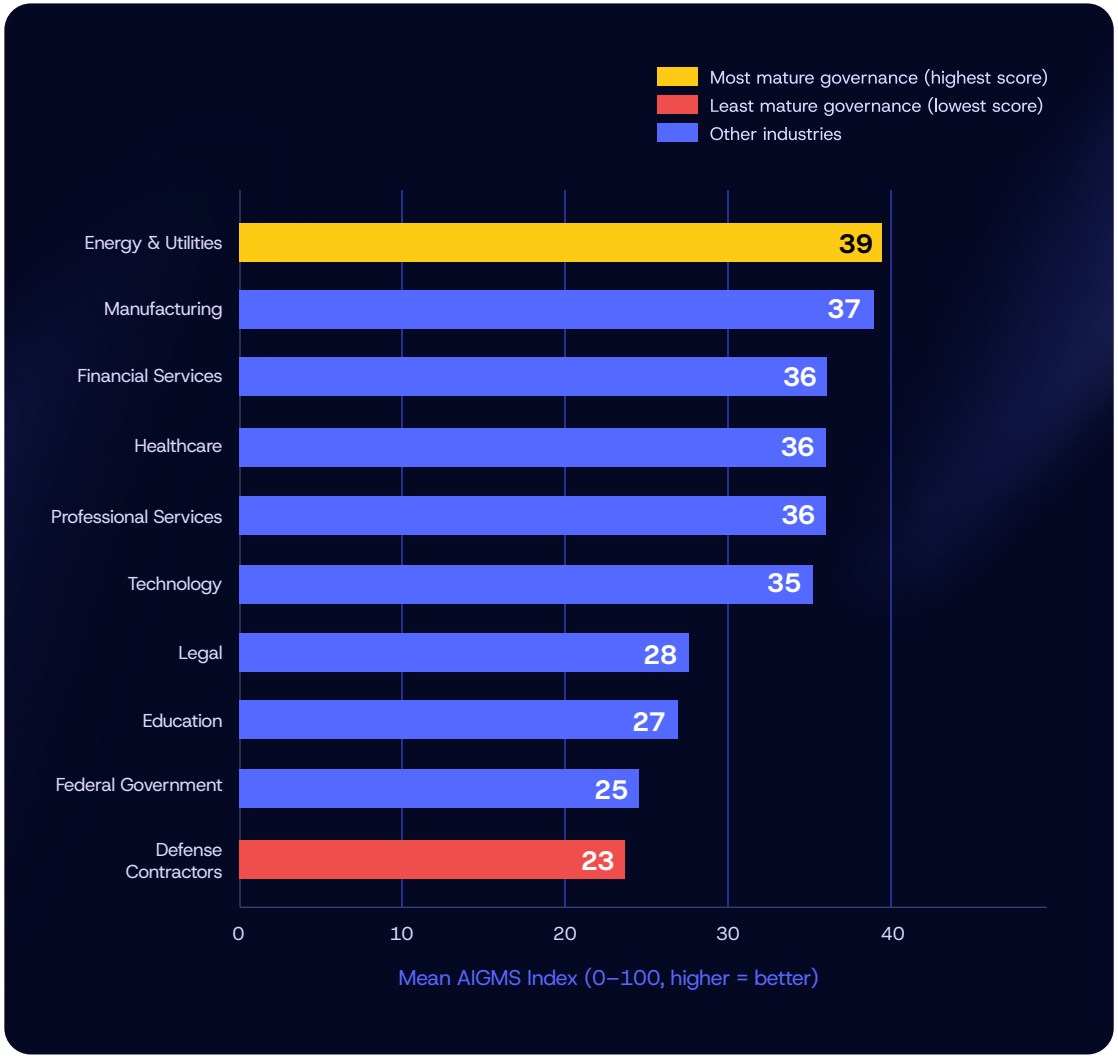

FIGURE 9

Mean AI Governance Maturity Score (AIGMS) by Region

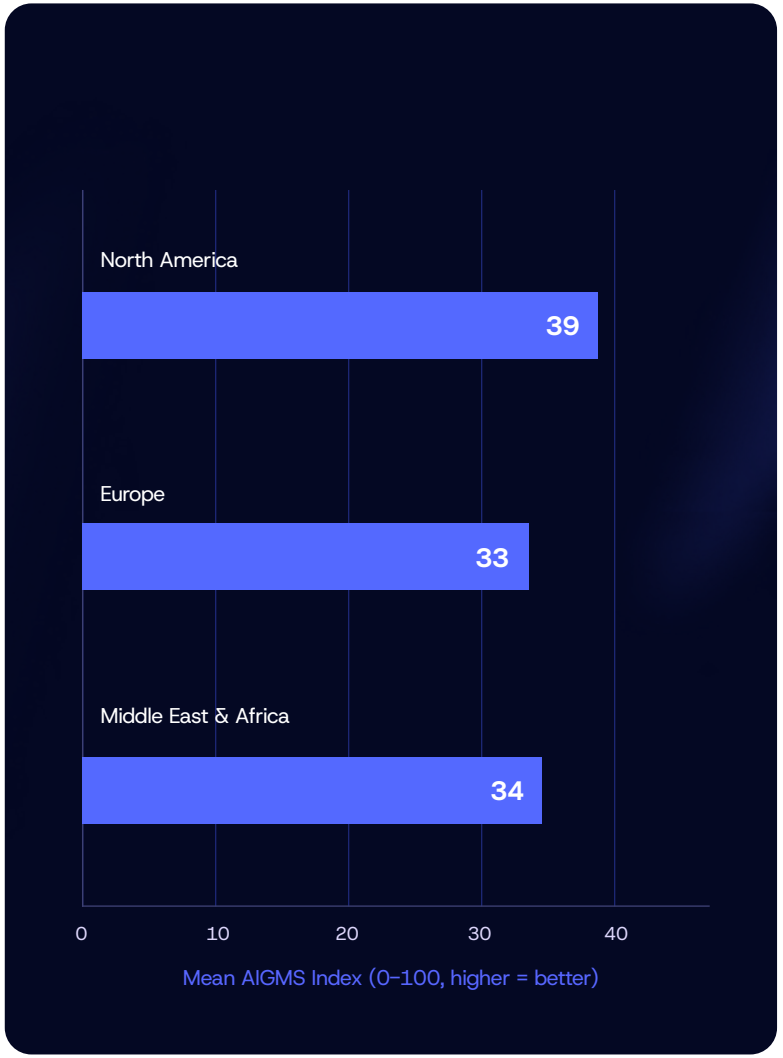

How DSMS and AIGMS Interact

DSMS and AIGMS are complementary but distinct measurements. DSMS measures operational data security controls — the foundational infrastructure an organization has deployed to govern how sensitive data moves, who can access it, and whether that access is logged. AIGMS measures AI data governance specifically — the technical and operational controls that determine what data AI systems can access, process, output, and store. An organization can carry moderate DSMS alongside low AIGMS — a combination found across the Technology sector, where AI deployment has substantially outpaced AI data governance. This is the highest-risk profile in the survey: organizations with enough security infrastructure to feel confident, but insufficient AI data governance to control what their AI systems are doing with sensitive data.[6]

The two scores define four distinct organizational risk profiles. Mapping DSMS against AIGMS creates a diagnostic framework that is more operationally precise than either score alone — and provides a direct prioritization signal for remediation investment.

	Low AIGMS (below 50) — AI Governance Gaps Present	**High AIGMS** (50+) — AI Governance in Place
High DSMS (55+) — Security Foundation Strong	**"FALSE CONFIDENCE"** Strong general security, no AI governance. Technology sector profile. Feels protected. Is not.	**"AI-READY"** Security foundation and AI governance aligned. 11% of organizations. The destination.*
Low DSMS (below 45) — Security Foundation Weak	**"DUAL EXPOSURE"** Both security and AI governance absent. Federal Government, Defense Contractors, MEA. Incidents are not a risk. They are a schedule.	**"FOUNDATION FIRST"** AI governance ahead of general security. Rare — typically AI-native organizations that skipped foundational infrastructure.

* Quadrant thresholds (DSMS ≥55, AIGMS ≥50) are independent analytical cutoffs used in this diagnostic framework. The scatter plot representation analyses in Figures 33, 36, and 39 use DSMS ≥45 and AIGMS ≥45, consistent with the Tier 3 entry boundary for each index.

A derived metric — the DSCRI, computed as DSMS × (AIGMS/100) — quantifies how far AI governance lags security maturity. The survey mean DSCRI is 16.2 (computed from individual respondent scores; the product-of-means estimate of 13.6 understates the true mean due to positive covariance between DSMS and AIGMS): The average organization's AI governance maturity substantially lags its general security posture.[7,8] The survey median DSCRI is 11.5 — half of all organizations score below 12.[9] The gap between

the mean (18) and the median (12) reflects a right-skewed distribution: A small number of high-performing organizations pull the average up, while the majority cluster at the low end.[10] By sector, the governance gap — the spread between security maturity and combined readiness — 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.[11]

FIGURE 10

Data Security and Compliance Readiness Index (DSCRI) — Distribution Across Organizations

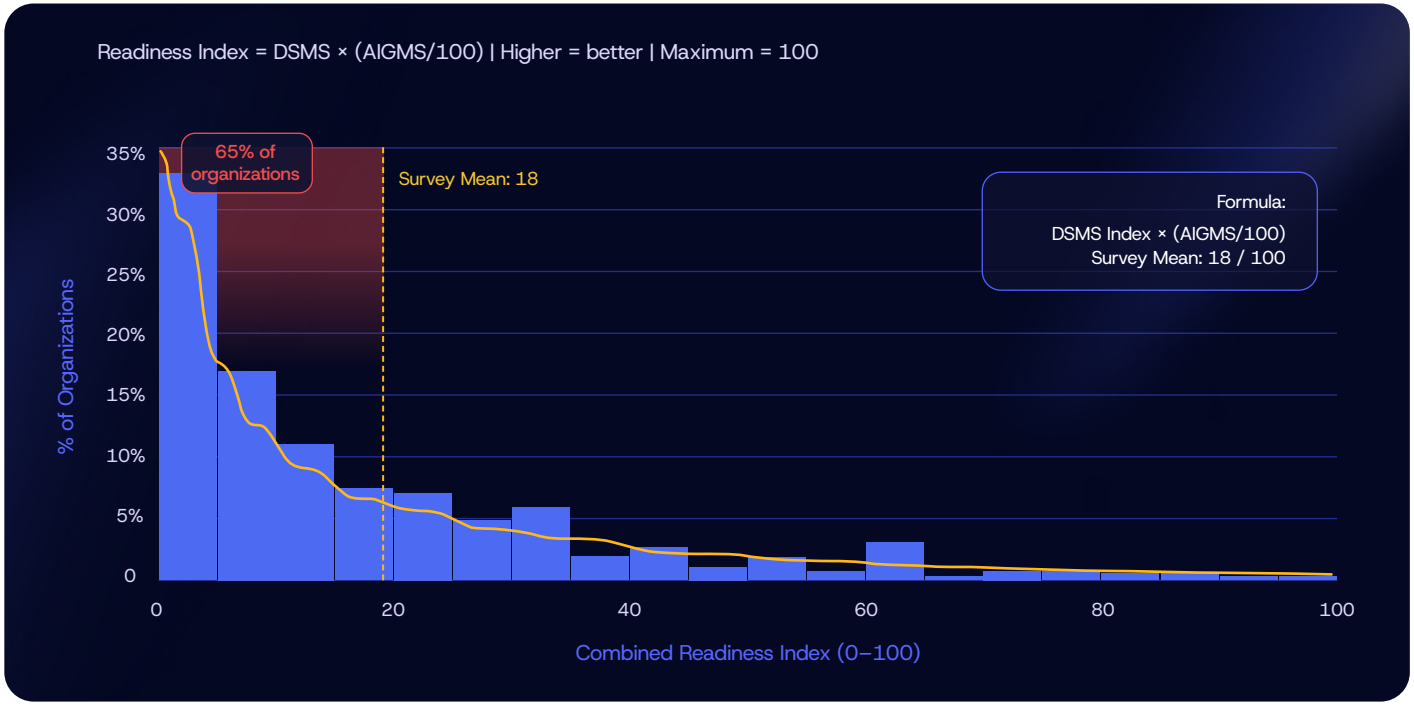

A small cohort of top performers illustrates what high readiness requires. The highest-scoring organization in this survey achieved a DSCRI of 84 — the product of a perfect DSMS (100) and an AIGMS of 84, meaning full deployment of all 11 security controls alongside deployment of 16 of 19 AI governance capabilities. Six respondents (1% of the sample) exceeded a DSCRI of 70. Breaking the midpoint threshold of 50 requires both DSMS and AIGMS to exceed 70 simultaneously — a standard only 19 respondents (4%) met.[12] 61% of organizations score below the survey mean. The index is not structurally depressed; it reflects where the market actually is.[13]

FIGURE 11

Data Security and Compliance Readiness Index (DSCRI) by Industry and Tier

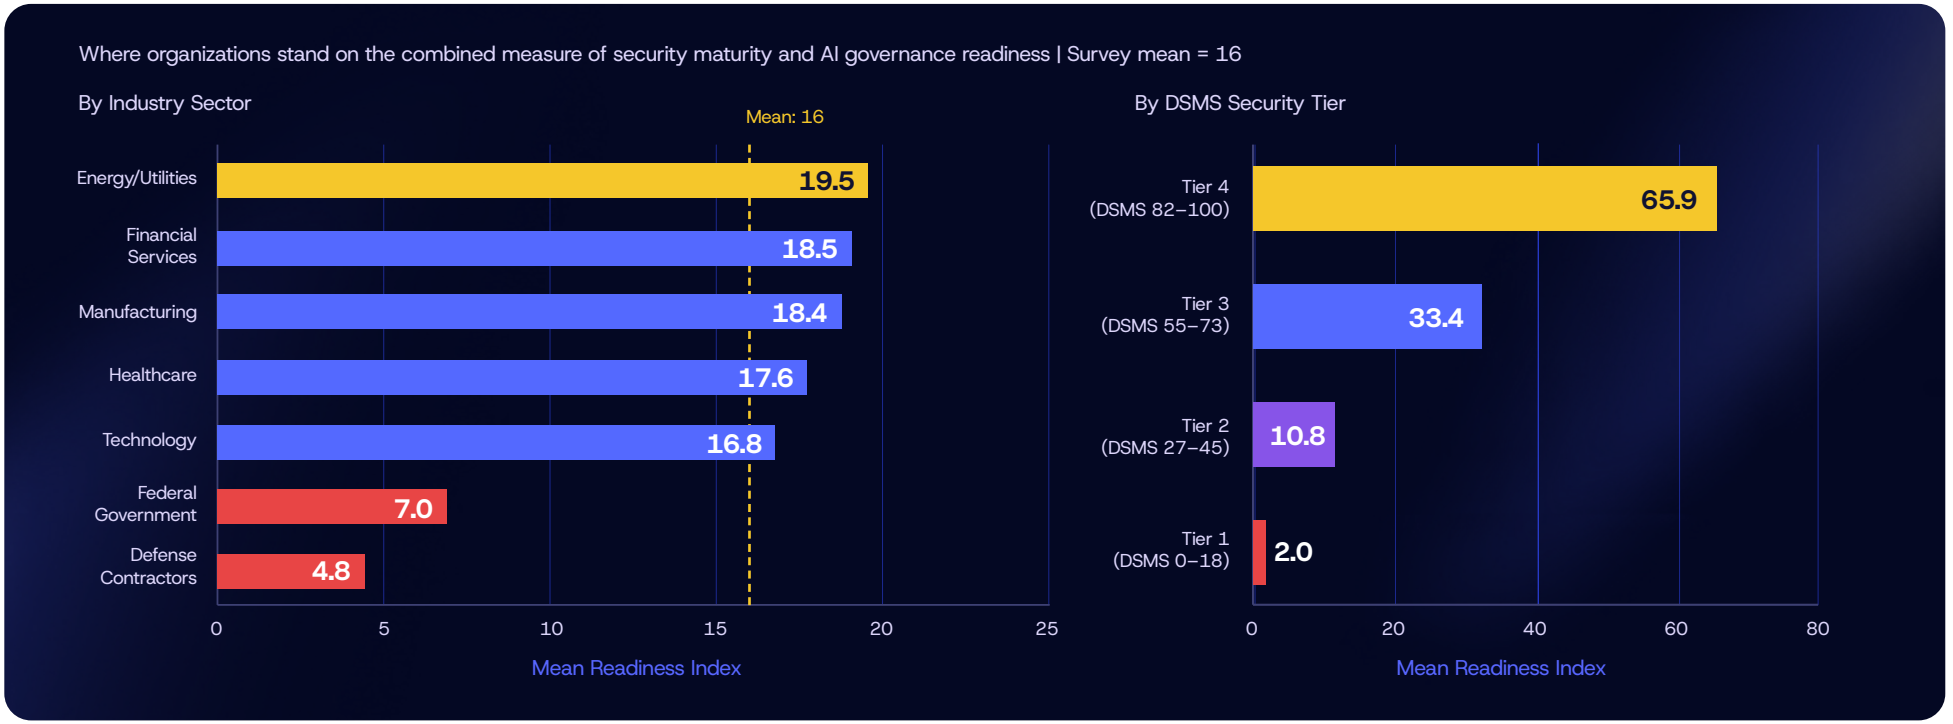

By Industry Sector — Energy & Utilities leads at 19.5, narrowly ahead of Financial Services (18.5) and Manufacturing (18.4), both sectors where regulatory pressure has driven parallel investment in security controls and AI governance. Healthcare (17.6) and Technology (16.8) cluster near the survey mean, though Technology's position conceals a structural risk: its DSMS of 41 is the second-highest in the survey, but its AIGMS of 35 matches the overall average — meaning AI deployment has outrun governance at the sector level. Federal Government (7.0) and Defense Contractors (4.8) sit well below the mean, reflecting simultaneous weakness in both dimensions.

By DSMS Security Tier — The DSCRI progression across tiers confirms that security maturity and AI governance co-develop: Tier 1 Nascent organizations average a DSCRI of 2.0; Tier 2 Developing, 10.8; Tier 3 Established, 33.4; Tier 4 Advanced, 65.9. The 20-fold difference between Tier 1 and Tier 4 is not explained by AI governance effort alone — Tier 4 organizations have the security infrastructure that makes AI governance controls operable. Organizations without that foundation cannot purchase their way to a high DSCRI through governance initiatives alone.

Key Findings

AI Data Governance and Containment

No AI containment control measured in this survey is deployed by more than 31% of organizations.[14]

To be precise: The modal enterprise has deployed AI into production and built almost none of the infrastructure that would tell it when something goes wrong. This is a governance posture. It is not a good one.

The scale of AI deployment makes the governance gap urgent rather than theoretical. 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.[15] Each production use case is a data access surface. Employee-facing chatbots, internal copilots, AI-powered code generation, customer-facing agents — these are not pilots. They are operational systems processing sensitive data in environments where, as the control adoption numbers show, the governance infrastructure has not kept pace.[16]

The control adoption data is unambiguous in its direction. AI-specific anomaly detection and behavioral monitoring is deployed by 31% of organizations — meaning 69% have no behavioral monitoring for AI systems. AI kill switch capability is deployed by 21% — meaning 79% have no automated mechanism to terminate a misbehaving AI agent. Additionally, human-in-the-loop review for high-risk AI actions is deployed by 30% — meaning 70% have no human checkpoint.[17] AI-specific DLP policies are deployed by 28% — meaning 72% have no technical control governing data flows into AI systems. Purpose binding for AI agents is deployed by 26% — meaning 74% do not restrict AI agents to authorized tasks and data scopes.[18] Board-level AI data governance reporting as a standing agenda item is present at 46% of organizations of organizations with AI systems deployed (or 30% of all organizations) — the highest-rated control. AI audit logs forwarded to SIEM are deployed by 33%.[19]

Forecast Check – Containment Controls

The Forecast report projected purpose binding gaps at 63% and kill switch gaps at 60% — and predicted 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. The Forecast Report was right that containment controls were the battleground. It was optimistic about how much ground organizations would take.

Each of these controls is a data control — not a policy framework. Kill switches stop data movement. SIEM integration makes data access visible. Purpose binding restricts data reach. AI-specific DLP blocks data transmission. The absence of these controls means AI systems are accessing, processing, and outputting sensitive data without any of the mechanisms that would make that access detectable, stoppable, or auditable.[20]

Kill switch capabilities deserve specific attention because they represent the last line of defense when an AI deployment behaves badly.[21] 23% of organizations with AI in production have never tested their AI agent termination capability. An untested kill switch is not a control. It is an assumption. The consequences of inadequate containment are already materializing: 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.[22, 23] 54% have no standing AI data governance agenda item at board level — meaning for the majority, AI data governance is either bundled into broader cybersecurity reporting, addressed only ad hoc, or absent from executive visibility entirely.[24]

Forecast Check – Board Governance

The Forecast report found that 54% of boards not engaged on AI governance and rated this the strongest correlation in the Forecast survey — organizations with board engagement scored 26–28 points higher on every AI maturity metric. The Annual Survey confirms the 54% figure precisely: 54% still have no standing AI data governance agenda item at board level. Six months after the Forecast was published, the board engagement gap has not moved.

AI output traceability is a fundamental requirement for regulatory audit response and incident investigation. 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.[25] This is not a marginal gap. It is a near-universal absence of the forensic infrastructure that compliance and incident response require.[26] The 80% incident rate reported in this survey across both general security and AI-specific incidents is a direct consequence of the control gaps documented here.[27]

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.

Adversary Speed Has Outpaced Containment Architecture

At current adversary speeds, the absence of machine-speed AI containment is not a governance gap — it is a structural vulnerability. CrowdStrike's 2026 Global Threat Report documents AI-enabled lateral movement in as little as 27 seconds.[28] Mandiant's M-Trends 2026 report corroborates this trajectory: The median time from initial access to secondary threat group handoff fell from more than eight hours in 2022 to 22 seconds in 2025.[29] An organization without an automated kill switch has no machine-speed response to a machine-speed threat. Human review processes that operate on minutes or hours timelines cannot close a 27-second exposure window.[30,31]

FIGURE 12

AI Containment and Governance Control Adoption

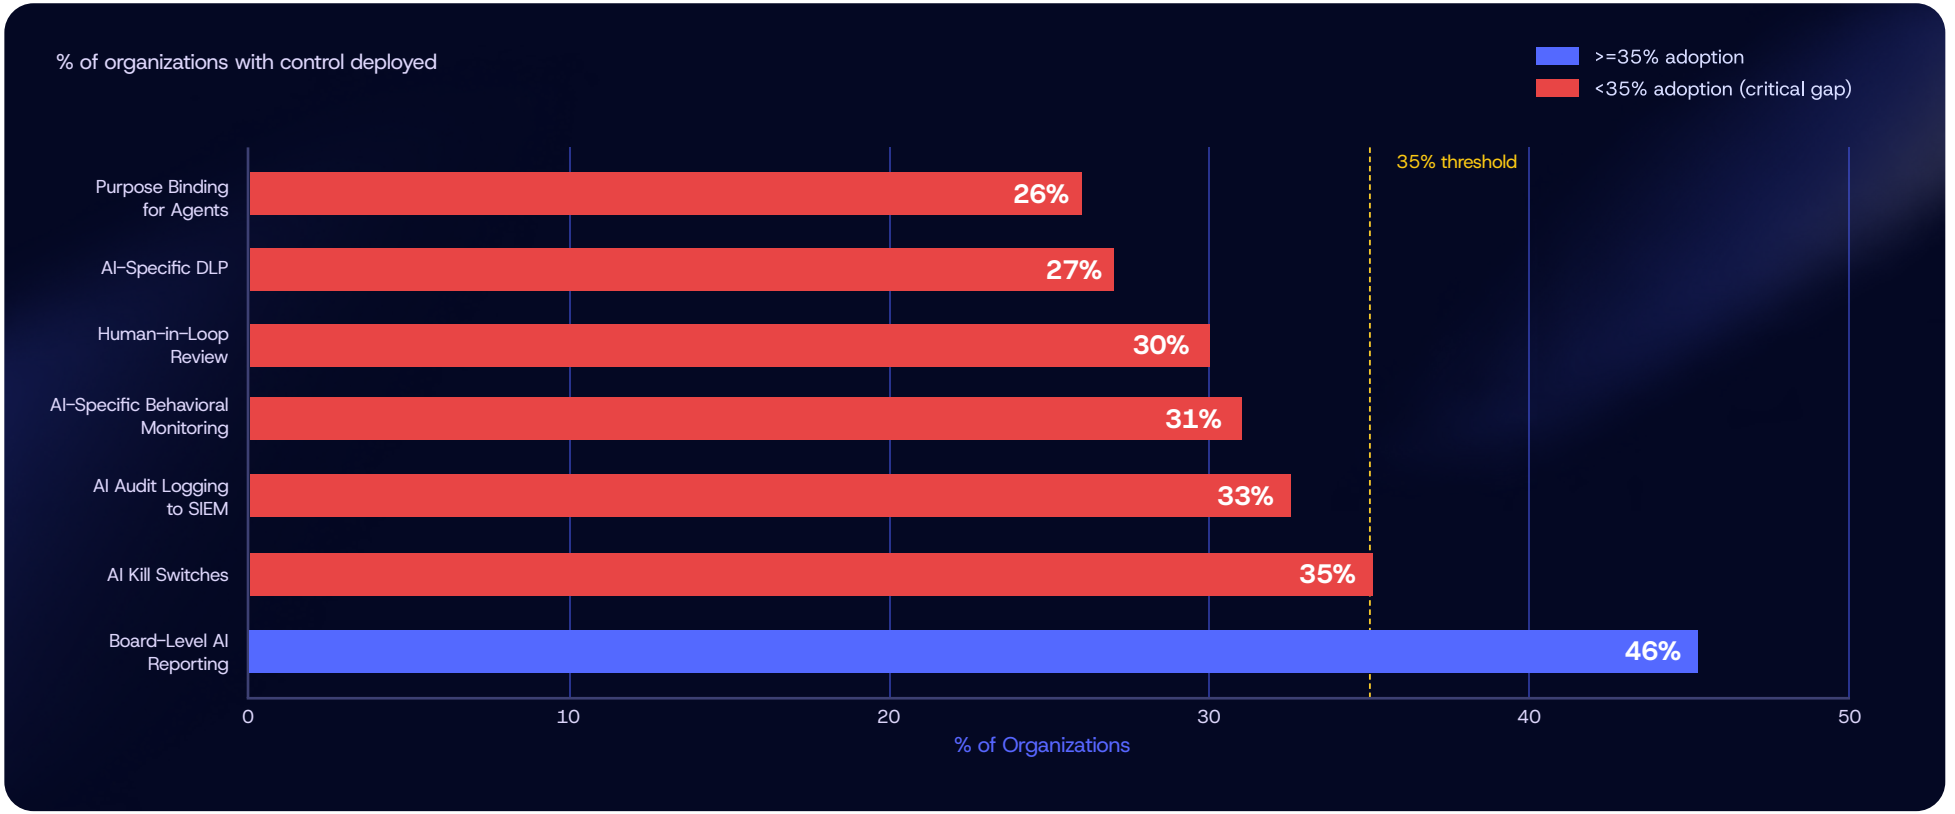

FIGURE 13

AI Kill Switch and Containment Capability

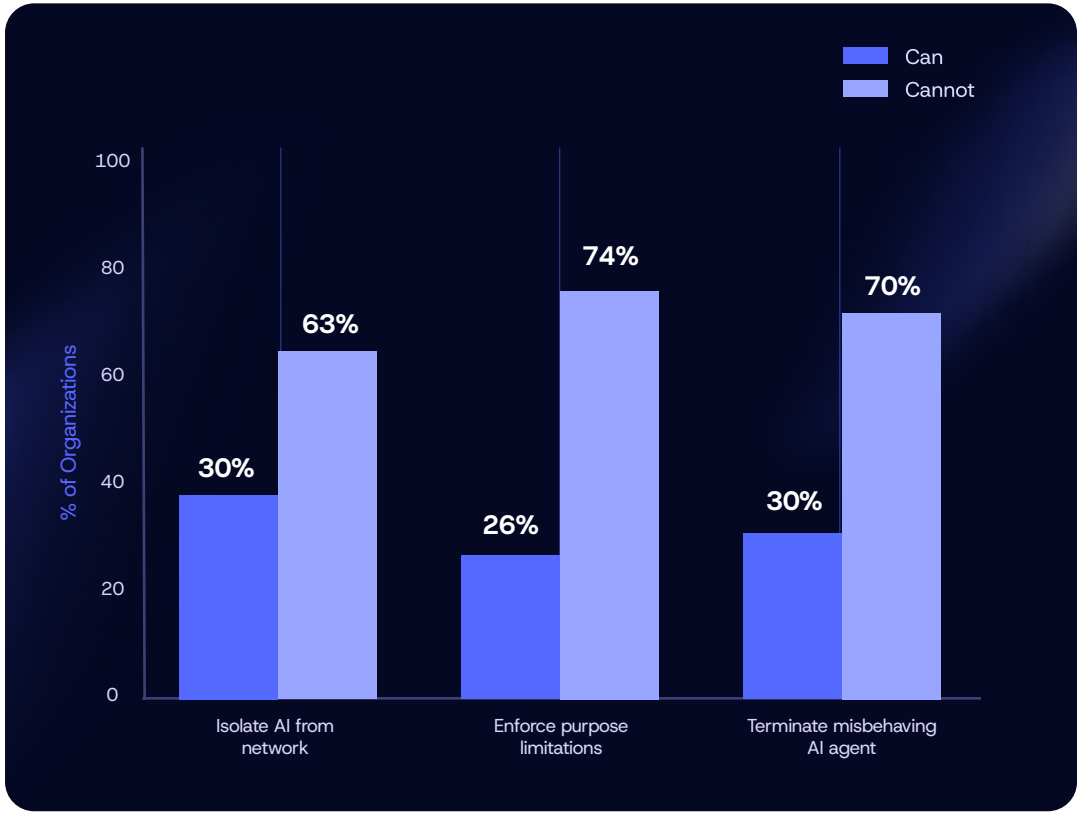

Measurement note: "AI kill switch" as used in this survey captures the formal documented capability to terminate AI agent operations on demand — a policy and architecture requirement distinct from automated session termination, which is measured separately. Organizations may have technical session timeout capabilities without a formal kill switch protocol, or a documented kill switch policy without tested termination automation. Both gaps represent control deficiencies; they represent different control deficiencies. The 30% with a deployed kill switch refers to the formal documented capability.

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.[32]

The channel inventory across the surveyed organizations spans every major category of sensitive data exchange. Cloud-based file sharing (SharePoint, Box, Dropbox, Google Drive) is used by 78%. Email (including encrypted email) is used by 73%. Managed file transfer (MFT) is used by 55%.[33] 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.[34] API-based integrations are used by 48%. Web forms to collect private data are used by 43%.[35] Physical media — USB drives and shipped storage — are used by 29%. AI tools for data processing or exchange are used by 38%.[36] Only 4% operate a single unified platform for sensitive data exchange.

The governance problem starts before channel selection. Only 39% of organizations validate their sensitive data inventory on a defined quarterly-or-more schedule. 44% have included AI systems in their data inventory at all.[37] The 56% who have not inventoried their AI systems cannot govern what those systems access — because they have not established what those systems are or where they sit relative to sensitive data. You cannot classify data you have not found. You cannot enforce controls on channels you have not mapped.

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.[38] Same structural failure, still unresolved.[39]

The approved-channel enforcement figure is the most operationally significant data point in this section. Only 27% of organizations have purpose binding in place — technically restricting AI agents to authorized tasks and data scopes.[40] The remaining 73% operate without this control, meaning agents can access data beyond their intended scope with no architectural enforcement. Classification coverage mirrors this gap: 55% have not achieved automated classification across all major systems, and only 47% have classification that enforces downstream controls rather than applying labels without behavioral effect.[41]

The incident data confirms that channel governance failures are producing real outcomes.[42] 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.[43] 22% experienced data exposure through web forms.[44] The fragmentation consequence follows a clear pattern: Tier 1 organizations average 5+ sensitive data exchange platforms; Tier 4 organizations average approximately 3.[45] The Kiteworks 2025 MFT Survey Report finds that 63% of MFT users have not integrated their MFT infrastructure with SIEM — meaning the highest-volume channel for regulated data exchange operates outside the visibility perimeter for most users.[46] This Annual Survey corroborates that finding: 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. The security posture of each channel is examined in detail on the following pages, drawing on the Kiteworks 2025 MFT Survey Report and 2025 Data Forms Survey Report alongside this Annual Survey's corroborating findings.

KEY FINDING

54% use 4 or more sensitive data exchange platforms. 73% lack purpose binding controls for AI agents.[47] 55% have not achieved automated classification across all systems.[48] 63% of MFT users are not SIEM-integrated (MFT Survey Report).[49,50] This Annual Survey finds only 26% of MFT users stream logs to SIEM in real time.

FIGURE 14

Sensitive Data Exchange Channels in Use

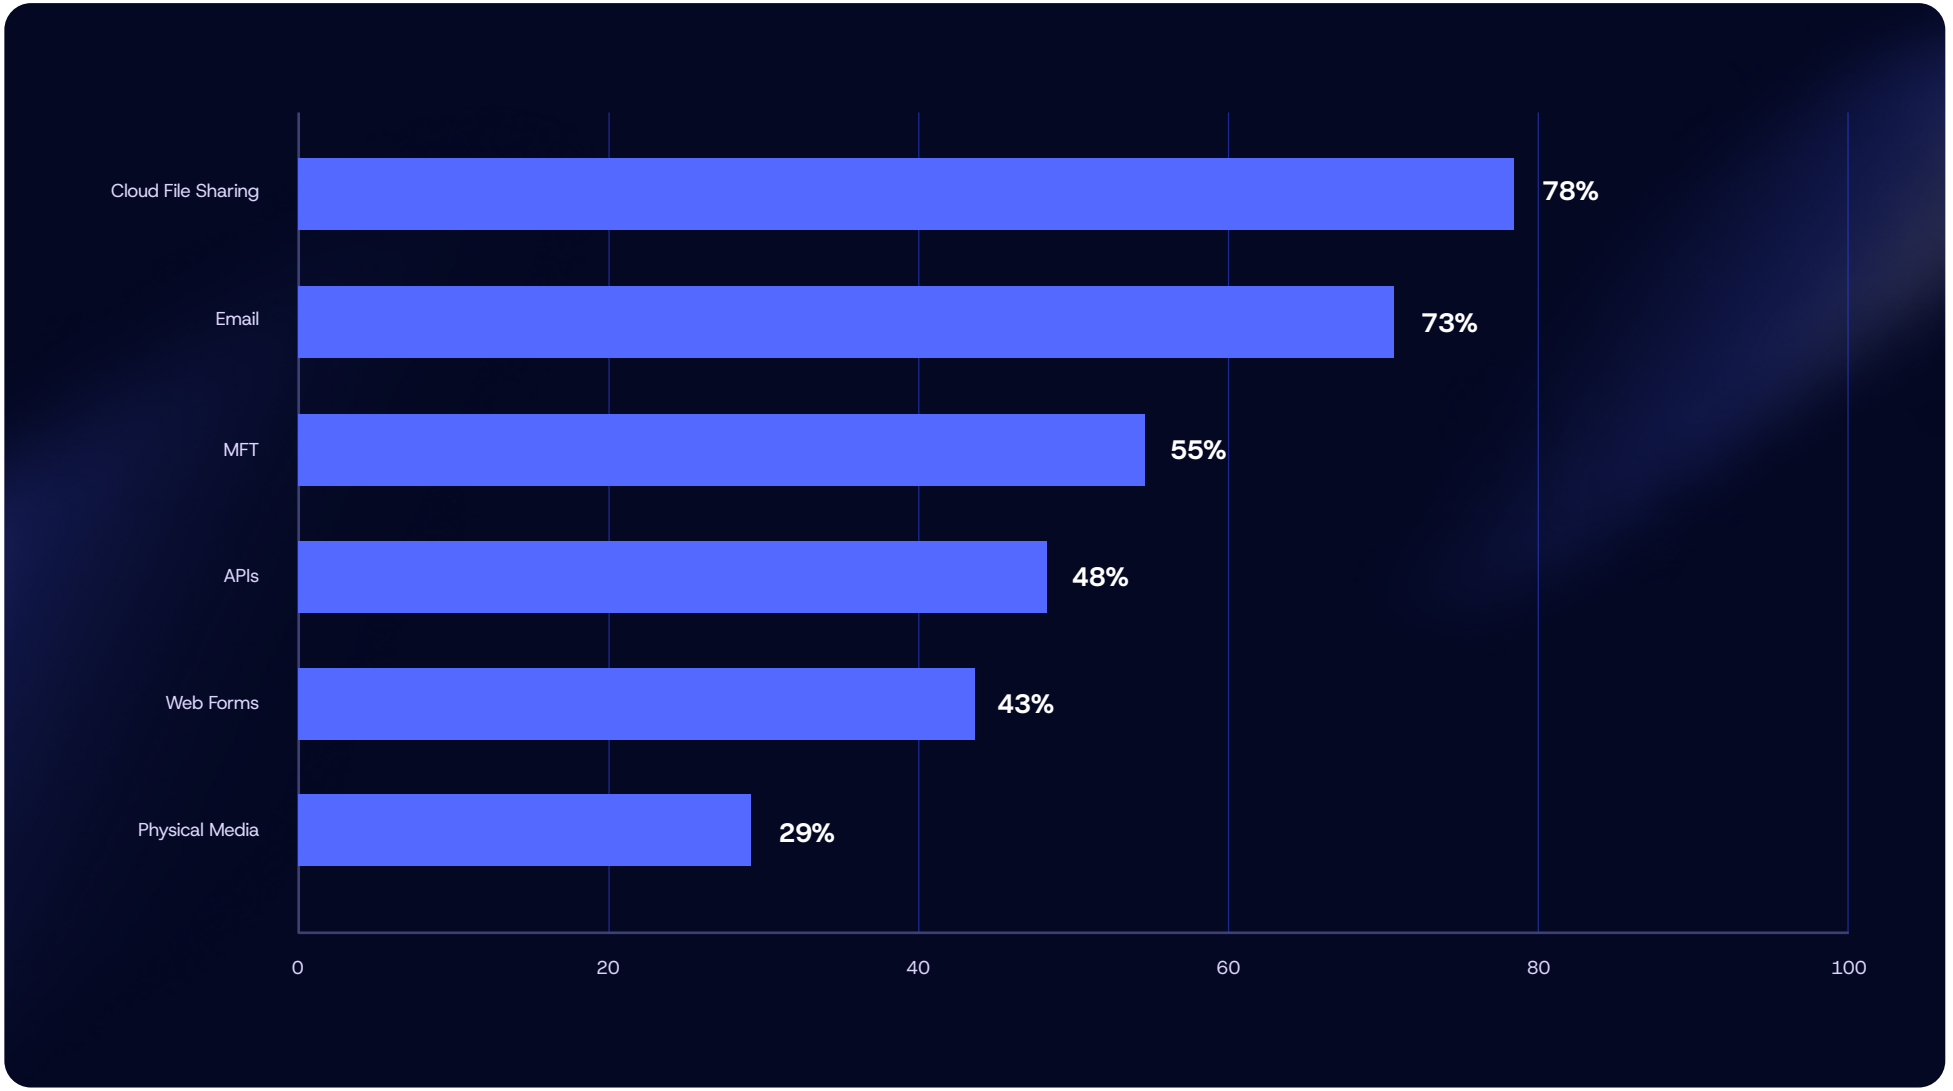

FIGURE 15

MFT Encryption and Security Posture

Source: Kiteworks 2025 Data Security and Compliance Risk: MFT Survey Report.

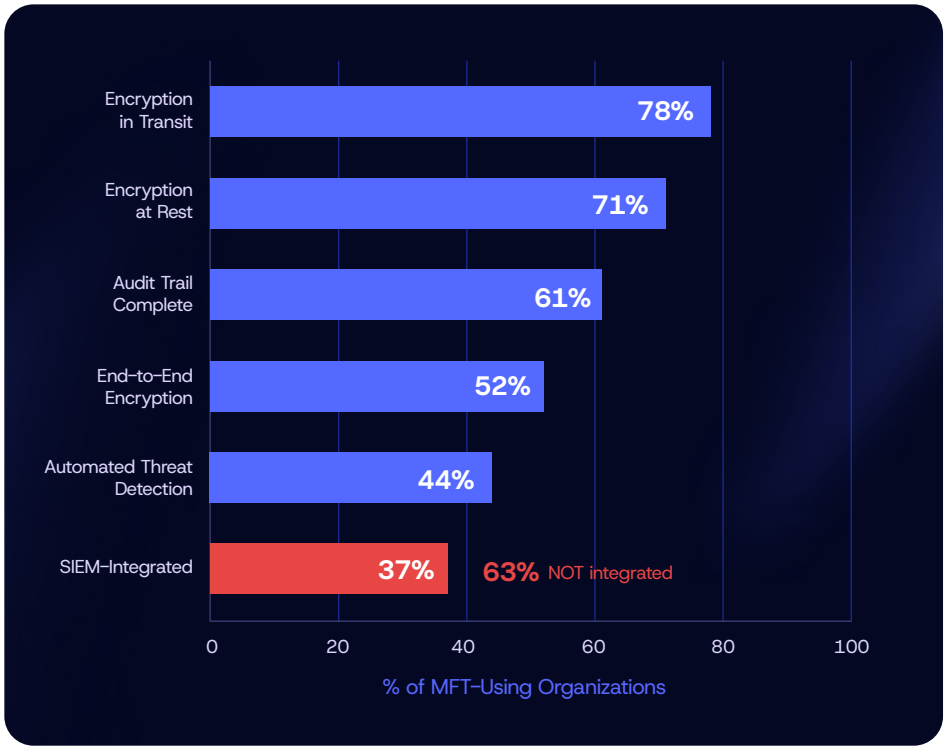

MFT (n=253 Annual Survey MFT users)

- 26% stream MFT logs to SIEM in real time — 74% have no real-time detection coverage on their highest-volume regulated data channel
- Only 39% have tamper-evident audit trails — 61% cannot produce tamper-proof evidence for MFT activity
- 31% experienced data exfiltration via email, file sharing, or MFT in the past 12 months

FIGURE 16

Web Form Data Exposure: Attack Vectors and Incidents

Source: Kiteworks 2025 Data Security and Compliance Risk: Annual Data Forms Survey Report.

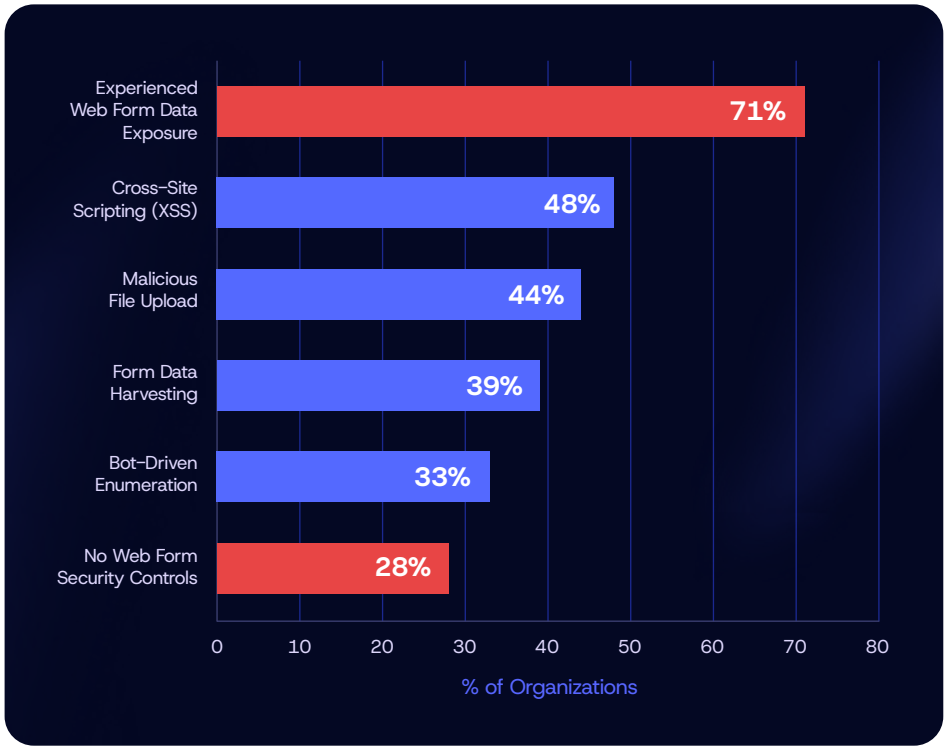

Web Forms (n=195 Annual Survey web form users)

- 27% experienced a PII breach — 13% higher relative rate than the all-respondent average of 24%
- Only 41% have compliance evidence centralized in a GRC platform
- 57% maintain a channel-level visibility inventory — more awareness than average, but still minority coverage

FIGURE 17

Security Tool Fragmentation

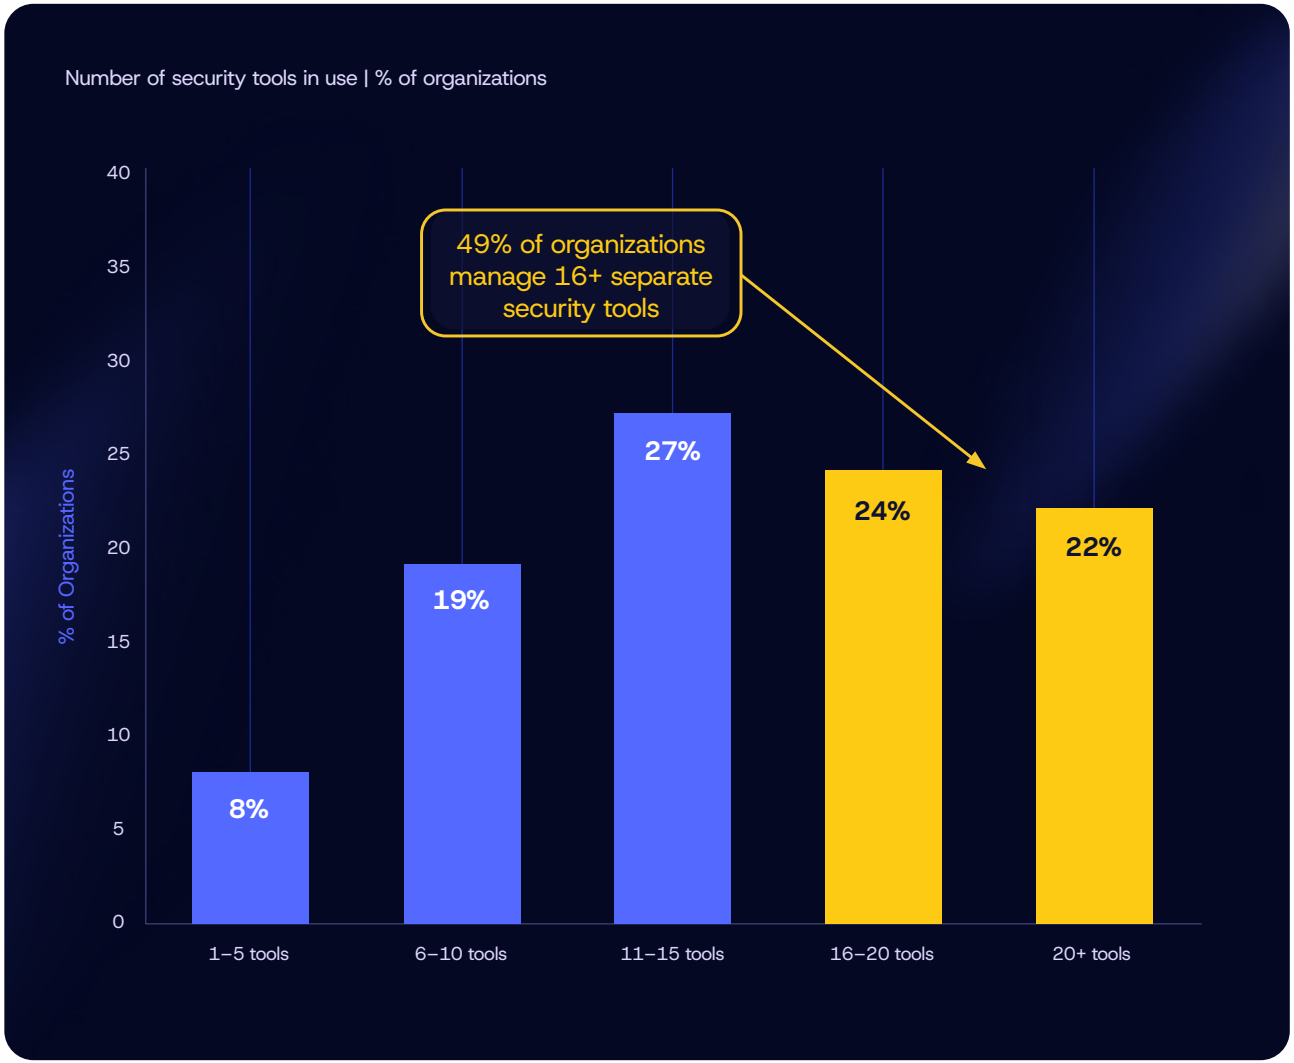

Detection and Response

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 — meaning one in four organizations has no systematic capability to identify when an AI system is accessing data it should not be accessing.[51]

One in four organizations has deployed AI systems into environments handling sensitive data and made no formal arrangement to notice when those systems access something they should not. The incidents are finding them anyway.[52]

Coverage inconsistency compounds the detection problem. 57% of organizations do not have consistent AI log coverage across all systems — they have gaps they know about, or no mapping at all.[53] 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%.[54]

One data quality anomaly in this section warrants disclosure. The survey finds 33% of organizations forward AI access logs to a SIEM or centralized platform, while 29% generate any AI system data access logs at all. Because forwarding to a SIEM presupposes log generation, these figures are logically inconsistent — an artifact likely produced by question framing: Respondents may have interpreted the SIEM integration question as asking about organizational intent or policy rather than actual system behavior. Both figures should be treated as upper-bound estimates of the underlying capability. The operationally important conclusion stands regardless of which figure is more accurate: The majority of organizations have neither capability in place.

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. DSMS measures the foundation. AIGMS measures what the foundation cannot see.

The DSMS tier differential on AI incidents is the starkest finding in this section. The AI incident rate drops meaningfully with security maturity tier — confirming that the controls the DSMS measures are the same controls that distinguish organizations that detect and contain AI incidents from those that do not. The difference is not luck or sector. It is the presence or absence of the specific detection and containment controls the DSMS measures. Organizations in Tier 1 are not experiencing fewer incidents because AI represents less risk to them. They are documenting fewer incidents because they have less capability to see them.

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. Third-party breach data makes the vendor governance risk concrete. 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.[55,56]

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.

FIGURE 18

The Detection-Response Gap

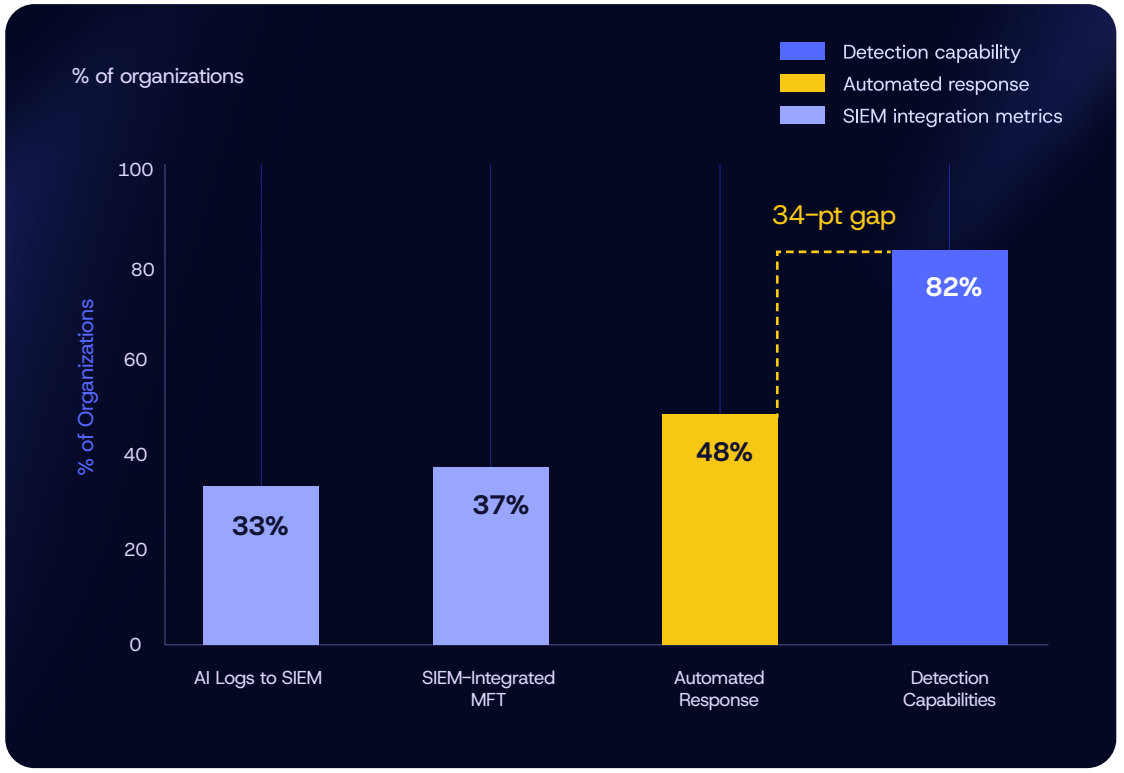

FIGURE 19

AI Data Access Log Coverage and Detection Capabilities

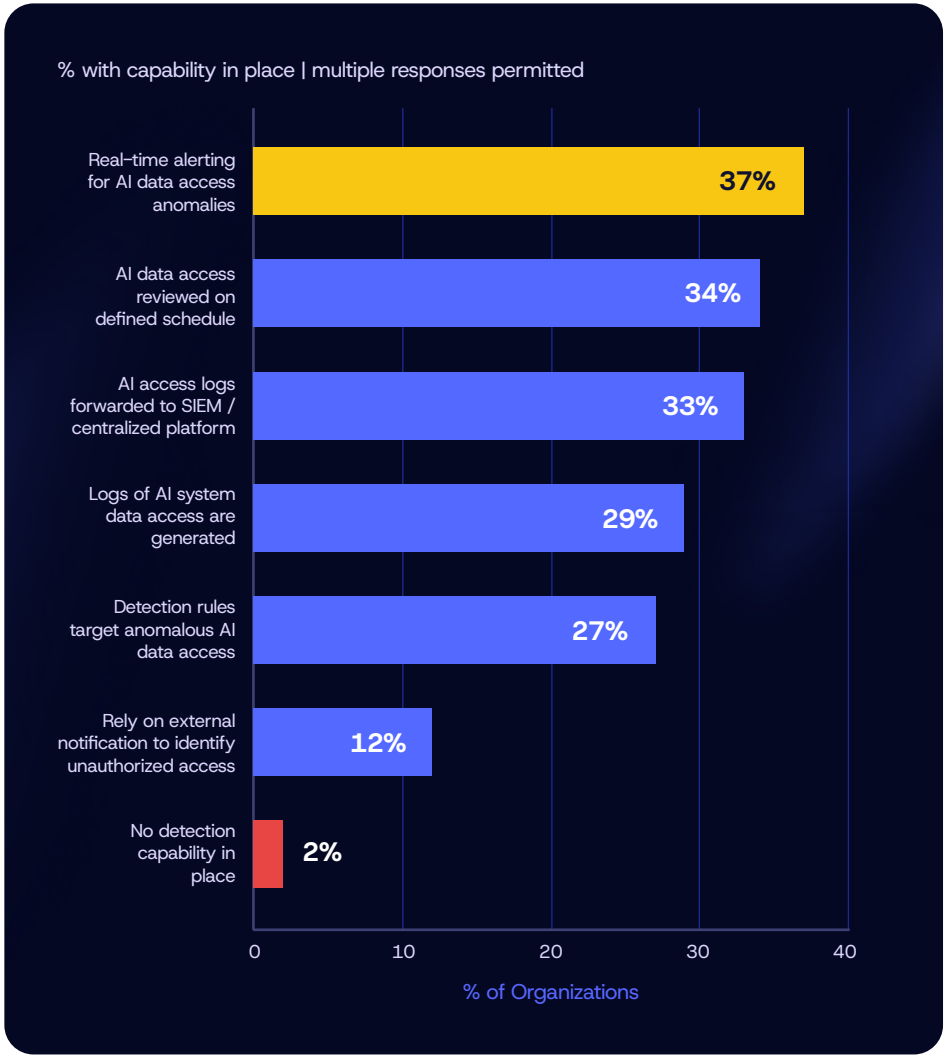

FIGURE 20

Greatest Perceived Security Risk (Next 12 Months)

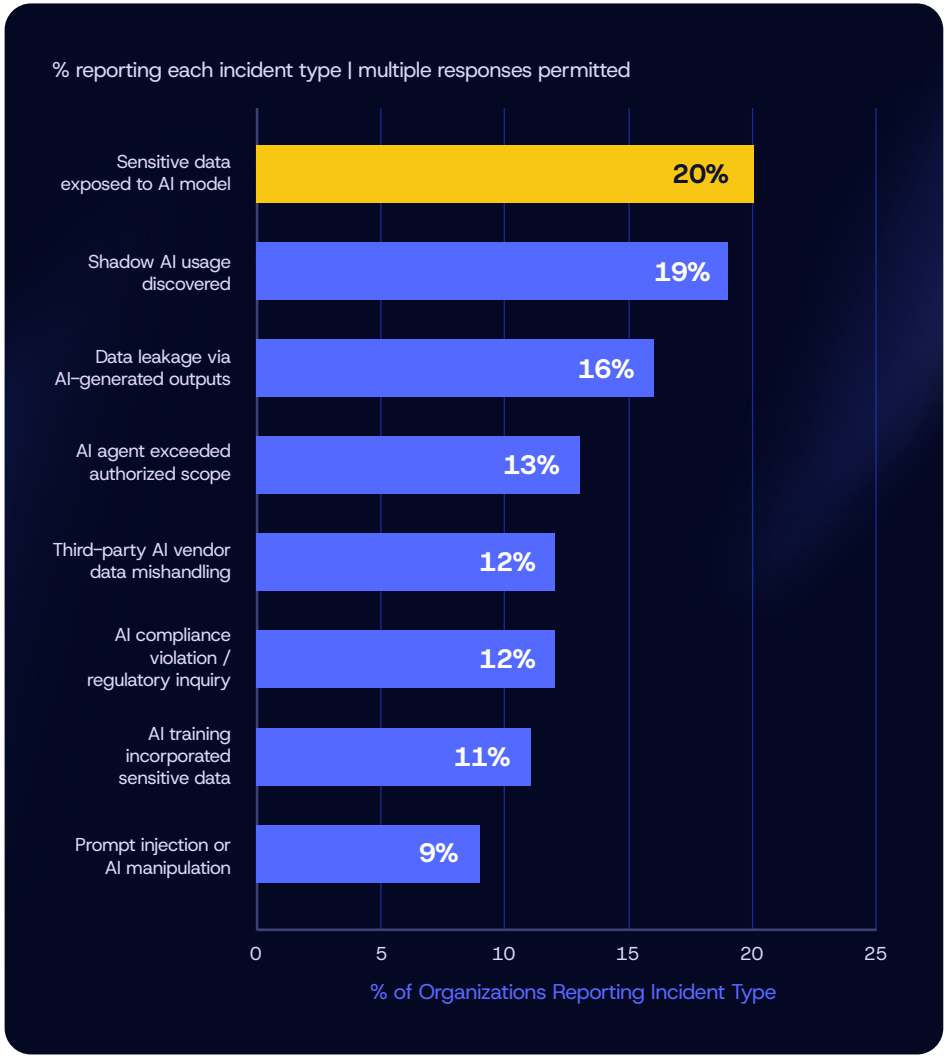

FIGURE 21

Greatest Perceived Security Risk (Next 12 Months)

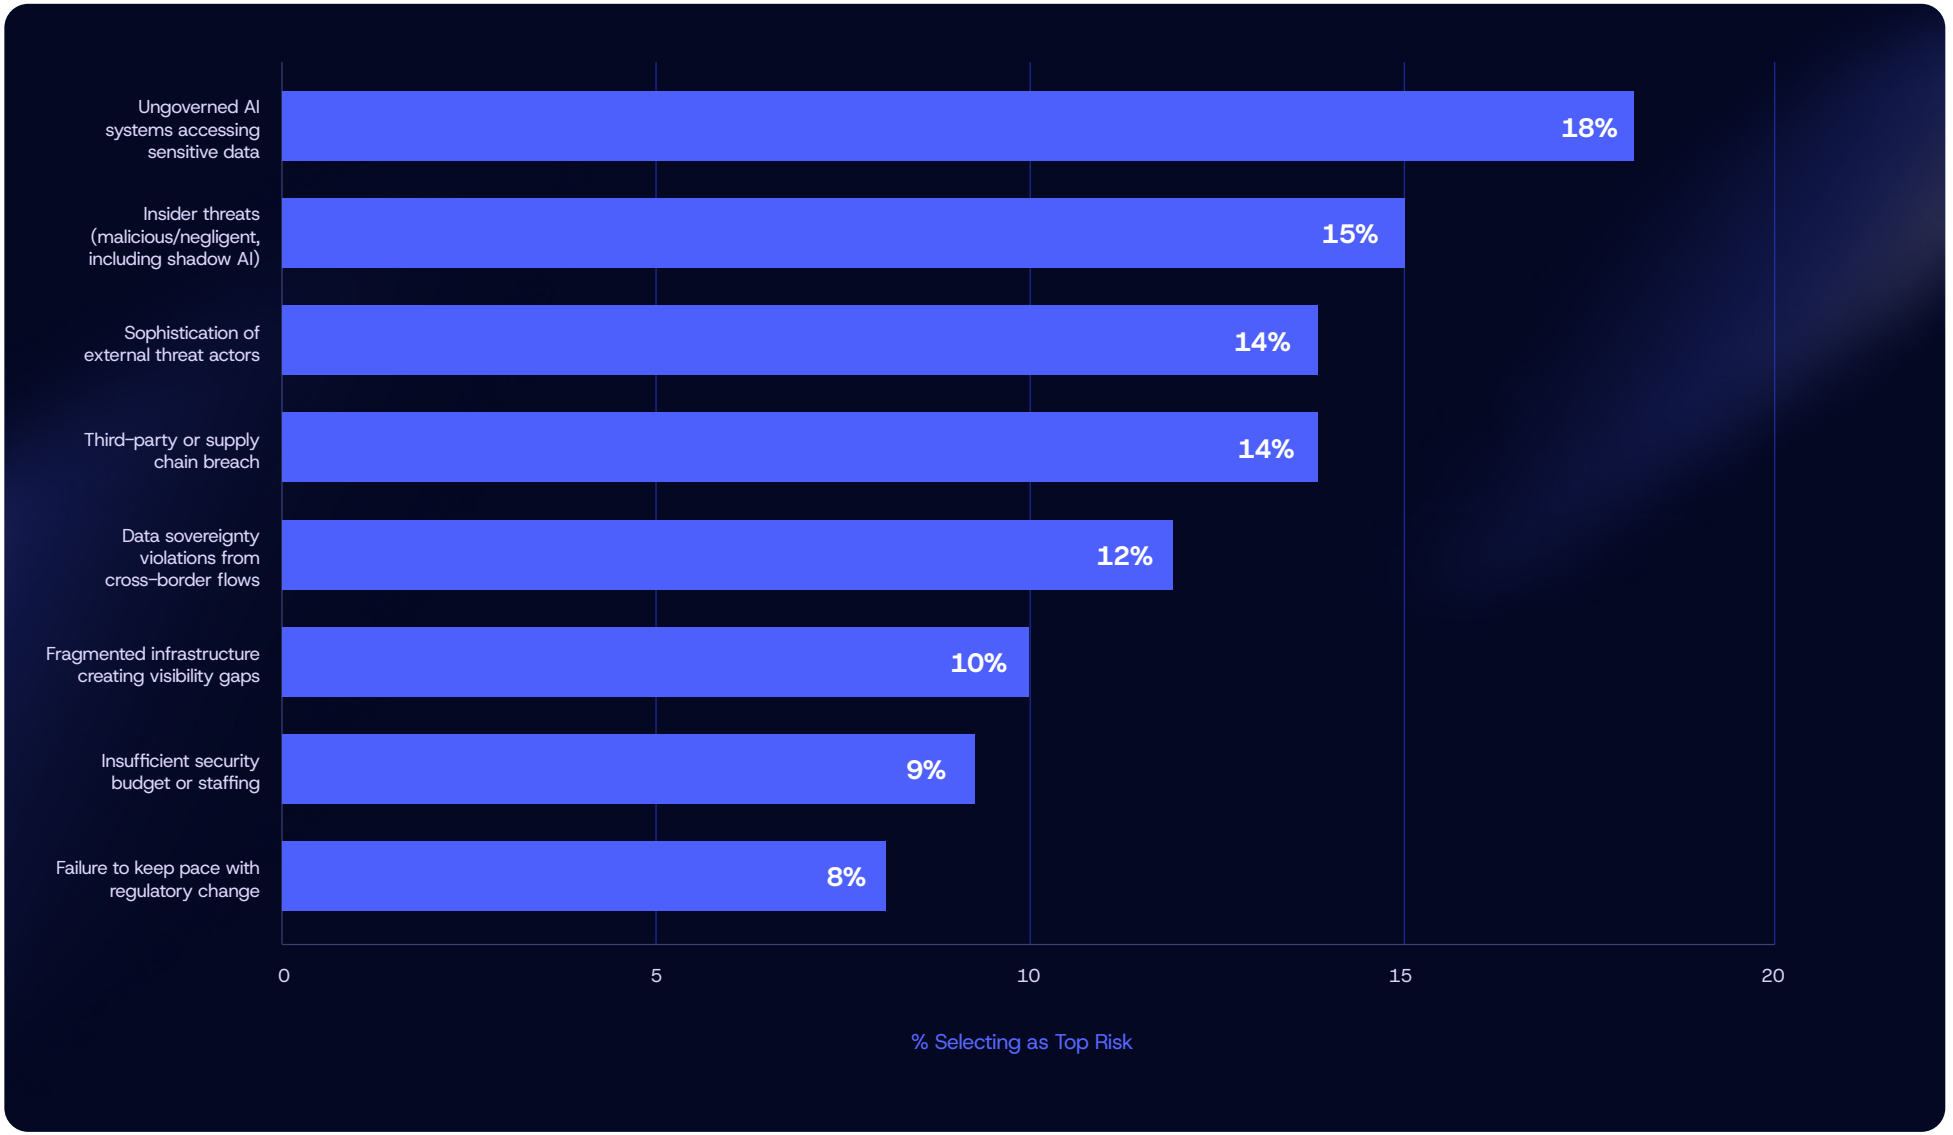

Compliance and Regulatory Readiness

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.[57]

The audit readiness data is the most operationally specific finding in this section. 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.[58] This is not a capability that can be built reactively. An organization that learns on a Thursday that a regulator wants an AI access audit record by Friday does not have time to build the infrastructure that week.[59]

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. The Forecast report projected DSPM would become the default baseline by end of 2026 but enforcement would lag. The data confirms the lag: 39% can enforce tagging and classification consistently across channels.

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.[60] The 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%.[61]

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.

Evidence consistency across compliance frameworks is a structural failure in most organizations. Only 40% apply a consistent evidence approach across all required frameworks.[62] 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.[63] 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. The 67% without them are exposed on any audit that requires evidence integrity.[64]

EU AI Act Obligations Are Already Active

ISO/IEC 42001 compliance is required for 40% of respondents. DORA operational resilience obligations for financial entities require audit-quality evidence on short notice. The EU AI Act imposes transparency, logging, and human oversight obligations that are enforceable now for high-risk system categories. Organizations that cannot produce an AI data access record within one business day are not compliant with these frameworks' audit obligations — regardless of what their policy documents say.[65]

FIGURE 22

AI Audit Record Producibility

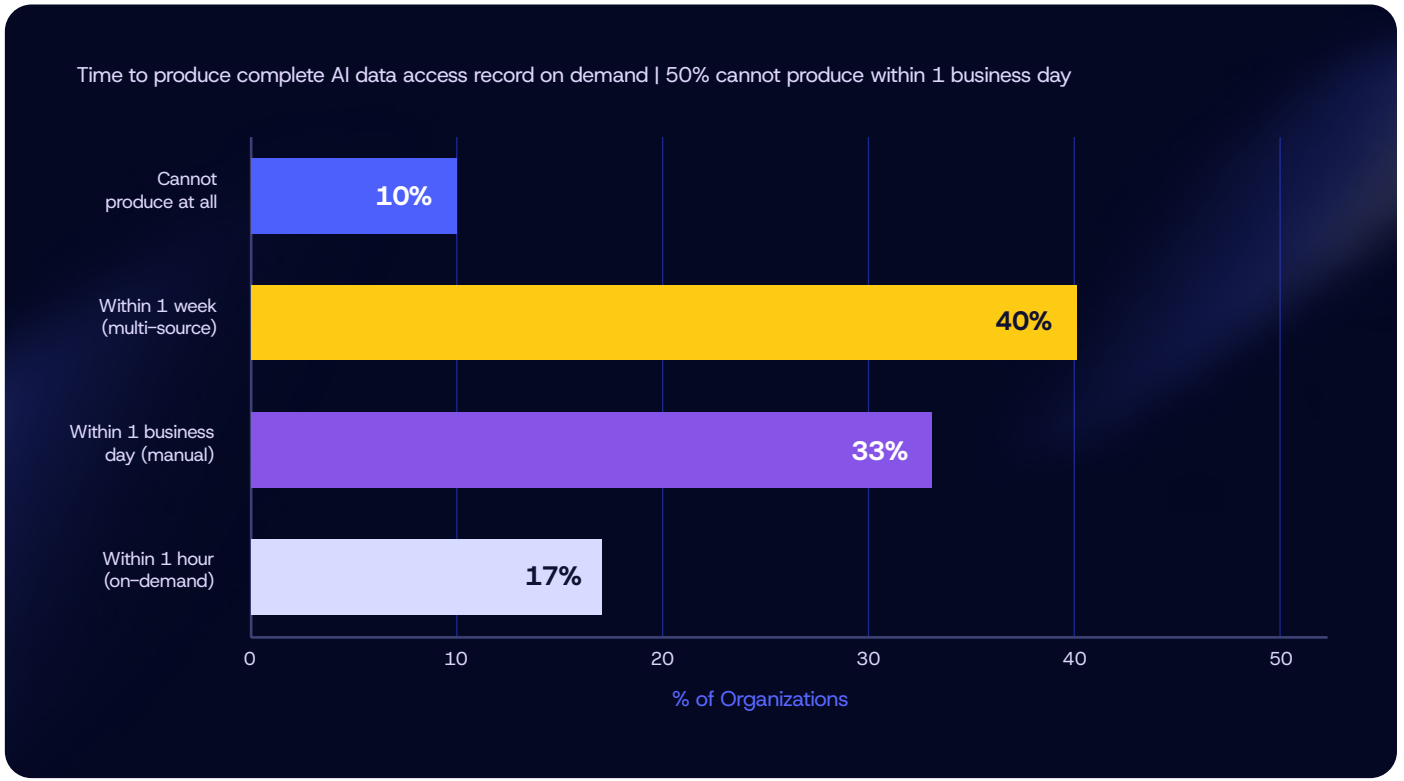

FIGURE 23

Top Compliance Challenges

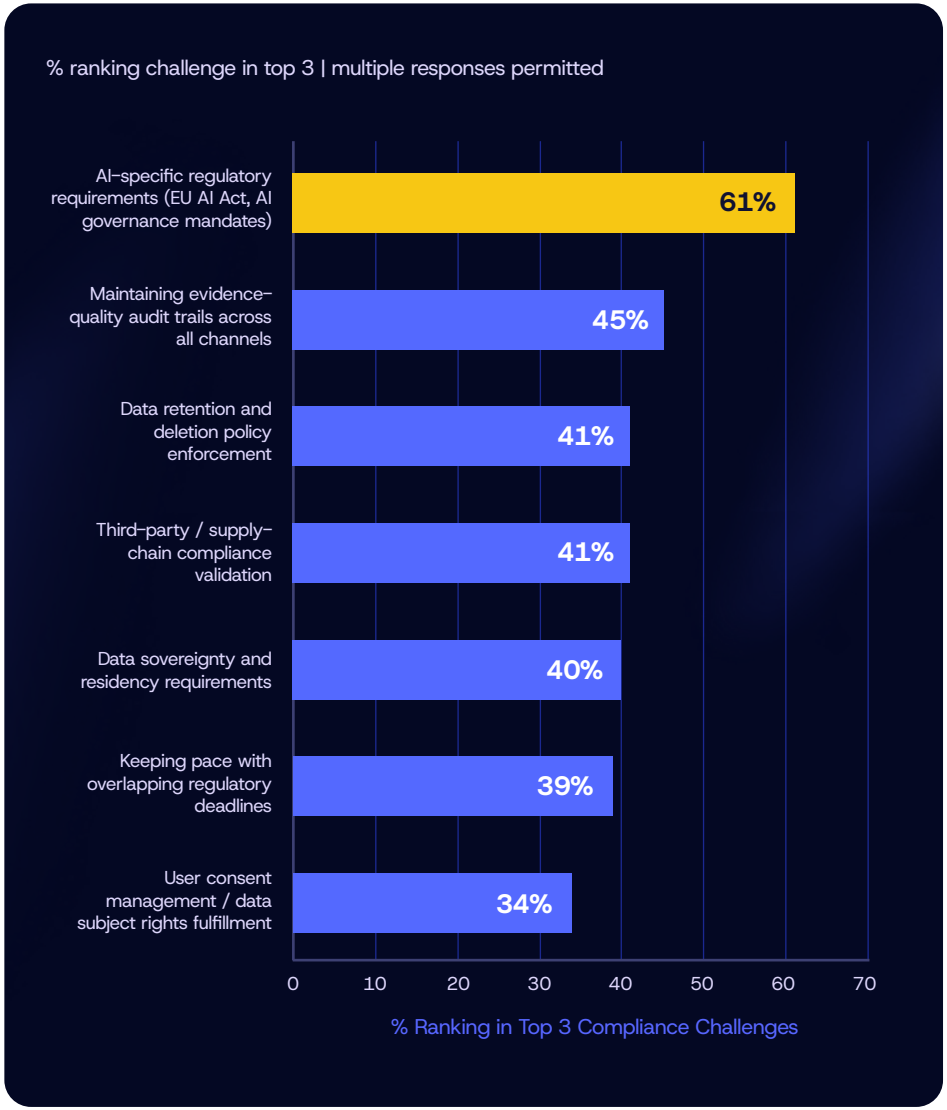

FIGURE 24

Compliance Frameworks Currently in Skope

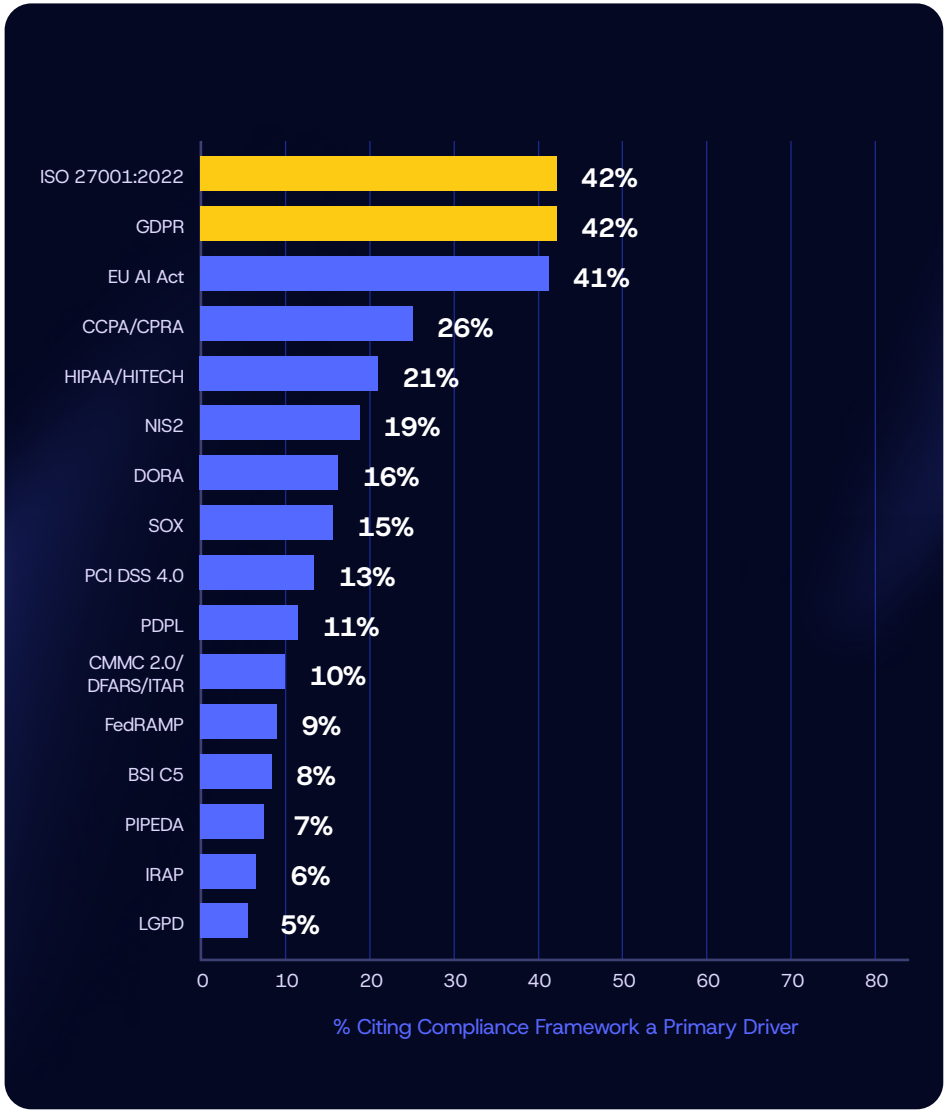

FIGURE 25

Compliance Consequences Experienced (Past 12 Months)

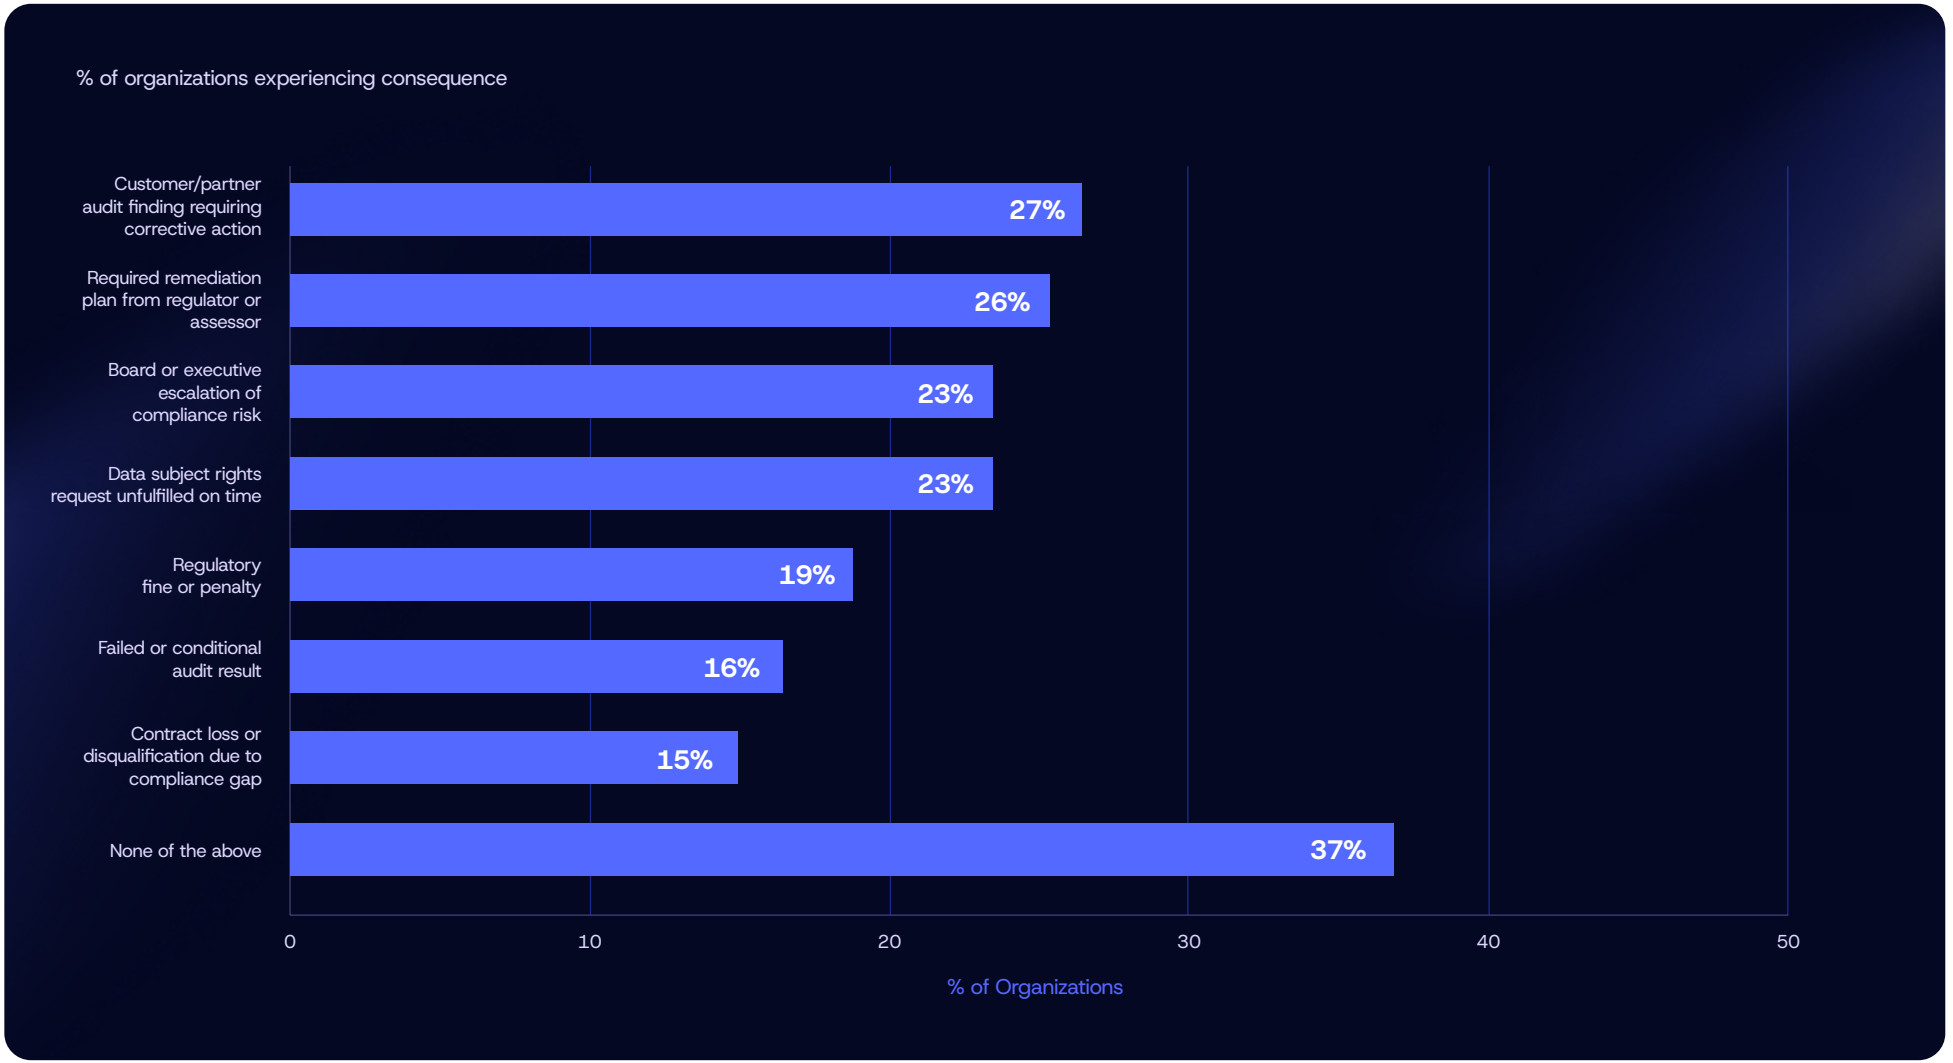

The Human Element: Shadow AI

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.[66]

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 35% who report no discovery are not necessarily operating clean environments — they are more likely operating environments with insufficient detection capability to see what is happening.[67] 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 unapproved AI channels. Average annual insider risk cost has reached $19.5 million per organization.[68] The pattern is consistent across regions: North America reports the highest shadow AI discovery rate (57% combined yes), followed by MEA (59% combined) and EU/UK (34% combined).

The data flowing through those 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.[69]

Those are not categories of data that organizations intend to share with AI systems. They are categories of data that employees are sharing with AI tools because the tools are useful and the governance infrastructure to prevent it — AI-specific DLP, approved-channel enforcement, purpose binding — is deployed in fewer than 30% of organizations. The data is moving. The controls are not. 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.[70]

The 35% who report no shadow AI discovery are not necessarily operating shadow-AI-free environments. Among organizations reporting no discovery, a significant share carry the detection profile that would prevent discovery even if shadow AI were present: no real-time alerting for AI data access anomalies, no AI system data access logs, and no scheduled review of AI access activity. Detection gaps look identical to clean environments from the outside. The regional pattern reinforces this reading: North America — which carries the most mature detection infrastructure of any region — also reports the highest shadow AI discovery rates (57% combined yes). EU/UK organizations, with

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.

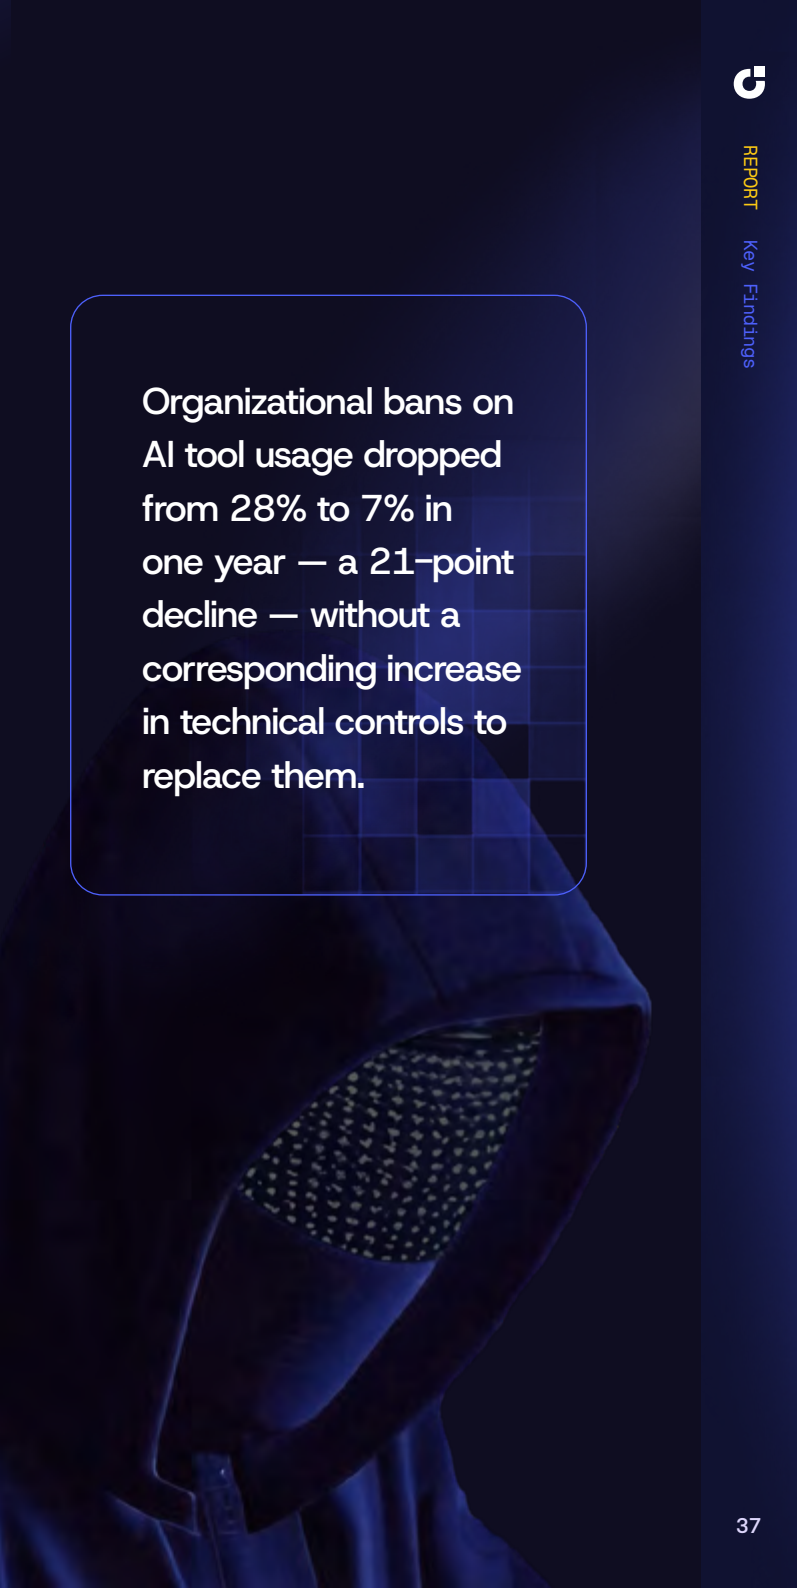

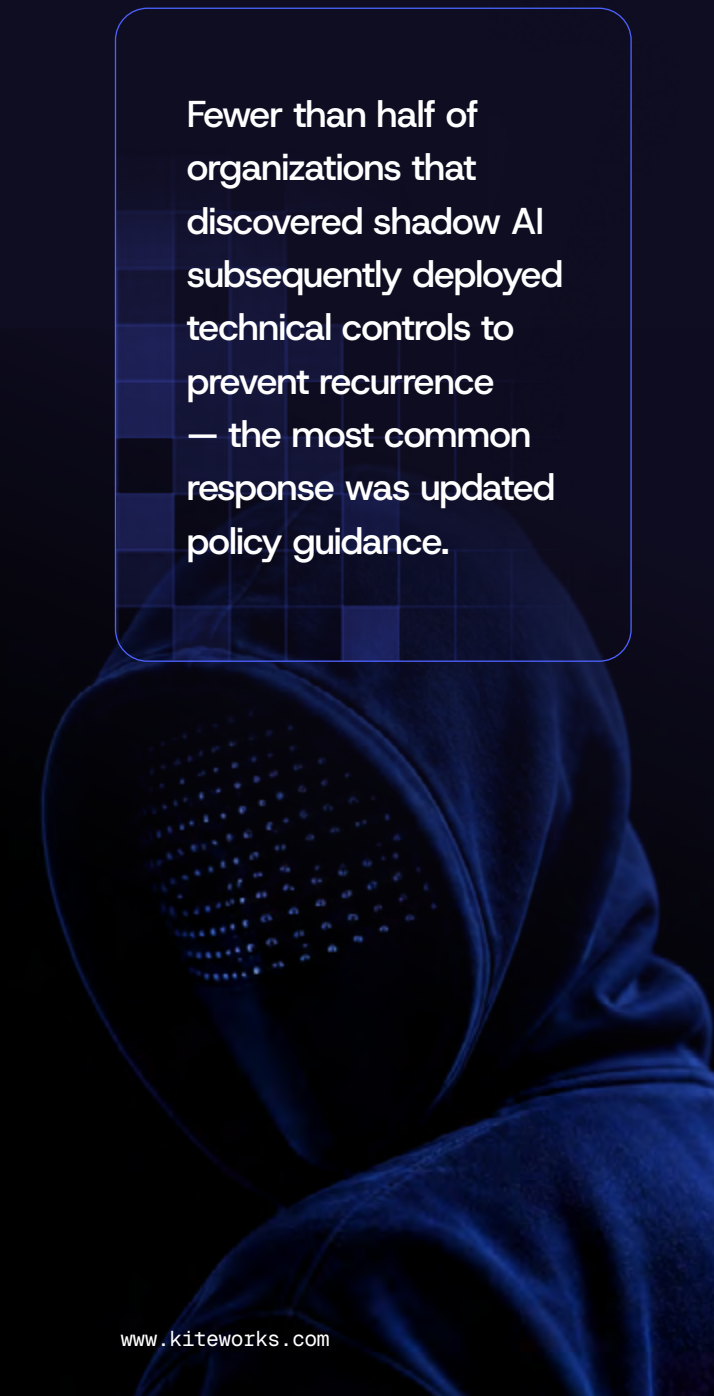

Fewer than half of organizations that discovered shadow AI subsequently deployed technical controls to prevent recurrence — the most common response was updated policy guidance.

stronger regulatory pressure and audit requirements, report 46% combined discovery. MEA, with the lowest DSMS scores and least mature detection infrastructure, reports the lowest discovery rates — not because shadow AI is less prevalent, but because the infrastructure to find it is least developed. Discovery rate is a function of detection capability, not shadow AI prevalence.

The control gap behind shadow AI is architectural. Only 28% of organizations have AI-specific DLP — the technical control that can block sensitive data from reaching unapproved AI tools at the transmission layer, not just document the policy violation after the fact. The Cloud Security Alliance's 2026 research on identity and access gaps for autonomous AI found that organizations without AI-specific DLP relied on general-purpose network controls that failed to detect the majority of AI-bound sensitive data transmissions in test environments. The 72% without AI-specific DLP are not operating with reduced protection — they are operating with tools that were not designed for the data flow patterns AI tools generate. Only 27% have purpose binding deployed — restricting AI agents to authorized tasks and preventing data from reaching systems outside their defined scope. Third-party AI vendor risk is governed even more weakly: 27% either have not evaluated whether their AI vendors use organizational data for model training, or rely solely on paper attestations with no technical verification.[71] The same Cisco study found that AI usage in enterprises more than doubled in 2025, with 62% of workers now using AI at work — but organizational governance has not kept pace.[72]

The investment signal is partially aligned with the architecture problem. 36% cite workforce training on AI data security and compliance as a top investment priority. But training alone cannot close a structural architecture gap. An employee who knows the policy but whose data can still reach a consumer AI tool because there is no technical enforcement layer is still a governance failure. The organizations that have solved shadow AI have done it with technical controls, not behavioral change programs.

Discovery without remediation is documentation of a problem, not resolution of one. The survey finds that among organizations that have discovered shadow AI usage, the most common organizational response is issuing updated policy guidance — a behavioral intervention applied to an architectural problem. Fewer than half of organizations that discovered shadow AI subsequently deployed technical controls to prevent recurrence at the transmission layer. The pattern is consistent with the investment data: 36% cite workforce training on AI data security as a top investment priority, making training the most commonly planned response to a problem that training cannot structurally solve. An employee who has been trained on the policy but whose data can still reach a consumer AI tool through an unblocked channel has learned about the violation they are still able to commit. The organizations that have reduced shadow AI incidence in the data 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.

FIGURE 26

AI-Related Security Incidents Experienced (Past 12 Months)

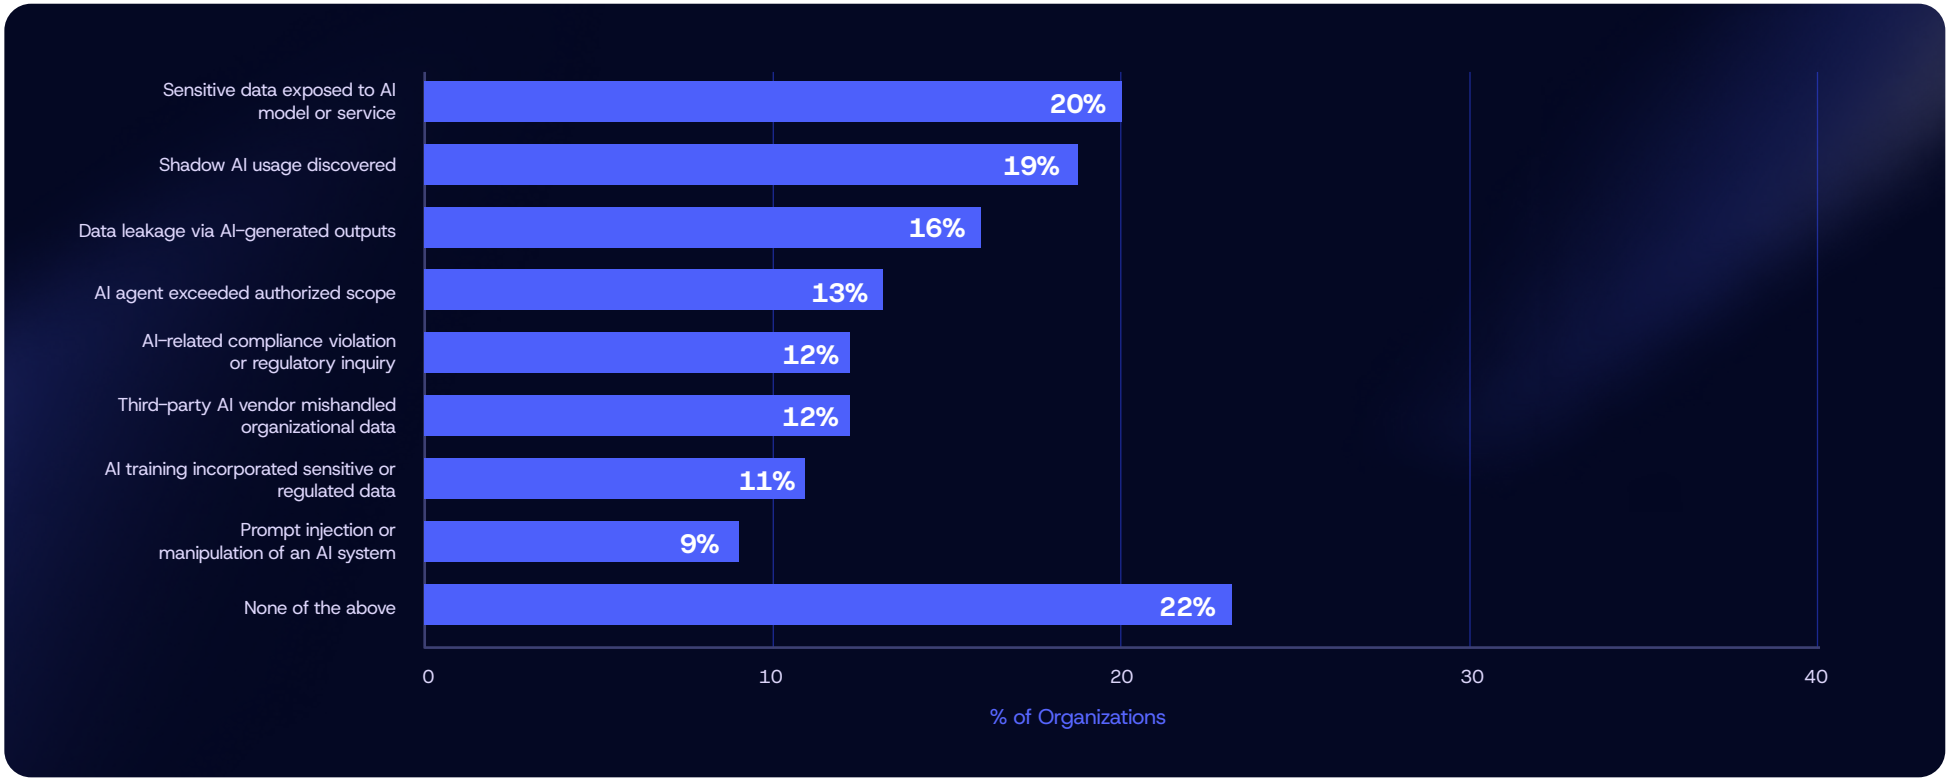

FIGURE 27

Data Security Incidents Experienced (Past 12 Months)

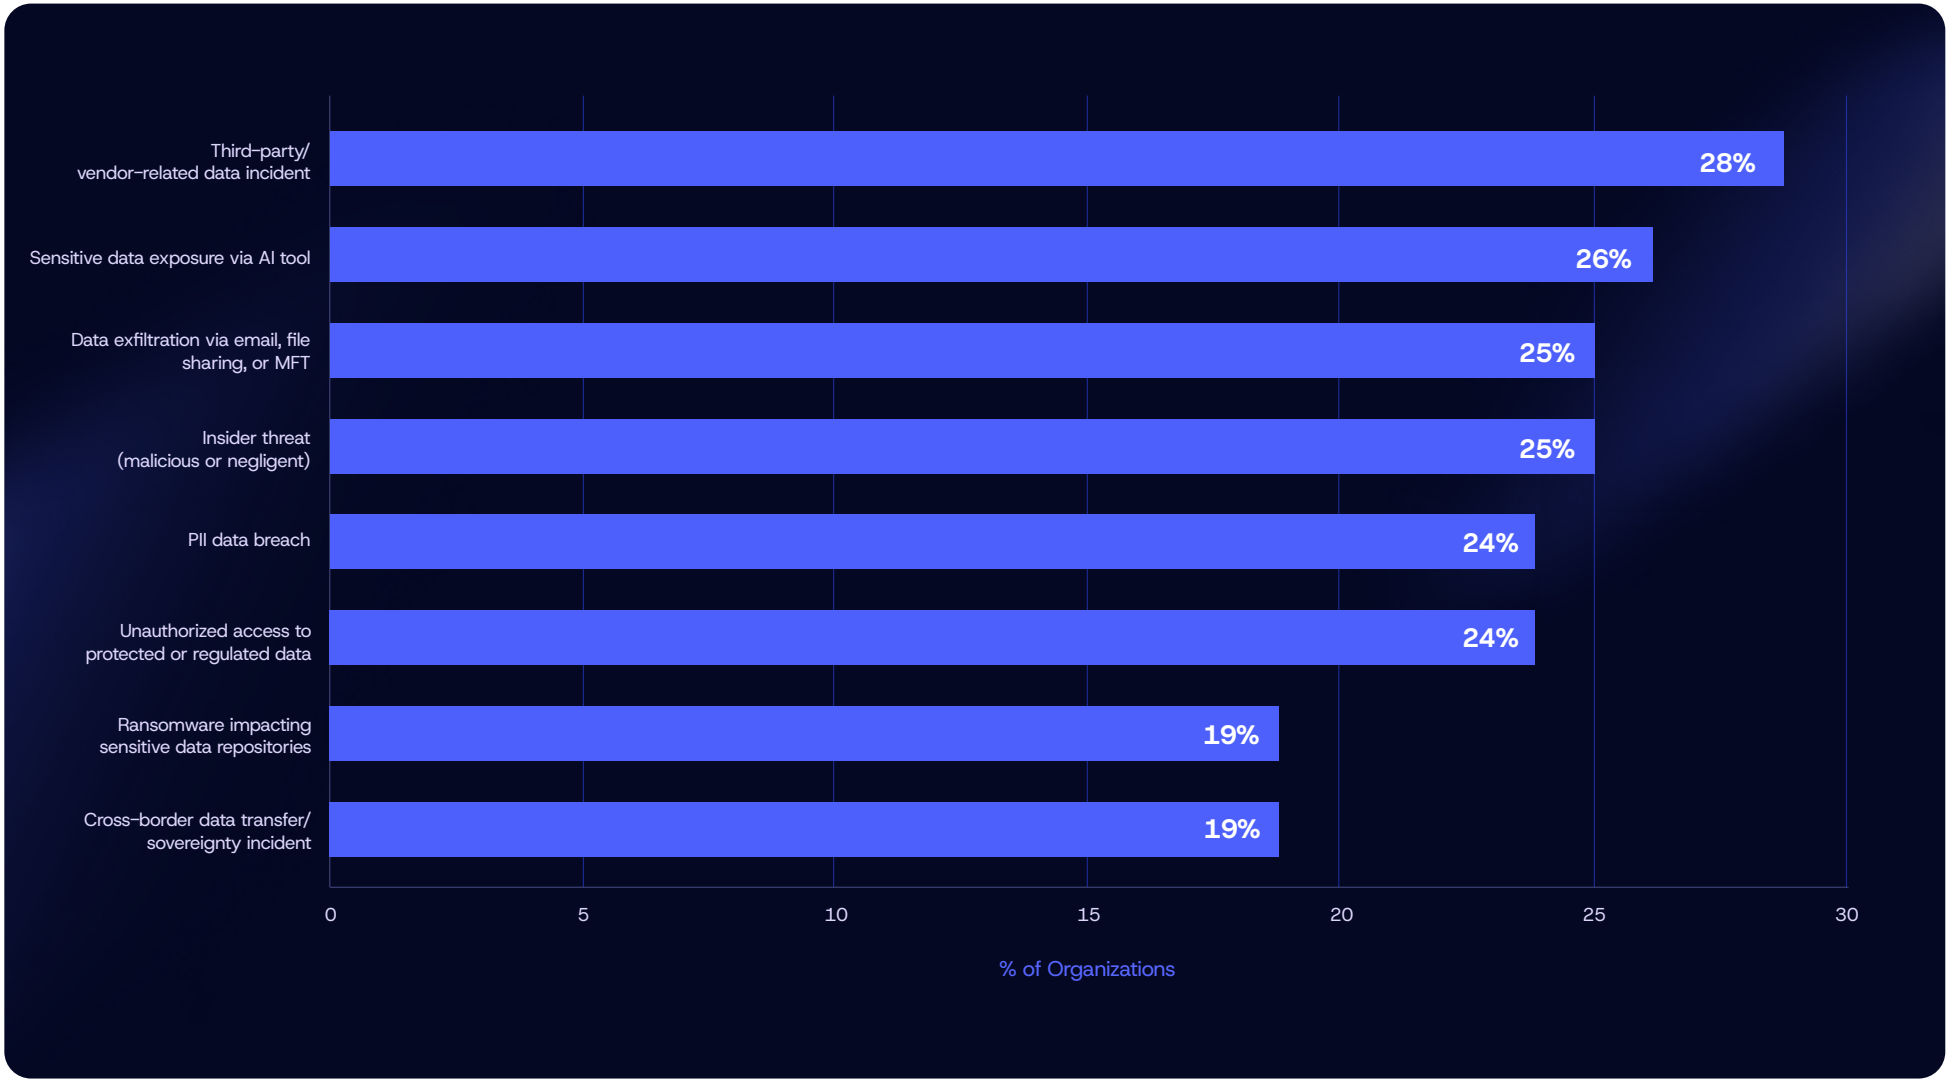

Cross-Analyses: Variations by Key Factors

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 that these averages conceal substantial variation — variation that has direct implications for how organizations in each cohort should prioritize their remediation investments.

Industry Variations

Financial Services and Energy and Utilities lead the DSMS distribution with mean indices of 43.3 and 43.2, respectively. The driver in Financial Services is regulatory mandate: DORA and model risk management requirements have operationalized control deployment in ways that voluntary frameworks have not.[73] Energy and Utilities organizations face similar operational technology security requirements that translate directly into DSMS components.[74] Manufacturing follows at 43.0, Healthcare at 42.3, and Technology at 40.7.[75]

Defense Contractors records the lowest DSMS of any sector at 15.2, though its sample size of three respondents limits generalizability. Federal Government is the lowest sector with a meaningful sample size at 22.7 — more than 16 points below the survey mean of 39. The pattern is direct: executive AI policy at the national level has not translated to operational control deployment at the organizational level. The gap between stated AI data governance policy and deployed security controls is wider in Federal Government than in any other sector with a statistically meaningful sample.

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 (22.7 — more than 16 points below the survey mean of 39.1) and the Defense Contractors records the lowest AIGMS of any sector (22.8) and Federal Government the second lowest (25.4) — not tied, but both well below the survey AIGMS mean of 34.7. The Forecast report described the problem with high confidence. Six months later, the Annual Survey found it unchanged.

Defense Contractors: Mandatory Framework, Minimum Maturity

Defense Contractors records the lowest DSMS Index of any sector surveyed (15) — 16 points below the survey mean and 28 points below the Financial Services sector. It ties Federal Government for the lowest AIGMS (25). These scores exist within a sector operating under CMMC — a mandatory compliance framework with contractual enforcement that Federal Government's policy-level mandates do not replicate.

The gap between regulatory mandate and operational control deployment is widest in the sector where the mandate is most legally enforceable. 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 audit findings.

The DSCRI for the Defense Contractors cohort — DSMS (15) × AIGMS (25) / 100 = 3.75 — is the lowest composite risk index of any sector. An organization at DSCRI 3.75 has not achieved the security foundation required to govern AI, has not operationalized AI governance, and is operating under a compliance framework whose audit timeline it cannot currently meet. The path from DSCRI 3.75 to the survey mean of 16.2 requires foundational remediation — classification, SIEM integration, complete audit trails — before AI-specific governance can be meaningfully layered on top.

On AI governance maturity, the same two sectors rank lowest. Defense Contractors records the lowest AIGMS of any sector at 22.8 — nearly 12 points below the survey mean of 34.7. Federal Government is the second lowest at 25.4, approximately 9 points below the mean. Both are separated from every other sector by a substantial margin; the next lowest, Education, sits at 26.6.

Technology presents a risk profile that its aggregate scores obscure. Its 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. Technology organizations have deployed technical security controls but have not built the institutional governance structures that AI-specific risk requires. The infrastructure exists; the governance perimeter around AI systems does not.

Professional Services (33.2), Legal (27.4), and Education (25.1) fall below the survey mean, joining Federal Government and Defense Contractors in the lower half of the DSMS distribution — sectors where neither regulatory mandate nor operational necessity has driven security control deployment to mean levels.

The DSCRI (= DSMS × (AIGMS/100)) makes the cross-industry risk differential visible in a single number. Energy and Utilities leads on DSCRI at 19.6, the product of the second-highest DSMS in the survey (43.2) and the highest AIGMS of any sector (39.2). Financial Services is second at 18.5, driven by a DSMS of 43.3. Technology sits just above the survey mean DSCRI of 16.2 (Technology DSCRI 16.8, DSMS 40.7, AIGMS 35.3), with neither dimension distinguishing it meaningfully from the average organization. Federal Government (DSCRI 7.0) and Defense Contractors (DSCRI 4.8) sit at the bottom of the distribution. Federal Government, Defense Contractors, and Education are the only sectors where both DSMS and AIGMS fall below 27 — placing the majority of their respondents in the lower-left laggard quadrant.

Representation Analysis: Which Industries Lead and Which Lag

Aggregate industry means tell part of the story. To understand which industries are producing high-maturity organizations at rates above or below their share of the survey, the data require a representation index: each industry's rate of appearing in the upper-right leadership quadrant (DSMS ≥ 45 and AIGMS ≥ 45 simultaneously) divided by the overall rate of 21.8%. An index above 1.25 indicates meaningful overrepresentation; below 0.75 indicates meaningful underrepresentation.

Healthcare produces leaders at 1.57 times the expected rate — the highest of any sector. The finding is counterintuitive relative to Healthcare's reputation as a security laggard. What the data shows is that Healthcare organizations that have moved past reactive compliance posture into proactive governance architecture have done so more consistently than most sectors. They appear in the lower-left laggard quadrant at 0.81 times the expected rate.

Energy and Utilities follows at 1.38 times — meaningfully above proportional — reflecting OT security requirements that translate into DSMS components and, distinctively, produce above-average AI governance maturity as well.

Technology, despite its reputation as a technically sophisticated sector, produces leaders at only 1.04 times the expected rate — barely above proportional. Its laggard representation is 0.99 times the expected rate, effectively proportional. The sector is not producing leaders at scale, and it is not systematically avoiding the laggard quadrant either.

The most stark 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. Education appears in the laggard quadrant at 1.34 times the expected rate with a leadership index of 0.23 — the second lowest of any sector with a meaningful sample. Federal Government's laggard representation index is 1.70, the highest of any sector with a meaningful sample size, and its leadership index is 0.25. For Defense Contractors (n=3), zero respondents appear in the leadership quadrant, and the laggard representation index is 1.27.

FIGURE 28

Mean Data Security Maturity Score (DSMS) by Industry

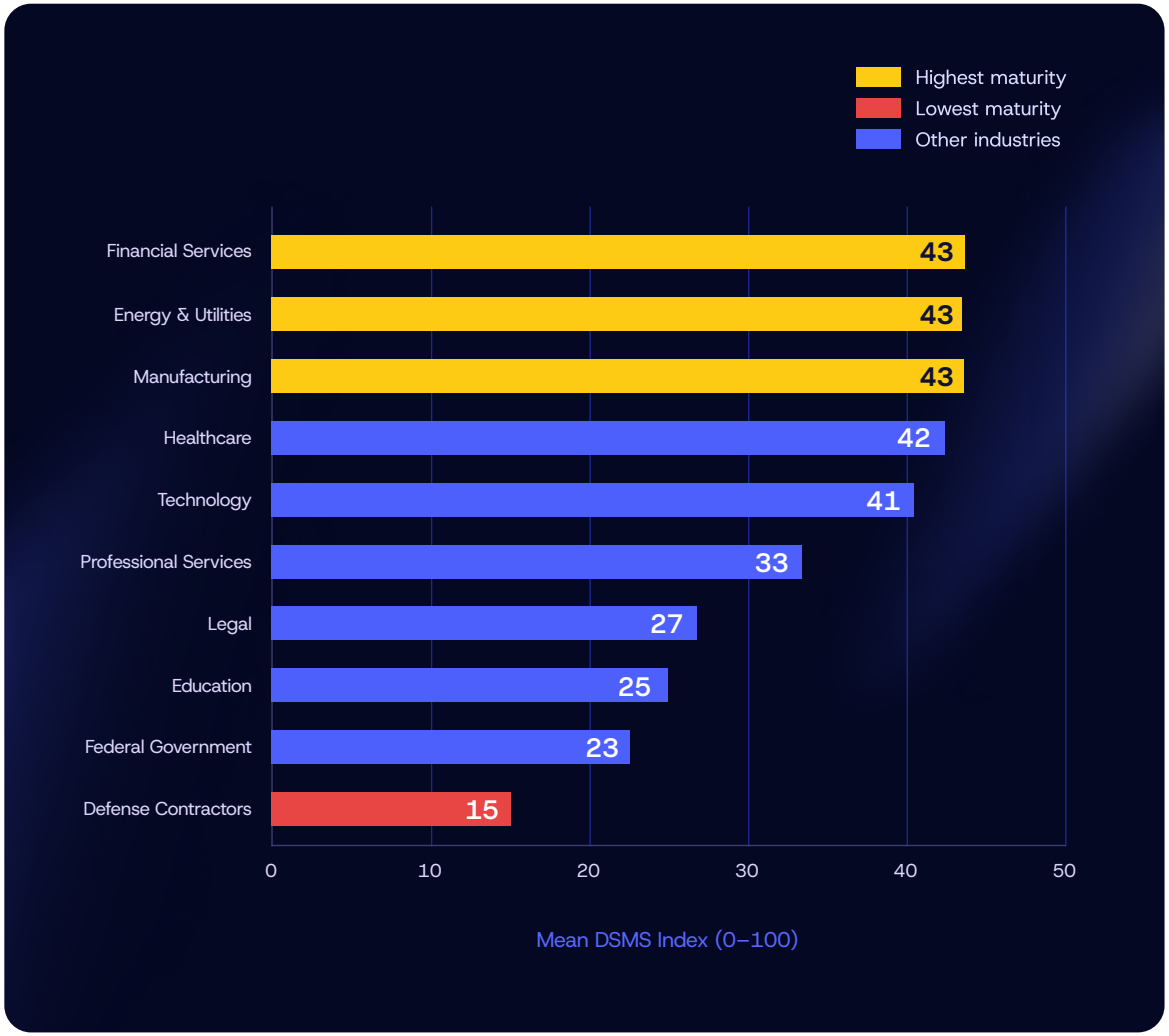

FIGURE 29

Mean Data Security Maturity Score (DSMS) by Region

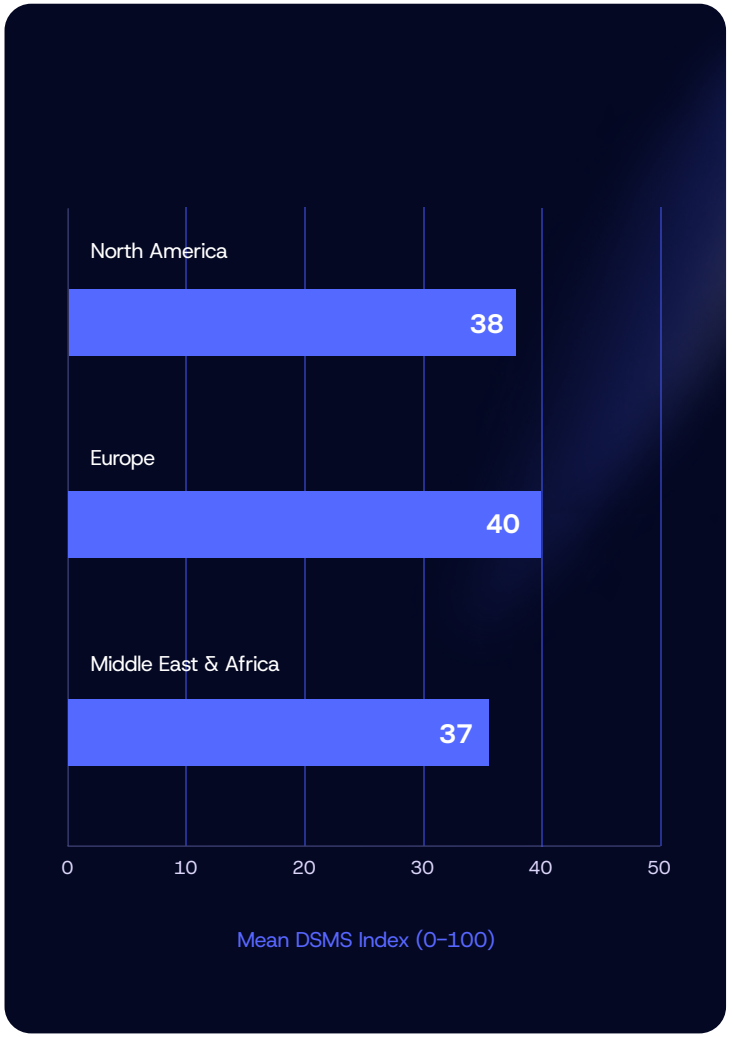

Regional Variations

North America (32% of the sample) reports the highest frequent shadow AI discovery rate: 24% of North American organizations discover shadow AI monthly or more, compared to 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. North America has the highest AI incident rate of any region at 84%, followed by the Middle East at 78% and Europe at 75%.

EU/UK organizations (56% of the sample) carry the highest rates of AI regulatory concern. EU/UK respondents rank AI-specific regulatory requirements as the top compliance challenge at higher rates than any other region — 29% place it first, versus 24% in North America — reflecting the concrete enforcement environment created by GDPR, the EU AI Act, and NIS2 notification timelines. Notably, however, 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. Whether this reflects more mature compliance programs, lower enforcement incidence, or underreporting is not resolvable from the survey data alone.

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 out of 100). MEA also 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 an actual incident response, versus 35% in North America. MEA also carries the highest compliance consequence rates of any region at 76%, suggesting an organization profile that is responding to incidents and enforcement actions rather than building preventively.

North America (32% of the sample, n=147) is overrepresented in the leadership quadrant at 1.19 times the expected rate [at the 45/45 threshold] or 1.32 times [at the 50/50 threshold]. It is slightly underrepresented in the laggard quadrant at 0.96 times. The regional distribution is more dispersed than Europe's — NAM produces more leaders proportionally but does not produce excess laggards, with its laggard rate close to proportional.

Europe (56% of the sample, n=258) is slightly underrepresented in the leadership quadrant at 0.94 times and slightly underrepresented in the laggard quadrant at 0.96 times.

Middle East (12% of the sample, n=54) is underrepresented in the leadership quadrant at 0.76–0.80 times the expected rate [depending on threshold]. More significantly, it is overrepresented in the laggard quadrant at 1.27 times — the highest laggard overrepresentation of any region. Middle East organizations are disproportionately concentrated at the bottom of the maturity distribution, not in the middle tiers. This is consistent with the region's lowest DSMS mean (35.7) and highest compliance consequence rate (76%), and with the finding that MEA organizations show the highest third-party revocation testing rates driven by actual incidents — an organization profile that is responding to breaches rather than building preventively.

Small organizations (under 1,000 employees, n=9) show the highest variance of any size band: a representation index of 1.53 in the leadership quadrant and 1.06 in the laggard quadrant. With only 9 respondents the small-organization finding should be interpreted cautiously — the directional pattern suggests high-variance outcomes but the laggard rate (55.6%) is not meaningfully above the overall rate (52.3%).

Representation Analysis: Regional Distribution Across Maturity Quadrants

North America shows the highest AI incident rates and highest frequent shadow AI discovery rates. MEA shows the highest compliance consequence rates and lowest DSMS scores. EU/UK ranks AI regulatory requirements as the top compliance challenge at higher rates than any other region.

North America (32% of the sample, n=147) reports the highest AI incident rate of any region at 84%, and the highest 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%. North America also shows the highest AI tool data exposure incident rates (36% in North America, 20% in EU/UK).

EU/UK organizations (56% of the sample, n=258) rank AI-specific regulatory requirements as the top compliance challenge at higher rates than any other region — 29% place it first, versus 24% in North America and 7% in the Middle East — reflecting the concrete enforcement environment created by GDPR, the EU AI Act, and NIS2 notification timelines. EU/UK organizations report compliance consequence rates of 57%, the lowest of any region. Whether this reflects more mature compliance programs, lower enforcement incidence, or survey composition effects is not resolvable from the data.

MEA organizations (12% of the sample, n=54) carry the highest compliance consequence rate of any region at 76%, and the lowest DSMS scores (mean 35.7). MEA also reports the highest third-party access revocation testing rates driven by actual incidents: 50% of MEA organizations have tested third-party revocation through an actual incident response, versus 35% in North America. The MEA pattern suggests an organization profile that is learning from breaches rather than building preventively.

Region-level means reflect aggregate DSMS and AIGMS scores. The representation analysis asks a different question: Are organizations from each region appearing in the extreme quadrants — leadership or laggard — at rates above or below their share of the sample?

North America (32% of the sample, n=147) is overrepresented in the leadership quadrant at 1.32 times the expected rate. It is slightly underrepresented in the laggard quadrant at 0.95 times. NAM produces more leaders than expected and slightly fewer laggards — a distribution skewed toward the upper tiers rather than dispersed across extremes.

Europe (56% of the sample, n=258) is underrepresented in the leadership quadrant at 0.86 times and nearly proportional in the laggard quadrant at 1.00 times. The EU/UK regulatory environment is producing compliance activity but not, at this stage, systematically better governance architecture. The region is producing leaders at below-expected rates despite having the strongest regulatory pressure of any region in the survey.

Middle East (12% of the sample, n=54) is underrepresented in the leadership quadrant at 0.80 times the expected rate and overrepresented in the laggard quadrant at 1.15 times — the highest laggard overrepresentation of any region. Middle East organizations are disproportionately concentrated at the bottom of the maturity distribution. This is consistent with the region's lowest DSMS scores (mean 35.7), highest compliance consequence rates (76%), and highest rate of third-party revocation testing triggered by actual incidents — an organization profile that is responding to breaches rather than building preventively, with incident-driven investment that has not yet produced the governance architecture needed to exit the laggard quadrant.

FIGURE 30

Mean AI Governance Maturity Score (AIGMS) by Industry

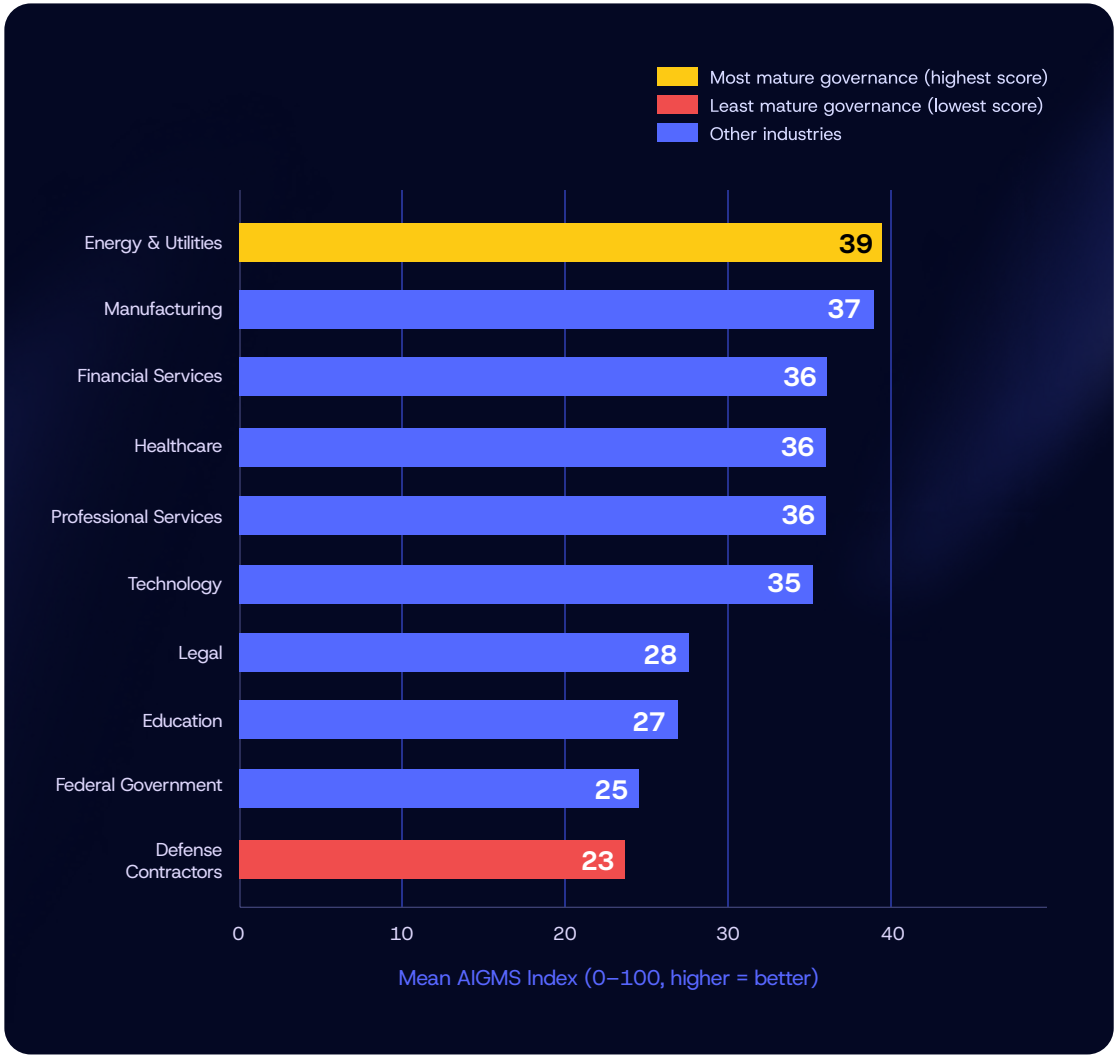

FIGURE 31

Mean AI Governance Maturity Score (AIGMS) by Region

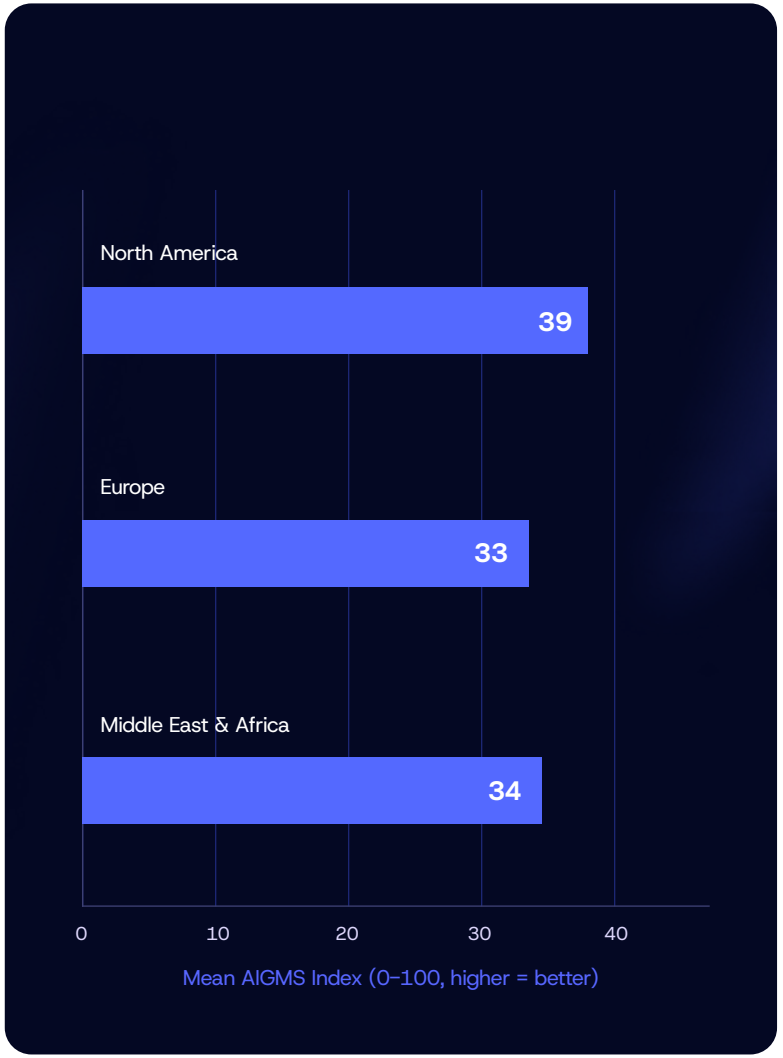

Organization Size Variations

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, n=211) 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, n=63) report the highest shadow AI discovery rate of any size band at 70% — higher than the 1,000–4,999 cohort (66%), Large (62%), or Global (62%). Shadow AI discovery scales with organizational complexity, not just resources.

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. The relationship between size and security control deployment is not linear — Global organizations face a coordination complexity that erodes the deployment advantage that scale otherwise provides. AIGMS tells a similar story: Large organizations (35.2) score above Mid (34.9), but Enterprise (34.6) and Global (33.4) decline, with Global carrying the lowest AI governance maturity of any band with substantial sample size.

The budget data provides the most important caution in this section. 54% of all respondents allocate 25% or more of their IT budget to cybersecurity. Despite this, the mean DSMS Index across the full sample is 39 — Tier 2, developing maturity. Cybersecurity spending does not automatically produce control deployment. The organizations that have reached Tier 4 are not necessarily the biggest spenders. They are the ones that have directed spending at the specific controls the DSMS measures.

The representation index asks which size bands are producing leadership-quadrant organizations, and which are producing laggards, at rates adjusted for their share of the sample.

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. They are underrepresented in the laggard quadrant at 0.94 times. The pattern is consistent with organizations that face sufficient regulatory scrutiny to mandate governance investment without the full coordination complexity that affects Global organizations.

Large organizations (5,000–10,000 employees, n=111) also overrepresent in the leadership quadrant at 1.16 times, with near-proportional laggard representation at 1.04 times.

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. The finding documents what the DSMS and AIGMS data imply: at Global scale, coordination complexity prevents governance from keeping pace with AI deployment. Global organizations are not failing at security investment. They are failing to convert security spending into governance architecture, and the representation index makes the magnitude of that failure visible.

Small organizations (under 1,000 employees, n=9) are overrepresented in the leadership quadrant at 1.59 times the expected rate and proportionally represented in the laggard quadrant at 1.01 times. With only 9 respondents the finding should be interpreted cautiously, but the directional pattern suggests that small organizations which invest in governance do so decisively — producing above-expected leader rates — while those that do not land in the laggard quadrant at roughly the same rate as the broader population.

FIGURE 32

Cybersecurity Spend Does Not Buy Security Maturity

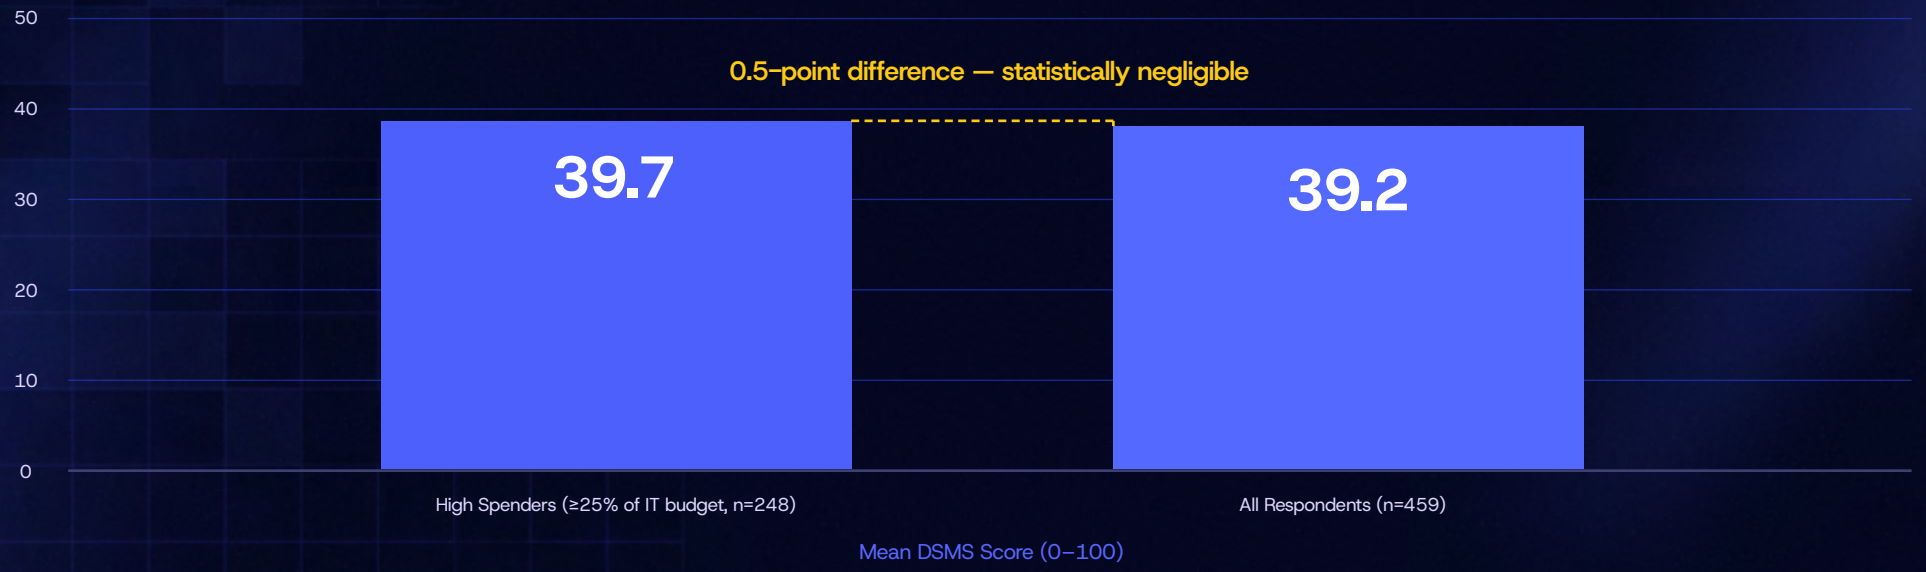

Cybersecurity Spend Does Not Predict Security Maturity

54% of respondents allocate 25% or more of their IT budget to cybersecurity. The mean DSMS for this high-spend cohort is 39.7 — less than one point above the full-sample mean of 39.2. Tier 4 (Advanced) organizations — those with the controls that produce 34% AI incident rates versus 91% in Tier 1 (Nascent) — are not distinguished by budget allocation. They are distinguished by architectural discipline: classification that enforces downstream behavior, SIEM integration that covers every sensitive data channel, kill switches that are tested before they are needed, and audit trails that exist before a regulator requests them.

The implication for security investment strategy is direct: The organizations that have closed the governance gap in this survey did not do so by spending more. They did so by directing spending toward the specific binary controls the DSMS measures — controls that are definable, deployable, and testable rather than aspirational.

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 the difference between containment within 31 days and containment after 91 days translates directly to millions in cost differential. 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. Classification that does not enforce behavior is not governance.

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. A Data Policy Engine closes the gap between policy documentation and policy enforcement: The 73% of organizations without purpose binding controls are operating on the honor system.

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 EU AI Act audit obligations operate on timelines that 50% of organizations cannot currently meet.[76]

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.

Security Maturity Readiness Checklist

The controls below are organized by implementation tier — the progression from Tier 1 to Tier 4 that the DSMS framework measures. Each tier builds on the previous. Organizations should complete Tier 1 controls before investing in Tier 3 controls. 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).

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 — not policy only • All sensitive data exchange channels unified under a single governance and audit framework	DSMS controls c_1–c_3: 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 c_4–c_7: 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 — not attestation only • Dedicated AI data governance ownership assigned — not an add-on to an existing role • AI regulation compliance externally validated (not self-attested only) • Compliance evidence generated continuously from operational systems — not assembled for audits	DSMS controls c_8–c_{11}: AI governance, third-party controls, continuous evidence

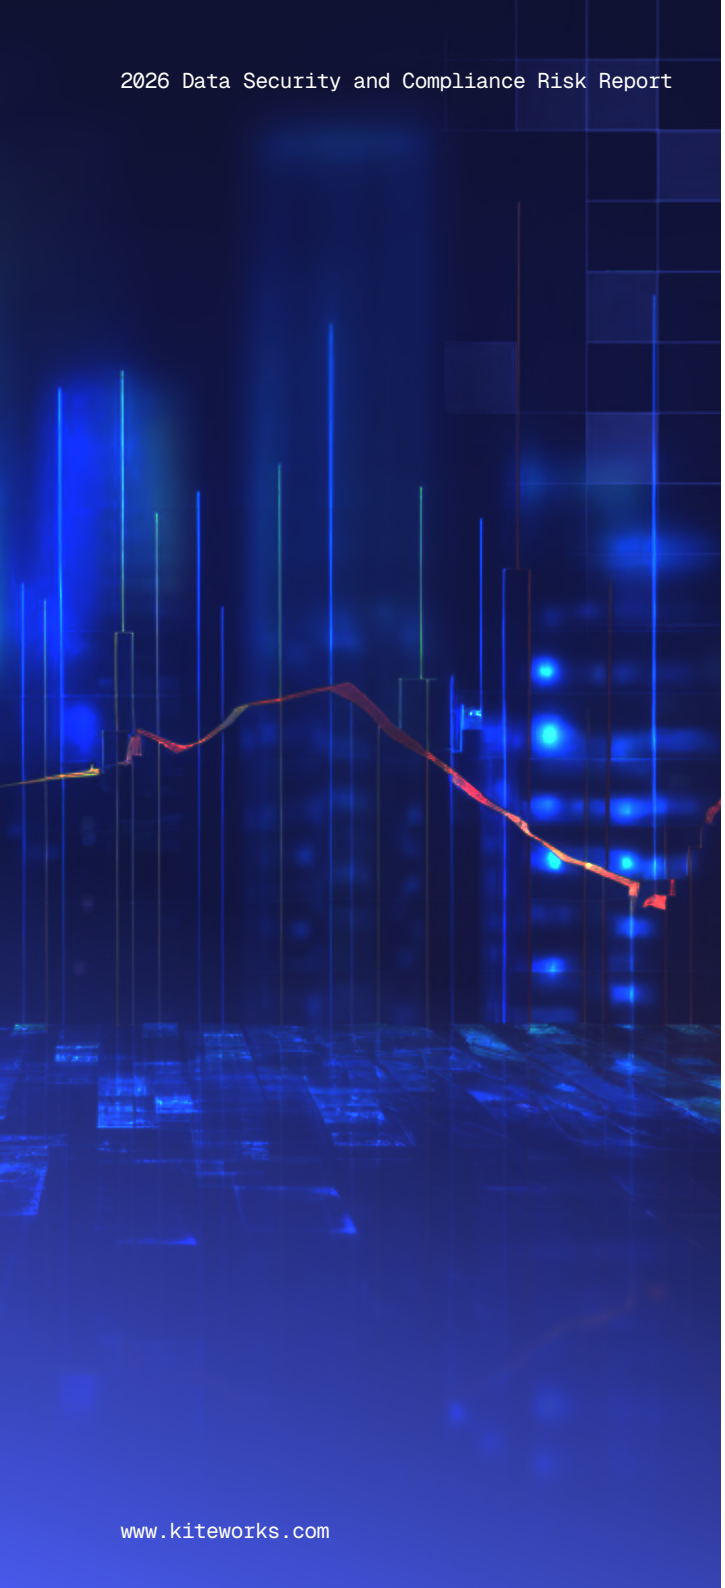

The tier progression table above describes what organizations should build. The three figures below show where 459 organizations actually are — plotted simultaneously on both maturity dimensions. Each respondent appears as a single dot positioned by their DSMS score (horizontal axis) and AIGMS score (vertical axis), inside the 4×4 tier grid defined by the thresholds above.

The curved dashed lines are DSCRI isocurves — lines of equal readiness. Two organizations can sit in the same tier cell and carry dramatically different DSCRI scores depending on whether their maturity is balanced or concentrated in one dimension. An organization with DSMS 70 and AIGMS 20 sits in Tier 3 on the security axis but carries a DSCRI of 14 — below the survey mean of 16.2, and lower than an organization with DSMS 40 and AIGMS 45 (DSCRI 18). The DSCRI's multiplicative structure penalizes single-dimension excellence. Moving up the AIGMS axis produces more readiness improvement per unit of effort than adding further security controls at current organizational baselines — because AI governance is the scarcer resource.

The black box marks the upper-right leadership quadrant: DSMS ≥ 50 and AIGMS ≥ 50 simultaneously. Eighty-six of the 459 respondents — 19% — occupy this space. Their mean DSCRI is 46, versus 16.2 for the full sample. The red box marks the lower-left laggard quadrant: DSMS < 50 and AIGMS < 50 simultaneously. Three hundred and three respondents — 66% — occupy this space. Their mean DSCRI is 8. The distance between those two populations is not primarily budget or headcount. It is the decision to treat AI governance as an architectural requirement rather than a compliance checkbox.

FIGURE 33

Mean Data Security and Compliance Readiness Index (DSCRI) by Region

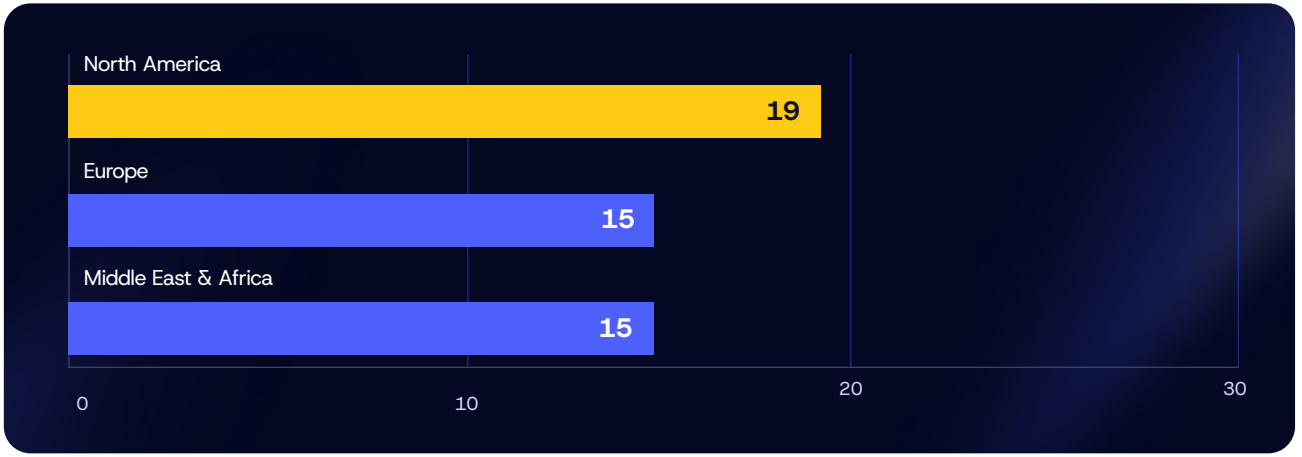

Each scatter plot organizes the 459 respondents across four positions defined by the intersection of DSMS and AIGMS at the 50-point threshold on each axis. The quadrant labels name what those positions mean operationally.

Resilient (upper-right: DSMS ≥ 50, AIGMS ≥ 50) — strong security controls and strong AI governance. 19% of the sample occupies this quadrant. Their mean DSCRI is 45, versus 16.2 for the full sample.

Fortified (lower-right: DSMS ≥ 50, AIGMS < 50) — strong security controls, AI governance gaps. These organizations have invested in the technical security layer but have not extended that discipline into AI-specific governance: kill switches, audit trails, AI access visibility, and board-level reporting remain incomplete.

Policy-Led (upper-left: DSMS < 50, AIGMS ≥ 50) — AI governance frameworks in place, but core security controls lag. These organizations have policies, reporting structures, and regulatory awareness without the technical controls that enforce them. The policy is ahead of the architecture.

Exposed (lower-left: DSMS < 50, AIGMS < 50) — security control gaps and AI governance gaps simultaneously. 66% of the sample is here. This is not a population of organizations that have not started. It is a population where the gap between stated intent and deployed control is largest.

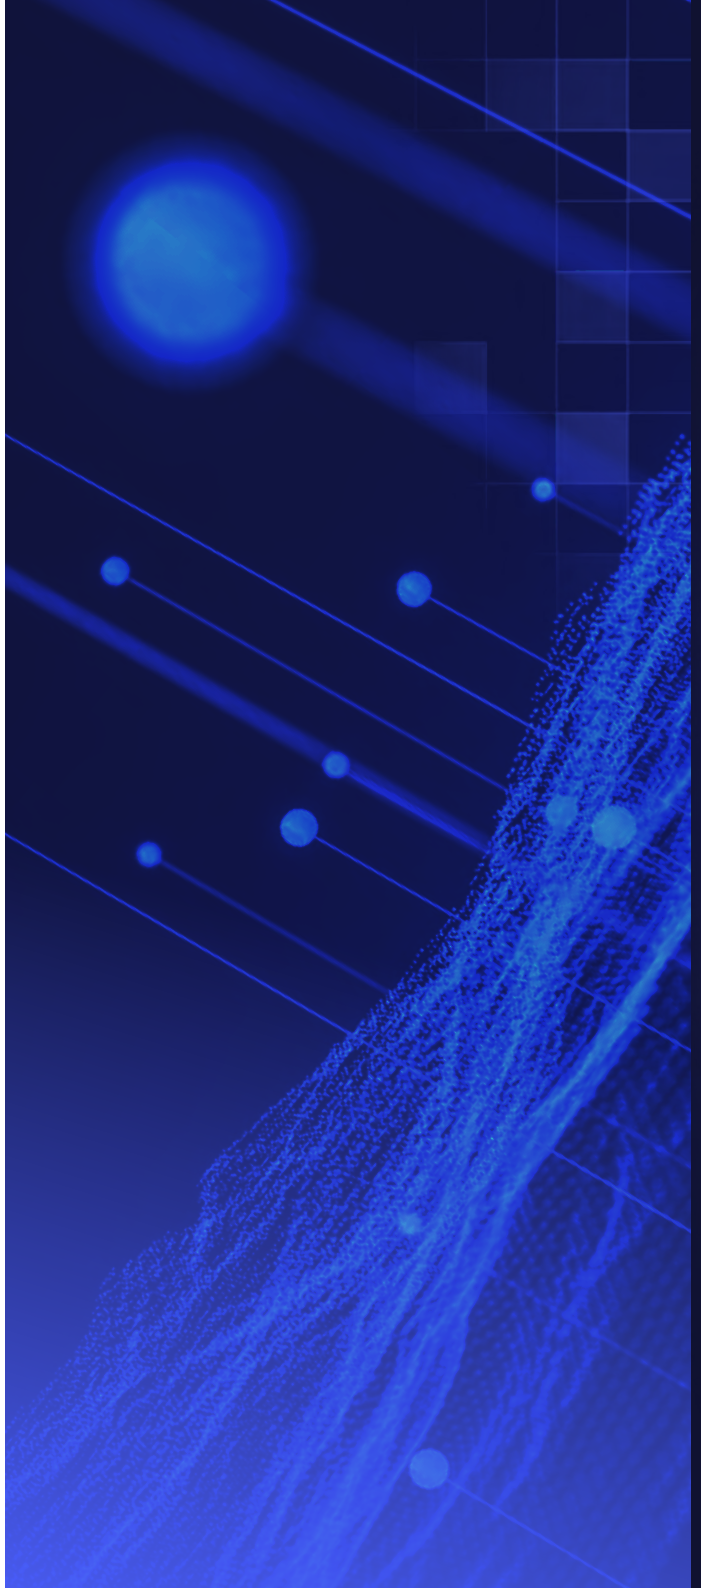

FIGURE 34

DSMS vs. AIGMS by Industry

Data Security Maturity Score vs. AI Governance Maturity Score — Segmented by Industry

Grid boundary: 50 / 50 | Yellow dashed lines = means

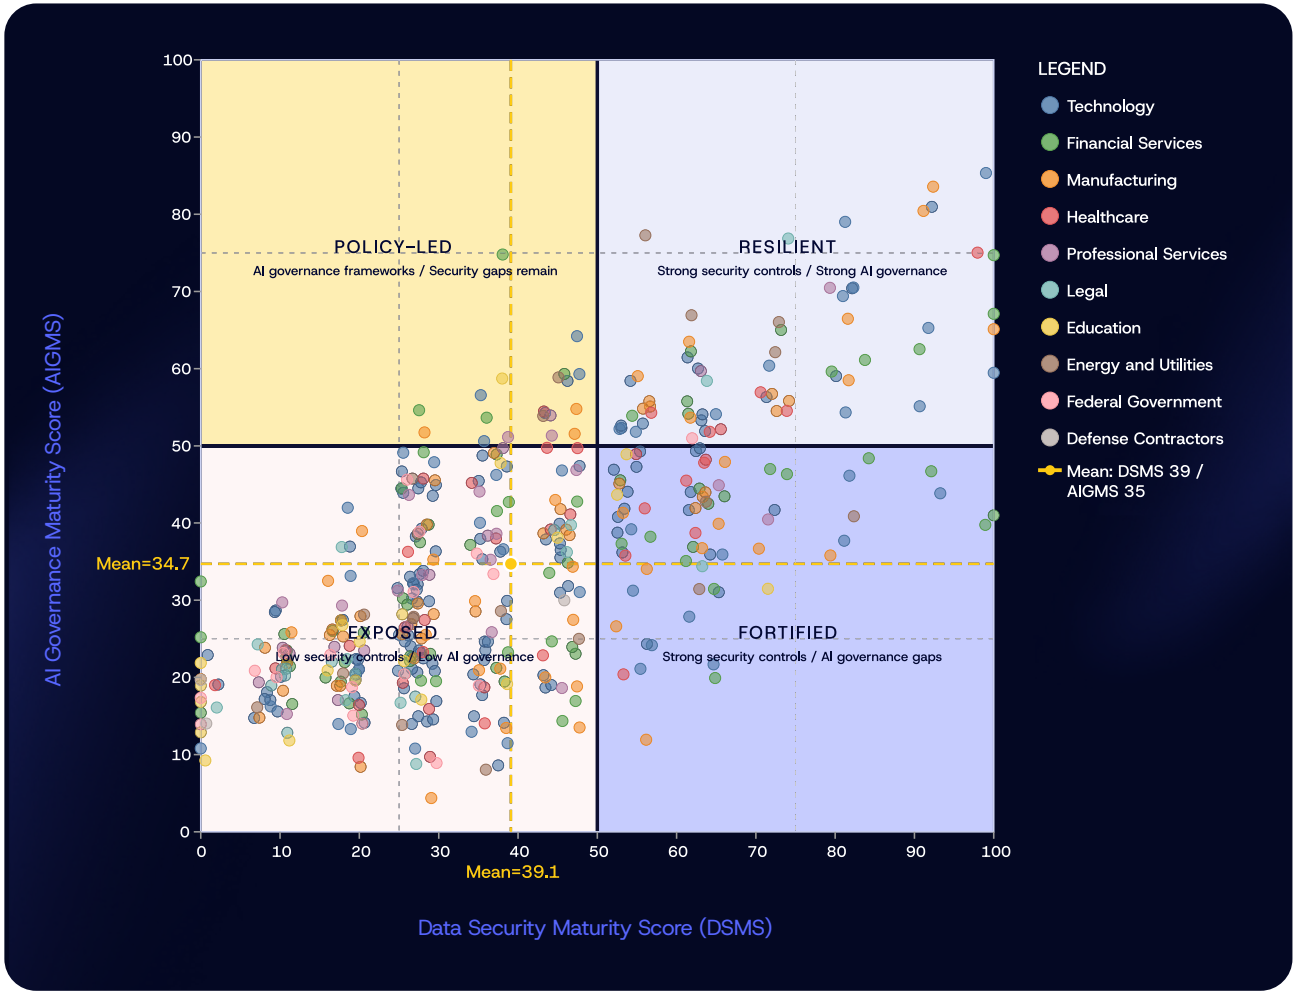

The industry plot shows that concentration in the Exposed quadrant is not evenly distributed. Legal, Education, Federal Government, and Defense Contractors are visually the most heavily weighted toward the lower-left — Defense Contractors respondents appear almost entirely in the Exposed quadrant. Technology, despite its sector's familiarity with the controls in question, shows its largest cluster in Exposed as well, with relatively few respondents reaching the Resilient quadrant.

Energy and Utilities is the exception visible at a glance: the sector produces some of the highest individual DSCRI scores in the survey, with several respondents positioned well into the Resilient quadrant. Financial Services shows broader distribution across quadrants than most sectors. The sidebar representation indices quantify these patterns: Energy and Utilities is overrepresented in the leadership quadrant at 1.43 times its sample share; Legal is overrepresented in the laggard quadrant at 1.31 times.

FIGURE 35

Industry Representation Index by Quadrant

Industry	Resilient	Fortified	Policy-Led	Exposed
Financial Services	1.37×	0.82×	2.04×	0.82×
Technology	1.10×	1.04×	0.91×	0.98×
Energy & Utilities	1.00×	1.00×	1.86×	0.90×
Professional Services	0.92×	0.69×	0.00×	1.19×
Manufacturing	0.91×	1.09×	1.19×	0.98×
Legal	0.91×	0.00×	0.00×	1.29×
Healthcare	0.88×	1.84×	0.65×	0.94×
Education	0.67×	1.00×	0.62×	1.12×
Federal Government	0.00×	1.11×	0.00×	1.33×
Defense Contractors	0.00×	0.00×	0.00×	1.50×

Financial Services and Technology are the only sectors with meaningful Resilient concentration — 1.37× and 1.10× respectively — reflecting both strong technical controls and AI governance maturity. Healthcare's dominant pattern is Fortified (1.84×): strong security infrastructure that has not been extended into AI-specific governance. Energy and Utilities shows the opposite split, leading in Policy-Led (1.86×) — governance frameworks and regulatory awareness ahead of the technical controls that enforce them. The most concerning finding is at the bottom of the table: Defense Contractors, Federal Government, and Legal carry zero Resilient representation and are heavily concentrated in the Exposed quadrant, at 1.50×, 1.33×, and 1.29× respectively. These are sectors that handle some of the most sensitive data in the economy, yet their maturity profiles are the weakest in the dataset.

FIGURE 36

DSMS vs. AIGMS by Organization Size

Data Security Maturity Score vs. AI Governance Maturity Score — Segmented by Organization Size

Grid boundary: 50 / 50 | Yellow dashed lines = means

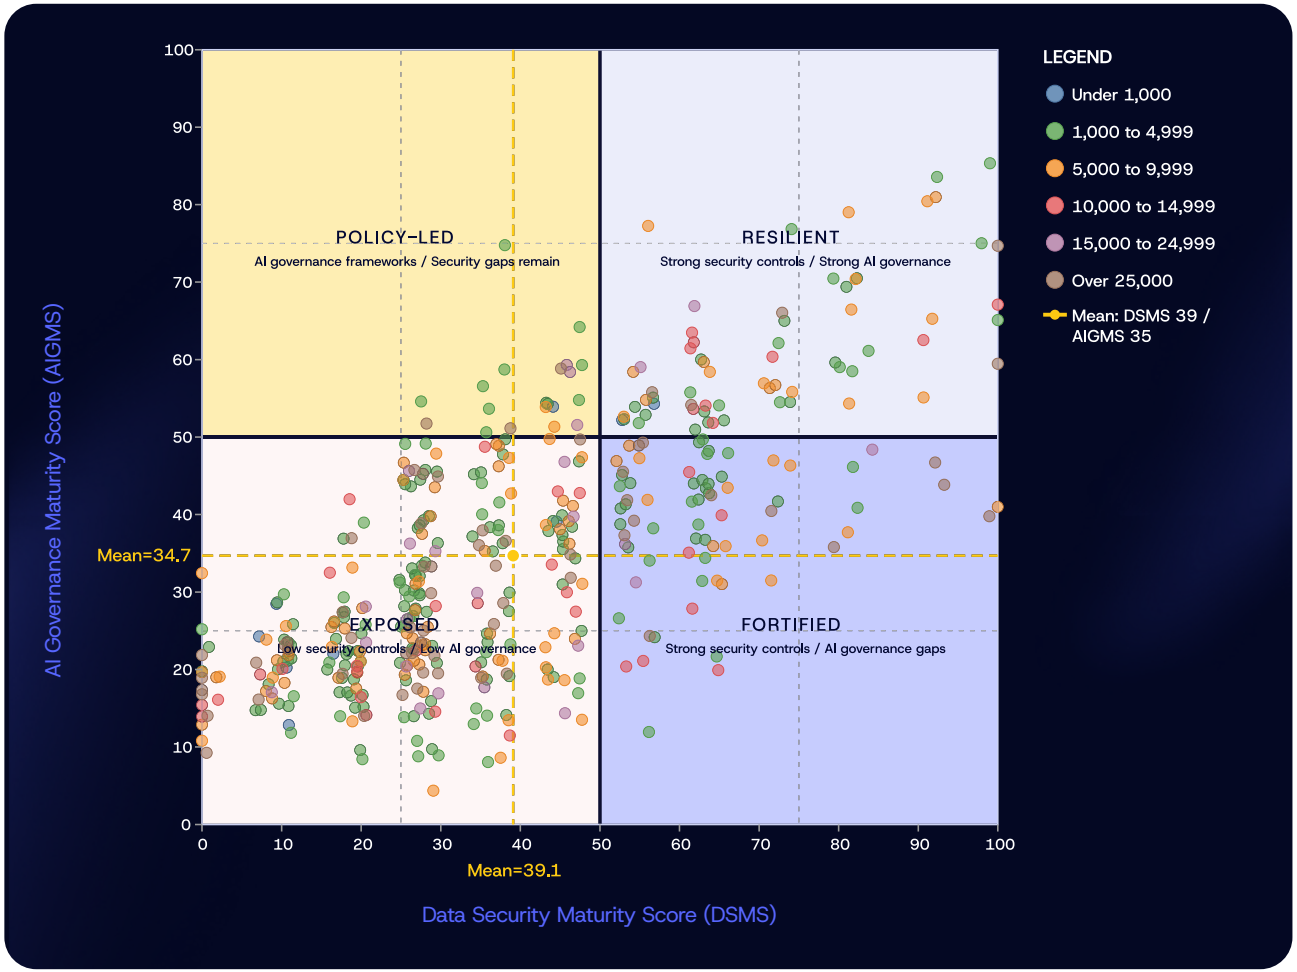

The size plot makes the Global organization (over 25,000 employees) pattern immediately visible: that population's dots cluster heavily in the Exposed quadrant, with very few reaching the Fortified or Resilient zones. Their leadership quadrant index is 0.55 — the lowest of any size band — despite carrying the highest mean security budgets in the survey. Budget allocation is not the constraint.

Enterprise organizations (10,000–25,000 employees) show the opposite pattern. Their dots reach into the Resilient quadrant at a higher rate than any other size band, producing a leadership index of 1.25. Mid-market organizations (1,000–4,999 employees) represent the largest share of the sample numerically and dominate the visual mass of the Exposed quadrant accordingly — their index of 0.95 reflects near-proportional distribution, concentrated below both mean lines. The Large band (5,000–9,999) sits between these poles at 1.16.

The Policy-Led quadrant is sparsely populated across all size bands, but the dots that do appear there are distributed across size categories — suggesting that governance-ahead-of-security is not a size-specific pattern but an organizational choice that occurs at multiple scales.

FIGURE 37

Organizational Size Representation Index by Quadrant

Organizational Size	Resilient	Fortified	Policy-Led	Exposed
Under 1,000	1.48×	0.00×	0.00×	1.16×
1,000–4,999	0.79×	1.04×	1.35×	1.00×
5,000–9,999	1.44×	0.90×	0.78×	0.94×
10,000–14,999	0.85×	1.79×	0.64×	0.96×
15,000–24,999	0.83×	0.42×	1.03×	1.12×
Over 25,000	1.02×	0.92×	0.57×	1.06×

Organizational size does not produce a monotonic maturity pattern — the data instead reveals distinct architectural profiles at each scale. Mid-market organizations (5,000–9,999 employees) show the strongest Resilient concentration at 1.44×, suggesting that this band has both the operational focus and the organizational agility to deploy technical controls and governance simultaneously. Large enterprises (10,000–14,999) are most concentrated in Fortified (1.79×) — significant security investment that has not translated into AI governance maturity, likely a function of complexity and organizational sprawl. The 1,000–4,999 band leads in Policy-Led (1.35×), indicating governance frameworks and compliance awareness ahead of the technical enforcement layer. The largest organizations (Over 25,000) are near-proportional across Exposed and Resilient, with no strong quadrant signature — scale alone neither protects nor disadvantages them. The Under 1,000 band shows high Resilient concentration (1.48×) but with only nine respondents, that result carries wide uncertainty.

FIGURE 38

DSMS vs. AIGMS by Region

Data Security Maturity Score vs. AI Governance Maturity Score — Segmented by Region

Grid boundary: 50 / 50 | Yellow dashed lines = means

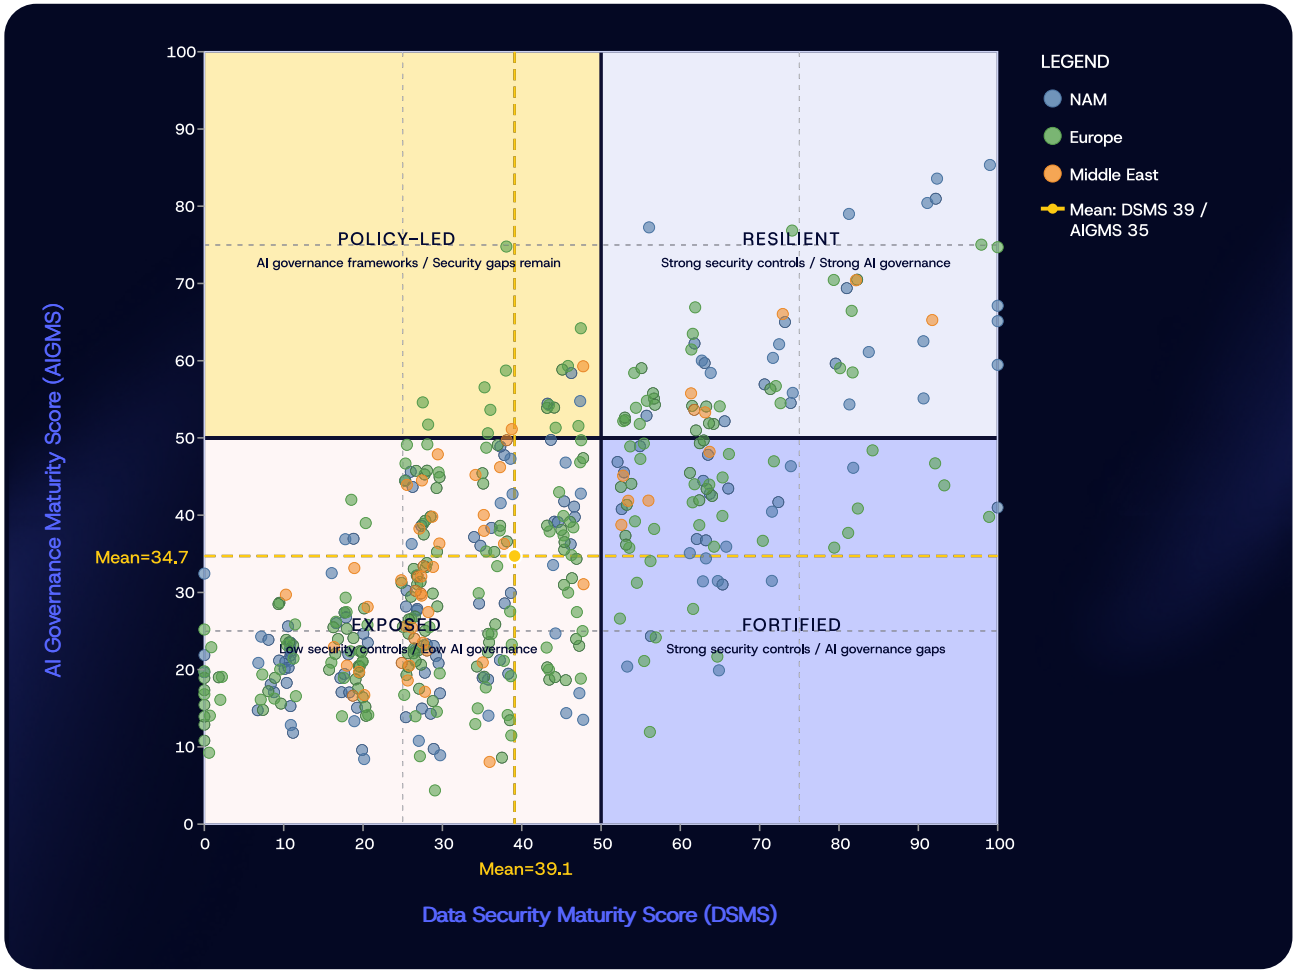

North America (blue) shows the widest spread of the three regions — more respondents reach the upper portions of the chart, and the NAM cluster in the Resilient quadrant is visually the densest of the three. North America is overrepresented in the leadership quadrant at 1.32 times its proportional share.

Europe (green) is the largest regional group at 56% of the sample, and its visual mass is concentrated in the Exposed quadrant — specifically in the lower-left and lower-center cells. Europe's representation index in the leadership quadrant is 0.86, slightly underrepresented. The region's governance gap is the dominant pattern: high DSMS scores are present, but most European respondents have not extended those controls into the AIGMS dimension.

Middle East (orange) is the smallest region at 12% of the sample, but its distribution is notable. Middle East respondents are overrepresented in the Exposed quadrant at 1.15 times their sample share — more concentrated in the lower-left than their proportional representation would predict. No Middle East respondents appear in the Policy-Led quadrant; the governance investment seen in parts of Europe is absent here.

FIGURE 39

Regional Representation Index by Quadrant

Region	Resilient	Fortified	Policy-Led	Exposed
NAM	1.45×	1.09×	0.84×	0.91×
Europe	0.85×	1.04×	1.11×	1.01×
Middle East	0.49×	0.55×	0.92×	1.19×

2.00

North America is the only region with meaningful Resilient concentration (1.45×), reflecting a combination of mature security tooling and more developed AI governance frameworks relative to its sample share. Europe's dominant pattern is Fortified (1.04×) with a secondary Policy-Led signature (1.11×) — strong security infrastructure and regulatory awareness, but AI governance maturity consistently lagging behind both. That pattern aligns with the region's regulatory environment: GDPR and NIS-2 have driven compliance investment and policy development, but technical AI governance controls — kill switches, AI-specific audit trails, access visibility — remain incomplete. The Middle East presents the most concentrated Exposed profile (1.19×) with the weakest Resilient representation (0.49×) of any region, meaning Middle East respondents reach the upper-right quadrant at less than half the rate their sample share would predict. No region reaches proportional representation in Policy-Led, confirming that governance-ahead-of-security is a pattern distributed across organizational choice rather than geography.

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 DSMS Tier 3 and Tier 4, and from AIGMS Tier 2 to AIGMS 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

FIGURE 40

Top Security Investment Priorities for Next 12 Months

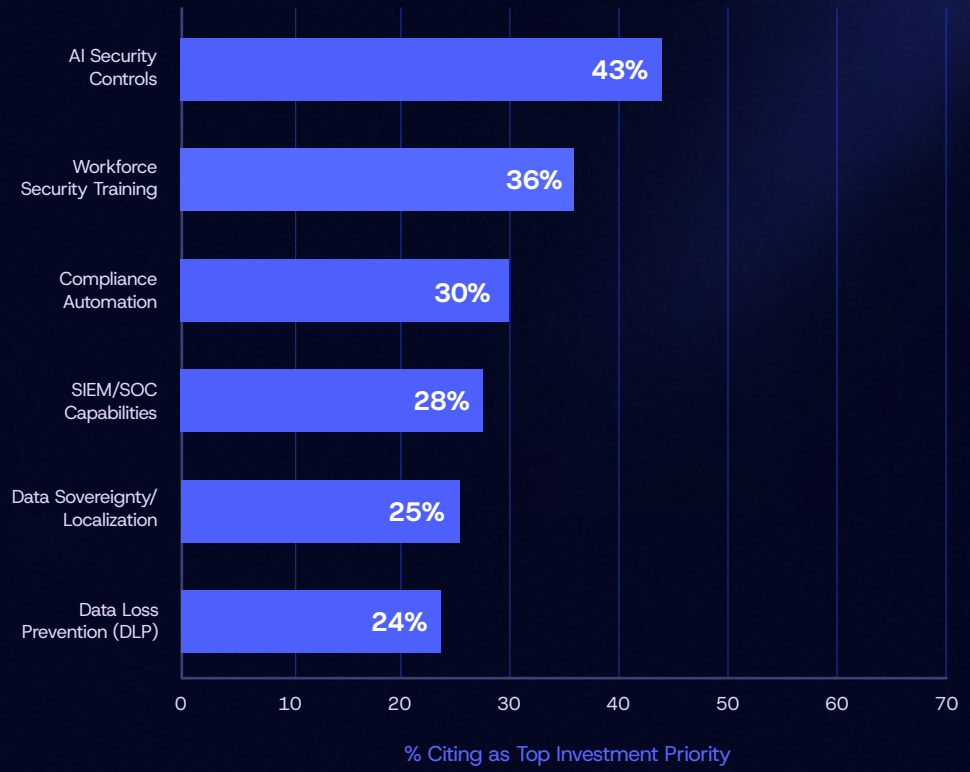

produce a complete audit record within one business day. None of these are aspirational. They are the specific controls that separate the 22% that avoided all incidents from the 80% that experienced at least one. The combined effect of progress on both dimensions is captured in the Data Security and Compliance Readiness Index (DSCRI = DSMS × AIGMS/100) — the survey mean of 16.2 reflects how far most organizations remain from the readiness level the current threat environment requires.

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 we run this survey again next year.

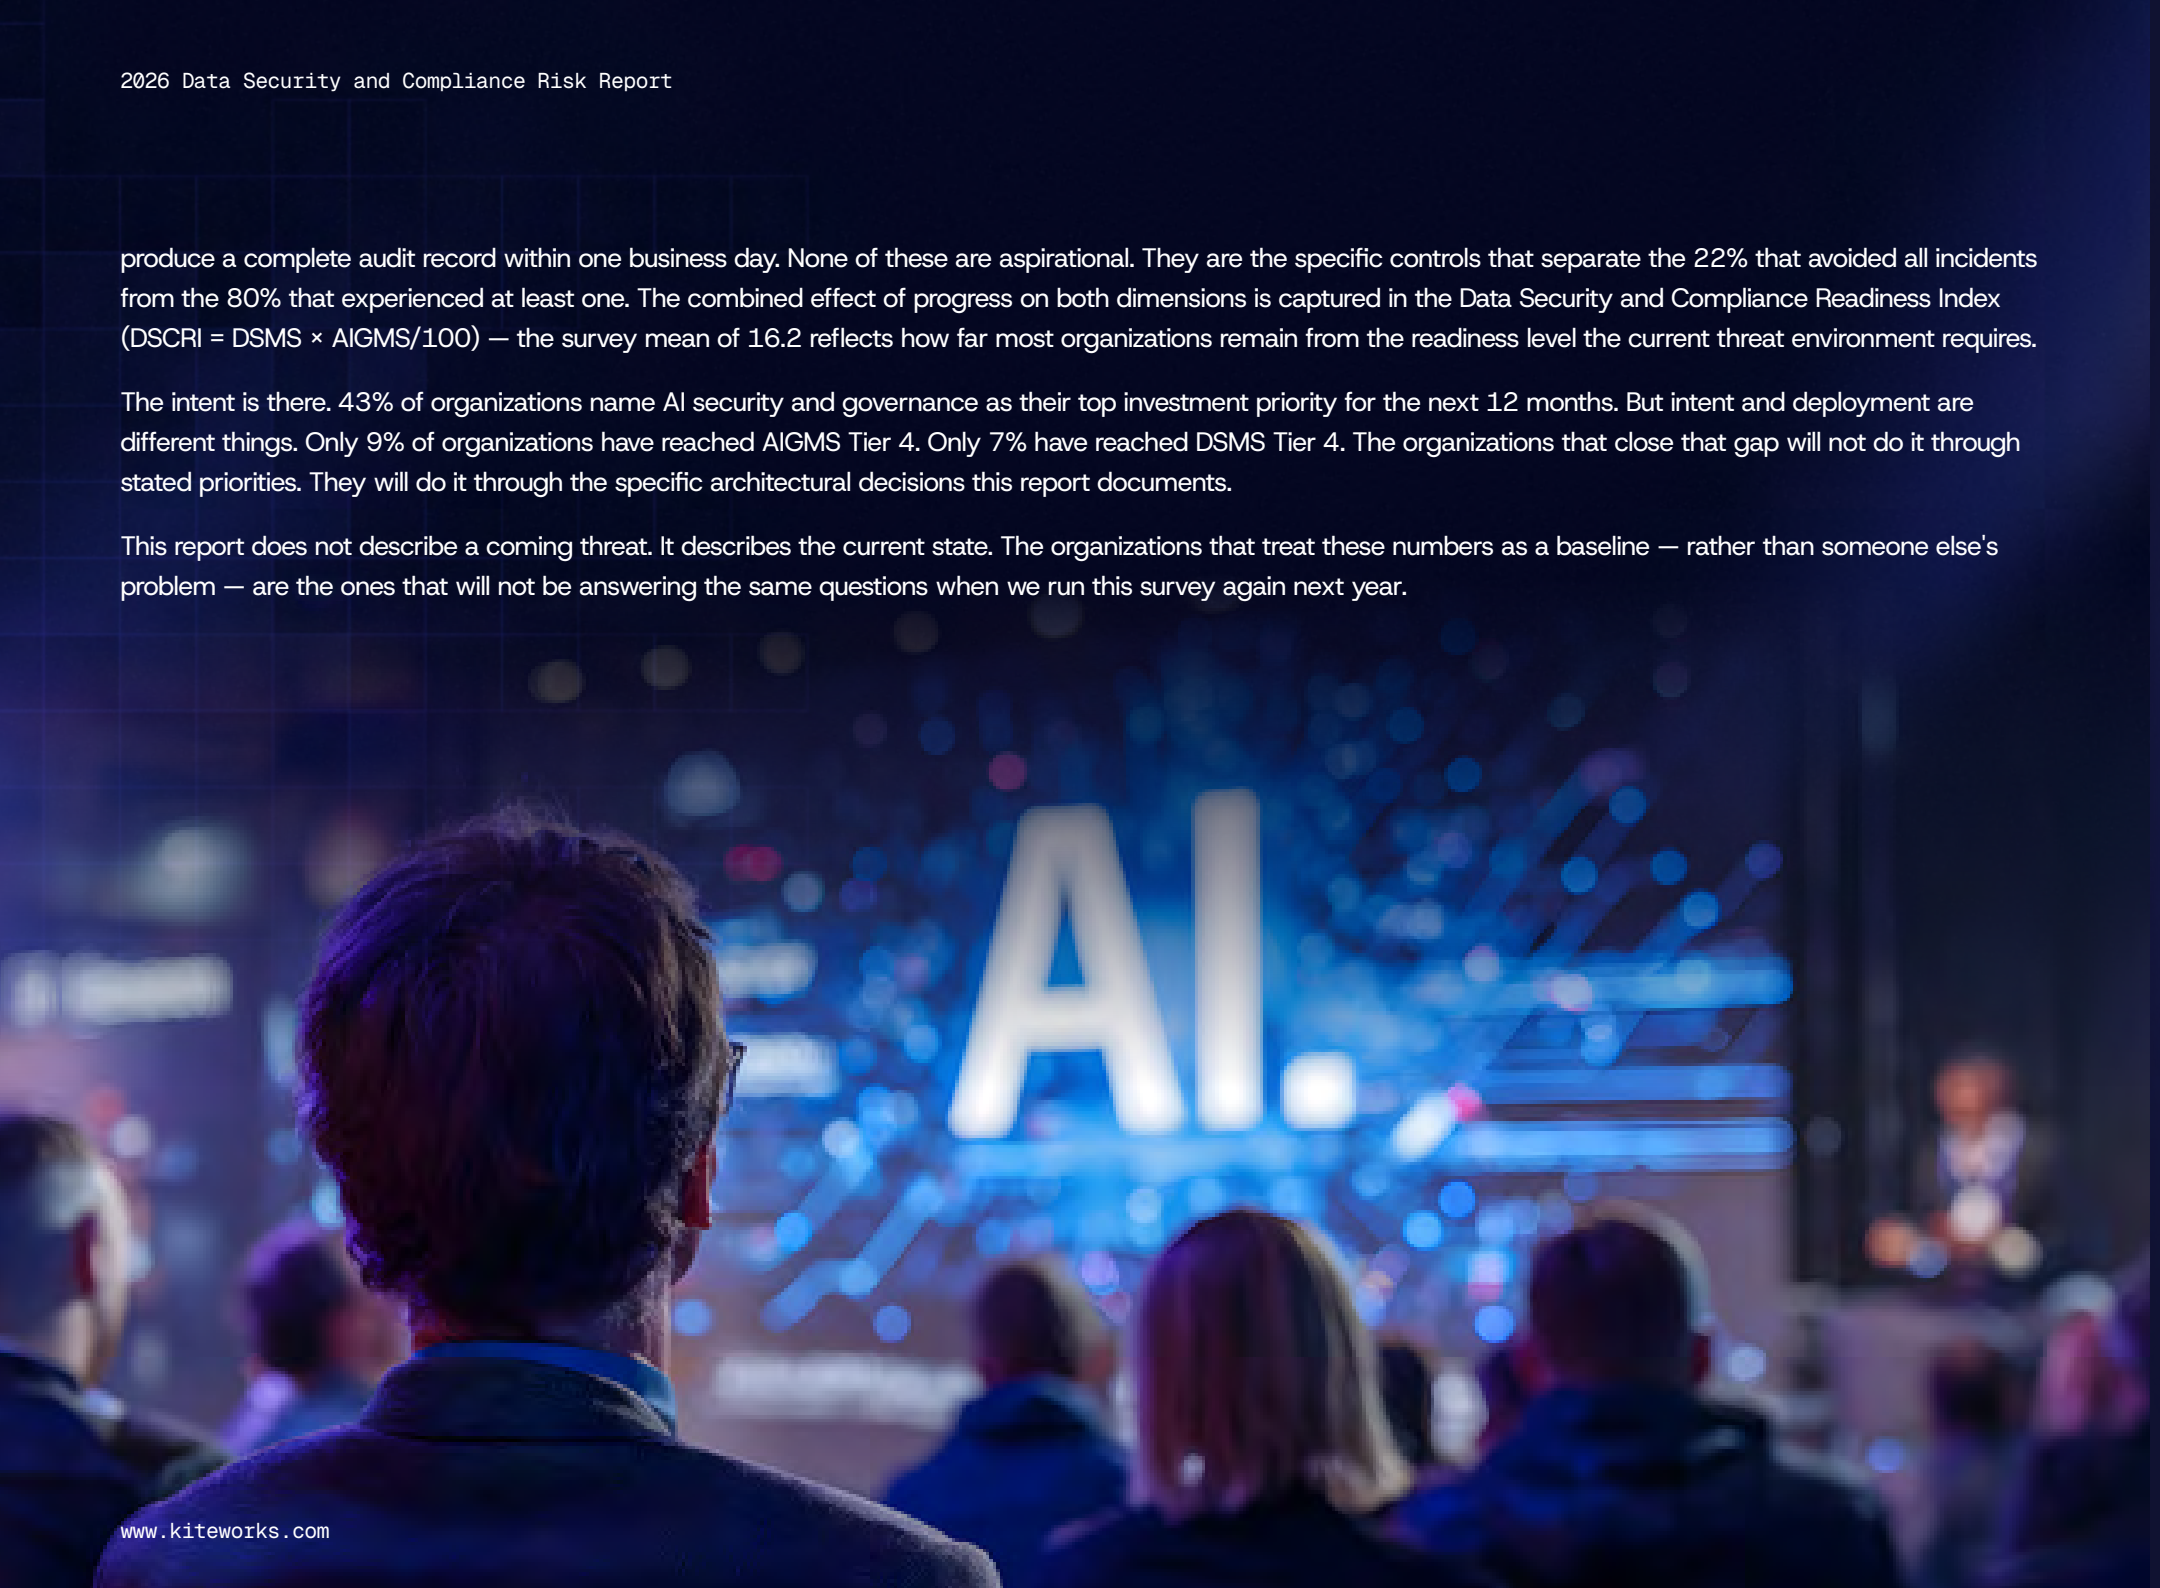

Legal Disclaimer

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.

Endnotes

1. World Economic Forum, "Global Cybersecurity Outlook 2026," January 2026.
2. IBM Security, Cost of a Data Breach Report 2025 (IBM Corporation, 2025), 47.
3. CrowdStrike, 2026 Global Threat Report (CrowdStrike Holdings, 2026), 6.
4. Verizon Business, 2026 Data Breach Investigations Report (Verizon Business, 2026), 11.
5. Kiteworks, AI Data Security and Compliance Risk Report (Kiteworks, 2025), 4.
6. Verizon Business, 2026 Data Breach Investigations Report, 13.
7. IBM Security, Cost of a Data Breach Report 2025, 12.
8. Kiteworks, 2025 Data Security and Compliance Risk: Annual Survey Report (Kiteworks, 2025), 18.
9. Verizon Business, 2026 Data Breach Investigations Report, 11.
10. Verizon Business, 2026 Data Breach Investigations Report, 40.
11. Verizon Business, 2026 Data Breach Investigations Report, 21.
12. Verizon Business, 2026 Data Breach Investigations Report, 36.
13. Verizon Business, 2026 Data Breach Investigations Report, 13.
14. Kiteworks, AI Data Security and Compliance Risk Report, 12.
15. Verizon Business, 2026 Data Breach Investigations Report, 88.
16. Kiteworks, 2025 Data Security and Compliance Risk: Annual Survey Report, 18.
17. Researchers at 13 institutions, Agents of Chaos: Agentic AI Security in Practice (February 2026), 8–14.
18. Researchers at 13 institutions, Agents of Chaos, 7.
19. IBM Security, Cost of a Data Breach Report 2025, 34–35.
20. Kiteworks, Data Security and Compliance Risk: AI Survey Report, 4.
21. CSA / Aembit, Identity and Access Gaps in the Age of Autonomous AI (Cloud Security Alliance, 2026), 12.
22. Kiteworks, "Data Security and Compliance Risk: 2026 Annual Survey Report," 2026.
23. Researchers at 13 institutions, Agents of Chaos, 22.
24. Researchers at 13 institutions, Agents of Chaos, 15–17.
25. Kiteworks, 2025 Data Security and Compliance Risk: Annual Survey Report, 24.
26. IBM Security, Cost of a Data Breach Report 2025, 35.
27. IBM Security, Cost of a Data Breach Report 2025, 44.
28. CrowdStrike, 2026 Global Threat Report, 6.
29. Mandiant / Google Cloud, M-Trends 2026 (Google LLC, 2026).
30. CrowdStrike, 2026 Global Threat Report, 37.
31. CrowdStrike, 2026 Global Threat Report, 7.
32. Kiteworks, 2025 Data Security and Compliance Risk: Annual Data Forms Survey Report (Kiteworks, 2025), 10.
33. Kiteworks, 2025 Data Security and Compliance Risk: MFT Survey Report (Kiteworks, 2025). Among organizations using MFT, 59% suffered a security incident in the past 12 months; 62% operate fragmented data exchange systems across MFT, email, file sharing, and web forms, creating inconsistent security policies and multiple points of vulnerability.
34. Thales Group, 2026 Thales Data Threat Report (Thales, 2026).
35. Kiteworks, 2025 Data Security and Compliance Risk: Annual Data Forms Survey Report, 6.
36. Kiteworks, Data Security and Compliance Risk: 2025 AI Survey Report, 9.
37. Kiteworks, 2025 Data Security and Compliance Risk: Annual Survey Report, 28.
38. Kiteworks, 2025 Data Security and Compliance Risk: MFT Survey Report, 11.
39. OPSWAT / Ponemon Institute, The State of File Security 2025 (OPSWAT, 2025), 14.
40. Verizon Business, 2026 Data Breach Investigations Report, 15.
41. Kiteworks, 2025 Data Security and Compliance Risk: MFT Survey Report, 12.
42. Ibid.
43. Verizon Business, 2026 Data Breach Investigations Report, 13.
44. Verizon Business, 2026 Data Breach Investigations Report, 36.
45. Check Point Research, State of Cyber Security 2026 (Check Point Software Technologies, 2026), 62.
46. CrowdStrike, 2025 Global Threat Report (CrowdStrike Holdings, 2025), 21.
47. Kiteworks, 2025 Data Security and Compliance Risk: Annual Data Forms Survey Report, 10.
48. Kiteworks, 2025 Data Security and Compliance Risk: Annual Data Forms Survey Report, 7.
49. Kiteworks, 2025 Data Security and Compliance Risk: MFT Survey Report, 11.
50. CrowdStrike, 2025 Global Threat Report, 22. The Verizon 2026 Data Breach Investigations Report confirms the visibility consequence: Organizations with SIEM-integrated infrastructure identify breaches significantly faster than those relying on fragmented or absent log coverage. The 63% of MFT users operating outside the SIEM perimeter are, by definition, operating outside their own detection perimeter.
51. DORA, Article 19(1).
52. DORA, Articles 28–44.
53. NIS2 Directive, Article 34(4).
54. U.S. Department of Defense, CMMC Program Final Rule, 32 CFR Part 170.
55. Black Kite Research, 2026 Third-Party Breach Report: Managing Risk Concentration in the Era of Cascading Failures (Black Kite, 2026).
56. Kiteworks, "Data Security and Compliance Risk: 2026 Annual Survey Report," 2026.
57. Ibid.
58. IBM Security, Cost of a Data Breach Report 2025, 19.
59. IBM Security, Cost of a Data Breach Report 2025, 42–43.
60. Kiteworks, 2025 Data Security and Compliance Risk: Annual Survey Report, 32.
61. Kiteworks, 2025 Data Security and Compliance Risk: Annual Survey Report, 36.
62. Verizon Business, 2026 Data Breach Investigations Report, 11.
63. Verizon Business, 2026 Data Breach Investigations Report, 21.
64. ENISA, Supply Chain Cybersecurity Report 2025 (European Union Agency for Cybersecurity, 2025), 4.
65. EU AI Act, Articles 6, 12–14, 43.
66. Kiteworks, 2025 Data Security and Compliance Risk: Annual Survey Report, 40.
67. Kiteworks, 2025 Data Security and Compliance Risk: Annual Survey Report, 41.
68. Ponemon Institute / DTEX Systems, 2026 Cost of Insider Risks Global Report, 8.
69. CrowdStrike, 2026 Global Threat Report, 6.
70. Cisco, "2026 Data Privacy Benchmark Study," March 2026.
71. CSA / Aembit, Identity and Access Gaps in the Age of Autonomous AI, 14.
72. Cisco, "2026 Data Privacy Benchmark Study," March 2026.
73. European Parliament and Council, Regulation (EU) 2022/2554 on Digital Operational Resilience for the Financial Sector (DORA) (Official Journal of the European Union, 2022), Article 2; Directive (EU) 2022/2555 (NIS2), Article 3.
74. Check Point Research, State of Cyber Security 2026, 28.
75. IBM Security, Cost of a Data Breach Report 2025, 6.
76. DORA, Article 19; NIS2 Directive, Article 23; U.S. Department of Defense, CMMC Program Final Rule, 32 CFR Part 170.

Kiteworks

July 2026

www.kiteworks.com