Kiteworks

EXECUTIVE SUMMARY

2026 Data Security and Compliance Risk Report — Europe Executive Summary

THE AI DATA GOVERNANCE GAP IN EUROPE: STRONG SECURITY CONTROLS, LAGGING AI GOVERNANCE, AND A REGULATORY DEADLINE TEN DAYS OUT

Overview

Europe (EU/UK) is the largest single region in the 2026 Data Security and Compliance Risk survey: 56% of the 459 respondents, n=258, more than North America (32%, n=147) and the Middle East and Africa combined (12%, n=54). What European organizations report is a specific, identifiable profile: security infrastructure that outperforms the rest of the world, paired with AI governance that lags it. That combination, not budget or headcount, is the story of Europe's exposure heading into the second half of 2026.

The three composite scores that anchor this report, computed directly from survey responses rather than self-reported estimates:

Data Security Maturity Score (DSMS)

an index from 0 to 100 built from 11 binary security controls. Europe's mean: **40/100**, ahead of the global survey mean of 39 and ahead of North America (38) and MEA (37).

AI Governance Maturity Score (AIGMS)

an index from 0 to 100 built from 19 binary AI data governance capabilities. Europe's mean: **33/100**, behind the global survey mean of 35 and behind North America (39) and MEA (34), the lowest AIGMS of the three regions.

Data Security and Compliance Readiness Index (DSCRI)

DSMS × (AIGMS/100). Europe's mean: **15/100**, below the global survey mean of 16.2 and below North America's regional mean of 19, despite Europe's stronger security foundation.

That last figure is the one to sit with. Europe has the best general security infrastructure of any region in the survey and still lands below the global average on combined readiness, because the multiplicative structure of the DSCRI punishes exactly the imbalance Europe carries: strong DSMS, weak AIGMS. A stronger lock on a door nobody taught the AI system it needed permission to open does not close the gap.

Not a Forecast, a Reported Outcome

The 2026 Data Security and Compliance Risk survey is the accountability document for the 2026 Data Security and Compliance Risk: Forecast Report, published by Kiteworks in December 2025. This survey of 459 respondents, collected in Q2 2026, tests that report's 15 predictions against what organizations actually experienced. Globally, 74% experienced at least one general security incident in the past 12 months, and 80% experienced at least one incident of either general or AI-specific type. In Europe specifically, 75% of organizations experienced an AI-related incident in the past 12 months, just below the Middle East's 78% and close to the global pattern. 63% experienced a compliance consequence globally; Europe's rate is the lowest of any region at 57%, versus 69% in North America and 76% in the Middle East. Whether that lower European rate reflects more mature compliance programs, less enforcement activity, or underreporting is not something the survey data can resolve on its own, and the report is explicit about that limit.

How the Scores Are Built

Both composite scores behind the Europe numbers above are computed from binary yes/no controls, not from confidence or self-assessment, which matters because the survey also finds that self-reported maturity runs well ahead of what organizations can actually demonstrate: 49% of respondents cannot identify their own SIEM or log management platform without consulting a colleague, and 41% cannot identify the audit standard or framework they are audited against.

The DSMS sums 11 controls: automated data classification, DLP policies (both general and technically enforced across channels), approved-channel restrictions enforced at the platform level, SIEM integration for sensitive data exchange, board-level AI data governance reporting, contractually and technically enforced third-party AI vendor controls, AI-specific DLP, AI access logs forwarded to SIEM, an AI kill switch, and continuously generated compliance evidence.

DSMS Tier	Index Range	Controls Deployed	% of Global Sample
Tier 1, Nascent	0–18	0–2 of 11	24%
Tier 2, Developing	27–45	3–5 of 11	46%
Tier 3, Established	55–73	6–8 of 11	23%
Tier 4, Advanced	82–100	9–11 of 11	7%

Europe's mean DSMS of 40 places the region solidly in Tier 2, at the upper end, just short of Tier 3.

The AIGMS sums 19 capabilities across four dimensions: technical data containment (kill switch, SIEM integration for AI logs, zero-trust access controls, network isolation, single control point for access revocation), data audit and traceability (complete audit log coverage, output-to-source traceability, one-business-day audit records, formal detection process, rollback capability, board-level reporting), data policy enforcement (AI-specific DLP, purpose binding, human-in-the-loop review, data minimization), and organizational data governance (dedicated ownership, documented responsibilities, contractually enforced third-party controls, externally validated compliance).

AIGMS Tier	Index Range	Capabilities Deployed	% of Global Sample
Tier 1, Nascent	0–26	0–5 of 19	48%
Tier 2, Developing	27–44	6–8 of 19	23%
Tier 3, Established	45–72	9–13 of 19	20%
Tier 4, Advanced	73–100	14–19 of 19	9%

Europe's mean AIGMS of 33 sits in Tier 1, Nascent, just below the Tier 2 threshold. Read the two tiers together and Europe's regional story is arithmetic before it is anything else: a Tier 2 security foundation carrying a Tier 1 AI governance layer.

Europe's Regulatory Environment

Europe's compliance burden is not hypothetical, and it is not evenly spread across the calendar. EU/UK respondents rank AI-specific regulatory requirements as their top compliance challenge at a higher rate than any other region: 29% place it first, versus 24% in North America and just 7% in the Middle East. That gap reflects a concrete enforcement environment built by GDPR, the EU AI Act, and NIS2 notification timelines, none of which are optional and none of which are new to European compliance teams. But ranking a challenge first is not the same as having closed it, and the survey's control-adoption data shows the distance between the two.

Across the full survey sample, the compliance frameworks organizations currently cite as primary drivers:

Framework	% Citing as Primary Driver
ISO 27001:2022	42%
GDPR	42%
EU AI Act	41%
CCPA/CPRA	26%
HIPAA/HITECH	21%
NIS2	19%
DORA	16%
SOX	15%
PCI DSS 4.0	13%

GDPR and the EU AI Act sit at the top of that list for the whole sample, and given that EU/UK organizations make up more than half the respondents, they are the main reason those two frameworks rank where they do. The timing matters here in a way it did not five months ago. The EU AI Act's transparency, logging, and human oversight obligations for high-risk systems (Articles 12–14) become enforceable on August 2, 2026, which at the time this survey was fielded was weeks away, not years. DORA's operational resilience obligations for financial entities already require audit-quality evidence on short notice. ISO/IEC 42001 compliance is required today for 40% of respondents. None of these are future-tense problems for a European compliance team.

Set against that calendar: 50% of organizations globally cannot produce a complete AI data access audit record within one business day, and 83% cannot produce one within one hour. The survey does not report this figure broken out by region, but there is no reason to expect European organizations sit meaningfully outside that global distribution, and Europe's own AIGMS score (33, the lowest of the three regions) argues the opposite. An organization that ranks AI regulatory risk as its top compliance concern and cannot produce an audit record on the timeline its own regulators require is carrying exposure that a compliance ranking does not fix.

The Governance Gap, in European Numbers

Europe's DSMS of 40 and AIGMS of 33 describe two different trend lines, and the gap between them is the central fact of Europe's risk profile. The DSCRI by region:

Region	Mean DSCRI
North America	19%
Europe	15%
Middle East and Africa	15%

Note that Europe and MEA post nearly identical DSCRI scores despite very different underlying profiles: MEA arrives there through low scores on both dimensions (DSMS 35.7, the lowest of any region), while Europe arrives there with the best DSMS of the three regions dragged down by the weakest AIGMS. Same destination, opposite roads. That distinction matters for what each region should do next: MEA needs to build security infrastructure and governance together; Europe already has the infrastructure and needs to extend it into AI-specific governance.

The four-quadrant framework the report uses to map organizations, Resilient (strong security, strong AI governance), Fortified (strong security, weak AI governance), Policy-Led (weak security, strong AI governance framing), and Exposed (weak on both), gives Europe's specific signature:

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

(Figures are representation indices: an organization's rate of appearing in a quadrant relative to what its regional share of the sample would predict. 1.00 is proportional.)

Europe's dominant pattern is Fortified, with a secondary Policy-Led signature. Strong security infrastructure and regulatory awareness are both present; AI governance maturity consistently lags both. Europe is underrepresented in the leadership (Resilient) quadrant, at roughly 0.85 to 0.94 times the expected rate depending on which of the report's two analytical thresholds is applied, and close to proportional in the Exposed quadrant. That is a meaningfully different profile than North America, which overrepresents in Resilient at 1.32 to 1.45 times, or the Middle East, which underrepresents in Resilient even more sharply (0.49 to 0.80 times) while overrepresenting in Exposed.

Europe is also the largest single cluster of respondents in the report's DSMS-versus-AIGMS scatter analysis (Figure 38 in the full report), and its visual mass sits concentrated in the Exposed quadrant, specifically the lower-left and lower-center cells of the chart. The report's own reading of that pattern: "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." That sentence is the single most important line in the report for a European reader.

Shadow AI in Europe: A Figure Worth Flagging

The report's shadow AI findings for Europe are the one area where the source data does not agree with itself, and it is worth being direct about that rather than picking a number silently. Three different sections of the full report state three different figures for EU/UK shadow AI discovery:

- The Human Element section states overall (any-frequency) discovery for EU/UK at **34% combined.**
- The same section, one paragraph later, states EU/UK discovery, attributed to "stronger regulatory pressure and audit requirements," at **46% combined.**
- The Regional Variations cross-analysis states overall shadow AI discovery for Europe at **57%**, with the Middle East highest at 76% and North America at 73%.

What is consistent across all three citations, and therefore reliable regardless of which overall figure is correct: EU/UK organizations discover shadow AI monthly or more at 11%, well below North America's 24%. Europe's shadow AI problem, whatever its exact overall discovery rate, is less frequent than North America's but not absent, and the region's AIGMS score (the lowest of the three regions) suggests Europe is not obviously better positioned to find it than the raw discovery numbers imply on their own. Given that Europe's detection infrastructure is not the strongest in the survey, a lower discovery rate is at least as likely to reflect weaker visibility as it is to reflect less shadow AI activity, a pattern the full report makes explicit for MEA and cautions against reading discovery rate as a clean proxy for prevalence anywhere in the dataset.

The data flowing through unsanctioned AI tools generally (not broken out by region in the survey) is not incidental: among organizations using employee-facing AI chatbots, 36% route customer and client data through them, 33% route IT credentials and access requests, 31% route employee personal and HR data, and 30% route financial data. The governance infrastructure to prevent this, AI-specific DLP, approved-channel enforcement, purpose binding, is deployed at fewer than 30% of organizations globally, and Europe's below-average AIGMS gives no reason to assume European organizations are meaningfully ahead of that baseline.

North America also reports the highest AI tool data exposure incident rate of any region at 36%, compared to 20% in EU/UK. That is one of the few places where Europe's numbers look comparatively better in absolute terms, though the same detection-capability caveat applies: a lower reported exposure rate is consistent with either less exposure or less visibility into it.

The Sector Parallel Worth Knowing

Europe's regional pattern, strong security controls paired with weak AI governance, is not unique to the region. The report finds the same imbalance at the sector level in Technology globally: Technology's DSMS (40.7) sits above the survey mean and its AIGMS (35.3) tracks the mean almost exactly, numbers that look average in aggregate, yet 65% of Technology respondents fall in the Exposed quadrant and 16% sit in Fortified, meaning the sector's infrastructure exists without a matching governance perimeter around AI systems. The report does not cross-tabulate industry against region, so it is not possible to say how much of Europe's regional pattern and Technology's sector pattern reflect the same respondents. But Technology is the single largest industry in the sample (36%) and EU/UK the single largest region (56%), so the two groups necessarily overlap to some degree, and the parallel is worth noting for that reason alone.

Key Findings Relevant to Europe

1 AI Governance and Containment

No AI containment control measured in this survey, globally, is deployed by more than 35% of organizations. Purpose binding, the technical restriction of AI agents to authorized tasks and data scopes, is deployed by 26% of organizations (74% do not have it). AI kill switch capability is deployed by 21% (79% have no automated mechanism to stop a misbehaving AI agent). These are global figures; the survey does not report a Europe-specific breakout for containment controls. But Europe's AIGMS of 33, the lowest of the three regions and eight points below North America's 39, is built from exactly these capabilities, so there is no basis in the data for assuming Europe sits meaningfully above the global containment baseline. If anything, the AIGMS gap argues Europe sits at or below it.

The Forecast report projected purpose binding gaps at 63% and kill switch gaps at 60% globally, predicting these would narrow through 2026. The Annual Survey found purpose binding absent at 74% and kill switch absent at 79%, globally. The gap widened everywhere, and nothing in the regional data suggests Europe was the exception.

2 Sensitive Data Exchange and Fragmentation

54% of organizations globally use 4 or more separate platforms for sensitive data exchange, and 73% have no technical enforcement restricting which of those channels employees can use. The channel inventory spans cloud file sharing (78% of organizations), email (73%), managed file transfer (55%), APIs (48%), web forms (43%), and physical media (29%). Only 4% operate a single unified platform. The survey does not report this fragmentation figure by region, but it does report that only 43% of organizations globally have achieved centralized AI gateway control (against a Forecast projection that only 57% would remain non-centralized), and only 39% can enforce tagging and classification consistently across channels. Neither figure is one where Europe's stronger-than-average DSMS gives grounds for assuming the region is meaningfully ahead, since classification and channel enforcement are DSMS components Europe has only partially deployed to reach its Tier 2 score.

Among the 253 MFT-using organizations in the survey, only 26% stream logs to SIEM in real time and 39% have tamper-evident audit trails, corroborating the separate Kiteworks 2025 MFT Survey Report finding that 63% of MFT users have not integrated their MFT infrastructure with SIEM at all. Given that MFT is one of the primary channels European organizations use to move regulated data internationally under GDPR and sectoral rules, an unintegrated MFT deployment is also a gap in the audit evidence a European compliance team would need to produce on request.

3 Detection and Response

574% of organizations globally experienced at least one general security incident in the past 12 months, and 80% experienced at least one incident of either general or AI-specific type. The detection gap behind that number is structural: only 37% have real-time alerting for AI data access anomalies, only 33% forward AI access logs to a SIEM, and only 29% generate any logs of AI system data access at all. 25% have no formal detection process whatsoever, meaning one in four organizations globally has no systematic way to notice when an AI system accesses data it should not.

The report's most operationally significant cross-cutting finding applies directly to Europe's specific imbalance: low AIGMS predicts AI incident rates regardless of DSMS tier. An organization with a strong security foundation and weak AI governance, Europe's regional profile, does not get a discount on AI-specific incident exposure for having good general security. The AI incident rate drops from 91% in DSMS Tier 1 organizations to 34% in Tier 4, but that reduction tracks the AI-specific controls the AIGMS measures, not general security maturity on its own. Europe's own AI incident rate of 75% in the past 12 months is consistent with a region whose AIGMS (33) has not caught up to its DSMS (40).

4 Compliance and Audit Readiness

This is where Europe's regulatory environment and its operational reality are furthest apart. 61% of all respondents rank AI-specific regulatory requirements in their top 3 compliance challenges, and EU/UK respondents drive a disproportionate share of that ranking given their sample weight and their 29% first-place rate. Yet 50% of organizations globally cannot produce a complete AI data access audit record within one business day, 67% lack tamper-evident audit trails, and only 40% apply a consistent evidence approach across all required frameworks. Europe's own compliance consequence rate, 57%, the lowest of any region, is the one data point that could be read as reassuring. It is also the data point the report explicitly declines to interpret with confidence, since a lower consequence rate is consistent with better compliance, less enforcement, or underreporting, and the survey cannot distinguish among the three.

5 The Human Element

36% of organizations globally cite workforce training on AI data security as a top investment priority, making it the most commonly planned response to shadow AI, a problem the report argues training cannot structurally solve on its own. Fewer than half of organizations that discovered shadow AI subsequently deployed technical controls to prevent recurrence; the most common response was updated policy guidance. Outright organizational bans on AI tool usage dropped from 28% of organizations in 2025 to 7% in 2026 globally (Cisco's 2026 Data and Privacy Benchmark Study), without a corresponding increase in technical controls replacing those bans. Europe's regulatory intensity has not, on this evidence, translated into a materially different governance posture than the rest of the survey.

Where Global Investment Priorities Point

The survey asked all 459 organizations, including Europe's 258, to name their top investment priorities for the next 12 months:

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

Data sovereignty and localization sit lower on this list, at 25%, than the rank EU/UK respondents give AI-specific regulatory requirements as a compliance challenge (29% naming it their top single concern). That gap between investment planning and stated regulatory concern is not unique to Europe (the report notes the same pattern globally: 43% name AI security and governance as a top priority, yet only 9% of organizations globally have reached AIGMS Tier 4), but it lands with particular weight in a region where the regulatory clock is already running.

Implications for Europe

Europe's specific exposure is not "insufficient security." It is unconverted security. The region has the security infrastructure to support AI governance controls and has not yet built most of them. That is a narrower, more solvable problem than the one facing MEA (which needs to build both dimensions from a lower base) or the parts of the Technology sector globally that show the same Fortified pattern Europe does at a sector level.

Three implications follow directly from Europe's numbers:

1

The AIGMS gap, not the DSMS gap, is where European remediation investment should concentrate.

Europe's DSMS (40) already exceeds the global mean. Its AIGMS (33) does not. Given the report's broader finding that AI governance investment returns more readiness improvement per unit of effort than additional security spending once an organization's security foundation is already adequate, and Europe's foundation is comparatively strong, European organizations are close to the point where AIGMS investment produces disproportionate DSCRI gains. The report's language for this generally: an organization with DSMS 70 and AIGMS 20 sits below one with DSMS 40 and AIGMS 45 on the combined index, because the DSCRI's multiplicative structure penalizes single-dimension excellence.

2

The EU AI Act deadline is not a planning horizon anymore.

Articles 12 through 14's transparency, logging, and human oversight obligations for high-risk systems become enforceable August 2, 2026. An organization that cannot produce an AI data access record within one business day, a capability only 27% of organizations hold globally, is not positioned to meet that obligation on the day it lands, regardless of what its policy documents say.

3

A lower European compliance consequence rate should be tested, not trusted.

At 57%, Europe's rate is the lowest of any region, but the report is explicit that this could reflect stronger compliance programs, softer enforcement, or underreporting. European organizations in a position to benchmark their own audit and incident history against this figure should do so before treating it as evidence of relative safety.

Operational Recommendations, Adapted for Europe's Profile

1

Extend existing security infrastructure into AI-specific governance, rather than building AI governance from scratch.

Europe's above-average DSMS means most of the underlying data security infrastructure, classification, SIEM, access controls, already exists. The gap is AI-specific: purpose binding, AI-specific DLP, AI access logs to SIEM, kill switches. A Data Policy Engine that extends existing classification and access-control investment into AI-specific enforcement closes this gap faster than parallel infrastructure builds.

2

Build AI audit trails against the EU AI Act's clock, not a generic best-practice timeline.

50% of organizations cannot currently produce a complete AI data access record within one business day. With Articles 12 through 14 enforceable from August 2, 2026, this is the single highest-priority control for any organization operating under the Act's high-risk system categories.

3

Treat GDPR, the EU AI Act, DORA, and NIS2 as one evidence requirement, not four.

Only 40% of organizations apply a consistent evidence approach across required frameworks. European organizations juggling all four of these regimes simultaneously carry the most to gain from consolidating audit evidence generation into continuous, automated collection rather than framework-by-framework manual assembly.

4

Do not treat a lower compliance consequence rate as a lower risk signal.

Pair the reported 57% consequence rate with an internal audit of actual incident and near-miss history before using it to justify slower AI governance investment.

5

Assign dedicated AI data governance ownership.

39% of organizations globally treat AI data governance as an add-on to an existing role (CISO or CIO) rather than a dedicated responsibility. Given the number of overlapping regulatory regimes a European organization is likely managing simultaneously, fragmented ownership is a more expensive failure mode in Europe than in regions facing a single dominant framework.

Conclusion

Europe enters the second half of 2026 with the best general security infrastructure of any region in this survey and the weakest AI governance maturity of the three regions measured. Those two facts, not budget, not headcount, and not the strength of GDPR or the EU AI Act on paper, are what set Europe's risk profile apart. The region has already done the harder, slower work of building security infrastructure. What remains is narrower and more specific: purpose binding, AI-specific DLP, audit trails that can produce a complete record within one business day, and governance ownership that is not an add-on to an existing role.

The EU AI Act's high-risk obligations do not wait for organizations to close that gap on their own schedule. They become enforceable August 2, 2026. European organizations that treat their existing security investment as a foundation to build AI governance on top of, rather than a substitute for it, are the ones positioned to meet that deadline. The ones that do not will be explaining the same gap when this survey runs again next year, this time against a regulator's timeline rather than a survey's.

FOR THE COMPLETE 2026 DATA SECURITY AND COMPLIANCE RISK REPORT WITH FULL METHODOLOGY, GLOBAL COMPARISONS, AND REGIONAL BREAKDOWNS, DOWNLOAD THE REPORT NOW.

DOWNLOAD THE REPORT