Key Takeaways From Kiteworks' Data Security and Compliance Risk: 2026 Forecast Report

2026 Data Security Forecast Findings: The 100% Problem

Kiteworks' Data Security and Compliance Risk: 2026 Forecast Report reveals a fundamental disconnect: Every organization surveyed has agentic AI on their roadmap, yet most lack the controls to govern it. Based on a survey of 225 security, IT, and risk leaders across 10 industries and 8 regions, the research delivers 15 predictions for enterprise data security and exposes the critical gaps that will define which organizations thrive and which learn hard lessons through incident response.

The Core Problem

2026 is the year AI data security moves from "emerging concern" to "operational reality." Organizations created a perfect storm by investing in monitoring while neglecting containment. 63% cannot enforce purpose limitations on AI agents. 60% cannot terminate misbehaving agents quickly. 55% cannot isolate AI systems from broader network access. The governance-containment gap—a 15-20 point divide between watching and stopping—is the central security challenge heading into 2026. Organizations that close it first will be demonstrably more resilient. Those that don't will learn the same lessons the hard way.

100% of organizations have agentic AI on their roadmap. **The majority cannot govern it.**

The 15 Predictions at a Glance

1. DSPM becomes the default baseline	61% can't enforce tagging
2. Data governance goes "managed-by-default"	37% below managed maturity
3. Centralized AI gateways become the control plane	57% non-centralized; 33% of government has no dedicated AI controls
4. Agentic AI goes mainstream	100% on roadmap; only 37-40% have containment
5. Containment controls become the battleground	63% lack purpose binding; 60% lack kill switch

6. AI risks dominate the security agenda	30% cite third-party AI; only 36% have visibility
7. Supply chain expands to AI attestations	72% no SBOM; legacy MFT can't support AI
8. Third-party risk pivots to visibility	89% never practiced IR with partners
9. IR becomes AI-infused	60% lack AI anomaly detection
10. Audit trails become the keystone	33% lack trails; 61% fragmented logs
11. Training-data controls becomeregulatory requirements	78% can't validate; 53% can't recover
12. AI governance hits every boardroom	54% of boards not engaged
13. EU AI Act creates a global template	22-33 point control gap between impacted and non-impacted
14. PQC moves to mainstream	84% haven't implemented
15. Data sovereignty becomes AI imperative	29% cite cross-border AI exposure

Three Critical Gaps

1. The Governance-Containment Disconnect

Organizations invested in watching AI systems but not stopping them. 59% have human-in-the-loop oversight and 58% have continuous monitoring, but only 37% have purpose binding and 40% have kill-switch capabilities. This 15-20 point gap means most organizations can observe AI agents doing something unexpected but cannot prevent them from exceeding authorized scope or quickly shut them down. The forecast: Containment controls become the AI security battleground in 2026. Pipelines are the largest in the survey—39% for purpose binding, 34% for kill switch—but even aggressive execution leaves 24-36% still missing basic containment at year-end.

2. The Audit Trail Foundation Gap

33% lack evidence-quality audit trails entirely, and 61% have fragmented logs scattered across systems. Organizations without audit trails show 20-32 point lower maturity on every AI metric—including training data recovery, human-in-the-loop controls, and purpose binding. The forecast: Audit trails become the keystone capability for AI governance. They predict everything else better than industry, region, or organization size. Fragmented infrastructure cannot support AI governance—you cannot build accountability on archaeology.

3. The Training Data Blindspot

78% cannot validate data before it enters training pipelines. 77% cannot trace where training data came from. 53% cannot recover training data after an incident. The forecast: Training-data controls and "unlearning-ready" architectures become regulatory requirements. The "right to be forgotten" is extending to AI through GDPR, CCPA/CPRA, and the EU AI Act. When regulators ask "how do you know there's no PII in your model?"—78% cannot answer.

Industry Patterns

Government is a generation behind: 90% lack purpose binding, 76% lack kill-switch capabilities, and 33% have no dedicated AI controls—while handling citizen data and critical infrastructure. 71% of government boards aren't engaged on AI governance. The forecast: Government requires transformation, not incremental improvement.

Healthcare faces severe IR gaps: 77% haven't tested RTO/RPO, 64% lack AI anomaly detection, and 68% run manual playbooks despite handling PHI. The forecast: Healthcare's first serious AI incident will expose capabilities that can't survive it.

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

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

Regulatory Trajectory

The EU AI Act is creating a two-tier market. Organizations not impacted are 22-33 points behind on every major AI control: 74% lack AI impact assessments (vs. 41%), 72% lack purpose binding (vs. 46%), 84% haven't conducted AI red-teaming (vs. 61%). The forecast: EU AI Act compliance creates a global governance template whether organizations recognize it or not.

82% of U.S. organizations don't feel pressure yet—but the regulation spreads through supply chains, multinational operations, and competitive benchmarking.

Data sovereignty has expanded from storage to processing. 29% cite cross-border AI transfers as exposure, but only 36% have visibility into where data is processed, trained, or inferred. The forecast: Data sovereignty becomes an AI governance imperative. Knowing where data resides isn't enough—organizations must know where it's processed, trained, and inferred.

Post-quantum cryptography remains largely unaddressed: 84% haven't implemented PQC while "harvest now, decrypt later" attacks are already active. The forecast: PQC moves from early adopter to mainstream, but most organizations are already behind the migration timeline.

The Board Effect

54% of boards don't have AI governance in their top five topics. Those organizations trail 26-28 points on every AI metric. This is the strongest correlation in the survey—stronger than industry, region, or size.

The forecast: AI governance hits every boardroom in 2026. Organizations without board engagement are half as likely to conduct AI impact assessments (24% vs. 52%) and are 26 points behind on purpose binding. When boards don't ask about AI governance, organizations don't build it.

Breaking the Cycle

Organizations avoiding the widest gaps implement four fundamental controls:

- **Close containment gaps first:** Deploy kill switch and purpose binding before incidents force it—60%+ can't terminate or constrain AI agents today
- **Build keystone capabilities:** Evidence-quality audit trails and training-data governance predict everything else (+20-32 point advantages)
- **Get board engagement:** Organizations with board attention lead by 26-28 points on every metric—put AI governance on the agenda
- **Consolidate infrastructure:** Unified data exchange enables audit trails; fragmented systems prevent them—61% are running disaggregated systems that can't support AI governance

Extend sovereignty controls from storage to AI processing. Establish joint IR playbooks with critical vendors—87% lack them, 89% have never practiced. Require third-party AI attestations in contract renewals—questionnaires won't cut it for AI.

Reality Check

This isn't about sophisticated threats. Organizations deploy AI aggressively while leaving fundamental governance gaps exposed. Every missing kill switch, every untraced training dataset, every fragmented audit log compounds risk daily.

The bifurcation is accelerating, and leaders compound advantage while laggards fall further behind. The next wave of AI incidents will come from organizations rushing to deploy without governance infrastructure.

The 15 predictions identify where the market is headed. The gaps identify where you're exposed. The question isn't whether your organization can achieve better AI governance. It's whether you'll act before joining the majority learning these lessons through incident response.

Kiteworks

2026

Data Security and Compliance Risk Forecast Report

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

REPORT

Kiteworks 2025. All Rights Reserved.

For the complete report with detailed methodology, industry breakdowns, and regional analysis, download it now.

Download the Report

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