Kiteworks

SOLUTION BRIEF

FOR CISOS AND COMPLIANCE OFFICERS

WHAT THE NEW FEDERAL AI INCIDENT GUIDANCE ASKS FOR, AND WHERE KITEWORKS FITS

NIST, the NSA, and CISA published ten voluntary practices for AI incident readiness and response. This brief maps each practice to what the Kiteworks platform does, and shows where your own process must carry the load.

The Challenge

The guidance, released October 16, 2026, is voluntary and written for developers of highly capable AI, though its authors say the practices may also apply to deployers. It asks for evidence most enterprises do not keep: what an agent did, what it touched, and whether the log can be trusted. It sets no retention period and does not mention litigation, so your sector rules and legal exposure must fill those gaps.

The Solution

Decide access to sensitive data in one place, and the record follows. The Kiteworks data policy engine evaluates every request from a person or an AI agent, and the platform logs each decision in one audit trail. That covers the data layer of the guidance. It does not cover model weights, training environments, or the decision to disclose publicly.

The Readiness Gap Behind The Guidance

Sources: Schellman 2026 State of AI Governance (525 U.S. professionals); ISACA (1,888 members, not a random sample); Kiteworks Data Security and Compliance Risk: 2026 Annual Survey Report (459 respondents, self-reported).

65% had an AI incident or near-miss in 24 months (Schellman)

44% have documented AI-specific incident response procedures (Schellman)

22% have defined thresholds for human review of agent work (Schellman)

8% regularly run AI-specific incident response exercises (ISACA)

33% have tamper-evident audit trails (Kiteworks 2026 Annual Survey)

The Ten Practices And Where Kiteworks Fits

Ten practices across five stages. Fit is Kiteworks' own assessment of the data layer, not a NIST, NSA, or CISA rating. Strong means the platform enforces or records the practice. Partial means it supplies one piece. Supports means its output feeds the practice.

Practice	How the Platform Fits	Fit
1. Design and Implement Technical Controls PREPARE	The data policy engine applies role-based and attribute-based policies to every request for sensitive data from a person or an AI agent, with least privilege as the default. The Secure MCP Server checks each agent tool call against granted OAuth scopes, and administrators can limit which tools agents see.	Strong
2. Build Organizational Readiness PREPARE	Default administrator roles for compliance, policy, audit, and CISO views support separation of duties between who sets agent data policy and who reviews the evidence. Training, tabletop exercises, and runbooks remain your team's work.	Partial
3. Monitor for Unexpected Activity DETECT	Audit events stream to your SIEM through syslog or Splunk Forwarder, and event subscriptions push selected events to a webhook in real time. The CISO Dashboard can filter activity by AI agent. This watches data access, not model behavior.	Strong

4. Validate and Secure Monitors DETECT	Webhook payloads can carry an HMAC signature so the receiver can verify integrity, failed deliveries are logged as audit events, and the platform runs as a hardened virtual appliance. Testing your own detection rules is still on you.	Partial
5. Collect, Retain, and Analyze Evidence RESPOND	A tamper-evident audit trail records who or what touched which data and what the policy decided. Retention is configurable, and event data is kept indefinitely by default. eDiscovery reports (DLI license required) cover files, activity, and email for a user and date range. Model weights, checkpoints, and sandbox logs are out of scope.	Strong
6. Mitigate and Shut Down Unwanted Activity RESPOND	Administrators can disable destructive tools, restrict which tools agents may use, manage the OAuth application an agent connects through, and revoke tokens to end a session. Rate limiting contains runaway requests. Stopping a model's own runtime is outside the platform.	Partial
7. Verify and Strengthen Containment RECOVER	Policy impact preview lets administrators simulate a change before applying it, and the policy audit report shows whether policies fired as intended. Rebuilding from verified baselines happens in your own infrastructure.	Partial
8. Coordinate With Affected Parties REPORT	Exported audit logs and compliance reports give legal, security, and partner teams one factual record, and secure email and file sharing carry it to affected parties under the same policies.	Supports
9. Notify Government REPORT	Scheduled and on-demand activity reports, plus compliance reports such as CMMC 2.0, GDPR, and HIPAA (some separately licensed), supply the data-access facts a notice needs. What to report, and when, is a decision for your counsel.	Supports
10. Share Information Publicly REPORT	No Kiteworks capability applies. Public disclosure is a communications and legal decision.	None

Agent tool restrictions and token revocation depend on the Kiteworks release. Practice names follow secondary coverage of the guidance.

What Kiteworks Does Not Do

Kiteworks does not preserve model weights, intermediate checkpoints, or training and evaluation environments, and it does not stop a model's own runtime. It does not replace your incident response plan, and no vendor can certify that you meet a voluntary federal document. It addresses the question most enterprises face first, which is what an agent did with the sensitive data it could reach.

See governed AI data access in action

www.kiteworks.com

Kiteworks

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

© 2026 Kiteworks LLC. All rights reserved.