Kiteworks

COMPLIANCE BRIEF

Kiteworks Supports NIST AI Risk Management Framework

Compliant AI and MCP Server Operationalize AI RMF's Governance Functions

Solution Highlights

- Compliant AI
- MCP Server for agentic-AI governance
- Data Policy Engine with risk policy management
- Secure Data Forms
- Immutable audit log with SIEM feeds
- Repositories Gateway for third-party repository governance

The National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF 1.0), released on January 26, 2023, establishes a voluntary framework to help organizations design, develop, deploy, and use artificial intelligence systems in ways that manage AI risks and promote trustworthy AI development. The framework is voluntary, rights-preserving, non-sector-specific, and use-case agnostic, providing flexibility to organizations of all sizes and in all sectors. The framework describes outcomes across four core functions: GOVERN, MAP, MEASURE, and MANAGE. The GOVERN function calls on organizations to establish AI governance structures, policies, and accountability mechanisms. The MAP function describes activities for documenting AI system purposes, categorizing AI capabilities, and characterizing potential impacts. The MEASURE function describes activities for implementing metrics and evaluation processes to assess AI trustworthiness characteristics. The MANAGE function describes activities for prioritizing risks, maximizing benefits, managing third-party dependencies, and implementing post-deployment monitoring. Kiteworks supports organizations working toward compliance with the NIST AI RMF through its secure file sharing and governance platform, which provides protection of unstructured data, advanced data governance and compliance capabilities, and specialized AI risk management and monitoring through Compliant AI and the Kiteworks MCP Server. Here's how:

GOVERN: AI Governance Structures, Policies, and Accountability

The NIST AI RMF GOVERN function calls on organizations to establish AI governance structures, policies, and accountability mechanisms across the full AI system life cycle. Kiteworks addresses the GOVERN function through integrated governance and contingency capabilities. For ongoing monitoring and role definition (GOVERN 1.5), the persistent audit event log captures all file, folder, user, mail, attachment, and comment events with entity resolution and detailed descriptions, and admin activity audit trails provide list, detail, and CSV export capabilities for forensic review and periodic assessments. Compliance Reports from the Data Policy Engine document adherence to regulatory controls based on individual policy configurations, supporting systematic review of data governance outcomes. Admin role management allows organizations to define custom roles with fine-grained per-component permissions, ensuring clear assignment of responsibilities for data governance oversight. For third-party contingency (GOVERN 6.2), high-availability cluster administration performs rolling reboots and OS upgrades with per-node prechecks and postchecks, database cluster failure recovery provides admin-triggered recovery with status tracking, maintenance mode toggles block user traffic during upgrades or incident response, and cross-region replication with consistency checks maintains system resilience when third-party components fail.

MAP: AI System Context, Categorization, and Risk Identification

The MAP function describes activities for documenting AI system purposes, categorizing AI capabilities, characterizing potential impacts, and identifying risk controls. Kiteworks addresses the MAP function through its data access governance and risk policy capabilities. For specifying and documenting application scope (MAP 3.3), attribute-based folder and file permission evaluation combines user, role, parent container, attachment context, and source attributes to produce allow/deny decisions that define which data an AI system is permitted to reach; policy-based feature entitlements gate access and establish clear operational context at the data layer; and content tagging and metadata through REST APIs ensure classification labels are carried forward across file versions and used for ABAC lookups, providing an auditable record of the data scope made available to AI systems. For identifying and documenting internal risk controls (MAP 4.2), the risk policy engine within the Data Policy Engine enables administrators to define risk policies with rules, directives, approvers, tag and form triggers, and automated actions on content; multi-engine antivirus, ATP, and ICAP scanning integrates multiple scan backends with quarantine and unquarantine capabilities; ICAP-based content scanning sends files to third-party ICAP scanners with configurable service URLs, request and response headers, chunked encoding, and result parsing; and Repositories Gateway integrates back-end ECM and storage systems both cloud and on-premises, ensuring comprehensive risk control coverage across all AI system components.

MEASURE: AI Risk Assessment, Metrics, and Transparency Documentation

The MEASURE function describes activities for implementing metrics and evaluation processes to assess AI trustworthiness characteristics. Kiteworks addresses this through layered monitoring, reporting, and feedback capabilities. For risk measurement (MEASURE 1.1, 1.2), the Compliance Report – Audit Log, risk policy reporting dashboard, and CISO risk dashboard with geolocation map surface violations, scan details, and communication flows across the platform. For data-layer validity documentation (MEASURE 2.5), Compliance Reports, asset classification rules, and content tagging with Microsoft Information Protection, Titus, and custom sensitivity labels document the data boundaries within which AI systems operate. For transparency and responsible use governance (MEASURE 2.8, 2.9), the persistent audit event log captures all events with entity resolution, and every file accessed through Compliant AI is identity-verified, ABAC-evaluated, FIPS 140-3 encrypted, and logged in a tamper-evident audit trail. For emergent risk tracking (MEASURE 3.1, 3.2), the CISO Dashboard and Splunk App surface anomalies across all data exchange channels, while Compliant AI enforces consistent data-layer controls regardless of model or agent framework. For feedback and expert consultation (MEASURE 3.3, 4.1), Secure Data Forms, web forms, and commenting workflows capture input from end-users, affected communities, and domain experts, with all submissions documented in the audit log.

MANAGE: AI Risk Response, Third-Party Oversight, and Post-Deployment Governance

The MANAGE function describes activities for risk response, third-party oversight, and post-deployment monitoring. Kiteworks addresses this through data-layer response and governance capabilities. For responding to unknown risks (MANAGE 2.3), the Data Policy Engine applies new risk policies immediately to quarantine or block affected content, maintenance mode controls block user and agent traffic during remediation, and database cluster failure recovery provides admin-triggered recovery with status tracking. For third-party risk monitoring (MANAGE 3.1), Repositories Gateway enables governed access to external ECM repositories within the Kiteworks security and policy framework, with per-source authentication tokens controlling access. For pre-trained model monitoring (MANAGE 3.2), multi-engine antivirus and ATP scanning detects corruption or tampering, file versioning tracks model iterations, and Compliant AI maintains a policy-enforced record of every operation on pre-trained model artifacts. For post-deployment monitoring and incident communication (MANAGE 4.1, 4.3), the MFT Server Operations Console surfaces anomalies in automated data flows, the Data Policy Engine generates incident records fed to SIEM, and email notification templates and advisory contacts management ensure timely, documented communication to all relevant AI actors.

Kiteworks provides organizations with the technical controls needed to implement the NIST AI RMF's outcomes across all four functions: GOVERN, MAP, MEASURE, and MANAGE. The GOVERN function is addressed through persistent audit logging, admin role management, high-availability cluster administration, and cross-region replication for third-party contingency. The MAP function is addressed through attribute-based access governance and content tagging for application scope documentation, and the Data Policy Engine's risk policy engine and multi-engine scanning for internal risk control identification. The MEASURE function is addressed through the CISO risk dashboard, Compliance Report – Audit Log, Compliant AI's tamper-evident data access record, Secure Data Forms feedback intake, and domain expert consultation workflows. The MANAGE function is addressed through Data Policy Engine risk policy updates and maintenance mode controls for incident response, Repositories Gateway for third-party risk monitoring, pre-trained model scanning and versioning, and MFT Operations Console post-deployment governance with email notification and advisory contacts for incident communication. As organizations adopt AI systems across their operations, Kiteworks serves as the unified platform that transforms NIST's voluntary guidance into operational reality through automated controls, continuous monitoring, and comprehensive documentation of AI risk management activities.