ripple treasury
Powered by GTreasury

GSmart AI

# Financial Services AI Risk Management Framework (FS AI RMF) — Compliance Alignment Guide

# Executive Summary

On March 1, 2026, the U.S. Department of the Treasury released the Financial Services AI Risk Management Framework (FS AI RMF), the most concrete federal guidance to date on how AI must be governed, evaluated, and deployed within financial institutions. The framework includes 230 control objectives mapped across the AI lifecycle, adapted specifically for the regulatory and operational realities of financial services.

This guide documents how Ripple Treasury's SmartR AI capabilities align with the FS AI RMF's four core function areas: GOVERN, MAP, MEASURE, and MANAGE. It is designed to give two audiences what they need:

- CFOs and treasury leaders: confidence that the AI powering SmartR Risk Insights and SmartR Forecast Insights was built to meet the governance standards regulators, auditors, and boards now expect.
- IT, security, and compliance teams: a structured, control-by-control reference mapping specific SmartR AI technical controls to FS AI RMF requirements, grounded in Ripple Treasury's AI Security and Compliance documentation.

Ripple Treasury's AI governance posture was built before the FS AI RMF existed. The framework's requirements describe controls SmartR AI already implements -- not a roadmap we are working toward.

ripple treasury
Powered by GTreasury

## About the FS AI RMF

The FS AI RMF was developed through the Artificial Intelligence Executive Oversight Group (AIEOG), a public-private partnership coordinated by the U.S. Treasury, the Financial Services Sector Coordinating Council (FSSCC), and the Cyber Risk Institute (CRI). It adapts the NIST AI Risk Management Framework for the specific context of financial services institutions.

The framework is organized around four function areas:

| Function | Focus | What It Requires |
|---|---|---|
| **GOVERN** | Policies & accountability | AI governance frameworks, regulatory compliance, human oversight, and third-party risk management |
| **MAP** | Risk identification | AI use case classification, data provenance, and tenant data isolation |
| **MEASURE** | Monitoring & audit | Audit logging, model monitoring, drift detection, anomaly detection, and third-party validation |
| **MANAGE** | Controls & response | Data lifecycle, access control, encryption, network segmentation, and restrictions on model training with client data |

While the FS AI RMF is currently voluntary guidance, it is expected to shape examiner expectations, auditor standards, and procurement evaluation criteria as AI adoption in financial services accelerates. Financial institutions and their technology vendors that align proactively will be better positioned when regulatory expectations harden.

ripple treasury
Powered by GTreasury

## How to Use This Guide

The compliance mapping table on the following pages is organized by FS AI RMF function area. For each control category, the table documents:

- The FS AI RMF requirement in plain language
- The specific Ripple Treasury control or capability that satisfies it
- The section of Ripple Treasury's AI Security and Compliance Overview that serves as the underlying source
- A status indicator confirming alignment

CFOs and treasury leaders should focus on the GOVERN and MAP sections, which address the questions most likely to arise in board, audit committee, and vendor evaluation conversations. IT and security teams should review all four function areas, particularly MEASURE and MANAGE, which contain the technical control specifics.

*This guide references Ripple Treasury's internal AI Security and Compliance Overview document (formerly GTreasury), which is available upon request for clients and prospects engaged in formal vendor evaluation or security review processes.*

ripple treasury
Powered by GTreasury

# GSmart AI — FS AI RMF Alignment Mapping

The table below covers 16 control categories across all four FS AI RMF function areas.

| Function | Control Category | FS AI RMF Requirement | Ripple Treasury Control | Whitepaper Source | Status |
|---|---|---|---|---|---|
| **GOVERN** | **AI Governance Framework** | *Establish policies, roles, and accountability for AI risk management across the organization.* | Dedicated AI governance framework in place. Proactive risk assessments and privacy impact evaluations conducted before any AI deployment. | *Section 4.3: AI Governance* | ✓ **Aligned** |
| | **Regulatory Compliance Alignment** | *Align AI governance with applicable laws, regulations, and internationally recognized standards.* | Compliant with EU AI Act (limited-risk classification), ISO/IEC 42001, ISO/IEC 27001, GDPR (EU clients), and CCPA (U.S. clients). | *Section 5: Compliance Alignment* | ✓ **Aligned** |
| | **Human Oversight & Accountability** | *Ensure human oversight of AI decisions, particularly for high-stakes or sensitive outputs.* | Human-in-the-Loop (HITL) checkpoints embedded across all agentic AI workflows. Critical decision-support outputs require confirmation by authorized personnel. Adjustable HITL thresholds based on risk level and context. | *Section 2.1: Human-in-the-Loop (HITL) Oversight* | ✓ **Aligned** |
| | **Third-Party & Vendor AI Risk** | *Manage risks introduced by third-party AI models and services.* | Only enterprise-grade licensed models from Microsoft Azure and Amazon AWS are used. Strict execution boundaries ensure client data never leaves Ripple Treasury's controlled environments. No reliance on public AI endpoints. | *Section 4.2: Infrastructure and Network Security* | ✓ **Aligned** |
| **MAP** | **AI Use Case Identification & Risk Classification** | *Identify AI use cases and classify them by risk level before deployment.* | All AI use cases classified according to the EU AI Act's risk framework. Solutions are categorized as limited-risk applications within financial services. Risk assessments conducted pre-deployment. | *Section 4.3: AI Governance* | ✓ **Aligned** |

| | | | | | |
|---|---|---|---|---|---|
| | **Data Provenance & Integrity** | *Document data sources, lineage, and quality controls for AI inputs.* | Every AI-generated response is linked back to specific source documents within the client's own dataset. Each action is recorded as a trace with rich metadata: session ID, durations, token usage, model used, and filters applied. | *Section 6: Transparency and Client Assurance* | ✓ **Aligned** |
| | **Tenant Data Isolation** | *Prevent cross-contamination of data between organizational units or clients in shared environments.* | Complete data isolation enforced at the application, storage, and AI indexing layers. No co-processing or data sharing across tenants. Client-specific indexes are fully segregated. | *Sections 3 & 4.1: Responsible Use of Client Data / Data Security* | ✓ **Aligned** |
| **MEASURE** | **Auditability & Logging** | *Maintain comprehensive logs of AI interactions and decisions to support audit and review.* | All agentic actions include full audit logging. Access to client data and AI components is logged and monitored by the Security Operations Center. Response traceability is available at the session level. | *Sections 2 & 4.1: Agentic AI / Data Security* | ✓ **Aligned** |
| | **Model Behavior Monitoring** | *Continuously monitor AI model performance, drift, and known limitations including hallucination risk.* | Known limitations and safeguards of LLMs, including hallucination risk, are documented and actively monitored. Ongoing monitoring is in place to control and improve AI service efficiency. | *Section 6: Transparency and Client Assurance* | ✓ **Aligned** |
| | **Security Monitoring & Anomaly Detection** | *Monitor AI infrastructure for security events, intrusions, and anomalous behavior.* | Automated intrusion detection, anomaly detection, and continuous monitoring across the full AI stack and core platform infrastructure. Security Operations Center monitors logs in real time. | *Section 4.2: Infrastructure and Network Security* | ✓ **Aligned** |
| | **Independent Assessment** | *Engage third-party validation of AI practices to support compliance and stakeholder assurance.* | Third-party assessments may be periodically engaged to validate AI practices and demonstrate transparency and compliance to stakeholders. | *Section 6: Transparency and Client Assurance* | ✓ **Aligned** |
| **MANAGE** | **Data Lifecycle & Retention Controls** | *Implement controls governing data retention, deletion, and reindexing aligned to contractual obligations.* | Client data and AI indexes retained only for the duration specified in the client contract and data processing agreement. Data lifecycle controls and reindexing policies are implemented. | *Sections 4.1 & 6: Data Security / Transparency* | ✓ **Aligned** |

| | | | | | |
|---|---|---|---|---|---|
| | **Access Control & Zero Trust** | *Enforce least-privilege access principles and strong identity validation for AI components.* | Zero trust access principles applied to all AI components. Role-based access control (RBAC) enforced across all components. Multi-factor authentication (MFA) mandatory for all privileged access. Conditional access policies restrict access to AI data and models. | *Sections 4.1 & 4.2: Data Security / Infrastructure Security* | ✓ **Aligned** |
| | **Encryption & Data-in-Transit Security** | *Protect AI data at rest and in transit using validated cryptographic standards.* | Data in transit encrypted using TLS 1.2+. Data at rest encrypted with AES-256 and FIPS 140-2 validated cryptographic modules. Applied consistently across all AI infrastructure layers. | *Sections 4.1 & 4.2: Data Security / Infrastructure Security* | ✓ **Aligned** |
| | **No Training on Client Data** | *Prevent client data from being used to train or retune AI models without explicit authorization.* | Guaranteed: client data is never used to train Generative AI models. AI models operate in inference-only mode. AI-powered features are opt-in and protected by feature flags, ensuring no client data is passed to AI unless explicitly triggered by the client. | *Sections 3 & 6: Responsible Use of Client Data / Transparency* | ✓ **Aligned** |
| | **Network Segmentation** | *Deploy AI components within segmented network zones to limit lateral movement risk.* | All key components, including APIs, serverless functions, and storage, are deployed in segmented network zones. Only explicitly authorized communication paths are permitted between layers. | *Section 4.2: Infrastructure and Network Security* | ✓ **Aligned** |

ripple treasury
Powered by GTreasury

# Key Differentiators for Compliance-Sensitive Evaluations

When treasury and compliance teams evaluate AI-powered platforms against the FS AI RMF, three Ripple Treasury controls consistently differentiate GSmart AI from alternatives in the market:

### 1. No Training on Client Data — Guaranteed

Many AI-powered platforms offer vague commitments around data privacy. Ripple Treasury makes a categorical guarantee: client data is never used to train generative AI models. AI operates in inference-only mode. This eliminates one of the most significant data governance risks the FS AI RMF addresses and is directly verifiable through contractual and technical documentation.

### 2. Full Response Traceability

Every AI-generated output in the GSmart platform is traceable back to the specific source documents within the client's own dataset. Each action is logged with rich metadata including session ID, model used, token usage, and filters applied. This level of traceability is a direct match to the FS AI RMF's auditability and explainability requirements and supports the kind of board-level and audit committee accountability CFOs require.

### 3. Human-in-the-Loop by Design, Not by Exception

HITL controls in GSmart AI are not an add-on or a fallback. They are embedded into agentic AI workflows by default, with adjustable thresholds based on risk level and context of use. Critical decision-support outputs require confirmation by authorized personnel before execution. This directly satisfies the FS AI RMF's human oversight requirements and reflects how treasury teams actually need to operate.

***For organizations currently conducting FS AI RMF gap assessments or preparing for AI governance reviews, Ripple Treasury's team is available to walk through these controls in detail. Contact your account team or reach out at treasury-sales@ripple.com.***