Persistent

The AI Governance Debt:

Navigating the Financial Services AI Risk Management Framework

A practical guide to understanding the FS AI RMF, its implications for financial institutions and the path to operationalizing AI governance at scale.

Section 01

Executive Summary

Artificial intelligence has become embedded in the operating fabric of financial services. Lending models, fraud-detection engines, claims-processing workflows and client-facing chatbots now run on AI systems that most institutions adopted faster than their governance functions could keep pace with. The result is a widening gap between what organizations have deployed and what they can demonstrably control, explain and defend under supervisory scrutiny.

In February 2026, the U.S. Department of the Treasury, in coordination with the Cyber Risk Institute (CRI) and the Financial Services Sector Coordinating Council (FSSCC), released the Financial Services AI Risk Management Framework (FS AI RMF). 108 financial institutions contributed to its development. Structurally aligned with the NIST AI Risk Management Framework and operationalized through 230 control objectives, this framework provides financial institutions with a concrete, auditable blueprint for governing AI across the enterprise.

This whitepaper examines the framework's structure and its implications for boards, risk committees and technology leaders, and considers where the real implementation challenges lie for institutions that will need to translate 230 control objectives into operating reality. It also sets out a practical view on how financial institutions can begin operationalizing the framework in a way that strengthens compliance posture and AI capability at the same time.

Section 02

The AI Governance Gap

The banking sector's AI spending is forecast to reach $53 billion in 2026 and grow at a 27% compound annual rate through 2028, according to Statista estimates based on IDC data. Over 80% of U.S. and Canadian banks now deploy AI or machine-learning tools in at least one critical function.

However, a sizable portion of that adoption happened organically. Data science teams have built models on infrastructure that was never designed for AI-scale governance. Business units have procured vendor tools with embedded AI that bypassed centralized risk review. Third-party providers introduced AI capabilities into existing platforms through routine software updates. The cumulative effect is that many financial institutions today cannot produce a complete inventory of the AI systems that are operating within their environment, let alone demonstrate consistent oversight across all of them.

Many financial institutions still run core transaction processing on COBOL systems written decades ago where original developers have retired, documentation is incomplete or interdependencies are only partially mapped. And yet these systems process trillions of dollars daily. AI governance in financial services is converging toward the same condition on a compressed timeline, with decision-making systems whose training data, validation history, authorization chains and operational boundaries are understood by few people. The difference is that COBOL took 40 years to reach this state. AI is getting there in five. The FS AI RMF is the industry's attempt to intervene before the knowledge gap becomes irreversible.

Section 03

Understanding the Financial Services AI Risk Management Framework

The FS AI RMF is a voluntary, industry-led framework tailored for financial services and aligned structurally with the NIST AI RMF. It was developed by the AI Executive Oversight Group (AIEOG), a public-private partnership between Treasury's Financial and Banking Information Infrastructure Committee (FBIIC) and the FSSCC, with the Cyber Risk Institute leading the execution. The 108 financial institutions that contributed included national and multinational banks, community banks, credit unions, insurers, investment firms and payment processors.

The framework is delivered through four interconnected components:

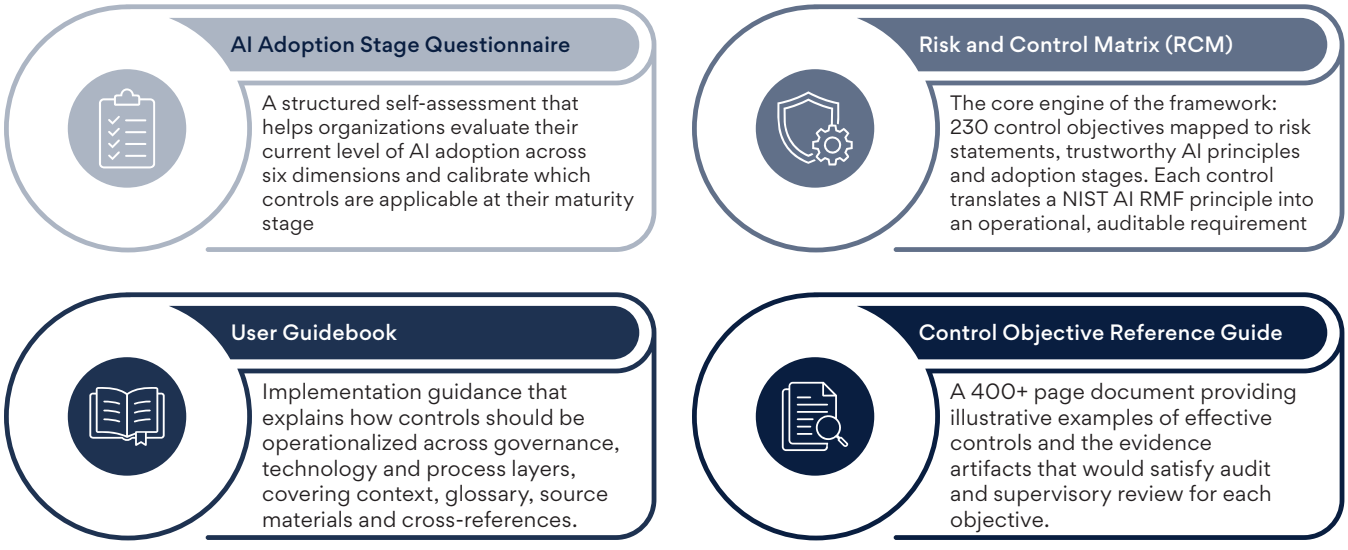

Key Distinction

The FS AI RMF is designed to complement and integrate with existing enterprise risk management programs. It builds on SR 11-7 (model risk management), the FFIEC IT Examination Handbook, OCC third-party risk management guidance (2023-17) and the NIST AI RMF. It augments them as a financial-services translation layer that connects generic AI risk principles to the specific operational realities of running AI in a regulated institution.

Section 04

The Four Functions: Govern, Map, Measure and Manage

The 230 control objectives are organized around four NIST-aligned functions. Together, they form an iterative AI risk management lifecycle, with Govern providing the foundation for accountability and Map, Measure and Manage translating that accountability into operational controls.

Govern

Govern establishes the organizational infrastructure for AI oversight: board accountability, AI strategy, risk appetite, roles, decision rights and governance committees. It answers the leadership question every institution must resolve first: who is accountable for AI risk, and how is that accountability enforced?

Map

Map focuses on context and risk identification. It requires institutions to maintain a comprehensive AI use-case inventory, document data lineage, categorize risks, assess stakeholder impact and bring vendor-embedded AI into the same governance perimeter as internally developed models.

Measure

Measure evaluates whether AI systems meet defined standards and remain within acceptable risk thresholds. It covers validation, bias and fairness testing, performance metrics, explainability and security testing, extending traditional model risk disciplines to machine learning, generative AI and agentic AI environments.

Manage

Manage brings governance into production. It covers continuous monitoring, model drift, incident response, remediation, escalation and decommissioning so that AI systems remain safe, accurate and aligned with institutional objectives after deployment.

Iterative by Design

These functions are meant to operate continuously and concurrently, not as a sequential checklist. An institution might begin with Govern, build out Map as it discovers AI systems across the enterprise, apply Measure as validation capabilities mature and layer in Manage as production deployments stabilize. Cross-referencing between functions is expected and encouraged.

Section 05

Adoption Model and Risk Domains: Maturity Before Scale

The FS AI RMF, via its staged adoption model, recognizes that institutions operate at different levels of AI maturity and should not all attempt to implement 230 controls at once. Instead, the framework helps organizations assess their current stage, prioritize the right controls and mature incrementally.

Stage	Description	Approx. Controls
Initial	Early-stage AI exploration. Limited deployments with basic governance structures. Foundational controls that every institution deploying AI should have in place regardless of scale.	~21
Developing	Growing AI portfolio. Controls expand to address data integrity, model validation and initial third-party oversight.	~70–90
Advanced	Mature AI operations with established model lifecycle management, embedded monitoring and structured incident response.	~150–180
Enterprise-Scale	AI integrated across the enterprise at scale. Full operationalization of all 230 control objectives including advanced supply-chain governance and systemic risk considerations.	230

The AI Adoption Stage Questionnaire helps institutions determine their current stage and define a target state. The framework recognizes that effective implementation begins with foundational controls and maturity-based progression rather than attempting all 230 controls simultaneously.

Seven Risk Domains

The 230 control objectives span seven risk domains that reflect the realities of AI in financial services:

1	**Governance and Accountability**	Board-level oversight structures, AI governance committees, risk appetite statements specific to AI, clear lines of accountability and policies governing AI use, procurement and decommissioning.
2	**Data Integrity and Privacy**	Data quality, lineage, provenance, consent, privacy management and bias detection across training and inference data.
3	**Model Development and Validation**	Model documentation, validation, testing, performance benchmarks, version control and lifecycle governance.
4	**Cybersecurity and Resilience**	Adversarial testing, model extraction and inversion protections, data poisoning defenses and AI infrastructure resilience.
5	**Third-Party and Supply Chain Risk**	Due diligence requirements for AI embedded in vendor systems, contractual provisions for AI transparency and auditability, supply-chain risk assessment for foundation models and pre-trained components and ongoing monitoring of third-party AI.
6	**Consumer Protection and Fairness**	Bias testing, explainability, disclosure obligations and complaint handling for AI-driven outcomes.
7	**Monitoring and Incident Response**	Continuous performance monitoring, model drift detection, alert thresholds and escalation procedures, incident response protocols for AI failures, remediation workflows and model decommissioning criteria. This domain ensures that AI governance extends beyond deployment into the full operational lifecycle.

Work with Persistent to accelerate your FS AI RMF readiness

Persistent brings a purpose-built AI Governance and Controls Assessment Dashboard that helps financial institutions rapidly evaluate their current AI adoption stage, benchmark control maturity against the FS AI RMF, identify coverage gaps across risk domains and prioritize remediation actions. Rather than starting with spreadsheets and manual assessments, institutions can leverage Persistent's accelerator to establish a baseline view of AI governance readiness and build a targeted roadmap for improvement.

Section 06

Why This Matters Now: Getting Ahead of the Findings

The FS AI RMF is voluntary, but the precedent is instructive. When NIST published the Cybersecurity Framework in 2014, it was also voluntary. Within a few years, the FFIEC, OCC, FDIC and NCUA had incorporated CSF concepts into examination procedures, vendor assessment requirements and supervisory letters. "Voluntary" became "expected" and institutions that had not aligned found themselves on the wrong side of examination findings.

State regulators use industry frameworks and guidance as benchmarks when evaluating institutional practices. The FS AI RMF gives them a concrete, industry-endorsed standard against which to evaluate institutional practices, particularly in areas like consumer fairness and data privacy where state regulatory activity has been expanding.

The framework's scope extends beyond institutions themselves to their third-party providers. Technology partners, platform vendors and service providers whose products embed AI sit within the framework's risk perimeter. Financial institutions will expect their vendors to demonstrate alignment with the FS AI RMF, creating a ripple effect across the financial services supply chain.

Comparison with the EU AI Act

While the EU AI Act defines what organizations must comply with (classification of high-risk AI systems, mandatory conformity assessments, prohibited practices), the FS AI RMF provides the how: a detailed, implementation-focused guide that helps institutions operationalize the governance outcomes that regulations like the EU AI Act mandate. For institutions operating across jurisdictions, the FS AI RMF and the EU AI Act are complementary and alignment with the former supports compliance with the latter.

Boards and risk committees should treat the FS AI RMF as an operational architecture standard for AI. The framework functions as a remediation blueprint, one that forces institutions to address the accumulated governance debt in their AI infrastructure before supervisory expectations harden into regulatory criteria.

Section 07

How Persistent Operationalizes the FS AI RMF

Operationalizing the FS AI RMF is hard because the framework exposes gaps that sit inside the engineering and operating model rather than the policy layer, which is where most institutions have concentrated their AI governance work to date. For example, a complete AI inventory sounds simple until an institution must identify vendor-embedded AI, shadow AI adopted by business teams, inherited models and automated decision systems that predate the current wave of machine learning. The same complexity appears in data lineage, identity and access governance, model lifecycle management, third-party contracts and audit evidence.

Persistent Systems brings the BFSI domain depth, AI engineering capability and modernization experience needed to help institutions operationalize the FS AI RMF. To accelerate this journey, Persistent also provides governance assessment accelerators that help institutions establish a baseline view of AI maturity, identify control gaps and prioritize remediation efforts.

- **AI management and planning:** Maturity assessment, governance design, risk appetite definition, committee structures and board-ready reporting.
- **Data foundation and lineage:** Automated data profiling, data quality controls, reconciliation, observability and lineage capabilities that support the Map function.
- **Model lifecycle management:** Secure development pipelines, validation, bias testing, explainability, monitoring and decommissioning workflows.
- **Agentic AI with embedded controls:** Governance-by-design for multi-LLM orchestration, human-in-the-loop processes, Model Context Protocol connectors and audit trails.
- **Third-party AI risk management:** Vendor AI discovery, contract review, transparency requirements and ongoing monitoring aligned to the framework's supply chain expectations.
- **Continuous monitoring and evidence automation:** Drift detection, control testing, remediation workflows and examination-ready evidence repositories across cloud-native data platforms.

A pragmatic implementation roadmap should move in three waves: assess current maturity, build the priority operating capabilities and operationalize monitoring, evidence and continuous improvement.

Phase 1

Assess

(Weeks 1–6)*

- Conduct the AI Adoption Stage Questionnaire to determine current maturity
- Perform enterprise-wide AI inventory including vendor-embedded and shadow AI
- Execute gap analysis against the Risk and Control Matrix at the applicable adoption stage
- Identify quick wins (controls that can be satisfied with existing infrastructure) and critical gaps (controls requiring new capabilities)
- Deliver an FS AI RMF readiness report to the board or risk committee

Phase 2

Build

(Weeks 7–18)*

- Establish or formalize AI governance structures: committees, charters, accountability assignments and risk appetite statements
- Implement data lineage and provenance capabilities across priority AI systems
- Deploy model lifecycle management infrastructure: version control, validation pipelines, performance monitoring
- Operationalize priority control objectives from the gap analysis beginning with the highest-risk domains
- Initiate third-party AI governance: vendor inventory, contract review, governance maturity assessment

Phase 3

Operationalize

(Weeks 19–30)*

- Stand up continuous monitoring and drift detection across production AI systems
- Build audit-trail generation and evidence management to support examination readiness
- Conduct mock supervisory review using the Control Objective Reference Guide as a benchmark
- Extend governance to cover emerging AI modalities (generative AI, agentic AI, multi-model architectures)
- Define the roadmap for progressing to the next adoption stage with control targets and milestones

Conclusion and References

The FS AI RMF is not fundamentally about compliance. It is about resilience. It gives financial institutions a sector-specific, operationally detailed way to build trust, accountability and operational discipline into their AI ecosystem before AI risk becomes a business, regulatory or reputational event. Institutions that succeed will treat the FS AI RMF as an information governance and engineering program, which means building the underlying data infrastructure, model management platforms and evidence repositories that make controls demonstrable at examination time rather than reconstructable after the fact. Persistent Systems helps financial institutions move from framework alignment to operating reality by combining BFSI domain depth with AI engineering and modernization capabilities built for that transition.

References

1. Cyber Risk Institute, "Financial Services AI Risk Management Framework," February 2026. cyberriskinstitute.org/artificial-intelligence-risk-management

2. U.S. Department of the Treasury, "Treasury Releases Two New Resources to Guide AI Use in the Financial Sector," Press Release SB0401, February 19, 2026. home.treasury.gov/news/press-releases/sb0401

3. National Institute of Standards and Technology, "AI Risk Management Framework (AI RMF 1.0)," January 2023. nist.gov/itl/ai-risk-management-framework

4. Financial Services Sector Coordinating Council, "AI Executive Oversight Group Deliverables." fsscc.org/AIEOG-AI-deliverables

About Persistent

Persistent Systems (BSE: 533179 and NSE: PERSISTENT) is a global services and solutions company delivering AI-led, platform-driven Digital Engineering and Enterprise Modernization to businesses across industries. With over 27,500 employees located in 21 countries, the Company is committed to innovation and client success. Persistent offers a comprehensive suite of services, including software engineering, product development, data and analytics, CX transformation, cloud computing, and intelligent automation. The Company is part of the MSCI India Index and is included in key indices of the National Stock Exchange of India, including the Nifty Midcap 50, Nifty IT, and Nifty MidCap Liquid 15, as well as several on the BSE such as the S&P BSE 100 and S&P BSE SENSEX Next 50. Persistent is also a constituent of the Dow Jones Best-in-Class World Index. The Company has achieved carbon neutrality, reinforcing its commitment to sustainability and responsible business practices. Persistent has also been named one of America's Greatest Workplaces for Inclusion & Diversity 2025 by Newsweek and Plant A Insights Group. As a participant of the United Nations Global Compact, the Company is committed to aligning strategies and operations with universal principles on human rights, labor, environment, and anti-corruption, as well as take actions that advance societal goals. With 468% growth in brand value since 2020, Persistent is the fastest-growing IT services brand in 'Brand Finance India 100' 2025 Report.

USA

Persistent Systems, Inc.
1731 Technology Drive
Suite 700, San Jose
CA 95110
Tel: +1 (650) 481 9180
Email: Info@persistent.com

India

Persistent Systems Limited
Bhageerath, 402
Senapati Bapat Road
Pune 411016
Tel: +91 (20) 6703 3000
Fax: +91 (20) 6703 6003

Persistent
Re(AI)magining the World