SIA

sia-partners.com

The Emerging Landscape of AI Regulation in U.S. Financial Services

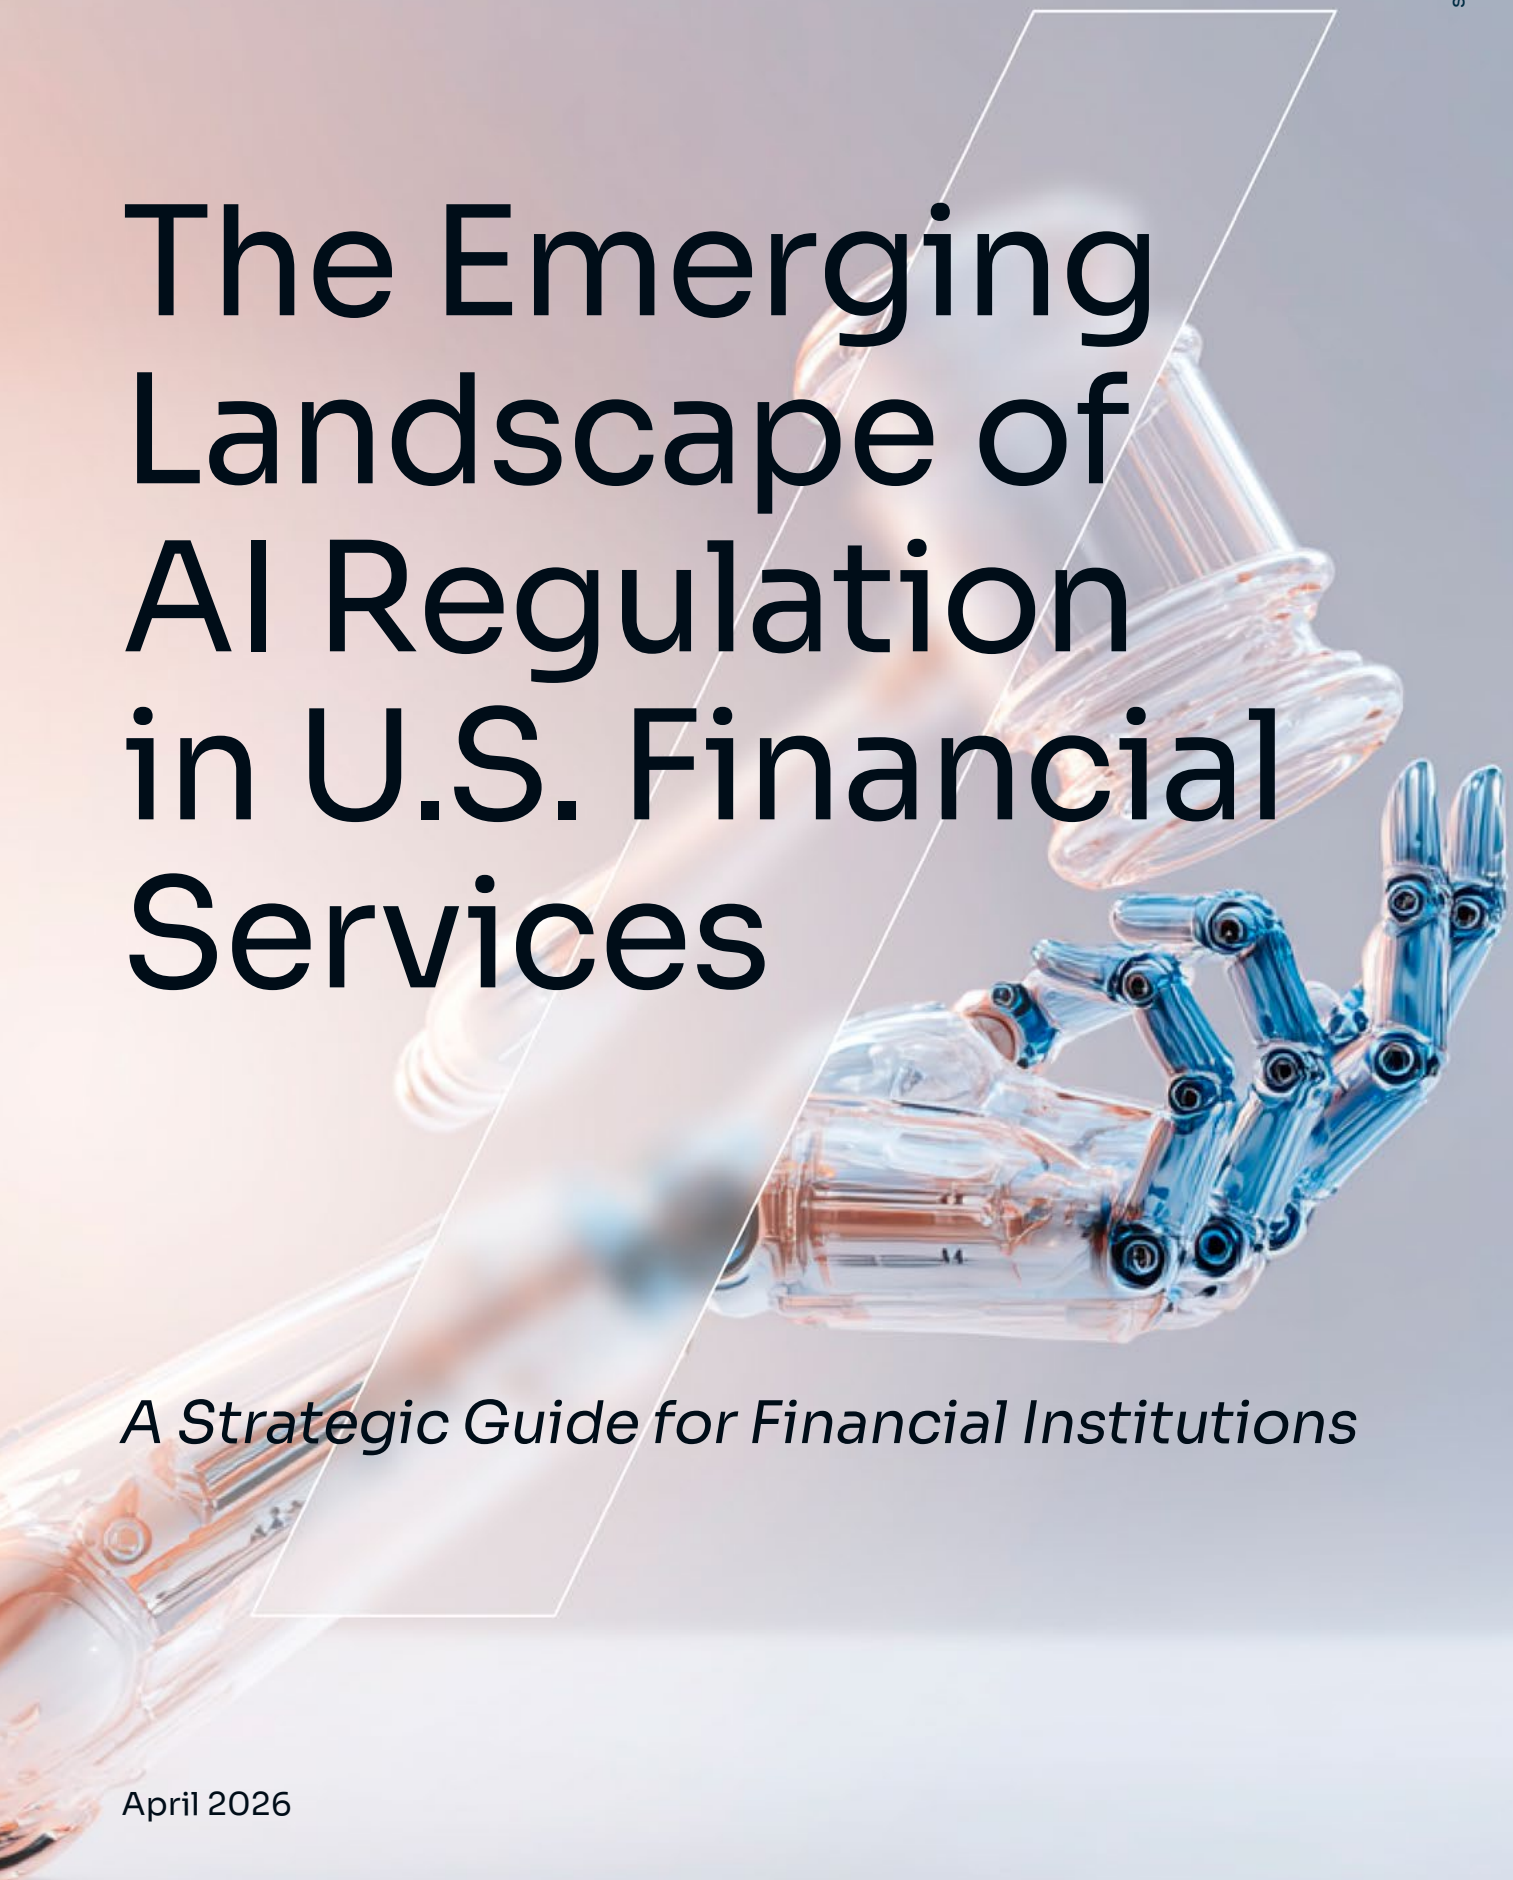

A Strategic Guide for Financial Institutions

April 2026

/ Table of contents

/ Executive Summary

- / Existing regulations already govern AI — SR 26-2, ECOA, BSA/AML, UDAAP, and FINRA rules apply now, regardless of model complexity.
- / The FS AI RMF (February 2026), developed with 108 institutions, is the de facto governance standard and will harden into examination benchmarks by 2028.
- / Four pillars define the U.S. landscape: federal guidance, reinterpreted regulatory authority, industry self-governance, and accelerating state legislation.
- / The EU AI Act — now in phased implementation through August 2026 — sets the international benchmark; U.S. and EU frameworks are converging on core principles.
- / Institutions must act now: build governance, conduct gap analyses, and document controls while the standards are still being shaped.

/ Introduction

The regulation of artificial intelligence in U.S. financial services has reached a defining moment. AI is no longer theoretical, it has become embedded in how financial institutions design products, make credit decisions, run operations, detect fraud, and manage risk. This fundamental shift makes AI governance an immediate executive imperative, even as the regulatory landscape continues to evolve without a single comprehensive federal law.

The United States has charted a fundamentally different course from the European Union's sweeping legislative approach. Rather than pursuing omnibus federal legislation, the U.S. has embraced a public-private partnership model, which was crystallized in February 2026 when the Department of the Treasury, the Financial Services Sector Coordinating Council, and the Cyber Risk Institute released the Financial Services AI Risk Management Framework (FS AI RMF) through the AI Executive Oversight Group (AIEOG). Developed with input from 108 financial institutions, this voluntary framework now stands as the centerpiece of American AI governance in financial services.

The most critical insight many institutions have been slow to internalize is this: existing regulations already apply to AI. Federal regulators possess substantial authority over activities that increasingly involve AI, and they are enforcing these mandates with growing AI-specific rigor. Institutions that wait for comprehensive federal legislation will find themselves significantly behind.

The Current Reality — AI Is Already Regulated

Before examining emerging frameworks, financial institutions must recognize that AI-related scrutiny can arise today under existing regulatory and supervisory frameworks. This is not a future risk, it is a present obligation.

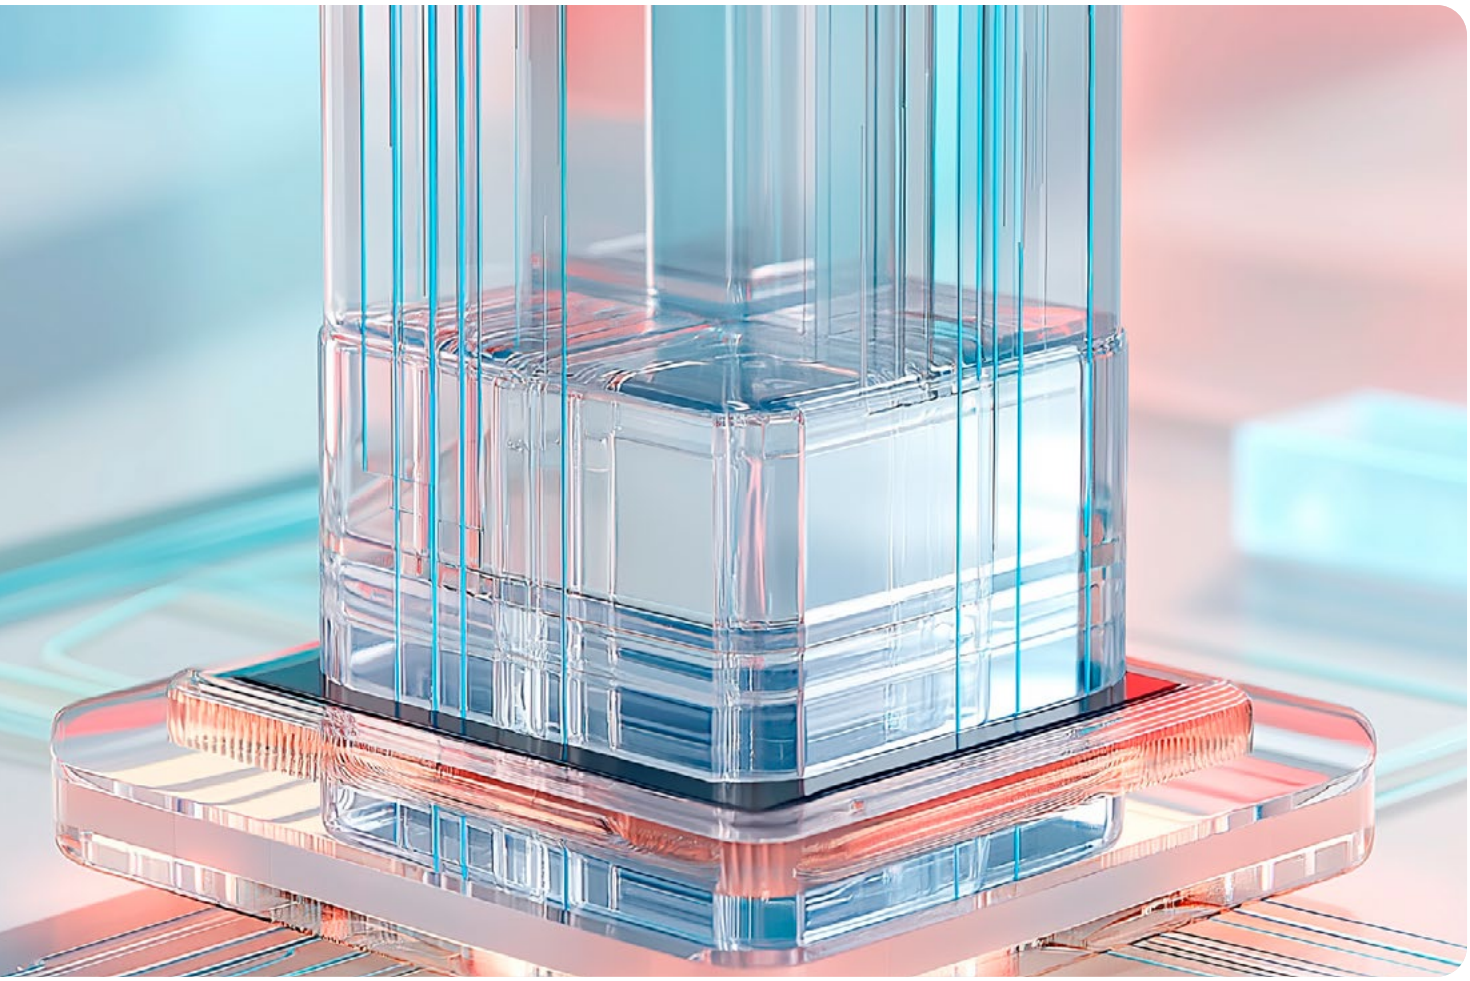

Model Risk Management

SR 26-2 (April 17, 2026) supersedes SR 11-7 and governs AI and ML models used in credit decisioning, fraud detection, and risk management, with oversight calibrated to model materiality . Critically, generative AI and agentic AI are explicitly excluded from SR 26-2's scope — a structural governance gap institutions must address immediately through parallel frameworks drawing on the NIST AI RMF, FS AI RMF (March 1, 2026), and applicable state AI laws . The agencies have signaled a forthcoming RFI on AI governance, indicating rapid regulatory evolution ahead .

Fair Lending and Consumer Protection

The Equal Credit Opportunity Act (ECOA) and the Fair Housing Act demand that institutions test AI credit models for disparate impact and provide meaningful adverse action explanations. The Consumer Financial Protection Bureau (CFPB) has explicitly stated that lenders using AI must still provide specific and accurate reasons for adverse actions — sophisticated technology does not diminish consumer rights or institutional obligations. UDAAP prohibitions extend to AI chatbots that mislead consumers, algorithmic pricing that exploits financial vulnerabilities, and opaque automated decisions that harm retail customers.

BSA/AML, Third-Party Risk, and Securities Supervision

Bank Secrecy Act and Anti-Money Laundering requirements extend to AI-driven transaction monitoring systems. Third-party risk management has become particularly consequential as many institutions deploy AI through vendor relationships: the 2023 interagency guidance makes clear that using a third party does not reduce a financial institution's responsibility to operate safely and in compliance with applicable law. In securities markets, FINRA has reinforced that existing rules apply to generative AI. Member firms are expected to supervise enterprise use, assess integrity and accuracy, and address bias, data quality, and cybersecurity risks.

“Key Takeaway”

/ If AI is influencing a regulated activity, firms must be prepared to explain it, evidence the controls around it, and defend the governance decisions behind it — today, under existing authority.

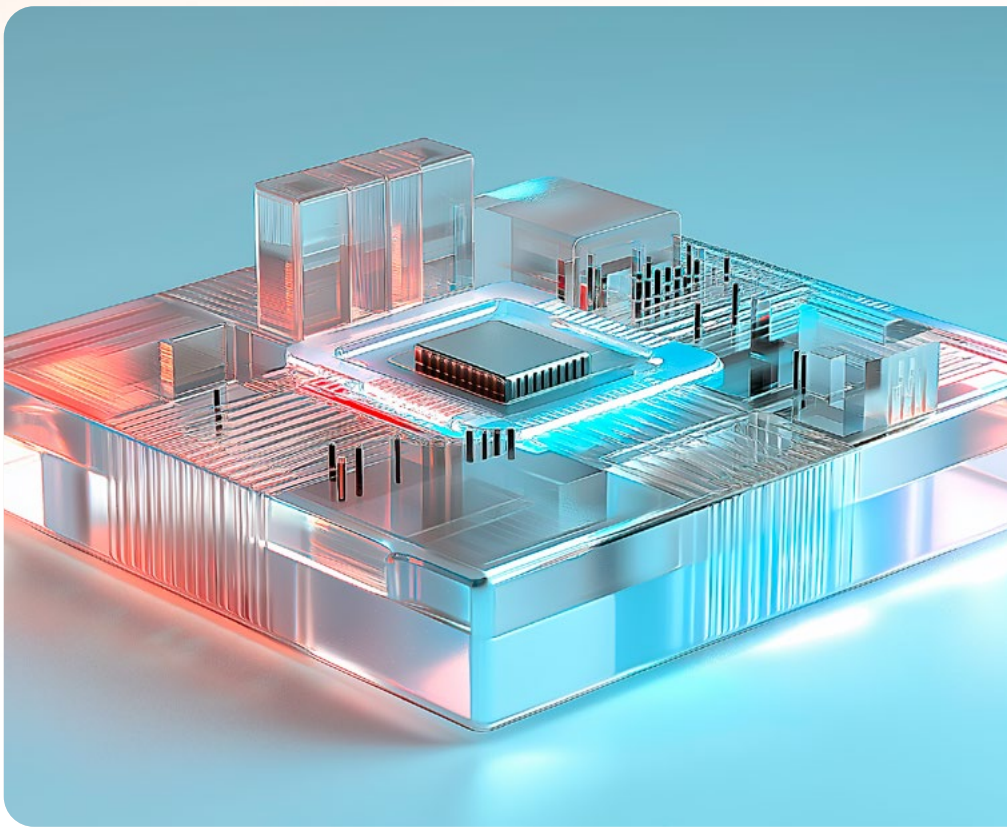

The Emerging Framework — A Layered Governance Model

While existing regulations already apply to AI, the landscape is evolving to address AI-specific risks more directly. What has emerged is a layered governance model resting on four reinforcing pillars.

Pillar 1 — Federal Guidance and Voluntary Frameworks

The White House AI Action Plan, released in July 2025, made the federal posture explicit: prioritize innovation, remove barriers to AI development, and rely on sector-specific expertise rather than horizontal regulation. Within this posture, the FS AI RMF and the NIST AI RMF 1.0 function as de facto standards, mirroring the trajectory of the NIST Cybersecurity Framework, which began as voluntary guidance and became an examination benchmark.

The NIST AI RMF provides the conceptual foundation, organized around four core functions: Govern, Map, Measure, and Manage. The FS AI RMF builds upon this foundation with 230 control objectives, cross-mapping to existing regulatory standards, and an AI Lexicon that establishes common terminology across the industry.

Pillar 2 — Reinterpretation of Existing Regulatory Authority

Federal regulators (the OCC, Federal Reserve, FDIC, CFPB, SEC, and CFTC) already possess substantial authority over activities that involve AI. Rather than waiting for new legislation, these agencies are applying existing mandates with growing AI-specific rigor, including through examination findings, supervisory letters, and public guidance.

Pillar 3 — Industry Standards and Self-Governance

The AIEOG's public-private partnership model (108 institutions collaborating directly with Treasury) represents something genuinely new in financial regulation. It signals that the industry has both the opportunity and the responsibility to shape its own governance future, ahead of formal rulemaking.

Pillar 4 — State-Level Legislation

State legislation creates the most immediate compliance complexity. Key developments include:

- **New York RAISE Act (signed December 19, 2025; subsequently amended via chapter amendment signed March 27, 2026; effective January 1, 2027).** Establishes the nation's first comprehensive AI safety governance regime for frontier model developers with $500M+ in annual revenue. Requires the development and maintenance of a frontier AI framework, 72-hour incident reporting, and safety-related disclosures. *(Note: The original bill's requirement to designate senior compliance personnel as a standalone obligation was removed in the March 27, 2026 chapter amendment.)*
- **Colorado AI Act (effective June 30, 2026):** Focuses on high-risk AI systems and algorithmic discrimination.

- **Utah AI Policy Act:** Imposes disclosure requirements for generative AI in consumer interactions.
- **Texas TRAIGA:** The Texas Responsible Artificial Intelligence Governance Act, enacted 2025 (signed June 22, 2025; effective January 1, 2026).

Over a dozen states now have AI-specific laws. This patchwork will only grow, and may ultimately serve as the catalyst for federal preemptive legislation, though the timing and form of any such federal response remains uncertain and will depend on the pace of state activity, the severity of any AI-related market disruptions, and the evolving priorities of Congress and the Administration.

“Legislative Outlook”

- Comprehensive federal AI legislation specific to financial services is unlikely within the next two to three years. More probable is targeted legislative action on algorithmic bias in lending, AI-generated deepfakes in financial fraud, and AI-related systemic risk — supplementing rather than replacing the layered governance model.

Regulatory Hardening — The FS AI RMF Trajectory

Regulators will use the FS AI RMF in four ways: as an examination benchmark (assessing institutional practices against control objectives), as a gap analysis tool (revealing compliance shortfalls), as a common language (the AI Lexicon becoming the lingua franca of supervisory dialogue), and as a maturity model (with larger, more complex institutions held to higher expectations).

Timeframe	2026
Phase	**Reference Phase**
Key Regulatory Milestones	Regulators cite FS AI RMF in speeches, examination guidance, and supervisory communications
Timeframe	2026–2027
Phase	**Assessment Phase**
Key Regulatory Milestones	Examination teams systematically assess practices against control objectives; MRAs and MRIAs cite specific gaps
Timeframe	2027+
Phase	**Enforcement Phase**
Key Regulatory Milestones	Formal supervisory guidance incorporates FS AI RMF principles; enforcement actions cite alignment failures

The trajectory is consistent with how the NIST Cybersecurity Framework evolved: from voluntary reference to supervisory expectation to enforcement standard. Institutions that wait for formal rulemaking will be significantly behind.

International Context and U.S.–EU Convergence

The European Union offers a critical contrast: it has moved further and faster toward comprehensive AI regulation than the United States. The EU AI Act, which entered into force on August 1, 2024, is a risk-tiered legislative framework that has a phased implementation plan:

- Prohibited AI practices: applied February 2025
- Obligations for general-purpose AI models: applied August 2025
- Full applicability for most high-risk AI systems (Annex III): August 2, 2026. Full applicability for high-risk AI systems embedded in products regulated under existing EU product safety legislation (Annex I/II): August 2, 2027. (The August 2026 deadline does not apply to all high-risk systems — Annex I/II product-embedded systems have a one-year extension.)

High-risk AI systems

Including those used for creditworthiness assessment and credit scoring, are subject to conformity assessments, risk management systems, data governance standards, human oversight requirements, and ongoing monitoring obligations. Non-compliance with high-risk AI system requirements (Articles 6–49) carries penalties of up to €15 million or 3% of global annual turnover, whichever is higher.

The higher penalty tier of €35 million or 7% of global annual turnover applies exclusively to violations of prohibited AI practices under Article 5 (e.g., social scoring, real-time biometric surveillance in public spaces), not to high-risk system non-compliance.

For U.S.-based institutions, EU developments matter because many U.S. firms have European subsidiaries, cross-border data flows, and global vendor relationships that bring them within scope of the EU AI Act. Further, EU requirements are shaping global vendor practices and international standard-setting bodies, such as the Bank for International Settlements Basel Committee. Finally, stakeholder expectations about responsible AI governance are being defined by the EU framework.

U.S.–EU Comparison

Dimension	U.S. Approach	EU Approach
Legislative Model	Sector-specific guidance; voluntary frameworks	Horizontal omnibus legislation (AI Act)
Legal Basis	Existing regulatory authority re-applied to AI	New binding obligations; risk-tiered categories
Enforcement	Examination benchmarks; MRAs / MRIAs; enforcement actions	Fines up to €35M or 7% of global turnover
Compliance Style	Principles-based; flexibility for institutions	Prescriptive requirements; conformity assessments
Timeline	Hardening trajectory: 2026–2028	Full applicability: August 2026
Convergence Points	Risk governance, transparency, human oversight, accountability	Risk governance, transparency, human oversight, accountability

Convergence Implications

Despite stark differences in approach, the two regimes are converging in substance on core principles: risk-based governance, transparency, human oversight, and accountability. The FS AI RMF's 230 control objectives address many of the same risk categories as the EU AI Act — governance structures, data quality, model validation, explainability, incident response, and third-party risk management.

Institutions building robust AI governance programs aligned with U.S. frameworks will be well-positioned to satisfy international requirements; however, jurisdiction-specific modules (e.g., EU conformity assessments and database registration) will still be required. The recommended architecture: identify the common core of requirements across regimes, develop jurisdiction-specific modules, and build governance structures flexible enough to accommodate future regulatory developments in both regions.

A Practical Roadmap for Financial Institutions

Firms should not wait for a final, unified U.S. AI rulebook before acting. The smarter course is to build an operating model for responsible AI governance now that assigns accountability, sets escalation paths, and applies rigorous review to high-risk use cases.

Phase	Priority Actions
Immediate (0–6 Months)	/ Conduct an enterprise-wide AI maturity assessment (models, governance structures, policy gaps, skills, third-party dependencies) / Establish a cross-functional AI Governance Committee with board reporting lines and authority to set AI risk appetite / Adopt the FS AI RMF AI Lexicon as institutional standard terminology / Build a comprehensive AI inventory: all models (including third-party embedded), use case descriptions, risk classifications, data inputs, and validation status
Medium Term (6–12 Months)	/ Conduct a systematic gap analysis against all 230 FS AI RMF control objectives; prioritize remediation by risk severity / Implement comprehensive pre- and post-deployment lifecycle controls: data classification, privacy impact assessment, performance testing, validation, monitoring for drift and bias / Extend third-party AI due diligence beyond procurement: security posture, explainability, audit rights, incident notification, and exit planning / Implement tailored controls for generative AI: hallucination risk, data privacy exposure, IP concerns, prompt injection vulnerabilities / Develop examination-ready documentation: governance charters, risk appetite statements, model inventories, validation reports, fair lending results, and AI incident response plans
Long Term (12–24 Months)	/ Build a regulatory convergence monitoring function tracking federal guidance, state legislation, international developments, and evolving industry standards / Invest in governance infrastructure: AI monitoring platforms, bias detection tools, explainability solutions, and specialized talent / Develop a multi-jurisdictional compliance architecture identifying the common core of requirements across U.S. and international regimes

“Generative AI — Specific Controls Required”

/ Generative AI demands attention beyond traditional model risk frameworks: hallucination risk, data privacy exposure, intellectual property concerns, prompt injection vulnerabilities, and output inconsistency require dedicated control layers. Institutions should develop generative AI policies, usage guidelines, and monitoring processes as a distinct workstream.

The Strategic Imperative

The regulation of AI in U.S. financial services is not a future event, it is a present reality. The FS AI RMF, existing regulatory authority, and accelerating state legislation together create a governance landscape that demands immediate action. Institutions that build strong AI governance now will realize five compounding advantages:

- Regulatory resilience: as voluntary guidance hardens into supervisory expectations, firms with mature programs avoid costly remediation
- Competitive advantage: confident AI deployment with appropriate controls, enabling faster innovation with lower regulatory risk
- Operational excellence: clear governance pathways give innovation teams the frameworks they need to move quickly and responsibly
- Global readiness: modular compliance architectures accommodate both U.S. and international requirements efficiently
- Strategic optionality: the ability to move faster while standards are still being set, shaping supervisory expectations through demonstrated practice

Responsible AI governance is not just a compliance exercise. It is a mechanism for improving execution, building resilience, and preserving strategic optionality. The AIEOG's public-private partnership model signals that the industry has both the opportunity and the responsibility to shape its own governance future.

The institutions that move decisively now will define the standard of practice. Those that wait will find the standard has been defined for them.

/ Conclusion

/ Comprehensive U.S. AI regulation is still emerging, but meaningful expectations already exist and will continue to develop through existing regulatory frameworks, voluntary guidance that hardens into supervisory standards, and state legislation. Financial institutions must act now to build governance capabilities that satisfy current obligations while positioning for future requirements.

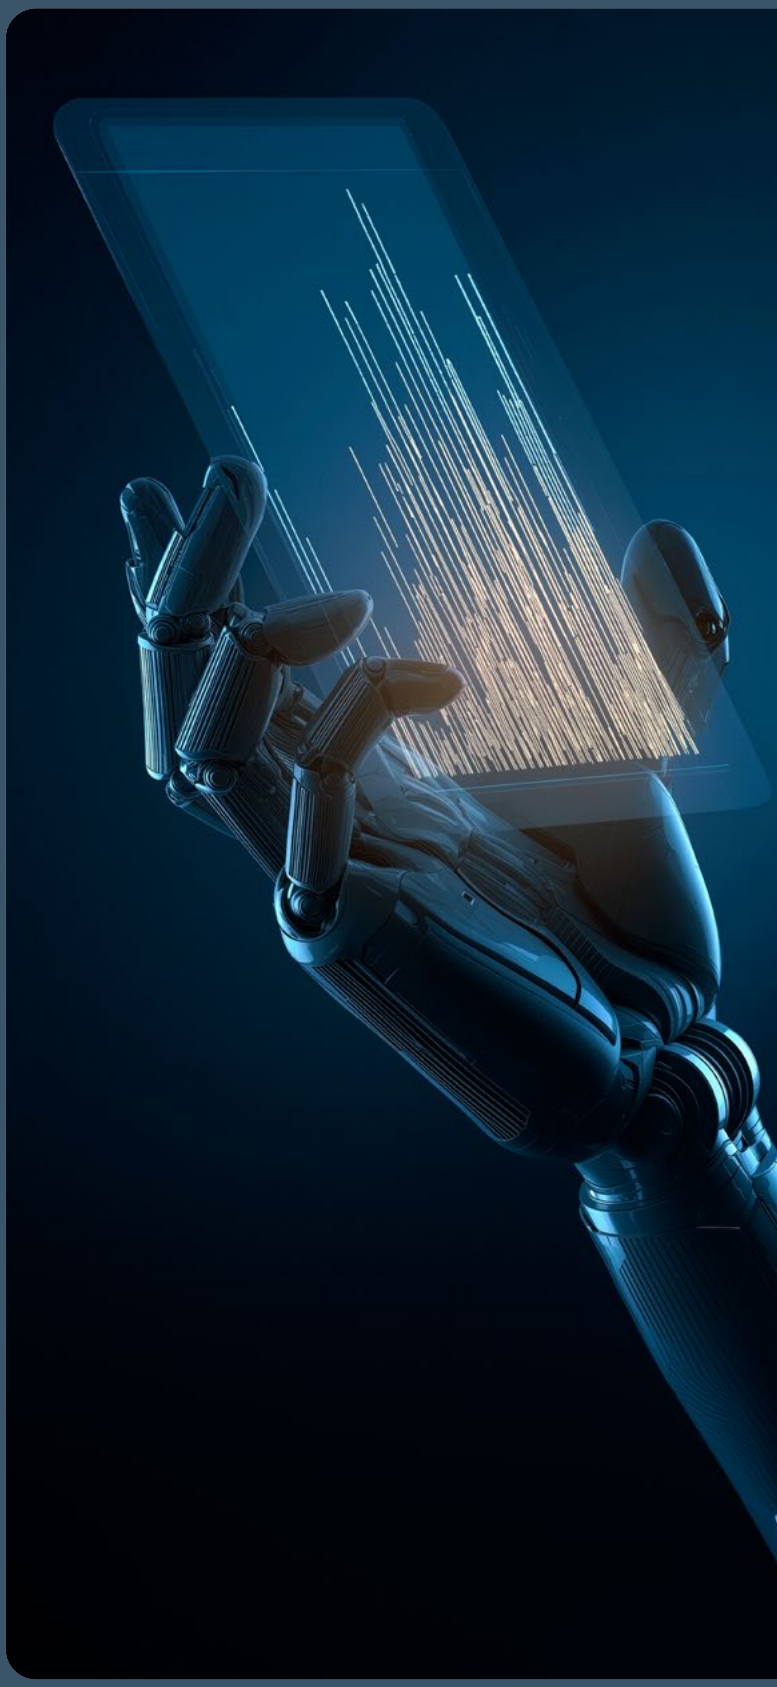

Optimists for change

Sia is a next-generation, global management consulting group—born digital, augmented by data, enhanced by creativity, and driven by responsibility. We partner with clients to resolve challenges and capitalize on opportunities. We believe that in today's world of change and disruption, optimism is a force multiplier.