FINOS

Fintech
Open Source
Foundation

AI GOVERNANCE FRAMEWORK

ABOUT AIGF

A practical, open source AI Governance Framework that banks can adopt, adapt, and extend. It provides a comprehensive collection of risks and mitigations that support on-boarding, development of, and running Generative AI solutions

Learn More:
air-governance-framework.finos.org

The FINOS AI Governance Framework

Open, peer-reviewed AI governance for financial services that is ready to adopt.

A guide for CTO Office · Head of AI · Model Risk Management

The Challenge

Every financial institution deploying AI and Agentic solutions is solving the same governance problem in parallel. Each firm identifies the same risks, satisfies the same second-line functions, and answers to the same regulators, behind closed doors, at full cost, and with only its own blind spots for the company. Competitive advantage in AI will not come from treating governance itself as proprietary. It will come from how quickly, safely, and effectively firms can adopt AI at scale.

What is the AI Governance Framework?

The **FINOS AI Governance Framework** (v2, October 2025) is an open, industry-developed catalog of the risks of deploying generative and agentic AI in financial services, along with practical mitigations to address them. It was developed by practitioners from member banks, asset managers, and technology firms under FINOS (the Fintech Open Source Foundation, part of the Linux Foundation). This coverage includes:

- **11 operational risks**: hallucination and inaccurate outputs, non-determinism, bias, explainability gaps, model overreach, availability, data quality, multi-agent trust, and more
- **9 security risks:** prompt injection, data poisoning, model tampering, vector store leakage, tool manipulation, supply chain compromise, credential harvesting

> **“A GenAI Governance Framework for FSIs by FSIs”**
>
> – Johnna Powell, Managing Director, and Head of Technology, Research and Innovation at DTCC

- **3 regulatory and compliance risks:** information leakage, regulatory oversight gaps, compliance failures
- **2 threat modeled reference architectures**

Each risk includes a detailed description, contributing factors, financial-services-specific examples, and links to research. Mitigations are classified as preventative or detective, and cross-referenced to standards your second line already recognizes: NIST, OWASP Top 10 for LLMs, FFIEC, MITRE ATLAS, and applicable regulation (EU AI Act, GDPR, SEC/FINRA/OCC guidance). A use-case-driven approach is then taken to filter and prioritize risks and mitigations. One of the most important aspects of the AI Governance Framework is its peer review process, through which financial institutions collectively assess and validate relevant AI standards for industry use. This is not just valuable input for those developing, controlling, and auditing AI, but serves as industry-validated deterministic input for agents.

Why adopt an industry framework rather than build your own?

- **Peer review.** Risks and mitigations are scrutinized by practitioners across multiple institutions. That is a higher bar than any single internal review, and a credible reference point when engaging boards and regulators.
- **Coverage**. Contributors from different institutions bring different deployment experiences, threat models, and regulatory contexts, so the catalog includes risks your internal team hasn't hit yet.
- **Cost**. Maintaining an equivalent risk catalog in-house is a permanent, duplicated expense. Sharing that investment across the industry frees your teams to apply governance rather than author it.
- **Speed**. Agentic AI, MCP, and multi-agent risks are already covered because the community updates the framework as the technology moves.
- **Human and Agent Access**: Available to both Humans in a clean UI and Agents through the MCP server

This does not prevent you from adding additional controls and variants, but provides an industry baseline on which you can build.

Addressing the skepticism

- **"Contributing exposes our internal governance processes."** - Adoption requires no disclosure at all. Contribution happens at the level of generalized risks and mitigations, not your control implementations or risk appetite. Member firms have contributed across two major versions without exposing proprietary processes.
- **"It won't integrate with our existing risk framework."** - It is designed to slot into existing enterprise risk management, not replace it. Risks map onto your existing taxonomies, and mitigations reference the standards (NIST, FFIEC, OWASP) your control library is already built on. Most firms adopt it as an AI-specific overlay on current MRM and operational risk processes. The MCP server also allows your Agentic tools to work at enhancing your proprietary controls using the AIGF resources, such as mitigation and reference architecture, as a deterministic input.
- **"How credible is it as a reference?"** - It is published by FINOS under the Linux Foundation, the same neutral, governed home as FDC3, CDM, and Legend, with named contributions from major financial institutions and an open change history. That provenance is stronger than an internal document no one outside your firm has examined and thus is more referenceable to auditors and regulators.

How do you apply it?

The framework comes with a defined adopter journey: the integration path taught in the FINOS / Linux Foundation Education training (see below). The guiding principle, taken directly from that training: *AIGF expands your existing governance capabilities; it does not replace them.*

1. **Start with existing governance.** Use your current ERM, MRM, security, legal/compliance, vendor-risk, and audit structures. Do not create a separate AI-governance silo.
2. **Introduce a standard AI intake.** Capture business intent, accountable sponsor, data sensitivity, provider model, decision impact, and autonomy level, then route each use case to the right governance lane (full MRM, AI-specific, or hybrid).
3. **Use AIGF as the AI-specific layer.** The risk/mitigation catalog for assessment, the reference architecture library for design patterns and threat models, and evals & benchmarks for measurable assurance evidence.

4. **Standardize review and sign-off.** Consistent evidence packs, approval forums, KRIs, and post-deployment monitoring expectations within your existing decision bodies.
5. **Operationalize and scale.** Reuse patterns, risk mappings, evaluation packs, and evidence requirements across use cases. Reassess when models, data, providers, or architectures change.

To accelerate adoption, FINOS and Linux Foundation Education have developed a structured training scheme. It includes LFE7000, "FINOS AI Governance Framework — Leading Strategically", an executive half-day workshop covering the regulatory imperative (EU AI Act, NIST AI RMF, ISO/IEC 42001, SR 26-2), the framework itself, and hands-on integration of AIGF into existing risk management. It is designed for the same audience as this paper.

Learn More and Get Involved

- Explore the framework: air-governance-framework.finos.org
- Use the print-ready single page for internal circulation: air-governance-framework.finos.org/single-page.html
- Join the FINOS AI Governance Framework working sessions to shape future versions: zoom-lfx.platform.linuxfoundation.org/meetings/ai-governance-framework)
- Enroll your leadership team in training: finos.org/ai-governance-framework-training

finos.org
info@finos.org

