Is Your Financial Institution Governing AI, or Is It Governing You?

Rafael DeLeon

Artificial intelligence (AI) isn't waiting for your financial institution (FI) to catch up. It's already embedded in vendor platforms, fraud detection systems, and loan origination software, to name a few areas. The real question: Are you governing it effectively?

The good news is that you don't have to start from scratch. You can add an AI governance framework to your FI's existing risk management program, with a few key additions.

The Current State of AI Regulations

Although there's no single AI rulebook for FIs, regulatory pressure is mounting from multiple directions, and waiting to be told what to do isn't a risk FIs can take. Three federal frameworks already apply, whether they explicitly mention AI:

- **Interagency Guidance on Third-Party Relationships**: If a vendor is using AI on your behalf, you're responsible for overseeing it.
- **NIST AI Risk Management Framework**: Examiners are referencing it, and it provides safe harbor protections under emerging state laws, such as the Texas Responsible AI Governance Act.
- **Fair lending laws (ECOA and the Fair Housing Act):** An algorithm doesn't get a pass. If it makes or influences a credit decision, the same rules apply.

At the state level, Colorado, California, Utah, and Texas have all passed AI-specific legislation. If you operate across state lines, your framework needs to meet the most restrictive requirements you face.

The federal signal is getting louder, too. The U.S. Department of the Treasury recently announced a major public-private initiative on AI risk management across the financial sector. For mortgage lenders, Freddie Mac's Single-Family Seller/Servicer Guide now includes a formal AI and machine learning governance requirement, effective March 3, 2026. Notably, it covers vendor-embedded AI, not just tools your FI built.

The Critical Components of an AI Governance Framework

Most AI governance failures aren't the result of bad intentions. They're the result of FIs that moved fast, didn't look closely, and assumed someone else had it covered.

A strong program addresses that issue with four key areas: inventory, risk classification, vendor oversight, and documentation.

1. Make an AI Inventory

You can't govern what you can't see, and right now, there's a good chance AI is being used across your FI in ways no one has formally approved.

Shadow AI is any tool employees are using without sign-off from IT, risk, or compliance. It often starts small, such as an employee using an AI chatbot to handle customer or member complaints. Nobody flags it, nobody reviews it, and before long, team members across your FI are using it without the proper guardrails.

Vendor-embedded AI is the other blind spot. Many of the platforms your FI already relies on have quietly added AI capabilities without proactive disclosure. A tool your team has used for years may be making or influencing decisions in ways you haven't evaluated.

That inventory is where governance must start. Survey every department, pull your vendor contracts, and search for terms like "machine learning," "artificial intelligence," and "automated decisioning." For every tool you find, ask: What decisions does it influence? What customer data does it touch? Can its outputs be explained?

What your FI can do: Take the time to build a cross-functional AI inventory. Once you have a foundation, you can add to it over time as you implement new tools and grow your vendor inventory.

2. Focus Your Oversight Where It Counts

Not all AI carries the same risk — and treating it like it does is a fast way to burn resources in the wrong places. A tiered approach lets you concentrate oversight where the exposure is highest.

- **High risk** includes credit underwriting, fair lending decisions, BSA/AML monitoring, and regulatory capital calculations. These warrant full model risk management treatment: independent validation, ongoing monitoring, and thorough documentation.
- **Medium risk** includes customer service chatbots, marketing analytics, and operational efficiency tools. Meaningful oversight, calibrated to a different risk profile.
- **Low risk** includes grammar checkers, scheduling assistants, and basic data organization. It's still AI, but it doesn't need a model validation program.

What your FI can do: Apply a risk tier to every tool in your AI inventory. Use that classification to guide where you invest oversight resources and to show examiners your program is risk-based, not checkbox-driven.

3. Don't Forget About Your Vendors' AI Usage

When a vendor uses AI, the risk doesn't stay with them; it transfers to you. Regulators will hold your FI accountable for vendor AI failures just as they would your own.

AI doesn't come with the same transparency as conventional software. With a traditional tool, you get documentation, specs, maybe even source code. With AI, you may never see how the model makes decisions — and those models learn, drift, and get retrained on data you can't see.

That's why your third-party risk management due diligence needs to reflect this reality. For any vendor using AI, ask: How was this model developed and validated? What data was it trained on? How are updates handled? What bias testing do you perform? What audit rights do we have?

Some vendors will push back on proprietary grounds, but you can push back harder. Build a standardized AI due diligence questionnaire and use it consistently.

What your FI can do: If a vendor updates their model, you need to know before it goes live, not after. Require advance notice of material model changes in every vendor contract — and consider a vendor management platform that centralizes due diligence

and tracks those provisions across your entire portfolio. Manual oversight at scale isn't sustainable.

4. Build the Paper Trail Now

When an examiner asks about AI governance, they're looking for evidence that your program is real and operational.

A mature program produces a clear paper trail: an inventory with risk classifications, a board-approved risk appetite statement, committee meeting minutes, vendor due diligence files with AI-specific elements, bias testing results, model validation reports, and staff training records.

That last one is often underestimated. Who in your FI has been trained on AI risks and policies, and on what specifically? The answer needs to be documented, not assumed.

What your FI can do: Audit your current AI governance documentation. If you can't hand it to an examiner tomorrow and feel confident, that's your starting point.

The Bottom Line

AI governance isn't a future-state project. The regulatory pressure is real, examiner expectations are already forming, and the FIs that move now will be better positioned than the ones playing catch-up later.

About Ncontracts

Ncontracts provides integrated risk management, compliance, and third-party risk management solutions to over 5,500 organizations worldwide, including 4,500 U.S. financial institutions, mortgage companies, and fintechs. The flagship Ncontracts IRM suite combines AI-powered software with expert services, helping financial institutions streamline risk, compliance, and vendor management through an intuitive, cloud-based platform. Ncontracts' Venminder solution is trusted by enterprise financial companies and other large organizations to strategically manage third-party risk across the entire vendor lifecycle.

Visit http://www.ncontracts.com or follow the company on LinkedIn and X for more information.

More Resources

- Whitepaper: Policies as a Power Tool: Creating Policies that Get the Job Done
- Whitepaper: Navigating Financial Institution Change Management
- Monthly Regulatory Compliance Update: Regulatory Brief