

# AI Agent Governance Framework for Banks

How banks and financial institutions can govern AI agents in production across model risk, autonomous decisions, MAS alignment, and audit evidence requirements.

**Use this asset to assess readiness, start a governance conversation, and identify candidates for a 6-week Aethryx design partner pilot.**

**Primary audience:** Banking, fintech, compliance and risk leaders

# Why agent governance matters for banks

AI agents can retrieve data, call tools, trigger workflows, draft customer responses, and support decisions. Governance must cover what the agent is allowed to do, when it must ask for approval, and how every action is evidenced.

## Banking agent control model

| Layer               | Banking risk  | Runtime control  |
|---------------------|---|--|
| Identity and access | Agent acts beyond approved role or context.             | Agent identity, permission boundary, tool allowlist.           |
| Customer impact     | Agent produces advice, eligibility, or service outcome. | Risk tiering, human approval, disclosure checks.               |
| Data and privacy    | Agent retrieves sensitive or unnecessary data.          | PII redaction, retrieval filters, purpose limitation.          |
| Decision autonomy   | Agent executes actions without oversight.               | Approval gates, spending/action limits, kill switch.           |
| Auditability        | Inability to reconstruct decision path.                 | Temporal trace, inputs, tool calls, outputs, reviewer actions. |

## Implementation sequence

- Inventory agents and classify them by autonomy, customer impact, and data sensitivity.
- Define bank-specific policies for customer advice, credit, AML/KYC, complaints, and operations.
- Instrument runtime logs across prompts, context, tools, approvals, and final actions.
- Generate evidence packs for risk committees, internal audit, and regulators.

# Readiness Call CTA

## Ready to assess your organization's AI governance maturity?

Book a 30-minute Aethryx AI Governance Readiness Call. We will help identify runtime AI risks, governance gaps, audit evidence requirements, and potential controls for your AI systems.

| Recommended next step       | Outcome   |
|-----------------------------|---|
| 30-minute readiness call    | Confirm priority AI systems, risk areas, and governance gaps.                       |
| 6-week design partner pilot | Map controls, instrument evidence, and validate runtime governance workflows.       |
| Executive debrief           | Provide a concise risk and control roadmap for risk, technology, and audit leaders. |