CedarBridge Financial

01 | Enterprise AI Governance Framework

Financial-services AI governance engagement

Engagement type	Simulated portfolio case study
Organization	CedarBridge Financial, fictional mid-sized consumer lender and digital bank
Scope	Enterprise AI and model governance across credit, fraud, customer service, marketing, and operations
Primary scenario	CreditPath consumer credit underwriting model and ClearSignal fraud-AI vendor

This artifact uses fictional systems and synthetic data. It is designed to show governance judgment, model-risk discipline, consumer-protection awareness, and operational evidence rather than to represent any real company or vendor.

Purpose

This framework establishes how CedarBridge governs AI systems and models used across consumer credit, fraud detection, marketing, customer service, and internal operations. It connects model risk management, consumer compliance, privacy, security, third-party oversight, and business accountability into one operating model.

Scope

- All statistical models, machine-learning models, scorecards, rules engines with adaptive components, LLM tools, vendor AI services, and internally configured AI workflows that influence consumers, employees, operations, compliance, or risk management.
- Includes model development, acquisition, implementation, monitoring, material change, incident response, retirement, and audit evidence.
- Excludes ordinary office productivity tools only when no confidential, consumer, financial, or regulated data is entered.

Governance principles

Accountability stays with CedarBridge	A vendor may provide a model, but CedarBridge owns the business use, consumer impact, controls, and monitoring.
Risk-based rigor	Critical and high-risk systems receive formal assessments, validation, approval, monitoring, and board visibility. Low-risk systems receive baseline controls.
Consumer harm is a governance risk	Fair lending, adverse-action explainability, complaints, privacy, accessibility, and operational disruptions are treated as risk indicators.
Human accountability	AI may inform decisions, but accountable owners must document how decisions, overrides, and exceptions are handled.
Evidence over policy language	Every required control must have an owner, evidence artifact, testing frequency, and escalation trigger.

Three lines of accountability

First line: Business and model owners	Own use case, design intent, implementation, monitoring, and remediation.	Credit Risk owns CreditPath use, adverse-action reason quality, override process, and performance thresholds.
Second line: Model Risk, Compliance, Privacy, Security	Challenge, approve, validate, and set standards.	Independent validation, fair-lending review, privacy impact review, third-party risk assessment, risk acceptance.
Third line: Internal Audit	Provides independent assurance over governance design and effectiveness.	Audit inventory completeness, approval evidence, validation independence, monitoring quality, and board reporting.

Approval authority by risk tier

Critical	Credit decisioning, loan pricing, high-scale fraud blocks, models with direct regulatory or financial statement reliance.	Business owner, Model Risk, Compliance/Fair Lending, Privacy/Security as applicable, AI Risk Committee, Executive Risk Committee. Board notice required.
High	Fraud alerts with account restrictions, collections prioritization, automated customer segmentation affecting terms or offers.	Business owner, Model Risk, Compliance, Risk Committee.
Moderate	Customer-service chatbot without decision authority, marketing propensity model, operations copilot with internal data.	Business owner, Privacy/Security, AI Governance Office.
Low	Internal drafting using approved tools and no consumer data.	Manager approval plus acceptable-use controls.
Prohibited	Black-box credit denial with no accurate reasons, protected-class inference, deceptive chatbot use, unsupported biometric/emotion inference.	Do not deploy. Escalate for remediation or disablement.

Required lifecycle gates

Intake and inventory	Business purpose, owner, model/system type, data categories, vendor status, downstream decisions, consumer impact.
Classification	Risk tier score, prohibited-use screen, regulatory exposure, consumer impact, decision influence, data sensitivity.
Assessment	Model risk assessment, privacy/security review, fair-lending or consumer-impact review, data/feature review.
Validation and approval	Independent validation proportional to tier, implementation test, limitation register, approval memo.
Production monitoring	Performance, drift, fairness/outcomes, complaints, adverse-action reason quality, override rates, vendor changes.
Change and retirement	Material-change review, retesting, rollback plan, decommissioning evidence, data retention and deletion.

Board and executive oversight

The Executive Risk Committee receives quarterly AI/model risk reporting. The board risk committee receives semiannual summaries and immediate escalation of critical incidents, prohibited-use findings, critical model validation failures, fair-lending concerns, or unresolved high residual risk.

Risk appetite statement

Consumer decision risk	No deployment of critical consumer-decision models without approved validation, adverse-action reason testing, fair-lending review, monitoring thresholds, and manual fallback.
Vendor transparency	No critical or high-risk vendor model may launch without documented model functionality, data use restrictions, change notice, audit rights, and exit plan.
Generative AI	No use of public or unapproved generative AI with confidential consumer data, nonpublic financial information, customer complaints, or regulatory filings.
Fairness and compliance	No knowingly unresolved disparate-impact indicator, unvalidated proxy feature, or inaccessible consumer process may be accepted without executive risk acceptance and remediation timetable.

Reference basis

This simulated package is informed by the 2026 Revised Guidance on Model Risk Management issued by the Federal Reserve, OCC, and FDIC, NIST AI RMF 1.0, the CFPB circular on adverse-action notices for complex credit algorithms, CFPB guidance on specificity of adverse-action reasons, and the 2023 interagency guidance on third-party relationships. These references are used as design inputs for a portfolio artifact, not as legal advice.