CedarBridge Financial

Enterprise AI and Model Governance Assessment

Executive / Board AI Risk Briefing

7
AI / model systems in enterprise inventory

2
Critical systems needing executive visibility

90
Day remediation roadmap

Simulated portfolio case study using fictional systems and synthetic data.

Why this matters

- Financial-services AI governance cannot stop at a general AI policy.
- Credit, fraud, servicing, marketing, and operations models create different consumer, compliance, and operational risks.
- Third-party systems still require CedarBridge-owned controls, evidence, monitoring, and escalation.
- The board needs visibility into risk appetite, high-impact decisions, validation gaps, and remediation progress.

Governance posture

Core framing

AI model risk is governed as a lifecycle: purpose, data, validation, implementation, use, monitoring, change, and retirement.

Enterprise AI landscape

1	1	2	3
Critical credit model	Critical AML model	High-risk systems	Moderate / low systems

Risk tier by domain

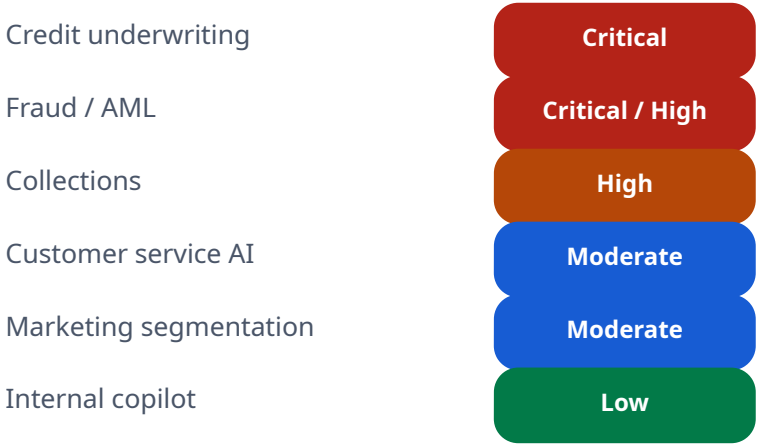

Key takeaway

The highest-risk exposures sit where AI materially influences consumer eligibility, pricing, account access, regulatory risk, or complaint outcomes.

Top enterprise risks

1 CreditPath fair-lending / proxy risk

Synthetic testing indicates a materially lower approval rate for one group and thin-file applicants.

2 Adverse-action reason reliability

The model cannot be used for negative credit decisions unless CedarBridge can provide specific, accurate reasons.

3 Vendor model transparency

FraudGuard documentation and audit rights are incomplete for a high-impact third-party model.

4 Inventory completeness

Embedded AI features in vendor platforms may bypass intake if procurement and business attestations are weak.

CreditPath model decision

Governance decision

Conditional approval for controlled parallel run only. No full production use for automated decline or pricing until remediation is closed and second-line approval is documented.

NO FULL LAUNCH

PARALLEL RUN ONLY

RETEST REQUIRED

Blocked until

- Fair-lending and outcomes remediation
- Adverse-action reason QA live
- Thin-file alternative documentation path
- Monitoring feeds and rollback tested
- Executive Risk Committee acceptance

CreditPath risk controls

Data & features

Prohibit ZIP-level and device-behavior underwriting features; restrict transaction-derived signals.

Validation

Conceptual soundness, performance, calibration, sensitivity, segment outcomes, implementation testing.

Explainability

Specific adverse-action reasons mapped to validated model drivers and QA sampled monthly.

Oversight

Human override taxonomy, exception review, alternative documentation route, documented recourse.

Monitoring

Drift, approval rates, impact ratios, complaints, reason accuracy, override patterns.

Third-party AI vendor risk: ClearSignal FraudGuard

Conditional approval score

2.95 / 5

Approved with conditions for 90 days only.

Required before renewal

- Regulator and CedarBridge audit/access rights
- Material model-change notice and approval rights
- Incident notice within 24 hours for high-impact events
- Data-use restrictions and subprocessor disclosure
- Exportable scores, reason indicators, and exit plan

Control gaps and evidence needed

High	**CreditPath fairness remediation**	Fair-lending report, feature log, validation retest
High	**Adverse-action reason QA**	Reason-code test results, QA sampling evidence
High	**Vendor audit rights**	Executed contract addendum and evidence access
Moderate	**Inventory completeness**	Quarterly business and procurement attestation
Moderate	**Override governance**	Reviewer taxonomy and monthly QA report

Risk appetite decisions

- No critical consumer-decision model without approved validation, adverse-action reason testing, fair-lending review, monitoring thresholds, and manual fallback.
- No high-impact vendor model without model documentation, change notice, audit rights, data-use restrictions, and exit plan.
- No public or unapproved GenAI use with nonpublic consumer information.
- No unresolved High residual risk without executive acceptance and a dated remediation plan.

Decisions requested

1

Approve risk appetite statement

2

Authorize CreditPath remediation

3

Require vendor addendum

Monitoring and incident triggers

Monthly / quarterly monitoring

- Performance and calibration by segment
- Feature and population drift
- Approval, pricing, and false-positive outcomes
- Adverse-action reason accuracy
- Overrides, complaints, and disputes
- Vendor changes and outages

Pause or rollback if

- Adverse-action reason engine fails
- Material disparity remains unexplained
- Vendor makes unapproved model change
- Consumer-impact metric breaches threshold
- Data leakage or unauthorized data use occurs
- Complaint cluster indicates AI-driven harm

12-month governance roadmap

0-30 days

CreditPath remediation design; reason-code QA; thin-file alternative route

31-60 days

ClearSignal contract addendum; BridgeAssist pilot controls; inventory attestation

61-90 days

Validation retest; parallel run dashboard; executive risk decision

3-6 months

All critical/high models on monitoring dashboard; audit evidence repository

6-12 months

Independent assurance review; board reporting maturity; annual reassessment

Decisions for the committee

Approve

- Enterprise AI risk appetite statement
- Risk-tier methodology and approval gates
- 90-day remediation roadmap

Require

- CreditPath remediation before launch
- ClearSignal vendor addendum before renewal
- Quarterly inventory attestations

Escalate

- Any prohibited AI use
- Any critical incident or unresolved High residual risk
- Any material customer harm indicator

Reference basis: 2026 Revised Model Risk Management Guidance; NIST AI RMF 1.0; CFPB adverse-action guidance for complex algorithms; 2023 interagency third-party risk guidance.