

THIRD-PARTY MODEL AND VENDOR GOVERNANCE REVIEW

AI Governance Program | Vendor Risk and Model Oversight

Document Name	Third-Party Model and Vendor Governance Review
AI System	Meridian Automated Loan Underwriting System
Vendor	CrediSure AI [Fictional — Portfolio Use Only]
Model	CrediSure Credit Decision Engine
Organization	Meridian Financial Services [Fictional — Portfolio Use Only]
Document Owner	Oyinjadesola Edu
Review Contributors	AI Governance Lead, Model Risk Reviewer, Compliance Reviewer, Security / Privacy Reviewer
Review Status	Draft — Pending Vendor Evidence Review
System Status	Pre-Production Pilot
Version	1.0
Date	May 23, 2026

NOTICE: This document is a fictional, synthetic governance artifact created for portfolio and educational purposes. It does not represent any real organization, system, vendor, or individual. All data, scenarios, and decisions are illustrative only.

1. EXECUTIVE SUMMARY

Meridian Financial Services is evaluating CrediSure AI as the third-party provider of the CrediSure Credit Decision Engine an AI-enabled credit decisioning model that powers the Meridian Automated Loan Underwriting System. The model generates credit risk scores, confidence scores, approval and denial recommendations, manual review flags, and denial reason codes for small business loan applications.

Because the model may influence or directly produce credit decisions, it creates significant third-party AI governance risk. Meridian remains accountable for how the system is used in lending, even when the model is supplied by an external vendor.

Review Status: Vendor and model review is incomplete. Production approval must not proceed until required vendor documentation, validation evidence, security evidence, model change obligations, and monitoring commitments are reviewed and accepted.

2. VENDOR OVERVIEW

Vendor Name	CrediSure AI [Fictional — Portfolio Use Only]
Vendor Type	Third-party AI credit decisioning provider
Hosting Provider	CrediSure AI Cloud
Integration Method	API
Primary Service	AI-enabled small business credit risk scoring and underwriting decision support
Vendor Relationship Status	Under pre-production review

Contract Status	Pending final risk, legal, compliance, and security review
------------------------	--

CrediSure AI is expected to provide sufficient evidence to support responsible deployment, including model documentation, validation results, fairness testing evidence, reason code methodology, security documentation, incident notification procedures, and model change controls.

3. MODEL OVERVIEW

Model Name	CrediSure Credit Decision Engine
Model Version	v2.3 — Pre-production Pilot
Model Type	AI-enabled credit risk scoring and decisioning model
Use Case	Small business loan underwriting
Deployment Status	Pending production approval
Associated System	Meridian Automated Loan Underwriting System

The model evaluates structured application, financial, credit, banking, business, and repayment data to produce underwriting outputs. Primary model outputs include:

- Credit risk score and confidence score
- Approval recommendation
- Denial recommendation
- Manual review flag
- Denial reason codes

4. BUSINESS USE CONTEXT

Meridian intends to use the CrediSure Credit Decision Engine to improve speed and consistency in small business loan underwriting. The proposed operating model routes approximately 94% of applications through automated processing.

Decision Outcome	Approximate Share	Human Review Required?
Auto-approve	52%	No — unless applicant appeals
Auto-deny	42%	No — unless applicant appeals
Manual review	6%	Yes — routed to human underwriter

Vendor Risk Note: A third-party model may influence or directly process 94% of underwriting outcomes. This creates elevated vendor dependency risk and requires robust contractual controls, monitoring, and change management obligations.

5. INTENDED MODEL CAPABILITIES

The model is expected to support the following capabilities. Each capability is relevant to high-risk decisioning because it may affect whether an applicant receives access to credit.

Capability	Description
Credit Risk Scoring	Evaluates application data to produce a numerical risk score

Underwriting Recommendation	Generates approval or denial recommendations based on risk thresholds
Automated Approve / Deny	Routes qualifying or disqualifying applications to automated outcomes
Manual Review Flagging	Identifies borderline or uncertain applications for human review
Confidence Scoring	Provides a confidence level associated with each decision recommendation
Denial Reason Codes	Generates reason codes to support adverse action communication
Risk Segmentation	Segments applications by risk profile for portfolio monitoring
Monitoring Support	Supports portfolio-level performance and drift monitoring

6. KNOWN MODEL LIMITATIONS

Based on the current governance review, the following limitations must be addressed before production deployment is approved.

Limitation	Detail
Limited Transparency into Training Data	Meridian does not yet have full visibility into the model's training data sources, representativeness, exclusions, or historical bias controls.
Limited Visibility into Model Architecture	Meridian does not yet have sufficient information about model design, feature weighting, performance boundaries, or explainability methodology.
Reason Code Dependency	Meridian depends on vendor-generated reason codes to support adverse action explanations, appeals, and compliance review. The accuracy and adequacy of these codes has not been independently validated.
Model Update Dependency	Vendor model updates could affect scoring, approval rates, denial rates, fairness outcomes, or reason code behavior without advance notice to Meridian.
Potential Performance Drift	Model performance may degrade due to economic shifts, portfolio changes, borrower behavior changes, industry volatility, or changes in credit conditions.
Limited Independent Validation	Meridian has not yet completed independent validation of the model's fairness, accuracy, stability, explainability, and monitoring thresholds.

7. DATA DEPENDENCY REVIEW

The model may process or rely on the following data categories. The model must not directly use protected characteristics. However, proxy-bias risk remains because several variables may correlate with protected or disadvantaged groups.

Data Category	Data Elements
Loan Application Data	Requested loan amount, loan purpose, loan term
Business Financial Data	Annual revenue, monthly revenue, monthly cash flow, existing debt, debt service coverage ratio
Credit Data	Business credit score, guarantor credit score, credit utilization, prior defaults, recent delinquencies
Banking Data	Bank transaction data, average account balance, overdraft activity, revenue volatility indicators
Business Profile Data	Industry classification, business age, geographic region (state / ZIP3 level only)
Risk and Compliance Data	Fraud screening indicators, compliance flags, missing documentation flags
Historical Data	Historical repayment and default data used for model training and scoring

Proxy Bias Risk: Variables such as geography, industry classification, business age, credit history, and cash-flow volatility may correlate with protected or disadvantaged groups. Proxy-bias analysis is required before production deployment.

8. BIAS AND FAIRNESS CONCERNS

The following bias and fairness concerns require review and documented mitigation before production deployment. Meridian must require evidence of fairness testing, disparate impact analysis, proxy-bias review, and segment-level outcome monitoring.

- Historical lending data may reflect prior disparities in credit access
- Geography may act as a proxy for protected or disadvantaged groups
- Industry classification may disadvantage businesses concentrated in certain communities
- Business age may disadvantage newer businesses with limited credit history
- Thin credit files may result in higher denial rates for underbanked applicants
- Cash-flow volatility may disadvantage seasonal or smaller businesses
- Credit bureau data may reflect historical inequities or incomplete financial context
- Automated denials may create customer harm if reason codes are inaccurate or incomplete

9. EXPLAINABILITY REVIEW

The model is expected to generate denial reason codes for adverse decisions. Reason code governance is critical because Meridian must be able to explain automated denial decisions to applicants, underwriters, compliance reviewers, and auditors.

The following explainability concerns remain open:

- Whether reason codes accurately reflect the model factors that drove the decision
- Whether reason codes are specific enough to support applicant understanding
- Whether reason codes are consistent across similar decisions
- Whether reason codes can be independently validated by Meridian
- Whether human reviewers can understand and meaningfully challenge model outputs
- Whether limitations in explainability are clearly documented by the vendor

Production Gate: Production deployment must be delayed if denial reason codes cannot support adverse action notices, applicant appeals, and governance review.

10. SECURITY AND PRIVACY REVIEW

The security and privacy review is not yet complete. The following items must be confirmed before production approval is granted.

Review Area	Requirement
API Security	Authentication, authorization controls, and API key management
Data Encryption	Encryption of data in transit and at rest
Access Controls	Least privilege practices and role-based access to model outputs and logs
Data Retention	Retention and deletion requirements for Meridian data processed by the vendor
Logging and Monitoring	Vendor access logging and monitoring for unauthorized activity

Incident Notification	Timelines and procedures for notifying Meridian of security incidents
Vulnerability Management	Vendor practices for identifying and remediating security vulnerabilities
Privacy Impact	Privacy impact considerations for data transmitted to the vendor API
Data Processing Terms	Contractual terms governing how vendor processes and stores Meridian data
Secondary Use Restriction	Prohibition on using Meridian data for vendor model training without approval

Current Status: Security and privacy assessment is pending completion. Production approval must not be granted until this review is documented and accepted.

11. VENDOR DOCUMENTATION REQUIRED

Before production deployment, Meridian must obtain and review the following documentation from CrediSure AI. All items are currently outstanding.

Document	Description	Status
Model Card / Model Documentation	Intended use, out-of-scope use, known limitations, and version history	Outstanding
Training Data Summary	Data sources, representativeness, exclusions, and historical bias controls	Outstanding
Feature Category Summary	Categories of input variables used by the model	Outstanding
Validation Results	Model performance metrics, accuracy, stability, and validation methodology	Outstanding
Fairness Testing Summary	Disparate impact analysis, segment-level outcomes, and bias mitigation approach	Outstanding
Reason Code Methodology	How reason codes are generated, validated, and mapped to model factors	Outstanding
Monitoring Approach	Drift detection, performance monitoring, and alerting processes	Outstanding
Model Update Process	How model changes are managed, tested, and communicated to deployers	Outstanding
Incident Notification Procedure	Timelines, contacts, and escalation process for security or model incidents	Outstanding
Security Documentation	API security controls, encryption, access management, and audit logging	Outstanding
Privacy and Data Processing Docs	Data processing terms, retention, deletion, and secondary use restrictions	Outstanding
Audit Support Commitments	Vendor obligations to support Meridian audits and regulatory inquiries	Outstanding

12. REQUIRED VENDOR OBLIGATIONS

The vendor agreement must include contractual obligations for the following areas. Meridian must not rely solely on vendor claims key controls and outputs must be independently reviewed by AI governance, compliance, model risk, and security teams.

Obligation Area	Requirement
-----------------	-------------

Documentation Access	Right to receive and review model documentation, validation evidence, and fairness testing results
Validation Evidence	Obligation to provide model performance and validation results on request
Fairness Testing Support	Obligation to provide disparate impact analysis and bias mitigation evidence
Reason Code Transparency	Obligation to explain and validate denial reason code methodology
Model Update Notification	Advance notice of material model changes before promotion to production
Material Change Review	Meridian right to review and approve material changes before deployment
Incident Notification	Defined timelines and procedures for notifying Meridian of security or model incidents
Security Requirements	Contractual security standards for API, data transmission, and access controls
Data Protection	Restrictions on data retention, secondary use, and vendor model training use
Audit Support	Obligation to support Meridian audits and regulatory or supervisory inquiries
Service-Level Commitments	API availability, response time, and support commitments
Unauthorized Data Reuse	Explicit prohibition on using Meridian data for vendor model training without approval

13. MODEL CHANGE MANAGEMENT REQUIREMENTS

CrediSure AI must notify Meridian before any material change that could affect the following model behaviors or outputs. No material model change may be promoted into production without review and approval by the designated Meridian reviewers.

Notification Triggers — Material Changes Requiring Advance Notice

- Model scoring logic or risk threshold changes
- Approval rate, denial rate, or manual review rate changes
- Reason code behavior or mapping changes
- Fairness outcomes or segment-level performance changes
- Model confidence score distribution changes
- API behavior, input schema, or output format changes
- Data processing practices or retention policy changes
- Security posture or access control changes

Review Responsibilities — Material Change Review Team

- AI Governance Lead — overall governance impact assessment
- Credit Risk Owner — underwriting policy and threshold impact
- Model Risk Reviewer — model performance and validation impact
- Compliance Reviewer — fair lending and regulatory impact
- Vendor Risk Owner — contractual and vendor relationship impact
- Security / Privacy Reviewer — security and data protection impact

14. MONITORING REQUIREMENTS

Meridian must monitor vendor model performance at launch, during pilot deployment, and throughout production use. The following metrics and indicators must be tracked.

Monitoring Area	Metrics and Indicators
Decision Outcomes	Approval rate, denial rate, manual review rate — tracked over time and by segment
Appeals and Overrides	Appeal rate, appeal success rate, and manual override rate
Credit Performance	Default rate and portfolio performance by decision cohort
Complaints	Complaint volume, categories, and resolution outcomes
Model Quality	Confidence score distribution, reason code distribution, and drift indicators
Fairness Metrics	Segment-level outcome differences and disparate impact indicators
Vendor Changes	Model version changes and observed impact on decision outcomes
API and Security	API availability, response times, and security incident reports
Model Update Frequency	Frequency and nature of vendor model updates

15. KEY VENDOR AND MODEL RISKS

The following risks are directly relevant to this vendor and model review. Each risk must be addressed through controls, contractual obligations, or documented risk acceptance.

Risk	Description
Proxy Bias Through Credit and Business Variables	Vendor model inputs may create unfair outcomes through variables correlated with protected or disadvantaged groups. Fairness testing and proxy-bias analysis are required.
Weak Explainability and Incomplete Reason Codes	Vendor-generated reason codes may not fully or accurately explain automated decisions, undermining adverse action obligations and applicant appeal rights.
Accountability Gaps in Third-Party AI Deployment	Responsibility for model outcomes may become unclear between Meridian and CrediSure AI without explicit contractual role definitions and governance controls.
Inaccurate Automated Denials	Model errors, data quality issues, or overly strict thresholds may cause qualified applicants to be incorrectly denied access to credit.
Lack of Meaningful Human Oversight	Vendor outputs may be accepted without sufficient challenge or independent review, particularly given the 94% automation rate.
Security and Privacy Risk	Sensitive financial and credit-related data transmitted to the vendor API may be exposed or misused if security controls, data processing terms, or access restrictions are inadequate.

16. REVIEW FINDINGS

The table below summarizes the current status of each review area. All items are open. No area has been completed or accepted at this stage of the review.

Review Area	Current Finding	Status
Model Documentation	Not fully reviewed — model card and documentation not yet received	Open
Training Data Summary	Not yet received — vendor has not provided training data documentation	Open
Fairness Testing Evidence	Not yet complete — disparate impact analysis and bias review outstanding	Open

Reason Code Validation	Pending — reason code methodology and accuracy not yet validated	Open
Security Assessment	Pending — API security, encryption, and data protection review not complete	Open
Vendor Audit Rights	Not confirmed — audit support obligations not yet contractually established	Open
Model Update Notification	Needs contractual confirmation — advance notice obligations not yet agreed	Open
Monitoring Thresholds	Not finalized — performance metrics, drift thresholds, and alerting criteria pending	Open
Incident Notification	Needs contractual confirmation — notification timelines not yet agreed	Open
Production Approval	Not ready — full production approval must not be granted at this stage	Open

17. PRODUCTION READINESS ASSESSMENT

Current Assessment: The vendor and model are not ready for full production approval at this stage. A limited or controlled pilot may be considered only after all conditions below are satisfied.

A limited pilot may be considered only if Meridian confirms all of the following:

- Fairness testing is completed and results are documented
- Reason codes are validated and confirmed adequate for adverse action
- Human review triggers are defined and approved
- Appeal process is documented and tested
- Security review is completed and accepted
- Vendor documentation is reviewed and accepted
- Vendor model change obligations are contractually confirmed
- Monitoring thresholds and alerting criteria are defined
- Incident notification requirements are documented and accepted
- Governance committee sign-off is obtained

18. RECOMMENDED GOVERNANCE DECISION

Recommendation: Conditionally approve for limited production use only after required vendor evidence, security review, reason code validation, fairness testing, monitoring controls, and contractual obligations are completed.

If the vendor cannot provide sufficient transparency, validation evidence, security evidence, model change notification, or audit support, deployment must be delayed until these requirements are satisfied.

19. EVIDENCE MAPPING

This document supports the following governance frameworks and competency areas.

Framework	Governance Competencies Demonstrated
-----------	--------------------------------------

EU AI Act	Provider/deployer role analysis; technical documentation dependency; human oversight and explainability concerns; data governance and bias concerns; post-deployment monitoring and incident escalation
ISO/IEC 42001	Supplier and third-party AI governance; AI system lifecycle responsibilities; model documentation requirements; operational controls; monitoring and continual improvement
NIST AI RMF	Govern: accountability, vendor oversight, and role clarity Map: model context, dependencies, and intended use Measure: limitations, validation needs, and fairness concerns Manage: vendor risk treatment, monitoring, and escalation