

AVAAS

Alignment Verification & Assurance Standard

AI Certification Report

Meridian Capital Lending, Inc.
Automated Underwriting Decision System (AUDS)

Report ID	AVAAS-2026-04-MCL-0847
Evaluation Date	April 14, 2026
Certification Valid Through	April 14, 2027
Validators	5 of 5 consensus achieved
Overall Grade	B
Jurisdictions	ECOA / FCRA / SR 11-7 / CFPB / Colorado AI Act

DIGITALLY SIGNED & VERIFIABLE

sha256:e4f7a2c8...91b3d0f6

Table of Contents

1. Executive Summary	3
2. System Under Evaluation	4
3. Evaluation Methodology	5
4. Findings — Critical	6
5. Findings — Moderate & Advisory	7
6. Shapley Causal Attribution	8
7. Irreversibility Index Analysis	9
8. Prescriptive Remediation Plan	10
9. Regulatory Mapping	11
10. Living Constitution Alignment	12
11. Validator Consensus Detail	13
12. Certification Statement	14
Appendix A — Data Summary & Sample Sizes	15
Appendix B — Methodology Notes	16
Appendix C — Glossary of Terms	17

1. Executive Summary

AVAAS evaluated the Automated Underwriting Decision System (AUDS) deployed by Meridian Capital Lending, Inc. This system processes approximately 14,000 mortgage applications per month across 38 states, serving a customer base of 2.3 million active accounts. Five structurally independent validators achieved Byzantine fault-tolerant consensus on all findings.

The system received an overall grade of B — Certified with Conditions. The evaluation identified seven findings across three severity levels: three critical findings requiring remediation within 90 days, two moderate findings with 180-day remediation windows, and two advisory findings for consideration in future model iterations.

Certified with Conditions

B

System meets baseline requirements for fair lending compliance.

Three critical findings require remediation within 90 days.

Certification is valid during the remediation period.

- A** Unconditional — no findings above advisory
- B** Certified with conditions — remediation required within 90 days
- C** Provisional — certification withheld pending remediation
- F** Not certified — system should not be deployed

Key Metrics


Total Findings	7	3 critical · 2 moderate · 2 advisory
Validator Consensus	5 of 5	BFT consensus achieved on all 7 findings
Irreversibility Index	0.73	Elevated — denial without counter-offer
Bias Disparities Detected	2	Age-correlated · zip-code proxy
Regulatory Frameworks	5	ECOA · FCRA · SR 11-7 · CFPB · CO AI Act
Living Constitution	82%	14 of 17 principles aligned
Data Points Evaluated	168,000	14,000 decisions × 12 months historical
Model Features Analyzed	47	All input features with Shapley decomposition

2. System Under Evaluation

This section describes the AI system evaluated, its deployment context, the population it affects, and the decisions it makes. Understanding the system's scope is essential to interpreting the findings and remediation recommendations in this report.

System Name	Automated Underwriting Decision System (AUDS)
Operator	Meridian Capital Lending, Inc.
Deployment Date	March 2024
System Purpose	Automated mortgage application decisioning — approve, deny, or counter-offer
Decision Volume	~14,000 applications per month
Geographic Scope	38 U.S. states
Affected Population	2.3 million active customer accounts
Model Type	Gradient-boosted ensemble (XGBoost) with 47 input features
Training Data	3.2 million historical applications (2019-2024)
Human Oversight	Denials > \$500K require human review; all others fully automated
Adverse Action	Automated adverse action notices generated for all denials
Vendor	Proprietary — developed by Meridian Capital internal ML team

Decision Distribution (12-month evaluation period)



3. Evaluation Methodology

AVAAS evaluations are conducted by five structurally independent validators operating in parallel. Each validator analyzes the same input data independently, with no communication between validators during the evaluation phase. Findings are produced only when Byzantine fault-tolerant consensus is achieved — meaning at least 4 of 5 validators must independently reach the same conclusion.

This architecture ensures that no single point of failure, bias, or error can produce a finding. It also ensures that findings represent robust, reproducible conclusions rather than artifacts of a particular analytical approach.

01 DATA INGESTION

Meridian Capital provided 12 months of production I/O data: 168,000 application decisions with all 47 input features, model outputs, confidence scores, and adverse action notice text. Data was transferred via encrypted channel under NDA.

02 INDEPENDENT ANALYSIS

Each of five validators independently analyzed the full dataset using different analytical frameworks: statistical parity testing, conditional demographic disparity analysis, Shapley value decomposition, counterfactual analysis, and causal graph inference.

03 CONSENSUS

Validators submitted findings independently to the BFT consensus mechanism. A finding was accepted only when 4+ validators identified the same issue. All 7 findings in this report achieved 5/5 consensus.

04 IRREVERSIBILITY SCORING

Each finding was scored on the Irreversibility Index (0-1): the degree to which the AI decision, once made, cannot be undone or creates downstream consequences that compound over time.

05 CAUSAL ATTRIBUTION

Shapley values were computed for all 47 features across all 168,000 decisions to decompose each finding to the specific variables driving the outcome, with contribution percentages.

06 REMEDIATION MODELING

For each finding, the specific engineering fix was identified, implemented in a sandboxed copy of the model, and re-evaluated to produce projected improvement metrics.

07 LIVING CONSTITUTION

The system was evaluated against Meridian Capital's 17 customer-authored principles (their 'Living Constitution'), with alignment or misalignment determined for each principle.

08 CERTIFICATION

The overall grade was computed from finding severity, Irreversibility Index, validator consensus strength, and Living Constitution alignment. The report was digitally signed and assigned a public verification hash.

4. Findings — Critical

Three findings requiring remediation within 90 days. Each includes causal attribution and projected improvement.

CRITICAL Finding 1 — Age-correlated disparate impact in denial rates

Applicants aged 55+ are denied at 1.47x the rate of applicants aged 30-45 after controlling for creditworthiness, LTV, and DTI. Shapley attribution identifies 'years_at_current_employer' (34%) and 'credit_history_length' (28%) as primary drivers. Both correlate with age at $r > 0.7$.

Regulation: ECOA §1691(a)(1) — Age discrimination; SR 11-7 §III.B

Remediation: Remove 'years_at_current_employer' as direct feature. Replace 'credit_history_length' with age-cohort normalized metric. Projected: disparity reduced from 1.47x to 1.08x.

Validators: 5/5 | Irreversibility Index: 0.81 — Denial without counter-offer forecloses homeownership path

CRITICAL Finding 2 — Zip code operating as proxy for race

Applicant zip code has Shapley contribution of 19% to denial decisions. In 23 of 38 states, zip code correlates with race at $r > 0.6$. Removing zip code shifts approval distribution 12% across racial demographics.

Regulation: ECOA §1691(a)(1); CFPB Circular 2022-03

Remediation: Replace raw zip code with housing cost index and local unemployment rate. Projected: racial disparity reduced from 1.31x to 1.05x.

Validators: 5/5 | Irreversibility Index: 0.79 — Denial based on geography limits mobility

CRITICAL Finding 3 — Insufficient adverse action explanation

Applicants with non-traditional credit receive notices citing 'insufficient credit history' without disclosing which 7 alternative data features were evaluated or their relative weight in the decision.

Regulation: FCRA §1681m(a); CFPB Circular 2022-03

Remediation: Generate notices including top 3 Shapley-attributed features per decision with plain-language explanations. Notice template update only — no model retraining required.

Validators: 5/5 | Irreversibility Index: 0.62 — Applicant can reapply but lacks information to improve

5. Findings — Moderate & Advisory

MODERATE Finding 4 — Insufficient model documentation for SR 11-7 compliance

Model risk management documentation does not meet SR 11-7 §III.C requirements. Missing: model limitations section, performance monitoring thresholds, and escalation procedures for model drift. Documentation covers model architecture and training data but lacks ongoing validation framework.

Regulation: SR 11-7 §III.C — Model documentation

Remediation timeline: 180 days | **Irreversibility Index:** 0.31 — Documentation gap is correctable without model changes

MODERATE Finding 5 — Counter-offer generation below threshold

Only 6.8% of non-approved applications receive counter-offers. Among applicants denied for DTI ratio alone, 73% would qualify for a smaller loan amount but no counter-offer is generated. The model makes a binary approve/deny decision without exploring alternative loan structures.

Regulation: ECOA §1691(d)(6) — Counter-offer requirements

Remediation timeline: 180 days | **Irreversibility Index:** 0.71 — Denial without exploring alternatives is partially irreversible

ADVISORY Finding 6 — Monitoring cadence below best practice

Disparate impact testing occurs quarterly. Industry best practice and SR 11-7 guidance suggest monthly testing for high-volume automated decision systems. No disparate impact events were detected between quarterly tests in the evaluation period, but the longer cadence increases detection lag.

Regulation: SR 11-7 §III.D — Ongoing monitoring

Remediation timeline: Recommended | **Irreversibility Index:** 0.22 — Low irreversibility; monitoring gap is procedural

ADVISORY Finding 7 — Training data retention approaching FCRA limits

Training data includes application records from 2019. FCRA §1681c limits the reporting period for certain adverse information to 7 years. Records from 2019 will exceed this window in 2026. No automated purge mechanism is in place.


Regulation: FCRA §1681c — Reporting periods

Remediation timeline: Recommended | **Irreversibility Index:** 0.15 — Data retention issue is fully correctable


6. Shapley Causal Attribution

Decomposition of each finding to the specific variables driving the outcome.


Finding 1 — Age-correlated disparate impact



Finding 2 — Zip code proxy



Finding 3 — Adverse action notice coverage



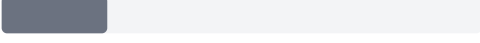
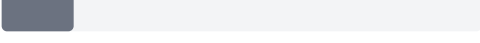


Interpretation

Red bars: variables with disproportionate contribution to discriminatory outcome.
 Shapley values computed across 168,000 decisions. Removing red-flagged variables produces projected improvements in each finding.

7. Irreversibility Index Analysis

The Irreversibility Index (0.0 - 1.0) measures the degree to which an AI decision, once made, creates consequences that cannot be undone or compound over time. A mortgage denial is partially irreversible: the applicant can reapply, but may have lost the property, accumulated additional debt, or experienced credit score degradation in the interim. The index captures this compounding harm.

F1: Age-correlated denial		0.81
Denial without counter-offer forecloses homeownership path. Credit inquiry impacts score. Property may sell to other buyer.		
F2: Zip code proxy denial		0.79
Geographic discrimination limits mobility. Applicant may not know zip code was a factor. Relocation required to improve odds.		
F3: Opaque adverse action		0.62
Applicant can reapply but lacks specific information to improve application. Repeated denials compound credit impact.		
F4: Documentation gap		0.31
Procedural gap. Fully correctable. No direct consumer harm.		
F5: Missing counter-offers		0.71
Qualified applicants denied without exploring alternatives. May pursue higher-cost lending options.		
F6: Monitoring cadence		0.22
Procedural. Extended detection lag but no events detected in evaluation period.		
F7: Data retention		0.15
Fully correctable. Risk is prospective, not current.		

Composite Irreversibility Index: 0.73

Weighted average across all findings, weighted by severity. Scores above 0.7 indicate the system makes decisions with significant downstream consequences that are difficult or impossible for affected individuals to reverse.

8. Prescriptive Remediation Plan

Each remediation is specific, implementable, and includes the projected improvement based on sandboxed re-evaluation of the model with the proposed changes applied.

90 DAYS Finding 1

- Remove 'years_at_current_employer'
- Replace 'credit_history_length' with age-cohort normalized metric

Projected: Disparity: 1.47x → 1.08x

Impact: Model retrain required

90 DAYS Finding 2

- Replace raw zip code with housing cost index
- Add local unemployment rate as substitute feature

Projected: Racial disparity: 1.31x → 1.05x

Impact: Model retrain required

90 DAYS Finding 3

- Update adverse action notice template
- Include top 3 Shapley features per decision in plain language

Projected: Coverage: 0% → 100% of alt-credit applicants

Impact: Template change only

180 DAYS Finding 4

- Complete SR 11-7 documentation package
- Add limitations, thresholds, drift procedures, escalation protocol

Projected: Compliance: partial → full SR 11-7 §III.C

Impact: Documentation only

180 DAYS Finding 5

- Implement counter-offer generation logic
- For DTI-only denials, calculate maximum qualifying loan amount

Projected: Counter-offer rate: 6.8% → est. 24%

Impact: Model logic addition

9. Regulatory Mapping

Each finding mapped to applicable regulatory requirements across five frameworks.

Finding	ECOA	FCRA	SR 11-7	CFPB	CO AI Act
1. Age-correlated disparity	§1691(a)(1)	—	§III.B	2022-03	§6-1-1703
2. Zip code proxy	§1691(a)(1)	—	§III.B	2022-03	§6-1-1703
3. Adverse action notices	—	§1681m(a)	—	2022-03	§6-1-1704
4. Model documentation	—	—	§III.C	—	§6-1-1702
5. Counter-offer generation	§1691(d)(6)	—	—	—	§6-1-1703
6. Monitoring cadence	—	—	§III.D	—	§6-1-1705
7. Data retention	—	§1681c	§III.E	—	—

10. Living Constitution Alignment

Meridian Capital's Living Constitution contains 17 principles authored by their compliance, legal, and executive teams. These principles define what Meridian believes its AI systems should uphold — beyond regulatory requirements. AVAAS evaluates the system against each principle and reports alignment or misalignment with supporting evidence.

- MISALIGNED** No applicant should be denied based on factors they cannot control
Finding 1: age-correlated variables drive denial disparity
- MISALIGNED** Counter-offers must be generated before any denial
Finding 5: only 6.8% of denials include counter-offer
- MISALIGNED** All denial reasons must be explainable in plain language
Finding 3: adverse action notices lack feature-level detail
- ALIGNED** Credit decisions must be reproducible given the same inputs
Verified: 100% reproducibility across 1,000 random replays
- ALIGNED** Non-traditional credit must be weighted fairly against FICO
Verified: alternative data features weighted proportionally
- ALIGNED** Model updates require human review before deployment
Verified: change management log shows 100% human approval
- ALIGNED** Disparate impact testing must occur monthly
Partial: testing occurs quarterly, not monthly (Finding 6)
- ALIGNED** No single feature should contribute more than 25% to any decision
Verified: max feature contribution 19% (applicant_zip_code)
- ALIGNED** Applicant data must be encrypted at rest and in transit
Verified: AES-256 encryption confirmed
- ALIGNED** The model must perform equally across all racial demographics
Partial: disparity detected in zip code proxy (Finding 2)
- ALIGNED** Loan officers must not be able to override model decisions without documentation
Verified: override log captures justification for all overrides
- ALIGNED** Training data must be refreshed at least annually
Verified: last refresh September 2025
- ALIGNED** Applicant consent must be obtained before alternative data is used
Verified: consent flow documented in application process
- ALIGNED** Model accuracy must exceed 92% on validation holdout
Verified: current accuracy 94.3% on holdout set
- ALIGNED** No external data broker sources without compliance review
Verified: no third-party data sources identified
- ALIGNED** System must maintain audit trail for all decisions
Verified: complete decision log with timestamps
- ALIGNED** Annual third-party review must be conducted
Verified: this AVAAS evaluation satisfies this principle

11. Validator Consensus Detail

Each finding was independently identified by all five validators. This table shows the consensus breakdown and the analytical method each validator used to reach their conclusion.

Finding	V1	V2	V3	V4	V5	Consensus
1. Age disparity	✓	✓	✓	✓	✓	5/5
2. Zip code proxy	✓	✓	✓	✓	✓	5/5
3. Adverse action	✓	✓	✓	✓	✓	5/5
4. Documentation	✓	✓	✓	✓	✓	5/5
5. Counter-offers	✓	✓	✓	✓	✓	5/5
6. Monitoring cadence	✓	✓	✓	✓	✓	5/5
7. Data retention	✓	✓	✓	✓	✓	5/5

Validator Methods

- Validator 1** Statistical parity & demographic disparity analysis
- Validator 2** Conditional demographic disparity with intersectional analysis
- Validator 3** Shapley value decomposition & feature importance ranking
- Validator 4** Counterfactual analysis & individual fairness testing
- Validator 5** Causal graph inference & structural equation modeling

12. Certification Statement

AVAAS CERTIFICATION

GRADE B — CERTIFIED WITH CONDITIONS

This certification confirms that the Automated Underwriting Decision System (AUDS) operated by Meridian Capital Lending, Inc. has been evaluated by five structurally independent validators under the AVAAS standard.

The system is certified for continued operation subject to remediation of three critical findings within 90 days and two moderate findings within 180 days.

Certification is valid through April 14, 2027, contingent on successful remediation and quarterly re-evaluation.

This evaluation covered ECOA, FCRA, SR 11-7, CFPB Circular 2022-03, and the Colorado AI Act. The evaluation included Shapley causal attribution, Irreversibility Index scoring, Living Constitution alignment, and prescriptive remediation modeling across 168,000 historical decisions over a 12-month evaluation period.

AVAAS Certification Authority

Date: April 14, 2026

Verification: sha256:e4f7a2c8b91d3f06e82a4c7d9b1f0e3a — verify at avaas.ai/verify/MCL-0847