

ANNEX IV

Technical Documentation Evidence Package

SYSTEM credit-decision-engine-v2.1

ORGANIZATION Meridian National Financial Group

PERIOD 2026-03-01 to 2026-03-31

REPORT ID ANNEX4-MNFG-CDE-202603-A7F

142,867

ATTESTATIONS GENERATED

High Risk

ANNEX III(5)(B) — CREDIT SCORING

100%

SIGNED & LEDGER-ANCHORED

REGULATORY BASIS

EU AI Act — Regulation (EU) 2024/1689, Article 11 and Annex IV. Article 12 record-keeping obligations for high-risk AI systems. Article 14 human oversight requirements.

United States — Federal Reserve SR 11-7 model risk management supervisory guidance. OCC Bulletin 2011-12. FDIC FIL-22-2017. Equal Credit Opportunity Act / Reg B (12 CFR § 1002). Fair Credit Reporting Act (15 U.S.C. § 1681).

Cryptographic integrity — every assertion in this document is backed by an ML-DSA-65 (NIST FIPS 204) signed attestation anchored to the Hedera Consensus Service mainnet topic 0.0.10416909, independently verifiable without operator cooperation.

SECTION 1 · EXECUTIVE SUMMARY

Scope and Findings

This evidence package documents the operation of **Credit Decision Engine v2.1**, an AI-assisted credit decisioning system operated by Meridian National Financial Group, during the period **2026-03-01** through **2026-03-31**. The package is intended to satisfy the technical documentation requirements of Article 11 and Annex IV of the EU AI Act for the referenced period, and to support examiner review under SR 11-7.

Scope of evidence

During the reporting period, the system processed 142,867 credit applications across consumer installment loan and residential mortgage product lines. Every decision — including the model version in production, the input features considered, the output classification, any human-in-the-loop review, and any adverse action notice issued — was captured as a Rubric attestation at the moment it occurred. Approval rate: **67.3%**. Adverse action notices issued: **46,782**. Human-in-the-loop reviews: **3,892**.

Cryptographic integrity of the record

All 142,867 attestations in the reporting period were signed using ML-DSA-65 (NIST FIPS 204, post-quantum) and aggregated into a hierarchical Merkle forest whose roots were anchored to the Hedera Consensus Service public distributed ledger (mainnet topic 0.0.10416909). The evidence below is independently verifiable by any party in possession of the public keys and the Hedera transaction identifiers listed in Section 7 — without the cooperation, continued existence, or trustworthiness of Rubric Protocol or Meridian National Financial Group.

Key findings

1. Coverage: 100% of credit decisions in the reporting period were captured and signed at the time of decision. No decision-to-attestation gap greater than 47 microseconds was observed.
2. Human oversight: of 3,892 decisions routed to human reviewers under the bank's established exception handling policy, 2,341 resulted in counter-offers and 1,551 were ultimately approved or declined in line with the model's recommendation.
3. Adverse action compliance: all 46,782 adverse action notices were generated with the model-driven principal reason codes required by ECOA Reg B § 1002.9, with the principal reasons themselves attested as a separate leaf type and linked to the decision record.
4. Model drift: no manifest drift events were detected during the reporting period. The model version in production was `cde-v2.1.7-stable` throughout.

Prepared by: Rubric Protocol automated evidence generator, invoked on 2026-04-03 14:22 UTC.

Intended recipients: Meridian Chief Risk Officer; Model Risk Committee; external auditor; national competent authority upon request.

Retention: this evidence package and the underlying attestations are retained for no less than seven (7) years. Hedera anchoring is permanent.

SECTION 2 · AI SYSTEM DESCRIPTION

System and Risk Classification

Intended purpose

Credit Decision Engine v2.1 is an AI-assisted decision support system that evaluates consumer credit applications for installment loans (up to \$75,000) and residential mortgages (up to \$2.5M). The system produces one of three outputs for each application: *Approve*, *Decline*, or *Counter-Offer*, together with a set of principal reason codes and a natural-language adverse action explanation (where declined).

Model architecture

Ensemble of gradient-boosted decision trees (XGBoost v2.0) trained on 11 years of the bank's historical underwriting data, augmented by a fine-tuned language model (Meridian-Explainer-v3, a Llama 3 70B derivative) used solely to generate consumer-facing adverse action notices in plain English. The language model does not influence the credit decision itself; it translates principal reason codes emitted by the primary model into Reg B-compliant prose.

Risk classification

Classification: High-Risk AI System under Annex III(5)(b) of Regulation (EU) 2024/1689 (AI systems intended to evaluate the creditworthiness of natural persons or establish their credit score).

Rationale: the system makes or materially influences decisions about access to credit for natural persons. Per Article 6(2), such systems are categorically High-Risk and subject to the full obligations in Chapter III, Section 2 of the AI Act.

Corresponding US framework: the system is a "Model" as defined by SR 11-7 (§1) and is subject to Meridian's enterprise Model Risk Management program. It is additionally subject to ECOA Reg B for consumer protection and the FCRA for any use of consumer credit report data.

Human oversight

Per Article 14, human oversight is implemented through a dedicated exception handling workflow. Any application meeting any of the following criteria is automatically routed to a credit analyst for human review before a final decision is communicated to the applicant: (a) model confidence below 0.80; (b) principal amount above \$500,000; (c) protected-class pattern detected by the fairness monitor; (d) applicant request for human review. During the reporting period, 3,892 of 142,867 applications (2.7%) were routed for human review.

SECTION 3 · EVIDENCE STATISTICS

Decisions and Attestations

142,867 TOTAL DECISIONS	96,143 APPROVED	44,383 DECLINED	2,341 COUNTER-OFFERED
-----------------------------------	---------------------------	---------------------------	---------------------------------

46,782 ADVERSE ACTION NOTICES	3,892 HUMAN REVIEWS	47µs AVG SIGNING LATENCY	0 DRIFT EVENTS
---	-------------------------------	------------------------------------	--------------------------

Weekly decision volume

Period	Decisions	Approved	Declined	Counter-Offered
Week of Mar 02	32,481	21,906	9,932	643
Week of Mar 09	35,124	23,618	11,014	492
Week of Mar 16	36,902	24,895	11,521	486
Week of Mar 23	38,360	25,724	11,916	720
Period total	142,867	96,143	44,383	2,341

Attestation leaf type distribution

Rubric captures each decision as a set of structured attestation leaves. The distribution below reflects every leaf generated during the reporting period.

Leaf type	Description	Count
0x01 decision_record	Core decision: inputs, model version, output, confidence	142,867
0x02 principal_reasons	Reg B principal reason codes for declined or counter-offered applications	46,724
0x03 adverse_action	Adverse action notice issued under ECOA §1002.9	46,782
0x04 hitl_review	Human-in-the-loop review attestation (reviewer ID, timestamp, outcome)	3,892
0x05 model_version	Cryptographic binding to the model version in production at decision time	142,867
0x07 reproducibility	Seed, temperature, top_p, prompt for any LLM-generated content	46,782

SECTION 4 · REGULATORY COVERAGE MATRIX

Where each obligation is evidenced

Every regulatory obligation is mapped to one or more sections of this evidence package and to specific Rubric attestation leaf types. This table is intended to let an auditor locate the evidence for any particular obligation in one pass.

Obligation	Citation	Evidence location	Leaves
Automatic recording of events (period of use, input data, persons in verification)	EU AI Act Art. 12(2)(a)–(d)	Section 5, Section 3	0x01, 0x04, 0x05
Technical documentation covering Annex IV elements	EU AI Act Art. 11 + Annex IV	Section 2, this document	A11
Human oversight — mechanism, effectiveness, interventions	EU AI Act Art. 14	Section 2 & Section 5 (record #2)	0x04
Post-market monitoring — drift, performance, events	EU AI Act Art. 72	Section 3 (drift); logged in 0x06	0x06
Model development, implementation, use, validation, governance	Fed SR 11-7 § V–VII	Section 2 + attestations 0x05	0x05
Notice to applicant — adverse action, principal reasons	ECOA Reg B 12 CFR § 1002.9	Section 5 (record #3), Section 3	0x02, 0x03
Accuracy of consumer credit report information relied upon	FCRA 15 U.S.C. § 1681e(b)	Section 3 + 0x01 input features	0x01
Cryptographic integrity of the record; retention	EU AI Act Art. 12(3); ECOA § 1002.12	Section 6 (signatures)	A11

Reading this table. Each obligation points to a document section and the Rubric attestation leaf types that bear the cryptographic evidence. The leaf-type bytes correspond to the Rubric Protocol v1.0 specification and are preserved under the Deprecation Policy described in the Rubric MSA (180 days' notice for any breaking change).

SECTION 5A · SAMPLE ATTESTATION RECORD

Standard approval decision

The record below is representative of 96,143 approval decisions made during the reporting period. Fields have been sanitized but the structure, signature, and anchor are real. Personal data appears only as a pseudonymous hash; no raw PII is written to the ledger.

attestation_id	cde-202603-a7f-000012847
leaf_type	0x01 decision_record
decision	APPROVE
product_line	consumer_installment_loan
principal_amount_usd	24,500.00
applicant_id_hash	sha3-256:4c2a...8b91e7
model_version	cde-v2.1.7-stable
model_confidence	0.947
input_feature_hash	sha3-256:9f17...2d83c4
hitl_reviewed	false
decision_timestamp_utc	2026-03-14T11:42:06.184729Z
signing_latency_us	43
signature_alg	ML-DSA-65 / NIST FIPS 204
sig_head	3a91f7d2b8e4c6a0cd...
merkle_leaf_index	12847
merkle_batch_root	0x7f3a9b2c8e4d1a6f...
hcs_tx_id	0.0.3923341@1710416426.091238475

✓ VERIFIED ON-LEDGER

SECTION 5B · SAMPLE ATTESTATION RECORD

Human-in-the-loop counter-offer

This record reflects a case where the model confidence fell below the 0.80 HITL threshold and the application was routed to a named credit analyst for review. The reviewer's identity, timestamp, and decision are independently attested as a linked leaf type (0x04).

attestation_id	cde-202603-a7f-000047121
leaf_type	0x01 decision_record + 0x04 hitl_review
decision	COUNTER_OFFER
product_line	residential_mortgage
principal_amount_usd	485,000.00
applicant_id_hash	sha3-256:2e8b...7c14a9
model_version	cde-v2.1.7-stable
model_confidence	0.742
model_recommendation	DECLINE
hitl_reviewer_id	analyst.mcauley.j / NMLS-284715
hitl_review_timestamp_utc	2026-03-18T14:23:51.447219Z
hitl_outcome	counter_offer_issued: \$445,000 at 30yr-fixed
hitl_rationale_hash	sha3-256:6b4f...a928d3
decision_timestamp_utc	2026-03-18T14:24:02.098432Z
signature_alg	ML-DSA-65 / NIST FIPS 204
hcs_tx_id	0.0.3923341@1710771842.447219000
linked_attestations	2 (decision + hitl review)

✓ VERIFIED ON-LEDGER

SECTION 6 · CRYPTOGRAPHIC SIGNATURE BLOCK

Verifiable without us

Every attestation in this package is signed under the same federation public key set and anchored to a single Hedera Consensus Service topic. With the public keys and the topic ID, any party can verify the integrity of this evidence package and every attestation it references — permanently, without Rubric's participation, cooperation, or continued existence.

Public key set

Attestations were signed by a 3-of-5 threshold across the following federation nodes. Compromise of the signature chain would require the simultaneous compromise of at least three of these nodes, located across five legal jurisdictions.

Node	Jurisdiction	Public key (ML-DSA-65, head...tail)
rubric-us	United States (Virginia)	pk:mlds65:3a91f7d2b8e4...6a0cd...9e2b4d81f06c3a57
rubric-ca	Canada (Ontario)	pk:mlds65:b82c4091fe7a...d7e1...4a1f9b302e6c8d11
rubric-eu	Germany (Frankfurt)	pk:mlds65:5f1a8e37c2b4...9d02...7c3e0a8b41d96f25
rubric-sg	Singapore	pk:mlds65:91e3b7f4a602...3c81...2f8d94b106a53e47
rubric-jp	Japan (Tokyo)	pk:mlds65:c42a9f1b387e...054d...8e16b03fa927c4d9

Hedera anchoring

Ledger	Hedera Consensus Service (Mainnet)
Topic ID	0.0.10416909
First anchor in period	0.0.3923341@1709294522.184729301
Last anchor in period	0.0.3923341@1711886842.091238475
Total anchors in period	31 (daily batches)
Period-end Merkle root	0x7f3a9b2c8e4d1a6f05c938b1e7a4d2f09c6b3e8a1d7f4c209b5e3a8d1f67c4b29
Signature scheme	ML-DSA-65 (NIST FIPS 204)
Aggregation	SHA3-256 Merkle forest, hierarchical (1M:1 compression)
Threshold	3 of 5 federation signatures required

✓ VERIFIED ON-LEDGER

SECTION 7 · INDEPENDENT VERIFICATION

How to verify this package without us

Below are the exact commands an auditor can run to independently verify the cryptographic integrity of this package. No Rubric account, API key, or network call to Rubric-operated infrastructure is required. The only external dependencies are a Hedera mirror node (publicly operated) and an open-source Rubric Protocol reference implementation.

Step 1 — Retrieve the on-ledger anchors

```
curl https://mainnet-public.mirrornode.hedera.com/
  api/v1/topics/0.0.10416909/messages
  ?timestamp=gte:1709294400.000000000
  &timestamp=lte:1711886399.999999999
```

Step 2 — Verify a specific attestation signature

```
rubric-cli verify \
  --attestation cde-202603-a7f-000012847 \
  --pubkey-set federation-keys.json \
  --merkle-batch-root 0x7f3a9b2c8e4d1a6f... \
  --hcs-tx 0.0.3923341@1710416426.091238475
```

Step 3 — Re-derive the period Merkle root

```
rubric-cli forest \
  --bundle annex4-mnfg-cde-202603-a7f.bundle \
  --expect-root 0x7f3a9b2c8e4d1a6f...
```

What verification proves: that every attestation in this package was signed by a 3-of-5 quorum of federation keys at the timestamp claimed, that the batch Merkle roots were anchored to the Hedera mainnet at the times claimed, and that nothing in the package has been altered since. **What verification does not prove:** the correctness of the underlying credit decision, the fairness of the model, or the quality of the human review — those are attested facts, not validated ones.

APPENDIX · DEFINITIONS AND CITATIONS

Glossary and legal citations

Key terms

Attestation. A structured, cryptographically signed record of an AI decision event, captured at the moment the event occurs.

Leaf type. A byte identifier designating the kind of attestation (decision, principal reasons, HITL review, etc.). Defined in Rubric Protocol v1.0.

Merkle forest. A hierarchical Merkle tree arrangement permitting efficient aggregation of large volumes of leaf records under a single root.

ML-DSA-65. Module-Lattice-Based Digital Signature Algorithm with NIST security category 3 parameters; the post-quantum signature standard published as NIST FIPS 204.

Threshold signing. A scheme in which a signature is valid only if produced by a specified threshold of distinct signers; Rubric uses 3-of-5 Shamir threshold signing across five federation nodes.

Hedera Consensus Service. A public distributed ledger service providing ordered, immutable, timestamped message submission; used by Rubric to anchor Merkle roots.

Citations

- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).
- Board of Governors of the Federal Reserve System, SR Letter 11-7: Supervisory Guidance on Model Risk Management (April 4, 2011).
- Office of the Comptroller of the Currency, Bulletin 2011-12: Sound Practices for Model Risk Management (April 4, 2011).
- Federal Deposit Insurance Corporation, FIL-22-2017: Adoption of Supervisory Guidance on Model Risk Management (June 7, 2017).
- Equal Credit Opportunity Act, 15 U.S.C. § 1691 et seq.; Regulation B, 12 CFR Part 1002.
- Fair Credit Reporting Act, 15 U.S.C. § 1681 et seq.
- NIST FIPS 204, Module-Lattice-Based Digital Signature Standard (August 13, 2024).
- NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0), January 2023.

About this sample. This evidence package is an illustrative artifact demonstrating the structure, content, and cryptographic integrity properties of an Annex IV evidence package generated by Rubric Protocol. The institution and system described are fictitious; the numerical figures are representative but synthetic. In a production deployment, every field in this document would be backed by a real, independently verifiable attestation on the Hedera Consensus Service. Contact scott@rubric-protocol.com for a walk-through keyed to your organization's actual AI systems.