

EU AI Act Article 9 Risk Management File

Loan Scoring Model

Version 2.1

Provider	Acme Financial Services (fictional)
Assessor	Alice Chen, AI Governance Lead
Assessment Date	04 October 2026
Review Date	29 September 2027
RMF Schema	v1.0.0
Generated	04 October 2026 08:19 UTC

SHA-256:
991ccc67bfcd2238a1b0bd0a672b93c6623da2f2e1cbdd6096c7e588f272cef2

1. Executive Summary

Purpose: Automated credit scoring for retail loan applications.

Annex III Category: Essential Services

Self-Classification Documented: Yes

Total Risk Items: 10

Open Items (above appetite threshold of 9): 0

Knowledge Gaps: 1

2. System Description

Field	Value
System ID	42565357-74c8-45c5-bfd8-f6e79dd89077
Name	Loan Scoring Model
Version	2.1
Provider	Acme Financial Services (fictional)
Provider Contact	Not specified
Intended Users	Not specified
Inputs	Not specified
Outputs	Not specified
Deployment Context	Not specified

3. Risk Register

Risk appetite threshold: 9 (scores above this require mitigation or acceptance).

3.1 Health Safety

Could the system's outputs directly influence a clinical, safety, or physical decision without mandatory human review? LOW (3)

Consider decisions about medication dosing, surgical planning, vehicle operation, equipment control, or emergency dispatch.

Likelihood	rare (1)	Severity	moderate (3)
Residual Likelihood	rare	Residual Severity	moderate
Source	question_bank	Regulatory Status	settled

Article refs: Art.9(2)(a), Art.14(1)

NIST AI RMF: MAP 1.5 | ISO/IEC 42001: Clause 6.1

3.2 Fundamental Rights

Has a fundamental rights impact assessment (FRIA) been conducted prior to deployment? HIGH (12)

Article 9(2)(a) requires identification of known and foreseeable risks to fundamental rights. A FRIA is the standard mechanism for this.

Likelihood	possible (3)	Severity	major (4)
Residual Likelihood	unlikely	Residual Severity	moderate
Source	question_bank	Regulatory Status	settled

Article refs: Art.9(2)(a)

NIST AI RMF: GOVERN 1.1 | ISO/IEC 42001: Clause 6.1

Mitigations:

- **Preventive** [planned] — Complete a fundamental rights impact assessment before deployment and repeat it for each major model version

3.3 Discrimination

Demographic bias in credit scoring outputs HIGH (12)

Credit scoring models trained on historical data may encode systemic biases against protected groups (age, gender, ethnicity). Article 9(9) requires consideration of impacts on vulnerable groups.

Likelihood	possible (3)	Severity	major (4)
Residual Likelihood	unlikely	Residual Severity	moderate
Source	pattern	Regulatory Status	settled

Article refs: Art.9(2)(a), Art.9(9), Art.10(2)(f)

NIST AI RMF: MEASURE 2.9 | ISO/IEC 42001: Clause A.7

Mitigations:

- **Preventive** [planned] — Remove postcode feature; block release when demographic parity difference exceeds 0.10

Has the training dataset been audited for representation imbalances across protected characteristics (gender, ethnicity, CRITICAL (20)

Biased training data produces biased outputs. Article 10(2)(f) requires data governance measures to address biases.

Likelihood	likely (4)	Severity	critical (5)
Residual Likelihood	unlikely	Residual Severity	moderate
Source	question_bank	Regulatory Status	settled

Article refs: Art.10(2)(f), Art.9(2)(a)

NIST AI RMF: MEASURE 2.9 | ISO/IEC 42001: Clause A.7

Mitigations:

- **Preventive** [planned] — Commission an independent bias audit of the training data before each retrain; block release on a failed audit

3.4 Privacy

Does the system process personal data or special category data (health, biometric, ethnicity, religion)? HIGH (12)

Processing of special category data triggers additional obligations under GDPR Article 9 and EU AI Act Article 10(5).

Likelihood	possible (3)	Severity	major (4)
Residual Likelihood	possible	Residual Severity	major
Source	question_bank	Regulatory Status	settled

Article refs: Art.10(5)

NIST AI RMF: GOVERN 1.6 | ISO/IEC 42001: Clause A.8

Accepted: Residual within appetite after controls.

3.5 Transparency

Lack of explanation for credit decisions MEDIUM (9)

Automated credit decisions must be explainable to affected individuals under GDPR Article 22. The system may not currently provide this.

Likelihood	possible (3)	Severity	moderate (3)
Residual Likelihood	possible	Residual Severity	moderate
Source	pattern	Regulatory Status	settled

Article refs: Art.13

NIST AI RMF: GOVERN 1.7 | ISO/IEC 42001: Clause A.6

Are users notified that they are interacting with or being assessed by an AI system? MEDIUM (9)

Article 13 requires providers to ensure a sufficient level of transparency to enable users to interpret the system's output. Notification is a baseline.

Likelihood	possible (3)	Severity	moderate (3)
Residual Likelihood	possible	Residual Severity	moderate
Source	question_bank	Regulatory Status	settled

Article refs: Art.13(1)

NIST AI RMF: GOVERN 1.7 | ISO/IEC 42001: Clause A.6

3.6 Human Oversight

Can a human operator override or stop the AI system's output in real time without significant technical barriers? MEDIUM (8)

Article 14(4)(e) requires the ability to intervene on or interrupt AI system operation. This must be technically implemented, not just policy-stated.

Likelihood	unlikely (2)	Severity	major (4)
Residual Likelihood	unlikely	Residual Severity	major
Source	question_bank	Regulatory Status	settled

Article refs: Art.14(4)(e)

NIST AI RMF: MANAGE 1.1 | ISO/IEC 42001: Clause A.9

3.7 Robustness

Has the system been tested against out-of-distribution inputs, edge cases, and adversarial perturbations? HIGH (12)

Article 9(7) and Article 15 require testing of AI systems against the intended purpose and reasonably foreseeable misuse.

Likelihood	possible (3)	Severity	major (4)
Residual Likelihood	possible	Residual Severity	major
Source	question_bank	Regulatory Status	settled

Article refs: Art.9(7), Art.15

NIST AI RMF: MEASURE 2.5 | ISO/IEC 42001: Clause A.9

Accepted: Residual within appetite after controls.

3.8 Data Governance

Is the provenance of all training, validation, and test datasets documented (source, collection date, license, version)? MEDIUM (6)

Article 10(2) requires data governance practices including documentation of dataset provenance. TraceForge can automate this.

Likelihood	unlikely (2)	Severity	moderate (3)
Residual Likelihood	unlikely	Residual Severity	moderate
Source	question_bank	Regulatory Status	settled

Article refs: Art.10(2), Art.10(3)

NIST AI RMF: GOVERN 1.7 | ISO/IEC 42001: Clause A.8

Knowledge gap — test requirements derived. See Section 5.

4. Article 9(7) Test Requirements

Risk Item	Metric	Threshold	Article
2e699798	demographic_parity	<= 0.10	Art.9(7)
aa98503e	explanation_coverage	>= 0.90	Art.13
dd03681f	fundamental_rights_impact_score	= 0	Art.9(2)(a)
b3e0a068	demographic_parity	<= 0.10	Art.9(7)
34effdcf	pii_leakage_rate	= 0.00	Art.9(7)
627f04c6	explanation_coverage	>= 0.90	Art.13
1641d6e8	faithfulness	>= 0.85	Art.9(7)
7a141cbf	data_quality_score	>= 0.95	Art.10

5. Audit Trail

This document includes a SHA-256 hash-chained audit trail stored in .riskforge/

audit.jsonl. The export hash links this document to a specific audit log entry.

Export

Audit Entry Hash 11a08ed65588c3eadc26b763358c4640bad8b26b1bcd52565bf3ade52102fc0e

Document SHA-256 991ccc67bfcd2238a1b0bd0a672b93c6623da2f2e1cbdd6096c7e588f272cef2

Signed By Unsigned

Verify integrity at any time: riskforge verify

6. Disclosure

This document was produced using RiskForge v1.1.3, question bank version 1.0.0. It represents the team's documented risk assessment and has not been reviewed by a qualified legal professional. It does not constitute legal advice under the EU AI Act or any other regulation.