AI COMPLIANCE DIAGNOSTIC REPORT

72

OUT OF 100

HIGH

Acme AI Labs

AI Compliance Diagnostic Report

Prepared for: Sarah Chen, Chief Technology Officer

Generated: 4/16/2026 | Ref: AP-2026-9263

TABLE OF CONTENTS

Risk Score: 72/100 HIGH

EXECUTIVE SUMMARY

Executive Summary

Acme AI Labs presents a High overall compliance risk profile based on three confirmed gaps identified during this diagnostic. While the organisation demonstrates meaningful controls — including clean training datasets, active bias detection thresholds, and end-user explainability — two critical structural gaps create material regulatory exposure. The absence of immutable data lineage records for the core inference model directly violates ISO 42001 traceability requirements and undermines SOC 2 CC6.1 audit defensibility. More critically, the lack of human-in-the-loop oversight for high-risk automated decisions is a direct violation of EU AI Act Article 14, carrying potential fines of up to €30M or 6% of global annual turnover. Immediate board-level remediation of these gaps is recommended before the EU AI Act enforcement deadlines take effect.

FRAMEWORK SCORECARD

EU AI ACT
48 /100
FAIL

NIST AI RMF
61 /100
REVIEW

ISO 42001
55 /100
REVIEW

SOC 2
70 /100
PASS

KEY REGULATORY FINDINGS

KEY REGULATORY FINDINGS

[CRITICAL] No Human-in-the-Loop Oversight for High-Risk Automated Decisions

Citation: EU AI Act Art. 14; NIST AI RMF GOVERN 1.2

Acme AI Labs has confirmed that high-risk automated decisions are executed by AI systems without any mandatory human review or intervention capability. EU AI Act Article 14 explicitly mandates that high-risk AI systems be designed and deployed with effective human oversight measures, enabling qualified individuals to monitor, intervene, override, or halt system outputs. Operating customer-facing or consequential AI decision-making without this safeguard exposes the organisation to enforcement action from the moment a high-risk use case is in production.

Required Action:

1. Conduct an immediate inventory of all automated decision workflows to identify those meeting the EU AI Act high-risk classification criteria under Annex III

2. Design and implement a mandatory human review checkpoint within each identified high-risk decision pipeline before outputs are actioned or communicated to end-users

3. Define escalation thresholds — such as confidence score minimums or decision category flags — that automatically route edge cases to a qualified human reviewer

4. Document the human oversight procedure in a formal operational runbook and assign named role ownership for each decision category

5. Validate the oversight controls through a structured walkthrough exercise and record evidence of human intervention capability for audit purposes

REGULATORY EXPOSURE
Up to EUR 35M or 7% of global turnover - EU AI Act Art. 83

KEY REGULATORY FINDINGS

[CRITICAL] Absence of Immutable Data Lineage Records for Core Inference Model

Citation: ISO 42001:2024 §8.4; SOC 2 CC6.1

Acme AI Labs does not maintain an immutable record of data lineage for its core inference model currently in production. ISO 42001:2024 and SOC 2 CC6.1 both require that organisations establish traceable records of data origins, transformations, and usage throughout the AI system lifecycle. Without this record, the organisation cannot demonstrate to auditors, regulators, or customers where its model training data originated, whether it was processed lawfully, or how data changes between versions affected model behaviour.

Required Action:

1. Select and deploy a data lineage tooling solution — such as Apache Atlas, OpenLineage, or a native feature of your ML platform — capable of write-once, tamper-evident logging

2. Retroactively document the known data sources, preprocessing steps, and transformation logic for your current production model to establish a baseline lineage record

3. Integrate lineage capture into your CI/CD and model training pipelines so that every future training run automatically generates a versioned, immutable lineage artifact

4. Define a data lineage retention policy specifying minimum retention periods aligned with ISO 42001 and your existing SOC 2 evidence requirements

5. Include data lineage records in your next SOC 2 Type II audit evidence package and validate completeness with your external auditor

REGULATORY EXPOSURE
Material - customer contract liability and Certification revocation risk

KEY REGULATORY FINDINGS

[FLAGGED] Formal AI Risk Assessment Not Conducted Within the Past 12 Months

Citation: SOC 2 CC7.3; EU AI Act Art. 9

Acme AI Labs confirmed that its last formal AI risk assessment was conducted more than 12 months ago, meaning the documented risk posture no longer reflects the current operational environment, model versions, or deployment scope. SOC 2 CC7.3 requires continuous risk monitoring and periodic formal reassessment to ensure that identified risks remain accurately characterised and that controls remain effective.

Required Action:

1. Schedule and initiate a formal AI risk assessment within the next 30 days, scoped to all AI systems currently in production or in staged deployment

2. Use a structured risk assessment framework — such as NIST AI RMF or ISO 42001 Annex A — to ensure consistent coverage across technical, ethical, and operational risk dimensions

3. Produce a risk register capturing identified threats, likelihood ratings, current control effectiveness, and residual risk for each model or use case assessed

4. Present the completed risk register to executive leadership and the board with a prioritised remediation plan and owner assignments

5. Establish a recurring calendar cadence — minimum annually — for AI risk assessment reviews with assigned ownership and board-level visibility

STRATEGIC REMEDIATION ROADMAP

01

Implement human-in-the-loop review checkpoints across all high-risk automated decision workflows to satisfy EU AI Act Article 14 requirements.

CRITICAL | HIGH EFFORT | 4-6 WEEKS | EU AI ACT ART. 14

STRATEGIC REMEDIATION ROADMAP

02

Deploy an immutable data lineage tracking system for the core inference model and retroactively document existing training data provenance.

FLAGGED | HIGH EFFORT | 6-8 WEEKS | ISO 42001:2024 / SOC 2 CC6.1

03

Conduct a formal AI risk assessment using NIST AI RMF or ISO 42001, covering all production models, and establish an annual reassessment cadence.

FLAGGED | MEDIUM EFFORT | 30 DAYS | SOC 2 CC7.3

04

Update internal AI governance documentation to reflect new lineage, oversight, and risk assessment controls, and cross-reference these in the SOC 2 evidence package.

FLAGGED | LOW EFFORT | 2-3 WEEKS | ISO 42001 / SOC 2

05

Establish a semi-annual AI compliance review cycle with named executive ownership to prevent recurrence of assessment lapses and maintain continuous regulatory alignment.

FLAGGED | LOW EFFORT | ONGOING | EU AI ACT / NIST AI RMF / ISO 42001