AI COMPLIANCE DIAGNOSTIC REPORT

68

OUT OF 100

ELEVATED

Meridian Software Group

AI Compliance Diagnostic Report

Prepared for: James Okafor, Chief Information Security Officer

Generated: 6/8/2026 | Ref: AP-2026-7832

TABLE OF CONTENTS

Risk Score: 68/100 ELEVATED

EXECUTIVE SUMMARY

Executive Summary

Meridian Software Group's AI stack - a mid-market B2B SaaS platform serving 340 enterprise customers across financial services and legal sectors - presents an Elevated overall compliance risk profile based on four confirmed gaps identified during this diagnostic. While the organisation demonstrates strong baseline controls including active bias detection, end-user explainability for automated outputs, and a current AI risk assessment, two Critical findings create material regulatory exposure that requires immediate board-level attention. The absence of human oversight on high-risk automated decisions is a direct violation of EU AI Act Article 14, and the lack of immutable data lineage records leaves the organisation unable to demonstrate ISO 42001 traceability or defend SOC 2 CC6.1 audit requirements to its enterprise customer base. Meridian's sector exposure - financial services and legal - places its AI systems within EU AI Act Annex III high-risk categories, meaning enforcement risk is not hypothetical but immediate upon the August 2026 deadline. The two Flagged findings represent operational gaps that, if unaddressed, will compound regulatory exposure during a supervisory inquiry. Remediating all four findings within the recommended timeframes will reduce the overall risk score to below 40 and position Meridian to meet enterprise procurement compliance requirements without exception.

FRAMEWORK SCORECARD

EU AI ACT
52 /100
REVIEW

NIST AI RMF
67 /100
REVIEW

ISO 42001
59 /100
REVIEW

SOC 2
74 /100
PASS

KEY REGULATORY FINDINGS

[CRITICAL] No Human-in-the-Loop Oversight on High-Risk Automated Decisions in Financial Services Workflows

Citation: EU AI Act Art. 14; NIST AI RMF GOVERN 1.2; ISO 42001:2023 §8.5

Meridian Software Group confirmed that AI-driven automated decisions affecting financial services customers - including credit-adjacent recommendations, document classification outcomes, and risk-scoring outputs - are executed without any mandatory human review checkpoint. EU AI Act Article 14 mandates that high-risk AI systems include effective human oversight measures enabling qualified individuals to monitor, intervene, override, or halt system outputs before they are actioned. Meridian's customer base in financial services and legal sectors places these decision workflows squarely within EU AI Act Annex III high-risk categories. Deploying consequential automated decisions in these sectors without documented oversight controls is an immediately enforceable violation and creates direct enterprise customer contract liability across Meridian's 340-customer book.

Required Action:

1. Conduct an immediate inventory of all automated decision workflows serving financial services and legal sector customers and classify each against EU AI Act Annex III high-risk criteria

2. Design and implement mandatory human review checkpoints within each high-risk decision pipeline - no output should be actioned or surfaced to customers without a qualified reviewer confirming or overriding it

3. Define escalation thresholds including confidence score minimums, decision value limits, and customer category flags that automatically route edge cases to a named human reviewer

4. Document the oversight procedure in a formal operational runbook with named role ownership per decision category and maximum response SLA before escalation

5. Conduct a structured walkthrough exercise of the oversight controls and retain evidence of human intervention capability for the EU AI Act technical documentation file

REGULATORY EXPOSURE

Up to EUR 35M or 7% of global annual turnover under EU AI Act Art. 71(4); enterprise customer contract termination risk across financial services segment

KEY REGULATORY FINDINGS

[CRITICAL] Absence of Immutable Data Lineage Records for Core Inference Model Undermines ISO 42001 and SOC 2 Audit Defensibility

Citation: ISO 42001:2023 §8.4; SOC 2 CC6.1; EU AI Act Art. 10

Meridian does not maintain an immutable record of data lineage for its production inference model. ISO 42001:2023 Section 8.4 requires that organisations document and maintain AI data management activities, including the provenance, preprocessing steps, and transformation history of training data. SOC 2 CC6.1 requires traceable audit records demonstrating how sensitive inputs are handled. Without lineage records, Meridian cannot demonstrate to its enterprise customers, external auditors, or regulators where training data originated, whether it was processed lawfully, or how data changes between model versions affected outputs. Given that Meridian processes client data from regulated financial services firms, the absence of lineage records is also a direct breach of standard enterprise data processing agreement clauses requiring audit trail capability.

Required Action:

1. Select and deploy a data lineage solution - such as Apache Atlas, OpenLineage, or a native MLflow Lineage feature - configured for write-once, tamper-evident logging of all training data operations

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

3. Integrate lineage capture into the CI/CD and model training pipelines so every future training run automatically generates a versioned, immutable lineage artifact with cryptographic integrity verification

4. Define a data lineage retention policy specifying minimum 5-year retention aligned with ISO 42001 and the evidence requirements of Meridian's enterprise customer SOC 2 audits

5. Include complete data lineage records in the next SOC 2 Type II evidence package and validate completeness with the external auditor before the audit window opens

REGULATORY EXPOSURE

Material - SOC 2 Type II audit finding and enterprise customer contract liability; ISO 42001:2023 certification revocation risk

[FLAGGED] AI Incident Response Plan Does Not Cover Model-Specific Failure Modes Including Hallucination, Bias Drift, and Adversarial Input

Citation: NIST AI RMF RS-1.1; SOC 2 CC7.4; ISO 42001:2023 §10.1

Meridian's existing incident response plan covers standard cybersecurity incidents but does not address AI-specific failure modes including model hallucination at scale, bias drift between retraining cycles, adversarial prompt injection, or cascading failure from third-party model API downtime. NIST AI RMF Response function (RS-1 through RS-4) requires that organisations maintain documented response procedures specific to AI system failures. For a platform serving regulated financial services customers, the absence of AI-specific incident response creates significant exposure - a hallucination event or biased output affecting a financial services customer's decision workflow could trigger regulatory notification obligations under multiple jurisdictions simultaneously.

Required Action:

1. Conduct a failure mode inventory for each production AI model identifying hallucination risk, bias drift indicators, adversarial input vectors, and third-party API dependency failure scenarios

2. Draft an AI-specific incident response playbook with defined detection triggers, triage procedures, customer notification thresholds, and escalation paths for each failure mode category

3. Integrate AI incident response procedures into the existing SOC 2 incident response framework with clear hand-off points between AI operations and security operations teams

4. Define regulatory notification obligations by jurisdiction - particularly for financial services customers in the EU - and document the decision tree for when a model failure triggers mandatory reporting

5. Run a tabletop exercise simulating a hallucination event affecting a financial services customer workflow and use findings to refine the playbook before the next SOC 2 audit period

REGULATORY EXPOSURE

SOC 2 audit finding risk; regulatory notification liability in financial services jurisdictions; enterprise customer contract breach risk

KEY REGULATORY FINDINGS

[FLAGGED] Third-Party AI Model API Dependencies Not Included in Vendor Risk Assessment Programme

Citation: NIST AI RMF MAP 5.1; ISO 42001:2023 §8.6; EU AI Act Art. 28

Meridian's production AI stack incorporates third-party model APIs - including foundation model providers - that are not currently included in the organisation's vendor risk assessment programme. NIST AI RMF MAP 5.1 and ISO 42001:2023 Section 8.6 both require that organisations assess and manage risks arising from AI supply chain dependencies. Foundation model providers represent a unique category of vendor risk: a change to model weights, safety filters, output formatting, or API availability can silently alter the behaviour of Meridian's downstream products in ways that affect end-customer outcomes and regulatory compliance posture, without Meridian receiving advance notice.

Required Action:

1. Inventory all third-party AI model API dependencies in production and classify each by risk tier based on criticality to customer-facing workflows, data exposure, and substitutability

2. Extend the existing vendor risk assessment programme to include AI-specific assessment criteria: model change notification policies, output stability commitments, data processing terms, and regulatory compliance posture

3. Establish monitoring for model API change logs and deprecation notices, with a defined internal review process triggered by any upstream model update that could affect customer-facing outputs

4. Include AI API vendor assessments in SOC 2 CC9.2 evidence and document supplier controls in ISO 42001 §8.6 records

5. Draft contractual addenda or confirm existing DPA coverage for all AI model API providers processing Meridian customer data as part of inference requests

REGULATORY EXPOSURE
ISO 42001:2023 non-conformance finding; SOC 2 CC9.2 audit finding; EU AI Act Art. 28 deployer obligations exposure

STRATEGIC REMEDIATION ROADMAP

01

Implement human-in-the-loop review checkpoints across all high-risk automated decision workflows serving financial services and legal sector customers.

CRITICAL | HIGH EFFORT | 0–6 WEEKS | EU AI ACT ART. 14

02

Deploy immutable data lineage tracking for all production inference models and retroactively document existing training data provenance.

FLAGGED | HIGH EFFORT | 0–8 WEEKS | ISO 42001 §8.4 / SOC 2 CC6.1

03

Draft and validate an AI-specific incident response playbook covering hallucination, bias drift, adversarial input, and third-party API failure scenarios.

FLAGGED | MEDIUM EFFORT | 2–5 WEEKS | NIST AI RMF RS-1.1 / SOC 2 CC7.4

04

Extend vendor risk assessment programme to include all third-party AI model API dependencies with AI-specific assessment criteria.

FLAGGED | MEDIUM EFFORT | 3–6 WEEKS | NIST AI RMF MAP 5.1 / ISO 42001 §8.6

05

Update AI governance documentation across all four findings and cross-reference in the next SOC 2 Type II evidence package.

FLAGGED | LOW EFFORT | ONGOING | ISO 42001 / SOC 2 / EU AI ACT