AI COMPLIANCE DIAGNOSTIC REPORT

81

OUT OF 100

CRITICAL

VERIFIED AUDIT

Series B AI Company (Anonymised)

AI Compliance Diagnostic Report

Prepared for: Michael Torres, Chief Technology Officer

Generated: 4/16/2026 | Ref: AP-2026-6532

TABLE OF CONTENTS

DIAGNOSTIC REPORT

Risk Score: 81/100 CRITICAL

EXECUTIVE SUMMARY

Executive Summary

This Verified Audit identifies six material compliance gaps across the organisation's AI stack, representing significant regulatory exposure under the EU AI Act, NIST AI RMF, ISO 42001:2024, and SOC 2 Type II. The overall risk score of 81/100 places this organisation in the Critical Risk category. The most severe findings relate to the absence of human oversight on customer-facing automated decisions, a complete lack of immutable data lineage records, and AI systems operating outside the scope of the CMDB lifecycle. Combined, these gaps create estimated fine exposure of up to EUR 35M under the EU AI Act and material liability under enterprise customer contracts. Immediate board-level remediation is required across all six findings. This report has been reviewed and attested by the AuditPulse Verified Audit methodology team.

FRAMEWORK SCORECARD

EU AI ACT
38 /100
FAIL

NIST AI RMF
52 /100
REVIEW

ISO 42001
44 /100
FAIL

SOC 2
61 /100
REVIEW

Legal Exposure Summary

Finding	Framework	Max Exposure
No Human Oversight	EU AI Act Art. 83	**EUR 35M / 7% turnover**
No Data Lineage	ISO 42001 / SOC 2	**Contract liability**
AI Systems not in CMDB	EU AI Act Art. 83	**EUR 15M / 3% turnover**
No DPIA Conducted	GDPR Art. 83	**EUR 20M / 4% turnover**
Stale Risk Assessment	SOC 2 CC7.3	**Material**
No AI Cyber Insurance	SOC 2 CC9.1	**Unquantified**

Combined maximum regulatory exposure: EUR 70M+

KEY REGULATORY FINDINGS

[CRITICAL] No Human Oversight on Customer-Facing Automated Decisions

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

The organisation's AI systems make automated decisions that directly affect customers — including credit-adjacent recommendations, content personalisation affecting access, and automated account actions — without any mandatory human review checkpoint. EU AI Act Article 14 requires that high-risk AI systems enable qualified persons to monitor, intervene, override, or halt system outputs. Customer-facing automated decision systems that affect people's access to services qualify as high-risk under Annex III. Operating these systems without documented human oversight is a direct violation of Article 14 and exposes the organisation to enforcement action.

Required Action:

1. Conduct an immediate classification of all customer-facing AI decision workflows against EU AI Act Annex III high-risk criteria

2. Design and implement mandatory human review checkpoints for all confirmed high-risk decision pipelines before outputs are actioned

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

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

5. Validate override capability through a structured exercise and retain evidence for audit

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

[CRITICAL] Absence of Immutable Data Lineage for Core Inference Models

Citation: ISO 42001:2024 §8.4; SOC 2 CC6.1; EU AI Act Art. 10

No immutable record exists of data lineage for the organisation's production inference models. The organisation cannot demonstrate to regulators, auditors, or enterprise customers where training data originated, whether it was processed lawfully, or how data changes between versions affected model behaviour. ISO 42001:2024 Section 8.4 requires documented AI data management activities. SOC 2 CC6.1 requires traceable audit records of how sensitive inputs are handled. Without lineage records the organisation cannot mount a credible defence in a regulatory inquiry or model bias dispute.

Required Action:

1. Deploy a data lineage tool such as OpenLineage, Apache Atlas, or MLflow Lineage capable of write-once tamper-evident logging

2. Retroactively document known data sources, preprocessing steps, and transformation logic for all current production models

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

4. Define a lineage retention policy with minimum 5-year retention for audit defensibility

5. Include lineage records in the next SOC 2 Type II evidence package

REGULATORY EXPOSURE
Material - contract liability and certification revocation risk

KEY REGULATORY FINDINGS

[CRITICAL] AI Systems Not Registered as Configuration Items in CMDB

Citation: ISO 42001:2024 §6.4; EU AI Act Art. 11; NIST AI RMF GOVERN 1.1

Production AI models are not represented as configuration items in the organisation's CMDB. There is no technical documentation linked to any AI system's record covering intended use, known limitations, training data characteristics, version history, or approved operating scope. This means AI systems are invisible to change management, impact analysis, and audit. ISO 42001 Clause 6.4 requires documented information about AI system design and intended use. EU AI Act Article 11 requires technical documentation for high-risk AI systems. Neither requirement is currently satisfied.

Required Action:

1. Define a CI class for AI systems in the CMDB with attributes covering model version, training data reference, intended use scope, performance thresholds, and approved deployment contexts
2. Register all production AI models as CIs with complete attribute population
3. Design relationship mappings between AI CIs and the services, processes, and applications they influence
4. Establish a change management workflow for AI model updates including retraining, version deployment, and scope changes
5. Include AI CI documentation in the next ISO 42001 and SOC 2 audit evidence packages

REGULATORY EXPOSURE
Up to EUR 15M or 3% of global turnover - EU AI Act Art. 83

[CRITICAL] No Data Protection Impact Assessment Conducted for AI Systems

Citation: GDPR Art. 35; EU AI Act Art. 9; ISO 42001 §6.1

The organisation processes personal data through AI systems and has not conducted a Data Protection Impact Assessment. GDPR Article 35 mandates a DPIA prior to processing that is likely to result in a high risk to individuals. AI systems making automated decisions about people — particularly where those decisions affect access to services — are explicitly listed as processing types requiring a DPIA. The absence of a DPIA means the organisation has not formally assessed privacy risks, identified mitigations, or obtained the required sign-off before processing commenced.

Required Action:

1. Conduct a DPIA for each AI system that processes personal data or makes automated decisions affecting individuals
2. Use the structured DPIA template recommended by your lead supervisory authority
3. Document identified risks, proposed mitigations, and residual risk assessments
4. Obtain sign-off from the Data Protection Officer or equivalent privacy lead
5. Establish a trigger process for DPIA re-assessment when AI systems are materially updated

REGULATORY EXPOSURE
Up to EUR 20M or 4% of global turnover - GDPR Art. 83

[FLAGGED] AI Risk Assessment More Than 12 Months Out of Date

Citation: SOC 2 CC7.3; EU AI Act Art. 9(1); NIST AI RMF GOVERN 1.1

The organisation's last formal AI risk assessment was conducted more than 12 months ago. The documented risk posture does not reflect current model versions, expanded use cases, or the regulatory changes that have occurred since the assessment. SOC 2 CC7.3 requires continuous risk monitoring and periodic formal reassessment. EU AI Act Article 9 mandates that the risk management system be an ongoing iterative process throughout the AI system lifecycle.

Required Action:

1. Initiate a formal AI risk assessment within 45 days covering all production AI systems
2. Use NIST AI RMF Map function as the assessment structure
3. Document findings in a risk register with named owners and remediation timelines
4. Obtain executive sign-off at CISO or CTO level with board reporting
5. Establish a semi-annual assessment cadence with triggered re-assessment on major model updates

KEY REGULATORY FINDINGS

REGULATORY EXPOSURE
Material - undermines SOC 2 continuous monitoring requirement

[FLAGGED] No Cyber Liability Insurance Coverage for AI-Specific Incidents

Citation: SOC 2 CC9.1; ISO 42001 §6.1; NIST AI RMF GOVERN 6.2

The organisation does not have cyber liability insurance that explicitly covers AI-related incidents including model failure, adversarial attack, biased output causing harm, or regulatory enforcement action triggered by AI system behaviour. Standard cyber liability policies typically exclude AI-specific incidents. As the organisation operates customer-facing AI systems in high-risk categories, the absence of specific AI incident coverage creates unquantified financial exposure.

Required Action:

1. Review existing cyber liability policy with your broker to identify AI-specific exclusions
2. Obtain quotes for AI-specific endorsements or standalone AI liability coverage
3. Ensure coverage includes regulatory enforcement defence costs, model failure liability, and third-party harm from biased outputs
4. Document the insurance programme in the AI risk register
5. Review coverage annually as AI operations scale

REGULATORY EXPOSURE
Unquantified - exposure depends on incident type and scale

STRATEGIC REMEDIATION ROADMAP

STRATEGIC REMEDIATION ROADMAP

01

Implement human oversight checkpoints across all customer-facing automated decision workflows to satisfy EU AI Act Article 14.

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

02

Conduct DPIA for all AI systems processing personal data. Obtain DPO sign-off and document residual risks.

FLAGGED | MEDIUM EFFORT | 3-4 WEEKS | GDPR ART. 35

03

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

FLAGGED | HIGH EFFORT | 6-8 WEEKS | ISO 42001 §8.4 / SOC 2 CC6.1

04

Register all AI systems as CMDB configuration items with complete attribute population and relationship mapping to dependent services.

FLAGGED | MEDIUM EFFORT | 4-5 WEEKS | ISO 42001 §6.4 / EU AI ACT ART. 11

05

Conduct formal AI risk assessment using NIST AI RMF covering all production models. Establish semi-annual cadence.

FLAGGED | MEDIUM EFFORT | 45 DAYS | SOC 2 CC7.3 / EU AI ACT ART. 9

06

Review and extend cyber liability insurance to explicitly cover AI-specific incidents including regulatory enforcement defence costs.

FLAGGED | LOW EFFORT | 2-3 WEEKS | SOC 2 CC9.1 / ISO 42001 §6.1

07

Update AI governance documentation across all six findings and cross-reference in SOC 2 and ISO 42001 evidence packages.

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

Board Recommendations

1. Establish an AI Governance Committee with named executive ownership and board reporting cadence. Assign a designated AI Risk Owner accountable for remediation progress across all six findings.

2. Approve emergency remediation budget for human oversight implementation and data lineage infrastructure. The cost of remediation is materially lower than the cost of enforcement.

3. Commission a formal DPIA immediately. Do not wait for the next scheduled privacy review. GDPR Article 35 obligations are not discretionary.

4. Place EU AI Act compliance on the board agenda for the next scheduled meeting. The August 2026 enforcement deadline requires board-level sign-off on the remediation plan.

Verified Audit Attestation

This report has been generated following the AuditPulse Verified Audit methodology, which applies a 22-question diagnostic across four regulatory frameworks: EU AI Act, NIST AI RMF 1.1, ISO 42001:2024, and SOC 2 Type II.

The findings contained in this report represent material compliance gaps identified through structured assessment of the organisation's responses to the AuditPulse Verified Audit diagnostic.

This report is issued as a governance documentation artifact. It is not legal advice. It is intended for use in board reporting, investor due diligence, and enterprise procurement contexts.

AuditPulse Verified Audit reports are mapped against four independently published regulatory frameworks and are reviewed against the AuditPulse methodology standard prior to delivery.

AuditPulse Methodology Team

getauditpulse.io
damon@getauditpulse.io

Methodology Version: AuditPulse Verified Audit v1.0
Classification: CONFIDENTIAL