

AI COMPLIANCE DIAGNOSTIC REPORT

54

OUT OF 100

ELEVATED

VERIFIED AUDIT

NovaMind Technologies Pty Ltd

AI Compliance Diagnostic Report

Prepared for: Sarah Chen, Chief Technology Officer

Generated: 6/8/2026 | Ref: AP-2026-2710

TABLE OF CONTENTS

EXECUTIVE SUMMARY 3

FRAMEWORK SCORECARD 3

LEGAL EXPOSURE SUMMARY 4

KEY REGULATORY FINDINGS 5

STRATEGIC REMEDIATION ROADMAP 8

BOARD RECOMMENDATIONS 9

VENDOR RESOURCES 10

ATTESTATION LETTER 11





Risk Score: 54/100 ELEVATED

EXECUTIVE SUMMARY

Executive Summary

NovaMind Technologies operates a multi-model AI stack - including OpenAI GPT-4, Anthropic Claude, and a proprietary recommendation engine - across a global customer base spanning Australia, the European Union, and the United Kingdom. This Verified Audit identifies five material compliance gaps that represent meaningful regulatory exposure under the EU AI Act, GDPR, NIST AI RMF 1.1, ISO 42001:2023, and SOC 2 Type II. The three Critical findings - absence of structured audit-trail logging, an incomplete Data Protection Impact Assessment, and unresolved model versioning gaps - create direct regulatory exposure of up to EUR 35M under the EU AI Act and EUR 20M under GDPR, and would not withstand a formal regulatory inspection or enterprise procurement due diligence review in their current state. AI governance is presently singularly owned by the CTO with no distributed accountability structure, creating a key-person dependency that conflicts with ISO 42001:2023 Clause 5.3 and is increasingly scrutinised by enterprise procurement teams as a disqualifying risk. NovaMind's positive controls - active bias detection, human-in-the-loop oversight on high-risk decisions, and current AI risk assessment cadence - provide a strong foundation. Addressing the five findings within the recommended timeframes will materially reduce regulatory exposure and position NovaMind competitively in enterprise sales cycles where SOC 2 Type II and ISO 42001 certification evidence is now a procurement prerequisite.

FRAMEWORK SCORECARD



Legal Exposure Summary

Finding	Framework	Max Exposure
Unstructured Model Input/Output Logging Does Not Satisfy EU AI Act Article 12 Structured Audit Trail Requirements	EU AI Act Art. 12(1) - failure to maintain structured, queryable, tamper-evident audit logs for AI system operations across EU-resident user interactions	Up to EUR 35M or 7% of global annual turnover
Data Protection Impact Assessment for AI Systems Has Not Been Completed Despite Processing EU Resident Personal Data at Scale	GDPR Art. 35 / Art. 83(4) - mandatory DPIA not conducted prior to high-risk automated processing of EU resident personal data	Up to EUR 20M or 4% of global annual turnover
Data Subject Access Request Handling for AI-Generated Outputs Remains Entirely Manual, Creating GDPR Article 22 Scalability Risk	GDPR Art. 22(3) / Art. 83(4) - inability to reliably fulfil data subject rights within statutory 30-day window for AI-related processing disclosures	Up to EUR 20M or 4% of global annual turnover

Refer to full findings section for detailed remediation steps.

[CRITICAL] Unstructured Model Input/Output Logging Does Not Satisfy EU AI Act Article 12 Structured Audit Trail Requirements

Citation: EU AI Act Art. 12(1); NIST AI RMF MG-2.2; SOC 2 CC7.2

NovaMind confirmed that all model inputs and outputs are logged, but in an unstructured format without defined schema, query capability, or tamper-evidence controls. EU AI Act Article 12 mandates that logging systems for AI capture events in a structured, queryable manner sufficient to enable post-hoc audit, incident reconstruction, and regulatory inspection. Given that NovaMind processes requests from EU-resident users through GPT-4, Claude, and its proprietary recommendation engine, the EU AI Act applies in full to these interactions and the current logging architecture would not survive a regulatory audit or discovery request. This gap also undermines SOC 2 CC7.2, which requires that security events be logged in a format that enables investigation and response.

Required Action:

1. Define a structured logging schema for all model inputs and outputs including fields for timestamp, model version, session ID, input hash, output hash, and inference latency
2. Migrate existing logging infrastructure to emit structured JSON to a centralised, tamper-evident log store with append-only write controls
3. Implement a minimum 12-month log retention policy aligned with NovaMind's documented data retention schedule and EU AI Act requirements
4. Deploy log integrity verification such as cryptographic chaining or append-only object storage to satisfy EU AI Act Art. 12 immutability requirements
5. Validate the new structured logging pipeline against a simulated regulatory data request before the next SOC 2 Type II audit window

REGULATORY EXPOSURE

Up to EUR 35M or 7% of global annual turnover under EU AI Act Art. 71(4) for failure to maintain compliant logging systems

[CRITICAL] Data Protection Impact Assessment for AI Systems Has Not Been Completed Despite Processing EU Resident Personal Data at Scale

Citation: GDPR Art. 35; EU AI Act Art. 9(2)(a); ISO 42001:2023 §6.1

NovaMind has initiated but not completed a Data Protection Impact Assessment for its AI systems. Under GDPR Article 35, a DPIA is mandatory prior to processing that is likely to result in high risk to individuals, and AI systems making automated recommendations that affect access to services - as NovaMind's recommendation engine does - meet this threshold explicitly under the Article 35(3) categories. Operating production AI systems that process EU resident personal data without a completed DPIA means NovaMind cannot demonstrate prior risk identification to supervisory authorities. Any enforcement action or data breach occurring before DPIA completion would be treated as an aggravating factor in penalty calculations under GDPR Article 83.

Required Action:

1. Appoint a DPIA lead and formally scope the assessment to cover all AI models in production including the GPT-4 integration, Claude integration, and proprietary recommendation engine
2. Complete a necessity and proportionality assessment documenting what categories of personal data may transit through model inputs and what safeguards are in place for each
3. Identify and document residual risks with explicit named risk owners and mitigation decisions signed off at CTO and DPO level
4. Submit the completed DPIA for independent review by a qualified privacy professional before formal sign-off and filing with the lead supervisory authority if required
5. Establish an annual DPIA refresh cycle with triggered re-assessment on any material change to model stack, data sources, or processing geography

REGULATORY EXPOSURE

Up to EUR 20M or 4% of global annual turnover under GDPR Art. 83(4) for failure to conduct mandatory DPIA

[CRITICAL] No Documented Model Versioning or Rollback Process Exists Across GPT-4, Claude, and the Proprietary Recommendation Engine

Citation: EU AI Act Art. 9(5); ISO 42001:2023 §8.4; SOC 2 CC8.1

NovaMind confirmed that no documented model versioning register or tested rollback procedure exists across its production AI stack. Without a formal versioning register and validated rollback process, NovaMind cannot demonstrate controlled change management to auditors, enterprise customers, or regulators following a model-related incident. This gap directly undermines SOC 2 Type II Change Management controls (CC8.1) and conflicts with ISO 42001:2023 Clause 8.4, which requires documented AI system lifecycle controls. Under EU AI Act Article 9, organisations deploying high-risk AI systems must maintain a quality management system that governs model updates - the absence of versioning documentation means this obligation is not met.

Required Action:

1. Create a model version register documenting all production model identifiers, deployment dates, configuration snapshots, and prompt engineering versions for GPT-4, Claude, and the proprietary recommendation engine
2. Define and document a formal rollback procedure including trigger criteria, approval authority, and maximum tolerable rollback time for each production model
3. Integrate version tagging into NovaMind's CI/CD pipeline so every model weight, configuration, or prompt change is automatically logged with a unique version identifier
4. Conduct a tabletop rollback exercise within 60 days to validate that the documented procedure works under realistic incident conditions and retain evidence for SOC 2
5. Assign explicit ownership of the version register to a named engineering role and schedule quarterly reviews as part of AI governance cadence

REGULATORY EXPOSURE

SOC 2 Type II material finding: customer contract liability and audit remediation costs; ISO 42001:2023 certification revocation risk

[FLAGGED] AI Governance Is Concentrated Solely in the CTO Role Without a Distributed Accountability Structure or Formal Succession Plan

Citation: ISO 42001:2023 §5.3; NIST AI RMF GOVERN 1.1; SOC 2 CC1.3

NovaMind's AI governance is owned exclusively by the CTO, creating a single point of failure inconsistent with ISO 42001:2023 Clause 5.3, which requires defined and distributed AI governance roles across the organisation. In the event the CTO is unavailable, there is no documented succession or delegation framework for AI risk decisions, policy updates, or incident response authority. This concentration also makes it difficult for NovaMind to demonstrate independent oversight to enterprise customers during procurement due diligence, where separation of duties in AI governance is increasingly expected as a condition of contract.

Required Action:

1. Document a formal AI governance accountability matrix (RACI) assigning ownership of ethics policy, risk assessment, incident response, and vendor oversight to named roles beyond the CTO
2. Designate a deputy AI governance owner such as a Senior Engineer, Head of Product, or Legal Counsel with documented delegation authority during CTO unavailability
3. Establish a quarterly AI governance review meeting with at least two named participants and a written agenda circulated in advance
4. Incorporate AI governance responsibilities into relevant job descriptions to institutionalise accountability beyond a single role
5. Present the updated governance structure to the board within one quarter to ensure executive visibility and formal endorsement

REGULATORY EXPOSURE

ISO 42001:2023 certification revocation risk; enterprise customer procurement disqualification and contract loss risk

[FLAGGED] Data Subject Access Request Handling for AI-Generated Outputs Remains Entirely Manual, Creating GDPR Article 22 Scalability Risk

Citation: GDPR Art. 22(3); GDPR Art. 15(1)(h); EU AI Act Art. 13(1)

NovaMind's process for handling Data Subject Access Requests related to AI-generated decisions is entirely manual, which introduces both a GDPR Article 22 compliance risk and a practical operational bottleneck as the company scales across EU markets. GDPR requires that data subjects have the right to obtain meaningful information about automated processing logic, and a manual process cannot reliably guarantee the 30-day statutory response window as request volumes grow. This gap creates inconsistency risk - manual handling is prone to incomplete disclosures - which could expose NovaMind to supervisory authority complaints from EU residents and associated enforcement action.

Required Action:

1. Map all data flows where AI model inputs or outputs could be linked to an identifiable EU-resident individual and document the categories of personal data held per interaction
2. Build or procure a lightweight DSAR intake workflow that captures requestor identity, verification, and request type and routes it to the appropriate team with SLA tracking
3. Create standardised response templates for AI-related DSARs that explain model logic in plain language aligned with EU AI Act Article 13 transparency obligations
4. Set a 20-business-day internal SLA for AI-related DSARs with escalation to the CTO if a response is likely to breach the 30-day statutory deadline
5. Review the DSAR process quarterly and automate data retrieval steps as logging infrastructure matures to structured format

REGULATORY EXPOSURE

Up to EUR 20M or 4% of global annual turnover under GDPR Art. 83(4); reputational risk from supervisory authority complaints in EU markets

01

Migrate model input/output logging to structured, tamper-evident format with defined schema and 12-month retention.

CRITICAL

HIGH EFFORT

0-8 WEEKS

EU AI ACT ART. 12 / SOC 2 CC7.2

02

Complete Data Protection Impact Assessment for all AI systems in production. Obtain DPO sign-off and file with supervisory authority if required.

FLAGGED

MEDIUM EFFORT

0-6 WEEKS

GDPR ART. 35 / EU AI ACT ART. 9

03

Document and test model versioning and rollback procedures for all production models with CI/CD integration and tabletop exercise.

FLAGGED

MEDIUM EFFORT

0-10 WEEKS

ISO 42001 §8.4 / SOC 2 CC8.1

04

Distribute AI governance accountability via a formal RACI matrix and designate a deputy AI governance owner with documented delegation authority.

FLAGGED

LOW EFFORT

0-4 WEEKS

ISO 42001 §5.3 / NIST AI RMF GOVERN 1.1

05

Build a structured DSAR intake and response workflow with 20-business-day internal SLA and AI-specific plain-language response templates.

FLAGGED

MEDIUM EFFORT

4-8 WEEKS

GDPR ART. 22 / EU AI ACT ART. 13

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 42001:2023, 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

Want help implementing these remediations?

AuditPulse can connect you with verified specialists. Vendor matching is free.

getauditpulse.io/marketplace