iDharma

SAMPLE

INDEPENDENT · EXPERT-LED · STANDARDS-BASED

AI RISK & READINESS ASSESSMENT

AI Audit Report

Client: Klaro (illustrative)

Audit tier: Quick Scan

Prepared by: iDharma — Brijesh Patel, Founder & CEO

Date: 25 May 2026

This is a sample report. It was produced for a fictional company, ‘Klaro,’ to show exactly what a real iDharma Quick Scan report looks like — its structure, depth, and tone. Klaro is not a real client and the findings are illustrative. A real report follows this same shape, built entirely from your own systems and documents.

1 Executive summary

Klaro is in a common position: three AI systems shipped quickly, all useful, none yet checked. The systems themselves work — but the controls around them have not kept pace, and one of them carries real regulatory exposure.

The single most important finding: **the Applicant Summary tool ranks job applicants**, which places it in the **high-risk category under the EU AI Act**. Combined with applicants not being told AI screens them, and a recruiter who “mostly trusts the ranking,” this is the area to address first.

Across all three systems, customer and applicant data is being sent to third-party AI providers with no confirmed data-processing agreement — a gap that matters both for security and for GDPR, given Klaro’s EU users.

None of this requires rebuilding the AI. It requires putting governance, disclosure, and oversight around it — work measured in days, not months. The five priorities are in Section 4.

2 Scope & approach

What was assessed: three AI systems — Klaro Assist (customer support), Applicant Summary (hiring), and Smart Search.

What was not assessed: no live security or penetration testing, no bias testing of the hiring model’s outputs (that is Risk Audit scope), and no review of systems not listed above.

Standards referenced: EU AI Act, NIST AI Risk Management Framework, ISO/IEC 42001.

Method: the iDharma AI Audit methodology — Scope, Gather, Assess, Report, Readout — working a structured checklist across eight domains.

Information basis: the completed intake form and documents provided by Klaro on 25 May 2026 (privacy policy and product description; AI policy, model documentation, data-flow diagrams and AI-vendor agreements were not available).

3 Findings

F-01 Applicant Summary tool is high-risk under the EU AI Act

CRITICAL

SYSTEM: APPLICANT SUMMARY · DOMAIN: LEGAL & REGULATORY EXPOSURE

What we found. The tool summarises and ranks job applicants to inform shortlisting. Under the EU AI Act, AI used in recruitment and candidate evaluation is classified high-risk, which carries obligations around risk management, transparency, human oversight, record-keeping, and accuracy. Klaro has EU users and likely EU applicants, so this applies.

Why it matters. High-risk obligations are substantive. Operating a high-risk AI system without them is a regulatory exposure — and exactly the kind of thing an enterprise procurement questionnaire is probing.

Recommendation. Treat the Applicant Summary tool as high-risk. Confirm EU applicant exposure with counsel, and either bring it up to the EU AI Act's high-risk requirements or restrict its use until you can.

F-02 Customer and applicant data sent to AI providers with no confirmed DPA

CRITICAL

SYSTEM: ALL THREE · DOMAIN: DATA & PRIVACY

What we found. Customer messages and résumés are sent to third-party AI APIs. Klaro is "not sure" whether a data-processing agreement (DPA) is in place with either provider.

Why it matters. Personal data is leaving Klaro's control with no confirmed contractual protection. For Klaro's EU users this is a direct GDPR gap, and it is a question any serious enterprise customer will ask.

Recommendation. Confirm and put in place a DPA with both AI providers (both offer them), and confirm whether your data is excluded from model training. Document it.

F-03 Job applicants are not told AI screens their application

HIGH

SYSTEM: APPLICANT SUMMARY · DOMAIN: TRANSPARENCY & DISCLOSURE

What we found. There is no indication applicants are told that an AI tool summarises and ranks their application.

Why it matters. This is both an EU AI Act transparency obligation for high-risk AI and a basic fairness expectation. Its absence compounds F-01.

Recommendation. Add a clear notice in the application process that AI assists in reviewing applications, and how.

F-04 Weak human oversight of the hiring tool

HIGH

SYSTEM: APPLICANT SUMMARY · DOMAIN: HUMAN OVERSIGHT

What we found. The recruiter "reads the summary but mostly trusts the ranking." There is no defined step where a human meaningfully reviews or can overturn the AI's ranking.

Why it matters. For a system that affects who gets a job, "mostly trusts the AI" is not meaningful human oversight — and the EU AI Act requires it for high-risk systems.

Recommendation. Define a real review step: the recruiter assesses applicants on documented criteria, with the AI summary as input only, and records the decision.

F-05 No AI policy and no one accountable

HIGH

SYSTEM: ALL THREE · DOMAIN: GOVERNANCE & ACCOUNTABILITY

What we found. No written AI policy or principles. No one is formally accountable for AI — the CTO handles it "by default."

Why it matters. “Everyone, informally” means no one. Without an owner and a policy, gaps like the ones in this report have nobody whose job it is to catch them.

Recommendation. Name an accountable owner for AI, and adopt a short AI policy covering approved use, data handling, disclosure, and human oversight. This can be one to two pages.

F-06 No quality measurement or monitoring after launch

MEDIUM

SYSTEM: ALL THREE · DOMAIN: MODEL & PERFORMANCE

What we found. None of the systems has a measure of accuracy or reliability, and none is monitored after launch. The known incident — Klaro Assist giving a wrong refund answer — surfaced only via a customer complaint.

Why it matters. Without monitoring, the business learns about AI failures from customers, not before them.

Recommendation. Define a simple quality check for each system and a lightweight way to monitor and log failures.

F-07 No guardrails or defined response for incorrect output

MEDIUM

SYSTEM: KLARO ASSIST · DOMAIN: SAFETY & GUARDRAILS

What we found. Klaro Assist gave a customer incorrect refund information. There are no guardrails on sensitive topics (e.g. refunds, billing) and no defined response when the assistant is wrong.

Why it matters. A support assistant confidently stating wrong policy can create real disputes and cost.

Recommendation. Constrain the assistant on sensitive topics (route to a human, or answer from approved text only) and define how errors are caught and corrected.

F-08 Users may not know they are interacting with AI

MEDIUM

SYSTEM: KLARO ASSIST, SMART SEARCH · DOMAIN: TRANSPARENCY & DISCLOSURE

What we found. It is not confirmed that customers are clearly told Klaro Assist is an AI assistant, or that Smart Search results are AI-ranked.

Why it matters. The EU AI Act requires people be told when they are interacting with AI. It is also a low-cost trust signal.

Recommendation. Add a clear “AI assistant” label to Klaro Assist; note that search results are AI-ranked.

Note. AI service credential and key security could not be assessed from the intake. We recommend confirming API keys are stored securely and access is logged — flagged for a quick check.

4 Prioritised action list

#	ACTION	CLOSES	SEVERITY	SUGGESTED OWNER
1	Put DPAs in place with both AI providers; confirm data isn’t used for training.	F-02	Critical	CTO

#	ACTION	CLOSES	SEVERITY	SUGGESTED OWNER
2	Treat Applicant Summary as EU AI Act high-risk; confirm exposure with counsel; restrict use until compliant.	F-01	Critical	CTO + counsel
3	Add a real human-review step for the hiring tool; tell applicants AI is used.	F-03, F-04	High	Hiring lead
4	Name an AI owner; adopt a short AI policy.	F-05	High	Founders
5	Add AI disclosure labels; add guardrails and monitoring on Klaro Assist.	F-06, F-07, F-08	Medium	CTO

5 Framework snapshot

FRAMEWORK	OVERALL STANDING	NOTES
EU AI Act	Gaps	Applicant Summary is high-risk and not yet meeting high-risk obligations; AI-interaction disclosure missing.
NIST AI RMF	Partial	“Govern” is the weak point — no policy, no owner. “Measure” / “Manage” are weak — no monitoring.
ISO/IEC 42001	Gaps	AI is run informally, not as a managed practice. A basic policy and owner would move this materially.

6 Next steps

Immediate. Actions 1 and 2 — the two Critical items.

Getting help. If it would help to have these fixed rather than just listed, iDharma can match Klaro with an independently verified AI consultant for the remediation work.

Deeper assurance. Given the hiring tool’s high-risk status, a Compliance Audit would be the logical next step before relying on it at scale — and a Risk Audit would add the bias and fairness testing Klaro asked about.