Security, GDPR, and EU AI Act Compliance Report

Prepared for: Ardona AI **Date:** February 2026

---

1. Executive Summary

This document outlines the data processing, AI integration, and regulatory compliance posture of the Ardona AI application. It specifically addresses compliance with the General Data Protection Regulation (GDPR) regarding data privacy, and the European Union Artificial Intelligence Act (EU AI Act) regarding the deployment of AI systems.

Ardona AI operates as a B2B SaaS platform that generates strategic workflow automation audits. By design, the application minimizes the collection of Personally Identifiable Information (PII) and strictly segregates user authentication data from AI processing pipelines.

---

2. System Architecture & Data Flow

2.1 The Tech Stack

- **Frontend:** Next.js (React)
- **Backend:** Next.js Serverless API Routes
- **Authentication & Database:** Firebase Authentication & Firestore (via `firebase-admin` SDK)
- **AI Engine:** Google Gemini API (Paid Service Tier)

2.2 Data Segregation Principle

The core security architecture relies on strict data segregation.

- **Auth Data & PII:** User email addresses, names (from onboarding and auth), LinkedIn profiles, and encrypted passwords never touch the application's backend AI logic. They are handled securely by Firebase Authentication and Firestore.
- **Business Data:** Information regarding a user's company (industry, software tools used, team size, bottlenecks, operational challenges) is stored in Firestore, tied to an anonymized Firebase User ID (`uid`).
- **AI Payload Data:** Only public or generic business data is transmitted to the AI engine.

---

3. GDPR Compliance Posture

3.1 What Data We Collect

1. **Personally Identifiable Information (PII):** Email Address and Name (via Auth), plus **First Name, Last Name, and LinkedIn Profile URL** (collected during the onboarding questionnaire).
2. **Corporate / B2B Data (Collected during Onboarding):** Company Website URL, Industry, Product/Service Description, Team Size, Desired "Magic Hire", Time-Consuming Processes, Revenue Bottlenecks, Mistake-Prone Areas, Avoided Tasks, Client Complaints, Current Software Tools, AI Usage/Wishlist, and Past Automation Investments.

3.2 What Data We Send to the AI (Google Gemini)

When a user generates an "AI Audit", the application sends specific prompts to the Google Gemini API. **The payload sent to Google INCLUDES:**

- The target public Website URL.
- Publicly scraped website content (Title, Meta Description, visible body text).
- Generic business metrics provided by the user (e.g., "SaaS Industry", "10-50 employees", "Marketing department").
- System instructions on how to format the JSON response.

The payload sent to Google STRICTLY EXCLUDES:

- User Email Addresses.
- User Names.
- Passwords or Auth Tokens.
- Firebase User IDs (UIDs).

3.3 GDPR Legal Assessment

Under GDPR, "Company Name" and generalized business metrics are recognized as **Corporate Data**, not *Personally Identifiable Information (PII)*. Because Ardona AI does not transmit PII to third-party AI processors, the risk of a GDPR data breach via the AI pipeline is effectively zero. The processing of public business data for service delivery falls under the "legitimate interest" legal basis.

3.4 Google API Data Processing & Training

Ardona AI accesses the Gemini AI through a Google Cloud Project with an active billing account, qualifying it as a **"Paid Service"** under Google's Terms of Service.

Under the Google Paid Services agreement:

1. **Google acts as a Data Processor:** They process the prompts solely to generate the response.
2. **Zero Training Policy:** Google explicitly **does not** use your prompts or the generated responses to train, fine-tune, or improve their AI models.
3. **Data Sovereignty:** Temporary logging of prompts is done strictly for automated safety/abuse detection (e.g., preventing illegal content generation) and is not accessible for commercial training.

---

4. EU AI Act Compliance Posture

4.1 Risk Classification

The EU AI Act classifies AI systems into four risk tiers (Unacceptable, High, Limited, Minimal).

Ardona AI utilizes Generative AI to analyze business structures and output strategic recommendations (documents, matrices, roadmaps).

- It does **not** process HR/Recruiting data to make automated hiring decisions.
- It does **not** operate in critical infrastructure (medical, legal, law enforcement).

Therefore, Ardona AI is classified as a **Limited / Minimal Risk AI System**. It is exempt from the heavy regulatory burdens (Conformity Assessments, Quality Management Systems) required of High-Risk systems.

4.2 Transparency & Bot Disclosure Obligations

For Limited Risk generative AI systems, the EU AI Act mandates **Transparency**. Users must be explicitly informed that they are consuming content generated by an AI, not a human.

How Ardona AI Complies:

1. **Brand Identity:** The application is explicitly named "Ardona AI".

2. **Feature Naming:** The core deliverable is labeled an "AI Readiness Audit".
3. **Explicit Disclaimers:** The final generated PDF report includes a hardcoded disclaimer on the cover page:

   > *"Disclaimer: This audit was generated using Ardona's AI engine. While we strive for accuracy, AI-generated content can contain inaccuracies. This report is a strategic starting point and should be verified before final implementation."*

These measures comprehensively satisfy the "Bot Disclosure" transparency requirements of the EU AI Act.

---

5. General Security Measures

1. **Authentication:** All private API endpoints are secured using `verifyAuth()`, which cryptographically validates the Firebase JWT.
2. **No IDORs (Broken Access Control):** All database queries are strictly scoped to the authenticated user's `uid`, preventing cross-tenant data access.
3. **Rate Limiting:** Public onboarding AI endpoints are protected by an IP-based rate limiter (capped at 10 requests per minute) to prevent Denial of Service (DoS) and financial resource exhaustion attacks by malicious bots.
4. **No Client-Side Exposed Secrets:** Sensitive keys (like `GEMINI_API_KEY`) remain securely on the server and are never exposed to the browser.