Tucker Payment Processing · Standards self-assessment

EU AI Act (Regulation (EU) 2024/1689)

Where Tucker sits in the EU's risk tiers, and the duties that apply · Self-assessment · November 24, 2026

3 met **1 partly met** **3 not applicable**

About this document

This is a self-assessment written by Michael Cunningham, who designed and built the site. It is not legal advice. Tucker Payment Processing is a fictional US company with no EU customers, so the AI Act would not apply to it in practice. This document asks a more useful question: if a real payment company deployed this same chatbot to customers in the EU, what would the Act require, and does the design already meet it?

The Act sorts AI by risk. Some practices are banned outright. "High-risk" systems face the heaviest duties. Systems that talk to people, or generate content, have transparency duties. The rest have few or none. The dates reflect the Digital Omnibus on AI (Regulation (EU) 2026/1744), in force since July 27, 2026, which pushed back some deadlines.

Classification

Prohibited practices (Article 5)

NOT APPLICABLE

Neither AI does anything on the banned list, such as manipulative techniques that cause harm, social scoring or exploiting people's vulnerabilities. Tucker is a support assistant; he cannot persuade anyone to spend money, and he has no tool that could.

High-risk systems (Article 6, Annex III)

NOT APPLICABLE

Payment companies need to watch this one. Annex III lists AI used to assess a person's creditworthiness or set their credit score as high-risk (fraud detection is excluded). Tucker makes no decision about anyone: no credit, eligibility, pricing or fraud decisions. He only reads a verified customer's own records and explains policies. If a real company later let a chatbot like this one decide disputes or offer credit, it would become high-risk. Under the Digital Omnibus, the Annex III duties now apply from December 2, 2027.

Transparency: telling people they are talking to an AI (Article 50(1))

MET

This is the duty that applies, from August 2, 2026: people must be told they are dealing with an AI unless it is obvious. Tucker's greeting in every chat says he is an AI assistant, his label on the home page reads "AI Support Agent", and his instructions require him to say he is an AI if asked. The AI SOC Squirrel is labeled "AI Security Analyst", and its incident report names it as the author.

Evidence: `site/index.html`, `site/treats.html`, `site/checkout.html`, `app/chat.py`.

Transparency: marking AI-generated content (Article 50(2))

PARTLY MET

Providers of AI that generates text must also mark it as AI-generated in a machine-readable way, where technically feasible, from August 2, 2026 (December 2, 2026 for systems already on the market before then). Every reply Tucker's model writes carries an "ai_generated: true" mark in the data the website receives; the app's own fixed messages, such as "Tucker's taking a nap", are marked false. The AI SOC Squirrel's incident report PDF names the AI as its author in the file's metadata. There is no settled industry standard yet for marking chat text, so a real company would adopt the European Commission's code of practice on marking AI-generated content once it is final.

Evidence: `app/chat.py`, `tests/test_guardrails.py`, `site/assets/report-pdf.js`.

Other duties

AI literacy (Article 4) — MET

Companies that provide or deploy AI should make sure the people working with it understand it. The Digital Omnibus softened the original wording, but the expectation remains. The site's builder and only operator wrote its AI policies, designed its guardrails and documents how each AI works and fails.

Evidence: `policies/ai-assistant-policy.md`, `site/about.html`.

General-purpose AI model duties (Articles 53 to 55) — NOT APPLICABLE

These fall on the company that makes the underlying model (Google, for Gemini), not on a company that builds a chatbot on top of it. A real deployer would keep Google's published documentation for the model on file.

Data protection alongside the Act — MET

The AI Act sits on top of the GDPR, which would apply to a real EU deployment. The design already practices data minimization: no IP addresses are sent to the AI, chat text is not stored with IP addresses, web logs are kept for at most two days, the only cookie is strictly necessary (so no cookie banner is needed), and the model provider does not train on the chats.

Evidence: `policies/data-retention.md`, `deploy/logrotate-apache2`, `app/limits.py`.

The bottom line

As built, Tucker is a limited-risk AI system. Its main duty, telling people they are talking to an AI, is met in every chat. The content-marking duty is partly met while the EU finalizes its marking standards. The design choice that keeps it out of the high-risk tier is the same one that makes it secure: the AI decides nothing about anyone.

Source: https://github.com/mcunningham-infosec/tucker-payment-processing/blob/main/docs/standards/eu-ai-act.md · Live site: https://tuckerpaymentprocessing.com
Written by Michael Cunningham with AI assistance (Claude). A self-assessment, not an audit, certification or legal opinion. Tucker Payment Processing is a fictional company and processes no real payments.