PAYOMETRY

The EU AI Act Merchant Checklist

Is the AI already running your fraud, chargeback, and risk decisions exposing you?

If any part of your payment stack uses AI or machine learning — your own models or a vendor's — the EU AI Act may already apply to you, even if you're a U.S.-based merchant or processor with European customers or card volume. This checklist walks through where AI hides in a typical payments operation, what's likely to be classified high-risk, and the concrete steps to get ahead of it before enforcement lands.

Step 1: Does the AI Act Apply to You?

[ ] You process card payments, chargebacks, or credit decisions for customers based in the EU — even occasionally.

[ ] You (or a vendor you use) run any model — rules-based, machine learning, or otherwise — to score transactions, flag fraud, predict chargebacks, or evaluate creditworthiness.

[ ] You've never asked a processor or fraud-tool vendor whether their AI features are classified as high-risk under the Act.

If you checked any of these, keep reading — the Act reaches beyond EU-headquartered companies to anyone whose AI system's output is used in the EU.

Step 2: Where AI Hides in a Typical Payments Stack

Area	Common AI Use
Fraud & risk scoring	Real-time transaction scoring, velocity checks, device fingerprinting models
Chargeback management	Dispute prediction, auto-response drafting, representment recommendation engines
Underwriting & onboarding	Merchant risk tiering, credit/creditworthiness scoring, KYB automation
Customer support	Chatbots, automated dispute intake, AI-drafted customer communications
Vendor-embedded AI	"Smart routing," auto-optimization, or scoring features bundled into your gateway or processor

A model built by a vendor and dropped into your checkout flow counts the same as one you built yourself — the Act follows the use case, not who wrote the code.

Step 3: High-Risk or Not? Quick Reference

Use Case	Likely Status
Credit scoring / creditworthiness of a natural person	**High-risk (Annex III)**
Pure fraud detection (narrowly scoped to identifying fraud)	Carved out — but easy to lose if the model also scores trustworthiness
Chargeback / dispute-outcome prediction affecting a consumer	Likely high-risk — treat as in-scope until confirmed otherwise
Customer-facing chatbots and automated support	Limited-risk — transparency disclosure required (Art. 50)

The fraud-detection carve-out is narrower than most merchants assume: the moment a "fraud" model starts influencing a customer's credit standing or account access, it can slide into high-risk territory. When in doubt, scope each model by exactly what decision it makes.

Step 4: The Compliance Checklist

Inventory & Classification

[ ] List every AI/ML model touching your payment flow — fraud scoring, chargeback prediction, dispute resolution, underwriting, chatbots, even vendor "smart routing."

[ ] For each model, document what decision it makes and who it affects (merchant-only vs. consumer-facing).

[ ] Confirm whether any model touches creditworthiness or credit scoring of a natural person — that's presumptively high-risk.

[ ] Confirm whether any "fraud detection" model has drifted into scoring customer trustworthiness — the carve-out is narrow and easy to lose.

Documentation & Data Governance

[ ] Maintain technical documentation for each high-risk system: purpose, logic, training data sources, performance metrics.

[ ] Document data lineage — where training, validation, and test data came from and how it's kept current.

[ ] Log model versions and material changes over time.

Human Oversight

[ ] Confirm a human can review, override, or halt an automated decision (a declined transaction, a held chargeback, a frozen account).

[ ] Define who that person is, and make sure they're trained on what the model does and doesn't do.

Vendor Due Diligence

[ ] Ask every processor, gateway, and fraud-tool vendor whether their AI features are classified high-risk under Annex III.

[ ] Get written confirmation of what compliance work the vendor has done versus what falls on you as the deployer.

[ ] Add AI Act compliance language to new and renewed vendor contracts.

Transparency

[ ] Disclose to customers when they're interacting with an AI system — chatbots, automated decisioning — per Article 50.

[ ] Update privacy notices to reflect AI-driven decisions in your payment or dispute flow.

Monitoring & Ownership

[ ] Set calendar reminders for the enforcement deadlines below.

[ ] Assign one internal owner for AI Act compliance — don't leave it homeless between legal, risk, and engineering.

Key Dates to Know

Date	What Happens
August 2, 2026	Original deadline for high-risk Annex III obligations and Article 50 transparency rules. A May 2026 "Omnibus" proposal would push high-risk enforcement to December 2027, but until that delay is formally adopted, treat this date as live.
December 2, 2027	Proposed new deadline for high-risk Annex III systems under the Omnibus deal — not yet final law.

What Non-Compliance Costs

High-risk non-compliance can carry fines of up to €15 million or 3% of global annual turnover, whichever is higher. National regulators also have the power to pull a non-compliant AI system out of the EU market entirely — which, for a payments company, can mean losing the ability to process EU transactions with that tool at all.