A Systems Architecture Guide for Founders, PMs & Developers

Privacy & AI Law — S06E05

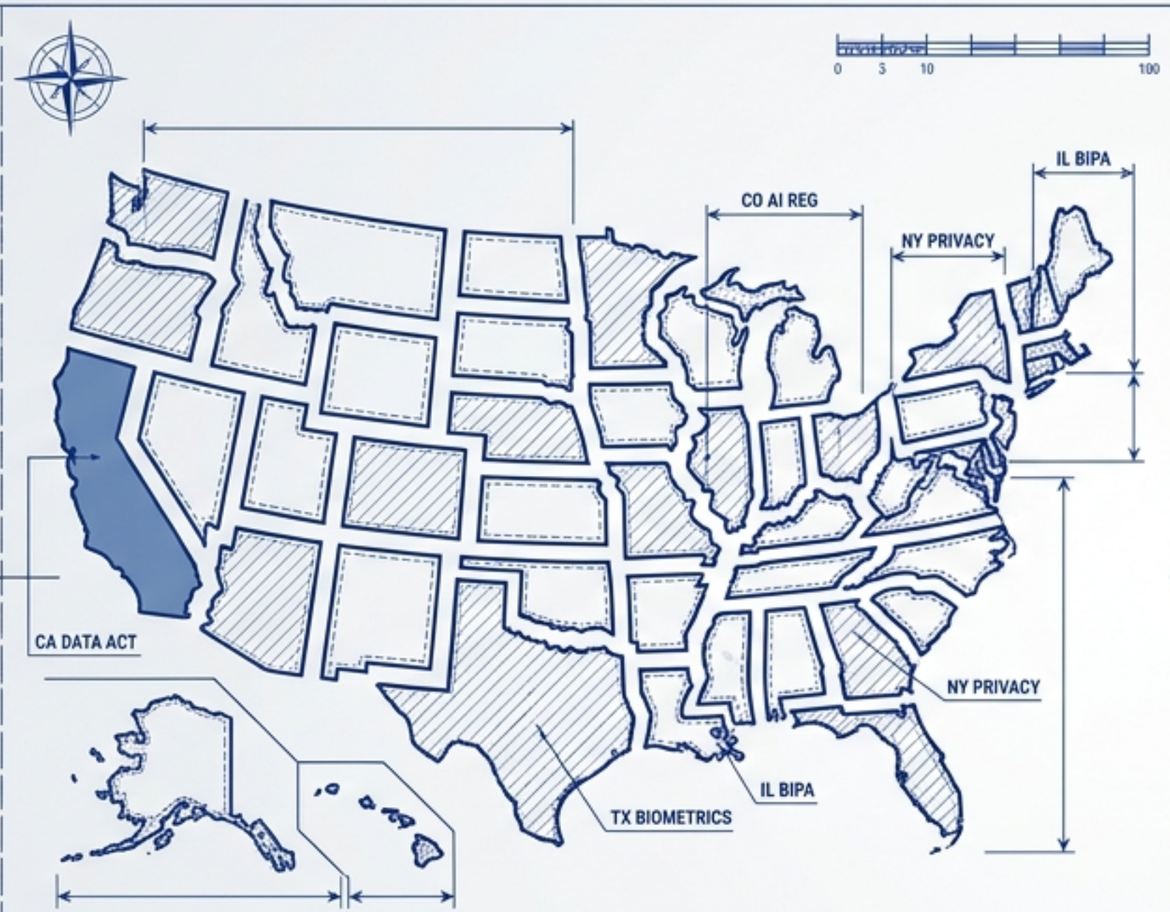
The EU AI Act Explained

[Risk Tiers] [GPAI] [Nov 2026 Phased Status]

⚠ General information, not legal advice — current as of November 2026.

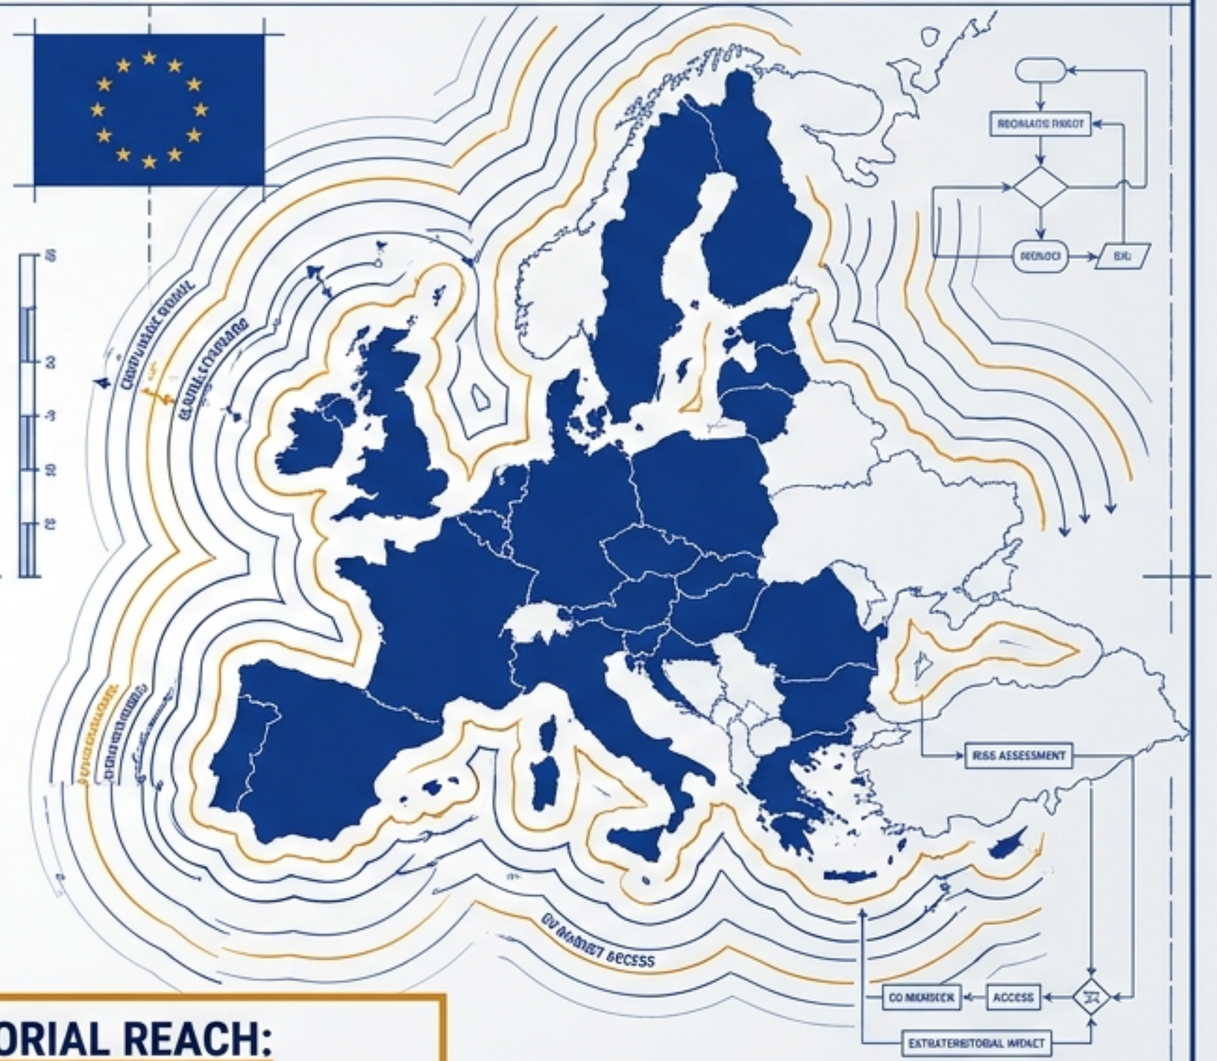
Gemini Notebook

THE US PICTURE - S06E04 RECAP



A fragmented, state-by-state patchwork focused primarily on data privacy rather than systemic AI product risk.

THE EU REALITY - REG (EU) 2024/1689 & 2026/1744



Comprehensive, use-case risk regulation. The EU does not ban LLMs or neural networks; it regulates how they are applied.

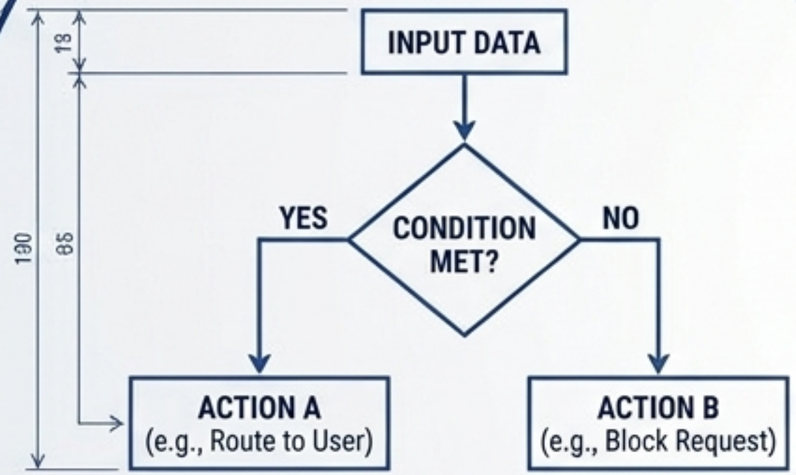
EXTRATERRITORIAL REACH:

If your system's outputs are used in the EU, you are in scope—even if your servers and headquarters are in Silicon Valley.

Gemini Notebook

What is an ‘AI System’? (Art. 3(1))

Out of Scope
Deterministic / Rule-Based

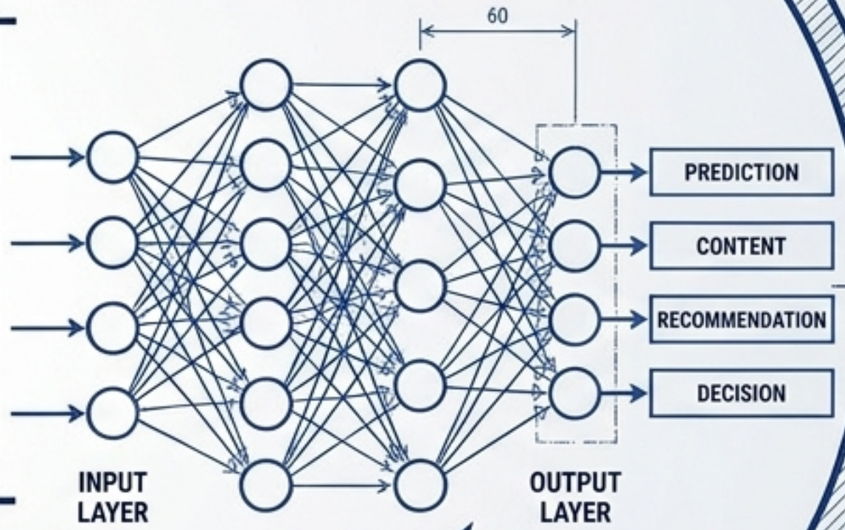


Standard search filters, basic SaaS routing

Commission Guidelines trigger:

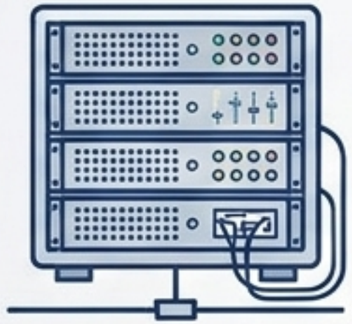
A machine-based system that infers, from the input it receives, how to generate outputs (predictions, content, recommendations, decisions) that influence physical or virtual environments.

In Scope - Art. 3(1)
AI System
Inferential & Adaptive



Machine learning, deep learning, complex pattern recognition

The AI Supply Chain: Provider vs. Deployer



The Model Creator

[Art. 3 'Provider']

Builds the AI system or puts its own brand/trademark on it.

Carries the heaviest conformity and registration burdens.

API Call / Integration



The SaaS Integrator

[Art. 3 'Deployer']

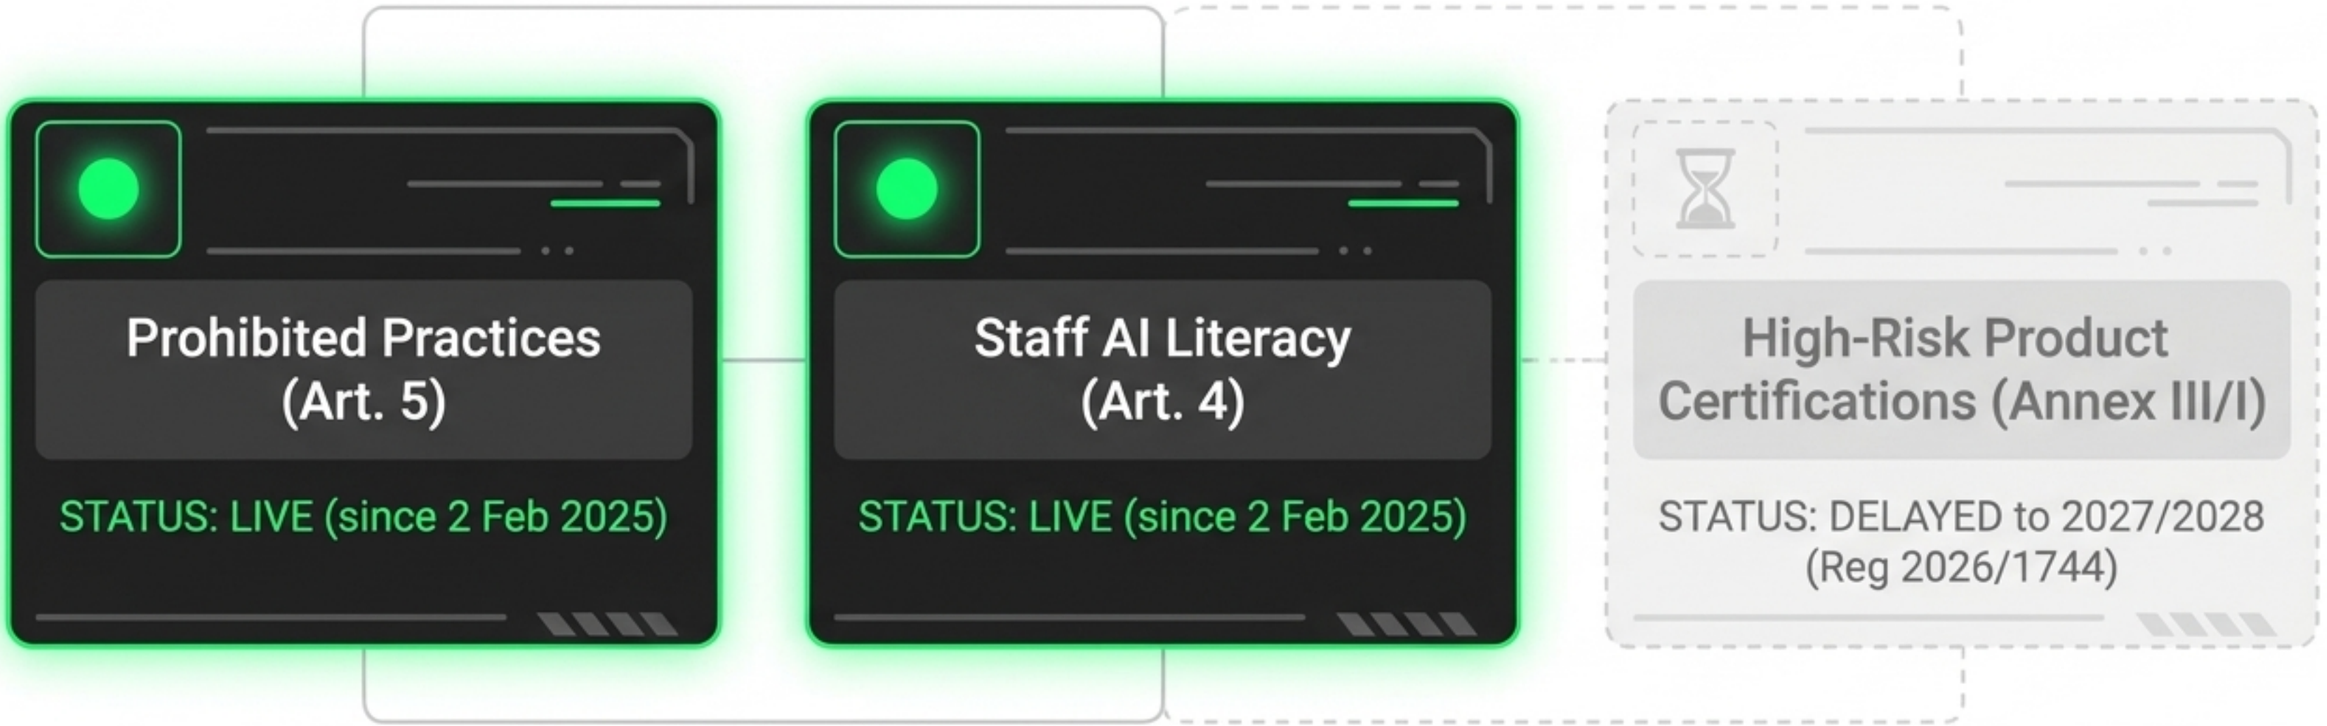
Integrates or uses the AI system in a professional business workflow.

Carries lighter oversight, logging, and transparency duties.

WARNING: Substantially modifying a model or changing its intended purpose legally flips you from a Deployer into a Provider.

Gemini Notebook

The Nov 2026 Reality Check: Don’t Wait for 2027



While high-risk product certifications were delayed by the July 2026 Digital Omnibus, the absolute bans and internal team requirements are actively enforced today.

Gemini Notebook

The 8 Outlawed Red Lines (Art. 5)



Subliminal/Manipulative Harm
(Bypassing user consciousness to materially distort behavior.)



Exploiting Vulnerabilities
(Targeting age, disability, or socio-economic distress.)



Social Scoring
(Detrimental treatment based on accumulated social behavior.)



Criminal-Risk Profiling
(Predictive policing based solely on personality traits.)



Untargeted Facial Scraping
(Building DBs from CCTV or mass internet scraping.)



Emotion Recognition
(Prohibited in workplaces and educational institutions.)



Biometric Categorisation
(Inferring sensitive traits: race, politics, sexual orientation.)



Real-Time Public Biometric ID
(Banned for law enforcement without strict narrow exceptions.)

(Plus the immediate ban on AI 'nudification' / CSAM ge

Gemini Notebook

Universal AI Literacy Duties (Art. 4)

Providers and Deployers must take measures to ensure staff dealing with AI have sufficient AI literacy.

The Digital Omnibus Update: Regulation 2026/1744 clarified that no specific ‘level’ or formal certificate is mandated, but training must be contextual to the risk of the tools used.

- **Assess role:** Are we providing or deploying?
- **Map risk:** What AI is used, and what are the specific failure modes?
- **Train contextually:** Do staff understand model hallucinations, data inputs, and when to override outputs?
- **Document:** Keep internal records of training initiatives and guidelines.

Gemini Notebook

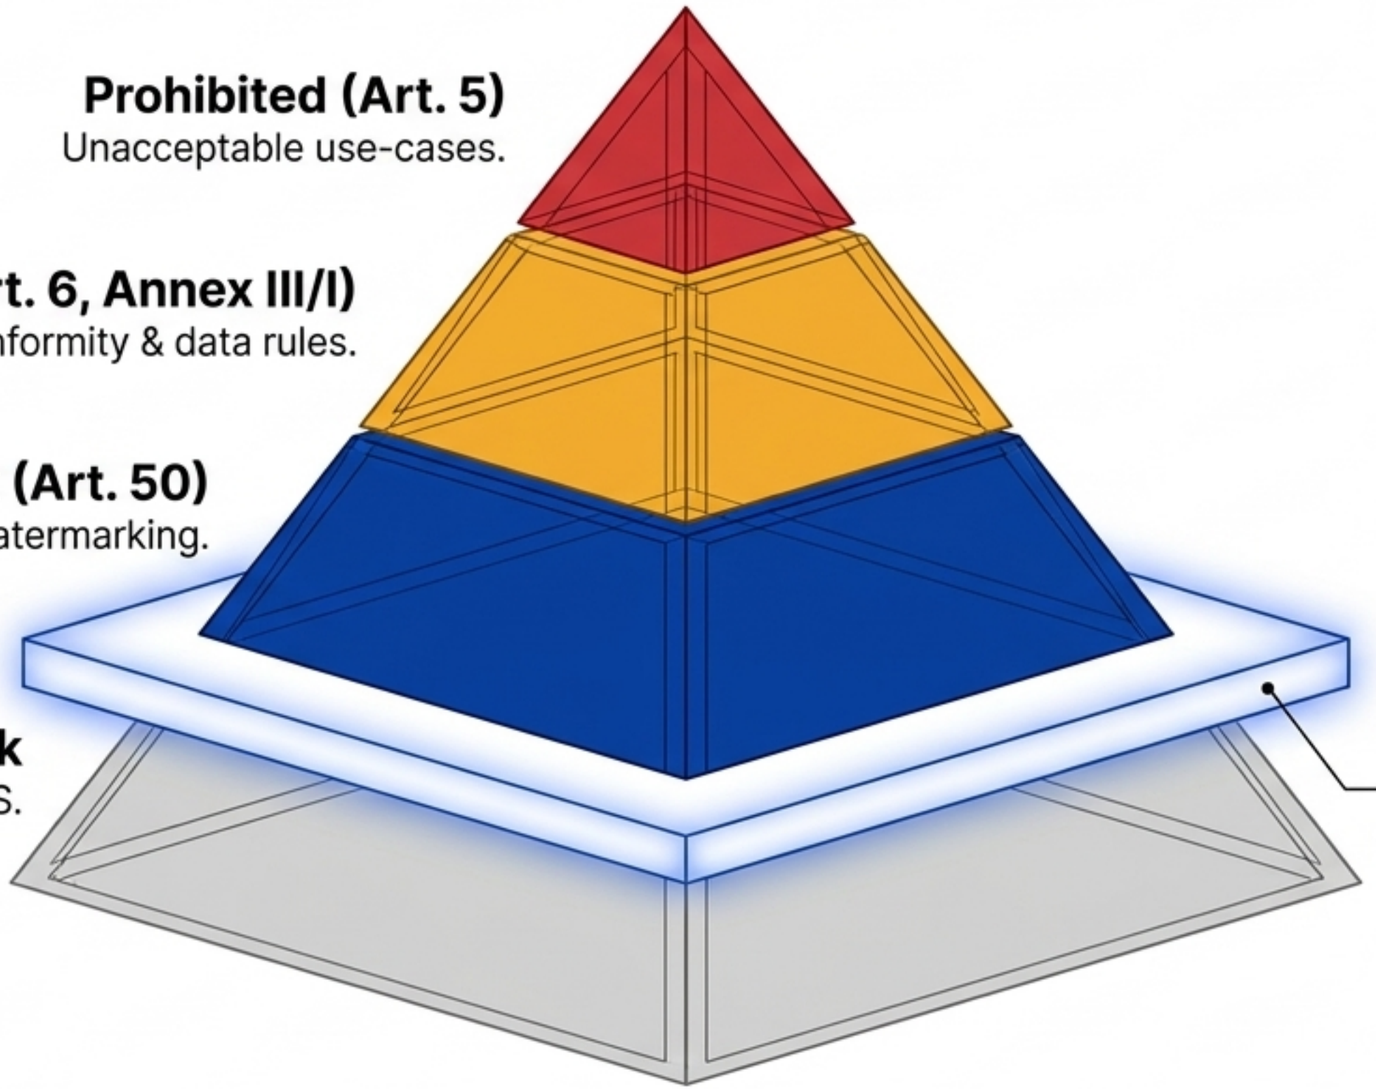
The EU AI Act Risk Pyramid

Prohibited (Art. 5)
Unacceptable use-cases.

High-Risk (Art. 6, Annex III/I)
Strict conformity & data rules.

Transparency Risk (Art. 50)
Disclosure and watermarking.

Minimal Risk
Vast majority of everyday SaaS.



General-Purpose AI (GPAI, Arts. 51-55)

The underlying foundation models that power the tiers above them.

Gemini Notebook

The Master Phased Timeline & Fine Matrix (Nov 2026 Status)

LIVE

2 Feb 2025
Art. 5 Bans & Art. 4 Literacy.

LIVE

2 Aug 2025
Arts. 51–55 GPAI Rules.

LIVE

2 Aug 2026
Art. 50 Transparency & Art. 57 Sandboxes.

2 Dec 2027
Annex III Stand-Alone High-Risk (per Reg 2026/1744).

2 Aug 2028
Annex I Regulated-Product High-Risk (per Reg 2026/1744).

Art. 99 Fine Caps

Category	Fine / Penalty
Banned Practices	Up to €35M or 7% global turnover.
High-Risk/GPAI Breaches	Up to €15M or 3% global turnover.
Incorrect Information	Up to €7.5M or 1% global turnover.

(SME proportionality applies).

Gemini Notebook

What Makes a Feature ‘High-Risk’? (Annex III & I)

Everyday SaaS becomes ‘High-Risk’ if applied in these specific domains:



Employment
CV screening, hiring, worker monitoring, task allocation.



Education
Student grading, admissions, automated proctoring.



Essential Services
Creditworthiness scoring, health/life insurance pricing.



Biometrics
Remote biometric identification (where permitted).



Critical Infrastructure
AI safety components in water, gas, electricity, traffic.



Law Enforcement/Migration
Assessing evidence, visa examination.

Gemini Notebook

The Art. 6(3) Carve-Outs (Not High-Risk)



THE CATCH: To claim this exemption, Providers must still document this assessment and register the system in the EU database before launching! (Unless it performs profiling, which is always High-Risk).

- An AI feature in a High-Risk domain is exempt if it does not materially influence human decision-making:
  1. Performing a narrow, non-material procedural task.
  2. Improving the result of a previously completed human activity.
  3. Detecting patterns without replacing human assessment.

Gemini Notebook

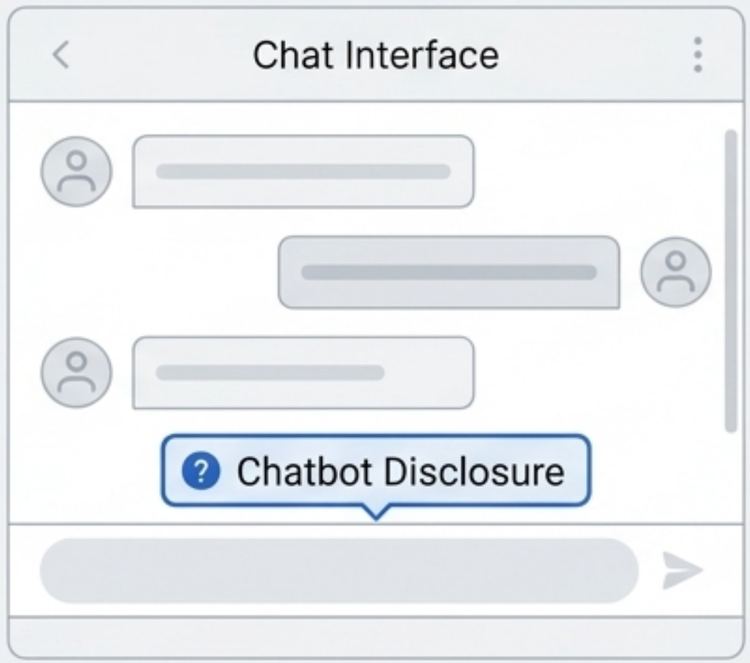
High-Risk Division of Labor: Provider vs. Deployer

High-Risk Provider (Arts. 8–25)	High-Risk Deployer (Art. 26)
• Establish a **continuous Risk Management System**. • Art. 10 **Data Governance** (bias mitigation, representative training data). • Draw up **technical documentation** and **instructions** for use. • Undergo **conformity assessment** and affix the **CE mark**.	• Implement Human Oversight (empower staff to override the AI). • Follow the Provider's instructions for use strictly. • Retain automatically generated logs for at least 6 months. • Notify workers/unions before deploying AI in the workplace.

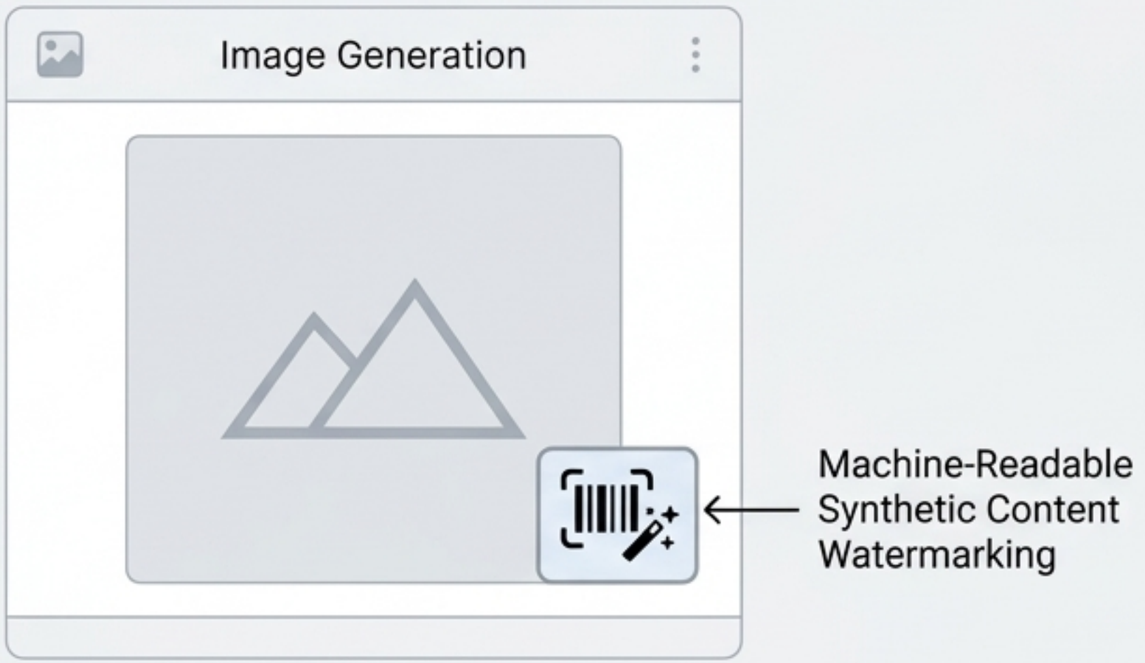
Gemini Notebook

Live Now: Transparency Duties (Art. 50)

Active since 2 August 2026 for all risk tiers.



Users must be clearly informed they are interacting with an AI system (unless obvious from context).



AI-generated audio, image, video, or text must be marked in a machine-readable format as artificially generated (aligning with the Code of Practice). Visible labels required for deepfakes.

Gemini Notebook

Minimal Risk AI: The Vast Majority of SaaS



If your AI feature is not in Annex III/I, and is not a prohibited practice, it is Minimal Risk.

Examples: AI-enabled spam filters, video game NPC logic, basic inventory forecasting, internal code-completion tools.

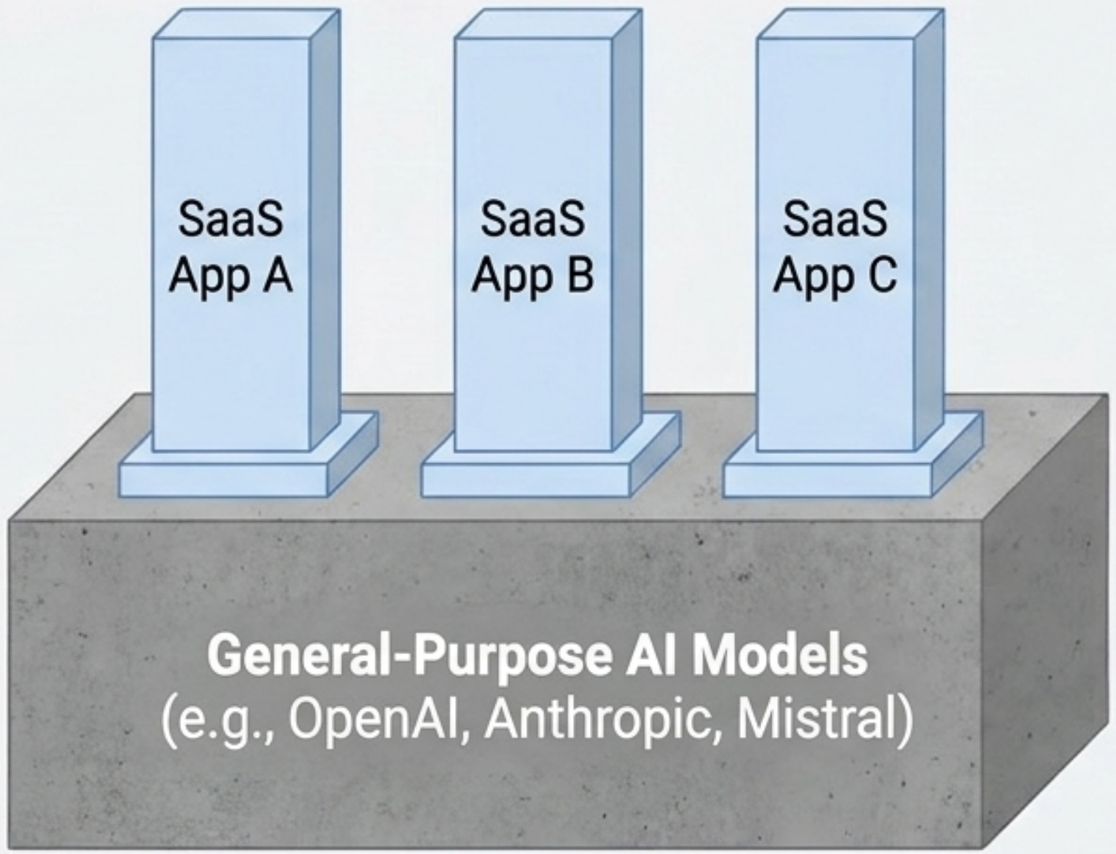
Legal Burden: No strict conformity or registration rules under the AI Act (though GDPR still applies to personal data).

Recommendation: The EU encourages adherence to voluntary Codes of Conduct (Art. 95) to build user trust without regulatory penalty.

Gemini Notebook

The GPAI Layer (Arts. 51–55)

Live since August 2025.



Definition: General-Purpose AI Models are models trained on massive data at scale that display “significant generality” and can competently perform a wide range of distinct tasks, designed to be integrated into downstream systems.

Scope: This layer regulates the creators of the foundation models, not necessarily the SaaS apps calling them via API.

Gemini Notebook

Upstream Duties & Systemic Risk

Standard GPAI (Baseline Art. 53 Duties)

- Provide technical documentation to downstream Deployers.
- Implement an EU copyright compliance policy.
- Publish a detailed public summary of training data (using the Art. 53(1)(d) template).

Narrow exemptions exist for true open-source models, except for copyright and data summaries.

Systemic Risk GPAI (> 10^25 FLOPs)

- **Triggered at 10^25 floating-point operations** (FLOPs) (Art. 55).
- **Mandatory adversarial testing** (Red-teaming).
- **Mandatory tracking and reporting of serious incidents.**
- Ensure **physical and digital cybersecurity** of the model infrastructure.

Gemini Notebook

Scenario: Downstream LLM Fine-Tuning

Startup calls a foundation LLM API to summarize internal meeting notes.

Status: Deployer / Minimal Risk (Liability rests upstream).

Startup uses the API to build an **automated HR CV-screening tool.**

Status Change: Startup is now a **High-Risk Deployer**. (Must implement human oversight and logging).

Startup puts its own brand on the HR tool and markets it externally to other companies.

Status Change: Startup is now a **High-Risk Provider**. (Needs CE mark and Art. 10 Data Governance).

Startup substantially modifies the weights and architecture of an open-source LLM.

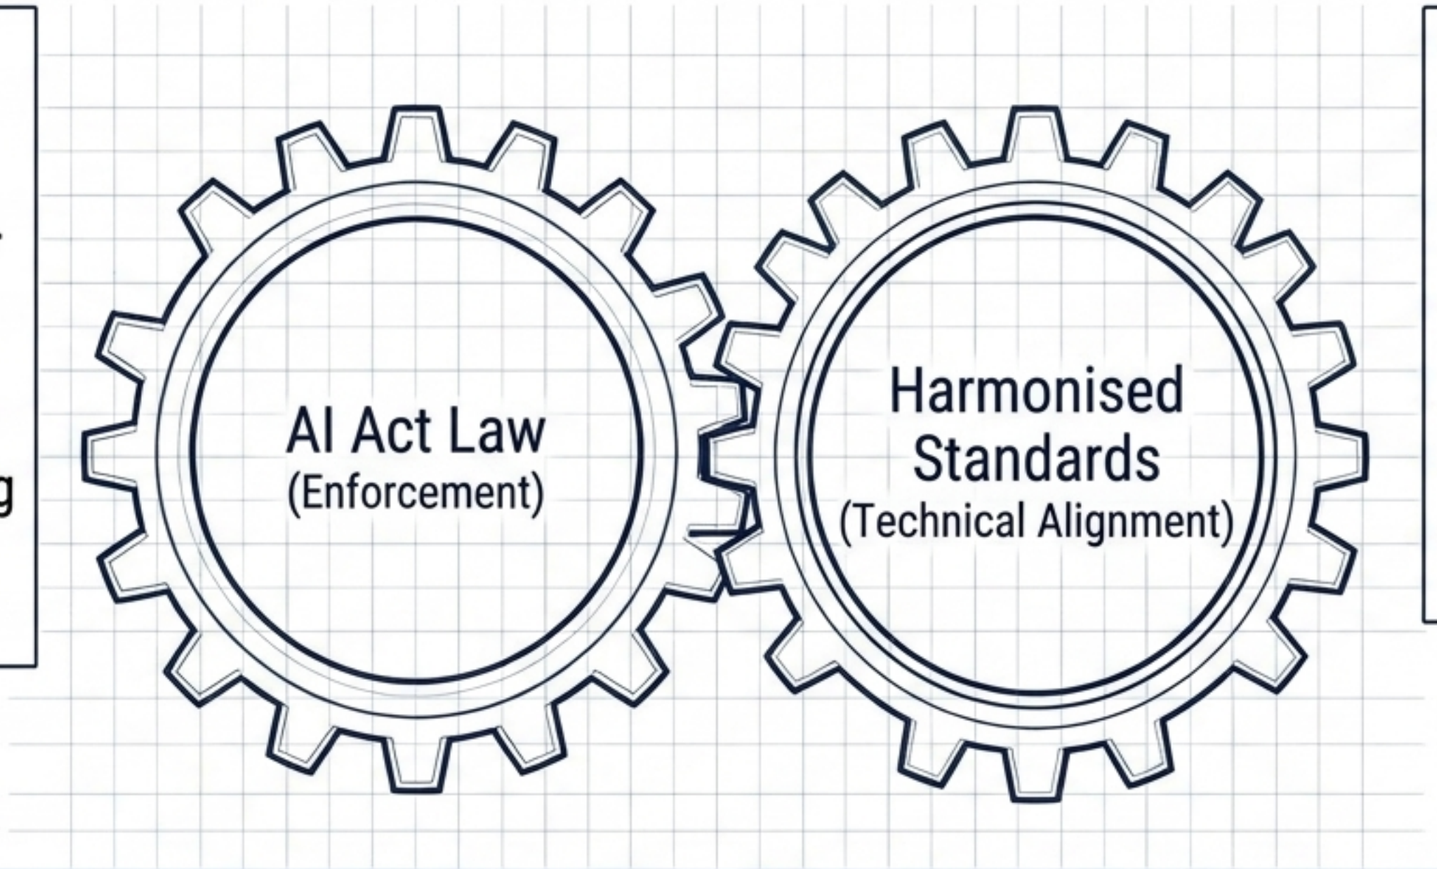
Status Change: Startup is now a **GPAI Provider**. (Arts. 51-55 apply).

Gemini Notebook

The July 2026 Digital Omnibus Shift (Reg 2026/1744)

The Problem

The original 2026 enforcement date for High-Risk systems arrived before the necessary technical standards and testing bodies were fully established.



The Solution

Regulation 2026/1744 delayed Annex III enforcement to **2 Dec 2027** and Annex I to **2 Aug 2028**.

The Strategy

This gives the market time to align with finalized harmonised standards, while keeping immediate pressure on Art. 5 Bans and Art. 50 Transparency (which remain live today).

Gemini Notebook

Boundary Questions & Grey Areas

Emotion vs. Sentiment (Art. 5)

Inferring emotions (anger, fatigue) in a workplace is banned. General sentiment analysis (routing angry tickets) is generally permitted. Where is the exact line on voice-tone analysis?



GPAI Fine-Tuning Thresholds

At what point does a downstream SaaS company's fine-tuning constitute a 'substantial modification' that flips them into being a GPAI Provider? The market relies on evolving Commission GPAI Guidelines.

Gemini Notebook

Regulator Overlap & The AI Pact

Enforcement
(EDPB Statement 3/2024)

National Data Protection Authorities (DPAs) will frequently overlap with new AI Market Surveillance Authorities. Dual-jeopardy on GDPR and AI Act fines is a live issue.

Safe Harbors (Art. 57)

Member states are launching SME AI Regulatory Sandboxes to test pre-market compliance.



The AI Office

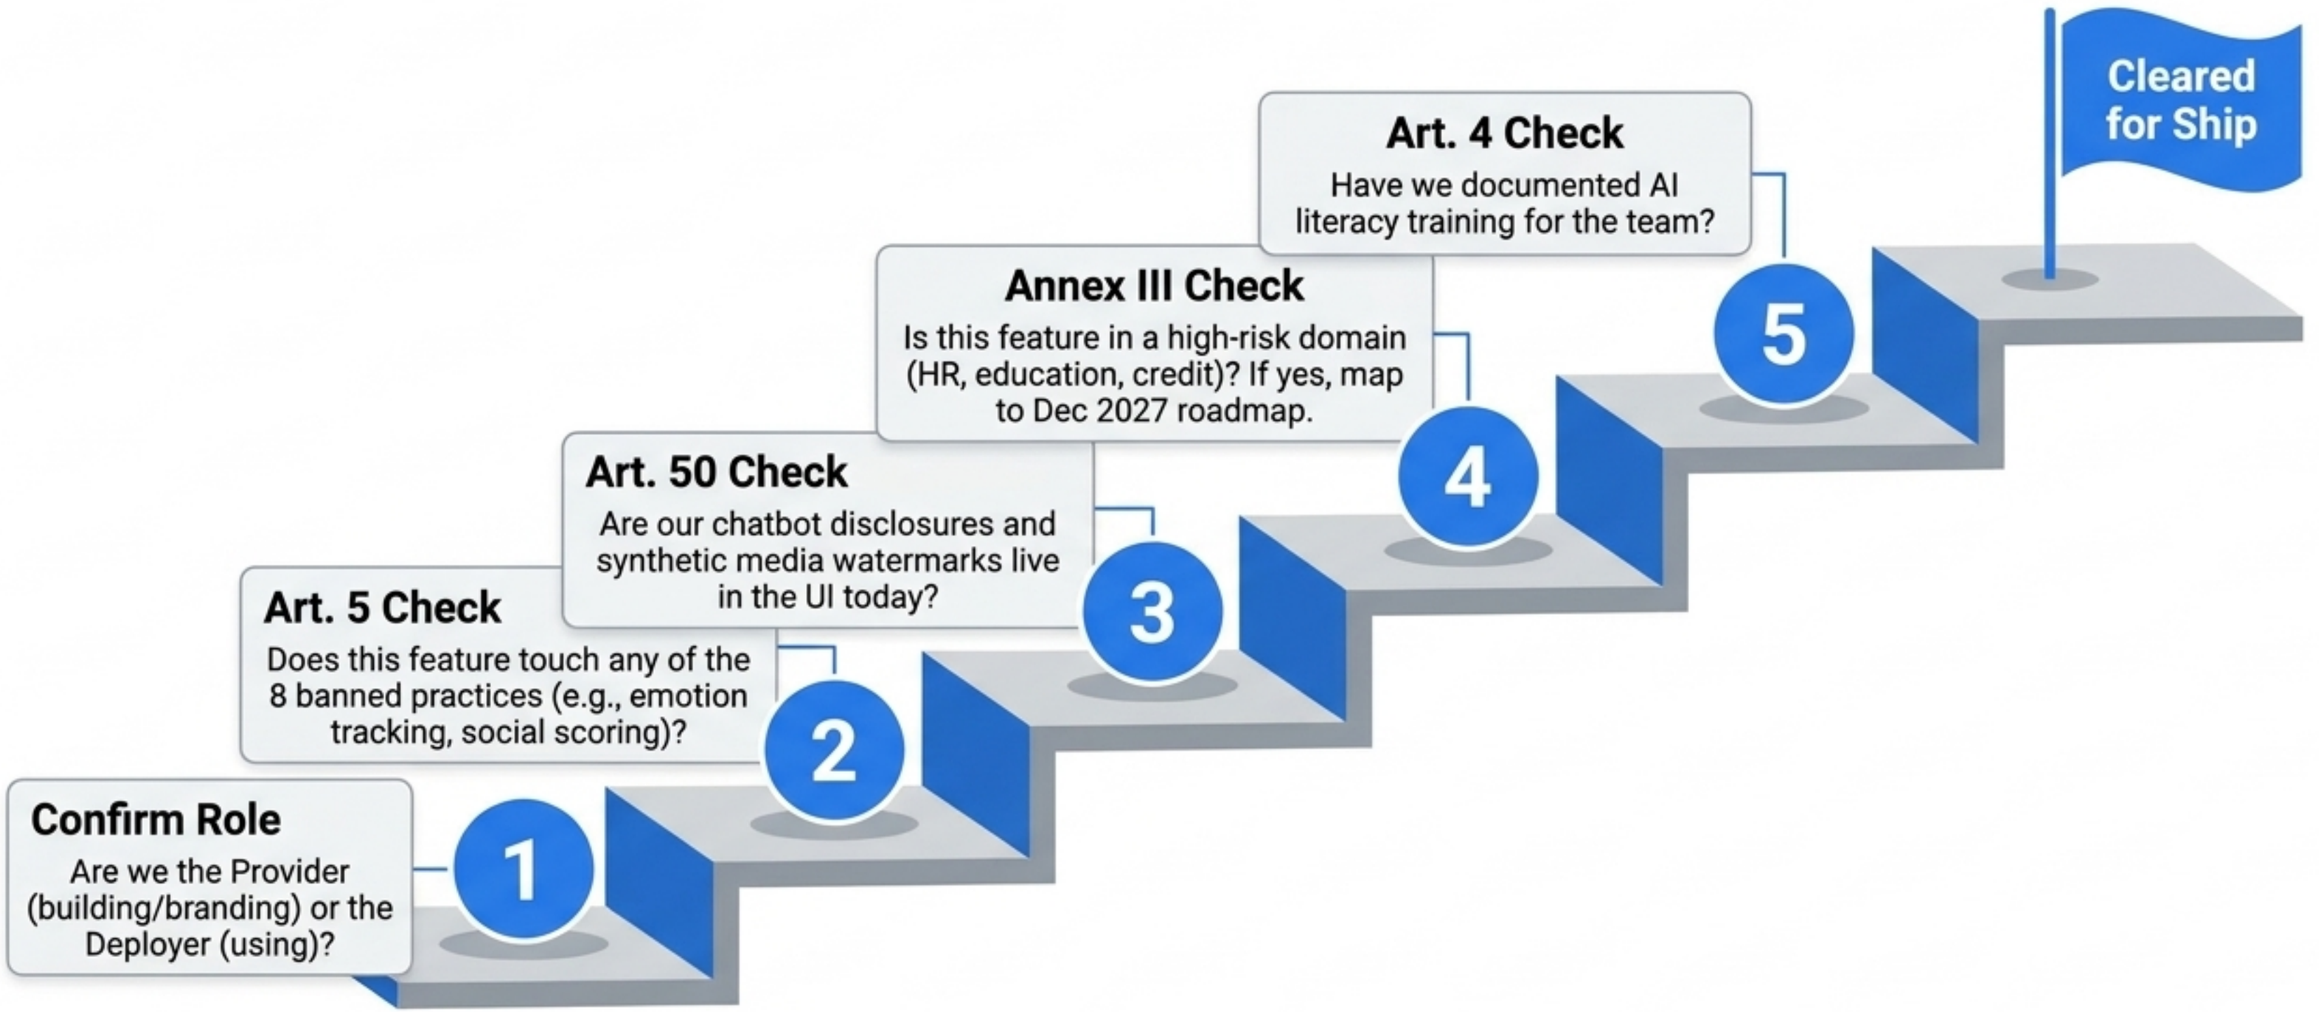
The central European body enforcing GPAI rules and systemic risk.

The AI Pact

A voluntary initiative by the Commission for early compliance, allowing companies to share best practices while standards finalize.

Gemini Notebook

Product Team AI Act Triage (Nov 2026)



Tool: Use the official EU AI Act Compliance Checker tool.

Gemini Notebook

Counsel Red Lines (When to Stop & Call Legal)



Escalate to specialist conformity/legal counsel **immediately** if your roadmap includes:

- Building or branding any feature in an Annex III or Annex I High-Risk domain (you will need a CE mark).
- Substantially modifying the architecture or weights of an open-source model.
- Any feature touching borderline Art. 5 practices (specifically biometrics, scraping, or emotion/intent inference).
- Deploying an AI system where a failure could result in physical injury or denial of essential services.

Gemini Notebook

Up Next in S06E06: “AI, Privacy and Consent”

How does the AI Act overlap with the GDPR? We explore the complex interplay of Article 22 (Automated Decision Making), the right to be forgotten inside an LLM, and navigating personal data in training sets.

⚠️ General information, not legal advice — current as of November 2026.

Gemini Notebook