RODER

EU AI ACT · READINESS GUIDE · JULY 2026

The EU AI Act, translated into controls.

What applies today, what the 2026 Omnibus actually changed, and the evidence your organisation needs to have in place before 2 December 2027.

ABOUT THIS GUIDE

The deadline moved. The obligations did not.

In June 2026 the EU formally agreed the Digital Omnibus on AI, pushing the high-risk obligations of the AI Act from August 2026 to December 2027. Much of the coverage stopped at the headline. This guide covers what the postponement actually changed, what already applies to your organisation today, and how to convert eighteen months of runway into controls and evidence rather than a longer pause.

2 Aug 2026
ARTICLE 50 TRANSPARENCY DUTIES APPLY, UNCHANGED BY THE OMNIBUS

2 Dec 2027
HIGH-RISK OBLIGATIONS FOR ANNEX III SYSTEMS, AFTER THE POSTPONEMENT

€35M / 7%
MAXIMUM PENALTY FOR PROHIBITED PRACTICES, WHICHEVER IS HIGHER

Who this is for. Security, compliance, data protection and platform leaders at European enterprises that deploy AI: approved assistants, AI features inside SaaS, internal copilots, agents and MCP-connected tools. Most of this guide addresses the deployer role, because that is what most enterprises are. The provider chapters matter if you build or substantially modify AI systems.

How to read it. Each obligation is translated three ways: what the regulation asks, what that means operationally, and what evidence you should be able to produce when asked. Compliance work that does not end in evidence is opinion.

The core argument. Nearly every duty in the Act resolves to the same five capabilities: an inventory of AI in use, enforceable policy, runtime logging, human oversight, and retained evidence. Build those once and the article-by-article mapping becomes bookkeeping instead of a programme restart.

This guide is general information, current as of July 2026. It is not legal advice. Figures reference Regulation (EU) 2024/1689 as amended by the Digital Omnibus on AI, adopted by the European Parliament on 16 June 2026 and by the Council on 29 June 2026.

01 · WHERE THE ACT STANDS

What the 2026 Omnibus changed, and what it left alone.

The AI Act entered into force on 1 August 2024 with staggered application dates. The Digital Omnibus on AI, agreed in May 2026 and formally adopted in June 2026, moved some of those dates. It did not soften the substance of the high-risk requirements, and it added a new prohibition.

POSTPONED

High-risk, stand-alone systems

Annex III systems (employment, credit, education, essential services, law enforcement and similar) now face the full high-risk regime from **2 December 2027** instead of 2 August 2026. The dates are fixed, not conditional on standards being ready.

POSTPONED

High-risk, embedded in products

AI embedded in regulated products under Annex I (medical devices, machinery, vehicles) moves from 2 August 2027 to **2 August 2028**. Regulatory sandboxes shift to August 2027.

UNCHANGED

Everything already in force

Prohibited practices (Article 5) and AI literacy (Article 4) have applied since **2 February 2025**. General-purpose AI model duties (Articles 51-56) have applied since **2 August 2025**. Penalty ceilings are untouched.

NEW

A new prohibition, and a near deadline

The Omnibus adds a ban on AI that generates non-consensual intimate imagery or child sexual abuse material, transitioning until **2 December 2026**. And Article 50 transparency still lands **2 August 2026**.

THE FOUR RISK TIERS, PLUS ONE OVERLAY

Prohibited

Social scoring, manipulation, untargeted face scraping, emotion recognition at work and school, NCII and CSAM generation.

High-risk

Annex III use cases and Annex I products. Risk management, data governance, logging, oversight, conformity.

Limited risk

Systems that interact with people or generate content. Disclose, label and mark under Article 50.

Minimal risk

Everything else. No new AI Act duties; GDPR, NIS2, DORA and sector rules still apply.

Overlay: general-purpose AI models carry their own duties (Articles 51-56) regardless of tier.

02 · TIMELINE

The dates that survive the headlines.

✓

2 FEBRUARY 2025 · IN FORCE

Prohibitions and AI literacy

Article 5 bans apply, with fines up to €35M or 7% of worldwide turnover. Article 4 requires a sufficient level of AI literacy in staff who operate AI systems. If your people use AI at work, this one already applies to you.

✓

2 AUGUST 2025 · IN FORCE

General-purpose AI models

Documentation, copyright policy and training-data summaries for GPAI providers; extra duties for models with systemic risk. Relevant to enterprises that fine-tune or redistribute models, not only to the big labs.

!

2 AUGUST 2026 · WEEKS AWAY

Article 50 transparency. Not postponed.

People must be told when they interact with an AI system. AI-generated and manipulated content must be machine-readably marked, and deepfakes disclosed. Systems already on the market get a watermarking grace period to 2 December 2026. The new NCII and CSAM prohibition also completes its transition on 2 December 2026.

27

2 DECEMBER 2027 · POSTPONED FROM AUG 2026

High-risk obligations, Annex III

The full high-risk regime for stand-alone systems: Articles 8-15 for providers, Article 26 for deployers, conformity assessment, registration. Seventeen months from now.

28

2 AUGUST 2028 · POSTPONED FROM AUG 2027

High-risk obligations, Annex I products

AI embedded in regulated products follows a year later, aligned with sector conformity regimes.

Why December 2027 is not a 2027 problem

The high-risk regime is evidenced by history: logs, risk reviews, oversight records, vendor documentation. Regulators will ask how systems behaved, not how they behaved after you noticed the deadline. Runtime evidence cannot be backfilled. An organisation that starts logging and governing in 2026 walks into December 2027 with eighteen months of proof; one that starts in autumn 2027 walks in with a slide deck.

03 · ROLES

Provider or deployer: the question that decides your workload.

The Act assigns duties by role, per AI system. A **provider** develops an AI system or has one developed and places it on the market under its own name. A **deployer** uses an AI system under its own authority. Most enterprises are deployers for most systems, and the deployer duty set is far lighter. But the role is decided system by system, and it can flip.

ROLE	CORE DUTIES FOR HIGH-RISK SYSTEMS	TYPICAL EXAMPLES
Provider	Articles 8-15 in full: risk management, data governance, technical documentation, logging capability, transparency to deployers, oversight design, accuracy and cybersecurity. Plus quality management (Art. 17), conformity assessment, CE marking, registration, post-market monitoring (Art. 72) and incident reporting (Art. 73).	You build an AI product, sell a model-backed feature, or white-label someone else's system under your brand.
Deployer	Article 26: use per the provider's instructions, assign competent human oversight, control input data, monitor operation, keep the system's logs at least six months, inform workers before workplace use, report serious incidents.	You roll out an AI assistant to staff, use an AI screening tool in hiring, or wire an agent into internal systems.

ROLE FLIP

When a deployer becomes a provider

Put your name or trademark on a high-risk system, substantially modify one, or repurpose a general system into a high-risk use, and you inherit the provider duty set. Fine-tuning a general-purpose model and shipping it internally at scale can have the same effect. Decide and document the role for every system in your inventory.

REALITY CHECK

"We only buy AI" is not an exit

Buying does not delegate Article 26. You still owe oversight, input-data discipline, log retention and worker information for the systems you deploy. And you can only fulfil those duties for systems you know exist, which makes discovery of unsanctioned AI use the practical first step of any programme.

04 · THE HIGH-RISK CORE, PART ONE

Articles 8–15, translated into work.

These provider requirements now apply from December 2027 for Annex III systems. They are also the clearest statement of what “well-governed AI” means in EU law, and buyers already use them as a due-diligence checklist. Article 8 is the umbrella: comply with all of the below, considering the system’s intended purpose.

ARTICLE 9

Risk management system

A continuous, iterative process across the whole lifecycle: identify risks, estimate them, mitigate, test, repeat. Not a one-off assessment at go-live.

EVIDENCE TO KEEP

A living risk file per system with versions, review dates, accepted residual risks and the sign-offs behind them.

ARTICLE 10

Data and data governance

Training, validation and test data must be relevant, sufficiently representative and managed for bias, with documented provenance and preparation choices.

EVIDENCE TO KEEP

Dataset documentation: sources, lineage, representativeness checks, bias examinations and the mitigations applied.

ARTICLE 11

Technical documentation

An Annex IV technical file, drawn up before market placement and kept current: architecture, capabilities, limitations, development choices and performance.

EVIDENCE TO KEEP

A maintained tech file with change history. If the system changed and the file did not, the file is wrong.

ARTICLE 12

Record-keeping

High-risk systems must automatically log events across their lifetime so operation is traceable: inputs, results, timing, the identities involved and the versions in play.

EVIDENCE TO KEEP

Tamper-evident runtime logs of every interaction, retained per policy. This is the article that makes an AI gateway a compliance asset rather than infrastructure.

The pattern to notice

Articles 9 through 12 are one loop: know what could go wrong, control the data going in, describe the system honestly, and record what it actually did. Every serious framework you already run, ISO 27001, SOC 2, DORA, has the same skeleton. The AI Act adds the model to the scope; it does not invent a new discipline.

04 · THE HIGH-RISK CORE, PART TWO

Oversight, robustness, and the duties that follow the sale.

ARTICLE 13

Transparency to deployers

Instructions for use that let deployers actually comply: intended purpose, accuracy levels, known limitations, oversight measures and expected lifetime.

EVIDENCE TO KEEP

Instructions of use per system version. Deployers: file these for every system you run; they anchor your Article 26 duties.

ARTICLE 14

Human oversight

Systems must be designed so natural persons can effectively oversee them: understand outputs, decide not to use them, intervene, or stop the system.

EVIDENCE TO KEEP

Named oversight owners, documented intervention and kill procedures, and records showing overrides actually happen.

ARTICLE 15

Accuracy, robustness, cybersecurity

Appropriate accuracy declared and tested, resilience to errors and misuse, and protection against AI-specific attacks: data poisoning, adversarial inputs and prompt injection for LLM-backed systems.

EVIDENCE TO KEEP

Evaluation results per release, red-team and injection test reports, and security incident records tied to the system.

ARTICLES 17, 43, 49, 72-73

The wrapper: QMS to incident reports

Providers also owe a quality management system, conformity assessment and CE marking, registration in the EU database, post-market monitoring, and serious-incident reports to authorities within 15 days, faster for the most severe cases.

EVIDENCE TO KEEP

QMS records, the declaration of conformity, registration entries and an incident log with report timestamps.

8-15

ARTICLES THAT DEFINE THE HIGH-RISK BASELINE, UNCHANGED IN SUBSTANCE BY THE OMNIBUS

15 days

STANDARD DEADLINE TO REPORT A SERIOUS INCIDENT UNDER ARTICLE 73

≥ 6 months

MINIMUM DEPLOYER LOG RETENTION UNDER ARTICLE 26, LONGER WHERE OTHER LAW APPLIES

05 · WHAT APPLIES TO ALMOST EVERYONE

Deployer duties, literacy, and the transparency deadline that did not move.

ARTICLE 26 · DEPLOYERS OF HIGH-RISK SYSTEMS, FROM DEC 2027

DUTY	WHAT IT MEANS IN PRACTICE
Use per instructions	Operate the system within the provider's stated intended purpose. Off-label use puts you in provider territory.
Assign human oversight	Named, trained, empowered people per system. An inbox is not oversight.
Control input data	Ensure inputs are relevant and representative for the purpose, which requires seeing what users actually send.
Monitor and escalate	Watch operation, suspend on risk signals, inform the provider and authorities where required.
Keep the logs	Retain the system's automatically generated logs at least six months, longer where sector law says so.
Inform workers	Tell affected workers and their representatives before deploying high-risk AI in the workplace.

APPLIES NOW

Article 4 · AI literacy

Staff who operate AI systems need a sufficient level of AI literacy, proportionate to their role and the systems in use. The Omnibus softened the framing, but a documented training programme with attendance records remains the credible answer.

2 AUG 2026

Article 50 · Transparency

Chatbots and AI interfaces must disclose they are AI. Synthetic audio, image and video must be machine-readably marked, and deepfakes labelled. If you ship AI-generated content to customers, this is your nearest hard deadline.

The shadow AI problem is a legal problem

Article 26 duties attach to the systems your organisation actually uses, not the ones IT approved. Staff pasting customer data into unsanctioned tools creates obligations, and GDPR exposure, that nobody is discharging. Discovery is not an optional maturity step; it is the precondition for every duty on this page.

06 · EVIDENCE CHECKLIST

Sixteen artefacts a regulator, auditor or enterprise buyer will ask for.

Print this page. Anything unticked is your programme backlog.

INVENTORY

- [ ] Register of every AI system and GPAI model in use, sanctioned or discovered · updated continuously, not annually
- [ ] Risk-tier classification per system · prohibited, high, limited, minimal, plus GPAI overlay
- [ ] Provider or deployer role decided and documented per system
- [ ] Named business owner and oversight owner per system

GOVERNANCE

- [ ] AI acceptable-use policy that is enforced technically, not only signed
- [ ] Provider instructions of use on file for every deployed system
- [ ] Data protection impact assessments where AI processes personal data
- [ ] Vendor due-diligence records covering Articles 8-15 posture

RUNTIME

- [ ] Every prompt, completion and tool call logged with user identity, model and version
- [ ] Logs tamper-evident, EU-resident, retained six months or longer
- [ ] DLP controls blocking or redacting personal data and secrets in prompts
- [ ] Incident process with an AI Act reporting path and tested timelines

PEOPLE AND TRANSPARENCY

- [ ] AI literacy training delivered, with attendance records
- [ ] Workers informed before workplace deployment of high-risk systems
- [ ] Chatbot and AI-interaction disclosure live in every user-facing surface
- [ ] Machine-readable marking of AI-generated media, ready before 2 December 2026 grace expiry

07 · WHERE RODER FITS

One control plane, most of the runtime evidence.

Roder sits inline between your organisation and every model, agent and MCP server, applying policy as traffic happens. That position is what makes it useful here: the Act's hardest evidence is runtime evidence, and runtime evidence has to be captured where the traffic flows.

CAPABILITY	WHAT IT DOES	SUPPORTS
AI Gateway	Every AI request through one governed, provider-agnostic endpoint with sub-second policy decisions.	Art. 12, 26 · monitoring, traceability
Policy Engine	Composable CEL rules that block, redact, modify or require, evaluated inline on prompts, completions and tool calls.	Art. 9, 14, 26 · enforced mitigations
Agent & MCP Security	Maps every agent and MCP server, scores identities, pins trust, and stops exfiltration paths in tool calls.	Art. 12, 15 · agent traceability, attack resilience
Shadow AI Control	Discovers unsanctioned AI via browser and gateway telemetry, then brings it under the same policy.	Art. 4, 26 · the inventory precondition
Observability	Immutable, EU-resident audit trail with full decision context, streamed to your SIEM, retained up to seven years.	Art. 12, 26, 73 · logs, retention, forensics
Cost Governance	Token-level attribution and ceilings. Not a compliance control, but it keeps the programme observable and funded.	Budget for all of the above

SCOPE, HONESTLY

What Roder does not do

Roder does not write your risk management system, perform conformity assessments, produce Annex IV documentation or replace counsel. It is the runtime control and evidence layer those processes depend on. Vendors claiming a platform alone makes you "AI Act compliant" are selling you an audit finding.

WHY EU RESIDENCY MATTERS

Sovereign by architecture

The evidence trail for EU regulation should not itself create a transfer problem. Roder is EU-hosted and EU-operated, with self-hosted and air-gapped options for defence and regulated sectors, so the audit trail stays inside the jurisdiction it serves.

08 · THE FIRST 90 DAYS

Ninety days to defensible. Seventeen months to comfortable.

You do not need a transformation programme to start; you need three thirty-day phases with visible output. This sequence front-loads the two things that cannot be rushed later: discovery and log history.

DAYS 1-30

See

- Discover AI in use across browser and network, including unsanctioned tools
- Build the system register; classify each entry by risk tier
- Decide provider or deployer role per system, in writing
- Name a business owner and an oversight owner per system

DAYS 31-60

Govern

- Route sanctioned AI through one gateway; retire or block the rest
- Turn the acceptable-use policy into enforced rules: DLP on personal data and secrets, allow and block lists
- Collect instructions of use from every provider; file the gaps as vendor actions
- Launch AI literacy training; record attendance

DAYS 61-90

Prove

- Logging on for all AI traffic: identity, model, version, decision
- Retention configured: six months minimum, seven years where sector rules demand
- Article 50 disclosures and content marking live before 2 August 2026
- Run one incident drill through the AI Act reporting path
- Generate the first evidence report; put it in front of your auditor

Then hold the cadence

Quarterly: refresh the inventory, review risk files, re-test oversight, and archive an evidence snapshot. Six quarters of that discipline is what December 2027 readiness looks like. The postponement is only an advantage for organisations that treat it as runway rather than relief.

RODER · THE CONTROL PLANE FOR ENTERPRISE AI

See every AI interaction. Govern it. Prove it.

Roder gives European enterprises one inline control plane for AI: gateway, policy, agent and MCP security, shadow AI discovery, observability and cost. EU-hosted, GDPR-native, with the audit trail this guide keeps asking you for.

Book a demo →

Apply for the 14-day trial →

RODER.AI/CONTACT

TALK TO US

sales@roder.ai

Book a scoping session

COMPLIANCE & SECURITY

compliance@roder.ai

security@roder.ai

RODER OY

Helsinki, Finland

EU-hosted · EU-operated