

THE

AI Governance Weekly

• 5-MIN READ • 7 DEVELOPMENTS

PUBLISHED BY COGNITA • AI GOVERNANCE THAT PROVES IT

• LEAD STORY

• EU

Brussels blinks: the EU agrees to delay the AI Act's high-risk deadlines.

In a provisional Digital Omnibus agreement reached 7 May 2026, the European Parliament and the Council moved Annex III high-risk obligations from 2 August 2026 to 2 December 2027, with product-embedded (Annex I) systems shifting to 2 August 2028. The deal is provisional and awaits formal adoption.

EDITOR'S READ ANALYSIS, NOT REPORTING

The clock moved; the expectations didn't. The deferral buys roughly 16 months – but the documentation, testing, and impact-assessment evidence the Act demands still takes that long to build. Treat the new date as runway, not relief.

[S-01] consilium.europa.eu

— THE WEEK IN ONE MINUTE

If you read nothing else.

Seven developments worth your attention — each one a named actor, a dated action, and a source you can check. Standing facts are marked ○ CONTEXT, not news. The full briefing follows.

■ EU

EU agrees to delay high-risk AI Act deadlines. On 7 May 2026 Parliament and the Council reached a provisional *Digital Omnibus* deal moving Annex III obligations to 2 Dec 2027 (Annex I to 2 Aug 2028). PROVISIONAL

[S-01] consilium.europa.eu ↗

■ EU

Same deal tightens two more rules. Watermarking (Art. 50(2)) is now due 2 Dec 2026, and a new Article 5 ban on AI “nudifiers” and CSAM is added.

[S-02] europarl.europa.eu ↗

■ US · FED

Federal agency AI inventories run under the 2025 OMB memos. M-25-21 and M-25-22 require use-case inventories and chief AI officer governance. ○ CONTEXT

[S-03] whitehouse.gov ↗

■ US · STATE

Colorado's SB 24-205 sets the US state benchmark. The first comprehensive US state AI law imposes reasonable-care duties on developers and deployers of high-risk systems. ○ CONTEXT

[S-04] leg.colorado.gov ↗

■ ISO · STD

ISO/IEC 42005 pairs with the 42001 management system. The AI system impact-assessment standard, published 2025, is run alongside ISO/IEC 42001 for repeatable evidence. ○ CONTEXT

[S-05] iso.org ↗

■ UK · GLOBAL

UK AI Security Institute continues model evaluations. The body renamed from the AI Safety Institute in 2025 keeps publishing frontier-model evaluation findings. ○ CONTEXT

[S-06] gov.uk ↗

How to use this: each item names who acted, what they did, and when — verifiable against the canonical source on page 7. Items marked **Context** are standing facts, not this-week events. **Provisional** flags a status that is not yet law. Framing is advisory: **readiness, not certification.**

What moved this week.

■ EU

Provisional agreement · 7 May 2026

EU agrees to delay the AI Act's high-risk deadlines under the Digital Omnibus.

On 7 May 2026, the European Parliament and the Council reached a provisional political agreement on the *Digital Omnibus on AI* (first tabled 19 November 2025). **Stand-alone Annex III high-risk obligations move from 2 August 2026 to 2 December 2027; product-embedded Annex I systems move from 2 August 2027 to 2 August 2028.** It takes legal effect only on formal adoption.

WHAT IT MEANS FOR YOUR AI PROGRAM

Plan against 2 Dec 2027 for Annex III systems — but confirm your role and Annex with counsel, since the deal is not yet law.

EDITOR'S READ ANALYSIS, NOT REPORTING

The obvious read is relief; what it misses is that the deferral resets the date, not the workload — 16 months is about how long the Act's evidence takes to build, and the co-legislators have signalled they will not move again.

[S-01] consilium.europa.eu

■ EU

7 May 2026

Watermarking deadline tightened; a new Article 5 prohibition added.

The same agreement compresses the Article 50(2) synthetic-content marking grace period from six months to three — now due 2 December 2026 — and inserts a prohibition on AI that generates non-consensual intimate imagery and CSAM. Other Article 50 transparency duties still apply from 2 August 2026.

WHAT IT MEANS FOR YOUR AI PROGRAM

If you generate synthetic media, the marking obligation arrives sooner than the headline delay suggests — scope watermarking against 2 Dec 2026, not 2027.

[S-02] europarl.europa.eu

■ US · FED

○ CONTEXT

Federal agency AI inventories and CAIO governance under the OMB memos.

Under OMB memoranda M-25-21 and M-25-22 (2025), federal agencies maintain AI use-case inventories and chief AI officer governance. This is a standing requirement, not a this-week event.

WHAT IT MEANS FOR YOUR AI PROGRAM

If you sell to government, expect inventory and governance questions in procurement. A maintained use-case register maps directly to these expectations.

[S-03] whitehouse.gov

— TOP DEVELOPMENTS · 2 OF 2

And the rest that matter.

■ US · STATE

○ CONTEXT

Colorado's SB 24-205 sets the US state high-risk benchmark.

Colorado enacted SB 24-205, the first comprehensive US state AI law — reasonable-care duties on developers and deployers of high-risk systems against algorithmic discrimination.

WHAT IT MEANS FOR YOUR AI PROGRAM

If you operate in Colorado, inventory consequential-decision systems and assemble impact-assessment evidence — the duties center on documentation and reasonable care.

[S-04] leg.colorado.gov ↗

■ ISO · STD

○ CONTEXT

ISO/IEC 42005 pairs impact assessment with the 42001 system.

ISO/IEC 42005, published 2025, gives guidance on AI system impact assessments. Organizations run it alongside an ISO/IEC 42001 AI management system for a defensible, repeatable process.

WHAT IT MEANS FOR YOUR AI PROGRAM

Mapping impact assessments to a recognized standard makes them reusable evidence across the AI Act, NIST AI RMF, and customer due diligence.

[S-05] iso.org ↗

■ UK · GLOBAL

○ CONTEXT

UK AI Security Institute continues its model-evaluation work.

The UK's AI Safety Institute was renamed the AI Security Institute in 2025. Its evaluation programme publishes findings on frontier-model capabilities and risks.

WHAT IT MEANS FOR YOUR AI PROGRAM

Public evaluation methods are a useful template for your own pre-deployment testing — adapt the categories to your use cases.

[S-06] gov.uk ↗

■ US · FED

○ CONTEXT

NIST AI RMF Generative AI Profile remains a common baseline.

NIST's Generative AI Profile (AI 600-1) companions the AI Risk Management Framework with GenAI-specific actions. It is cited in programs and contracts as a shared reference.

WHAT IT MEANS FOR YOUR AI PROGRAM

Mapping controls to the RMF functions (Govern, Map, Measure, Manage) gives you a vocabulary that travels across frameworks and buyers.

[S-07] nist.gov ↗

— JURISDICTION ROUNDUP

Where the rules are moving.

Five lanes, colour-coded to match the live wall. A short read on each region this week.

● EU

AI Act high-risk deadlines deferred — Annex III to 2 Dec 2027 — under the provisional Digital Omnibus deal.

consilium.europa.eu ↗

Watermarking due 2 Dec 2026; a new **Article 5** prohibition on “nudifiers” / CSAM added.

europarl.europa.eu ↗

GPAI rules (Art. 51–55) in force since Aug 2025 — untouched by the Omnibus.

europa.eu ↗

● US · FED

OMB memos (M-25-21/22) drive agency AI inventories + CAIO governance.

whitehouse.gov ↗

NIST AI RMF + GenAI Profile cited as a common baseline.

nist.gov ↗

Procurement increasingly asks for governance evidence.

whitehouse.gov ↗

● US · STATES

Colorado SB 24-205 — reasonable-care duties for high-risk AI.

leg.colorado.gov ↗

First-in-nation law anchors **developer + deployer** obligations.

leg.colorado.gov ↗

State bills converge on inventory + disclosure themes.

leg.colorado.gov ↗

● UK · GLOBAL

AI Security Institute publishes model evaluations.

gov.uk ↗

Sector-regulator approach continues over a single AI act.

gov.uk ↗

International testing cooperation deepens across institutes.

gov.uk ↗

● ISO · STD

ISO/IEC 42001 AI management system adoption widening.

iso.org ↗

ISO/IEC 42005 impact-assessment guidance in uptake.

iso.org ↗

Mappings to NIST + AI Act reduce duplicate effort.

nist.gov ↗

— ON THE HORIZON

Dates worth a calendar hold.

Real, sourced dates only — each marked **PROVISIONAL** (pending the Omnibus adoption) or in-force. Confirm exact triggers with counsel; phasing varies by role and Annex.

Aug 2, 2026

EU AI Act — Article 50 transparency duties apply **IN FORCE**

Most Article 50 transparency obligations (including deployer-facing duties) are **not** delayed by the Omnibus and apply from this date. Formal adoption of the Omnibus is expected before it.

europarl.europa.eu ↗

Dec 2, 2026

EU AI Act — synthetic-content watermarking due **PROVISIONAL**

The Article 50(2) marking grace period was compressed from six months to three under the provisional deal. **Scope watermarking against this date**, not the high-risk timeline.

europarl.europa.eu ↗

Dec 2, 2027

EU AI Act — Annex III high-risk obligations **PROVISIONAL**

Deferred from 2 August 2026 under the Digital Omnibus. **Confirm your role** (provider vs. deployer) — this date applies to stand-alone high-risk systems listed in Annex III.

consilium.europa.eu ↗

Aug 2, 2028

EU AI Act — Annex I product-embedded obligations **PROVISIONAL**

Deferred from 2 August 2027. Covers high-risk AI embedded in regulated products (medical devices, machinery, and the like) under EU product-safety law.

consilium.europa.eu ↗

— SOURCES & CITATIONS

Every item, independently verifiable.

Canonical URLs for everything cited this week, numbered and grouped by section. This page is the point: we cite everything.

— LEAD & TOP DEVELOPMENTS

- S-01 **Council of the EU** — Press release: provisional agreement on the Digital Omnibus on AI, 7 May 2026.
consilium.europa.eu/en/press/press-releases/2026/05/07
- S-02 **European Parliament** — Press release: AI Act amendments (watermarking grace period; Article 5 prohibition), May 2026.
europarl.europa.eu/news/en/press-room
- S-03 **White House / OMB** — Federal AI use guidance (M-25-21, M-25-22).
whitehouse.gov/omb
- S-04 **Colorado General Assembly** — SB 24-205, Consumer Protections for AI.
leg.colorado.gov/bills/sb24-205
- S-05 **ISO** — ISO/IEC 42001 (AI management system) & ISO/IEC 42005 (impact assessment).
iso.org/standard/81230.html
- S-08 **European Commission** — Digital Omnibus on AI proposal, tabled 19 Nov 2025 (backs the lead).
digital-strategy.ec.europa.eu/en/library/digital-omnibus

— ROUNDUP & HORIZON

- S-06 **UK Government** — AI Security Institute (renamed from AI Safety Institute, 2025).
gov.uk/government/organisations/ai-security-institute
- S-07 **NIST** — AI Risk Management Framework & Generative AI Profile (AI 600-1).
nist.gov/itl/ai-risk-management-framework
- S-09 **European Commission** — GPAI Code of Practice / AI Office (Articles 51–55).
digital-strategy.ec.europa.eu/en/policies/ai-office

— METHOD & HONESTY

- M-1 Each item names a **dated action, an actor, and an instrument**. We paraphrase for flow and quote verbatim for load-bearing claims; read the canonical source for exact wording.
- M-2 Items marked **Provisional** are not yet law; those marked **Context** are standing facts, not this-week events. The most recent dated development wins.
- M-3 Every cited claim is re-checked against its source in a separate **verification pass** before publication; any claim the source contradicts or outdates is cut.
- M-4 “What it means” notes and “Editor’s read” are **advisory analysis, not legal advice** — readiness, mapped to the NIST AI RMF, never certification.

FROM THE PUBLISHER

Run a continuous AI risk program, not a one-time checklist.

The deadlines on page 6 won't wait for an annual review. Cognita Compass maps your systems to ISO 42001, the NIST AI RMF, and the EU AI Act — and keeps the evidence current as the rules move.

START THE 100-DAY PATH

cognitagrc.io/compass

A continuous AI risk program in 100 days — readiness, not certification. The first phase is free to start.



◆ Forward this to a colleague who owns AI risk · **Subscribe free at cognitagrc.io/weekly**

Honest framing: The AI Governance Weekly summarizes public developments and links every source. It is not legal advice. Cognita helps you build readiness mapped to the NIST AI RMF and prepare for ISO 42001 — it does not issue certifications or guarantee regulatory outcomes.