CSAI Foundation | Cloud Security Alliance

EU AI Act GPAI Enforcement Goes Live

The Grace Period Ends and the AI Office Gains Teeth

2026-08-03

AI-assisted Rapid Research

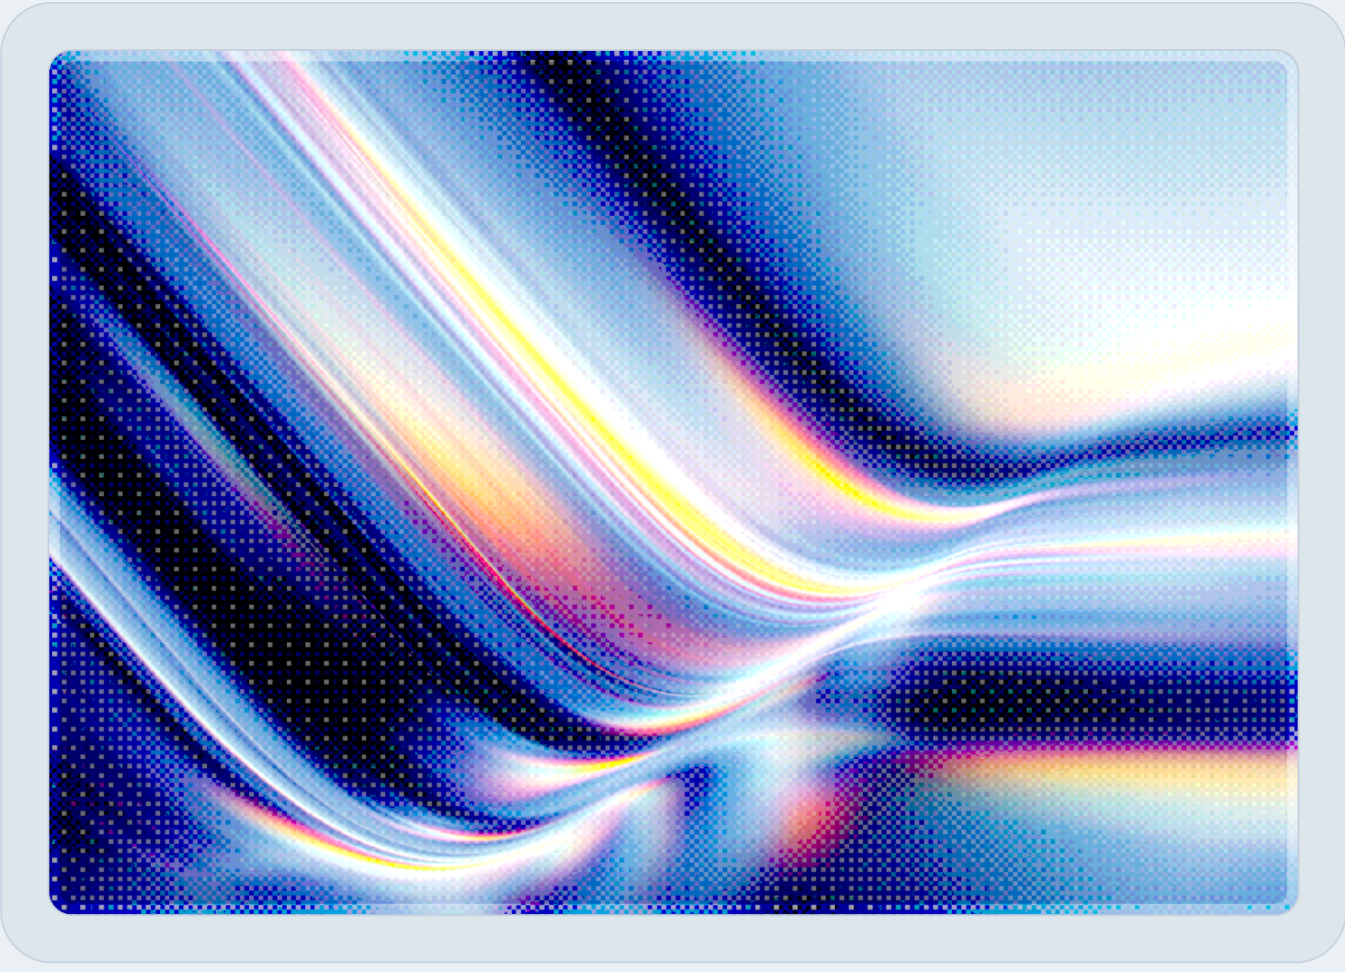

CSAI

CSA cloud security alliance®

Key Takeaways

August 2, 2026 closes the one-year window during which providers of general-purpose AI (GPAI) models were legally bound by the EU AI Act's Articles 51 through 56 but effectively unpoliced. As of yesterday, the European Commission's AI Office holds the authority to demand documentation, commission independent model evaluations, order corrective measures up to market withdrawal, and impose fines of up to 3 percent of worldwide annual turnover or €15 million, whichever is higher, for GPAI non-compliance [1][2].

This enforcement regime is structurally different from the rest of the AI Act. Where national market surveillance authorities police most obligations, Chapter V enforcement against GPAI providers sits exclusively with the Commission under Article 88, giving the bloc a single, centralized enforcement point [2]. The Digital Omnibus package that deferred Annex III high-risk system deadlines to December 2, 2027 left this GPAI enforcement timeline untouched, just as it left Article 50's transparency duties untouched — a pattern CSA's prior research has already flagged as a recurring source of enterprise confusion about what, exactly, the Omnibus deferred [3][4].

The more consequential story on the day enforcement formally begins is not the legal deadline itself but whether the AI Office can act on it. A recent analysis of the Office's staffing and budget, drawing on a Pour Demain report, describes a unit that remains small relative to its mandate, with an even thinner pool of qualified technical evaluators, and recommends the Office scale toward roughly 160 staff by 2030 to match its enforcement responsibilities [5]. The coming months function as a credibility window: early choices about whether the Commission proactively requests documentation, and whether it treats Code of Practice signatories differently from holdouts like Meta, will signal whether GPAI enforcement becomes a genuine compliance driver or remains largely symbolic [5][6].

CSA's own research anticipated this milestone in a pre-deadline readiness note published in May 2026 [7]. This note picks up where that one left off: enforcement authority is no longer prospective, and organizations that treated the twelve-month gap between binding obligation and enforceable obligation as slack now need to treat GPAI compliance as an active regulatory exposure rather than a future planning item.

Background

The EU AI Act's general-purpose AI provisions became legally binding on August 2, 2025. From that date, providers of GPAI models have been required to maintain technical documentation reflecting the model as actually deployed, publish a summary of training content using the AI Office's template, adopt and publish a policy for complying with EU copyright law, and, for providers established outside the EU, designate an authorized representative within the bloc [1][6]. A smaller set of providers — those whose models are trained using more than 10^25 floating-point operations, a threshold intended to capture only the most capable systems in production — carry additional systemic-risk obligations: model evaluation and adversarial testing, documented risk assessment and mitigation measures, cybersecurity protections for the model and its infrastructure, and serious-incident reporting to the AI Office [6][8]. For a full year, these obligations existed on paper without a mechanism to compel compliance: providers were bound to the letter of the law, but no EU authority had the power to enforce it [1].

Alongside the binding statutory obligations, the European Commission published the final General-Purpose AI Code of Practice on July 10, 2025, a voluntary instrument built around three chapters covering transparency, copyright, and safety and security for systemic-risk models [8][9]. Signing the Code carries a practical incentive: regulators treat adherence as a presumption of conformity with the underlying statutory obligations, reducing a signatory's administrative burden and legal uncertainty even though the Code itself creates no new legal requirements [8]. OpenAI, Anthropic, Microsoft, Google, and Mistral AI indicated they would sign [8], while Meta publicly declined on July 18, 2025, arguing the Code introduced obligations that exceeded the AI Act's own text and created legal uncertainty for the company [10][11]. xAI took a middle path, signing only the safety-and-security chapter and opting to demonstrate transparency and copyright compliance through other means [8]. Critically, declining to sign does not exempt a provider from enforcement or fines; it simply removes the mitigating weight the AI Office may give to a documented Code of Practice commitment when it calculates a penalty [6].

The Digital Omnibus negotiations that concluded with Council approval on June 29, 2026 reshaped part of the AI Act's timeline, pushing the compliance deadline for standalone high-risk systems under Annex III out to December 2, 2027 in response to concerns that harmonized technical standards were not ready [3][4]. That deferral, however, never touched Chapter V. GPAI obligations and enforcement, like the Article 50 transparency duties CSA analyzed in its prior note, sit outside the risk-tiered Annex I/Annex III structure the Omnibus renegotiated, and both continued on their original schedule throughout the Omnibus debate [3][4]. The practical effect is that August 2, 2026 arrived as scheduled for GPAI enforcement even as much of the public conversation about the EU AI Act in mid-2026 focused on what had been delayed.

Security Analysis

The technical substance of what the AI Office can now demand is more consequential than the fines headline suggests. Article 91 gives the Commission the power to request documentation and information directly from GPAI providers, with penalties attached to responses that are incorrect, incomplete, or misleading, which functionally requires providers to maintain technical documentation as a continuously accurate artifact rather than a point-in-time compliance exercise completed once in August 2025 [2][6]. For systemic-risk providers, the Commission can also commission independent evaluations — including, per some legal analyses, arrangements for evaluator access to model internals — a capability that goes beyond a documentation review and starts to resemble the kind of technical audit security teams undergo from an external assessor rather than a typical regulatory paperwork check [2][5].

The gap security teams should track most closely, though, is the one between the AI Office's legal authority and its operational capacity to use it. Analysis of the Office's staffing situation describes the team responsible for systemic-risk oversight as small, its pool of qualified technical evaluators as thinner still, and its current trajectory as inadequate relative to the enforcement demand it will face without substantial growth toward the roughly 160-person staffing level recommended for 2030 [5]. This is not simply a bureaucratic curiosity; it shapes the near-term threat model for compliance programs. A resource-constrained enforcer is more likely to act on visible, low-effort signals — public complaints, high-profile incidents, downstream provider escalations under Article 88(2), or qualified alerts from the Act's scientific panel — than to conduct proactive, comprehensive sweeps of every GPAI provider's documentation [2][5]. Organizations should not read early enforcement quiet as evidence that the regime lacks teeth; Lawfare's analysis of the Office frames the coming months explicitly as a "credibility window" in which the Office's first visible actions, not its statutory text, will establish the actual level of scrutiny providers should expect [5].

A second structural point worth internalizing is that the systemic-risk threshold itself — 10^25 FLOPs of training compute — is a static number applied to what many industry observers describe as a fast-moving target for training compute. A model that falls just under the threshold today, and therefore avoids the safety-framework, red-teaming, and incident-reporting obligations that come with systemic-risk classification, should not be assumed to remain a marginal case indefinitely, given the pace at which frontier training runs have scaled in recent years, and providers operating near that boundary should treat their classification as something to reassess on a recurring basis rather than a one-time determination made at initial deployment [6][8]. Downstream deployers face a parallel, less technical but equally consequential risk: fine-tuning a base model substantially enough to constitute a significant modification under the Commission's guidelines can convert a downstream deployer into a provider in the AI Office's eyes — a status shift that is easy to overlook when an organization assumes GPAI compliance is solely the foundation model vendor's responsibility [12].

Recommendations

Immediate Actions

Organizations should inventory every GPAI model in use across the enterprise, recording the provider, deployment date, and whether any internal team has fine-tuned or substantially modified the model in a way that could shift the organization's own status from deployer to provider under the Act's definitions [6][12]. For each model, confirm that the provider's technical documentation and training-content summary reflect the currently deployed version rather than an earlier snapshot completed to meet the August 2025 baseline, and verify that non-EU providers in the supply chain have designated an EU authorized representative as required [1][6].

Short-Term Mitigations

Enterprises that rely on third-party GPAI models as downstream integrators should formally request the provider's current compliance documentation package, including copyright policy and, where applicable, systemic-risk safety framework summaries, and should escalate through Article 88(2) channels — which allow downstream parties to prompt Commission review — if a provider is unresponsive [2][6]. Compliance and legal teams evaluating vendor relationships should factor Code of Practice adherence into vendor risk assessments, not because non-signatories are automatically non-compliant, but because signatory status affects the mitigating weight regulators apply and offers a useful, publicly verifiable signal of a provider's own EU AI Act posture [8][10].

Strategic Considerations

Security and governance functions should treat GPAI compliance as a recurring assessment rather than a completed project, revisiting systemic-risk classification, documentation currency, and fine-tuning-driven provider status at a defined cadence rather than only at initial deployment [6][8]. Programs should track the AI Office's early enforcement actions over the coming months as the clearest available signal of practical enforcement intensity, given the gap between the Office's statutory powers and its current staffing, and should align GPAI compliance evidence with CSA's AI Controls Matrix (AICM) v1.1 so that documentation built for EU AI Act purposes strengthens a broader, auditable AI governance program rather than existing as a standalone regulatory checklist [13].

CSA Resource Alignment

CSA's May 2026 research note, "EU AI Act GPAI: Security Compliance Before August 2026," is the direct predecessor to this analysis, having laid out the same enforcement milestone, the systemic-risk threshold, and a phased pre-deadline compliance plan while the Commission's powers were still prospective [7]. This note is the intended follow-through: the compliance window that note described as three phases of preparation has closed, and the recommendations here shift accordingly from readiness planning to ongoing operational monitoring of a now-active enforcement regime.

CSA's analysis of the EU AI Act Digital Omnibus recalibration is the most directly relevant prior publication for understanding why GPAI enforcement proceeded on schedule despite widespread reporting about EU AI Act delays, since it specifically documented that the Omnibus deferral applied to Annex III high-risk systems and left GPAI obligations, along with Article 5 prohibited practices and Article 50 transparency duties, on their original timelines [3]. CSA's companion note on Article 50 transparency obligations illustrates the identical structural pattern now playing out for GPAI: a legal obligation that survived the Omnibus negotiations untouched, arriving on schedule while public attention remained fixed on the deadlines that did move [4]. Read together, these three notes describe a consistent theme in the EU AI Act's 2026 rollout — the obligations attached to foundation models and their outputs have so far remained on their original timelines, while the risk-tiered obligations attached to downstream, sector-specific high-risk systems have not.

Finally, organizations building or updating GPAI compliance evidence should map that work directly to CSA's AI Controls Matrix (AICM) v1.1, whose domains addressing AI governance, model documentation, and third-party/supply-chain risk provide a control-level structure that turns EU AI Act compliance artifacts into reusable governance evidence rather than a jurisdiction-specific side project [13].

References

[1] ComplianceHub.Wiki. "EU AI Act GPAI Enforcement Goes Live August 2, 2026: A Readiness Guide for AI Governance Teams." ComplianceHub.Wiki, 2026.

[2] artificialintelligenceact.eu. "Enforcement of Chapter V Under the EU AI Act." artificialintelligenceact.eu, 2026.

[3] Cloud Security Alliance. "EU AI Act Digital Omnibus: Enterprise Risk Recalibration." CSA AI Safety Initiative, June 2026.

[4] Cloud Security Alliance. "EU AI Act Article 50: Transparency Obligations Take Effect." CSA AI Safety Initiative, July 2026.

[5] Massaro, Christoph. "How Much Power Does the EU AI Office Actually Have?" Lawfare, 2026.

[6] Latham & Watkins. "EU AI Act: GPAI Model Obligations in Force and Final GPAI Code of Practice in Place." Latham & Watkins, 2025.

[7] Cloud Security Alliance. "EU AI Act GPAI: Security Compliance Before August 2026." CSA AI Safety Initiative, May 9, 2026.

[8] European Commission. "The General-Purpose AI Code of Practice." Shaping Europe's Digital Future, 2025.

[9] artificialintelligenceact.eu. "An Introduction to the Code of Practice for General-Purpose AI." artificialintelligenceact.eu, 2025.

[10] CNBC. "Meta Says It Won't Sign Europe AI Agreement, Calling It an Overreach That Will Stunt Growth." CNBC, July 18, 2025.

[11] Euronews. "Meta Rebuffs EU's AI Code of Practice." Euronews, July 18, 2025.

[12] European Commission. "Guidelines for Providers of General-Purpose AI Models." Shaping Europe's Digital Future, 2025.

[13] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.