CSAI Foundation | Cloud Security Alliance

GPAI Enforcement Is Live: What Security Teams Must Do Now

2026-08-29

AI-assisted Rapid Research

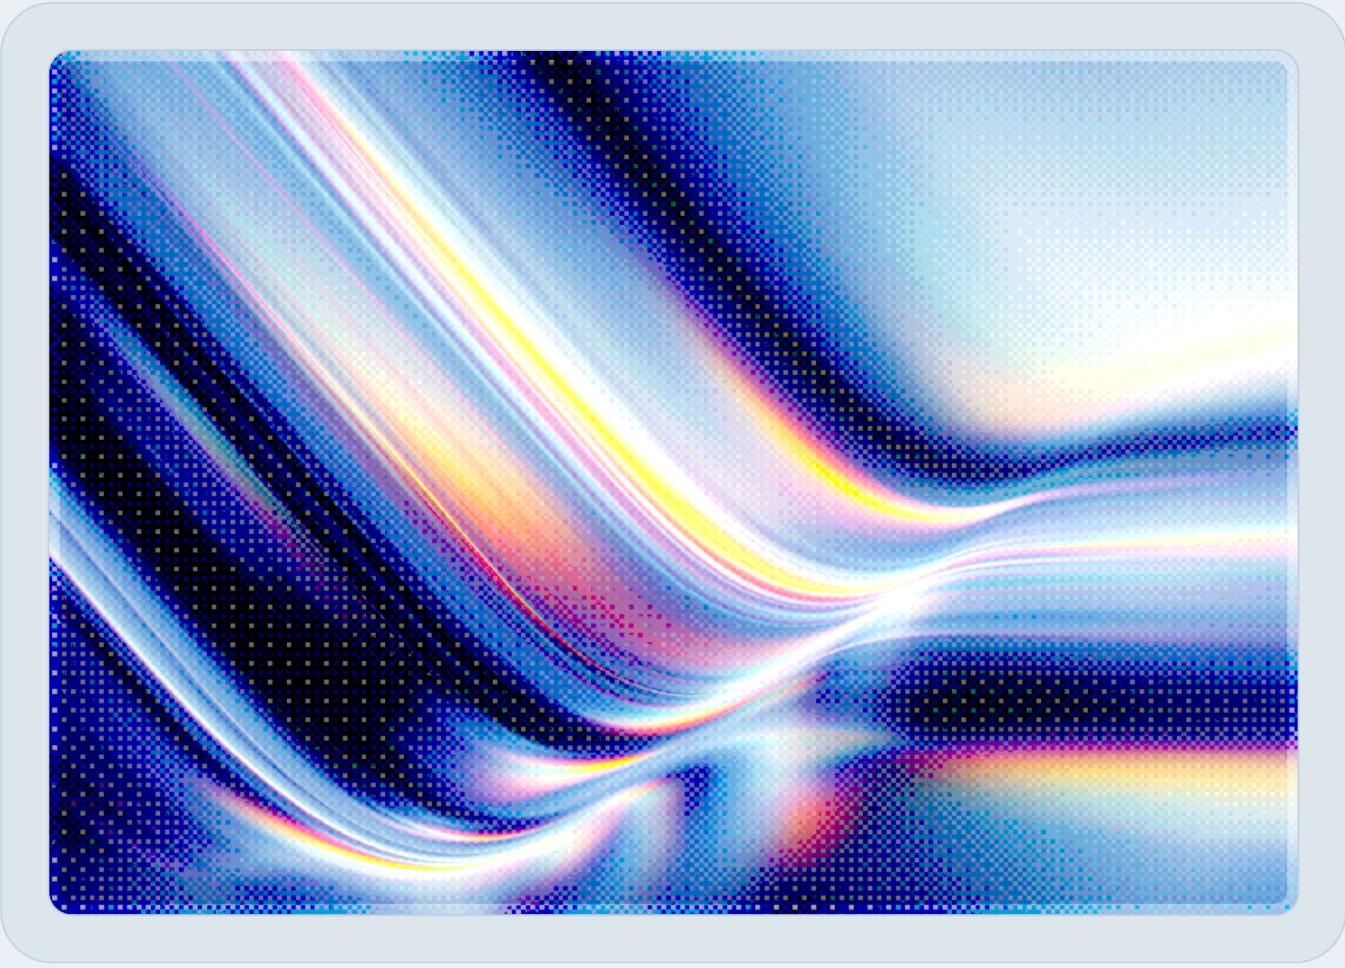

CSAI

CSA cloud security alliance®

Key Takeaways

Since August 2, 2026, the European Commission's AI Office has held full enforcement authority over general-purpose AI (GPAI) model providers, ending a one-year period in which GPAI obligations were legally binding but functionally unenforceable [1][2]. The Office can now compel documentation, run its own model evaluations, order mitigation measures, restrict or withdraw models from the EU market, and impose fines of up to €15 million or 3 percent of a provider's global annual turnover, whichever is higher [3][4]. Because the underlying obligations under Articles 53 through 55 of the Regulation took effect a full year earlier, on August 2, 2025, the Commission's opening enforcement actions can reach back over conduct that occurred before enforcement power existed — the compliance record built during the "grace" year now matters retroactively [1][5]. For security teams, the practical exposure is less about direct regulatory liability, since most enterprises are GPAI deployers rather than providers, and more about vendor concentration risk: a restriction or withdrawal order against a systemic-risk model an organization depends on operationally would function as an abrupt, regulator-driven service disruption with no commercial notice period. This note updates CSA's earlier pre-enforcement guidance with what the mechanism actually does now that it is live, and lays out what security and vendor-risk functions should be doing in the first weeks of the enforcement era.

Background

The EU AI Act's treatment of general-purpose AI models has always run on two separate clocks. The substantive obligations in Articles 53 through 55 — transparency documentation, copyright policy statements, training-data summaries, and, for models presumed to carry systemic risk, adversarial testing and cybersecurity requirements — became legally binding on August 2, 2025 [6][7]. But the Commission could not fine anyone for missing them until a second date, August 2, 2026, when its dedicated GPAI enforcement mechanism under Chapter V, Articles 88 through 94, entered into application [3][8]. That one-year gap functioned as an adjustment period, giving the market time to operationalize compliance through soft-law instruments, chiefly the Code of Practice for General-Purpose AI, which the AI Office finalized in July 2025 as a voluntary route to demonstrating compliance [9][10][12]. Most frontier developers, including the providers behind GPT, Claude, and Gemini model

families, signed the Code; Meta and the major Chinese labs did not [10]. Signing the Code creates a rebuttable presumption of conformity with the underlying legal obligations, which suggests the AI Office is likely to apply greater scrutiny to non-signatories once enforcement begins.

August 2, 2026 also happens to be the date several other EU AI Act obligations came due, which compounds the compliance load for any organization operating in the region. Article 50 transparency duties for chatbots, synthetic media, and deepfake-adjacent systems became enforceable the same day, independent of whether the underlying system qualifies as high-risk [11]. The separate track for prohibited AI practices under Article 5 also gained teeth on this date, carrying steeper penalties of up to €35 million or 7 percent of global turnover [1][2]. The May 2026 Digital Omnibus agreement pushed back the compliance deadline for standalone high-risk systems under Annex III by sixteen months, to December 2027, but it left the GPAI enforcement date and the Article 50 transparency date untouched, which means the regulatory calendar organizations built around this summer is now largely accurate for the GPAI track specifically even as the high-risk track shifted underneath it [18]. Security and governance teams tracking EU AI Act readiness through a single master timeline should treat GPAI enforcement as the one date in that calendar that did not move.

The threshold that determines which GPAI models fall into the higher-scrutiny "systemic risk" tier is defined by training compute: any model trained using 10^25 floating-point operations or more is presumed to have high-impact capabilities, and its provider must notify the Commission within two weeks of reasonably foreseeing that the model will cross that threshold [10]. Systemic-risk providers carry additional obligations under Article 55, including state-of-the-art model evaluation, adversarial testing, reporting of serious incidents to the AI Office, and cybersecurity protections covering the model and its weights across its lifecycle [7]. The Commission has already operationalized part of that regime by publishing a standardized reporting template that systemic-risk providers use to document serious incidents [13]. This is the population of models — the ones large enough to matter systemically — where enforcement actions, once they begin, are most likely to land first.

Security Analysis

What actually changed on August 2, 2026 is the AI Office's posture, not the underlying legal text. Before that date, the Office could observe, request voluntary cooperation, and build a case file, but it lacked the formal power to compel anything or issue a penalty. After it, the relationship between the Office and any GPAI provider becomes one the Office itself has described in terms of moving "from persuasion to compulsion" [1]. Article 91 gives the Commission the power to request the full technical documentation a provider was already required to produce under Articles 53 and 55, plus any additional information it deems necessary to assess compliance, and providers who respond with incomplete, incorrect, or

misleading information face fines on that basis alone, independent of whatever the underlying evaluation finds [14][15]. Article 89 authorizes ongoing monitoring, and Article 90 creates a formal channel for the independent scientific panel supporting the AI Office to flag emerging systemic-risk concerns for investigation, which means the Office does not have to wait for a public incident to open a file [19][20]. Where an evaluation raises serious and substantiated concerns, the Office can require risk-mitigation measures and, in the most severe cases, order that a model be restricted, withdrawn, or recalled from the EU market [1][2]. Article 101 caps the resulting fines at the higher of €15 million or 3 percent of worldwide annual turnover for GPAI infringements [4].

In CSA's assessment, the retroactivity question is the one most likely to be underestimated by organizations that treated the 2025-2026 window as a compliance grace period rather than a compliance record-building period. Because the Article 53-55 obligations were already legally binding from August 2025, the Commission is not limited to policing conduct that occurs after August 2, 2026; it can examine whether a provider's documentation, risk assessments, and incident-reporting practices met the legal bar throughout the preceding year, and act on what it finds [1][5]. For a GPAI provider, this means the technical documentation package assembled in mid-2025 is now the evidentiary record the Office will test first. For an enterprise deploying GPAI models rather than providing them, the practical takeaway is different but no less material: it clarifies why a provider's public claims of "compliance since 2025" said nothing about enforcement exposure, and it argues for treating any provider's Code of Practice signatory status, or lack of it, as a genuine risk signal rather than a marketing detail.

That signal matters most because of where the actual exposure sits for most CSA member organizations. Most CSA member organizations are deployers, not providers, in the Regulation's sense — building products and internal tooling on top of models from a small number of frontier labs. The AI Office's enforcement toolkit is aimed at providers, but its consequences flow downstream. An organization's compliance posture, and increasingly its operational continuity, is now partly a function of its vendors' compliance posture. If the AI Office restricts or withdraws a systemic-risk model from the EU market — the most severe tool in its kit, and one the Office's toolkit is designed to reach for where mitigation measures fail — every deployer with production workloads built on that model faces an abrupt, regulator-driven service disruption with none of the notice period a commercial deprecation would normally carry. This converts EU AI Act enforcement from a legal-and-compliance topic into a vendor-concentration and business-continuity topic that belongs on the same risk register as a cloud provider outage or a critical open-source dependency going unmaintained. Security teams that have not yet mapped which production systems depend on which specific GPAI model, and whether that model's provider is a Code of Practice signatory with a clean documentation record, do not currently have visibility into this exposure.

A second, quieter risk sits in the accountability chain for non-EU providers. Article 54 requires any GPAI provider established outside the Union to appoint an EU-based authorized representative before placing a model on the market; that representative is responsible for verifying the provider's technical documentation, retaining it for ten years, and cooperating with the AI Office on the provider's behalf [16] [17]. Fines for Article 54 non-compliance are not limited to the provider's substantive obligations — failures in the representative relationship itself are independently sanctionable [17]. For security and procurement teams conducting vendor due diligence on a non-EU AI provider, confirming that a valid, functioning authorized representative exists is now a concrete, checkable control, and its absence is a documented compliance gap rather than a theoretical one.

Finally, the AI Office has signaled that its default opening move with providers and large deployers will be what it calls "technical compliance dialogues" — structured conversations intended to assess status and resolve ambiguity before any formal proceeding begins [1]. This is a meaningful operational detail: it means the first contact from the Office, when it comes, is more likely to look like a detailed information request than a fine, and organizations that have never rehearsed a response to a regulator's documentation request risk mishandling that first interaction badly enough to escalate it.

Recommendations

Immediate Actions (0-30 days)

Security and AI governance teams should build or refresh an inventory of every GPAI model in production and pre-production use, tagged by provider, deployment context, and — critically — whether the model meets or is likely to meet the 10^25 FLOP systemic-risk presumption. For each provider on that list, confirm Code of Practice signatory status and request current Article 53 technical documentation and cybersecurity attestations directly; a provider that resists a documentation request from its own customer is unlikely to have handled the AI Office's equivalent request any better. For any non-EU provider in the inventory, verify that a functioning Article 54 authorized representative is in place, since this is now a discrete, auditable compliance gap rather than a formality.

Short-Term Mitigations (30-90 days)

Organizations with material dependence on a single GPAI provider for critical workflows should build a documented contingency plan for a scenario in which that provider's model is restricted or withdrawn by the AI Office, including an identified fallback model or provider and an estimate of the switching cost and downtime involved. Internal incident detection and escalation workflows should be reviewed against the

reporting expectations that apply to systemic-risk models, so that an AI-related anomaly is routed to the team responsible for external reporting without delay rather than sitting in a general IT ticket queue. Where the organization is itself building or fine-tuning models that could approach the systemic-risk compute threshold, legal and security teams should establish the two-week notification workflow to the Commission now, before it becomes urgent.

Strategic Considerations

Enforcement precedent will accumulate through the AI Office's early technical compliance dialogues and any resulting public actions, and organizations should treat the next several months as the period in which the practical enforcement bar gets set. GPAI vendor governance is better framed as a continuing supply-chain risk discipline, integrated into the same vendor-risk-management program that already covers cloud and software dependencies, than as a one-time compliance sprint tied to a single deadline. Because the GPAI, Article 50 transparency, and Article 5 prohibited-practices enforcement tracks all activated on the same date but carry different penalty structures and different triggers, legal, privacy, and security functions should coordinate on a shared compliance calendar rather than tracking each obligation in isolation. Longer term, the retroactive-exposure dynamic argues for building documentation and evidentiary discipline into AI vendor contracts now, so that whatever regulatory attention a provider draws in the future does not surface gaps the organization only discovers after the fact.

CSA Resource Alignment

This note extends CSA's own pre-enforcement analysis, "EU AI Act GPAI: Security Compliance Before August 2026" (May 2026), which laid out a twelve-week readiness framework built around the same August 2, 2026 activation date this note addresses; where that note focused on preparation, this one addresses what the mechanism actually does now that the date has passed and enforcement is live. CSA's "EU AI Act Article 50: Transparency Obligations Take Effect" (July 2026) covers the parallel transparency-obligation track that activated on the identical date, and security teams managing both tracks should read the two notes together rather than treating them as separate compliance projects. Finally, CSA's "EU AI Act Digital Omnibus: Enterprise Risk Recalibration" (June 2026) provides the broader regulatory-timeline context needed to understand why the GPAI enforcement date held firm even as the high-risk system deadline moved, which matters for any organization trying to reconcile multiple, differently-timed EU AI Act workstreams into a single governance calendar.

References

[1] Wilson Sonsini. "EU AI Act Enforcement Phase Begins." Wilson Sonsini Goodrich & Rosati, August 2026.

[2] Taylor Wessing. "GPAI obligations under the EU AI Act: Enforcement has started 2 August 2026." Taylor Wessing, August 2026.

[3] EU Artificial Intelligence Act (tracker). "Enforcement of Chapter V under the EU AI Act." artificialintelligenceact.eu, 2026.

[4] EU Artificial Intelligence Act (tracker). "Article 101: Fines for Providers of General-Purpose AI Models." artificialintelligenceact.eu, 2026.

[5] Enterprise DNA. "EU AI Act Enforcement Is Live: Fines Now Real." Enterprise DNA, August 2026.

[6] EU Artificial Intelligence Act (tracker). "Section 2: Obligations for Providers of General-Purpose AI Models." artificialintelligenceact.eu, 2026.

[7] EU Artificial Intelligence Act (tracker). "Article 53: Obligations for Providers of General-Purpose AI Models." artificialintelligenceact.eu, 2026.

[8] EU Artificial Intelligence Act (tracker). "Article 88: Enforcement of the Obligations of Providers of General-Purpose AI Models." artificialintelligenceact.eu, 2026.

[9] European Commission. "Guidelines for providers of general-purpose AI models." Shaping Europe's Digital Future, 2026.

[10] Wilson Sonsini. "EU Releases Final Code of Practice for General-Purpose AI Models." Wilson Sonsini Goodrich & Rosati, 2025.

[11] Greenberg Traurig. "Deepfakes, Chatbots, AI-Generated Text: European Commission Details Transparency Obligations Under the AI Act." Greenberg Traurig, June 2026.

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

[13] European Commission. "AI Act: Commission publishes a reporting template for serious incidents involving general-purpose AI models with systemic risk." Shaping Europe's Digital Future, 2026.

[14] EU Artificial Intelligence Act (tracker). "Article 91: Power to Request Documentation and Information." AI Act Service Desk, European Commission, 2026.

[15] JD Supra. "The EU AI Act: New Investigative Powers for the European Commission – What AI Providers Should Know." JD Supra, 2026.

[16] EU Artificial Intelligence Act (tracker). "Article 54: Authorised Representatives of Providers of General-Purpose AI Models." artificialintelligenceact.eu, 2026.

[17] Stephenson Harwood. "EU: Obligations on providers of GPAI models under the EU AI Act." Stephenson Harwood, 2026.

[18] Gibson Dunn. "EU AI Act Omnibus Agreement – Postponed High-Risk Deadlines and Other Key Changes." Gibson Dunn, May 2026.

[19] EU Artificial Intelligence Act (tracker). "Article 89: Monitoring Actions." artificialintelligenceact.eu, 2026.

[20] EU Artificial Intelligence Act (tracker). "Article 90: Alerts of Systemic Risks by the Scientific Panel." artificialintelligenceact.eu, 2026.