FEATURE INVESTIGATION | JUNE 2026

Supply Unchained

The Accountability Gap

How Big Tech's AI Safety Promises Measure Up Against the Law That Will Judge Them

We applied the EU AI Act's article-by-article risk and control framework to the responsible-AI documentation published by four leading model providers — the companies whose models now power most of the world's enterprise AI. What we found should concern every procurement director, ESG officer, and board member who has deployed their products.

By Nayantara Sriram · Founder & CEO, Supply Unchained | June 20, 2026

The AI era has a trust problem — and it is structural, not incidental. The four companies that have effectively become the infrastructure of global enterprise AI have built impressive, well-intentioned responsible-AI frameworks. They have published policies, hired ethicists, conducted red-team exercises, and written governance principles with evident sincerity. But when you apply a consistent, article-by-article compliance methodology — the same framework used to assess any other regulated industry from pharmaceuticals to financial services — a disquieting picture emerges.

None of the four providers would pass an EU AI Act compliance audit today. Not one has registered in the EU AI Act's public database, despite an enforcement deadline of August 2, 2026. Not one has published a training-data copyright compliance policy with an opt-out mechanism for EU rights holders — an obligation that has been legally in force since August 2, 2025. Not one has an Article 73-compliant incident reporting protocol that would notify national competent authorities within the 15-day window required when an AI system is linked to death or serious harm.

> *The companies whose AI systems power most of the world's enterprise decisions are, collectively, in a state of structured non-compliance with the legal framework designed to govern them.*

This is not a story about bad actors. It is a story about an industry that has moved faster than its own governance structures, and a regulatory framework that has moved faster than the industry expected. The collision point is now. With the EU AI Act fully applicable in August 2026 and the AI Office beginning enforcement monitoring of General-Purpose AI providers simultaneously, the gap between what companies have published and what the law requires is about to become very expensive.

This article does three things. First, it explains the EU AI Act's risk-based framework and why it matters for responsible-AI governance beyond Europe. Second, it lays out the methodology — a structured, article-by-article risk and control assessment matrix — that we applied consistently

across all four providers. Third, it presents our findings: what they are doing well, where the critical gaps lie, and why those gaps matter not just as regulatory exposure but as questions of trust, safety, and human rights.

I. The Regulation That Changed Everything

The EU Artificial Intelligence Act — Regulation (EU) 2024/1689 — entered into force on 1 August 2024. It is the world's first comprehensive legal framework governing artificial intelligence, and it has no meaningful precedent in any prior technology regulation. Its scope is extraordinary: it applies to providers, deployers, importers, and distributors of AI systems that affect people in the European Union, regardless of where those companies are headquartered. A model built in California and deployed via API to a Warsaw HR department is subject to the Act. An open-source model released in the United States and fine-tuned elsewhere for EU credit assessment is subject to the Act.

A Risk-Based Architecture

The Act's genius — and its complexity — lies in its risk-proportionate structure. Rather than regulating AI as a monolith, it creates four tiers of obligation calibrated to the potential for harm:

- **Unacceptable risk (Article 5):** Absolute prohibitions. These include real-time biometric identification in public spaces, social scoring systems, and AI that exploits psychological vulnerabilities. In force since February 2, 2025. In May 2026, the EU added non-consensual intimate deepfakes ('nudifiers') to this list — prompted in part by high-profile generative-image incidents on a major social platform in late 2025. Maximum penalty: €35 million or 7% of global turnover.
- **High risk (Articles 6–27):** Permitted, but subject to comprehensive pre-market obligations — risk management, data governance, technical documentation, audit logging, transparency, human oversight, and accuracy/robustness requirements. Obligations apply to specific use cases in employment, credit, healthcare, education, and law enforcement under Annex III.
- **Limited risk (Article 50):** Chatbots and AI-generated content must disclose their AI nature. Synthetic content must be watermarked. Full application from August 2, 2026.
- **Minimal or no risk:** No specific obligations. Covers the vast majority of AI applications.

Cutting across these tiers is a parallel set of obligations targeting General-Purpose AI (GPAI) models — the large foundation models that power downstream applications across every sector. These obligations (Articles 53–55) entered force on August 2, 2025 and require technical documentation, training-data transparency, copyright compliance, and — for systemic-risk models above the 10^{25} FLOP threshold — adversarial testing, EU AI Office incident reporting, and safety frameworks.

> **A note on VLOPs.** The November 2025 AI Omnibus proposal introduces an important new layer — centralising oversight of AI systems embedded in Very Large Online Platforms (VLOPs) and Very Large Online Search Engines (VLOSEs) with the EU AI Office. This means that major search, video, and social platforms face convergent AI Act and Digital Services Act obligations. Several of the companies in our assessment are not merely AI developers; they are, in some cases, among the most regulated platform entities in the EU regulatory universe.

Why This Matters Beyond Brussels

The EU AI Act is not merely a European compliance exercise. Its extraterritorial reach — established through the nexus test (systems deployed to EU users or affecting EU persons) — means that every multinational company using AI tools procured from these four providers is, knowingly or not, a deployer subject to the Act's obligations. Under Article 26, deployers of high-risk AI systems carry independent legal obligations regardless of provider compliance. Under Article 29, they face civil liability for failures to take 'all appropriate measures' to prevent harm.

For supply chain and procurement professionals, this creates a direct and immediate exposure. If you have deployed an AI system for supplier risk screening, employee background assessment, logistics optimisation, or contract management — and that system produces discriminatory or harmful outputs — your organisation is a liable deployer under EU law. The quality of your AI provider's governance framework is no longer just a vendor evaluation criterion. It is a legal risk factor.

II. The Methodology: A Consistent Yardstick

We applied a structured EU AI Act Risk and Control Matrix to each of the four providers — the same framework, the same scoring criteria, the same 25 control areas, assessed against the same Gold Star standard for each article. The methodology is derived from established compliance governance frameworks, including the DOJ's Evaluation of Corporate Compliance Programs and the OECD Due Diligence Guidance for Responsible Business Conduct, adapted specifically for EU AI Act obligations.

How the Framework Works

For each article, we assess five dimensions:

- **Evidence:** What has the company actually published — policies, system cards, governance frameworks, model documentation — that addresses the specific obligation?
- **Gap analysis:** What is missing relative to the EU AI Act's Gold Star standard — the aspirational, fully-effective implementation that regulators would expect of a compliant, mature provider?
- **Inherent risk:** What is the raw exposure if no controls exist? Scored on a Likelihood × Severity matrix (0–5 × 0–5 = 0–25).
- **Residual risk:** What exposure remains after accounting for the company's current controls and their effectiveness?
- **Compliance score:** A 0–100 scale reflecting how far the company's published evidence meets the Gold Star standard for that article.

Scores are then aggregated to an overall compliance rating: Strong Compliance (80+), Partial Compliance (60–79), Limited Compliance (35–59), or Non-Compliant (below 35). Critically, we score on published evidence only. We cannot assess what companies do privately. A control that exists but is not documented provides no legal defence and no public accountability — and is therefore treated as absent for scoring purposes.

> **On fairness.** We have applied the same standards to each provider. We are assessing published responsible-AI frameworks against a consistent regulatory standard — not the companies' intentions, internal culture, or private practices. Where we find strengths, we name them. Where we find gaps, we name those too.

One important calibration: the four providers are not comparably positioned in the EU AI Act's regulatory architecture. The three closed-API providers (Providers A, B, and C) can, in principle, enforce compliance controls across all deployments of their models. The open-weights provider (Provider D, which releases model weights publicly) cannot — once weights are distributed, it has no runtime enforcement capability. This creates a structurally different compliance challenge for Provider D that the framework acknowledges explicitly.

III. The Findings

The headline numbers are informative, but the story is in the detail.

Provider	Score	Rating	Strongest Area	Critical Gap	GPAI CoP
Provider A(closed-API)	**48/100**	**LIMITED (upper)**	Adversarial testing & content watermarking	Training-data provenance (Art 10)	Not signed
Provider B(closed-API, agentic)	**43/100**	**LIMITED**	GPAI Code of Practice signed (Art 95)	Copyright compliance & DB registration	**Signed Aug 2025**
Provider C(closed-API)	**37/100**	**LIMITED**	Systemic-risk safety framework	EU DB registration & Art 73 reporting	Not signed
Provider D(open-weights)	**29/100**	**NON-COMPLIANT**	Deployer guidance	Open-weights enforcement impossibility	Not signed

Source: Supply Unchained EU AI Act Risk and Control Matrix assessment, June 2026. All scores based on published documentation only.

Provider A: Industry-Leading Practice, Structural Gaps

Provider A scores highest in our assessment — 48/100 — but this should not be read as a passing grade. It means that Provider A, the best-performing of the four, addresses roughly half of its EU AI Act obligations at the level a regulator would expect. The remaining half comprises genuine compliance gaps, not theoretical risks.

Where Provider A genuinely leads the industry: its adversarial testing programme is world-class. Its internal red team conducted over 350 exercises in 2025 alone, spanning text, audio, images, video, and agentic AI. A panel of independent external evaluators — including a national AI security institute — provides independent validation. This goes well beyond what any other company in our assessment has published.

Equally significant is its content-provenance work. The combination of imperceptible watermarking across all AI-generated modalities, a public verification portal, and native camera-level content-credential integration on its flagship device represents the most comprehensive approach to synthetic content provenance of any AI provider. This matters enormously for the EU AI Act's Article 50 transparency obligations — which come into full force in August 2026 and will require AI-generated content to be labelled and detectable.

But Provider A's most significant compliance gap is also the one that matters most for human rights due diligence: training-data governance. Its most recent responsible-AI progress report describes research on AI bias and output quality through a public factuality benchmark, but publishes nothing about how its models' training datasets were audited for demographic representation, provenance tracked, or copyright compliance verified. For an organisation of this data scale, this is the single

largest gap between its published responsible-AI posture and what Article 10 of the EU AI Act requires.

> **The gap that concerns us most for Provider A.** A deployed medical-imaging screening AI — described as supporting nearly one million screenings — has obtained a CE marking and deployed globally. It may already qualify as a high-risk AI system under Annex III of the EU AI Act (medical-device AI), requiring registration in the EU database. That registration deadline is August 2, 2026. No registration plan is mentioned anywhere in publicly available documentation.

Provider B: The Strategic Move That Changes the Calculus

Provider B's overall score of 43/100 places it second in our assessment, but this understates its strategic position relative to the others. In August 2025, it signed the EU AI Office's GPAI Code of Practice — the single most consequential EU AI Act compliance action taken by any company in this assessment series.

The Code of Practice, finalised by independent experts in July 2025 and endorsed by the European Commission, provides a presumption of conformity with GPAI obligations. Signatories commit to transparency, copyright compliance, and safety frameworks as binding undertakings. The AI Office will treat Code adherence as compliance with the Act's GPAI obligations — and will focus enforcement monitoring on signatories from August 2026. By signing, Provider B has placed itself in a collaborative rather than adversarial relationship with the EU's primary AI regulator.

This strategic move is inseparable from its obligations. Provider B has committed — publicly and bindingly — to publish a training-data summary using the EU AI Office's mandatory template, to implement a copyright opt-out mechanism for EU rights holders under the DSM Directive, and to establish an incident reporting pipeline to the AI Office. None of these has been implemented yet. The AI Office begins enforcement monitoring in approximately 14 months. The gap between commitment and demonstrated implementation is Provider B's primary compliance risk.

Provider B's computer-use agent product — capable of autonomously browsing the web, making purchases, filling forms, and executing multi-step tasks — deserves particular scrutiny. No other product in our assessment creates such concentrated Article 14 (human oversight) and Article 5 (prohibited practices) exposure. An agent taking real-world consequential actions on a user's behalf, operating across the EU's digital infrastructure, may qualify as a high-risk AI system under multiple Annex III categories simultaneously. Yet no published Annex III classification assessment for the agent product exists.

Provider C: The Most Sophisticated Safety Framework, The Most Regulatory Gaps

Provider C presents a paradox. Its systemic-risk safety framework (current version effective May 2026) is, in our assessment, the most intellectually sophisticated internal AI safety framework published by any of the four companies. Its capability thresholds for CBRN risks, misaligned AI, and automated AI R&D are more demanding than the EU AI Act's Article 55 systemic-risk provisions. Its governance structures — a dedicated scaling-oversight officer, an independent benefit trust, and external review obligations — are more robust than anything the other three have established publicly.

And yet Provider C scores 37/100 — the second-lowest of the four. Why? Because the framework was designed to address a different problem: catastrophic AI risk, defined as existential threats and systemic civilisational harm. The EU AI Act was designed to address a broader, more prosaic problem: harm to individual people arising from AI systems deployed in everyday commercial contexts. These are complementary but not coextensive concerns, and the framework's sophistication in addressing the former leaves significant gaps in addressing the latter.

The core tension in Provider C's position is this: an organisation that has thought more carefully about AI safety than perhaps any other commercially operating AI lab has not mapped its thinking to the regulatory framework that will govern it. Its CBRN capability threshold does not cite Article 5. Its governance structures do not reference Article 9's risk management requirements. Its external review commitment does not engage with Article 55's systemic-risk assessment obligations. These are not the same thing, and the gap between them is where Provider C's regulatory exposure concentrates.

Provider C's framework demonstrates that it thinks hard about safety. What it does not demonstrate is that it has thought hard about compliance.

Provider D: The Open-Source Paradox

Provider D's responsible-use guidance scores 29/100 — Non-Compliant — in our assessment. But the score, and the finding, requires careful contextualisation. Its compliance challenge is structurally different from the other three, and in important ways more difficult than anything a compliance programme can fully resolve.

The open-weights distribution model means that once model weights are released, Provider D cannot enforce downstream compliance, implement runtime controls, monitor production performance, or respond to serious incidents. The EU AI Act was designed with closed-API providers in mind. Its Article 12 (audit logging), Article 14 (human oversight), Article 15 (accuracy monitoring), Article 72 (post-market monitoring), and Article 73 (incident reporting) obligations all presuppose that a provider retains some degree of runtime visibility and control over its model's deployment. For an openly distributed model, that visibility and control disappear at the moment of weight release.

This is not a problem Provider D can solve through better compliance documentation. It is a fundamental tension between the open-source AI distribution model and a regulatory framework designed for centralised, controlled deployment. Provider D needs to engage the EU AI Office in a policy dialogue about what provider-level obligations are achievable for open-weights GPAI models — and it needs to do so urgently, because the AI Office begins enforcement monitoring of GPAI providers in August 2026.

Its responsible-use guide is, in one respect, the best deployer-facing guidance in our assessment. The guide's four-stage development framework (use-case determination, model-level alignment, system-level alignment, transparency mechanisms) is practical, detailed, and genuinely useful for developers building on the model. Provider D has created substantive responsible-AI tooling — input/output safety classifiers, a cybersecurity evaluation suite, and a code-security filter — that

many deployers use. But guidance and tooling are not controls. They are education. And education does not substitute for enforcement.

IV. The Universal Gaps: Where All Four Companies Fail

Beyond the company-specific findings, our assessment reveals five areas where all four providers have critical gaps — obligations that are already in force and that none of the assessed companies has demonstrably addressed. These are not speculative future risks. They are present non-compliance.

Article	A	B	C	D	Issue
Art 5 – Prohibited Practices	40	35	20	10	None of the four has published an explicit Art 5 compliance mapping. The open-weights model cannot enforce prohibitions post-distribution.
Art 10 – Training Data	35	35	30	35	No company publishes demographic bias audits or Art 10-compliant data-governance programmes. Universal critical gap.
Art 53 – GPAI Copyright	20	30	20	15	No opt-out mechanism or training-data summary published by any company. All face active copyright litigation. Provider B marginally ahead via CoP commitment.
Art 71 – EU DB Registration	15	15	10	10	Complete absence across all four. August 2026 enforcement monitoring is imminent. Provider A's healthcare AI likely already registrable.
Art 73 – Incident Reporting	30	30	25	20	No Art 73-compliant NCA notification protocol at any company. 3–15 day reporting timelines not operationalised. Agentic AI products create elevated exposure.

Source: Supply Unchained EU AI Act Risk and Control Matrix assessment, June 2026. Red = below 20 (Non-Compliant). Amber = 20–35 (Limited–Non-Compliant).

The Training Data Reckoning

The training-data transparency obligation is the one that will define the next phase of AI regulation. Article 53(1)(d) requires every GPAI provider to publish a 'sufficiently detailed summary' of their training content using the mandatory template published by the EU AI Office in July 2025. For models placed on the EU market from August 2, 2025 onwards, this obligation is already in force. For pre-existing models, providers have until August 2, 2027.

None of the four companies has published this summary. What makes this particularly significant is that the training-data summary template is not merely a transparency exercise. It requires providers to describe how they complied with EU copyright law, how they implemented opt-out mechanisms for rights holders, and what content moderation measures they used during data collection. For companies simultaneously facing active copyright litigation from publishers, authors, and news organisations, the training-data summary is a document with significant legal consequences — which may partly explain the collective reticence to publish it.

The OECD published its Due Diligence Guidance for Responsible AI in May 2026, explicitly connecting responsible AI development with the human rights due diligence obligations established by the OECD Guidelines for Multinational Enterprises. The Guidance takes a deliberately risk-agnostic approach — it does not prescribe which AI risks to prioritise, but it establishes that AI training-data sourcing, like any other value-chain activity, is subject to the same due diligence obligations as sourcing raw materials from conflict regions or manufacturing goods using third-party

labour. If a company would conduct a Tier 2 supplier audit for forced labour risks, it should apply equivalent rigour to the provenance of the data that trains its AI systems.

V. Trust, Safety, and Human Rights: The Framework Argument

The EU AI Act is regulatory compliance. But the underlying concern that animates it — and that connects it to the CSDDD, the OECD AI Principles, and the UN Guiding Principles on Business and Human Rights — is something more fundamental. It is the question of whether AI systems deployed at scale can be trusted to treat people fairly, to protect their rights, and to be accountable when they fail.

The HRDD Connection

The Corporate Sustainability Due Diligence Directive, even in its Omnibus-simplified form (effective March 2026, applying from July 2029 for companies with over 5,000 employees and €1.5 billion turnover), creates a powerful connection between AI governance and human rights due diligence. The CSDDD requires companies to identify, prevent, mitigate, and account for adverse human rights impacts across their chain of activities.

AI systems are not exempt from this obligation. If a company deploys an AI system for supply-chain risk screening that exhibits racial or geographic bias — systematically downweighting suppliers from the Global South while favouring European counterparts — that system is a potential source of adverse human rights impact under the CSDDD. If a company uses AI for employment screening that produces discriminatory outcomes by gender, disability, or ethnicity, that is not merely an employment-law issue. It is an HRDD issue, and it requires the same structured response: identification, prevention, mitigation, and accountability.

This is precisely why the training-data governance gap matters so much. An AI system trained on biased data produces biased outputs. Those outputs, when deployed in high-stakes decisions — lending, hiring, insurance pricing, supply-chain partner selection, medical triage — are the source of adverse human rights impacts that the CSDDD and the EU Charter of Fundamental Rights are designed to prevent. The gap between what these companies have published about their training-data governance and what Article 10 of the EU AI Act requires is not a technical compliance detail. It is the mechanism by which bias enters the system.

> *Responsible AI governance is not a compliance exercise layered on top of AI development. It is the condition under which AI development earns the right to deploy into human lives.*

The Accountability Architecture

What distinguishes a genuinely trustworthy AI governance framework from a sophisticated press release? We suggest four criteria drawn from the DOJ's Evaluation of Corporate Compliance Programs — the global gold standard for assessing compliance-programme effectiveness.

- **Autonomy and resources:** Does the compliance function have sufficient seniority, independence, and budget to challenge adverse business decisions? At Provider C, the scaling-oversight officer and independent benefit trust provide meaningful institutional

independence. At Provider A, the senior AI governance council includes parent-company board members. At Providers B and D, the evidence is thinner.

- **Continuous testing:** Are controls tested continuously, or only at launch? Provider A's red-team programme (350+ exercises per year) sets the standard. Provider C's 3–6 month risk-report cadence is structured but episodic. Provider D's testing is pre-release only — with no post-deployment visibility.
- **Evidence retention:** Are compliance records maintained in a form that would support a legal defence? None of the four companies has published a 5-year documentation retention commitment of the kind that Article 29 of the CSDDD requires.
- **Consequence management:** When a violation is found, is it investigated promptly and impartially? Provider C's non-compliance reporting processes are the most detailed. Provider A's trust-and-safety infrastructure is the most mature operationally. Providers B and D's mechanisms are less transparent.

Applying these criteria, the conclusion is uncomfortable: none of the four companies has a compliance programme that would satisfy the DOJ's own standard for effectiveness — and the DOJ standard is the international baseline. The EU AI Act is more demanding in several respects. The CSDDD is more demanding in others. The convergence of these frameworks, alongside the OECD AI Due Diligence Guidance, is creating a compliance architecture that no company in our assessment currently meets.

VI. What Deployers Need to Do Now

If you are a procurement director, a Chief Risk Officer, or an ESG lead at a company that has deployed AI tools built on any of these four providers' models, this assessment has direct implications for your organisation. You are a deployer under the EU AI Act. You carry Article 26 obligations that are independent of your provider's compliance status. You carry Article 29 civil-liability exposure under the CSDDD for any adverse human rights impacts caused by AI systems you deploy. And you carry reputational risk for every bias, every harmful output, and every privacy failure that your AI deployments produce.

Immediate Actions

- **Conduct an AI system inventory.** Identify every AI system your organisation uses or deploys, the provider, the model version, and the use case. Map each system against the Annex III high-risk classification criteria. If any system falls into employment screening, credit assessment, essential services, healthcare, or law enforcement categories, your organisation has active Art 26 compliance obligations.
- **Demand provider documentation.** Request from your AI provider: (a) a model card structured to Annex IV/XI requirements; (b) confirmation of GPAI Code of Practice signing status; (c) evidence of training-data copyright compliance; (d) the inference logging capabilities available for your deployment; and (e) their Art 73 incident reporting protocol. If they cannot provide these, that is material information for your vendor risk assessment.
- **Assess your own deployer obligations.** Under Article 26, you must: implement appropriate technical and organisational measures; ensure use within the provider's intended purpose; implement human oversight; and control input data quality. Under Article 27 (if high-risk), conduct a Fundamental Rights Impact Assessment. Under Article 73, establish your own incident reporting pipeline with Art 73-compliant timelines.

- **Integrate AI risk into your HRDD programme.** Identify every AI system used in decisions that affect people — employees, suppliers, customers, communities. Apply the same OECD due diligence logic you would apply to any other high-risk supply-chain activity: identify potential adverse impacts, prevent them, mitigate them, monitor for them, and document your response.
- **Register high-risk AI systems.** Before August 2, 2026, complete EU database registration for any AI system you have classified as high-risk. This is the deployer's obligation, not the provider's, for deployer-built systems. Check whether your provider has registered shared-responsibility systems.

The Longer Game

The EU AI Act's August 2026 full-application date is not the finish line. It is the starting gun for active regulatory supervision. The AI Office has committed to verification of GPAI provider compliance from that date. Finland became the first EU member state with operational AI Act enforcement powers in January 2026. Other NCAs are activating through Q1 and Q2 2026. The enforcement infrastructure is being built, rapidly, in parallel with the compliance timelines.

Companies that approach this moment with the goal of minimal documentation — the equivalent of buying a fire extinguisher after the building inspector has left — will not be adequately positioned. The AI Act's penalty structure, with compounding multi-article exposure reaching €35 million or 7% of global turnover per category, makes the compliance investment straightforward to justify. More importantly, the companies that build genuine, auditable, evidence-backed AI governance frameworks will be positioned to deploy AI in the ways that matter most — not just as productivity tools, but as instruments for addressing the problems that the world actually needs AI to solve.

Conclusion: The Trust Dividend

The companies in our assessment are not irresponsible. They employ thoughtful people who are genuinely trying to build AI systems that benefit humanity. The researchers who designed sophisticated alignment frameworks, the engineers who built industry-leading content-watermarking systems, the teams that wrote rigorous model specifications and responsible-use guides — these are people who understand the stakes. The problem is not intention. It is architecture.

Responsible AI governance requires the same structural rigour that we demand of any other high-stakes compliance domain — pharmaceuticals, financial services, aviation safety. It requires documented controls, not just stated principles. It requires independent testing, not just internal red teams. It requires external accountability, not just internal review. And it requires evidence retention — the auditable paper trail that proves, in a court of law or before a national competent authority, that the company took 'all appropriate measures' to prevent harm.

None of the four companies currently meets that standard in full. All four are making progress. One — Provider B, through its GPAI Code of Practice signing — has made a binding strategic commitment to a regulatory relationship that will hold it to account. The others will face a choice: engage proactively with the regulatory framework and help shape it, or wait for enforcement to create the forcing function.

For those of us in supply chain, procurement, and ESG — the practitioners who will carry the downstream compliance consequences of our providers' choices — the message is clear. The era of trusting AI providers' responsible-AI commitments at face value is over. The era of demanding evidence has begun.

Trust in AI is not given. It is earned through documentation, accountability, and the willingness to be audited. The EU AI Act is, at its core, a mechanism for making that earning process mandatory.

The companies that understand this — and that build the governance infrastructure to demonstrate it — will not merely survive the regulatory moment. They will earn the trust dividend: the competitive advantage that comes from being the AI provider that enterprise customers can rely on, that regulators can engage with, and that the public can hold to account. In a world where AI is becoming infrastructure, that dividend is worth more than any feature advantage or benchmark lead.

METHODOLOGY NOTE

This assessment is based on publicly available documentation reviewed in May–June 2026, including each provider's most recent published safety framework, responsible-AI progress report, preparedness framework, model specification, responsible-use guide, and system cards. The EU AI Act Risk and Control Matrix framework is available as a companion Excel workbook covering all four providers with full article-by-article scoring.

Scores reflect published evidence only. We cannot assess private practices. This is not legal advice. Organisations with specific EU AI Act compliance questions should consult qualified legal counsel in the relevant jurisdiction.

The framework cites Regulation (EU) 2024/1689 as amended by the Digital Omnibus (political agreement May 7, 2026), the EU AI Office's GPAI Code of Practice (July 2025), the Commission's Training Data Summary Template (July 2025), the OECD Due Diligence Guidance for Responsible AI (May 2026), CSDDD as amended by Omnibus I (in force March 2026), and the DOJ Evaluation of Corporate Compliance Programs.

— END —

Supply Unchained

AI Governance · Human Rights Due Diligence · Supply Chain Risk

supplyunchained.co.uk