Supply
Unchained

AI GOVERNANCE · HUMAN RIGHTS DUE DILIGENCE · SUPPLY CHAIN RISK

FEATURE INVESTIGATION · JULY 2026

The Golden Thread: Fundamental Rights Impact Assessments and the New Rules for AI

Human rights are becoming the single measure by which the world's digital laws judge harm — and Fundamental Rights Impact Assessments are how AI companies will be asked to prove they measured it first.

Nayantara Sriram — Founder & CEO, Supply Unchained

9 July 2026 · London

The argument at a glance

Five minutes' reading for a board; the evidence for each point follows in the sections indicated.

2 Dec 2026

NEW AI ACT PROHIBITION

AI systems that generate CSAM or non-consensual intimate imagery are banned outright — the Digital Omnibus's only new prohibition.

2 Dec 2027

HIGH-RISK TIER & FRIA

Annex III obligations — including the Article 27 Fundamental Rights Impact Assessment — deferred 16 months from August 2026.

€35M / 7%

TOP PENALTY TIER

The new prohibition sits in the Act's highest fine band, enforced by an AI Office that now holds inspection powers.

01 **The Digital Omnibus deferred deadlines but hardened child protection.** In a package designed to simplify, the only new prohibition targets AI-generated CSAM and NCII — the clearest signal of regulatory priorities. (Section 2)

02 **Human rights is the common measure across every digital regime.** The AI Act, DSA, UK OSA and US child-safety statutes all descend from the same due-diligence grammar: identify, prevent, mitigate, account. One harms-based engine can serve them all. (Sections 3–4)

03 **For VLOPs and VLOSEs, convergence is now structural.** The omnibus hands the EU AI Office supervision of AI systems integrated into designated platforms — so content moderation, consumer protection and AI safety are no longer separate compliance conversations. (Section 5)

04 **The FRIA is human rights due diligence, codified.** Article 27 converts two decades of HRDD practice into an enforceable pre-deployment gate — and deployers will demand FRIA-ready documentation from their AI providers long before the 2027 deadline. (Section 6)

05 **Undocumented controls are treated as absent.** The Accountability Gap scored four major TMT platforms 48, 43, 37 and 29 out of 100; re-scored under the new Art. 5 CSAM prohibition, every score falls — the best to 45. None passes. (Section 7.) The practical guide in Section 9 sets out what to do next, by deadline, across the EU, UK and US.

CONTENTS

01 First principles: what "fundamental rights" means

Before the deadlines and the acronyms, the term doing all the work. Every instrument in this article measures harm against the same object — so defining that object precisely is the thesis of the piece.

The EU term: fundamental rights

In EU law, fundamental rights are the rights codified in the **Charter of Fundamental Rights of the European Union**, legally binding since the Treaty of Lisbon in 2009 and organised under six titles — dignity, freedoms, equality, solidarity, citizens' rights and justice. They include privacy and data protection, non-discrimination, freedom of expression, consumer protection and, at Article 24, the rights of the child. When the AI Act uses the phrase — in its Article 1 purpose clause, in the Annex III logic that classifies AI as high-risk precisely where it touches these rights, and in the Article 27 Fundamental Rights Impact Assessment — it imports the entire Charter, with its case law, into product regulation.

The international term: human rights

Human rights is the same commitment in the international-law dialect. The line runs from the Universal Declaration of Human Rights (1948) through the twin UN covenants to the specialised conventions — above all, for this article, the **UN Convention on the Rights of the Child**, the most widely ratified human-rights treaty in existence. Article 34 obliges states to protect children from "all forms of sexual exploitation and sexual abuse," and the UN Committee's General Comment No. 25 (2021) extends the Convention explicitly to the digital environment. The Charter itself declares that where its rights correspond to the European Convention on Human Rights, their meaning and scope are the same. The two vocabularies are two enforcement architectures built over one rights corpus — what differs is who is bound and where you litigate, not what is protected.

	FUNDAMENTAL RIGHTS (EU DIALECT)	HUMAN RIGHTS (INTERNATIONAL DIALECT)
Source	EU Charter of Fundamental Rights; ECHR via Art. 52(3)	UDHR; UN covenants; CRC; regional conventions
Who is bound	EU institutions; Member States implementing EU law — and, through the AI Act and DSA, the companies they regulate	States as duty-bearers; companies through the UNGPs' responsibility to respect
Where enforced	EU and national courts; Commission, AI Office and national regulators	Treaty bodies and national implementation; HRDD legislation (CSDDD)
Child-protection anchor	Charter Art. 24 — the child's right to protection and care	CRC Art. 34 — protection from all forms of sexual exploitation and abuse; General Comment No. 25 (digital environment)

The corporate bridge: UNGPs and HRDD

For companies, the bridge between the two dialects is the **UN Guiding Principles on Business and Human Rights** (2011): states must protect rights, businesses must respect them, and victims must have access to remedy — operationalised through human rights due diligence: identify, prevent, mitigate, and account for adverse impacts on people. The OECD Guidelines made HRDD the global expectation; the CSDDD is making it European law; and the AI Act's FRIA renders it as product regulation. Same discipline, escalating enforceability.

KEY TERMS	
Fundamental rights	The EU Charter's binding rights catalogue — the AI Act's operative standard of harm.
Human rights	The same rights in international law: UDHR, UN covenants, CRC, ECHR.
UNGPs / HRDD	The corporate duty to respect rights, run as a due-diligence cycle: identify → prevent → mitigate → account.
FRIA	The AI Act's Article 27 assessment — HRDD codified as a pre-deployment legal gate, notified to the regulator.
CSAM / CSE	Child sexual abuse material, and the broader category of child sexual exploitation — grooming, sextortion and trafficking conducted or facilitated through a platform.

This is the thesis that connects everything that follows. Because the AI Act, the DSA, the UK Online Safety Act, COPPA and the rest all reference the same rights corpus, a harm to a person is measured on one instrument across every statute. Child sexual exploitation is the paradigm case: CRC Article 34 becomes the AI Act's new generation ban, the OSA's apex priority offence, the DSA's minors duties, and the US reporting and takedown regimes. One right, five statutes, one golden thread — the rest of this article follows it.

> *"Fundamental rights" and "human rights" are two legal dialects for the same promise. The AI Act's contribution is to make that promise auditable.*

02 What the Digital Omnibus actually changed

On 29 June 2026, the Council of the EU gave final approval to the Digital Omnibus on AI. Most coverage focused on the deadlines it moved. The more telling story is the one thing it made stricter.

The Digital Omnibus on AI was proposed by the European Commission on 19 November 2025, politically agreed on 6–7 May 2026, endorsed by the European Parliament on 16 June 2026, and given final Council approval on 29 June 2026. It rewrites the EU AI Act's compliance calendar in material ways, and any AI company still working from the 2024 timeline is now planning against the wrong dates.

The new prohibition: CSAM and NCII generation

The omnibus adds a new prohibited practice to Article 5 of the AI Act: **AI systems for generating child sexual abuse material (CSAM — depictions of the sexual abuse and exploitation of children) and non-consensual intimate imagery (NCII)**, including so-called "nudification" apps — prompted in part by high-profile generative-image incidents on a major social platform in late 2025. The prohibition applies from **2 December 2026**, after a transitional period, and it sits in the Act's highest penalty tier: up to **€35 million or 7% of global turnover**.

The drafting reaches much further than intent. For **providers**, the prohibition bites not only where generating CSAM or NCII is the system's intended purpose, but where it is a *"reasonably foreseeable and reproducible*

outcome without significant technical modification" and the provider has not implemented adequate safeguards. In plain terms: if your image or video model can be prompted into producing this material without being meaningfully re-engineered, and you cannot evidence the safeguards that prevent it, you are within the prohibition. For **deployers**, the prohibition is narrower — it targets deliberate use — but the provider-side test converts model-level safety engineering from good practice into a hard legal requirement. Accidental generation by a user does not excuse an unguarded model.

> *"In an omnibus designed to simplify and defer, the only new prohibition is a child-protection rule. Regulators are telling the industry, in statute, what they consider non-negotiable."*

EXHIBIT 1

The road to full application — five dates that matter

Source: Regulation (EU) 2024/1689 as amended by the Digital Omnibus on AI. Gold markers: nearest hard deadlines.

EXHIBIT 2

The enforcement calendar after the Digital Omnibus

OBLIGATION	OLD DATE	NEW DATE	WHAT IT MEANS
NCII/CSAM generation ban (new Art. 5 prohibition)	—	**2 Dec 2026**	New absolute prohibition; €35M / 7% tier; provider-side safeguards test
Art. 50 transparency (AI disclosure, deepfake labelling)	2 Aug 2026	**2 Aug 2026** (unchanged)	The nearest hard EU deadline; Art. 50(2) synthetic-content marking gets a grace period to 2 Dec 2026 for systems already on the market
Annex III high-risk duties (employment, education, credit, essential services, law enforcement) — incl. the Art. 27 FRIA	2 Aug 2026	**2 Dec 2027**	16-month deferral, linked to availability of harmonised standards
Annex I high-risk duties (AI in regulated products)	2 Aug 2027	**2 Aug 2028**	One-year deferral
Member-State regulatory sandboxes	2 Aug 2026	**2 Aug 2027**	One-year deferral
GPAI model duties (Arts. 53–55)	In force since 2 Aug 2025 — unchanged		Already live; AI Office enforcement monitoring begins August 2026

Source: Digital Omnibus on AI, final Council approval 29 June 2026; Regulation (EU) 2024/1689 as amended. Supply Unchained analysis.

Two structural changes matter as much as the dates. First, **enforcement is centralised**: the EU AI Office gains exclusive supervisory competence — with inspection powers — over AI systems built on a provider's own general-purpose AI models, and over AI systems integrated into Very Large Online Platforms and Search Engines under the DSA. For the major model providers, that means answering to Brussels directly, not to 27 national authorities. Second, the omnibus **did not soften** the GPAI obligations or Article 50 transparency: it deferred the high-risk tier and hardened the child-protection line. Civil society groups, it should be noted, have criticised the final text for diluting other fundamental-rights protections — a reminder that the omnibus was a political bargain, not a coronation of rights. But the direction of travel on children is unambiguous.

03 The golden thread: human rights as the common measure

Step back from the acronyms and a pattern emerges. The EU AI Act states in Article 1 that its purpose is to protect "health, safety and fundamental rights." The Digital Services Act requires Very Large Online Platforms to assess systemic risks to *fundamental rights* — alongside illegal content, civic discourse, public health, and the protection of minors — every year, and to have those assessments independently audited. The UK Online Safety Act is built on a register of priority harms to people, with child sexual exploitation and abuse (CSEA) at its apex. The US COPPA regime, the state age-verification statutes, and the federal TAKE IT DOWN Act all regulate a single thing: harm to identifiable rights-holders, most of them children.

These regimes were drafted by different legislatures, enforced by different regulators, and litigated under different constitutions. Yet they measure the same thing, in the same way, because they all descend — directly or indirectly — from the same source: the **UN Guiding Principles on Business and Human Rights**, the OECD Guidelines for Multinational Enterprises, the EU Charter of Fundamental Rights, and the UN Convention on the Rights of the Child. Human rights due diligence (HRDD) — identify, prevent, mitigate, account — is the methodological backbone of the DSA's fundamental-rights risk assessments, the analytical grammar of the OSA's risk-assessment duties, and, as of May 2026, the explicit subject of the **OECD Due Diligence Guidance for Responsible AI**, which establishes that AI training-data sourcing and model deployment are subject to the same due diligence obligations as sourcing minerals from conflict regions.

This is the golden thread. It means something practical and, for compliance teams, liberating: **you do not need a separate methodology per statute**. A harms-based assessment discipline — grounded in the UNGP severity criteria of *scale* (how grave), *scope* (how many people), and *irremediability* (can it be put right) — can serve the DSA Article 34 assessment, the OSA illegal-content and children's risk assessments, the AI Act's risk-management system, and the Fundamental Rights Impact Assessment, with jurisdiction-specific outputs layered on top. At Supply Unchained we score human-rights severity on a five-level scale built on exactly those criteria; the point is not our scale but the principle — one engine, many regulatory artefacts.

EXHIBIT 3

The golden thread — one measure runs through every regime

Human rights due diligence — identify, prevent, mitigate, account — is the shared grammar of all four regimes. Supply Unchained analysis.

> *"Human rights is not one compliance topic among many. It is the measuring instrument every digital regulator is now holding — and the AI Act, the DSA, the OSA and the US child-safety statutes are all reading from it."*

04 The child-protection stack: CSAM, CSE and the rules that now bind AI

Nowhere is the golden thread tighter than in the protection of children. Any AI company whose systems can generate, host, recommend, or fail to detect harmful content involving minors now faces an interlocking stack of obligations across jurisdictions. Terminology matters here — and so does using the full weight of the words. **CSAM** is **child sexual abuse material**: the statutory and professional term for any depiction of the sexual abuse or exploitation of a child, including — under multiple modern statutes — AI-generated depictions. **CSE/CSEA** is **child sexual exploitation and abuse**: the broader UK/EU category that CSAM sits within, covering sexual *exploitation* in all its platform-borne forms — grooming, sextortion, live-streamed abuse, and trafficking conducted or facilitated through a platform. An AI company's exposure is never to the material alone: recommender systems that surface children to adults, chat surfaces that enable grooming, and marketplaces that enable trafficking are all exploitation pathways the same statutes reach.

European Union

- **AI Act, Article 5 (from 2 Dec 2026):** the new prohibition on AI systems for generating CSAM and NCII, described above — the first time a horizontal AI statute has banned a category of generative capability outright.
- **DSA, Article 28:** online platforms accessible to minors must put in place appropriate and proportionate measures to ensure a high level of privacy, safety and security for minors; Article 28(2) prohibits ads targeted at minors based on profiling. The Commission's 2025 guidelines expect age-appropriate design, effective age assurance where risk warrants it, and default privacy for minors.
- **DSA, Articles 18 and 34–35:** hosting services must notify authorities of suspicions of offences threatening life or safety (which operationally captures CSAM); VLOPs must assess and mitigate

systemic risks to minors annually — explicitly including risks from generative-AI features integrated into designated services.

- **The proposed CSA Regulation ("chat control"):** still in negotiation. The Council's late-2025 position dropped mandatory detection/scanning, but the file remains live and contested; platforms should track it rather than build against it.

United Kingdom

- **Online Safety Act — illegal content duties (live since March 2025):** CSEA is a priority offence. Services must risk-assess against Ofcom's register of priority offences and operate proportionate systems to prevent, detect, and remove CSEA content — including grooming pathways and AI-generated CSAM, which UK law treats as illegal regardless of whether a real child was involved.
- **Children's duties (live since July 2025):** children's access and risk assessments, and compliance with Ofcom's Protection of Children codes — including **highly effective age assurance** for pornography and other primary-priority content (suicide, self-harm, eating-disorder promotion), safer recommender defaults for children, and controls on priority-harm content. Ofcom updated its children's codes and guidance in 2026 and, with the ICO, issued a March 2026 joint statement putting services on notice that age assurance must demonstrably work.
- **Enforcement:** fines up to £18M or 10% of qualifying worldwide revenue, business-disruption measures up to ISP blocking, and senior-manager criminal liability for information-notice failures. Ofcom's enforcement programmes on age assurance and illegal-content duties are live and producing investigations now.

United States

- **NCMEC CyberTipline reporting (18 U.S.C. §2258A, extended by the REPORT Act):** US providers must report apparent CSAM — mandatory, deadline-bound, with content-preservation duties. This is the hardest-edged reporting obligation in the US content stack, and it applies to AI companies whose services surface such material.
- **TAKE IT DOWN Act:** criminalises publication of NCII — explicitly including AI-generated deepfakes — and, since May 2026, requires covered platforms to operate a notice-and-removal process taking NCII down within 48 hours, enforced by the FTC.
- **COPPA, amended Rule (effective 23 June 2025):** separate opt-in verifiable parental consent for third-party disclosures including targeted advertising; expanded personal-information definitions (biometrics); written data-retention policies with no indefinite retention.
- **Age verification and app stores:** after *FSC v. Paxton* (June 2025), dozens of state age-verification statutes are operative; the Texas and Utah app-store acts (from January 2026) require age-category verification and parental consent for minors' downloads.
- **AI-specific state law:** Texas TRAIGA prohibits AI systems for CSAM generation outright; California SB 243 imposes disclosure, self-harm protocols, and minor-specific safeguards on companion chatbots.
- **The federal pipeline:** the KIDS Act package — bundling KOSA, COPPA 2.0, the SAFE BOTs Act on chatbots, and the SCREEN Act on age verification — passed the House on 29 June 2026. It is not yet law, but it is the design brief for where US law is heading. Notably, the December 2025 preemption Executive Order **explicitly carves out child-safety rules**: whatever happens to state AI laws in the courts, the kids-and-AI stack is the stable floor.

Australia completes the picture: the Social Media Minimum Age regime (in force 10 December 2025) requires age-restricted platforms to take reasonable steps to keep under-16s off, with penalties up to AUD ~$49.5M — and layered age assurance, not self-declaration, is the compliance standard.

EXHIBIT 4

The child-protection stack at a glance — instruments, status and exposure

JURIS.	INSTRUMENT	CORE DUTY FOR AI COMPANIES & PLATFORMS	STATUS	MAX EXPOSURE
EU	AI Act Art. 5 (as amended)	No AI systems for CSAM/NCII generation; safeguards must be adequate and evidenced	2 Dec 2026	€35M / 7% turnover
EU	DSA Arts. 18, 28, 34–35	Minors protection & no profiling ads to minors; CSAM notification; GenAI in annual systemic-risk assessments	Live	6% turnover
UK	OSA + children's codes	CSEA priority offence; highly effective age assurance; product-triggered risk assessments	Live 2025	£18M / 10%; criminal liability
US	18 U.S.C. §2258A + REPORT Act	Mandatory NCMEC CyberTipline reporting with preservation duties	Live	Federal penalties per failure
US	TAKE IT DOWN Act	48-hour NCII notice-and-removal, incl. AI deepfakes	Live May 2026	FTC §5 enforcement
US	COPPA amended Rule	Separate opt-in consent for third-party disclosure; written retention policies	Live Jun 2025	>$50k per violation, per child
AU	Social Media Minimum Age	Reasonable steps to exclude under-16s; layered age assurance	Live Dec 2025	AUD ~$49.5M

Status as at 9 July 2026. Supply Unchained analysis of statutory texts and regulator guidance.

> *"Across four continents, the same rule is crystallising: if your system can put a child in harm's way, you must prove — in advance and with evidence — that you assessed and mitigated that risk."*

05 Platforms in the crosshairs: VLOPs, VLOSEs and the triple convergence

For Very Large Online Platforms and Very Large Online Search Engines — the DSA designations for services with 45 million-plus average monthly EU users — the golden thread is not an analytical convenience. It is now the architecture of their supervision. These are the entities on which content regulation, consumer protection and AI safety converge hardest, because a single generative-AI feature on a designated service triggers all three regimes at once.

One feature, three regimes, one regulator

The Digital Omnibus makes this convergence structural. The EU AI Office now holds exclusive supervisory competence over AI systems integrated into VLOPs and VLOSEs — the same Commission body that already directly supervises those platforms' DSA obligations. A GenAI assistant embedded in a designated search

engine, or an AI recommender or image tool inside a designated platform, is simultaneously: a **GPAI-based AI system** under the AI Act (Art. 50 transparency now; the new Art. 5 CSAM/NCII prohibition from December 2026); a **platform feature** whose risks must be assessed and mitigated in the DSA Article 34–35 annual systemic-risk cycle — explicitly including generative-AI risks to minors, civic discourse and fundamental rights; and a **consumer-facing product** subject to the DSA's dark-pattern ban (Art. 25), advertising and recommender transparency (Arts. 26–27, 38–39), and the prohibition on profiling-based ads to minors (Art. 28(2)). What used to be three compliance conversations with three teams is now one conversation with one regulator holding both rulebooks — and inspection powers.

Content, consumer protection, AI safety — the same evidence chain

- **Content:** AI-generated media flows straight into the platform's moderation stack — every deepfake that Art. 50 requires to be labelled is also content the DSA requires to be risk-assessed, actioned with a statement of reasons, and filed to the public transparency database. Synthetic CSEA and NCII sit at the intersection of the new AI Act prohibition, DSA Article 18 notification duties, the UK OSA's priority offences and the US TAKE IT DOWN regime.
- **Consumer protection:** AI features are marketed, subscribed to, and cancelled — which drags in the EU consumer acquis and UCPD, the UK DMCCA (CMA fines up to 10% of global turnover, no court required), and FTC Act §5 with ROSCA for AI premium tiers. Overstated capability claims ("AI-washing") and dark patterns in AI consent or cancellation flows are the enforcement theories of choice on both sides of the Atlantic.
- **AI safety:** where the platform's parent is also the model provider, GPAI duties (Arts. 53–55), systemic-risk model obligations, and serious-incident reporting to the AI Office stack on top — with penalty exposure compounding across regimes: 6% of global turnover under the DSA, 7%/3% under the AI Act, 10% under the UK OSA and DMCCA.

EXHIBIT 5

The triple convergence on designated platforms

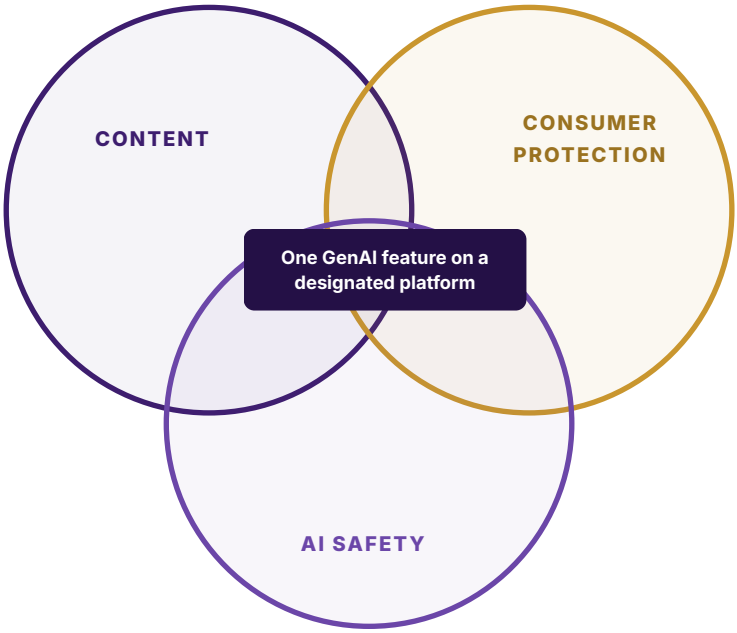

● **Content:** Art. 50 labels feed DSA moderation, statements of reasons, systemic-risk assessment; synthetic CSEA/NCII engages the Art. 5 ban, OSA and TAKE IT DOWN.

● **Consumer:** dark-pattern bans (DSA Art. 25, UCPD), ad transparency, DMCCA and FTC §5/ROSCA on AI marketing, subscriptions and cancellation.

● **AI safety:** GPAI and systemic-risk duties stack where platform and model provider are one undertaking; incidents report to the AI Office.

Supply Unchained analysis. Penalty exposure compounds: 6% (DSA) · 7%/3% (AI Act) · 10% (UK OSA, DMCCA).

> *"For a designated platform, a single generative-AI feature is a content-moderation object, a consumer product and a regulated AI system at the same time — and the evidence file must satisfy all three readings at once."*

The practical consequence: the DSA systemic-risk assessment, the OSA risk-assessment suite and the AI Act's risk-management and FRIA disciplines should be run off one platform-level harm taxonomy, refreshed on product change — because a new GenAI feature legally re-opens the DSA assessment, the OSA children's assessment and the Art. 50 compliance analysis simultaneously. Platforms that operate these as separate annual exercises will produce inconsistent artefacts; the Commission, which now reads all of them, will notice.

06 The FRIA: where the golden thread becomes a legal instrument

If human rights is the measure, the **Fundamental Rights Impact Assessment** is the instrument that operationalises it. Article 27 of the EU AI Act requires certain deployers of high-risk AI systems — public bodies, private entities providing public services, and deployers of AI for creditworthiness and life/health insurance pricing (Annex III points 5(b) and 5(c)) — to conduct a FRIA **before first use** of the system.

A compliant FRIA must set out: the deployer's processes in which the system will be used; the period and frequency of use; the categories of natural persons and groups likely to be affected; the specific risks of

harm to those groups; the human-oversight measures; and the measures to be taken if the risks materialise — including internal governance and complaint arrangements. The results are notified to the market surveillance authority on a template provided by the AI Office. Under the omnibus calendar, the obligation now applies with the rest of the Annex III regime from **2 December 2027**.

EXHIBIT 6

Anatomy of a Fundamental Rights Impact Assessment (Article 27)

Context
The deployer's processes in which the high-risk system will be used

Duration
Period and frequency of intended use

Affected people
Categories of persons and groups likely to be affected

Specific harms
Risks of harm to those groups — severity and likelihood

Human oversight
How humans monitor the system and can intervene

Mitigation & remedy
Measures if risks materialise; governance and complaint arrangements

Who: public bodies, private providers of public services, and deployers of credit-scoring and life/health insurance AI (Annex III 5(b)–(c)). **When:** before first use — results notified to the market surveillance authority on the AI Office template. **From:** 2 December 2027.

Source: Regulation (EU) 2024/1689, Article 27, as amended by the Digital Omnibus.

Read that element list again and you will recognise it: it is human rights due diligence, codified. Affected rights-holders. Specific harms. Severity and likelihood. Prevention, mitigation, remedy. The FRIA is the UNGPs translated into an enforceable pre-deployment gate — and it is the same intellectual structure as the DSA Article 34 assessment, the OSA children's risk assessment, and the CSDDD's chain-of-activities due diligence (in force March 2026, applying from July 2029 for the largest companies). A company that builds one rigorous, harms-based assessment capability can produce all of these artefacts from it. A company that treats each statute as a separate checkbox exercise will build four inconsistent processes and be impeached by their differences.

Why FRIAs matter to AI companies before December 2027

The 16-month deferral is not a reprieve; it is a construction window. Three reasons to act now:

- **The deferral did not move the thinking, only the deadline.** The FRIA's inputs — affected-group mapping, harm taxonomies, oversight design — are the same inputs the DSA and OSA assessments demand *today*, and the same evidence the new CSAM/NCII prohibition demands by **December 2026**: a provider that cannot show it assessed foreseeable misuse of its generative models is, from that date, arguably operating a prohibited system.
- **Deployers will demand it of providers.** Deployers carry Article 26 obligations and FRIA duties regardless of provider compliance status. They can only complete a FRIA with provider documentation — model cards structured to Annex IV/XI, known limitations, logging capabilities, incident protocols. Under the omnibus's strengthened Article 25, failing to share that documentation now carries fines up to 3% of worldwide turnover. Providers who cannot feed their customers' FRIAs will lose enterprise procurement.

- **It is the audit defence.** As our June investigation, *The Accountability Gap*, found when we applied an article-by-article EU AI Act risk-and-control matrix to the four leading model providers: not one would pass a compliance audit today, and a control that exists but is not documented provides no legal defence. The FRIA discipline — documented, dated, evidence-backed assessment before deployment — is precisely what converts "we take safety seriously" into something a regulator will accept.

07 The Supply Unchained HRDD Framework, applied across TMT

A Fundamental Rights Impact Assessment is only as good as the due-diligence discipline behind it. This is the framework we use — and what it found when we applied it to the technology, media and telecommunications companies whose AI now shapes most digital lives.

The Supply Unchained Human Rights Due Diligence Framework transfers two decades of supply-chain HRDD practice to AI and digital operations. For each right at risk — or each obligation — it assesses five dimensions: the **evidence** (what is actually documented), the **gap** against a gold-star standard (what a regulator would expect of a mature operator), the **inherent risk** (likelihood × severity, on people, before controls), the **residual risk** (after controls, honestly weighted for their effectiveness), and a **compliance score** on a 0–100 scale that makes progress measurable. Severity is scored on the UNGP criteria — *scale, scope, irremediability* — on a five-level scale, which is what makes a child-safety risk (irremediable, by definition) rank where it belongs: at the top, always. Scores aggregate to four grading bands: **Strong (80+)**, **Partial (60–79)**, **Limited (35–59)**, **Non-Compliant (below 35)**. And the framework scores published, documented evidence only — because a control that exists but is not documented provides no legal defence and no public accountability, and is treated by auditors, and courts, as absent.

EXHIBIT 7

The Supply Unchained HRDD Framework — five dimensions, UNGP severity

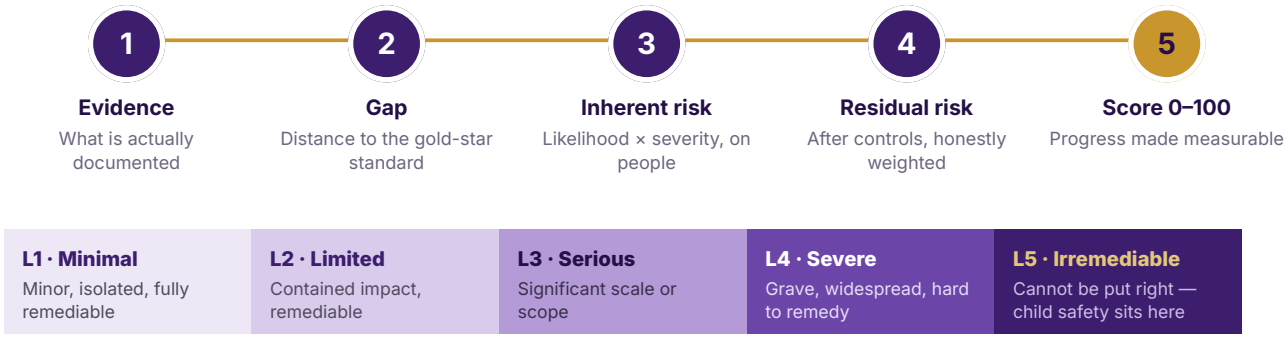

Severity assessed on the UNGP criteria — scale (how grave) · scope (how many people) · irremediability (can it be put right). Grading bands: Strong 80+ · Partial 60–79 · Limited 35–59 · Non-Compliant <35.

Four majors, one yardstick — the rights lens

In June 2026 we applied this framework, article by article, to the four companies whose models power most of the world's enterprise AI — our *Accountability Gap* investigation, which produced overall compliance scores of 48, 43, 37 and 29 out of 100. Exhibits 8 and 9 show those original scores — and what happens when the new Article 5 prohibition on generating CSAM and NCII is folded into the prohibited-practices control area, with its demand for evidenced, adequate safeguards against generation, grooming and exploitation misuse. Every score falls, and the open-weights provider falls furthest, because the prohibition

turns on exactly what an open distribution model cannot evidence: post-release control. Anonymised as in the original assessment.

EXHIBIT 8

The Accountability Gap scores — original, and re-scored under the new Art. 5 prohibition

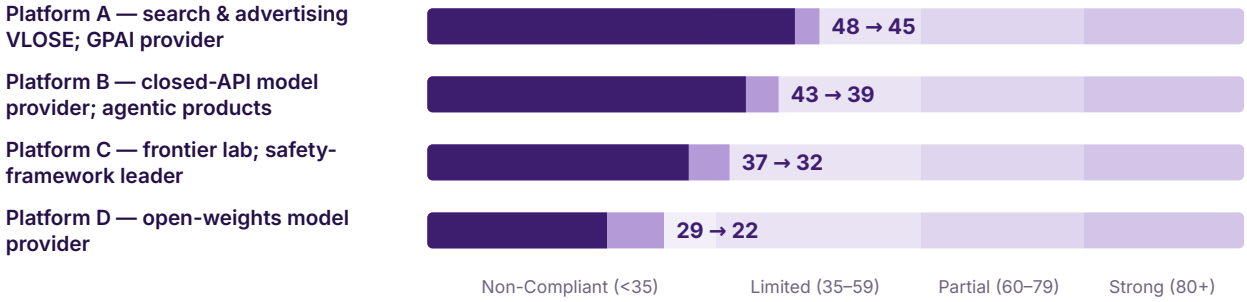

Light bar: original overall score, The Accountability Gap, June 2026. Dark bar: Supply Unchained July 2026 re-scoring, folding the new Art. 5 CSAM/NCII prohibition (evidenced safeguards against generation and exploitation misuse) into the prohibited-practices control area. Published evidence only.

EXHIBIT 9

Four major TMT platforms — original Accountability Gap scores and Art. 5 re-scoring

ENTITY (ANONYMISED)	SCORE (JUN → JUL)	GRADING	STRONGEST RIGHTS CONTROL	CRITICAL HRDD GAP
Platform A — global search & advertising VLOSE; GPAI provider	**48 → 45**	Limited (upper)	Adversarial testing (350+ red-team exercises/yr) and content provenance across all modalities	Training-data governance: no published demographic-bias audit or provenance documentation — the mechanism by which bias enters the system
Platform B — closed-API model provider; agentic products	**43 → 39**	Limited	GPAI Code of Practice signature — binding transparency, copyright and safety undertakings	Commitments not yet implemented; agentic product with no published human-oversight (Art. 14) or affected-group analysis
Platform C — frontier lab; most sophisticated safety framework	**37 → 32**	Non-Compliant on re-score	Catastrophic-risk governance exceeding Art. 55 systemic-risk demands	Framework addresses civilisational risk, not person-level rights harms — no FRIA-style mapping of who is affected and how they are remedied
Platform D — open-weights model provider	**29 → 22**	Non-Compliant	Best-in-class deployer-facing responsible-use guidance and safety tooling	The HRDD chain — identify, prevent, mitigate, account — breaks at weight release: no post-distribution visibility, control or remedy

Source: Supply Unchained HRDD Framework scoring of rights-relevant control areas, derived from the EU AI Act Risk and Control Matrix assessment, The Accountability Gap, June 2026. Published evidence only; anonymisation as in the original assessment.

What the Accountability Gap tells us about HRDD

The summary finding of that investigation bears directly on this article's thesis. None of the four would pass an EU AI Act compliance audit today: none had registered in the EU database, none had published an Art. 53-compliant training-data summary or copyright opt-out mechanism, and none had operationalised the Art. 73 incident-reporting timelines. But the deeper finding was about human rights due diligence. The universal gap — shared by all four, whatever their safety sophistication — is training-data governance: no company publishes how its training datasets were audited for demographic representation, provenance-tracked, or rights-cleared. An AI system trained on biased data produces biased outputs; those outputs, deployed in hiring, lending, insurance pricing or supplier screening, are precisely the adverse human-rights impacts the CSDDD, the EU Charter and the FRIA exist to prevent. That gap is not a technical compliance detail. It is the point where the industry's safety promises and its rights obligations come apart — and where the OECD's May 2026 guidance now demands the same rigour applied to a Tier 2 supplier audit for forced labour.

> *"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 companies that close these gaps first will collect what we called the trust dividend: the competitive advantage of being the provider that enterprise customers can rely on, regulators can engage with, and the public can hold to account. Trust in AI is not given. It is earned through documentation, accountability, and the willingness to be audited — and the FRIA is the instrument through which that earning process is about to become mandatory.

08 Who regulates what — and how

The golden thread has keepers. Knowing which regulator holds which end of it — and what powers they bring — is the difference between a compliance strategy and a compliance surprise.

Dispatches from Geneva: AI for Good 2026

As this article went to press, the UN's **AI for Good Global Summit** (ITU, with 53 UN agencies) was concluding in Geneva, 7–10 July. Three takeaways matter for this article's thesis. First, **standards are becoming the connective tissue between principle and enforcement**: the ITU's AI Standards Exchange and the harmonised-standards programme are where abstract duties — "adequate safeguards" against CSAM generation, machine-readable content marking, high-risk risk management — get their testable technical meaning; recall that the omnibus explicitly tied the Annex III start date to standards availability. Second, **the global conversation has shifted from principles to plumbing**: the Geneva agenda was dominated not by ethics declarations but by governance architecture — who inspects, who audits, who reports to whom — the same question this section answers for the hard-law regimes. Third, **the multilateral layer supplies the shared vocabulary**: the UN system frames AI governance in the language of human rights and the SDGs, which is precisely why regimes drafted in Brussels, London and Washington converge on the same measure of harm. Geneva writes the dictionary; the regulators below enforce the grammar.

EXHIBIT 10

The enforcement map — who regulates what, and with which teeth

REGULATOR (SEAT)	WHAT IT REGULATES	UNDER	HOW IT ENFORCES
European Commission (Brussels)	VLOPs and VLOSEs directly — systemic risk, minors protection, transparency	DSA	Art. 67 RFIs, formal proceedings, interim measures, fines to 6% of global turnover; enhanced supervision up to suspension
EU AI Office (Brussels, within the Commission)	GPAI models; post-omnibus, exclusive competence over AI systems built on a provider's own GPAI and AI integrated into VLOPs/VLOSEs	AI Act	Compels documentation, runs model evaluations, orders mitigation, **inspects**, can order market withdrawal; fines to 7%/€35M (prohibitions), 3%/€15M (duties)
Coimisiún na Meán (Dublin)	Ireland's Digital Services Coordinator — first-line EU regulator for the many TMT platforms established in Ireland; national online-safety code	DSA; Irish OSMR Act	Certifies trusted flaggers and out-of-court bodies, handles complaints, audits non-VLOP duties, coordinates with and refers to the Commission
National market surveillance authorities (27 capitals)	High-risk AI systems from Dec 2027; recipients of Art. 73 serious-incident reports and FRIA notifications	AI Act	Market surveillance powers; corrective orders; national penalty regimes within the Act's ceilings
Ofcom (London)	User-to-user and search services with UK links; illegal content, children's safety, categorised-service duties	UK OSA	Codes of practice, statutory information notices (criminal exposure for wrong answers), fines to £18M/10% of global revenue, business-disruption orders up to ISP blocking
Washington: FTC, State AGs, DOJ	Deception and unfairness in AI claims and design; children's privacy; NCII removal; state AI acts (CO, TX, CA); preemption litigation	FTC Act §5, COPPA, ROSCA, TAKE IT DOWN; state UDAP/AI statutes	6(b) studies and CIDs (certified answers), consent decrees with 20-year audit tails, per-violation civil penalties; State AG multistate coalitions; DOJ AI Litigation Task Force contests state laws — child safety carved out
eSafety Commissioner (Canberra)	Online safety expectations; social-media minimum age; removal notices	AU OSA, SMMA	Compulsory transparency notices, 24-hour removal orders, published non-compliance findings, penalties to AUD ~$49.5M
Multilateral layer (Geneva · Strasbourg · Paris)	The shared vocabulary and technical substrate: UN/ITU AI for Good and standards exchange; Council of Europe Framework Convention on AI; OECD AI due-diligence guidance	Soft law; CoE treaty	No fines — but standards define what "adequate" means in every hard-law regime, and the human-rights framing sets the measure the enforcers apply

Status as at 9 July 2026. Supply Unchained analysis. "Teeth" summarised; see the practical guide for deadline-sequenced actions.

Two features of this map deserve emphasis. First, **centralisation in Brussels**: for a TMT major whose platforms are designated and whose models are self-supplied, the Commission — wearing its DSA and AI Office hats — is now the single most consequential regulator on earth, holding content, AI safety and (with the consumer acquis) fairness files simultaneously. Dublin remains the first port of call for establishment-based supervision; London holds the sharpest criminal-exposure tool; Washington enforces through evidence demands and consent decrees rather than codes. Second, **the layers interlock**: a Geneva standard becomes a Brussels presumption of conformity, a London code measure, a Washington substantiation

benchmark. The company that builds one evidence chain can answer all of them; the company that answers each regulator separately will eventually contradict itself — and cross-regulator consistency is itself now a compliance discipline.

09 The practical guide: what AI companies must do next

What follows is the action checklist we would put in front of any AI provider or deployer board this quarter. It is sequenced by deadline, not by jurisdiction size.

EXHIBIT 11

The dates that now govern your roadmap

DEADLINE	JURISDICTION	OBLIGATION
Now / continuous	EU · UK · US	GPAI duties (live since Aug 2025); OSA illegal-content & children's duties (live); NCMEC CyberTipline reports; TAKE IT DOWN 48-hour NCII removal; COPPA amended Rule; DSA duties for designated platforms
2 Aug 2026	EU · US (CA)	AI Act Art. 50 transparency (AI disclosure, deepfake labelling) + California SB 942 provenance duties — one watermarking build serves both
2 Dec 2026	EU	NCII/CSAM generation prohibition fully applies; end of Art. 50(2) marking grace period for pre-existing systems
2 Dec 2027	EU	Annex III high-risk regime: risk management, data governance, conformity assessment, registration — and the Art. 27 FRIA
2 Aug 2028	EU	Annex I high-risk duties (AI in regulated products)

Source: Regulation (EU) 2024/1689 as amended by the Digital Omnibus; US federal and state statutes as at 9 July 2026. Supply Unchained analysis.

EUROPEAN UNION — THE COMPLIANCE ENGINE ROOM

- **Stand up Article 50 transparency before 2 August 2026.** AI-interaction disclosure on every conversational surface; machine-readable marking of synthetic audio, image, video, and text; deepfake-labelling flows. Verify the December 2026 grace period actually applies to each system before relying on it.
- **Treat the CSAM/NCII prohibition as a model-safety engineering deadline (2 December 2026).** Red-team every generative model for the "reasonably foreseeable and reproducible" test; implement and document input/output safeguards, abuse monitoring, and takedown loops; keep the evidence file — the prohibition turns on whether safeguards are adequate and demonstrable.
- **Close the GPAI gaps that are already in force.** Model Documentation Form kept current per release; a published copyright policy honouring TDM opt-outs; the public training-content summary on the Commission template. These have been legally required since August 2025, and the AI Office — now holding inspection powers — begins enforcement monitoring in August 2026.
- **Build the FRIA capability against December 2027 now.** Inventory every AI system against Annex III; for in-scope deployments, pilot the FRIA template (affected groups, specific harms, oversight, mitigation, remedy); wire it to the same harm taxonomy as your DSA/OSA assessments. Providers: prepare the Article 25 documentation pack your deployers' FRIAs will require.

- **If you run a VLOP or VLOSE:** run one platform-level harm taxonomy across the DSA Art. 34 cycle, the AI Act analysis and OSA assessments; cover generative-AI features explicitly; enforce minors protections under Article 28 and the 2025 guidelines; audit AI subscription, consent and cancellation flows against dark-pattern rules (DSA Art. 25, UCPD, DMCCA, ROSCA); and remember a new GenAI feature re-opens the assessments.

UNITED KINGDOM — EVIDENCE THAT YOUR PROTECTIONS WORK

- **Keep the OSA risk-assessment suite current and product-triggered.** A GenAI feature that can surface CSEA or primary-priority content to children re-opens the illegal-content and children's risk assessments. Document the re-assessment, not just the feature launch.
- **Implement highly effective age assurance where required — and prove effectiveness.** Ofcom and the ICO's March 2026 joint statement makes clear that deployed-but-ineffective age assurance is non-compliance. Measure circumvention and over-blocking; document the proportionality analysis where you decide age assurance is not required.
- **Map every Ofcom code measure to an implemented control** or a documented alternative-measures justification — and operate an information-notice response protocol with legal sign-off, because wrong answers to Ofcom carry criminal exposure for senior managers.

UNITED STATES — THE PATCHWORK, NAVIGATED

- **Operationalise the two hard-edged federal duties:** NCMEC CyberTipline reporting (with preservation) for any service that may surface CSAM, and the TAKE IT DOWN 48-hour NCII notice-and-removal process, FTC-enforced since May 2026.
- **Comply with operative state AI law as written — do not bank on preemption.** Colorado's AI Act (in force 30 June 2026: impact assessments, algorithmic-discrimination disclosure to the AG within 90 days), Texas TRAIGA and California SB 53/SB 243/AB 2013 (all in force since January 2026) are binding unless and until courts say otherwise. The preemption Executive Order explicitly carves out child safety.
- **Run a claims-substantiation gate.** Every public statement — "safe for teens," "industry-leading safety" — needs pre-existing evaluation evidence behind it. FTC deception cases are built from the delta between incidents and marketing; the September 2025 6(b) study into AI companion chatbots shows exactly what will be demanded: safety testing records, minor protections, disclosures.
- **For minors:** COPPA amended-Rule consent flows (separate opt-in for third-party disclosure), written retention policies, app-store age-signal readiness, and — for any conversational AI reachable by minors — default-safe configurations, crisis-referral protocols, and documented pre-launch safety evaluation.

CROSS-CUTTING — BUILD ONCE, EVIDENCE EVERYWHERE

- **One obligations register:** every instrument decomposed to article-level obligations, mapped to owner, control, evidence location, and effective date.
- **One harms-based risk engine** serving the FRIA, DSA Art. 34, OSA assessments, AI Act risk management, and Colorado impact assessments — UNGP severity criteria at the core, jurisdiction-specific outputs at the edge.

- **One incident process** routing to the right regulator clock: AI Office (systemic-risk models), market-surveillance authorities (Art. 73: 15 days, compressed to 2–10 for the gravest cases), NCMEC, Ofcom, state AGs — with a claims-review loop so public statements are re-checked after every incident.
- **One evidence repository per model and product.** The AI Office's new inspection power converts documentation quality into the front line of defence. If it is not written down, it did not happen.

10 Conclusion: prove it

The Digital Omnibus bought the industry time on the high-risk tier — sixteen months on Annex III, a year on Annex I. It bought no time at all on the things regulators care about most: transparency about what is synthetic, accountability for general-purpose models, and the absolute protection of children from sexual exploitation and abuse. The new Article 5 prohibition, the OSA's children's codes, the DSA's minors provisions, TAKE IT DOWN, COPPA, the KIDS Act pipeline — these are not separate compliance projects. They are one project, measured on one instrument: harm to human beings, assessed before deployment, documented well enough to survive an audit.

That is what the Fundamental Rights Impact Assessment is for. Not a form to be filed in December 2027, but the discipline — inherited from two decades of human rights due diligence in physical supply chains — of asking, before a system touches a person: *who could this harm, how badly, how many, and can it be undone?* Companies that answer that question with evidence will find that every regulator on the thread — Brussels, London, Washington, Canberra — is asking for a version of the same document. Companies that cannot will discover, as our Accountability Gap assessment showed, that sophisticated safety intentions without documented controls provide no legal defence and no public accountability.

> *"Every regime converges on the same demand: prove it. Human rights is the golden thread; the FRIA is where you show your stitching."*

METHODOLOGY & SOURCES

This article applies the Supply Unchained Human Rights Due Diligence Framework — an article-by-article assessment across 25 control areas scoring evidence, gap, inherent risk (likelihood × severity), residual risk, and compliance (0–100, graded Strong/Partial/Limited/Non-Compliant), derived from the DOJ Evaluation of Corporate Compliance Programs and the OECD Due Diligence Guidance for Responsible Business Conduct, with severity assessed on the UNGP-aligned five-level scale (scale, scope, irremediability). Exhibits 8–9 show the original overall scores from our EU AI Act Risk and Control Matrix assessment of four major TMT platforms (The Accountability Gap, June 2026) and a July 2026 Supply Unchained re-scoring that folds the new Article 5 CSAM/NCII prohibition — including safeguards against exploitation misuse — into the prohibited-practices control area; anonymised as in the original; published evidence only.

Regulatory positions are stated as at 9 July 2026 and include: Regulation (EU) 2024/1689 as amended by the Digital Omnibus on AI (European Parliament endorsement 16 June 2026; final Council approval 29 June 2026); the EU AI Office GPAI Code of Practice (July 2025); Regulation (EU) 2022/2065 (DSA) and the Commission's 2025 guidelines on the protection of minors; the UK Online Safety Act 2023 and Ofcom's Protection of Children codes and 2026 updates, including the Ofcom–ICO joint statement on age assurance (March 2026); the US TAKE IT DOWN Act, COPPA Rule as amended (effective 23 June 2025), 18 U.S.C. §2258A, and the KIDS Act package as passed by the House (29 June 2026); the OECD Due Diligence Guidance for Responsible AI (May 2026); and the CSDDD as amended by Omnibus I (in force March 2026). This is not legal advice. Organisations with specific compliance questions should consult qualified counsel in the relevant jurisdiction.

Supply Unchained

AI GOVERNANCE · HUMAN RIGHTS DUE DILIGENCE · SUPPLY CHAIN RISK

supplyunchained.co.uk

For engagements and enquiries: nayantara.sriram@supplyunchained.co.uk

JULY 2026