



Volume 23, Issue 01, June 2026

General-Purpose AI under the EU AI Act: A Conceptual Allocation of Duties across the Value Chain

*Dr. iur. Dr. rer. pol. Fabian Teichmann, LL.M. (London), EMBA (Oxford)**



© 2026, Fabian Teichmann

Licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license

DOI: 10.2218/scrip.12300

Abstract

This article examines how the final version of the EU Artificial Intelligence Act (“AI Act”, adopted 2024) allocates obligations across the AI value chain, with a focus on general-purpose AI (“GPAI”) or foundation models. It proposes a taxonomy of key actors – foundation model providers, fine-tuners, integrators, and deployers – and analyses the interfaces between them, including documentation tools (model cards, system cards) and logging requirements. Building on principles of control, foreseeability, benefit, and capability, the article argues for a principled distribution of compliance duties: those who design and train foundational AI models should bear upstream transparency and safety obligations, while those who adapt or deploy AI in specific contexts shoulder downstream risk

management and oversight duties. The analysis discusses the EU Act's enforcement model (market surveillance authorities and an EU AI Office) and contrasts it with data protection supervision, highlighting the need for coordination in areas of overlap. A comparative outlook considers the UK's pro-innovation, context-driven approach and the evolving US regulatory landscape. Finally, the article offers recommendations for regulatory guidance and standards development (ISO/IEC, CEN/CENELEC) to support effective implementation, accounting for recent developments including the Digital Omnibus Regulation, the AI Office's emerging guidelines, and the broader debate on EU regulatory competitiveness triggered by the Draghi Report. This approach aims to clarify how the EU AI Act's final provisions on general-purpose AI models can serve as a global benchmark for balanced AI governance.

Keywords

EU AI Act; General-Purpose AI; Value Chain Allocation; Foundation Models; Regulatory Compliance

* Fabian Teichmann, PhD (Zurich), LL.M. (London), EMBA (Oxford) is an attorney-at-law and managing partner at Teichmann International (Switzerland) AG. He holds a graduate degree in Information Management Systems from Harvard University. teichmann@teichmann-law.ch

1 Introduction

The advent of general-purpose AI systems – exemplified by large language models like GPT-4 – has prompted regulators to rethink how AI is governed across its value chain. Unlike narrowly-purposed AI, these foundation models can be adapted to myriad downstream tasks, from writing code and offering medical advice to generating images and driving autonomous vehicles. Their versatility and widespread deployment raise complex questions of accountability: Who should ensure a powerful model is trained responsibly, and who is answerable if its use causes harm? The European Union’s answer, in the newly adopted EU AI Act,¹ is to impose tiered obligations on different actors in the AI life cycle. This marks a shift from earlier drafts that focused only on end-use applications – the final Act squarely regulates general-purpose AI model providers (the upstream developers of foundation models) in addition to the providers and users of high-risk AI systems. The Act thereby inaugurates a hybrid governance approach: combining product safety-style market regulation with elements addressing fundamental rights, and extending compliance duties both upstream (at the model level) and downstream (at the application and deployment level).

The rapidly growing body of scholarship on the AI Act has begun to address the allocation of responsibilities across complex AI supply chains as one of the regulation’s central challenges. The first article-by-article commentaries – notably the volume edited by Necati Pehlivan, Forgo and Valcke (2024) – provide granular doctrinal analysis of each provision, while Hacker (2024) has offered an influential account of how the Act distributes regulatory burdens between

¹ Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L 1689/1 (“AI Act”).

providers and deployers, identifying residual ambiguities around fine-tuning thresholds and downstream liability.² Veale and Zuiderveen Borgesius (2021) laid early groundwork by demystifying the Commission's original proposal and critiquing its product-safety orientation, work that Laux, Wachter and Mittelstadt (2024) extended by examining whether the Act's risk classification genuinely delivers on its promise of "trustworthy" AI.³ Poncibó (2024) and Noto La Diega and Bezerra (2024) have interrogated the consumer protection and liability dimensions respectively, while Gstrein, Haleem and Zwitter (2024) focused specifically on how the Act's final text handles foundation models – an analysis on which the present article builds and to which it responds.⁴ The present contribution differs from this existing literature by developing an integrated conceptual framework grounded in four allocation principles (control, foreseeability, benefit, and capability), by systematically mapping these principles onto the Act's specific provisions, and by offering concrete recommendations for standards development and regulatory guidance that take account of the most recent political and regulatory developments.

² Ceyhun Necati Pehlivan, Nikolaus Forgo and Peggy Valcke (eds), *The EU Artificial Intelligence (AI) Act: A Commentary* (Wolters Kluwer, 2024); Hacker P, "Comments on the Final Trilogue Version of the AI Act" (2024), available at SSRN: <https://ssrn.com/abstract=4757603>, 2.

³ Johann Laux, Sandra Wachter and Brent Mittelstadt, "Trustworthy Artificial Intelligence and the European Union AI Act: On the Conflation of Trustworthiness and Acceptability of Risk" (2024) 18(1) *Regulation and Governance* 3-32, 6; Michael Veale and Frederik Zuiderveen Borgesius, "Demystifying the Draft EU Artificial Intelligence Act – Analysing the Good, the Bad, and the Unclear Elements of the Proposed Approach" (2021) 22(4) *Computer Law Review International* 97-112, 112.

⁴ Oskar Josef Gstrein, Noman Haleem, and Andrej Zwitter, "General-Purpose AI Regulation and the European Union AI Act" (2024) 13(3) *Internet Policy Review* 1-26, 18-19; Guido Noto La Diega and Leonardo CT Bezerra, "Can There Be Responsible AI Without AI Liability? Incentivizing Generative AI Safety Through Ex-Post Tort Liability Under the EU AI Liability Directive" (2024) 32(1) *International Journal of Law and Information Technology* 1-21, 2-3; Larry DiMatteo, Cristina Poncibó and Geraint Howells, *The Cambridge Handbook of AI and Consumer Law* (Cambridge University Press, 2024), pp. 116-120.

This article proceeds as follows. Section 2 delineates a taxonomy of the relevant roles – foundation model providers, fine-tuners, integrators, and deployers – and the key interfaces between these actors, including documentation and logging mechanisms. Section 3 introduces the four allocation principles and examines how each informs the distribution of obligations between upstream and downstream actors. Section 4 analyses how liability and enforcement are structured under the AI Act’s market surveillance model and contrasts this with the General Data Protection Regulation’s (“GDPR’s”) data protection supervision regime,⁵ highlighting coordination challenges. Section 5 broadens the perspective with a cross-jurisdictional analysis of the UK and US approaches, updated to reflect recent policy shifts. Section 6 presents recommendations for regulatory guidance and standardisation, taking account of the Digital Omnibus Regulation,⁶ the AI Office’s emerging guidelines,⁷ and the wider debate on EU regulatory competitiveness. Section 7 concludes.

2 Conceptual Taxonomy of Actors and Interfaces

2.1 Actors in the AI Value Chain

Modern AI systems involve a multi-tiered supply chain. In the context of general-

⁵ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) [2016] OJ L 119/1.

⁶ European Commission, *Proposal for a Regulation of the European Parliament and of the Council amending Regulations (EU) 2024/1689 and (EU) 2018/1139 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI)* COM(2025) 836 final (19 November 2025) <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025PC0836> accessed 17 June 2026.

⁷ European Commission AI Office, *Guidelines on the Definition of an Artificial Intelligence System Established by Regulation (EU) 2024/1689 (AI Act)* (6 February 2025) <<https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-ai-system-definition-facilitate-first-ai-acts-rules-application>> accessed 17 June 2026.

purpose AI (“GPAI”), four main categories of actors can be identified, each bearing distinct obligations under the AI Act.

Foundation Model Providers are the organisations or individuals that develop large-scale GPAI models and place them on the market. Under the Act’s definitions (Article 3), a provider of a GPAI model is a person or entity that develops a general-purpose AI model (or has it developed) and makes it available, whether for a fee or free. Examples include AI labs releasing models like GPT-4, Gemini, or open-source large models such as Llama. These actors are subject to Chapter V obligations specific to GPAI model providers, including documenting how the model was developed and tested, disclosing information about training data and model limitations, and implementing appropriate data governance measures (Article 53, Annexes XI-XII). The Act recognises that open-source models warrant differentiated treatment: if a model is released with full access to weights and sufficient documentation, certain documentation obligations may be adjusted (Recital 102). As Hacker (2024) has noted, the provision remains ambiguous on what degree of openness triggers this exception, and the AI Office will need to clarify this boundary through guidance.⁸ The stakes are high, since an excessively broad exception could create a regulatory gap for powerful models released under permissive licenses, while an excessively narrow one could deter beneficial open-source development.

Fine-Tuners (Downstream Model Providers) adapt a pre-existing foundation model to create a new model, often for a more specific function or domain. Fine-tuners take a base model and train it further on domain-specific data or with additional objectives. The Act (Recital 97) acknowledges that GPAI models can be “further modified or fine-tuned into new models,” and that

⁸ Hacker (n 2), 6–7.

downstream entities doing so may themselves become providers with corresponding obligations. Recital 109 attempts to calibrate the burden: a fine-tuner's documentation duties should focus on the changes made, relying on the original model's documentation for the remainder. As Laux, Wachter and Mittelstadt (2024) observe, the threshold for when modification converts a fine-tuner into a "provider" of a new model remains one of the Act's most significant grey areas, with practical implications for the open-source ecosystem and for commercial fine-tuning services.⁹ A start-up fine-tuning an open-source language model for legal advice, or a company adapting a general model for medical diagnostics, needs to know whether it has crossed the line into full provider obligations, including conformity assessment for a potentially high-risk system. Further guidance on this threshold – perhaps involving criteria such as the degree to which model behaviour is altered or whether it is repurposed for a high-risk use – is urgently needed.

Integrators (AI System Providers) embed an AI model into a final application or system, potentially combining it with other software or hardware, and determine the system's intended use. A medical device company that takes a pre-trained diagnostic model and embeds it in a clinical decision support tool is an integrator; a software firm that plugs a foundation model into an HR screening system is another. Under the Act, such an entity is the provider of an AI system and, if the system is high-risk under Annex III, it is subject to requirements including conformity assessment, risk management (Article 9), technical documentation (Annex IV), transparency, accuracy, and human oversight measures. The Act requires cooperation between the original GPAI model provider and the integrator: the provider must supply the information

⁹ Laux et al (n 3), 26–27.

listed in Annex XII to enable the integrator’s compliance (Article 53(1)(b)). As Veale and Zuiderveen Borgesius (2021) anticipated, this creates an information pipeline that depends heavily on the quality and completeness of upstream documentation. In practice, an integrator might receive a model card or technical dossier from the foundation model provider, use it to perform a risk assessment, and then produce system-level documentation covering how the model is used and controlled within the integrated application.¹⁰

Deployers are the natural or legal persons that use an AI system under their authority in a professional capacity (Article 3(4)). A hospital deploying an AI diagnostic tool, a bank using an AI system for credit scoring, or a public agency employing an AI system for automated decision-making are all deployers. Under the Act, deployers of high-risk AI systems bear their own set of obligations (Article 26): ensuring proper human oversight, providing training to personnel, monitoring the system’s performance, maintaining automatically generated logs for at least six months, respecting transparency towards affected individuals, and performing fundamental rights impact assessments in certain use cases. Deployers must also use the AI system in accordance with the provider’s instructions and must suspend use if the system is found non-compliant or poses a risk.

These roles can overlap in a single entity or be distributed across many. A large tech company might simultaneously serve as foundation model provider, integrator, and deployer. Conversely, in the open-source ecosystem, distinct independent actors may occupy each layer. The Act uses legal hooks like “placing on the market” and “putting into service” to attach obligations to

¹⁰ Veale and Zuiderveen Borgesius (n 3), 103.

whichever entity effectively controls the AI's introduction into the EU market or usage environment.

2.2 Interfaces and Information Flows

For this chain to function safely, several interfaces – points of exchange of information and control – are critical.

Technical Documentation and Model Cards. The Act requires GPAI model providers to prepare technical documentation containing detailed information about the model's design, testing, and limitations (Annex XI), and to provide transparency information to downstream providers integrating the model (Annex XII, Article 53(1)(b)). This requirement aligns closely with the concept of Model Cards originally proposed by Mitchell et al. (2019) as a transparency measure for AI models.¹¹ A standardised model card describes the model's architecture, training data characteristics, intended uses, performance benchmarks (including across subgroups), known risks and failure modes, and guidelines for safe use. By standardising this interface, the Act ensures that integrators receive the information needed to perform their own risk management and compliance tasks.

System Documentation. When an integrator builds a high-risk AI system using a GPAI model, it must produce instructions for use and technical documentation (Annex IV) that describe the specific context of use, how the model is integrated, system-level performance metrics, human oversight provisions, and constraints on deployment. There is a clear through-line: the model card feeds into the system documentation. Where gaps remain, the Act's

¹¹ Margaret Mitchell et al., "Model Cards for Model Reporting" (2019) *Proceedings of the Conference on Fairness, Accountability, and Transparency*, [https://doi-org.eux.idm.oclc.org/10.1145/3287560.3287596](https://doi.org.eux.idm.oclc.org/10.1145/3287560.3287596), 220–222.

cooperation requirement enables integrators to request further information from upstream providers. The emerging industry practice of “system cards” – exemplified by OpenAI’s system card for GPT-4 – provides a model for how this information can be structured and communicated to deployers, regulators, and affected persons.

Logging and Record-Keeping. The Act mandates that high-risk AI systems be designed with automatic logging capabilities (Article 12), and deployers must maintain those logs for at least six months to enable tracing of decisions and investigation of incidents. For GPAI models designated as having systemic risk (Article 51), providers must implement post-market monitoring and incident reporting, which implies their own collection of telemetry data. Logging is therefore a two-way interface: it helps deployers supervise AI in operation and provides feedback upstream to model creators and regulators. As commentators have rightly flagged, the Act’s logging requirements raise significant data protection questions, particularly regarding the storage of personal data in audit trails – a tension that requires coordinated guidance from AI and data protection authorities (see generally Noto La Diega and Bezerra 2024 on the interplay between AI regulation and liability frameworks).¹²

Contractual and Policy Interfaces. The Act effectively mandates certain contractual terms: GPAI providers must have a policy on compliance with EU data laws in training data (Article 53(1)(c)), and deployers must use AI systems in accordance with the instructions of use provided by the provider (Article 26(1)). These provisions ensure that information and constraints flow downstream in writing: the provider communicates intended use and limitations, and the deployer has a duty to respect them. Any integrator in

¹² Noto La Diega and Bezerra (n 4), 2–3, 7–8.

between must pass these on and potentially add further instructions specific to the integrated system. These contractual interfaces transform the regulatory obligation into a practical governance mechanism operating along the entire supply chain.

3 Allocation Principles: Control, Foreseeability, Benefit, and Capability

Allocating legal duties across different actors in the AI lifecycle should ideally follow principled criteria. The EU AI Act and broader regulatory-theoretical literature on risk governance reflect four interrelated principles – control, foreseeability, benefit, and capability. This section examines how each principle informs the distribution of obligations and applies them to the roles identified in Section 2.

3.1 Control

An overarching tenet is that the more control an actor exercises over the design or functioning of the AI, the greater their responsibility for preventing harm. In product safety law, manufacturers bear primary duties because they control product design. AI model developers exercise substantial control over model architecture, training data curation, and built-in safety mechanisms; accordingly, the Act requires them to implement appropriate data governance and risk mitigation measures before release (Article 53). They also control the upfront allocation of compute resources and know the scale of the model, which the Act leverages in identifying “systemic risk” models (e.g., providers know whether they exceed the 10^{25} FLOP training threshold that triggers enhanced obligations under Article 51).

Downstream, deployers control the operational context – who uses the AI, for what tasks, and whether there are human checks – and are obligated to ensure proper use and the ability to intervene (Article 26). They can control input data quality in real-time and must do so (Article 26 requires ensuring input data relevance and integrity for high-risk AI).

The Act's treatment of open-source models exemplifies the alignment of responsibility with control: where model weights are freely released, certain provider obligations are relaxed because direct control over downstream use is relinquished (Recital 102). The rationale is that once control is lost – the model is “in the wild” – holding the original provider to some obligations is impractical; instead, the burden shifts to whoever takes that model and deploys it. As Hacker (2024) argues, the control principle also explains the escalating obligations for systemic-risk providers: the greater the model's capabilities and potential reach, the more consequential the provider's design choices become, justifying heightened regulatory scrutiny.¹³

Human oversight requirements similarly track control. Whichever actor can exert real-time intervention – typically the deployer or the person operating the system – must ensure its effectiveness (Article 14). The foundation model creator cannot directly ensure a human is in the loop during deployment; that responsibility lies with the deployer, who controls the operational environment.

3.2 Foreseeability

The law assigns duties to the party that can reasonably foresee and prevent potential harms. Upstream providers can foresee certain generic risks of their models: a large language model's propensity for generating disinformation or

¹³ Hacker (n 2), 12–13.

toxic content, or its potential adaptation for hazardous purposes such as generating malware or lowering barriers to chemical weapon design. Article 53 accordingly requires GPAI providers to document “foreseeable unintended outcomes,” and providers of systemic-risk models must “assess and mitigate systemic risks, including novel risks” (Article 55), implying pre-release red-teaming and safety testing to discover foreseeable failure modes.

Downstream actors have a better vantage point for context-specific risks. An integrator or deployer knows the specific application and user population. An HR deployer should foresee the risk that an AI system could discriminate in hiring and must check for relevant biases. A hospital deploying an AI diagnostic tool can foresee that doctors might over-rely on it, necessitating training to avoid automation bias. Article 26 and the fundamental rights impact assessment provisions compel deployers to formally evaluate impacts in their particular setting.

The concept of “reasonably foreseeable misuse” (Article 3(24)) requires providers to consider not only intended use but also how the system could be misused through human behaviour or interaction with other systems. As Laux, Wachter and Mittelstadt (2024) note, this foreseeability mandate places a significant burden on upstream developers to anticipate downstream applications, even where these are inherently difficult to predict.¹⁴ Yet the principle also sets limits: the Act does not classify AI as high-risk in the abstract but by specific use (Annex III), recognising that a foundation model is not “high-risk” until it is deployed in a high-risk domain – at which point the foreseeability of high impacts activates for the integrator or deployer. In summary, foreseeability justifies requiring each actor to conduct risk assessments

¹⁴ Laux et al (n 3), 6, 26–27.

appropriate to their stage: model providers conduct broad hazard analysis (bias, cybersecurity, misuse potential) before release, and deployers conduct context-specific impact assessments and monitoring during use. This prevents a situation where every actor defers responsibility, claiming the harm was not within their purview.

3.3 Benefit

This principle holds that those who derive economic benefit from AI should bear commensurate responsibilities for its impacts, reflecting a “polluter pays” logic. Foundation model providers profit from distributing their models; integrators profit from AI-powered products; deployers gain efficiency from AI deployment. The EU legislature explicitly recognised that the impact of GPAI is too significant to leave to a few powerful players reaping economic benefits without oversight (Recital 105). By extending obligations to GPAI providers – including public disclosure of training data summaries (Article 53(1)(d)) – the Act ensures that companies at the top of the value chain contribute to risk reduction rather than pushing all compliance costs downstream.

Similarly, deployers derive direct value from AI use and are accordingly charged with investing in staff training, performance monitoring, impact assessments, and transparency measures. If using AI saves a bank money through automated credit scoring, some of that saving must be reinvested in compliance measures – a fair allocation of the burden. Had the Act placed all obligations only on providers, deployers might externalise risk by using AI recklessly because “it was certified by the provider.”

The benefit principle also underlies liability discussions. The proposed AI Liability Directive reflects the idea that the enterprise benefiting from deployment should bear responsibility for harms, especially where victims would otherwise go uncompensated. This alignment of “pay-off and

accountability” internalises the externalities of AI use, echoing basic fairness principles in regulation.

Benefit ties into power asymmetry as well. Foundation model providers are often major tech companies; deployers may be smaller organisations. The Act tries to be even-handed: obligations are assigned by role, not by size per se. However, Recital 109 and Article 56(5) introduce proportionality based on size and capacity, recognising that a small open-source developer should not bear the same burden as a large commercial AI lab. This tempering by the next principle – capability – ensures regulatory equity.

3.4 Capability

This principle considers the ability and resources of an actor to fulfil obligations. Foundation model developers possess specialised expertise (AI research talent, computing infrastructure) to conduct extensive testing – probing a model for biases, red-teaming for adversarial exploits, evaluating emergent capabilities. Downstream deployers typically lack these capabilities, treating the model as a given component. The Berkeley AI RMF Profile for GPAI explicitly notes that many risk controls – such as addressing emergent behaviours or installing safety mitigations – can only meaningfully be implemented by upstream developers who are “in a better position” than downstream actors.¹⁵ For example, if a language model occasionally produces personal data from its training set (a privacy risk), the model provider is capable of detecting and minimising that through training adjustments or filters far more than an end-user deployer.

¹⁵ Anthony Barrett et al, “AI Risk-Management Standards Profile for General-Purpose AI Systems (GPAIS) and Foundation Models” (2023) *CLTC Berkeley* (available at <https://cltc.berkeley.edu/wp-content/uploads/2023/11/Berkeley-GPAIS-Foundation-Model-Risk-Management-Standards-Profile-v1.0.pdf>) 7–8, 66–67.

Conversely, certain risk mitigations require local knowledge or real-time intervention that only downstream actors can provide. A hospital deploying an AI diagnostic tool is capable of performing clinical validation with its patient data and can impose usage protocols. The model maker cannot do this for every deployment. The Act accordingly requires deployers to monitor outcomes, keep logs, ensure human oversight, and suspend use if needed – tasks that leverage the deployer’s on-site capability.

Economic and organisational capacity matters too. Recital 109 calls for codes of conduct to “take due account of the size of the provider and allow simplified compliance for SMEs and start-ups.” Article 56(5) provides that reporting obligations for systemic-risk model providers should reflect differences in size and capacity. This capacity-based tailoring ensures that regulation does not inadvertently favour only the biggest players who can absorb high compliance costs. For instance, if a small company fine-tunes an open model for a niche application, it might not be realistic for it to run massive retraining to fix base model issues – instead, its obligation might reasonably be to document the limits and test the narrower use, while the heavy lifting of fundamental model improvements lies with the original provider.

Capability also involves technical access. A deployer using a black-box proprietary model via API cannot realistically explain how the model works internally – only the provider can (hence the Act requires GPAI providers to supply information and even source code to authorities upon request). If the model is open, a deployer might have more capability to inspect it and may then share responsibility for modifications.

3.5 Applying the Principles

In practice, these principles work together. The obligation to ensure “accuracy, robustness, and cybersecurity” of a high-risk AI system (Article 15) is distributed

precisely according to these criteria: the provider trains a robust model and discloses known accuracy limits (capability, control), the integrator validates for the specific application (capability, foreseeability), and the deployer monitors ongoing performance and secures the operational environment (control, capability). Each benefits from the system's accuracy (benefit) and thus has aligned incentives.

By aligning obligations with who holds the levers (control), who can anticipate issues (foreseeability), who gains from the AI (benefit), and who can act to mitigate (capability), the framework aims for both efficiency and fairness. No actor is overburdened with tasks beyond its reach, and none can evade responsibility by pointing fingers. This approach reflects what the regulatory literature calls shared responsibility in complex supply chains – fair precisely because it is grounded in these factors. It also helps interpret grey areas: if a question arises whether a fine-tuner or the original model provider should fix a particular issue, we can ask – who has the control over that aspect of the model and the capability to address it?

4 Liability and Enforcement under Market Surveillance vs. Data Protection Supervision

4.1 The AI Act's Enforcement Architecture

The enforcement of the AI Act is entrusted to national Market Surveillance Authorities (“MSAs”), coordinated at EU level by the European AI Board and the AI Office. This mirrors the New Legislative Framework for traditional product regulation, where each Member State has authorities ensuring that products on the market comply with EU safety rules. MSAs have powers including requiring information from providers, ordering corrective actions, banning or recalling non-compliant AI systems, and issuing administrative fines.

For general-purpose AI models specifically, enforcement is more centralised. The Commission and the AI Office directly enforce the Chapter V provisions (Articles 53-55) in cooperation with Member States. The AI Office can investigate foundation models and, if a provider of a systemic-risk model fails to cooperate, the Commission can impose fines up to 3% of worldwide turnover. This centralisation was a deliberate design choice, recognising that foundation model providers are few, often headquartered outside the EU, and that a unified approach avoids the bottlenecks experienced under GDPR. As Hacker (2024) observes, the EU appears to have learned from the GDPR experience – where Ireland’s DPA effectively set the enforcement pace for all of Europe through the one-stop-shop mechanism – by opting for direct Commission involvement for the highest-impact upstream actors.¹⁶

The Act’s penalty regime (Article 99) is tiered by severity: up to 7% of global annual turnover (or EUR 35 million) for prohibited practices, 3% for GPAI provider non-compliance, and lower tiers for other breaches. These thresholds exceed GDPR’s 4% ceiling, signalling the EU’s intent to make AI rule-breaking potentially more costly.

The European AI Board, comprising Member State representatives and the Commission, will facilitate consistent application and develop guidance. However, unlike the European Data Protection Board’s (“EDPB’s”) binding dispute resolution mechanism under GDPR Article 65, the AI Board has more limited powers – consistency will depend on coordination and Commission oversight. The AI Office itself has begun issuing operational guidance, including on prohibited practices (effective February 2025) and on GPAI model obligations,

¹⁶ Hacker (n 2), 7–8.

though stakeholders have raised concerns that some guidance has been published too close to compliance deadlines.

4.2 Data Protection Supervision and Overlap

The GDPR is enforced by independent Data Protection Authorities (“DPAs”) through a cooperation mechanism for cross-border cases. Because AI systems frequently process personal data, both the AI Act and GDPR will often apply simultaneously. A biometric identification system in a public space engages the AI Act (high-risk classification under Annex III) and GDPR (processing of biometric data) in parallel. A company deploying such a system could face scrutiny from an AI MSA checking conformity and accuracy requirements, and from a DPA checking lawfulness of data processing and purpose limitation.

The AI Act states it is “without prejudice to” GDPR and other data protection laws, meaning compliance with both is required simultaneously. Some Member States have signalled that their DPAs will handle certain AI Act aspects – the Netherlands’ DPA announced its expectation to be involved in AI Act enforcement, and German data protection commissioners have argued for a role given their experience in algorithmic assessments. Other States will establish separate AI authorities. The resulting institutional landscape requires robust coordination to avoid duplicated enforcement, inconsistent demands, and regulatory arbitrage.

Commentators have highlighted a particularly acute tension: the Act requires deployers to maintain logs for at least six months, while GDPR’s data minimisation principle may pull in the opposite direction where logs contain personal data. A related issue arises where the AI Act requires training data transparency summaries but GDPR’s storage limitation principle would call for deletion of personal data no longer needed. These conflicts are not irreconcilable – logging for legal compliance can be grounded in GDPR Article 6(1)(c) and data

minimisation can be addressed through pseudonymisation and aggregation – but they require coordinated regulatory guidance.

4.3 Civil Liability

The AI Act’s administrative fines do not directly compensate those harmed by AI. The proposed AI Liability Directive (“AILD”),¹⁷ though not yet enacted, would complement the Act by relaxing the burden of proof on causality: if a provider violates an AI Act requirement designed to prevent the kind of harm that occurred, a rebuttable presumption of causality would arise. This essentially binds regulatory obligations to civil accountability: breach of the AI Act eases the victim’s path to compensation. The updated Product Liability Directive,¹⁸ which explicitly covers software and AI, further clarifies producers’ strict liability for defective AI outputs.

The interplay is designed to close the accountability loop: the AI Act establishes duties (breach constitutes fault or defect), and the liability framework channels that breach into redress for victims. This ensures both deterrence (via administrative fines) and compensation (via civil liability), addressing what scholars including Vasudevan (2023) and Lior (2021) have identified as the critical gap in pre-Act AI governance.¹⁹ The allocation principles discussed in

¹⁷ Proposal for a Directive of the European Parliament and of the Council on adapting non-contractual civil liability rules to artificial intelligence (AI Liability Directive) COM(2022) 496 final <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52022PC0496> accessed 17 June 2026.

¹⁸ Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC [2024] OJ L, 2024/2853 <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024L2853> accessed 17 June 2026.

¹⁹ Anat Lior, “Artificial Intelligence and Tort Law: Who Should Be Held Liable When AI Causes Damages?” available at <https://il.boell.org/en/2021/12/24/artificial-intelligence-ai-tort-law-and-network-theory-who-should-be-held-liable-when-ai>; Amrita Vasudevan, “Addressing the Liability Gap in AI Accidents” (2023) *Centre for International Governance Innovation Policy Brief No. 177*, available at https://www.cigionline.org/static/documents/PB_no.177.pdf.

Section 3 – control, foreseeability – will inevitably influence courts and the application of the AILD: who controlled the relevant aspect of the AI, and who could have foreseen the harm, will determine where liability falls.

5 Cross-Jurisdictional Perspective

Regulation of AI is a global challenge, with different jurisdictions pursuing markedly different approaches. This section compares the EU’s framework with developments in the United Kingdom and United States, focusing on general-purpose AI governance and updated to reflect significant recent policy shifts.

5.1 United Kingdom

The UK, no longer bound by EU regulations post-Brexit, has chosen not to replicate the EU AI Act. Instead, the UK government set out its approach in a White Paper titled “A pro-innovation approach to AI regulation”,²⁰ characterised by non-statutory guidance, reliance on existing regulators and laws, and five cross-cutting principles: safety, transparency, fairness, accountability, and contestability. These principles are to be interpreted and applied by existing sectoral regulators – the Financial Conduct Authority (“FCA”), Medicines and Healthcare products Regulatory Agency (“MHRA”), Information Commissioner’s Office (“ICO”), Equality and Human Rights Commission, and others – within their respective domains.

The White Paper acknowledged foundation models as a “novel challenge” for governance, noting that only a “relatively small number of organizations” develop them and that “it can be challenging to identify and allocate

²⁰ Department for Science, Innovation and Technology, *A Pro-innovation Approach to AI Regulation* (CP 754, March 2023) <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper> accessed 17 June 2026

accountability for outcomes” when these models are used widely. However, the UK chose not to impose statutory obligations on foundation model providers, preferring a monitoring-and-evidence approach. Initiatives include the AI Safety Institute (established with GBP 100 million in funding, succeeding the Foundation Model Taskforce) and the Bletchley Declaration on frontier AI safety – cooperative risk governance rather than hard regulation.

The Department for Science, Innovation and Technology’s (“DSIT’s”) voluntary Code of Practice for Foundation Model Developers aligns substantively with many AI Act themes – data transparency, robust testing, information sharing with downstream users – but remains non-binding.²¹ The Competition and Markets Authority (“CMA”) published draft principles on foundation models addressing competition and consumer protection,²² while the ICO has issued AI-specific guidance on data protection, bias, and transparency.²³

The UK’s approach distributes accountability through existing legal frameworks: employers deploying AI must comply with the Equality Act,²⁴ financial institutions with FCA rules, healthcare providers with MHRA medical device regulations. This represents diffuse accountability buttressed by broad duties of responsible use. It lacks the EU’s sharp delineation of role-specific

²¹ Department for Science, Innovation and Technology, *AI Cyber Security Code of Practice* (31 January 2025) <https://www.gov.uk/government/publications/ai-cyber-security-code-of-practice> accessed 17 June 2026.

²² Competition and Markets Authority, *AI Foundation Models: Initial Report* (18 September 2023) <https://www.gov.uk/government/publications/ai-foundation-models-initial-report> accessed 17 June 2026.

²³ Information Commissioner's Office and The Alan Turing Institute, *Explaining Decisions Made with AI* (20 May 2020) <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/explaining-decisions-made-with-ai/> accessed 17 June 2026.

²⁴ Equality Act 2010.

obligations, though individual regulators are beginning to fill the gap with sector-specific guidance. The White Paper itself concedes that foundation models' "wide-reaching impact... means they are unlikely to be directly caught by any single regulator," and the UK has signalled that if voluntary measures prove insufficient, statutory regulation may follow.

5.2 United States

The US regulatory landscape for AI has undergone significant shifts. The Biden Administration's approach centred on the National Institute of Standards and Technology's ("NIST's") voluntary AI Risk Management Framework,²⁵ which provides a structured approach for identifying, assessing, and managing AI risks through four functions (Govern, Map, Measure, and Manage); the Blueprint for an AI Bill of Rights,²⁶ outlining five principles for the design and use of automated systems; and Executive Order 14110 on Safe, Secure, and Trustworthy AI,²⁷ which invoked the Defense Production Act²⁸ to require developers of very large models to notify the government and share safety test results.

However, the change of administration in January 2025 brought a sharp reversal. President Trump revoked Executive Order 14110 in his first days in office, signalling a clear preference for deregulation and innovation-led growth

²⁵ National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1, US Department of Commerce, January 2023) <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf> accessed 17 June 2026. (NIST AI RMF).

²⁶ White House Office of Science and Technology Policy, *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People* (October 2022) <https://www.govinfo.gov/app/details/GOVPUB-PREX23-PURL-gpo193638> accessed 17 June 2026.

²⁷ Executive Order No 14110, *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*, 88 Fed Reg 75191 (1 November 2023) <https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence> accessed 17 June 2026.

²⁸ Defense Production Act of 1950, 50 USC §§ 4501–4568.

over prescriptive safety mandates. This revocation removed the mandatory reporting requirements for developers of large models and weakened the executive branch's direct oversight role. The current US approach relies primarily on sectoral enforcement through existing agencies: the Federal Trade Commission ("FTC") enforces against deceptive or unfair AI practices under consumer protection law; the Equal Employment Opportunity Commission ("EEOC") holds employers responsible for discriminatory AI hiring outcomes under Title VII; the Food and Drug Administration ("FDA") reviews AI-based medical devices under existing device regulations; and state-level legislation continues to expand (e.g., the Colorado AI Act,²⁹ Illinois BIPA,³⁰ New York City's bias audit law for automated employment decision tools).³¹

NIST's AI RMF remains influential as a voluntary benchmark, including its GPAI Profile³² which provides detailed guidance for foundation model developers on risk governance, testing, and documentation. However, without the mandatory reporting requirements of the now-revoked Executive Order, the framework's reach is limited to voluntary adoption and market pressure from customers demanding assurance.

The US approach effectively places primary legal responsibility on deployers: employers, banks, and other end-users face liability under anti-

²⁹ Colorado Artificial Intelligence Act, Colo Rev Stat §§ 6-1-1701–6-1-1711 (2024) <https://leg.colorado.gov/bills/sb24-205> accessed 17 June 2026.

³⁰ Illinois Biometric Information Privacy Act, 740 ILCS 14/1–99 (2008) <https://law.justia.com/codes/illinois/chapter-740/act-740-ilcs-14/> accessed 17 June 2026.

³¹ New York City Local Law No 144 of 2021 (Automated Employment Decision Tools), codified at NYC Administrative Code §§ 20-870–20-874 <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page> accessed 17 June 2026.

³² National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile* (NIST AI 600-1, US Department of Commerce, July 2024) <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf> accessed 17 June 2026.

discrimination, consumer protection, and sector-specific statutes, and in turn pressure AI vendors through contracts to provide bias audits, documentation, and indemnities. This market-driven chain obligation partially mirrors the EU's value-chain allocation but without the statutory granularity of the Act's role-specific duties.

5.3 Convergence Through Standards

Despite divergent regulatory philosophies, convergence is emerging through international standards. ISO/IEC 23894 (risk management),³³ ISO/IEC 42001 (AI management systems),³⁴ and the work of CEN-CENELEC JTC 21³⁵ provide a common technical vocabulary. An AI developer following these standards may simultaneously satisfy EU harmonised standard requirements, UK recommended practice, and US NIST guidance. The anticipated “Brussels Effect” – whereby companies preparing for EU AI Act compliance extend those practices globally – is already visible, as it occurred with GDPR. Foundation model providers operating across jurisdictions will find it more efficient to implement a single, EU-Act-aligned governance system than to maintain separate regional processes.

The growing transatlantic divergence following the US deregulatory shift may, however, create pressure on EU policymakers from industry arguing for regulatory “parity.” The EU should resist a race to the bottom while remaining

³³ International Organization for Standardization and International Electrotechnical Commission, *ISO/IEC 23894:2023 Information Technology—Artificial Intelligence—Guidance on Risk Management* (ISO 2023) <https://www.iso.org/standard/77304.html> accessed 17 June 2026.

³⁴ International Organization for Standardization and International Electrotechnical Commission, *ISO/IEC 42001:2023 Information Technology—Artificial Intelligence—Management System* (ISO 2023) <https://www.iso.org/standard/81230.html> accessed 17 June 2026.

³⁵ CEN-CENELEC Joint Technical Committee 21, *Artificial Intelligence* <https://www.cenelec.eu/areas-of-work/cen-cenelec-topics/artificial-intelligence/> accessed 17 June 2026.

pragmatic about implementation timelines and compliance burdens – a balance explored further in the recommendations below.

6 Recommendations for Guidance and Standardisation

Drawing from the analysis above and accounting for recent regulatory and political developments, this section presents recommendations for clarifying and operationalising the allocation of AI governance duties across the value chain.

6.1 Develop Detailed Guidance on Role-Specific Obligations

Regulators should issue interpretative guidance concretely delineating the obligations of each actor under various scenarios. The AI Office has begun this process: its General-Purpose AI Code of Practice,³⁶ published in draft form in late 2024, represents an important first step. However, gaps remain. Guidance should clarify what constitutes “sufficiently detailed” documentation for a foundation model, perhaps aligning with the content of model cards (architecture, training data summary, intended uses and limits, known risks) and referencing Annexes XI-XII as checklists. It should clarify when a downstream modifier becomes a “provider” of a new model, providing thresholds or criteria – for instance, if model behaviour is significantly altered or if it is fine-tuned for a high-risk use, the fine-tuner likely assumes provider duties. This could prevent uncertainty that discourages beneficial adaptations. For integrators, guidance should outline how to fulfil system-level requirements using upstream documentation – a blueprint for creating Annex IV technical documentation for a system incorporating a

³⁶ European Commission AI Office, *First Draft of the General-Purpose AI Code of Practice published, written by independent experts* (14 November 2024) <https://digital-strategy.ec.europa.eu/en/library/first-draft-general-purpose-ai-code-practice-published-written-independent-experts> accessed 17 June 2026.

GPAI model, including how to perform a gap analysis between the model's Annex XII information and the system's risk assessment needs. For deployers, sector-specific guidelines on implementing Article 26 obligations should be produced in cooperation with industry associations – for instance, how to conduct a Fundamental Rights Impact Assessment in employment contexts, and how to set up logging and human oversight in practice, ideally aligned with existing frameworks (such as Data Protection Impact Assessments (“DPIAs”) under GDPR or safety management systems) to avoid duplication.

6.2 Foster Collaborative Interface Mechanisms

To ensure smooth information exchange between upstream and downstream actors, stakeholders should standardise the format and content of key documentation artefacts. Regulators should endorse and potentially mandate a standardised model card format for GPAI providers to use when complying with Article 53(1)(b) and (d). The format should include sections on model overview (architecture, size), intended purposes, training dataset overview, performance benchmarks, limitations and known failure modes, and usage guidelines. The AI Code of Practice (Article 56) can incorporate the requirement for signatories to publish model cards according to this template, enabling semi-automated compliance and ensuring downstream integrators know exactly where to find needed information.

Integrators should be encouraged to produce system cards documenting the specific context of use, how the model is integrated, performance in that context, human oversight and risk mitigations, and constraints on use. Making a summarised system card public could improve transparency and trust for users and affected persons.

Standardisation bodies (CEN-CENELEC JTC 21³⁷, ISO/IEC JTC1 SC 42³⁸) should prioritise developing a technical standard for AI system logging: what events should be logged, format requirements, retention periods, and data protection safeguards. The CEN-CENELEC “AI Trustworthiness Framework”³⁹ standard under development should be accelerated and designated as a harmonised standard once published. Collaboration with data protection experts is essential to ensure logging does not violate GDPR – for instance, by logging references or hashed values rather than raw personal data.

6.3 Leverage International Standards

The Commission should actively adopt relevant AI standards to confer presumption of conformity. CEN/CENELEC standards on AI Risk Management and Quality Management Systems should be aligned with ISO/IEC 23894:2023⁴⁰ and ISO/IEC 42001⁴¹ and cited in the Official Journal as harmonised standards for Article 9 and Article 17 compliance. Standards such as CEN/CLC ISO/IEC/TR 24027:2023 (bias measurement)⁴² and CEN/CLC ISO/IEC/TR 24029-1:2023 and EN ISO/IEC 24029-2:2023 (robustness assessment)⁴³ should be referenced for

³⁷ *Supra* (n 35).

³⁸ International Organization for Standardization and International Electrotechnical Commission, *ISO/IEC JTC 1/SC 42 – Artificial Intelligence* <https://www.iso.org/committee/6794475.html> accessed 17 June 2026.

³⁹ CEN-CENELEC JTC 21 (n 35).

⁴⁰ *Supra* (n 33).

⁴¹ *Supra* (n 34).

⁴² CEN-CENELEC, *CEN/CLC ISO/IEC/TR 24027:2023 Information Technology – Artificial Intelligence (AI) – Bias in AI Systems and AI Aided Decision Making (ISO/IEC TR 24027:2021)* (CEN-CENELEC 2023).

⁴³ CEN-CENELEC, *CEN/CLC ISO/IEC/TR 24029-1:2023 Artificial Intelligence (AI) – Assessment of the Robustness of Neural Networks – Part 1: Overview (ISO/IEC TR 24029-1:2021)* (CEN-CENELEC 2023). CEN-CENELEC, *EN ISO/IEC 24029-2:2023 Artificial Intelligence (AI) – Assessment of the Robustness of Neural Networks – Part 2: Methodology for the Use of Formal Methods (ISO/IEC 24029-2:2023)* (CEN-CENELEC 2023).

compliance with the Act's data governance (Article 10) and accuracy/robustness (Article 15) requirements. Standards for AI evaluation methodologies and auditor qualifications should be developed to ensure consistency among notified bodies. These technical recipes reduce ambiguity in the Act's open-textured obligations, giving companies clear methods for demonstrating compliance and enabling a degree of global alignment through dual ISO/CEN adoption.

6.4 Account for the Evolving Political and Regulatory Context

Several developments since the Act's adoption require attention in implementation.

The Digital Omnibus Regulation.⁴⁴ The Commission's proposed Simplification Omnibus (November 2025) would amend the AI Act alongside the Cybersecurity Act and Machinery Regulation, seeking to reduce compliance costs, particularly for SMEs. Proposals include extended timelines, simplified conformity procedures for lower-risk applications, and clearer guidance on the interplay between the AI Act and sectoral product legislation. These simplification measures should be welcomed insofar as they genuinely reduce unnecessary burden, but must not hollow out the Act's substantive protections – particularly the transparency and risk management requirements that form the backbone of the value-chain allocation described in this article.

The Draghi Report and Competitiveness Debate. The Draghi Report on European competitiveness (September 2024)⁴⁵ identified overregulation as a drag on innovation and urged the EU to shift toward an "innovation principle." While it did not specifically target the AI Act, its broader critique has fuelled calls

⁴⁴ *Supra* (n 6).

⁴⁵ European Commission, *The Draghi Report on EU Competitiveness*

https://commission.europa.eu/topics/competitiveness/draghi-report_en accessed 17 June 2026.

from industry associations for a moratorium on AI Act enforcement deadlines. Major groups including BusinessEurope and DigitalEurope have argued that enforcement should be delayed until sufficient guidance, standards, and compliance tools are available.⁴⁶ These calls warrant a calibrated response: compliance timelines should be realistic and accompanied by adequate guidance and tooling, but a blanket moratorium would undermine the Act's credibility and the EU's standard-setting position. The better approach is to accelerate the production of harmonised standards, templates, and guidance so that compliance becomes technically feasible before enforcement begins.

AI Office Guidelines.⁴⁷ The AI Office has been issuing guidance on prohibited AI practices, GPAI model obligations, and the Code of Practice. These are essential for translating the Act's legal text into operational requirements. Stakeholders have noted that some guidance has arrived late relative to compliance deadlines, creating uncertainty. The AI Office should publish a clear roadmap of forthcoming guidance, with deadlines aligned to enforcement milestones.

The US Deregulatory Shift. The Trump Administration's revocation of Executive Order 14110 creates a growing transatlantic divergence. EU policymakers should be aware that this divergence may intensify pressure from EU-based companies arguing for regulatory "parity." However, the EU's comparative advantage lies in offering a predictable, rights-protective framework that can serve as the global baseline.

⁴⁶ BusinessEurope, *Joint Industry Statement on the EU Artificial Intelligence (AI) Act* (2025) <https://www.buinessurope.eu/publications/joint-industry-statement-on-the-eu-artificial-intelligence-ai-act/> accessed 17 June 2026. DigitalEurope, *Joint Industry Statement on the AI Omnibus: Administrative Clean-up or a Boost for Europe's AI Competitiveness?* (12 March 2026) <https://www.digitaleurope.org/news/joint-industry-statement-on-the-ai-omnibus-administrative-clean-up-or-a-boost-for-europes-ai-competitiveness/> accessed 17 June 2026.

⁴⁷ *Supra* (n 7).

6.5 Enhance Coordination Between AI and Data Protection Regimes

The EDPB and the European AI Board should produce joint guidance on conducting combined DPIAs and Fundamental Rights Impact Assessments, reconciling logging with data minimisation, and delineating lead authority responsibilities. Memoranda of understanding between national AI authorities and DPAs should be established where they are separate bodies. A streamlined incident notification mechanism should be developed for events implicating both regimes. Joint guidance should also clarify that actions to comply with the AI Act (e.g., collecting sensitive attribute data for bias testing) can be legally grounded under GDPR where done proportionately, and that explanation requirements under both regimes can be aligned so that one disclosure satisfies both.

6.6 Encourage Sector-Specific Implementation

Different industries will apply the Act's requirements in domain-specific contexts. Sectoral guidance for healthcare AI (building on medical device regulations and ISO 81001-1:2021⁴⁸), financial services AI (through EBA/ESMA codes of conduct), and public sector AI (integrating human rights obligations with AI Act requirements) should translate general obligations into domain-appropriate practices. Regulatory sandboxes (Article 57) should be actively used to pilot compliance approaches. SME-friendly templates, checklists, and open-source compliance tools should be funded through the Digital Europe Programme and Horizon Europe.

⁴⁸ International Organization for Standardization and International Electrotechnical Commission, *ISO 81001-1:2021 Health Software and Health IT Systems Safety, Effectiveness and Security – Part 1: Principles and Concepts* (ISO/IEC 2021) <https://www.iso.org/standard/91523.html> accessed 17 June 2026.

6.7 Strengthen Upstream-Downstream Communication Through Contracts

The AI Office should publish recommended contractual clauses for agreements between GPAI providers and downstream actors, analogous to GDPR standard contractual clauses. These should cover obligations to supply Annex XII information, notify downstream parties of newly discovered risks, cooperate in remediation, and respect usage instructions. Model providers of systemic-risk GPAI should maintain awareness of significant downstream deployments to enable prompt risk communication, akin to product recalls. Third-party certification schemes and trust marks for GPAI models, leveraging Article 57, can reduce duplicated due diligence while maintaining accountability. An information clearinghouse – an online portal where GPAI providers post model cards, risk analyses, and usage guidelines – could particularly assist SMEs and open-source users.

6.8 Ensure Adaptive Governance

The Commission's mandated review (Article 112) should specifically assess whether the allocation of duties remains appropriate as technology evolves. An expert advisory group should continuously monitor AI developments and recommend adjustments. International regulatory dialogue through the EU-US Trade and Technology Council and the Global Partnership on AI should facilitate convergence. The AI Office should publish regular reports on the state of AI compliance and the risk landscape, flagging systemic implementation challenges and recommending solutions.

7 Conclusion

This article has developed a conceptual framework for understanding how the

EU AI Act allocates governance duties across the AI value chain. Four allocation principles – control, foreseeability, benefit, and capability – provide the normative foundation for the Act’s stratified approach: foundation model providers bear upstream transparency and safety obligations because they control model design and possess the capability for extensive testing; fine-tuners inherit and extend those obligations proportionate to their modifications; integrators shoulder system-level risk management and conformity assessment duties as the actors who determine intended use; and deployers assume responsibility for operational oversight, logging, and context-specific risk mitigation.

This allocation represents a significant evolution from the Act’s earlier drafts, which overlooked general-purpose AI entirely. It also reflects a broader regulatory maturation: the recognition that complex AI supply chains require shared responsibility, with each actor’s obligations calibrated to their actual position in the chain. The comparative analysis confirms that while the UK and US pursue different regulatory strategies – principles-based guidance and sectoral enforcement, respectively – the underlying problem of distributing accountability across the AI lifecycle is universal, and convergence through international standards is already under way.

The Act’s success will depend on implementation. The recommendations presented here – from harmonised standards and model documentation templates to regulatory coordination mechanisms and adaptive governance – aim to translate the legal framework into operational reality. Recent developments underscore the need for both rigour and pragmatism. The Digital Omnibus proposals, the AI Office’s emerging guidelines, the Draghi Report’s competitiveness critique, and industry calls for a moratorium on enforcement deadlines all reflect legitimate concerns about readiness. The appropriate response is not to delay the Act but to accelerate the production of the tools –

standards, guidance, templates, compliance toolkits – that make compliance feasible and proportionate, particularly for SMEs and open-source developers.

The EU AI Act offers a principled and potentially global benchmark for distributing AI governance duties. Its framework of accountable handoffs – where documentation, testing results, and risk assessments flow along the supply chain from upstream to downstream – creates a chain of trust in which each actor knows its role and can verify that others are fulfilling theirs. If effectively implemented, this approach can serve not only Europe but, through the Brussels Effect and international standard-setting, the wider global effort to ensure that the benefits of general-purpose AI are realised in a manner that is innovative, trustworthy, and respectful of fundamental rights.