cerre

INFORMATION SHARING AND COOPERATION ALONG THE AI VALUE CHAIN

ISSUE PAPER

June 2026

Daniel Schnurr

As provided for in CERRE's bylaws and procedural rules from its "Transparency & Independence Policy", all CERRE research projects and reports are completed in accordance with the strictest academic independence.

The CERRE Regulatory Lab series, of which this note is part, received the support and/or input from several CERRE members including Apple, Booking.com, and Microsoft. Each draft note is discussed in Chatham House meetings between the CERRE corporate members, national regulatory authorities, EU institutions and civil society organisations. However, they bear no responsibility for the contents of this report. The views expressed in this CERRE report are attributable only to the authors in a personal capacity and not to any institution with which they are associated. In addition, they do not necessarily correspond either to those of CERRE, or of any sponsor or of members of CERRE. AI-based tools were used to assist with text editing and language refinement in this paper. The author retains full responsibility for its content and conclusions.

info@cerre.eu – www.cerre.eu

Table of Contents

TABLE OF CONTENTS 1

1. INTRODUCTION 2

2. ILLUSTRATIVE AI VALUE CHAIN SCENARIOS 5

2.1 High-risk AI systems built on top of GPAI systems 5
2.2 Cloud computing services, model fine-tuning, and data inputs 5
2.3 Agentic AI systems, orchestration of GPAI models, and tool use of external services 6

3. GENERAL PRINCIPLES FOR INFORMATION SHARING UNDER ARTICLE 25 AIA 9

3.1 Information Flows under Article 25 AIA in Complex AI Value Chains 9
3.2 Conditions for Responsibilities under Article 25(4) AIA 11
3.2.1 Scope of Covered Input Suppliers 11
3.2.2 Scope of Legitimate Information Requests 12
3.3 Information Sharing Mechanisms 13

4. INFORMATION SHARING FOR SPECIFIC HIGH-RISK AI SYSTEM REQUIREMENTS 17

4.1 Article 9 AIA: Risk management system 17
4.2 Article 10 AIA: Data Governance 17
4.3 Article 15 AIA: Accuracy, robustness, and cybersecurity 18
4.4 Article 72 AIA: Post-market monitoring 19
4.5 Article 73 AIA: Serious incident reporting 20
4.6 Article 12 AIA: Record-keeping 21
4.7 Additional high-risk AI system obligations that may require information sharing 21

5. RECOMMENDATIONS FOR IMPLEMENTING THE AI ACT'S INFORMATION SHARING PROVISIONS 23

ABOUT CERRE 26

ABOUT THE AUTHOR 27

1. Introduction

The AI Act (AIA) establishes a horizontal framework for the regulation of artificial intelligence (AI).[1] By adopting a risk-based approach, it seeks both to promote the uptake of human-centric and trustworthy AI and to ensure a high level of protection for health, safety, and fundamental rights, while also supporting innovation.[2] Although the AIA entered into force in August 2024, the EU is still deep in the implementation process. This is particularly the case with regard to the rules governing high-risk AI systems and the role of input suppliers in supporting provider compliance. At the same time, the AI Act has become a focal point in wider debates about Europe's regulatory approach, and economic competitiveness, and the EU's poor record in commercialising digital innovation. These debates have culminated in the "Digital Omnibus on AI", a package of targeted amendments to the AIA that, among other changes, affect the application of high-risk requirements.[3]

A key issue in this context concerns the value chains through which AI products and services are developed and provided. These value chains, involving multiple independent economic actors, raise questions about how accountability for risk prevention and risk mitigation should be allocated, and what duties each actor must fulfil in light of its role in the value chain. The AI Act recognises this general challenge in the context of high-risk AI systems and establishes a regulatory value chain that mainly distinguishes between different operator roles of AI systems, with specific obligations assigned to each. Beyond high-risk AI systems, the AI Act also explicitly recognises the role and obligations of providers of general-purpose AI (GPAI) models, which may serve as inputs for downstream AI systems. However, in practice, AI value chains are often far more complex and dynamic than these pre-specified roles in the regulatory value chain.[4]

In this context, the AIA establishes general rules for cooperation along AI value chains, covering not only operators of high-risk AI systems, but also their input suppliers. Article 25(4) AIA sets out general responsibilities for any third party supplying an AI system, tools, services, components (including models), or processes that are used or integrated in a high-risk AI system. This is to ensure that providers of high-risk AI systems relying on inputs from third-party suppliers can fully comply with the obligations laid down in the AI Act. In particular, Article 25(4) requires that the parties specify by written agreement the necessary information, capabilities, technical access, and other assistance that the third party will provide to the provider of the high-risk AI system. Only third parties that provide their inputs under a free and open-source licence are exempted from this duty, unless the input they provide is a GPAI model.

The AI Act does not further specify the information, capabilities, access, or assistance exchanged between the parties beyond requiring that it be based on the generally acknowledged state of the art and enable compliance of the AI system provider. This lack of specificity may be viewed positively, as it allows for flexibility in the development of market practices and context-specific solutions.

[1] Regulation 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), http://data.europa.eu/eli/reg/2024/1689/oj.
[2] Recital 1 AIA.
[3] European Commission. (2025). Digital Omnibus Regulation Proposal. Available at https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-regulation-proposal; As of June 2026, the Digital Omnibus on AI had been politically agreed but not yet formally adopted or published in the Official Journal.
[4] Meyers, Z., Schnurr, D., & Larouche, P. (2025). The AI Act and Technological Neutrality. CERRE Issue Paper. Available at https://cerre.eu/publications/the-ai-act-and-technological-neutrality/. Engler, A. & Renda, A. (2022). Reconciling the AI Value Chain with the EU's Artificial Intelligence Act. Available at https://www.ceps.eu/ceps-publications/reconciling-the-ai-value-chain-with-the-eus-artificial-intelligence-act/.

Nonetheless, uncertainty about what forms of information sharing and cooperation between parties will be considered necessary and adequate could jeopardise investments in high-risk AI systems and their upstream inputs. That uncertainty may be further compounded by the need to ensure that information sharing and cooperation do not compromise the input supplier's intellectual property rights or trade secrets.[5] Worse, existing inputs may not be made available to providers of high-risk AI systems if third-party suppliers perceive the obligations to share information as uncertain or excessive. Because high-risk AI domains are also areas in which AI can generate significant social value (for example, medical treatment), the potential welfare losses would be particularly harmful.

Further regulatory guidance could help reduce this uncertainty. In particular, Article 25(4) states that the AI Office may develop and recommend voluntary model contract terms that can serve as templates for agreements between high-risk AI system providers and third-party input suppliers. However, such guidance must take into account the complexity, diversity, and fluidity of today's AI value chains in order to avoid being overly rigid or quickly outdated by technological and market developments. Because Article 25(4) governs the relationship between providers of high-risk AI systems and suppliers of more general inputs, it is important that implementation of the information sharing provisions preserves the AI Act's risk-tiered approach and avoids turning high-risk obligations into de facto general requirements for the provision of AI inputs. At the same time, assistance and information sharing along the AI value chain are vital for high-risk AI system providers to comply with the AIA requirements, and to implement effective risk management and mitigation.

To address these challenges and achieve the overarching goals of the AIA, the application and implementation of the AIA's information sharing rules should follow several high-level principles that also guide the analysis in this paper. To ensure effectiveness and promote innovation, implementation must ensure that rules can accommodate quick technological progress and evolving value chain configurations. Therefore, regulatory guidance should be principle-based and leave flexibility for context-specific compliance solutions. At the same time, testing and demonstrating how such principles operate in concrete use cases helps to reduce uncertainty arising from overly abstract rules and can provide market participants with one possible route to compliance. Furthermore, to ensure proportionality, the original risk-based approach of the AIA should be maintained. This is particularly important, as value chains and collaboration networks connected to high-risk AI systems will often extend beyond high-risk areas and include a wide range of suppliers and service providers that operate in the digital realm but may not even be directly concerned with AI technology.

Against this background, this paper takes a first step towards clarifying how cooperation and information sharing along the AI value chain, in relation to high-risk AI systems, could operate under the AI Act. To this end, the paper identifies major open questions that must be addressed for effective implementation, and develops recommendations for governing information sharing along the AI value chain under Article 25(4) for high-risk AI systems. In addition to setting out general principles for information sharing, the paper examines specific categories of information that may need to be shared in relation to selected requirements for high-risk AI systems.

The paper is part of the CERRE Regulatory Labs, which serve as forums to explore key regulatory and policy challenges related to the digital economy and society. The present Lab explored how to

[5] Recital 88 AIA.

implement selected provisions of the AI Act, specifically those on information sharing and cooperation along the AI value chain.

The remainder of this paper proceeds as follows. First, it discusses illustrative AI value chain scenarios to demonstrate the complex relationships that already emerge in relatively simple contemporary AI use cases. Second, it compares different information flow architectures for implementing information sharing under Article 25(4) AIA, examines which input suppliers fall within the scope of that provision and the limits of what they can be required to share, and suggests a combination of information sharing mechanisms to accommodate a diverse set of requirements that may be in tension with one another. Third, it discusses specific obligations for high-risk AI systems under the AI Act in order to derive categories of information that input suppliers may share to enable AI system providers to comply with those obligations. Finally, it concludes by summarising overarching recommendations on the implementation of the AI Act to support information sharing and cooperation along the AI value chain.

2. Illustrative AI Value Chain Scenarios

This section presents illustrative AI value chain scenarios underlying AI applications in high-risk domains. While these scenarios already go beyond the roles explicitly specified in the AI Act, they remain deliberately simple for the purposes of exposition and to focus the discussion. AI value chains in practice are often significantly more complex, increasing the number of relationships and the potential interactions between them.

2.1 High-risk AI systems built on top of GPAI systems

Many current AI applications rely on inputs from a GPAI model. Yet the model is not always directly integrated into the high-risk AI system. It is often embedded in an intermediate system that provides extended functionalities, user interfaces, safety measures and guardrails, or context-specific output tailoring. Where such a system, by virtue of integrating a GPAI model, can serve a variety of purposes, it qualifies as a *general-purpose AI system* under the AI Act. The provider of the GPAI model and the GPAI system may be the same entity (for example, OpenAI offering both the model GPT and the system ChatGPT) or different entities (for example, Microsoft offering the system Copilot using OpenAI's GPT model). This distinction is relevant for information sharing, because in the latter scenario the GPAI system provider may not have direct or full access to information about the GPAI model integrated into its system. Potential information flows in such cases are discussed in Section 3.1.

Scenario A: High-risk AI system built on top of a GPAI system that leverages a third-party GPAI model.

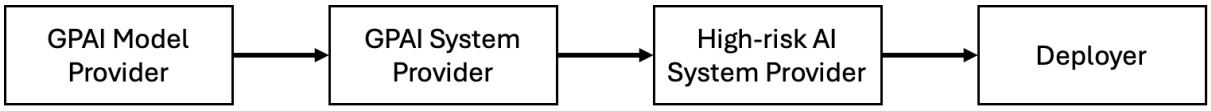

The same value chain scenario may arise where a GPAI system is deployed in a high-risk domain, even though the original system was not classified as high-risk. This would likely qualify as a modification of the AI system's intended purpose, in which case the deployer would become the provider of the high-risk AI system (Article 25(1) AIA). In that case, Article 25(2) AIA also requires the provider of the original system to cooperate closely with the new provider and to make available the necessary information, as well as provide the reasonably expected technical access and other assistance that are required for compliance with the AI Act's obligations.

2.2 Cloud computing services, model fine-tuning, and data inputs

Access to AI services, and to GPAI models in particular, is often provided through cloud computing services. Cloud service providers host a range of AI models, including models supplied by external providers.[6] Customers can then access GPAI models through application programming interfaces

[6] See, e.g., Microsoft Azure Foundry Models and Amazon Bedrock: Microsoft (2026). Foundry Models. Find the right model from exploration to deployment all in one place. Available at https://azure.microsoft.com/en-gb/products/ai-foundry/models; AWS (2026). Amazon Bedrock. The platform for building generative AI applications and agents at production scale. Available at https://aws.amazon.com/bedrock/.

(APIs). In this context, the cloud service provider becomes an intermediary input supplier that, in addition to providing access and hosting, may offer additional tools to augment the outputs of the GPAI model. In such scenarios, the provider of the high-risk AI system has a direct contractual relationship with the cloud computing provider(s). However, it may not have a direct contractual relationship with providers of other upstream inputs (such as the GPAI model or datasets used for fine-tuning).

This value chain scenario may become more complex if the customer uses the cloud computing service to fine-tune the original GPAI model using either internal business data or data from a specialised external data provider. In that case, the customer may become the provider of the fine-tuned GPAI model, which may be hosted by the same cloud service provider, moved to another provider, or hosted on-premise. Furthermore, the AI system provider may integrate a retrieval-augmented generation (RAG) system that leverages internal and external data sources during inference to produce more accurate and relevant outputs.[7]

In such cases, the risks associated with the final downstream AI system are potentially influenced by a large number of diverse inputs. Moreover, some inputs may be rather stable over time (for example, the fine-tuning of the GPAI model may only be conducted once), while others may be constantly changing (such as the data basis for the RAG system if it is updated on a continuous basis).

Scenario B: Cloud computing services, model fine-tuning, and data inputs

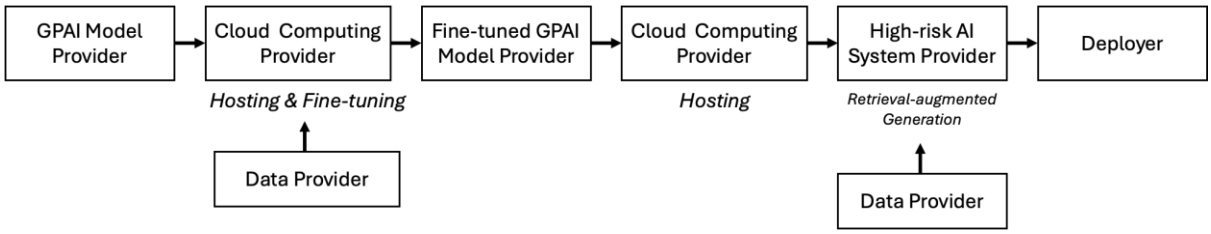

2.3 Agentic AI systems, orchestration of GPAI models, and tool use of external services

Agentic AI is widely regarded as a qualitative leap in the evolution of AI. Many major AI companies and cloud computing service providers now offer agentic AI systems that can be deployed directly in application contexts or integrated into downstream systems.[8] In contrast to more traditional AI systems, AI agents can operate with greater autonomy and interact with their environment, often through tools such as software, APIs, or external systems.

Most modern AI agents are powered by GPAI models as core inputs. In practice, agentic AI systems will often take the form of multi-agent GPAI systems that orchestrate and integrate specialised AI agents, each of which is powered by its own GPAI model (see illustration below).[9] For example, an AI

[7] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., ... & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. In: Advances in Neural Information Processing Systems 33 (NeurIPS 2020), 9459-9474.
[8] See, for example, Microsoft Azure. (2026). Azure AI apps and agents. Available at https://azure.microsoft.com/en-us/solutions/ai.
[9] Belcak, P., Heinrich, G., Diao, S., Fu, Y., Dong, X., Muralidharan, S., ... & Molchanov, P. (2025). Small language models are the future of agentic ai. Available at https://arxiv.org/pdf/2506.02153; Google (n.d.). Guide to multi-agent systems (MAS). Available at https://cloud.google.com/discover/what-is-a-multi-agent-system

agent serving as a financial advisor may integrate the inputs of a data analyst agent, a trading analyst agent, an execution agent, and a risk evaluation agent.[10]

In addition, agentic AI systems will typically access external APIs to execute actions or collect additional data during runtime. Online search engines are among the most popular external services called by agentic AI systems, as this allows the system to retrieve current information even post-deployment. These inputs can have significant effects on the risks associated with the agentic AI system, as they directly influence its behaviour and performance.

If such a multi-agent GPAI system is deployed to a high-risk domain or integrated into a high-risk AI system, the challenges to information sharing along the value chain may become even more complex than in Scenario A. The integration and interplay of several GPAI models as inputs for the multi-agent GPAI system further increase the number of potential relationships between the high-risk AI system provider and individual input suppliers.

At the same time, the risks associated with the multi-agent AI systems may be different from those associated with the individual GPAI models, as the interaction between subagents and their models may give rise to new behaviour and outputs. Therefore, to assess the risks of the AI system, a comprehensive analysis is typically needed. However, such an analysis is only required where the system is actually provided or deployed in a high-risk domain. In this case, Article 25(4) AIA raises the question to what extent the GPAI system provider would need to contribute information, technical access, or assistance to the high-risk AI system provider.

Moreover, the integration of API calls to service providers raises questions about the relationships of high-risk AI system providers and these service providers as third-party input suppliers. It appears particularly challenging for the high-risk AI system provider to obtain a written agreement with the web service provider. At the same time, the web service provider may not even be aware that its service is integrated into a high-risk AI system, as the connection between the AI system and the web service is only established during runtime.

An open question is whether such services may be exempted from the responsibilities under Article 25(4) AIA. However, while web services may often be accessed for free, they are typically not made available under an open-source licence.

[10] See the Financial Advisor Agent available in Google Vertext AI Agent Builder, available at https://console.cloud.google.com/vertex-ai/agents/agent-garden/samples/financial-advisor.

Scenario C: Agentic AI systems, orchestration of GPAI models, and tool use of external services

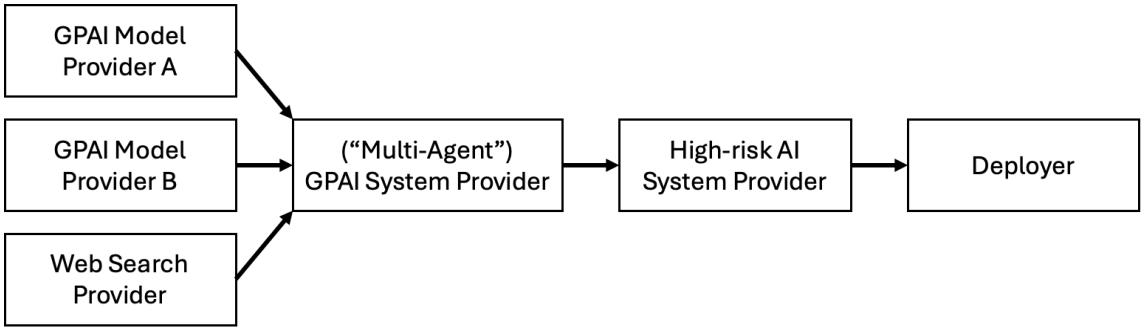

3. General Principles for Information Sharing under Article 25 AIA

Several observers have concluded that, in these more complex scenarios, providers of high-risk AI systems are not able to fully comply with the regulatory obligations imposed by the AI Act on their own.[11] To enable effective assurance across AI value chains, Allison et al. propose the concept of value chain accountability, understood as the processes and mechanisms through which all actors across AI value chains can be appropriately held to account by other actors and stakeholders. In this context, information sharing is viewed as a central means of improving value chain accountability. In particular, it is expected to mitigate the limited "accountability horizon" of individual actors, who may not even be aware of the full value chain in which they operate.[12]

At the same time, as illustrated in the previous section, the complexity and fluidity of relationships in many AI value chain scenarios create significant challenges for information sharing itself and, therefore, for the implementation of Article 25(4) AIA. This section therefore examines information flows in complex AI value chains (Section 3.1), as well as the conditions under which input suppliers fall within the scope of Article 25(4), the extent to which they can exclude the use of their inputs in high-risk AI systems, and the limits on the information requests they must satisfy (Section 3.2). In doing so, it addresses three central questions: **how information is to flow through complex AI value chains**, **which input suppliers are subject to the obligations laid down in Article 25(4) AIA**, and **how far those obligations extend in practice**. Finally, the section considers potential information sharing mechanisms, including their key requirements and design dimensions (Section 3.3).

3.1 Information Flows under Article 25 AIA in Complex AI Value Chains

A central issue in allocating responsibilities along the AI value chain is how information should flow through the value chain and reach the provider of a high-risk AI system so that it can comply with its obligations under the AIA. Even the relatively simple Scenario A outlined in Section 2 raises the issue of how information from different upstream input suppliers should flow to the provider of the high-risk AI system.

Two broad architectures are conceivable (see Figure 1): Under a **bilateral information flow architecture**, information flows directly between each third-party input supplier and the high-risk AI system provider. In Scenario A, for example, the high-risk AI system provider could receive information separately from the GPAI model provider and the GPAI system provider.

Under a **sequential (chain-based) information flow architecture**, information flows instead follow the structure of the value chain. Under this approach, upstream suppliers first share information with

[11] Engler & Renda (2022); Allison et al. (n.d.). MAGF –< A Call and Proposal for Assurance Information Sharing Standards Multi-Actor Governance Framework - Value chain communications for regulatory compliance and risk mitigation of AI systems. Global Digital Foundation White Paper. Available at https://static1.squarespace.com/static/63dce129c4a0240681746067/t/65abb1652c21fc1dc8859830/1705750887433/MAGF+Framework+White+Paper.pdf

[12] Cobbe, J., Veale, M., & Singh, J. (2023). Understanding accountability in algorithmic supply chains. In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (pp. 1186-1197).

downstream suppliers, and the information is then passed on to the high-risk AI system provider. In Scenario A, this would mean that the high-risk AI system provider receives information only from the GPAI system provider, which in turn collects relevant information from the upstream GPAI model provider.

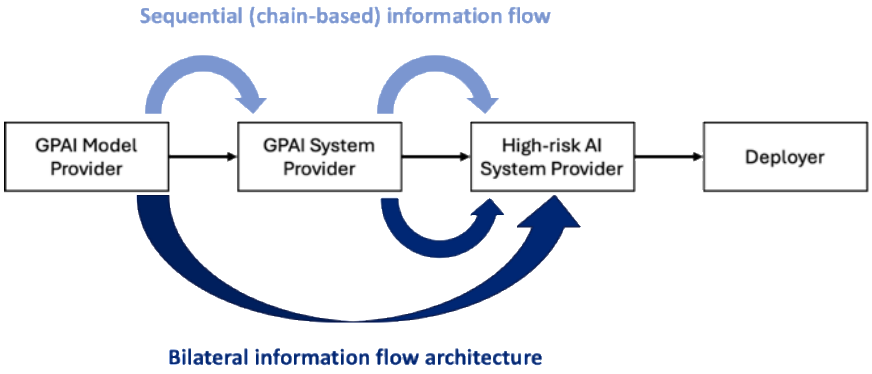

Figure 1: Bilateral and sequential (chain-based) information flow architectures, illustrated for Scenario A.

Article 25(4) AIA refers only to bilateral written agreements between the high-risk AI system provider and input suppliers. On its face, this suggests that Article 25(4) requires the GPAI model provider to share information with the high-risk AI system provider, but not with the GPAI system provider. This supports an interpretation of the AIA that favours bilateral information flows over sequential (chain-based) flows.

In practice, however, a fully bilateral architecture is difficult to implement where third-party input suppliers rely on upstream suppliers that have no direct relationship with the high-risk AI system provider. In such cases, a provider cannot conclude an agreement with a party with which it has no relationship, that it cannot identify, or that is unaware its input is being used. In Scenario A, for example, the GPAI model provider may have no direct commercial relationship with the high-risk AI system provider because it supplies only an intermediate input to the GPAI system provider. More generally, the high-risk AI system provider may not even know the identity of all relevant upstream suppliers.

This problem becomes particularly acute in scenarios with many upstream suppliers, such as Scenario C, where a GPAI system provider integrates multiple GPAI models and data sources to offer an agentic AI system. If such a system is deployed in a high-risk context (for example, to analyse job applications in HR), it may be very difficult, if not impossible, for the high-risk AI system provider to identify all relevant suppliers and conclude separate written agreements with each of them. Consequently, bilateral information flows where the high-risk AI system provider alone is responsible for collating information can place an excessive burden on these providers and potentially undermine the effectiveness of the AIA's risk-mitigation framework in more complex value chains.

Furthermore, the visibility problem also runs in the opposite direction. Upstream suppliers may not be aware that their inputs are incorporated into an AI system that is ultimately used in a high-risk domain. This is especially likely in Scenario C, where the AI system relies on function calls to APIs of external web service providers and the connection is established only at runtime.

In this context, it is noteworthy that Recital 88 AIA, in discussing responsibilities along the AI value chain, refers specifically to inputs "incorporated by the provider into the AI system." This wording suggests that a bilateral information sharing obligation may depend on the high-risk AI system provider actively and directly incorporating the relevant input into the system. Read from an implementation perspective, this may support a more sequential (chain-based) information flow architecture.

Under such an architecture, the high-risk AI system provider would conclude written agreements on information sharing, access, and assistance (for compliance with the AIA's high-risk requirements) with those input suppliers with which it directly interacts. Each of these input suppliers would, in turn, be responsible for exchanging information with its own upstream input suppliers to the extent necessary to fulfil its responsibilities towards the provider of the high-risk AI system.

This approach has several major advantages. It enables efficient information flows without inflating the number of bilateral relationships. It also makes use of actors' local knowledge about the value chain and can adapt relatively quickly to changes in the value chain structure. Importantly, it anchors the AIA's information sharing and cooperation duties in existing contractual relationships, thereby facilitating implementation while promoting effectiveness and proportionality. Finally, because it follows existing commercial relationships, it is more likely to preserve the incentives of the actors involved to participate in information sharing for risk management and mitigation.

Recommendation
The existing contractual relationships of the provider of the high-risk AI system should anchor the responsibilities set out in Article 25(4) AIA and form the basis of a sequential (chain-based) information sharing architecture.

A sequential (chain-based) information sharing architecture calls for further clarification on the duties to cooperate and share information in the upstream parts of the value chain. This is especially the case where information is ultimately destined for the downstream high-risk AI system provider, but the request for it is made not by that provider itself but by one of its direct input suppliers to its own upstream supplier. Because Article 25(4) explicitly references only the relationship between input suppliers and the provider of the high-risk AI system, it is largely silent on these upstream flows, which in practice rest on each supplier's own (back-to-back) contractual arrangements. The extent to which Article 25(4) requirements should be passed along the chain in a sequential architecture, and how, should therefore be the subject of further regulatory guidance.

3.2 Conditions for Responsibilities under Article 25(4) AIA

3.2.1 Scope of Covered Input Suppliers

For non-high-risk AI systems that later become inputs to high-risk AI systems through substantial modification or a change in intended purpose, Article 25(2) AIA provides that the initial provider's cooperation duty does not apply where that provider clearly specified that the system is not to be changed into a high-risk AI system.

Article 25(4) AIA does not contain a comparable rule. However, to avoid imposing Article 25(4) responsibilities on input suppliers that neither anticipate the use of their inputs in high-risk AI systems nor may even be aware of such integration (see, for example, Scenario C in Section 2), such responsibilities should arise only where the supplier has a direct contractual relationship with the provider of the high-risk AI system concerning that system. Otherwise, there is a risk that requirements designed specifically for high-risk AI systems could diffuse into the broader market for digital services, potentially extending far beyond the AI Act's risk-based framework.

Recommendation
To maintain the AIA's risk-based approach, Article 25(4) responsibilities should arise only for input suppliers that have a direct contractual relationship with the provider of the high-risk AI system concerning that system.

Once an input supplier enters into a direct contractual relationship with the provider of a high-risk AI system, Article 25(4)'s mandatory framing indicates that responsibilities under that provision cannot be waived by the parties once the supplier is within its scope. However, the provision still leaves room for the parties to define the scope and modalities of cooperation, as long as the arrangement enables the high-risk AI system provider to comply with the AIA. Such arrangements may differ in the information sharing mechanisms they implement, as discussed in Section 3.3.

To ensure proportionality and preserve input suppliers' freedom to conduct a business, third-party suppliers should be able to rule out the use of their inputs in high-risk AI systems where they clearly indicate this, for example in the terms of use of their service or in contractual agreements. Consistent with the scope of Article 25(4) set out above, such a declaration can serve as evidence that the supplier has not engaged in a contractual relationship concerning a high-risk AI system and therefore falls outside the provision. However, where the supplier is in an existing contractual relationship concerning the high-risk AI system and its own conduct (such as marketing its input for high-risk applications) or the surrounding circumstances show that it was aware of, and had accepted, that use, the declaration should not relieve it of its cooperation and information sharing duties under Article 25(4). This approach would preserve business freedom without allowing suppliers to evade the protective purpose of the high-risk regime through superficial declarations.

3.2.2 Scope of Legitimate Information Requests

Establishing that an input supplier falls within the scope of Article 25(4) does not mean that it must satisfy every request the high-risk AI system provider chooses to make. Specifically, Article 25(4) requires the supplier to provide the information, capabilities, technical access, and other assistance that are *necessary*, and *based on the generally acknowledged state of the art*, *to enable* the provider to comply with its obligations under the AI Act. A request therefore falls within the supplier's Article 25(4) duty only to the extent that the information sought is necessary for, and tied to, a specific obligation that the provider must fulfil in relation to the high-risk AI system. Section 4 discusses specific categories of information that may fall under this scope for individual high-risk requirements.

Requests that are not connected to this specific compliance purpose, or that demand more than the state of the art supports, fall outside the duty and may be declined. This limit matters in practice, because otherwise input suppliers could be exposed to a flood of broad or speculative requests. Responding to such requests is costly, and, especially for suppliers that serve many customers, the

cumulative burden could be significant, while legitimate requests may suffer from crowding-out and associated delay. To prevent such a situation, regulatory guidance may support implementation by clarifying which requests fall clearly outside the scope of the cooperation duties under Article 25(4), as well as the possible processes and conditions under which input suppliers can ask a high-risk AI system provider to substantiate its request.

Whether a particular request is legitimate will nonetheless depend heavily on the specific application and context. A useful substantive measure may be the influence of the input on the risk profile of the high-risk AI system: the more a given input shapes the system's risks, the stronger the justification for a request concerning it, and the wider the range of information the provider can legitimately seek. This benchmark directly aligns with the principle of proportionality, even though its precise operationalisation may present its own challenges. This is therefore an area where regulatory guidance could do valuable work, by indicating how the influence of an input on the risk profile of a high-risk AI system may be assessed and how this bears on the threshold at which an information or assistance request is considered necessary and reasonable, and therefore legitimate.

Recommendation
To support proportionate and effective implementation, regulatory guidance should clarify the limits on requests under Article 25(4) and how an input's influence on the system's risk profile bears on what may reasonably be requested.

3.3 Information Sharing Mechanisms

Information sharing mechanisms for operationalising Article 25(4) AIA must meet several **key requirements**, including:

- **Effectiveness in enabling compliance by the high-risk AI system provider**
  The information shared must be sufficient in both quantity and quality to enable the provider of the high-risk AI system to comply with its obligations under the AIA. At a minimum, the input supplier must provide all information relevant for compliance that the high-risk AI system provider cannot reasonably obtain by its own means. Because the requirements for high-risk AI systems in Chapter III AIA differ, the categories of information needed for compliance will also vary. In this context, Section 4 discusses specific AIA requirements to identify potential corresponding categories of information to be shared.
- **Protection of legitimate interests of the input supplier**
  Recital 88 emphasises that input suppliers' obligation to share information should not compromise their intellectual property rights or trade secrets. This may include source code, internal documentation, and other sensitive technical assets, such as model weights. In addition, input suppliers may regard performance information as a sensitive category of information, as its disclosure could have implications for business practices and competition. At the same time, this information may be relevant for the compliance of high-risk AI systems, in particular under Article 15 AIA on appropriate accuracy and robustness (see Section 4).13 To remain workable in practice, information sharing mechanisms under Article

[13] Schnurr, D. (2025). Effective Implementation of Requirements for High-risk AI Systems Under the AI Act: Transparency and Appropriate Accuracy. CERRE Issue Paper. Available at https://cerre.eu/wp-content/uploads/2025/02/Effective-Implementation-of-Requirements-for-High-Risk-AI-Systems-Under-the-AI-Act_FINAL-1.pdf.

25(4) AIA must balance such tensions.

- **Account for information asymmetries**
  Actors in AI value chains hold private knowledge that other parties lack. For example, providers of high-risk AI systems may have context-specific application knowledge and information about how different inputs interact within the AI system that upstream input suppliers do not possess. Both kinds of knowledge may be important for identifying what information is relevant to effective risk management and mitigation and, in turn, to compliance with the AIA. Conversely, input suppliers may hold knowledge about potential updates, changes, or emerging risks associated with their inputs that may also affect risks in the downstream AI system, but of which the high-risk AI system provider may be unaware.
- **Efficiency**
  To reduce compliance costs, information sharing mechanisms should be designed to minimise transaction costs. This calls for avoiding redundant information exchanges and for developing harmonised templates and best practices.

The listed requirements, particularly the need to protect the legitimate interests of input suppliers, may be in tension with one another. Selecting different **modes of information sharing** according to the sensitivity of the information can help reconcile these tensions.

- **Public disclosure:** General information about an input that is neither sensitive nor highly context-dependent may be made publicly available, for example on the input supplier's website or through publicly accessible documentation. Such information may include the purpose of the input, its main functions, and its general mode of operation. Public disclosure can significantly reduce transaction costs by limiting the need for repeated bilateral information exchanges with individual downstream system providers. As inputs evolve over time, it should always be clearly stated which version or versions of the input the published information covers.
- **Bilateral information exchange:** More sensitive information that cannot be publicly disclosed, as well as information that is specific to the application context of an individual high-risk AI system provider, may be shared bilaterally between the parties. This may include, for example, information on the performance of an input and the procedures used to test it. To protect confidentiality, the parties may rely on non-disclosure agreements and other appropriate safeguards. Bilateral information sharing arrangements may cover not only the categories of information to be shared, but also points of contact and communication channels through which high-risk AI system providers can raise specific questions relating to risk management and mitigation, particularly where timely exchange is important, such as in the context of post-market monitoring.
- **Information sharing with regulatory authorities or trusted third parties:** Some information held by the input supplier may be relevant for demonstrating that the supplier has taken risk management and mitigation measures that support the compliance of downstream high-risk AI systems, while at the same time being highly sensitive. This may include, for example, information on internal procedures, data sourcing, or security-relevant vulnerabilities. In such cases, the protection of the input supplier's legitimate interests may weigh against sharing the information directly with providers of high-risk AI systems. Where this is the case, the

information could instead be shared with regulatory authorities in order to support the compliance of the downstream high-risk AI system provider.[14]This would allow regulatory authorities to scrutinise the adequacy of the relevant risk management and mitigation measures while protecting sensitive information from broader disclosure.

Alternatively, the information could be shared with a trusted private third party that verifies it while ensuring its confidentiality. This may be particularly suitable where some parties do not regard regulatory authorities as the most appropriate custodians of confidential or competitively sensitive information.

For neither option does Article 25(4) AIA currently provide a stand-alone legal basis, since the provision frames the obligation as running to the provider of the high-risk AI system, to enable that provider's own compliance. The trusted-third-party channel may be accommodated through contractual design, as the third party mediates assurance that ultimately reaches the provider, and the voluntary model contractual terms could usefully standardise such arrangements. Routing information to regulatory authorities in place of the provider, by contrast, alters to whom the duty is owed. While this is more far-reaching, it could be enabled through regulatory guidance specifying when and how such a channel discharges the input supplier's responsibilities under Article 25(4).

A further key dimension of information sharing mechanisms is which party initiates the transfer of relevant information. Accordingly, Allison et al. distinguish between two mechanisms:[15]

- **Information push**: The originator of the information initiates the exchange and provides the information proactively to the other party. In the context of Article 25(4) AIA, this means that the third-party input supplier provides information proactively to the high-risk AI system provider. A push mechanism is particularly suitable where the receiving party may not know that the information exists or that it is relevant. This may be the case, for example, where the input is updated or otherwise changed by the supplier in a way that could affect risk prevention or mitigation. At the same time, such a mechanism requires the input supplier to maintain communication channels with all downstream providers of high-risk AI systems.
- **Information pull:** The party that intends to use the information initiates the exchange by requesting it from the originator. In the context of Article 25(4) AIA, this means that the high-risk AI system provider requests information from the third-party input supplier. A pull mechanism is particularly suitable where the downstream provider has context-specific information needs, or where downstream activities (such as system analysis, incident investigation, or corrective measures) create a need for additional upstream information.

To meet the key requirements for information sharing under Article 25(4) AIA and to accommodate potential tensions among them, agreements between input suppliers and providers of high-risk AI systems should combine different modes of information sharing with both push and pull mechanisms. Although the specific configuration of these mechanisms will need to be tailored to the characteristics of the input supplier, the high-risk AI system provider, and the particular application context of the

[14] Such information sharing with regulatory authorities would go beyond the direct exchange of information between input suppliers and providers of high-risk AI systems envisaged by Article 25(4) and Recital 88, but it provides a means of balancing competing requirements.
[15] Allison et al. (n.d.). Global Digital Foundation White Paper.

high-risk AI system, this combination provides a useful general template for operationalising information sharing under Article 25(4) AIA.

Recommendation

Written agreements between input suppliers and providers of high-risk AI systems should distinguish between different modes of information sharing (public disclosure, bilateral information exchange, and information sharing with regulatory authorities or trusted third parties) and combine push and pull mechanisms as key dimensions for operationalising responsibilities under Article 25(4) AIA. The AI Office may develop and recommend voluntary model contract terms that map different categories of information onto these information sharing mechanisms for each specific obligation applicable to high-risk AI systems.

4. Information Sharing for Specific High-Risk AI System Requirements

This section turns to specific requirements for high-risk AI systems and examines how information sharing can support their implementation. For selected requirements, it aims to identify categories of information that can be mapped onto the information sharing mechanisms presented in Section 3. It addresses the main question: **What information and assistance should input suppliers share with downstream AI system providers to enable compliance with specific AIA requirements for high-risk AI systems?**

The section proposes categories of information that suppliers of inputs to high-risk AI systems may need to share in relation to the selected AIA requirements discussed below. The listed categories of information, shown in italics, represent initial proposals that may serve as a basis for further discussion and refinement. In particular, the proposed categories of information need to be assessed with regard to their precise scope and limits, their applicability in different use cases and risk domains, as well as which of the mechanisms presented in Section 3 are suitable for sharing them.

4.1 Article 9 AIA: Risk management system

Providers of high-risk AI systems must establish, document, implement and maintain a continuous, iterative risk management process, spanning the entire lifecycle of the AI system. The risk management process comprises the identification, analysis, estimation, evaluation, and mitigation of known and foreseeable risks to health, safety or fundamental rights arising when the system is used in accordance with its intended purpose, or under conditions of foreseeable misuse.

To comply with Article 9 AIA, providers of high-risk AI systems must assess how the integration of upstream inputs affects the system's risk profile. Their risk management system should therefore be informed by what risk identification, testing, and mitigation measures may have been undertaken at upstream levels of the value chain.

Accordingly, input suppliers may provide information on:

- the intended purpose of the input,
- whether risks to health, safety, or fundamental rights have been identified and, if so, which types of risks,
- whether risk-mitigation measures have been implemented and, if so, which measures,
- the scope of any testing carried out to identify and mitigate risks.

4.2 Article 10 AIA: Data Governance

High-risk AI systems that use data for training models are subject to stringent data governance and data quality requirements intended to support reliability, safety, and compliance with EU law. Training, validation, and testing datasets must be relevant, representative, free of errors and complete to the greatest extent possible. Providers must also put in place data governance and management practices appropriate to the intended purpose of the high-risk AI system, including in relation to data

collection, preparation and processing, data quality assessment, bias detection and correction, and the identification of relevant data gaps. Furthermore, providers must retain records concerning dataset composition, data curation practices, and the results of bias testing.

To support compliance with Article 10 AIA, input suppliers that provide or integrate data should inform the provider of the high-risk AI system about the type and general content of that data, its quality, and the data governance and management practices applied at the supplier level. Where input suppliers use data to train models, these categories of information should, where relevant, be provided separately for the training, validation, and testing datasets used.

Accordingly, where input suppliers provide or integrate data, they may provide information on:

- the type and general content of the data,
- whether data quality checks, quality improvement measures, or bias correction methods have been implemented and, if so, which ones,
- whether data governance and data management practices have been implemented and, if so, which practices, with reference to the different stages of the data lifecycle,
- where available, relevant data quality metrics, records on dataset composition, identified data gaps, and the results of bias assessments,
- where data has been used to train an AI model at the input level, the information above in relation to the training, validation, and testing datasets concerned.

4.3 Article 15 AIA: Accuracy, robustness, and cybersecurity

Article 15(1) AIA requires that "[h]igh-risk AI systems shall be designed and developed in such a way that they achieve an appropriate level of accuracy, robustness and cybersecurity, and that they perform consistently in those respects throughout their lifecycle." The provision therefore makes clear that high-risk AI systems must meet both appropriate performance standards and appropriate non-functional requirements.[16] Before a high-risk AI system is placed on the market or put into service, system providers are required to provide detailed information about its capabilities and limitations in performance as part of the mandatory technical documentation. This information should include "the degrees of accuracy for specific persons or groups of persons on which the system is intended to be used and the overall expected level of accuracy in relation to its intended purpose". The technical documentation should also provide information on the validation and testing procedures as well as the metrics used to measure performance. In addition, high-risk AI systems must be resilient against errors, faults or inconsistencies that may occur within the system or the environment in which the system operates (*robustness*), as well as against attempts by unauthorised third parties to alter their use, outputs or performance by exploiting system vulnerabilities (*cybersecurity*).

In principle, performance testing can be carried out on the final AI system independently. However, information on the performance and testing of upstream inputs can facilitate that assessment and help the provider identify potential triggers and causes of performance limitations in the system as a

[16] Schnurr (2025). CERRE Issue Paper.

whole. Similarly, information shared by input suppliers on the robustness and cybersecurity of inputs can support the identification and mitigation of vulnerabilities that might otherwise remain undetected, even where the final AI system is itself subject to testing. At the same time, these categories of information are particularly sensitive, as they could be exploited by malicious actors. Because information on vulnerabilities and possible remedies may emerge after the AI system has been deployed, such information should be shared through agreed channels and in a timely manner (see also the discussion below on post-market monitoring).

Accordingly, input suppliers may provide information on:

- the average performance of the input across relevant performance dimensions (such as accuracy or timeliness),
- the extent to which performance may vary across different inputs, over time, and across relevant deployment or application contexts,
- potential triggers or events that may lead to performance deterioration,
- the procedures used to measure performance (possibly including relevant metrics) and the scope of testing undertaken,
- possible errors, faults, or inconsistencies that may arise within the input or in the environment in which it is intended to operate,
- possible vulnerabilities and any cybersecurity safeguards that have been implemented,
- relevant updates and patches when they become available.

4.4 Article 72 AIA: Post-market monitoring

Under Article 72(1) AIA, providers of high-risk AI systems must establish a post-market monitoring system, to collect and review experience gained from the use of their high-risk AI system after it was placed on the market or put into service, in order to identify "any need to immediately apply any necessary corrective or preventive actions" (Article 3(25) AIA).[17] Its main purpose is the active and systematic collection, documentation, and analysis of data to allow the provider to evaluate the performance of high-risk AI systems through their lifetime and the continuous compliance with the requirements for high-risk AI systems (Article 72(2) and Recital 155 AIA). The AI Act recognises that relevant information may be available only from other actors in the value chain and notes that such data may be "provided by deployers" or "collected through other sources" (Article 72(2) AIA).

Effective post-market monitoring therefore requires communication channels between input suppliers and providers of high-risk AI systems. These channels may support:

- the sharing by input suppliers of information on significant changes to the intended purpose of an input, identified risks, performance limitations, or security vulnerabilities,
- requests for information from providers of high-risk AI systems where risks, performance limitations, or security vulnerabilities are identified after deployment of the AI system,
- the timely distribution by input suppliers of relevant updates and patches when they become

[17] Meyers, Schnurr & Larouche (2025). CERRE Issue Paper.

available,

- where the input is an AI system or model that continues to learn after being put into service, updates from the input supplier on newly identified risks and performance limitations.

4.5 Article 73 AIA: Serious incident reporting

Article 73(1) AIA requires providers of high-risk AI systems to report any serious incident, as defined in Article 3(49) AIA, to the market surveillance authorities of the Member State in which the incident occurred. Such reporting is only the first step, as providers must subsequently perform "the necessary investigations in relation to the serious incident and the AI system concerned", including corrective action and a risk assessment of the incident (Article 73(6) AIA).

More broadly, the market surveillance regime of Regulation 2019/1020[18] applies to all AI systems falling within the scope of the AI Act, irrespective of which tier of the risk-based pyramid they fall under, i.e. whether they are high-risk AI systems or not (Article 74(1) AIA).[19] While that Regulation does not require the introduction of monitoring systems, it nevertheless puts system providers under a duty to inform authorities if they have reason to believe that their AI system presents a risk to health, safety or fundamental rights.[20]

While Article 73 AIA specifies that providers of high-risk AI systems must cooperate with the competent authorities, and where relevant the notified body concerned, it does not establish an equivalent duty of cooperation among actors in the value chain for the purposes of such investigations. Accordingly, any responsibility of input suppliers to support the identification, reporting, investigation, or mitigation of serious incidents arises to the extent that it is covered by Article 25(4) AIA.

Information sharing between input suppliers and providers of high-risk AI systems in relation to serious incidents should cover similar categories to those relevant for post-market monitoring, but it will often require more immediate exchanges of information. In addition, it may be necessary to share more detailed information with the provider of the high-risk AI system or regulatory authorities to identify and report serious incidents as well as to develop corrective measures. Information exchanges may therefore support:

- requests for information from providers of high-risk AI systems where serious incidents are suspected or identified after deployment of the AI system,
- the sharing by input suppliers of information on serious incidents connected to the input, including incidents identified in connection with other AI systems using the same input,
- the timely distribution by input suppliers of relevant updates, corrective measures or recall decisions once a serious incident connected to the input has been identified.

As discussed in more detail in a recent CERRE issue paper, new institutions and mechanisms for post-market information sharing on serious incidents and risks, extending beyond bilateral relationships between providers of high-risk AI systems and their input suppliers, could become an important

[18] Regulation 2019/1020 on market surveillance and compliance of products [2019] OJ L 169/1, as subsequently amended.
[19] Meyers, Schnurr & Larouche (2025). CERRE Issue Paper.
[20] Regulation 2019/1020, Article 4(3)(c), combined with AI Act, Article 79 (1).

element of effective AI risk mitigation.[21] This need is reinforced by the likelihood that serious incidents and risks will also arise outside the domains classified as high-risk under the AI Act.

Industry-driven initiatives could promote vertical and horizontal information sharing in AI industries by establishing a common database or information hub for AI-related incidents, following models that exist in other sectors such as aviation or cybersecurity. Such an institution could build on existing databases, which currently rely largely on user-contributed reports. Promoting standardised reporting schemas, while implementing safeguards for trade secrets, security, and competition law, could further improve the effectiveness of such an institution and its associated information sharing mechanisms.

4.6 Article 12 AIA: Record-keeping

Under Article 12(1) AIA, high-risk AI systems must be designed to allow for the automatic recording of events (logs) throughout their lifetime. This record-keeping must enable the recording of events relevant to identifying situations in which the system may present a risk within the meaning of Article 79(1) AIA or undergo a substantial modification. It must also facilitate post-market monitoring by the provider and monitoring of the system's operation by deployers.

Article 19(1) AIA further requires providers of high-risk AI systems to retain the logs automatically generated by those systems, to the extent that such logs are under their control. This may create practical challenges where the system depends on inputs accessed through APIs, in which case the provider typically has no access to logs generated at the level of the input. Similar issues may arise where the provider relies on inputs supplied as cloud computing services and has no, or only partial, access to the relevant logging capabilities.

Accordingly, information sharing between input suppliers and providers of high-risk AI systems in relation to record-keeping may support:

- transparency on whether, and to what extent, logs are generated at the level of the input in relation to the operation of the high-risk AI system,
- the sharing of relevant logs, where available and subject to appropriate safeguards for intellectual property rights and trade secrets, where the provider of the high-risk AI system suspects that the system may present a risk or requires the data for post-market monitoring,
- access by the provider of the high-risk AI system to relevant logs where the input is supplied as a cloud computing service.

4.7 Additional high-risk AI system obligations that may require information sharing

Further provisions applicable to high-risk AI systems that may require information sharing by input suppliers, but are not specifically addressed in this paper, include technical documentation (Article 11 AIA), transparency and the provision of information to deployers (Article 13 AIA), human oversight (Article 14 AIA), and other elements relating to corrective actions and duty of information (Article 20 AIA). For the implementation of these provisions, similar mappings of relevant categories of

[21] Meyers, Schnurr & Larouche (2025). CERRE Issue Paper.

information to be shared between input suppliers and providers of high-risk AI systems should be developed.

5. Recommendations for Implementing the AI Act's Information Sharing Provisions

This paper has explored how cooperation and information sharing along the AI value chain could operate under Article 25(4) AIA in relation to high-risk AI systems. It shows that several key open questions arise from the provision that can create considerable uncertainty in practice, especially where AI value chains are complex and quickly evolving. In particular, there is a need for clarity about who falls within the scope of Article 25(4) and how information should flow through value chains. Because the information sharing and cooperation duties connect the high-risk and non-high-risk tiers of the AI Act, as well as broader parts of the digital economy, these duties must be interpreted proportionately, so that the risk-based approach of the AIA is maintained. Finally, practical and harmonised guidance, drawing on templates and concrete use cases, could significantly reduce this uncertainty and avoid fragmentation across sectors and Member States. To these ends, the paper has made several specific recommendations in the individual sections, and concludes with the following overarching recommendations for governing information sharing and cooperation along the AI value chain under the AI Act:

Recommendation 1: *To preserve the AI Act's risk-based approach, the responsibilities under Article 25(4) should arise only for input suppliers that have a direct contractual relationship with the provider of the high-risk AI system concerning that system. Input suppliers should be able to rule out the use of their inputs in high-risk AI systems by clearly indicating this. However, such a declaration should not relieve a supplier of its responsibilities where it is in such a relationship and its conduct or the surrounding circumstances show that it was aware of, and had accepted, that use.*

This keeps the regime proportionate by tying responsibility to the supplier's role and conduct. Suppliers that do not engage with the high-risk domain are not drawn in merely because an input reaches a high-risk AI system somewhere downstream, while those that participate in high-risk value chains cannot use a superficial declaration to evade the protective purpose of the high-risk regime.

Recommendation 2: *Implementation should account for the complexity of AI value chains and information flows. Accordingly, a sequential (chain-based) information flow architecture is likely more suitable than a bilateral information flow architecture, for both feasibility and effectiveness. Moreover, regulatory guidance should set out criteria that help providers distinguish legitimate from illegitimate information and cooperation requests.*

Even relatively simple AI applications involve multiple actors and information flows, and a fully bilateral architecture quickly becomes impractical where upstream suppliers have no direct relationship with the high-risk AI system provider. In particular, the requirement of a written agreement appears infeasible for relationships that only manifest at runtime of a system, for example, when input services are accessed ad hoc through APIs. Anchoring information flows in the chain of existing contractual relationships is therefore both more feasible and more effective. To promote effectiveness and proportionality, guidance should also help providers classify which requests are legitimate and which fall outside the duty, so that input suppliers are not exposed to a flood of broad or speculative requests. The need is immediate: downstream AI providers are already seeking clarity from their suppliers on their high-risk obligations, a difficulty compounded by the delay in harmonised

standards envisaged under the AIA. To make these criteria workable, regulatory guidance should be grounded in real-world examples of AI value chains, of the kind illustrated through the stylised use cases in this paper.

Recommendation 3: *Cooperation under Article 25(4) should build on parties' existing contractual relationships and allow parties to allocate responsibilities through mutual agreement. This accommodates complex and diverse value chain scenarios and promotes flexibility. Such agreements should specify each party's duties and anticipate failure scenarios and how they are to be resolved. Because contractual arrangements alone may not resolve every issue, especially in cases of a breakdown or an incident, regulatory guidance may help establish additional mechanisms to backstop them.*

Existing commercial relationships are the natural basis for operationalising Article 25(4): they already connect the relevant parties, and building on them avoids creating a parallel web of relationships and obligations. Contractual agreements may not only allocate information sharing and cooperation duties between input suppliers and providers of high-risk AI systems but also provide for failure scenarios, for example where a supplier cannot or will not furnish needed information, and the means of resolving them. There are market incentives for high-risk AI system providers to contract with upstream suppliers that can support their compliance, which should encourage such arrangements to form. However, relying on contractual freedom and market-based arrangements has limits: power and information asymmetries between the parties may prevent them from resolving every issue by agreement, and voluntary cooperation may break down precisely when it matters most, once a serious incident occurs or a risk materialises. Guidance, and in some cases mechanisms beyond voluntary contractual agreements, may therefore be needed to ensure that cooperation does not fail at the point of greatest need. At the same time, the AI Office and other authorities need to anticipate that the interpretations of the AIA's rules and defaults they set through regulation will affect the positions of the contracting parties, and therefore the agreements reached as outcomes.

Recommendation 4: *To balance the divergent key requirements for information sharing and the interests of the involved parties, agreements under Article 25(4) may draw on a combination of information sharing mechanisms, matched to the sensitivity of the information: public disclosure for general information, bilateral exchange for context-specific or moderately sensitive information, and channels to regulatory authorities or other trusted parties for highly sensitive information. Information should also be able to flow upstream as well as downstream, and agreements may combine push and pull mechanisms. Where available and suitable, regulatory guidance may refer to harmonised templates and international standards.*

No single information sharing mechanism suits every category of information. General information that is neither sensitive nor highly context-dependent can be made publicly available, reducing the need for repeated bilateral exchanges. Information specific to a particular high-risk AI system, or too sensitive to publish, can be exchanged bilaterally under appropriate confidentiality safeguards. Finally, highly sensitive information, whose disclosure to the provider might harm the supplier's legitimate interests, may instead be shared through a channel to a regulatory authority or a trusted private third party that verifies it while preserving its confidentiality. As set out in Section 3.3, however, a channel to a regulatory authority or a trusted private third party may require additional

regulatory guidance, as Article 25(4) frames the information sharing duty as running directly to the provider.

Information flows should also not be assumed to run only downstream: upstream sharing can also be important, for example to alert an input supplier to incidents caused or influenced by its input, including those observed in other AI systems relying on the same input. Across these mechanisms, harmonised templates and conformity with international standards for information requests could, where available and suitable, improve effectiveness and help avoid fragmentation.

Recommendation 5: *The AI Office should provide guidance that reduces uncertainty about the scope and implementation of the information sharing and cooperation duties. Such guidance should be harmonised across sectors and Member States to avoid fragmentation and should be grounded in real-world use cases reflecting today's AI value chains. To support implementation, the guidance should clarify what categories of information are expected to be shared under the individual high-risk requirements. To preserve flexibility, it may offer concrete routes to compliance without mandating them as the only permissible ones.*

With the harmonised standards envisaged under the AIA delayed, guidance from the AI Office is particularly important in the near term, as it can clarify the scope and implementation of the duties. Guidance that maps categories of information to individual high-risk requirements would be especially valuable, provided it retains the flexibility that the dynamism of AI value chains demands. Such guidance could offer concrete routes to compliance without mandating them as the only permissible ones, thereby reducing uncertainty while leaving room for context-specific solutions. In this context, it could also make use of defined thresholds where reasonable and sufficiently broadly applicable, accepting that even imprecise proxies may be preferable to open-ended uncertainty.

Where the input is a GPAI model, the model documentation template developed under the General-Purpose AI Code of Practice to address the downstream transparency requirements under Article 53(1)(b) AIA offers a natural starting point for identifying what model information providers of high-risk AI systems may seek from their model suppliers. Because that documentation applies to all GPAI models, including those not used in high-risk domains, additional information may be required for specific high-risk obligations. To be useful in practice, guidance should build on harmonised templates, as recommended above, and be illustrated through real-world value-chain use cases, consistent with the AI Office's practice in its earlier guidance.

About CERRE

Providing high quality studies and dissemination activities, the Centre on Regulation in Europe (CERRE) is a not-for-profit think tank. It promotes robust and consistent regulation in Europe's network, digital industry, and service sectors. CERRE's members are regulatory authorities and companies operating in these sectors, as well as universities.

CERRE's added value is based on:

- its original, multidisciplinary and cross-sector approach covering a variety of markets, e.g., energy, mobility, sustainability, tech, media, telecom, etc.;
- the widely acknowledged academic credentials and policy experience of its research team and associated staff members;
- its scientific independence and impartiality; and,
- the direct relevance and timeliness of its contributions to the policy and regulatory development process impacting network industry players and the markets for their goods and services.

CERRE's activities include contributions to the development of norms, standards, and policy recommendations related to the regulation of service providers, to the specification of market rules and to improvements in the management of infrastructure in a changing political, economic, technological, and social environment. CERRE's work also aims to clarify the respective roles of market operators, governments, and regulatory authorities, as well as contribute to the enhancement of those organisations' expertise in addressing regulatory issues of relevance to their activities.

About the Author

Daniel Schnurr is a CERRE Research Fellow and a Professor of Information Systems at the University of Regensburg, where he holds the Chair of Machine Learning and Uncertainty Quantification. Previously, he led the Data Policies research group at the University of Passau. He received his Ph.D. in Information Systems from the Karlsruhe Institute of Technology, where he also completed his B.Sc. and M.Sc. in Information Engineering and Management. Daniel Schnurr has published in leading journals in Information Systems and Economics on competition and data sharing in digital markets, regulation of data-driven market power, and competition and cooperation in telecommunications markets. His current research focuses on the role of artificial intelligence in competition, privacy and data sharing in digital markets as well as regulation of AI, cloud computing and the data economy.

cerre

Avenue Louise 475 (box 10)
1050 Brussels, Belgium
+32 2 230 83 60
info@cerre.eu
www.cerre.eu

Centre on Regulation in Europe (CERRE)
CERRE Think Tank
CERRE Think Tank