06.2026 Article 25(4) AIA

From Policy to Practice

Prototyping the EU AI Act's value chain requirement

Kenniscentrum Data & Maatschappij

Artificiële Intelligentie Vlaanderen

About the Knowledge Centre Data & Society

[1] *See Knowledge Centre Data & Society.*

The Knowledge Centre Data & Society (KCDS) is the central hub in Flanders for the legal, social and ethical aspects of data-driven and AI applications.

The KCDS brings together knowledge and experience on this topic tailored to industry, policy, civil society and the general public.

Specifically, our objectives include:

- → **Disseminating information and knowledge** on the legal, social and ethical aspects of data-driven applications and AI. All publications are made publicly available and aim to foster a positive and proactive relationship between these innovations and our society.
- → **Promoting structural initiatives** that strengthen vision development and valorise the social and economic opportunities of data-driven applications and AI among governments, industry and other social actors.
- → **Stimulating public awareness and debate** on the benefits and drawbacks as well as the legal, social and ethical aspects of data-driven applications and AI, in all layers of society.
- → **Building and supporting a network and learning environment** for stakeholders and strengthening collaboration between different policy levels and actors.
- → **Contributing to the development of legal frameworks and guidelines** on the use and framing of data-driven applications and AI for policymakers, businesses, organisations and employees.

The policy prototyping project is one of our activities to achieve these objectives.

Please visit our website[1] for more information about the KCDS, our objectives and our activities.

Overview

Executive summary

1

Article 25(4) of the European AI Act describes the value chain management for high-risk AI systems. This article requires providers of such systems and third parties that supply AI systems, tools, services, components or processes used or integrated into a high-risk AI system to enter into a written agreement. This agreement must set out the necessary information, capabilities, technical access and other assistance in line with the generally acknowledged state of the art, in order to enable the provider to comply with its obligations under the AI Act.

For this policy prototyping project, we worked in **several phases**. In an initial substantive phase, we organised a design workshop in which we guided the participants to jointly draft prototyping clauses based on various use cases. In a later phase, we presented these prototyping clauses to experts and asked for their feedback in interviews.

This input has been incorporated into this report in the form of **best practices and policy recommendations**. The policy prototyping clauses and the experts' feedback have been integrated into the text of this report. The participants of the design workshop did not take part in the expert interviews, and vice versa.

The report starts with an introduction to policy prototyping (section 4) and elaborates on the various phases of this project (section 5). In the following section (section 6), we provide an overview of Article 25(4) of the AI Act. Next, the prototyping clauses from the design workshop are described, and the additional input from the experts is also discussed (section 7). Following this, the next section provides an overview of the experts detailed legal feedback on Article 25(4) of the AI Act (section 8). The final section of the report provides a conclusion that summarises the key findings (section 9).

Key findings related to prototype compliance documents and Article 25(4) of the AI Act

Below, we highlight some of the findings in relation to the prototype compliance documents and the text of Article 25(4), based on participant feedback.

1. **Qualification of the parties**: the written agreement between provider and third-party supplier should take into account the allocation of roles that applies in practice. It is equally important to consider other parties that may be involved in the value chain.

2. **Defining key terms**: the terms 'the necessary information, capabilities, technical access and other assistance' require clearer definition. Reviewers considered the terms to be overly vague and intangible, leading to a lack of understandability and a need for clarification. These could be specified by reference to Articles 11, 13 and 14 of the AI Act.

3. **Existing legal instruments**: Article 25(4) is considered to be an 'open concept' article. This flexibility makes practical implementation challenging as some of the terms such as 'capabilities, necessary information and technical access' are vague and require interpretation. Sector-specific guidelines and voluntary terms are essential to provide clear benchmarks for compliance, as providers and third-party suppliers otherwise have no way of measuring whether they sufficiently meet the requirements, nor do they have the legal certainty to do so. At the same time, however, implementation is also facilitated by the existence of already established legal frameworks or contracts that can help clarify those terms, such as support contracts and service agreements, etc. These frameworks can serve as guidelines, particularly for SMEs (small and medium-sized enterprises). Nevertheless, it is important to strike a balance between the provision, additional voluntary terms, and sector-specific guidelines, and existing legal instruments and benchmarks.

4. **Language & structure**: the language of the compliance documents must be tailored to the target audience(s) and the specific sector. In addition, the written agreement should incorporate further appendices, including, for example, an overview of the relevant technical documentation.

Introduction

2

2.1 Introduction to policy prototyping

Policy prototyping is an **alternative approach to policymaking**, comparable to product or beta testing. It can be understood as a form of user-centred policy design or applying the design thinking methodology to the legislative or policymaking process. Policy prototyping should enable policymakers to map the effects, strengths and limitations of a proposed policy, leading to more effective and evidence-based policymaking while avoiding the societal costs of 'bad policy' negatively impacting stakeholders. A policy prototyping project typically consists of multiple phases:

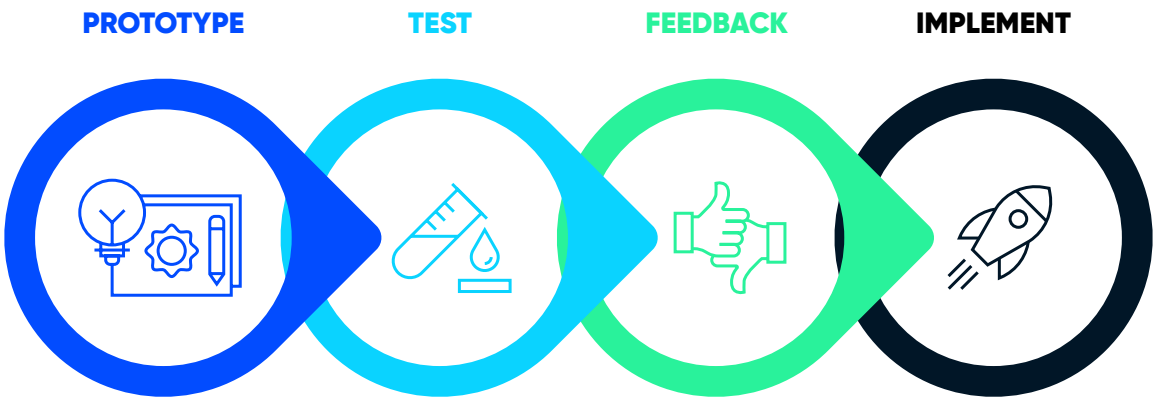

- → **Prototype**: prototyping involves the creation of basic models or designs for a machine or other product to test an idea or a concept in practice. In this context, prototyping entails drafting a new policy or law. Such prototypes can be elaborate or minimal, allowing the testing of specific features to find out 'what works' through several iterations.
- → **Test**: a group of stakeholders performs a mock compliance exercise and implements the envisaged legal requirements. In our policy prototyping project this group of stakeholders took part in the design workshop.
- → **Feedback**: participants provide feedback on the mock implementation of the policy prototype. During our policy prototyping project, we were joined by reviewers with expertise in a variety of fields related to AI, intellectual property (IP) and IT law, etc.
- → **Implement**: this feedback is used to evaluate if the law is effective and 'fit for purpose' and to complete and/or amend it accordingly, issue additional guidance and highlight ambiguities, etc.

In summary, using this approach, policymakers and stakeholders are able to create **tangible and practical prototypes** of proposed policies and related compliance documents. These prototypes allow them to **test and refine** the policy measures before committing to a full-scale implementation.

Policy prototyping can help identify potential gaps, challenges or unintended consequences at an early stage of the policymaking process. It gives policymakers the opportunity to make **necessary adjustments and improvements** to the policy, and allows stakeholders to prepare for future policy. In essence, policy prototyping may bridge the gap between policy design and actual implementation, enhancing the effectiveness, feasibility and acceptance of policies while minimising the risk of unanticipated policy mistakes or failures.

Policy prototyping projects should also consider some possible concerns for which they should ensure transparency or accountability. More specifically, the group of participants involved in a project should ideally reflect the diverse group of stakeholders affected by the envisaged policy, while public transparency regarding the participants also needs to be ensured. Additionally, policy prototyping projects will generally be conducted with small test groups. This may give casuistic results, reducing their representativity and scalability, as they may not be applicable on the large scale to which regulation usually applies. In section 5, we will explain in more detail how we applied this approach (including the concerns) in the policy prototyping project, which is the subject of this report.[2]

[2] *For more information on policy prototyping, see: B. Benichou, T. Gils and K. Vranckaert, Design thinking in the legislative process: the key to useable legislation?, April 2021. See also: T. Gils, K. Vranckaert and B. Benichou, Exploring Policy Prototyping – Some Initial Remarks, July 2021.*

2.2
Policy prototyping at the KCDS

Policy prototyping has been a focal point in the work of the KCDS for several years now.

Firstly, in 2023, we carried out a policy prototyping project that focused specifically on the EU AI Act's **transparency requirements**. The project concerned both transparency requirements for high-risk AI systems (Article 13 AI Act) and 'certain AI systems' such as AI-generated/ deep fake content (Article 50 AI Act). The main findings of this policy prototyping project can be found in our report 'From Policy to Practice: Prototyping The EU AI Act's Transparency Requirements'.[3]

Secondly, in 2025, we implemented a policy prototyping project on **the human oversight requirements** (Article 14 AI Act) that intended to test these requirements for high-risk systems. The main findings of this specific project can be found in our report 'From Policy to Practice: Prototyping The EU AI Act's Human Oversight requirements'.[4]

This current project was launched in 2025 and conducted in 2026, and focuses on Article 25(4) AI Act.

The following objectives have been pursued as part of this project:

1. A **detailed examination** of the envisaged value chain requirements of Article 25(4).
2. The **creation of operational guidance**, including prototype compliance documents for high-risk AI systems.
3. **Feedback gathering** on the value chain requirements of Article 25(4) and their practicability, desirability and feasibility through the reviewers.
4. Provision of our **findings and lessons learned** to policymakers and other stakeholders.

[3] *T. Gils, F. Heymans and W. Ooms (Knowledge Centre Data & Society), From Policy to Practice: Prototyping The EU AI Act's Transparency Requirements, January 2024.*

[4] *W. Ooms, L. Cools, T. Gils and F. Heymans (Knowledge Centre Data & Society), From Policy To Practice: Prototyping The EU AI Act's Human Oversight Requirements, March 2025.*

Policy prototyping: methodology and process

3

In this section we will take a closer look at the methodology used in this policy prototyping project. We believe that this is necessary in order to interpret and use the results of this project correctly.

The policy prototyping project described in this report **was launched in October 2025**, building on the initial phase in which the topic for the policy prototyping project was selected (i.e. value chain management for high-risk AI systems). A call for participants was subsequently issued and interested stakeholders were identified. This group was invited to a **design workshop** in Brussels **in February 2026**. During this workshop, attendees collaborated in small groups to draft model clauses based on different use cases, as required under Article 25(4) of the EU AI Act. Following the design workshop, participants were given the opportunity to continue working on the prototyping clauses, if necessary. Those prototyping clauses are included in the annex of this report.

The prototyping clauses and Article 25(4) of the AI Act were then reviewed in detail using feedback obtained from **experts** through qualitative online or in-person interviews. This report summarises the findings of those interviews. The visual below shows how our phases align with the theoretical phases mentioned in section 4.1.

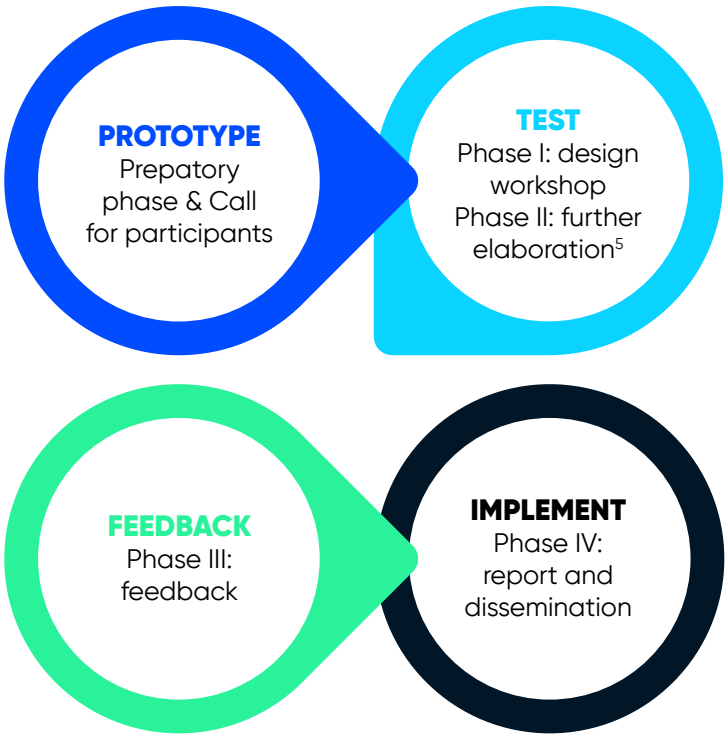

3.1

Preparatory phase: decision on legislative framework and practical considerations

The selection of Article 25(4) as the next policy prototyping project was based on various factors, namely our **own assessment of interesting topics** as well as **stakeholder input**.

Firstly, we drew up a **shortlist** of potential topics for prototyping, including various provisions from the AI Act. We then presented this list to various KCDS stakeholders and asked them to indicate their preferences. The feedback we received from the stakeholders clearly showed that Article 25(4) plays a crucial role in the implementation of AI systems. We therefore decided to focus this current prototyping project on the requirements regarding value chain management arising from Article 25(4) of the AI Act.

In addition, our choice for this topic relies on Article 25(4)'s emphasis on the importance of the **value chain lifecycle of AI**. Similar to other products, AI systems are products that must navigate the entire value chain. The responsibilities linked to the value chain start with the provider, who is the primary point of accountability, and end with the deployer, who plays a key role due to their close relationship with the end users. Another key player in this lifecycle is the third-party supplier. While this role is not explicitly defined in the AI Act, Article 25(4) contains a responsibility for the party that supplies certain elements of an AI system, including but not limited to the hardware and/or software. The relationship between the third-party supplier and the provider is regulated by the AI Act, which stipulates the requirement for a written agreement to be drawn up between the two parties. However, this provision could be regarded as open-ended and potentially vague. Although this issue is partly addressed by the involvement of the AI Office, which is responsible for drafting voluntary model terms, these have not yet been published. It is therefore crucial that the KCDS focuses its efforts on this provision so that it can provide policymakers with guidelines that are rooted in practice.

As with our other projects, budgetary and logistical considerations shaped our approach, which involved prioritising inclusive and practical engagement while maintaining a clear focus on the core obligations of Article 25(4). Although we welcomed international participants at all stages of the project, we could not reimburse international travel expenses. We relied on the voluntary commitment of participants and did not pay anyone for their participation. In the following section, we will outline the methodology and participant contributions in further detail.

3.2
Call for participants

[5] *Knowledge Centre Data & Society, Policy prototyping 2026: call for participants.*

The KCDS attaches great importance to ensuring that our group of project participants is diverse and representative. For this reason, similar to our previous policy prototyping projects, we combined an **open call** for participants with **targeted invitations** to organisations or stakeholders whom we believed could make a valuable contribution to our project.

The open call for participants contained the following information[5]:

→ Firstly, we looked for **providers and suppliers** in the AI value chain. Participating providers and suppliers could have their business serve as a use case for the prototyping project. Concretely, their high-risk AI system, or the tools, services or components they supplied, were to be used as the basis for prototype contractual clauses that comply with the AI Act requirements.

→ Secondly, we aimed to include **legal experts** who could help shape prototypes of the contractual clauses. They could provide feedback on the AI Act's requirements.

In terms of time investment, we estimated that participants would spend about two to three working days in total (a full day at the design workshop plus further elaboration on the prototypes).

3.3
Phase I: design workshop

[6] *Knowledge Centre Data & Society, KDM Policy Prototyping AI Act Supply Chain Management Questionnaire.*

To ensure meaningful personal interaction, we organised **a legal design workshop**. Eight participants and four facilitators worked together for an entire day in three different groups to shape first versions of different prototype compliance documents. Each group focused on applying the Article 25(4) requirements to one specific use case. The use cases were selected based on input from both the team and providers, suppliers or others who had submitted their use case through a form.[6] Based on feasibility for further development during the session, three different use cases were selected that centred on the **human resources (HR), insurance and medical industries**.

This approach ensured that our prototyping exercise was grounded in potential use cases from various relevant sectors. We did not expect participants to share technical or sensitive

details in relation to their use cases. It should be underlined that we looked for high-risk AI systems as defined by the AI Act. If there was any uncertainty about whether an AI system qualified as high-risk in a specific use case, participants were instructed to assume that the use case was high-risk for the purpose of the workshop.

The use cases are explained in depth below (section 7). All three groups worked on prototype documents detailing the value chain requirements as if they were part of the compliance teams of the provider and third-party supplier of the high-risk AI system. These prototype compliance documents can be found in the **annex** to this report.

The workshop followed a legal design methodology, building on our earlier experiences with legal design workshops.[7] In practice, this means that the workshop had **four parts**.

1. Empathise

The first part focused on **understanding** the technical use case and its environment. For each use case, this also included mapping out the affected stakeholders and their concerns.

2. Define

During the second part, participants defined **the problem(s)** that needed to be solved. This included considering questions such as: What needs to be in the prototype? Which (legal or practical) requirements may be difficult to include? Are there aspects of the system's environment or users that are an issue for the prototype? When will the prototype be used?

3. Ideation

The ideation part served to brainstorm **possible solutions** to the problems defined in the previous part, while taking into account the affected stakeholders and their concerns. At the end of this part, possible solutions were grouped and prioritised, and a choice was made regarding the prototype that would be developed.

4. Prototyping

During the last part, participants started to work on **an actual prototype**. As participants knew that the prototypes would be further developed during the next phase of the policy prototyping project, they focused on agreeing on the structure and substantive foundation of the prototypes.

[7] *Knowledge Centre Data & Society, Our first legal design workshop: five prototypes for more accessible legal processes.*

3.4

Phase II: further elaboration of prototype compliance documents

The design workshop was followed by a second phase, during which the prototype compliance documents created during the workshop were further developed by the respective team members. This phase took place throughout February 2026.
The idea of this phase was to create well-developed **prototype compliance documents**. The documents should approach, according to the evaluation of the participants and in so far as possible for the use cases, a final document that could have been created by a provider and third-party supplier. The document was subsequently presented to reviewers for comprehensive feedback.

3.5

Phase III: feedback

Once the prototype compliance documents were created, we launched phase III of the policy prototyping project: **feedback**. In order to diversify the potential feedback, we published a second call for participants.[8] This call did not distinguish between types of participants and aimed to attract professionals and experts in AI, as well as IT law, contract law, intellectual property rights and so on. People who registered their interest to participate during the first call for participants but could not attend the design workshop were also invited to contribute to this phase. Diverse profiles, including both technical and legal experts, responded to the call for feedback.

We gathered feedback on both the created prototype compliance documents and the related legal requirements from Article 25(4) AI Act. Participants were able to provide feedback on (i) how the prototypes implemented the requirements of the AI Act, and (ii) the practicability, feasibility and desirability of the legal requirements themselves.

Feedback was gathered through online or in-person interviews. A total of nine interviews were conducted using a predetermined questionnaire provided to the reviewers in advance, along with the prototype compliance documents.

[8] *Knowledge Centre Data & Society, Call for experts: Help us reviewing model clauses for article 25(4) of the EU AI Act.*

3.6

Phase IV: report – publication of feedback and lessons learned

This report is the **final stage** of our policy prototyping project. It contains our findings, based on aggregated participant feedback, and lessons learned regarding the implementation of the requirements of Article 25(4) AI Act.

This report is driven by **multiple objectives**. Primarily, it aims to **assist** stakeholders and professionals to effectively operationalise the Article 25(4) requirements by offering examples of compliance documents, coupled with best practices and valuable lessons learned. Additionally, it seeks to **convey** the insights gathered from this project to policymakers and authorities, providing them with a practical perspective that could be instrumental in improving the AI Act's future implementation. Lastly, the report **contributes** to the evolving conversation on policy prototyping, advocating for its significant value as a tool in the policy development process. Hence, **the intended audience of this report** is (i) policymakers involved with the AI Act and its implementation, (ii) supervisory authorities that will be involved in the enforcement of the AI Act, (iii) all stakeholders that will need to comply with or benefit from these requirements, and (iv) all other interested parties.

The AI Act – text of Article 25(4)

4

[9] *Article 6 of the AI Act*.

Article 25(4) of the AI Act stipulates that a written agreement must be in place between the provider of a high-risk AI system and the third party supplying, among other things, components (e.g. hardware, software and data), tools (e.g. the development suite) and services (e.g. support, storage and computing power) for integration into the AI system. This is a written agreement on the necessary information, capabilities, technical access and other assistance, based on the generally acknowledged state of the art, which enables the provider to fully comply with the obligations set out in the regulation.

An AI system is considered to be **high-risk** if:[9]

- → The AI system is used as a safety component or itself is a product covered by EU laws specified in **Annex I of the AI Act** and required to undergo **a third-party conformity assessment** under those laws, such as lifts, cableway installations, radio equipment, etc.; or
- → The AI system is listed under **Annex III** use cases, such as AI systems intended to be used to evaluate learning outcomes, **except** if:
  - → The AI system is intended to perform a **narrow procedural task**;
  - → The AI system is intended to perform **the result** of a **previously completed human activity**;
  - → The AI system is intended to **detect decision-making patterns or deviations** from prior decision-making patterns and is **not meant to replace or influence** the previously completed human assessment, without proper human review; or
  - → The AI system is intended to perform a **preparatory task** to an assessment relevant for the purposes of the use cases listed in Annex III.

If the AI system fits one of the scenarios described above but engages in **profiling**, it is still considered a high-risk system.

Below we include the text as used by the participants:

Article 25: Responsibilities along the AI value chain

Paragraph 4.

The **provider of a high-risk AI system and the third party that supplies an AI system, tools, services, components, or processes** that are used or integrated in a **high-risk AI system** shall, by **written agreement**, specify **the necessary information, capabilities, technical access and other assistance** based on the **generally acknowledged state of the art**, in order to enable the provider of the high-risk AI system to **fully comply with the obligations set out in this Regulation**. This paragraph shall not apply to third parties making accessible to the public tools, services, processes, or components, other than general-purpose AI models, under a free and open-source licence.

The **AI Office may develop and recommend voluntary model terms for contracts** between providers of high-risk AI systems and third parties that supply tools, services, components or processes that are used for or integrated into high-risk AI systems. When developing those voluntary model terms, the AI Office shall **take into account possible contractual requirements** applicable in **specific sectors or business cases**. The voluntary model terms shall be published and be available free of charge in an easily usable electronic format.

Results

5

First phase

In the first phase of the policy prototyping project, we invited professionals working in AI and/or with the AI Act to a design workshop. The objectives of the workshop were as follows:

- → Understand the use cases and concerns of the potential stakeholders.
- → Identify and brainstorm how required elements must be included in the prototype clauses.
- → Design the actual prototype, which will serve as the foundation for further development in a later phase.

During the workshop, participants selected elements and designed **prototype clauses** for the AI value chain in compliance with Article 25(4) of the EU AI Act. The designed clauses are based on three use cases, each of which were created by a different group of experts and professionals who attended the workshop.

Next phase

Regarding the next phase, in which we sought feedback from experts, we primarily focused on the following aspects:

- → The clarity and adequacy of the designed prototype, its potential implementation challenges and the general fit with the requirements of Article 25(4) (section 7).
- → Examining the text of Article 25(4) of the EU AI Act in detail, including its practicability, desirability and feasibility (section 8).

To ensure a proper understanding of the results, we have included **the prototype documents from the design workshop in the annex**. We encourage you to refer to these documents if any of the following points in this report are unclear.

We will now present the **feedback on the three prototypes** that were drafted during the design workshop. Each of the designed contractual clauses took into account a different use case (see sections below) and the requirements of Article 25(4) of the EU AI Act.

5.1

Use case 1 – Chatbot insurance industry

The first use case involves a company that offers various types of AI systems to insurance companies. The company aims to deploy an AI-powered chatbot consisting of **a risk-assessment tool** for customers of insurance companies to apply for **life insurance to pay off their mortgage**. The tool uses fully automated decision-making, whereby applications can be accepted or rejected via the system. The risk assessment consists of various elements that are provided by the customer, such as their age, salary, health, and whether they are taking out the insurance alone or with another person. The company licenses the chatbot from a third-party supplier for the core of the AI model as well as the software platform. The tool is categorised as a high-risk AI system under Annex III of the AI Act.

The workshop participants decided to design a clause that focuses on the following aspects: the licence granted by the supplier, the back-end access and technical cooperation, and the protection of intellectual property, trade secrets and data as well as confidentiality.

General overview

According to the reviewers, the first prototype provided a **sound basis** for implementation. However, it was also deemed too abstract and general, not encompassing the nuances of the insurance industry, particularly the significant impact that the chatbot could have on clients (individuals applying for a mortgage).

Included and found beneficial

- → Sound starting point
- → Clear language
- → Logical structure

Lacking or requiring further elaboration

- → Terms such as 'proportionate' need to be further clarified
- → Extra clauses or annexes on:
  - → Updated documentation needed for compliance
  - → An overview of compliance capabilities
  - → Support and assistance for e.g. security incidents and vulnerabilities
  - → Obligations relating to logging, traceability and evidence reporting
  - → Obligations on human oversight and fairness monitoring
  - → Regulatory cooperation and auditing
- → Clarity on the obligations resting on the supplier and the division of responsibilities between the parties

Prototype clarity and adequacy

Overall, the clause was considered to be a **solid starting point** for governance. Nevertheless, this standpoint was immediately nuanced by highlighting the lack of practicality of the prototype. As stated by one reviewer: 'The prototype is a sound starting point, but for a fully automated high-risk AI system, it is too thinly developed compared to what Article 25(4) requires in practice.' Several reviewers noted that the prototype lacked concrete details on human oversight, logging and technical documentation, etc. In addition, one of the reviewers emphasised that the contractual clause implicitly divides responsibilities in a manner that would be difficult for both parties to sustain. It was also noted that the contractual definitions of 'provider' and 'third-party supplier' were unclear. One of the reviewers remarked that the provider appeared to function more as a reseller than as a provider.

The response to the question of whether the clause could be implemented in the provider and third-party supplier relationship was negative. Two reviewers pointed out that there was **insufficient consideration** given to the **processing of special categories of personal data** in a high-risk context. The reviewers also stated that the clause primarily protects the supplier and fails to strike a balance from the provider's perspective. In practice the clause would benefit if it was part of a broader contractual framework and a data processing agreement (DPA), particularly for the health data processing. Lastly, one reviewer stated that the **responsibilities were not properly embedded** and that the division of responsibilities should start earlier, for example in the quotation or proof of concept phase. They also highlighted the issue of **intellectual property (IP) allocation**, stating that existing IP must be retained, while IP specifically developed for the client must be transferred to them.

The reviewers shared concerns about the **lack of sector-specific nature** of the prototype clause. The clause does not fully address the implications that input from a chatbot or automated assessments may have on individuals' access to services or their financial outcomes. One of the reviewers specifically stated the following: 'Every chatbot is biased. Technical documentation must be provided sufficiently, including information on the bias the chatbot may encounter.' Overall, the reviewers showed concern about a potential over-reliance on automation. It was therefore deemed crucial to **incorporate human oversight measures** to the clause. In addition, the feedback suggested adding explicit safeguards for contestation and justification of individual decisions. As this specific use case involves the processing of sensitive information such as health data, the prototype clause must also govern **how the system is being managed and monitored**. Finally, in light of the particular context of this use case, it is imperative that the prototype incorporates risk management measures such as bias detection, traceability, auditing, logging and the continuous monitoring of outcomes (for example a history of error rates), and so on.

On the clarity of the contractual clause, the reviewers agreed that the **language** is accessible and generally clear. However, some expressions such as 'proportionate' or 'appropriate' need further clarification. Additionally, the reviewers suggested that there were provisions lacking on incident reporting, a documentation package provided by the supplier and a clause for the third-party supplier to enable human oversight, etc. (see the table above).

In addition, the contractual clause's **structure** received positive feedback from the reviewers. However, while the structure was considered to be logical, the reviewers noted that some provisions were missing (see the table above). One of the reviewers stated that the clause would appear to be more logical if the contract between the parties were to be divided into the following phases: a quotation phase, a service agreement phase, a terms of use phase and a support agreement phase.

Finally, the reviewers expressed some concerns regarding the adequacy of the contractual clause as a means of **ensuring compliance** with Article 25(4) AI Act. On several occasions the reviewers stated that the clause provides a useful starting point. However, the division of responsibilities is considered unclear, in particular with regard to the identification of the provider, supplier or 'intermediary'. As mentioned previously, various additional provisions are required for the prototype to become clearer and more comprehensive.

Practical implementation challenges

As stated earlier, the reviewers reiterated their concerns that the prototype does not adequately address the potential risks associated with the use of chatbots in the insurance sector. The lack of provisions addressing **human oversight, the risk of bias and discrimination, and automated decision-making** by the chatbot were highlighted.

In addition, it was generally acknowledged that the prototype implicitly assumes the third-party supplier is a large and established supplier, rather than a small specialist vendor. AI vendors often do not operate alone, but rely on sub-processors or other critical parties. The current text of the contractual clause creates additional friction in practice, particularly **SMEs with less bargaining power, limited resources and greater difficulty in conducting legal reviews**. According to one of the reviewers, it should also be noted that cloud providers should not be approached as 'ordinary' suppliers, as they tend to apply their own contractual frameworks.

On the effective negotiation and implementation of the contractual clause, the reviewers identified several issues. Some of the obligations were considered to be **too broad** and not specific enough, resulting in a lack of clarity regarding the required deliverables. Furthermore, one of the reviewers mentions that there is a **misalignment between responsibility and control**, due to the supplier's power to restrict access to the AI system. The same reviewer argued that the wording of the prototype prevents the provider from acting with sufficient transparency regarding its own obligations. This leads to a contractual imbalance that may not be practical enough for effective implementation.

The prototype in light of the Article 25(4) AI Act requirements

In relation to the specific elements of Article 25(4), namely 'necessary information, capabilities, technical access or other assistance', the general consensus of the reviewers is that the contractual clause covers these partially. The current text of the contractual clause is deemed **too general** to fulfil the obligations under Article 25(4). The reviewers suggested that an annex should be provided to address the following elements: the relevant bias and error margins, the technical documentation that is expected and which tools, such as bias detection, traceability and human oversight, are required. The term 'other assistance' also remains too vague, particularly regarding the handling of incidents and other support obligations. Furthermore, while the clause refers to 'technical access', it is rather restrictive and appears to offer a greater advantage to the third-party supplier.

The responses to whether the provider can fulfil its obligations under the AI Act were predominantly negative. The reviewers claimed that the prototype offers **insufficient support** for the provider to effectively fulfil their obligations. The primary issues referred to were the limited technical access for testing and monitoring, as well as the lack of reference to auditing mechanisms and explainability tools. One of the reviewers also drew specific attention to the

difference between responsibility and control. In this specific use case, the provider may not be able to demonstrate that they exercised sufficient supervision, despite being the prime accountability party under the AI Act.

Finally, the reviewers highlighted the risk of discrimination and lack of individual explanations for decisions. In an **insurance context**, the chatbot may (indirectly) discriminate against a certain group of applicants. Furthermore, incorrect input can lead to disputes. With the risk of over-reliance on the automated decision-making by the chatbot, human oversight remains essential. It is also important to note that insurance laws may vary by state or jurisdiction, which can further complicate compliance with the AI Act.

5.2
Use case 2 – Medical device software

The second use case relates to the development of medical device software that (i) assists physicians in **prescribing the appropriate dosage of antibiotics**, and (ii) **predicts and stays ahead of emerging antibiotics resistance** to improve the treatment and prevention of serious infections. Drug-resistant infections are a growing global problem and many believe that superbugs can be defeated by supercomputers. In this regard, the provider and third-party supplier are entering into a partnership in order to achieve the previously mentioned targets. The third-party supplier provides the following to the provider: (i) specific AI models (including the algorithms for decision-making) and (ii) the hardware (GPUs/TPUs) for processing. The AI system is categorised as a high-risk AI system under Annex I of the AI Act.

The participants drafted a clause in which they focused primarily on data governance aspects, given the great importance of reliable data in the medical sector.[10]

[10] *Practical note: the design workshop for this specific use case involved three people, two of whom took turns participating, meaning that one of them missed the first part and the other the last part.*

General overview

In summary, the prototype is understood by the reviewers as **very general** in the sense that it reiterates the theory set forth in the AI Act. Furthermore, it fails to provide concrete details regarding the obligations under Article 25(4) of the AI Act.

Included and found beneficial

- → Clear language
- → Data governance is covered

Lacking or requiring further elaboration

- → General wording of the prototype clause
- → Definition of specific terms
- → More detailed description of the intended purpose
- → Not enough specific obligations that strike a fair balance between the parties, for example:
  - → Provision for human oversight
  - → Provision for periodic evaluation and review
  - → Provision for a right to renegotiate
- → Insufficient attention to liabilities

Prototype clarity and adequacy

The reviewers' initial impression of the prototype was that the language was **clear** but the text was rather general. The obligations were not specified in sufficient detail. It was also noted by one of the reviewers that the wording of the AI Act had been reproduced. Additionally, there was no reference made to **regulatory overlap** (General Data Protection Regulation or GDPR, Medical Device Regulation or MDR, and Data Act). Furthermore, one of the reviewers noted that insufficient attention had been paid to elements such as technical access and model development, which are, in essence, important aspects of the clause.

The reviewers considered that the prototype could not be applied in the relationship between the provider and the third-party supplier, as it lacks **concrete detail** and it is difficult to ascertain what parties are actually required to do. The reviewers also noted that the **balance of obligations** between the parties needs to be more equitable.

Additionally, the prototype lacked coverage of important aspects, such as audit rights, corrective actions, risk management, clarity around model development, cybersecurity and, ultimately, liability allocation and indemnification. **Incidents**, such as cybersecurity incidents, must be reported to the provider by the third-party supplier when needed so that adjustments can be made to the supplied components. If the supplier decides not to make certain updates to the system, at the very least a justification must be provided as to why. If there is no valid justification, the third-party supplier should be held liable in the event of a similar incident.

On top of that, the prototype should include measures that mitigate the risk of **bias**. Particularly in a medical context, attention must be paid to the technical aspects of avoiding bias, but ethical and social concerns must also be taken into account. In terms of the technical aspects of avoiding bias, this can be addressed in a data governance provision that ensures the quality and integrity of data.

Another point made by the reviewers was that the prototype placed greater emphasis on the use of data and less on the technical aspects. In addition, one reviewer underlined the importance of human oversight measures. The third-party supplier may need to assist the provider in drafting instructions for use where necessary. The significance of this lies in the fact that it is, naturally, the user who enters data into the system. Consequently, users must meet certain requirements regarding the quality and integrity of the data that may be entered into the system. It is and remains a system, and it only works as well as the data it is fed. In certain cases, it may be necessary to consult with the third-party supplier in order to assist in determining these requirements.

Relating to the **clarity** of the prototype, all reviewers agreed that there was a need for definitions of certain terminology or reference to certain **sector-specific regulations** in order to clarify exactly what was being set out in the contractual agreement. A specific example is the high-risk system in question; in this context, the prototype must also reference the specific applicable legislation (such as the MDR).
In the prototype, the intended purpose is described in a simplified manner as the two objectives of the proposed software, with no additional information. The reviewers recommend that it is explicitly stated that the software is used in the medical sector.

Another aspect that would benefit the involved parties in their respective relationships with each other is the addition of a **compliance clause** to the written agreement. This clause would see the parties declaring their intention to comply with all their legal obligations arising from different regulatory frameworks, such as GDPR, MDR and the Data Act.

In addition, the evolving landscape of legislation, as well as the potential reclassification of the high-risk system into a non-high-risk system, warrants the inclusion of a clause that stipulates **the right to renegotiate**. This right could ensure that specific provisions can be amended quickly without the necessity of drafting an entirely new contract. Even though Belgian law includes a doctrine regarding hardship and good faith that can trigger this right, it remains a good practice to include a clause for renegotiation explicitly in the contract as well.

Furthermore, one reviewer proposed to include a **provision for periodic evaluation and review**. Ideally, this would entail a cooperation and information-sharing obligation, which can be combined with an audit right.

Reviewers also mentioned that it is crucial to set out the **intellectual property rights** in the contract, given the constant improvement of algorithms. This is relevant for monitoring and post-market surveillance. As the third-party supplier may also hold vital information, it would be beneficial for them to pass this information on to the provider.

Finally, the importance of the issue of **liability allocation** should not be underestimated. Particularly in this use case, which involves patients, it is essential to reach agreements on this matter. It is not always clear who can be held responsible if things go wrong. One expert also noted that the Belgian law provides a general liability regime, under which one cannot be exonerated for physical injury, which further highlights the significance of this liability issue.

Practical implementation challenges

The experts unanimously agreed that, in theory, everything is open to discussion. Nevertheless, in practice, major players (in particular providers), impose their general terms and conditions on the other party. A prototype could therefore be a **solid starting point**, particularly for **SMEs** that may not have the resources to hire a legal team to draft contracts. One expert specified that a prototype is usually designed with small businesses in mind.

However, as the business expands, the importance of **solid contracts** increases, as there is a greater risk of things going wrong. This highlights the importance for the provider to have robust contracts in place, as they bear the end responsibility under the AI Act. It is essential that the third-party supplier maintains their involvement and supports the provider in the further implementation and launch of the AI system. It is for this reason one expert explicitly states that drafting a contract becomes a matter of liability.

It is also evident that the AI value chain may involve numerous actors. Therefore, it is crucial to clarify the roles of the provider and the third-party supplier. However, for the sake of completeness, reference should also be made to **other stakeholders**, as their involvement may affect the AI system and the obligations of the provider and third-party supplier.

The prototype in light of the Article 25(4) AI Act requirements

According to the reviewers, Article 25(4)'s reference to 'necessary information, capabilities, technical access and other assistance' is **unclear**. In order to comply with Article 25(4) of the AI Act, it is important that the parties clarify the aforementioned matters: what information must be provided, what form must the assistance take and how often it is provided, etc. One of the practical consequences of uncertainty is that the parties involved recognise the need to collaborate, but the manner in which this must be done is less clear. In the event of a negative outcome, there is a limited safety net.

Furthermore, the reviewers consider the individual terms 'necessary information', 'capabilities', 'technical access' and 'other assistance' to be overly broad. It is challenging to comprehend the intended meaning of these terms as stipulated by the legislator. Furthermore, it should be noted that the interpretation of these terms is subject to change over time. All reviewers therefore stress the importance of guidance from the European Commission to enable compliance.

5.3

Use case 3 – AI system for the ranking and screening of CVs

The AI system in this use case supports recruitment decisions by **ranking and screening candidates** on **the basis of CV data and other inputs**. The system is positioned by its provider as a decision-support tool, but in practice it produces outputs that **materially influence** whether a candidate is invited to interview or hired. The AI system can be categorised as a high-risk AI system under Annex III.

The prototype document developed for this use case focuses on incident handling between the parties (i.e. provider and third-party supplier) to the contract. It is structured in **three layers**: a high-level colour-coded overview, a light version covering the core obligations, and a full version that elaborates a points-based criticality classification. The full version distinguishes between low, medium, high and critical incidents, and sets out the cooperation expected between the parties at each level.[11]

[11] *Practical note: The choice for incident handling in particular was made during the design workshop by the participants. The main reason for said choice was that this was a matter that the participants agreed was feasible to complete in one day, being aware of the limitation that covering all supply-chain-related aspects would not be feasible in one day for the given use case.*

General overview

In summary, the prototype is understood by the reviewers to contain **a logical structure**. However, it does not address certain elements that are of particular importance to the HR sector, nor does it provide more detailed guidance on the key aspects of Article 25(4) AI Act.

Included and found beneficial

- → Layered approach (colour-coding, light and full version)
- → Criticality classification and points-based system
- → Detection, mitigation and remediation logic
- → Flexible timelines left to parties
- → Inclusion of bias-related deviations in the definition of incident
- → Logical and chronological structure of the clauses

Lacking or requiring further elaboration

- → Reference to a fully fledged AI Act addendum
- → Intended purpose of the system
- → Concrete HR terminology (discrimination, fundamental rights, minority protection)
- → Capabilities, necessary information and technical access (the literal Art. 25(4) elements)
- → A general information-exchange clause beyond incidents

Prototype clarity and adequacy

The reviewers' impressions of the HR prototype were a mixture of recognition and reservation. One reviewer praised the **layered approach**, including the light version, as clearly thought through, and described the criticality classification together with the points-based system as creative and useful.

However, other reviewers were more critical of the scoping of the prototype, in particular on the choice for the participants to focus on incident handling. One reviewer even stated that the prototype fell outside of the scope of Article 25(4) AI Act due to it focusing on the relationship between a provider and a deployer, rather than the relationship between the provider of a composite system and the provider of a component of a high-risk AI system. The above could show that the participants had an issue understanding the relationships to be covered by the prototype, or latent ambiguity in classifying different relationships between different upstream and downstream actors.

Reviewers also indicated that the prototype reads more as **an operational or technical document**, comparable to a service level agreement, than as a fully fledged legal addendum. The AI Act addendum to which such an operational document would normally be attached was missing from the package, which led one reviewer to conclude that some of the core legal obligations under the AI Act remain unaddressed by the prototype itself. The reviewers nonetheless agreed that the clauses on when an incident must be reported, the associated risks and the procedure to follow are realistic and the kind of provisions that would indeed feature in this type of agreement. One reviewer underlined that the parties will always need to adapt the clauses to their own service relationships, and that the prototype is best suited for **negotiated use within a service level agreement**. A recurring observation was that the document presupposes a one-to-one cooperative relationship between two medium-sized enterprises, which is **not necessarily representative** of how supply chain interactions actually unfold in practice.

Reviewers were cautiously positive about the prototype's potential as a **contractual instrument**, though each qualified their assessment in a different way. One reviewer adopted a wait-and-see position, noting that instruments that appear promising on paper do not always hold up when tested in practice. A second reviewer acknowledged the value of having model clauses available, particularly given the risk that parties might otherwise rely on informal or unreliable sources when drafting high-risk AI contracts, but pointed to gaps in the treatment of escalation, governance, audit rights and suspension mechanisms. Taken together, the responses suggest that the prototype offers a credible starting point, but requires further refinement before it can function as a robust contractual instrument in practice.

The reviewers also questioned whether the prototype sufficiently reflects the **specific realities of the HR sector**. Although a reference to employment law sensitivities is included in clause 1.4.5 and the definition of incident covers bias-related deviations, the reviewers found the document too generic for an HR high-risk use case. Specific concepts such as direct or indirect discrimination, fundamental rights in the recruitment context, and the protection of minority groups were largely absent. Reviewers observed that, in Belgium in particular where employment law is strictly regulated, the absence of these elements is a notable gap. One reviewer also remarked that the points-based system, while creative, contains values that have not been calibrated to HR-specific harms, which makes it difficult to assess how serious it actually is when, for example, five hundred candidates are affected rather than fifty.

The reviewers found the **language** of the prototype overall to be precise and appropriate for a contractual instrument. Some jargon is considered unavoidable in this kind of document, but the reviewers suggested supplementing the prototype with worked examples or sidebars in which a concrete use case is explained and linked back to the text. Such practical illustrations were considered helpful for both HR professionals who may have less experience with IT-style operational documents and non-tech-savvy stakeholders more generally.

The **definition of incident** emerged as a recurring point of concern. The current definition, which refers to any incident or malfunction affecting the AI solution, was seen as circular by one reviewer and too narrow for an HR context by another. In recruitment, an incident may involve discriminatory outcomes that only become visible across patterns and trends over time, rather than in individual decisions. Reviewers questioned whether the wording captures such events, and whether bias-related deviations are sufficiently reflected in the clauses on detection, mitigation and remediation. Concepts such as fundamental rights also need clarification, particularly in their interaction with employment law.

On **structure**, the reviewers found the prototype logical and chronological, with one reviewer noting that the document clearly tells the story of an incident from detection to remediation. There were nonetheless some observations on the layered approach. One reviewer questioned the added value of the first layer, which essentially functions as a colour code aimed at a C-suite audience, while the second layer was considered the substantive core of the document. The numbering of clauses was felt to go too deep, up to four levels, which could be simplified without losing meaning.

Whether the prototype provides **sufficient guidance under Article 25(4)** was assessed cautiously by the reviewers. One reviewer described the question itself as difficult to answer in binary terms, as it amounts to a question of overall legal compliance. Within the scope of incident handling specifically, the reviewers considered the guidance to be reasonable but not exhaustive. The scoring system was identified as inevitably open to interpretation, which was not seen as problematic in itself, provided the parties using the document are properly trained to apply it consistently.

Practical implementation challenges

Several risks particular to the HR use case were identified as inadequately addressed. The most prominent of these is **bias and discrimination**, which by their nature unfold over time and across multiple decisions and which are not easily captured by an incident management framework that focuses on discrete events. One reviewer noted that an HR system may function without any apparent incident in any individual case and yet produce systematically discriminatory outcomes when assessed over the course of a year. The same reviewer recommended that the existing clause on detection and escalation be complemented by a broader obligation to monitor continuously for incidents, rather than only to detect strange cases.

Suspension rights and the consequences of serious incidents under the AI Act, in particular under Article 73, were also considered insufficiently developed. The reviewers observed that critical incidents and serious incidents are currently placed under the same heading, with every serious incident being treated as a critical incident, while the AI Act itself attaches additional reporting and procedural obligations to serious incidents. A clearer separation of the two definitions, together with explicit guidance on when fundamental rights are at stake, was recommended. **Human oversight** was likewise felt to be missing: the prototype concentrates on what happens after an incident has materialised and does not address the extent to which the user retains the ability to override or influence the AI system's output before that point. **Remediation** itself was found to be loosely defined, with one reviewer observing that it is not entirely clear when a problem is considered resolved and when the provider has fulfilled all of its obligations.

On the **diversity of third-party suppliers**, the reviewers were unanimous in finding that the prototype implicitly assumes a particular type of supplier. The cooperative one-to-one relationship presupposed by the document is realistic between two medium-sized enterprises, but does not match the way in which large cloud or AI providers actually engage with their customers. Large providers, the reviewers observed, will issue general notices about incidents rather than addressing situations individually, and unless a customer has the highest tier of service the cooperative communication assumed by the prototype will not happen in practice. Conversely, the prototype is not the type of contract that could realistically be imposed on a Microsoft, Workday or another large HR or AI system provider.

The reviewers nonetheless considered that the prototype could be **negotiated and implemented in practice**, in particular between SMEs and as part of a service level agreement. The classification system was viewed as useful, although complex, and as requiring training and a shared understanding between the parties. Reviewers also anticipated that, in actual contract negotiations, providers would seek to push obligations towards the deployer, and that practical clarifications would be needed about entry points, the deployer's responsibility for its own user interface, and the allocation of monitoring duties between the parties.

The prototype in light of the Article 25(4) AI Act requirements

The reviewers found the prototype's coverage of the elements explicitly listed in Article 25(4) to be partial. One reviewer concluded that, judging from the document on **incident handling** alone, it is not clear what the 'capabilities' of the tool actually are or what is meant by the 'necessary information', and that 'technical access' is not addressed. Another reviewer observed that the prototype focuses on incident management while Article 25(4) also speaks to other compliance aspects such as **technical documentation, intended use and human oversight**. The reviewers suggested that the prototype could be reframed by starting from the literal text of Article 25(4) and then systematically spelling out what is meant by necessary information, with reference to Articles 11, 13 and 14 of the AI Act.[12]

Overall, the reviewers considered the prototype adequate for incident handling within its chosen scope, but **insufficient** if assessed against the full breadth of Article 25(4). One reviewer described the document as a small piece of a larger whole, and another stressed that the prototype does not contain a general obligation to exchange information, capabilities and technical assets between the parties even when no incident has occurred. One reviewer even went as far as to state that this prototype fell outside the scope of Article 25(4) to begin with.

Several reviewers suggested introducing a more **general cooperation clause** before the existing incident-specific provisions, stating that the parties shall collaborate and exchange information so as to enable each other to fulfil their obligations under Article 25(4). The starting point of an incident (occurrence versus detection) and the moment at which mitigation and remediation obligations are triggered were also flagged as requiring clarification, since an incident in an HR system may linger for months before it surfaces.

The reviewers agreed that the document is **helpful but not sufficient** on its own. Two reviewers considered that, within the scope of incident management, the prototype provides a workable basis. One reviewer noted that the contractual obligations are largely clear and that the provider would be able to fulfil their obligations on that basis, in particular with regard to the reporting of incidents and the cooperation expected with the deployer. The reviewers nonetheless cautioned that obligations such as the preparation of technical documentation or the implementation of human oversight fall outside the scope of the prototype and cannot be covered by it. One reviewer also pointed out that the suggestion in the light version that the document can be read as a standalone agreement of two pages is unrealistic given the missing elements, and recommended an explicit hierarchy clause stating that the overarching AI Act addendum takes precedence in case of conflict.

[12] *These articles relate to technical documentation, transparency and human oversight.*

5.4

Recommendations and lessons learned for Article 25(4) written agreement between providers and third-party suppliers

Based on the extensive discussion of reviewers' feedback above, this section of the report brings together **recurring recommendations, best practices and lessons learned** from the policy prototyping project that can be used by providers to enhance their own written agreements with third-party suppliers.

Recommendations for drafting written agreements according to Article 25(4)

Assess the qualification of each party in detail

The written agreement between provider and third-party supplier should take into account the **allocation of roles** that applies in practice and not just on paper. The written agreement should reflect the actual supply chain accurately and avoid setting unrealistic expectations of one of the two parties. It should be noted that it is also important to clearly define the obligations of the parties. This ensures that the parties understand what is expected of them. Attention must also be paid to striking a fair balance between the rights and obligations of the parties.

In addition, consideration must also be given to **other parties** that may be involved in the supply chain, such as developers that use third-party supplier elements or deployers that carry notification and detection duties, particularly for incidents that originate at the level of the deployer's user interface rather than the provider's back end.

Define key terms

In addition, written agreements should be structured by starting from the literal text of **Article 25(4)**. This would commit the parties to specifying 'the necessary information, capabilities, technical access and other assistance'. These terms could be spelled out by reference to Articles 11, 13 and 14 of the AI Act, which set out, respectively, the technical documentation, the instructions for use and the human oversight obligations applicable to high-risk AI systems. One reviewer also suggested drawing a parallel with the transparency obligations applicable to general-purpose AI systems under Article 50, by requiring the third-party supplier to provide transparency information, copyright sources and a security analysis to the extent that these exist.

Relating to the term '**assistance**', parallels are drawn with data processing agreements (DPAs). Typically, the DPA includes a clause that specifies the scope of the service and the assistance one would reasonably expect from the data processor to the data controller. The latter is subject to the many obligations under GDPR. In practice, however, it has been noted that the controller makes requests to the processor, the processor agrees, sends the requested

information, and an invoice follows. The controller often counters this by arguing that they are fulfilling their obligations under GDPR, and that the processor must assist in this regard. The fundamental question that arises from this perspective is: what may be charged additionally, and what constitutes reasonable assistance? By analogy, the same applies to Article 25(4) of the AI Act. In this instance, the provider and third-party supplier are obliged to define the scope of ongoing assistance, to the greatest extent possible.

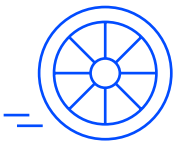

Don't reinvent the wheel

It is also recommended that existing materials or standard contracts be used. This may be of particular relevance to SMEs. Many aspects relating to the broad terminology of 'capabilities, necessary information and technical access' are already being implemented in practice. **Scoping** is of the utmost importance. It is imperative for companies to use quotations to safeguard their interests. During the quotation phase, the precise nature of the provider or supplier's role can be clearly defined. The subsequent service agreement will then allow everything to be finalised. The relevant technical specifications can be included in the terms of use. The support contract can concentrate on dealing with incidents, as well as other related issues. All of these frameworks already exist and can serve as guidelines, particularly for SMEs.

General best practices

Language and additional information

To ensure the prototype is user-friendly, the language used must be tailored to **the target audience(s)**. Depending on the context of the prototype – for example, in the field of HR – this must be clearly reflected in the language of the document, with explicit references to discrimination and employment law.

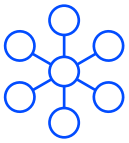

Structure

Great importance was attached to the structure of the prototypes. A clear and logical structure enhanced the **user-friendliness** of the documentation and was considered useful for service providers and third-party suppliers. Reviewers referred to **existing contractual structures** that could serve as a source of inspiration for devising a better structure for the document, such as data processing agreements, service level agreements and so on. Several reviewers have also noted that the written agreements should include an appendix containing technical documentation that provides further information on the source code, history access, history of tests for potentially discriminatory outcomes and so on.

Feedback on Article 25(4) of the AI Act

6

In addition to the questions on the prototype itself, the feedback interviews invited the reviewers to share **their views on Article 25(4) of the AI Act** as a regulatory mechanism. The questions covered the practicability and feasibility of the obligation, in particular for SMEs, the availability of existing contractual frameworks that the new agreements could build on, the desirability of using written contracts to manage value chain obligations along the AI value chain, and the elements of the article itself that the reviewers found unclear or difficult to operationalise. The reviewers' answers are summarised below, synthesised across all interviews.

Practicability
assesses whether the requirements can be implemented without excessive difficulty in real-life situations.

Feasibility
assesses whether the requirements can be operationalised given the resources and constraints available to a provider, such as budget, time, technology and workforce.

Desirability
assesses whether the requirements and their operationalisation are useful and valuable to the intended audience and users.

6.1 Practicability

Negotiating and implementing Article 25(4) agreements in practice

The reviewers were almost unanimous in their assessment that Article 25(4) agreements are **difficult to negotiate and implement** in practice, in particular for SMEs. The asymmetry of bargaining power between SMEs and large technology providers was identified across many interviews as one of the main practical barriers.

One reviewer described the article itself as extremely brief and generic, with no allocation of obligations between the parties. In that reviewer's view, meaningful implementation will only be possible if additional guidance is issued from a higher level, comparable to the role played by the Article 29 Working Party guidance under GDPR. The **AI Office** could fulfil that role here by publishing voluntary terms for contracts, which in addition could also be of great help for SMEs.

In addition, it must be taken into account that off-the-shelf solutions typically come with boilerplate terms, which leave SMEs little room to negotiate. Negotiations with large providers are essentially take-it-or-leave it. Where parties are similarly sized, neither side may have legal counsel available. In the reverse scenario (an SME provider dealing with a large purchasing entity) the large enterprise's strict procurement conditions create an equivalent imbalance. Realistic negotiation under Article 25(4) requires a multidisciplinary team combining technical and legal expertise, as well as an anticipatory understanding of how the AI model will be used in the future. This is rarely achievable for small (data-science) teams, which may in practice rely on AI tools to draft or negotiate the contract itself. One reviewer drew a parallel with the data processing agreements introduced under GDPR, where in practice the larger party typically imposes its own template. In any case, another reviewer points out that templates remain crucial for smaller businesses because they do provide some practical guidance.

Then there is also the '**flexbox problem**'. According to one reviewer, there is a clear asymmetry between the role of AI providers and third-party suppliers, with third-party suppliers relying heavily on the protection of trade secrets. This can result in access to essential information about models and training data remaining limited, while precisely that access is crucial for supervision, risk assessment and compliance. This contradiction becomes a fundamental problem, particularly with generative AI systems like chatbots. This asymmetry can also translate into an imbalance in negotiations: major providers dominate the market, leaving smaller players with little leverage to enforce transparency or access rights. Furthermore and as mentioned previously, SMEs often lack the legal and financial resources to negotiate complex contractual safeguards. This is worsened by structural constraints, such as limited budgets and understaffed legal departments. Additionally, the requirement of Article 25(4) does not take sufficient account of the costs and resources required for ongoing due diligence and compliance monitoring.

According to another reviewer, the complexity is further exacerbated by the overlap of **various regulatory frameworks**, including the AI Act, GDPR and sector-specific regulations. This leads to a fragmented and difficult to navigate landscape, especially for SMEs.

[13] *International Organization for Standardization.*

[14] *European Insurance and Occupational Pensions Authority, Guidelines on System of Governance.*

Existing contractual frameworks to build on

One of the obstacles to drafting contractual provisions is that the field of contractual frameworks can be regarded as genuinely new territory. The reviewers also indicate this. Large law firms may already be developing internal instruments for clients, but anything that exists is highly customised, confidential and not available in public sources.

One reviewer suggested that Article 25(4) agreements do not need to be invented from scratch. **European public procurement clauses** already cover comparable AI-related obligations, with a particular focus on requiring data that complies with the AI Act, and could be a useful starting point. More broadly, the type of agreement is conceptually close to existing IT contracts in which technical requirements such as uptime, capacity or interoperability are negotiated jointly by legal and technical teams. Another reviewer refers to **existing frameworks**, such as software-as-a-service (SaaS) agreements, data processing agreements and security standards (such as the ISO[13] standards) that can already provide a solid foundation. Reference is also made to the **EIOPA guidelines**[14], which can be used as inspiration for the insurance industry.

What is new, however, are the **AI Act-specific obligations** that must be integrated as additional clauses. In essence, therefore, these obligations act as add-ons to existing contractual structures, rather than being a completely new system. According to one reviewer, this also suggests that the possibility for a pragmatic approach exists: Article 25(4) can be implemented via standardised annexes to existing contracts. The role of the AI Office in the development of these clauses could however still be particularly useful in this regard.

In considering the particular obligations set out in the AI Act, it is important to recognise that, for **high-risk use cases**, there is an existing framework that can be built upon, while this is not the case for other scenarios. For example, with regard to AI systems listed in Annex I and covered by sectoral regulations, the requirement of Article 25(4) will be regarded as 'merely' an additional layer. However, for use cases listed in Annex III, no overarching framework already exists and it may therefore be necessary to start from scratch.

Unclear, ambiguous or difficult-to-operationalise elements

The reviewers identified specific elements of Article 25(4) that they consider unclear, ambiguous or difficult to operationalise, although the points raised differed across the interviews.

One range of comments related to a series of **open concepts** within the article. The notion of 'necessary information' raises the question of what is, in fact, necessary. Minimum requirements would need to be set out in guidance. The terms 'capabilities' and 'technical access' similarly require concrete elaboration. One of the most fundamental concerns raised was the absence of a default rule allocating obligations between the parties. In its current wording, Article 25(4) reads somewhat like a link between joint controllers under GDPR, whereas it should perhaps function more like a controller-type provision that identifies who is primarily responsible for what.

Specific scope ambiguities were also described. One is the **allocation of responsibility**: both the provider and the third-party supplier are required to include such clauses in their contracts, but it is unclear to what extent the article can in practice be invoked by a provider against a large supplier (for example, to argue that a company such as Microsoft should have included a particular clause). Other remarks concern the AI systems that fall within the scope: while general-purpose AI models receive a great deal of attention in the AI Act, it is less clear how far the article reaches in respect of AI models developed for specific applications. The added value of the provision was also questioned, because the reviewers assume that compliant providers would typically already include this kind of clause in their contracts.

One specific question was raised regarding the **cooperation of the third-party supplier**. Reviewers wondered if this could be provided free of charge, or if a reasonable fee might be charged. If a fee can be charged, it could be set prohibitively high. If cooperation must be free, this has an impact on the third-party supplier. A clarification on whether a reasonable fee may be asked would be very important.

Finally, one of the reviewers concludes on a more pragmatic note. The reviewer argues that it is still **too early to assess** whether the practical implementation of this article, as well as the AI Act as a whole, remains a challenge. In practice, lawyers often prove to be creative in finding workable solutions. Just as privacy experts have had to adapt in the past, the AI Act will also lead to practical adjustments.

6.2
Desirability

State of the art benchmark

As an additional question, the reviewers were asked how they would themselves define or interpret the 'generally acknowledged state of the art' benchmark referred to in Article 25(4). The reviewers converged on the view that the concept is open and consensus-driven, and that **sector- and use-case specificity** are essential to give it meaning. Assessing what is appropriate therefore requires attention to the concrete field of application and, often, technical expertise. The concept could for example include information on model documentation and performance indicators. Reference can also be made to recognised standards, such as the various ISO standards.

At the same time, the **general formulation of the benchmark** is useful because elements such as accuracy levels cannot be fixed once and for all and will continue to evolve. Contracts can nevertheless make the obligation more concrete, for example by requiring that accuracy improves over time, that data is filtered in a particular way, or that sector-specific concerns are explicitly addressed. The benchmark is therefore consensus-driven: a solution is state of the art because it is widely used and accepted in its context, not because it satisfies a single technical test. This reading is anchored in Article 8(1) of the AI Act.[15]

[15] *Article 8(1) AI Act: High-risk AI systems shall comply with the requirements laid down in this Section, taking into account their intended purpose as well as the generally acknowledged state of the art on AI and AI-related technologies. The risk management system referred to in Article 9 shall be taken into account when ensuring compliance with those requirements.*

Comprehensive framework and compliance costs

Another concern that has been raised relates to the way in which the legislation has been drafted. The AI Act itself comprises a certain number of pages, whereas the accompanying guidelines, taken together, are much more extensive. This raises the question of whether this imbalance is desirable and whether there may be a risk of **over-regulation**.

In addition, another area where one might ask whether the AI Act is desirable is illustrated by the following example. One of the reviewers noted that large companies employ specialists to interpret the AI Act, but that different firms sometimes reach entirely different conclusions. This represents a major **bottleneck** in the development of regulations. The reviewer acknowledges that it is not possible to regulate everything comprehensively, but emphasises that the minimum applicable standards must be clear. Ultimately, those responsible for implementing the provisions need to know what is expected of them, as the aim is not to increase compliance costs.

6.3 Feasibility

Written contracts as a regulatory mechanism

Article 25(4) was generally seen as addressing a real issue, but as doing so with too few points of reference to operationalise it effectively. The comparison with the **Digital Operational Resilience Act** (DORA) is instructive: under DORA, the legislative text was complemented by detailed technical documentation that allowed legal and technical teams to identify obligations precisely, and in some cases even shifted bargaining power towards the client, with providers arguing that the requirements were too stringent. Article 25(4) sits closer to the opposite end of the spectrum according to reviewers: it imposes an obligation on both parties, but without making clear who is in the driving seat, what information and capabilities must be supplied, or how far the responsibility of upstream suppliers extends when the provider itself integrates and further builds the AI system.

A more **explicit regime for AI components**, including separate categories for suppliers, might have offered greater clarity, but would likely be too granular and restrictive given the diversity of AI components on the market. As currently drafted, the provision risks functioning as a catch-all or 'shotgun' clause: large companies may continue to offer their own templates and dare counterparties to challenge them, meaning that the practical effect could remain limited and outcomes may stay broadly similar to current market practice.

At the same time, swinging too far in the other direction would also be problematic, as overly strict regulation may lead to contractual requirements that providers cannot realistically comply with. The key issue is therefore one of scope and calibration. **Guidelines** clarifying what exactly must be supplied upstream would be welcome, and a lex specialis approach, possibly combined with **sector-specific guidance** or **the voluntary model terms** expected from the AI Office, could help to find a workable middle ground. That said, guidance should not accumulate to the point where it becomes unmanageable for non-lawyers. Overall, Article

25(4) can be regarded as a sufficiently well-drafted starting point, and possibly one of the clauses most likely to persist over time, but it needs further clarification if it is to have effects beyond existing market practice.

Another reviewer explicitly states that written agreements are not sufficient on their own without additional implementation conditions. Effectiveness depends heavily on the **implementation** conditions, while access to information in order to effectively implement the requirements is often restricted by trade secrets or is untransparent because the data was collected through web scraping. Compliance must therefore be shared across the whole value chain, rather than being placed on one party. However, the **strength** of a contractual approach remains that it can be tailored to specific parties and sectors and create binding obligations on the different actors. At the same time, the **weakness** lies precisely in the fact that these contracts often reflect a power imbalance. Large suppliers can dictate the contractual terms and offer little transparency, while providers may have no access to models, software or data.

According to one reviewer, a **hybrid regulatory** approach seems appropriate. The European Commission could develop sector-specific model clauses. This could then create agreements that are legally binding both on paper and in reality.

6.4 Recommendations

The feedback gathered through the expert interviews points to a number of concrete areas where additional guidance and clarification are needed to support the practical implementation of Article 25(4). While the provision establishes an important obligation, its current formulation leaves too much room for divergent interpretation. The recommendations below reflect the recurring concerns raised by reviewers and are intended for policymakers, the AI Office and other relevant authorities.

Across the interviews, reviewers consistently pointed to the need for additional information and guidance to support the implementation of the AI Act. Central to this effort is the clear and unambiguous formulation of **minimum applicable standards**. As it stands, Article 25(4) is difficult to negotiate and implement in practice, particularly for SMEs, which hold less bargaining power than large technology providers. The voluntary terms that the AI Office is expected to develop could provide meaningful support in this regard. These terms should also address the significant power asymmetry that can arise between providers and third-party suppliers, especially where the supplier is a large entity that places considerable reliance on trade secrets.

Equally important to highlight is the risk of **over-regulation**, partly due to the volume of supplementary guidelines or legislations beside the core text of the AI Act. It remains, therefore, important to not cause too many deviations from the existing legal instruments. Companies could use existing frameworks or contracts such as the ISO standards and DPAs for inspiration. The obligations set out in the AI Act for high-risk AI system providers can also be used as a starting point for the provisions of the written contract based on Article 25(4).

In addition, the current texts of the AI Act as well as other digital regulations must remain clear and enforceable without always requiring specialist expertise. It is therefore essential that the terms developed by the AI Office take the **target audience** into account. It may be advisable to provide generic model terms that are broad enough to be subsequently tailored by the specific body to its own use case. Alternatively, to allow for greater divergence and inspiration for the contract drafters, examples of more sector-specific model contract terms could also be included in the document.

Finally, the reviewers emphasise the importance of **legal certainty**: stakeholders must have a clear understanding of what is expected of them when implementing the provisions. The current reference to 'necessary information, capabilities, technical access and other assistance', as well as 'generally acknowledged state of the art', causes confusion. It would therefore be beneficial to provide more comprehensive guidelines on this issue. Doing so will streamline the implementation process.

Conclusion

7

Article 25(4) AI Act plays **a vital role** in **the value chain management function** of the AI Act. At various points in the AI value chain multiple parties supply AI systems, tools and services, as well as components or processes. These are then incorporated by the provider into the AI system with various objectives in mind. These parties have a significant role to play in the value chain towards the provider of the high-risk AI system into which their AI systems, tools, services, components or processes are integrated. Accordingly, the legislator stipulates that the third-party supplier should give this provider the necessary information, capabilities, technical access and other assistance in writing, based on the generally acknowledged state of the art, to enable the provider to fully comply with the obligations set out in the AI Act.

This report provides **comprehensive insights, suggestions and practical guidance** for policymakers, supervisory authorities, stakeholders and professionals navigating the complexities of the AI Act's value chain requirements. Through the development of policy prototypes, the report sheds light **on both sector-specific and general insights** into the concrete implementation of Article 25(4) and the difficulties that providers and third-party suppliers may face. These prototypes act as **practical examples** to explore how the obligation can be applied effectively in diverse contexts, taking into account the unique needs and challenges of each sector. Subsequently, the report provides detailed **comments and feedback on Article 25(4)** itself, aiming to inform and guide policymakers and authorities on difficulties as evaluated by the reviewers.

The implementation of Article 25(4) of the AI Act faces **various challenges in practice**, in particular due to asymmetries in bargaining power, unclear standards and limited guidance of the European Commission. Although the provision addresses a relevant issue, it remains too abstract to function effectively without further clarification, particularly for SMEs, which often lack both legal and technical resources.

A first priority is the development of **additional guidelines** that define minimum standards and expectations. Individual concepts such as 'necessary information', 'technical access' and 'assistance' need to be defined in concrete terms, preferably through **cross-sectoral and sector-specific guidance**. Closely linked to this is the need for standardised **voluntary terms**. This can strengthen the negotiating position of SMEs and contribute to more consistent market practices, without completely restricting the flexibility of the contractual relationship between the respective provider and third-party supplier.

The **division of responsibilities** between providers and third-party suppliers must be set out more explicitly. A functional approach, based on the actual role of the parties in the AI value chain, is necessary to prevent disputes and ambiguity. At the same time, this should not mean starting from scratch. **Existing instruments** such as SaaS agreements, ISO standards and sectoral guidelines already provide a solid foundation on which AI-specific obligations can be built, avoiding unnecessary complexity and promoting practical applicability.

Access to **information and transparency** within the value chain remain further concerns. The current reliance on trade secrets may hinder effective compliance and oversight. Contractual agreements must therefore contain clear provisions on documentation, audit rights and information exchange, while still striking a balance between the protection of the business' interest and legal requirements. In addition, clearly defining the obligations is essential so that each party knows what is expected of them.

Ultimately, the effectiveness of Article 25(4) depends on a **balanced regulatory approach**. Over-regulation risks stifling the flexibility that makes the provision workable across sectors, yet without sufficient detail it lacks practical impact. A hybrid model, in which the legislative framework is complemented by workable guidelines and voluntary terms, appears best suited to deliver both legal certainty and flexibility in implementation.

Acknowledgements

8

We would like to thank all participants whose contributions made this policy prototyping project possible. Your enthusiastic engagement and valuable insights were pivotal in shaping this project. Special thanks are extended to those who participated in the design workshop, investing time and expertise to draft mock documents. The collaborative efforts of everyone involved have been instrumental in the production of this report. Your commitment to advancing the discourse on policy prototyping in the field of AI and data policy is genuinely appreciated.

Project participants

Prerna Chaudhary – AI Policy Advisor – Municipality of Utrecht
Lotte Cools – IT Lawyer – Timelex
Ana Carolina Costa – Legal Counsel – Randstad
Arnoud Engelfriet – IT Lawyer – ICTRecht
Hilal Ersahin – IT Law – Attorney – Ersa Legal
Giovanni Garofalo – Doctoral Researcher - Institut Supérieur de l'Aéronautique et de l'Espace (ISAE-SUPAERO)
Dino Gliha – IP Lawyer sub-specialisation in the field of AI
Ingrid Lambrecht – Digital & AI Domain Lead – Legile
Sophie Meszaros – Interoperability Strategy Lead and Researcher
Wannes Ooms – Legal Counsel
Raf Schoefs – Senior Counsel – KPMG Law
Stefanie Stappers – IT/IP Lawyer – Monard Law
Angeliki Tiligadi - Head of Privacy, DPO & AI Governance Officer – Qover
Wouter Torfs – Attorney – Timelex
Matthias Vandamme – Attorney/PhD Candidate – Academic assistant – UGent
Stan Vander Sande – Legal counsel – Legal Department LRD KU Leuven
Danny van Roijen - Compliance manager and EU Policy Expert

Knowledge Centre Data & Society team

Sultan Erdogan – Researcher – Centre for IT and IP Law – KU Leuven/KCDS
Frederic Heymans – Research Associate – imec-SMIT, VUB/KCDS
Shannen Verlee – Researcher – Centre for IT and IP Law – KU Leuven/KCDS
Koen Vranckaert – Researcher – Centre for IT and IP Law – KU Leuven/KCDS

Citation

Sultan Erdogan, Shannen Verlee, Frederic Heymans and Koen Vranckaert (Knowledge Centre Data & Society), “From Policy To Practice: Prototyping the EU AI Act’s value chain requirement”, July 2026.

Contact
sultan.erdogan@kuleuven.be and frederic.heymans@vub.be.

data-en-maatschappij.ai/en/ →

Kenniscentrum Data & Maatschappij

Artificiële Intelligentie Vlaanderen