

# Proposal Dutch ‘regulatory sandbox’ under the AI Act

Based on the results of the 2023 preliminary investigation and the 2024 pilot

## Date

13<sup>th</sup> of March 2025

**The AI Act<sup>1</sup> requires Member States to establish at least one AI regulatory sandbox from August 2026.<sup>2</sup>** Within a regulatory sandbox, supervisors provide support to providers of AI systems who, during development, experience difficulties with how they can comply with the AI Act; ensuring that they can place their product on the EU market in accordance with the AI Act. To lower the burden on compliance, supervisory authorities’ guidance should clarify the interpretation of the AI Act. In addition, supervisors gain better insights and understanding of the opportunities, emerging risks and the impacts of AI use. Supervisory authorities are innovation-friendly and flexible throughout the sandbox process whilst existing regulations remain in force.

**This document is a proposal for the Dutch implementation on the regulatory sandbox, published on behalf of the Dutch Authority for Digital Infrastructure (‘Rijksinspectie Digitale Infrastructuur’, RDI) and the Dutch Data Protection Authority (‘Autoriteit Persoonsgegevens’, AP).** This proposal describes the foundational principles of the Dutch sandbox and an effective process design. It is the result of an intensive collaboration between various supervisory authorities and ministries, under the coordination of the Ministry of Economic Affairs (EZ), the AP and the RDI. Within this collaboration, a sandbox-pilot was executed in which 40 questions from 25 different organizations were examined and discussed with stakeholders from the Dutch AI ecosystem. This published proposal is a simplified and shortened version of the final report drafted for the departments.

**The aim of this proposal is to provide a solid basis for the further collaboration and alignment needed to launch a final sandbox by August 2026, or earlier.**

**This proposal is closely aligned with the Dutch final recommendation on the regulatory governance for the AI Act.<sup>3</sup>**

**The proposal takes the different regulatory roles stemming from the AI Act into account.** In the Dutch context, it is to be expected that the implementing act will designate different market surveillance authorities for the different sub-areas of supervision based on the national context and the areas named in the AI Act. In addition, in order to facilitate collaboration, it is expected that one or more authorities will be designated with coordinating roles as ‘coordination market surveillance authorities’. Finally, AI Act supervision interplays with existing sectoral, domain-specific and fundamental rights supervision, which can therefore also be designated a role within AI supervision.

**This Dutch proposal is drafted while the implementing acts for the sandbox from the European Commission are not yet established.** These implementing acts will set out principles on the selection criteria, application and participation procedures and the conditions applicable to participants.<sup>4</sup> However, through the AI Board, The Netherlands will be consulted in the drafting process of these implementing acts. The Netherlands will contribute to this along the lines of the principles set out in this proposal. Should these implementing acts affect certain parts of this proposal, the final implementation of the sandbox will be aligned accordingly.

---

<sup>1</sup> The AI Act is a product regulation that imposes requirements on different categories of AI systems, depending on the risk. These requirements must be met before an AI system can be placed on the market.

<sup>2</sup> Article 57(1) AI Act.

<sup>3</sup> [Final recommendation on supervision of AI \(AP & RDI\)](#)

<sup>4</sup> Article 58(1) AI Act.

## Proposed principles Dutch sandbox

The regulatory sandbox under the AI Act should contribute to improving legal certainty, innovation and better regulation, and the development of good practices. It is important that this new tool in which supervisory authorities and AI developers meet functions well, and that everyone contributes to the goals of the sandbox through their designated roles. The following principles are therefore proposed for the Dutch sandbox:

1. **Every supervisory authority under the AI Act participates in the Dutch sandbox.**<sup>5</sup> The AI Act affects a large number of sectors and domains, and it is expected that a multitude of Dutch supervisors will have a role in supervision.<sup>6</sup> The proposal is to create a single regulatory sandbox in which all relevant supervisory authorities<sup>7</sup> participate, resulting in that questions from all sectors are answered in an efficient way. General questions about the AI Act are tackled centrally and questions that require specific expertise are dealt with by the relevant authority.
2. **The Netherlands will have one multi-sectoral sandbox. All relevant (market surveillance) authorities involved in the AI Act are involved in this sandbox.** Access to the sandbox is thus granted via a single access point. The sandbox and its processes must be easily accessible for AI-providers. Due to the involvement of all relevant supervisory authorities, AI providers do not have to figure out by themselves who they should contact. At the same time, as a result, supervisory authorities will jointly come to a consistent interpretation of the AI Act.<sup>8</sup>
3. **Roles of supervisory authorities within the sandbox are commensurate with their national role under the AI Act.** The role of a supervisory authority under the AI Act is the basis for their role within the sandbox. Supervisory authorities with a coordinating role under the AI Act will fulfil that same role in the sandbox. Market surveillance authorities<sup>9</sup> designated for a specific scope will be responsible for answering questions corresponding to the same scope within the sandbox. Supervisors who provide additional horizontal or domain-specific perspectives into AI supervision will also provide this into the sandbox.
4. **The Dutch sandbox focuses on supporting compliance with the AI Act. AI-providers remain solely responsible for the end-product compliance.** The NL sandbox focuses primarily on support through providing regulatory guidance. This aligns with the role of a supervisory authority and is an efficient use of resources. This support includes providing legal and/or technical guidance on activities carried out by a provider or developer to enable an AI system to comply with the requirements of the AI Act. Supervisors provide guidance, but do not help providers to effectively make the AI system comply with the requirements. The provider always remains responsible for achieving compliance of the AI product with the AI Act<sup>10</sup>. The Dutch Sandbox does not provide any form of certification.
5. **Supervisory authorities guide developers with their expertise and do not provide physical or technical infrastructure themselves.** Setting up such an infrastructure is capital intensive, influences eligibility and only helps a limited number of parties. Such services are not mandatory under the AI Act. In addition, there are several existing instruments that offer these services, such as the Testing and Experimentation Facilities (*see principle 6*).
6. **The sandbox should link as closely as possible to tools that address other needs of AI-providers, but remains independent of them.** The sandbox has a specific role vis-à-vis other innovation-enhancing instruments such as the Testing and Experimentation Facilities<sup>11</sup> or the European Digital Innovation Hubs<sup>12</sup>. Further investigation is

---

<sup>5</sup> Given the independence of the judiciary and its potential impact on market surveillance of the AI Act, it may be necessary to develop a separate regime for the judiciary. However, no further explanation can be given at the time of writing this proposal.

<sup>6</sup> The final recommendation on AI supervision identifies several supervisors that should be given a role in the supervision of the AI Act: [Final recommendation on supervision of AI \(AP & RDI\)](#).

<sup>7</sup> These are supervisory authorities supervising AI-providers based on the AI Act.

<sup>8</sup> In relation to this topic, the remark made in note 5 regarding a possible separate regime for the judiciary still applies.

<sup>9</sup> The authority responsible for supervision as referred to in Regulation (EU) 2019/1020. The term 'market surveillance authority' is consistent with the terminology used in Article 3(26) of the AI Act. In this document, it also refers to the 'national competent authorities' as referred to in Section VI of the AI Act, which will be tasked with setting up and operating a sandbox.

<sup>10</sup> In addition to the fact that such far-reaching support does not fit the role of supervisors, this could also create competition with existing private (legal) services.

<sup>11</sup> [Sectorial AI Testing and Experimentation Facilities under the Digital Europe Programme - Shaping Europe's digital future \(europa.eu\)](#)

<sup>12</sup> [European Digital Innovation Hubs – Shaping Europe's digital future \(europa.eu\)](#)

needed to determine the extent to which these instruments can provide supplementary services that the sandbox does not offer, such as a technical or physical infrastructure. It is important to communicate this clearly to the market and to refer to other instruments where necessary.

7. **The Dutch sandbox welcomes questions that are not eligible for an intensive sandbox-procedure.** The minimum requirement under the AI Act is that AI-providers can go through an intensive sandbox-procedure with a designated supervisor. In this procedure supervisors and AI-providers have a dialogue to explore and discuss the submitted question(s) in depth. Such a procedure is suitable for questions that are complex enough to justify a generous use of resources. During the investigation and pilot, however, we concluded that (a) there is a need by organizations to receive answers to questions that are not eligible for an intensive sandbox-procedure, (b) that these answers are useful for supervisory authorities in their general supervisory role of providing guidance and to understand which challenges the market faces and (c) that the additional resources these questions cost can be justified because of the added value for AI-providers, supervisors and the broader AI ecosystem. Furthermore, in order to determine which questions are relevant for an intensive sandbox-procedure, it is necessary for market surveillance authorities to carry out a selection. The Dutch sandbox therefore includes questions that are not eligible for an intensive sandbox-procedure, by answering these questions directly in writing. When operating under this principle, it is important that answers are shared publicly where possible. This will further contribute to the responsible development of AI and facilitate compliance with the law. In addition, the Dutch Sandbox will align with the broader goals of (future) market surveillance authorities regarding providing guidance on AI Act.
8. **Questions can be submitted continuously in the Dutch sandbox, but it is important to allow adjustments in order to meet expectations.** If the resources available are exceeded, it will be possible to limit submissions by, for example, temporarily closing the sandbox or by setting a limit on eligible questions.
9. **The Dutch sandbox communicates within which timeframe questions will be answered, depending on the complexity of the questions.** No rights can be derived from these estimated terms. To facilitate this, the Core Team will set requirements for submission to ensure effective and efficient answering of questions for AI-providers.<sup>13</sup>
10. **The Dutch sandbox is driven by the needs of the market.** At the start of the sandbox, as much as possible is kept open what kind of questions can be submitted, and how complex these questions are.
11. **The sandbox focuses on guidance on compliance with the AI Act.** Questions about other legislation are only considered if they are relevant in the context of the AI Act, or if the AI Act requires it to be so.<sup>14</sup>
12. **Over time, the design of the sandbox implementation can adapt to the needs and maturity of the market and regulators.** After an initial operational period, the functioning of the sandbox is (periodically) evaluated and, if necessary, adjusted. Needs and maturity are expected to develop on both sides, and the sandbox can be adapted accordingly.

---

<sup>13</sup> This includes the supervisory authorities involved in the sandbox. In addition, this depends on the implementing acts of the European Commission.

<sup>14</sup> This is the case for sandbox involvement of the data protection authority, in line with the obligation to do so in the AI Act.

## Proposed process design Dutch sandbox

Based on the principles and insights, supervisory authorities propose the sandbox-process as described in this chapter. This proposal consists of three pillars: (i) a Core Team for the coordination and processing of relatively simple questions, (ii) an expert pool to answer questions which require specialized expertise, (iii) and sandbox-procedures in which supervisory authorities and participants meet repeatedly to explore and discuss questions in depth. Within this process, the Dutch sandbox consists of six phases:

1. **Pre-registration**  
The Dutch sandbox provides the possibility to submit questions via an online portal, managed by a portal-team.
2. **Application, selection and first response**  
Questions are accepted, triage takes place to identify the different types of questions, and relatively simpler questions are immediately answered in writing.
3. **Preparation**  
The AI-provider and the relevant supervisor(s) prepare an intensive sandbox -procedure.
4. **Participation sandbox-procedure**  
The AI-provider and the relevant supervisor(s) will start the process to further explore the question on compliance with the AI Act.
5. **Review & completion**  
The supervisory authority, in consultation with the provider, drafts the final report containing the learning outcomes and discusses the form of disclosure of the results.
6. **Post-participation**  
The supervisory authorities publish the learning outcomes where possible.

The proposal for the Dutch sandbox is schematically presented below:

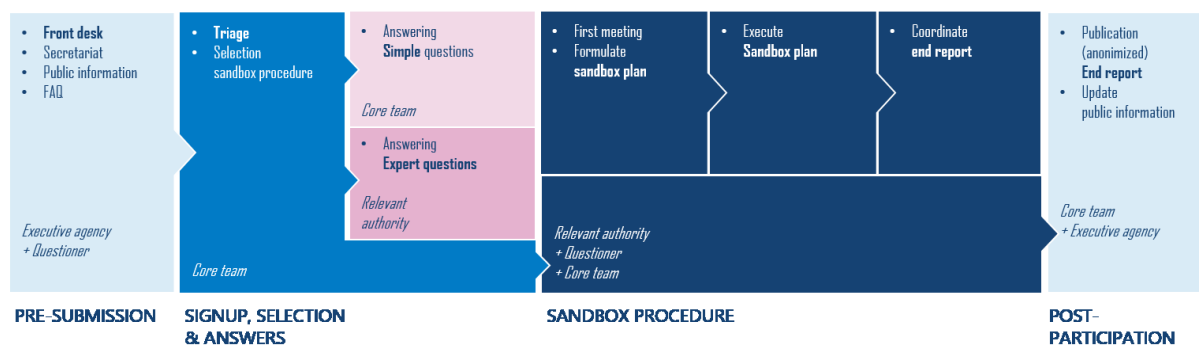


Figure1. Schematic overview Dutch sandbox.

Phases based on European Commission presentation 'Key phases of the sandbox process'.

The phases are described in more detail in the remainder of this chapter. The example below (Box 1) provides an insight on how a question can be answered within the Dutch sandbox.

### Example: How is a question addressed in the different sandbox phases?

An AI developer develops an AI system for a shipping company to make predictions about the maintenance of their ships.

#### *Pre-registration*

The developer submits a question in **the portal, managed by a portal-team**. The portal-team assesses whether the question meets the requirements: does the question the AI Act and whether the answer is already publicly available. If the question is relevant, the portal-team forwards the question to the Core Team.

#### *Registration, selection & first answer*

**The Core Team** determines whether they can answer the question themselves, whether sector-specific knowledge is necessary to answer the question, or whether the question is feasible for a intensive sandbox-procedure. If they classify the question as **an expert question**, they will contact the **relevant market surveillance authority and domain or sector specific supervisor(s)**. In this example, the Core Team contacts the Human Environment and Transport Inspectorate (ILT). The ILT will answer the question and, only if necessary, contact the AI-provider via the Core Team for clarification. The ILT formulates the answer in writing and forwards the result to the Core Team, after which they share the answer with the AI-provider.

#### *Preparation and participation sandbox process, evaluation and completion & post-participation*

If the question is considered feasible for an intensive **sandbox-procedure**, the Core Team will involve the relevant supervisor(s) (in this case at least the ILT) and after which the provider will be contacted. Subsequently, they will jointly carry out an intensive sandbox-procedure on the basis of a mutually drafted plan. Afterwards, the lessons learned are written down in a report.

#### *Box 1: Example process question regulatory sandbox*

## 1. Pre-registration

This phase functions as a portal, allowing AI-providers to approach the Dutch sandbox and access the public information. In essence, the portal serves as a platform, in the form of a website, containing information on the Dutch sandbox, the AI Act, available guidelines and FAQs. This portal (i.e. portal) is not necessarily maintained and operated by the supervisory authorities; it might better suit an executive agency. Supervisory authorities promote the sandbox through their own communication channels. In addition, the sandbox will be promoted to attract a diverse range of stakeholders. These promotional activities aim to raise awareness of the sandbox and increase the variety of participating AI developers, with a specific focus on SMEs and start-ups.

The submission of questions by AI-providers via the website falls under the pre-registration phase. These questions are processed by the portal-team before being forwarded to the supervisory authorities. Before forwarding, repeated questions or questions that obviously do not fall within the scope of the Dutch sandbox, are filtered out. These are, for example, questions that concern other legislation than the AI Act or questions which are not clearly enough related to specific AI product(s) being developed. Relevant questions are forwarded to the Core Team of supervisors.

Fundamentally, no costs will be charged for the submission of questions, nor for participation in the intensive sandbox process. However, in order for a question to be eligible, a question must meet certain requirements. Although the implementation fosters accessibility, some minimum requirements are necessary to ensure the quality and efficiency of the sandbox, see box 2.

## 2. Application, selection and first response

After the pre-registration phase, the portal forwards the questions to the Core Team. This team consists of the coordinating supervisory authorities. A characteristic of the Dutch sandbox is the extensiveness of the *application and selection* phase; questions that are forwarded by the portal, but that are not eligible for an intensive sandbox process, will receive guidance in writing by the supervisory authorities nonetheless (see Figure 1 for a schematic representation). Once questions are received by the Core Team, they will provide guidance on questions that are suitable for *easy answering*. In case of no objection by other supervisors involved in the sandbox implementation, the answer is forwarded back to the AI-provider. However, occasionally a question may not be suitable for immediate answering, this applies to two types of questions:

1. **Expert questions:** Some questions require more time or sector- or domain-specific knowledge. Experts at involved supervisory authorities will be consulted. Expert questions will be answered in writing. Short contact with the AI-provider can sometimes be part of the process.
2. **Sandbox-procedure questions:** Some questions are more complex and/or of societal relevance, so that it is more feasible to answer them within an intensive sandbox-process. This entails that the question is discussed and explored in several sessions together with the AI-provider. Eligible questions will continue to the *preparation* phase described in the next section.

An expert question may prove suitable for further deepening in a sandbox-procedure during the process. In that case, the supervisors will jointly discuss this together with the AI-provider. Furthermore, all provided guidance could be published, after anonymization, via the portal, for example by converting them into FAQs.

The process within the *application, selection and first response phase* is a form of triage by the Core Team that meets periodically. Questions that are not eligible for a intensive sandbox-procedure are identified and prepared to deal with in writing. Questions eligible for an intensive sandbox process are identified. Within the triage of questions, submission and selection criteria apply, see box 2 below.

### **Submission requirements**

In order to be eligible, a question must meet certain submission requirements. This ensures the quality and efficiency of the sandbox. Submission requirements shall be proportionate and shall ensure that access to the sandbox is as fair as possible. The requirements may change over time in line with trends and development of AI Act understanding of AI-providers. However, at all times, the submission requirements will be in line with the forthcoming implementing acts of the European Commission.

For example, AI-providers may be required to provide a clear description of the AI system and its context of use. A second requirement may be that the AI-provider provide their view on the AI Act themselves, and indicate how the question applies to the AI system under development.

### **Selection criteria**

In addition to submission requirements for eligibility of questions in the Dutch Sandbox, selection principles are necessary for situations where the amount of questions exceeds the available capacity of the supervisory authorities.

### **One-time answer or expert question**

Questions for a written procedure shall, in principle, be processed in order in which they are received. In case of a large number of questions, supervisors can allocate priority to questions based on the selection criteria that also apply to for an intensive sandbox-procedure. The sandbox can be temporarily closed if the number of questions exceeds the available capacity of the supervisors.

### **Sandbox-procedure**

The supervisory authorities are free to base their efforts on the transparent selection principles that will be published on the website of the Dutch sandbox. The selection principles are the objectives of the regulatory sandbox as set out in Article 57(9) of the AI Act, and the possible additions thereto in the European Commission's implementing act. In doing so, supervisors may grant priority access to SMEs established in the Union, including start-ups, on the basis of Article 62(1).

Questions dealt within a sandbox-procedure must relate to AI systems that must (possibly) meet requirements under the AI Act. These are high-risk AI, AI with transparency risks and general-purpose AI systems (to the extent that the AI Office does not supervise the latter).

In addition, supervisors should be able to take the maturity of the AI system into account in relation to the activities in the sandbox process, such as, assisting with the final question articulation during the preparation of the sandbox plan. Asking a good question, by matching it to the maturity of the product, is a prerequisite for effectively going through the process. Providers may, e.g., be required to have the technical and operational capacity to go through the process and to (want to) make efforts to ensure compliance with legal obligations.

*Box 2: Submission requirements & selection criteria*

### 3. Preparation

This third phase serves as preparation for the execution of the sandbox-procedure. To this end, the appointed supervisor responsible for the procedure contacts the AI-provider for an introduction, further in-depth examination of the question and discusses what the sandbox procedure may look like.

After the first interview, agreement could be reached to start a sandbox-procedure. In that case, the responsible supervisor and the AI-provider shall draft a sandbox plan, as defined in the AI Act. The sandbox plan shall contain at least the objectives, conditions, timetable, methodology and requirements for the activities to be carried out. The sandbox plan can also be used for further elaboration and clarification of the question asked in order to achieve a useful process for both the AI-provider, supervisor and the broader AI ecosystem. In addition, a link with other facilities from the AI ecosystem can be made (see box 3). Agreements are made about disclosing the insights gained.

It can also be determined that a question is not suitable for a sandbox process, or that it can be addressed in writing without an intensive process.

#### **Link with other initiatives in AI ecosystem**

When setting up the implementation of the Dutch sandbox, other organizations will be identified which can be involved incidentally or structurally in the regulatory sandboxes. Examples are the Testing and Experimentation Facilities, European Digital Innovation Hubs, European Data Spaces, the ELSA labs, or the so-called notified bodies. Such a link provides AI-providers the opportunity to be assisted, inside or outside the sandbox, with matters outside the scope of the sandbox. For example, by providing technical infrastructure or data through other initiatives.

*Box 3: Link with other initiatives in the AI ecosystem*

### 4. Participation Sandbox-procedure

In the fourth phase, the AI-provider and supervisor(s) execute a sandbox-procedure, in accordance with the agreed upon sandbox plan, in order to achieve a final report in which the question asked has been answered to the best extent possible.

The supervisor can provide legal and technical expertise to support the testing of the AI system against the requirements of the AI Act. The supervisory authorities do not test technical specifications. In addition, testing technical specifications that are not related to the AI Act fall outside the scope of the sandbox in all cases.

The regulatory sandbox does not provide exemptions to legal obligations. Supervisory authorities are obliged to be innovation-friendly. Based on this principle, AI-providers are expected to conduct corrective measures of possible infringements with Union or national legislation that are identified within the sandbox. If not executed to the satisfaction of the supervisor, the sandbox-procedure can be terminated. This agreement is recorded in the sandbox plan.

#### **Testing under real world conditions**

The AI Act facilitates the testing of high-risk AI systems listed in Annex III in real world conditions, without classifying the system as placed on the market or put into service. The aim is to collect reliable and robust data and to assess and verify whether the AI system is performing according to the requirements of the AI Act, before it is officially placed on the market. If testing performed under real world conditions within the framework of a sandbox-procedure, the AI-provider and the responsible supervisor shall agree on conditions and appropriate safeguards to protect fundamental rights, health and safety. In view of the conditions imposed by the AI Act on testing in real world conditions *outside* a sandbox-procedure, testing in real-world conditions *within* a sandbox-procedure should also provide the same level of protection.

*Box 4: Testing under real world conditions*

## **5. Evaluation & Completion**

The fifth phase is used to draft the final documents of the sandbox-procedure. This includes at least the final report describing the activities carried out, the results and any learning outcomes. A written proof of the activities carried out in the sandbox can also be drafted for the purpose of the subsequent conformity assessment for a product.

## **6. Post-participation**

The sixth phase ensures that the learning outcomes of the sandbox are accessible and available for further distribution. This supports the broader AI ecosystem. Supervisors make efforts to share the learning outcomes of the questions answered in writing in an anonymised form, for example by creating FAQs. The final reports of the sandbox trajectories will be published if agreed upon with the participant. An alternative is an anonymised public summary with the lessons learned.



## Proposed roles Dutch sandbox

The proper functioning of the Dutch sandbox requires different levels of commitment from the parties involved at each stage. Fundamentally, it aims to align with the mandates of supervisory authorities as will be laid out in the to be adopted implementing act. This results in that all parties must build up capacity in accordance with their designated role, to be able to fulfil their role effectively and responsibly.

### **Involved supervisors: expert questions & sandbox-procedures**

The Core Team proposes, in consultation, which organisation is most suitable for an expert question or sandbox-procedure, and be designated as responsible supervisor for that process. The responsible supervisor is responsible for drafting a response and the necessary coordination. In case of a sandbox-procedure, it is also the point of contact for the AI-provider. The responsible supervisor also ensures that other supervisor(s) are involved in the process, if domain-specific knowledge is necessary. The Core Team supports this process.

To assign the role of the supervisor responsible for the process, the Core Team adheres to the powers as laid down under the AI Act, as well as other legislation. The supervisor whose powers best align with the question, will be tasked with supervising the procedure. In practice, this results in that supervisors designated as market surveillance authorities under the AI Act can be designated as supervisor responsible for the process. Sector or domain-specific supervisors can be involved to support this, if their specific knowledge can assist market surveillance authorities explain the law, provided that they have the capacity to do so.

The answers to all types of questions, including those answered by the Core Team, will be accessible to all supervisors involved in the sandbox, before being sent to the AI-provider. This allows these supervisors to provide comments or remarks and suggest adjustments. This is important to ensure the quality and consistency of all answers given. Furthermore, the Core Team can contact the AI-provider at any time.

*Box 5: Supervisors involved in answering questions*

Below, the different roles within the sandbox are further explained. It also describes how they relate to the supervisors' recommendation on the supervisory structure for the AI Act.<sup>15</sup>

### **Core team**

The Core Team consists of the envisaged Coordinating Market Surveillance Authorities. They answer the simpler questions, coordinate the allocation of expert questions and the start of sandbox-procedures. In addition, from their coordinating role, they play a pivotal role within the sandbox. The Core Team has an important role in knowledge-sharing, both towards supervisors and towards the market: knowledge gained during sandboxing is proactively shared. In addition, the Core Team is responsible for the annual (public) reporting on the sandboxes. Within the Core Team, both legal and technical AI expertise is permanently present to facilitate this collaboration effectively.

### **Market surveillance authority**

Based on the recommendation of the supervisory authorities, a market surveillance authority is designated for specific product groups (as listed in Annex I and III of the AI Act).<sup>16</sup> In order to align as best as possible with the recommendation, the same market surveillance authorities will be responsible within the sandbox for answering expert questions and guiding sandbox trajectories.<sup>17</sup> They contact AI-providers, coordinate sandbox trajectories and read along with all written answers to ensure quality and consistency. Each market surveillance authority involved is expected to appoint a point of contact for the Core Team. This person can work on the questions themselves or forward them a colleague best suited to take on the question.

### **Sectoral & domain-specific supervisors**

Existing sectoral and domain-specific supervisors play an important role in the supervision under the AI Act, as they (possibly) already supervise (the effects of) AI-providers within their sector or domain. In the sandbox, they are invited when they can add relevant knowledge to answer a question. For example, market surveillance authorities can consult

<sup>15</sup> [Final recommendation on AI supervision: Sectoral and centrally coordinated Authority for Personal Data](#) (Autoriteit Persoonsgegevens)

<sup>16</sup> [Final recommendation on AI supervision: Sectoral and centrally coordinated - Dutch Authority for Digital Infrastructure](#)

<sup>17</sup> (Coordinating) market surveillance authorities are also the competent authorities as referred to in section VI of the AI Act.

them if the answer to a question requires specific sector or domain knowledge. The involvement of a supervisor in the sandbox thus reflects their broader role within the supervision under the AI Act. Therefore, the extent to which these supervisors can be involved depends on the their within a broader perspective of supervision under the AI Act .

#### **Online portal team**

The portal-team manages the sandbox's website, is responsible for setting up and managing a contacts relation management (CRM) system, and filters ineligible questions. Exact requirements and responsibilities are still being explored. For example, it is being investigated if and which organisation may be best suited to manage the portal.

#### **Fundamental Rights Authorities**

The AI Act assigns a specific role to fundamental rights authorities identified under Article 77 of the Act. In order to strengthen these fundamental rights authorities, specific powers and obligations on requests for information are included and market surveillance authorities should collaborate with those authorities if risks on fundamental rights are identified. For example, fundamental rights authorities may request market surveillance authorities to investigate an AI system, if they suspect possible violations of fundamental rights. Furthermore, market surveillance authorities should inform fundamental rights authorities if they identify fundamental rights risks or receive reports of serious incidents. However, at the time of writing, specific supervisory roles under the AI Act are not clear yet in the Netherlands. The premise is that the supervisors' roles in the sandbox will reflect their broader role within the supervision on the AI Act.

#### **Data Protection Authority**

Under the AI Act, the Dutch DPA is involved in the sandbox if participating AI systems involve the processing of personal data.<sup>18</sup>

#### **Interdisciplinary expertise**

It must be noted that for effective use of available capacity, interdisciplinary expertise is required. Although the majority of the capacity is of legal nature, the Dutch sandbox requires a good balance between legal and technical (AI) knowledge. In the sandbox-procedure, technical (AI) knowledge is necessary to facilitate a good functioning of the regulatory dialogue. This is necessary to comprehensively understand the AI product, and understand the underlying question. Besides the sandbox-procedure, technical (AI) may be required to answer simple and expert questions.

---

<sup>18</sup> Article 57(10) AI Act.