Beyond the AI Act

Responsible AI Legal & Ethical Guide — July 2026

Center for Responsible AI

VIEIRA DE ALMEIDA

Table of contents

Table of contents

III. Placing AI systems on the market

IV. Deploying AI systems

Deployment of the AI system – use conditions

Data to deploy the AI system

Governance

Compliance and effects of non-compliance

IP protection of the AI output

Particularities for organisations, sectors and consumers

Algorithmic collusion

V. Measures for AI innovation and uptake

I.

Inside this Guide

A. Why this guide now

The AI Act is moving from an abstract legislative framework to a concrete compliance reality, with organisations now preparing for its application. At the same time, the updated AI Act under the Digital Omnibus on AI reshapes key elements of the regime, adjusting the treatment of sectoral products and safety components, as well as core aspects of governance, timelines and penalties. Because AI also sits within a wider set of EU regimes on product safety, resilience and sustainability, data (personal and non personal), intellectual property, digital services, accessibility, cybersecurity, competition, contracts and consumer protection, any serious assessment of obligations for developing and using AI must go beyond the AI Act alone.

Organisations therefore face the challenge of applying dense, highly technical rules to real projects and, in many cases, revisiting assumptions built on the original AI Act. These developments, together with early implementation experience and this wider legal patchwork, have brought into focus recurring pain points: identifying which systems are in practice high risk, managing overlaps and tensions with sectoral product, data and cybersecurity rules, allocating responsibilities along complex value chains, and implementing demanding obligations in a coherent way.

This Guide is published now to address those practical challenges highlighted by the first applications of the AI Act and evolving approaches to AI regulation, while also offering a broad overview of the main regulatory obligations under the AI Act and other applicable laws. It comes at a particularly sensitive moment, when organisations are beginning to translate high level strategy into concrete implementation choices against the backdrop of a changing AI Act.

This Guide is thus meant to sit where law and practice now meet: it translates a moving and fragmented regulatory landscape into concrete choices, tradeoffs and questions that organisations must confront as they design, build and use AI systems.

B. Who should read this guide

This AI Guide is intended primarily for:

AI providers

Understood as organisations that develop AI systems (or have them developed) and place them on the market or put them into service under their own name or trademark.

AI deployers

Namely organisations that use an AI system under their authority in a professional context, excluding purely personal non-professional use.

Although the focus is on AI providers and deployers of AI systems, many other actors in the AI ecosystem may also find this Guide relevant, including:

- Product manufacturers.
- AI vendors, including sellers, licensors, importers, distributors and resellers.
- AI marketplaces.
- AI component and service suppliers.
- Data providers, data analytics providers and data marketplaces.
- Infrastructure providers, such as cloud and computing.
- Providers of other inputs, such as hardware, chips and raw materials.
- AI providers and deployers of AI models.

In addition, this Guide is useful for internal stakeholders involved in AI-related decision-making, such as product and engineering, technology and systems, quality, procurement, marketing, risk and security, data science and analytics, legal, compliance, ethics and management teams.

This Guide will help you:

- **Understand the key legal and regulatory rules that may apply to your activity and ensure legal compliance — with a focus on EU law.**
- **Identify the right questions to ask, strengthening your self-sufficiency and enabling a more structured and forward-looking approach to AI.**

By using this Guide, you will become better equipped to develop and use AI in a successful and compliant manner.

C. AI Covered in this guide

This Guide is limited to legal and regulatory guidance on **AI systems**. AI models are expressly excluded. While AI models, including general-purpose AI models, constitute a key technical component of AI systems, this Guide does not seek to provide a comprehensive treatment of AI models as stand-alone objects. References to AI models are included only as relevant for the guidance provided on AI systems.

AI system is defined in the AI Act as a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.

The European Commission Guidelines on the definition of an AI system (2025) further explains each of these elements.

An AI system designed to perform confined and specific (narrow) tasks is commonly labelled as **narrow-purpose AI or specific-purpose AI**. Over time, AI systems evolved and can now handle many different tasks. Such systems are commonly labelled as "general-purpose AI systems" ("**GPAI system**").

A new category of **agentic AI** is emerging to describe systems that act with a significant degree of autonomy, pursuing clear goals rather than executing isolated, step-by-step instructions from humans. Whereas AI agents are often conceived as single systems that autonomously carry out specific tasks, agentic AI goes further by planning, deciding and executing multi-step workflows, using tools, services or other systems, and by perceiving aspects of their environment and adapting over time. Although early deployments have already automated substantial parts of business processes, fully autonomous agentic AI systems capable of independently managing complex operations at scale remain largely experimental.

Prior to the creation of a system, it is necessary to develop an **AI model**: the essential part of an AI system, used to make inferences from inputs to produce outputs – where the AI algorithm lays. Indeed, as clarified by the AI Act, "AI models are essential components of AI systems, [but] they do not constitute AI systems on their own. AI models require the addition of further components, such as for example a user interface, to become AI systems. AI models are typically integrated into and form part of AI systems".

Types of AI models include generative models, foundation models, natural language processing (NLP) models, large language models (LLMs), image and voice generation models and computer vision models. An AI model designed to perform confined and specific (narrow) tasks is commonly labelled as **narrow-purpose or specific-purpose AI model**, while AI models that handle many different tasks labelled as "general-purpose AI models" ("**GPAI model**").

The AI Act defines a GPAI model as "an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development or prototyping activities before they are placed on the market".

This Guide covers all stages of the AI value chain, as indicated below.

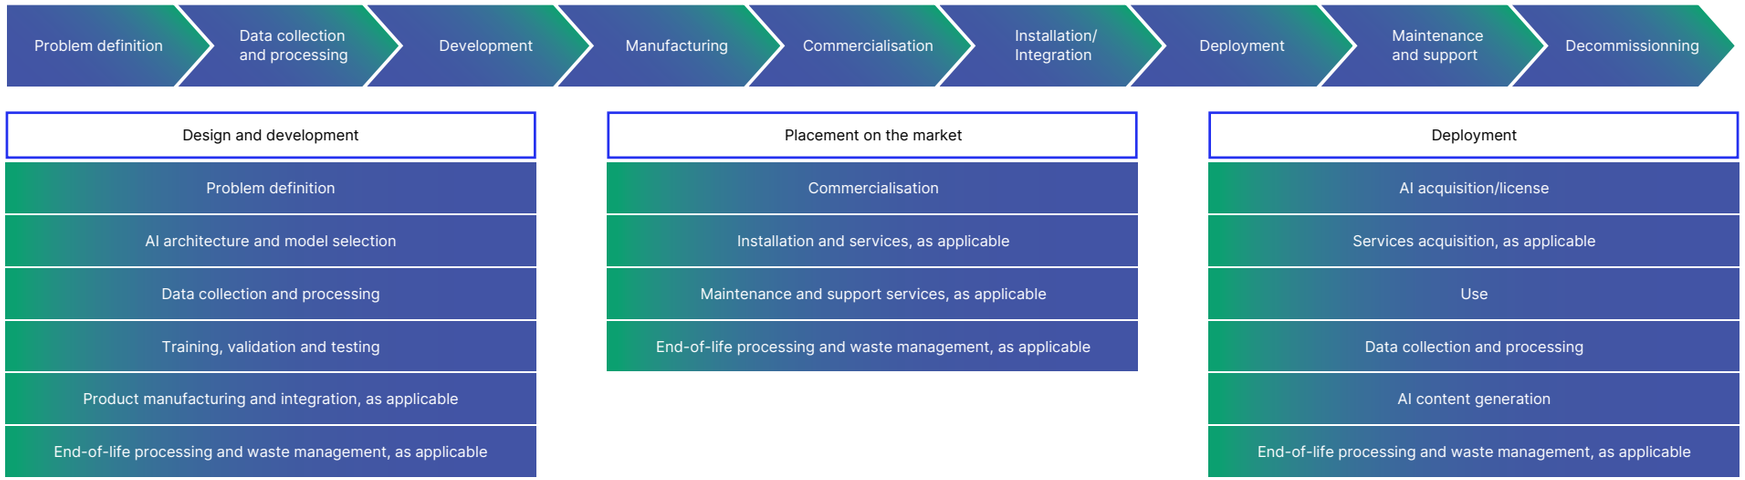

D. How to use this guide

This AI Guide is divided into three main chapters covering the main stages of the AI value chain:

- AI development
- AI commercialisation
- AI deployment

Each of the chapters presents a set of questions on legal and regulatory topics, and, where relevant, on ethics matters.

For each question, a **Checklist and Risk Mapping** is proposed to help organisations understand how to respond to the requirements addressed in each question. Each Checklist and Risk Mapping covers goals, an indicative list of questions organisations must answer, examples of the expertise required and the relevant risk mapping.

The risk mapping considers:

- The **likelihood of the risk** – how probable it is that a harmful effect will occur given the way the AI system is designed and deployed, and who is affected by the harmful effect. It thus relates to probability and exposure. Levels of likelihood can be rare, possible, likely or highly likely
- The **severity of the risk** – how serious the consequences would be if the harmful effect occurs (e.g., minor inconvenience vs. exclusion from education or loss of employment; brief interruption vs. sustained reduction of service quality; delayed reporting vs. complete absence of reporting; small cost overrun vs. major financial loss; slightly outdated data leading to marginally less relevant recommendations vs. severely inaccurate or biased data leading to systematically wrong decisions) and the level of effort to address such consequences. It thus relates to seriousness and effort. Levels of severity can be low, medium, high or very high.

The combination of these two dimensions provides a risk score for each risk in each question, assessed in accordance with the following matrix:

		Severity			
		Low	Medium	High	Very High
Likelihood	High	Medium	High	Very High	Very High
	Likely	Low	High	Very High	Very High
	Possible	Low	Medium	High	Very High
	Rare	Low	Low	Medium	Very High

Lists of risks are provided, but they are only illustrative. In practice, AI risk repositories have been developed and can be used as additional input for risk assessment. One example is the MIT AI Risk Initiative, a living database of more than 1,700 AI risks, organised using causal and 11 domain taxonomies. Another is the Atlas of AI Risks. Other relevant resources include security-focused catalogues (such as OWASP GenAI, the AI Vulnerability Database and MITRE ATLAS) and incident-based databases (such as the AI Incident Database), which document real-world AI failures and associated harms.

Throughout the Guide, icons will also guide you on the topics at hand:

This icon indicates that the topic is being developed in more detail

This icon indicates that additional information is being provided

Checklist and risk mapping

Goal

Questions

Expertise required

Risk mapping

Rsisk	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

You can read the entire AI Guide to get a complete overview of the main topics impacting AI, or you may choose to only read the chapters, questions or matters that interest you most.

E. Beyond this guide

This AI Guide is complemented by a set of AI products and services. These products and services aim to respond to the full spectrum of stakeholders' needs from a legal and regulatory point of view, ensuring that AI development and use is legally compliant in light of all laws that may apply.

The products and services are framed in a set of building blocks, which provide a blueprint guiding stakeholders throughout their path in the provision or use of AI.

This AI Guide was developed as part of the Center for Responsible AI initiative (CRAI), a flagship private project for AI developed under the Portuguese Recovery and Resilience Plan. Its aim is to create next-generation AI products based on the technologies and principles of Responsible Artificial Intelligence, as well as to undertake research on AI. For more information on CRAI, please visit https://centerforresponsible.ai.

Under the CRAI initiative, Responsible AI Legal, Ethical and Product Principles have been researched, taking into account the seminal EU AI Act, with the goal of:

- Developing a broad knowledge base around the legal, regulatory and ethical challenges and respective solutions to assist organisations in the creation and deployment of Responsible AI products;
- Providing legal and ethical feasibility studies related to the developed solutions and building principles, guidance, criteria and methodologies for Responsible AI, covering also privacy, data protection, cybersecurity, ethics, biases, and sustainability;
- Creating global impact by raising awareness about the importance of Responsible AI and influencing regulators to prioritise its adoption.

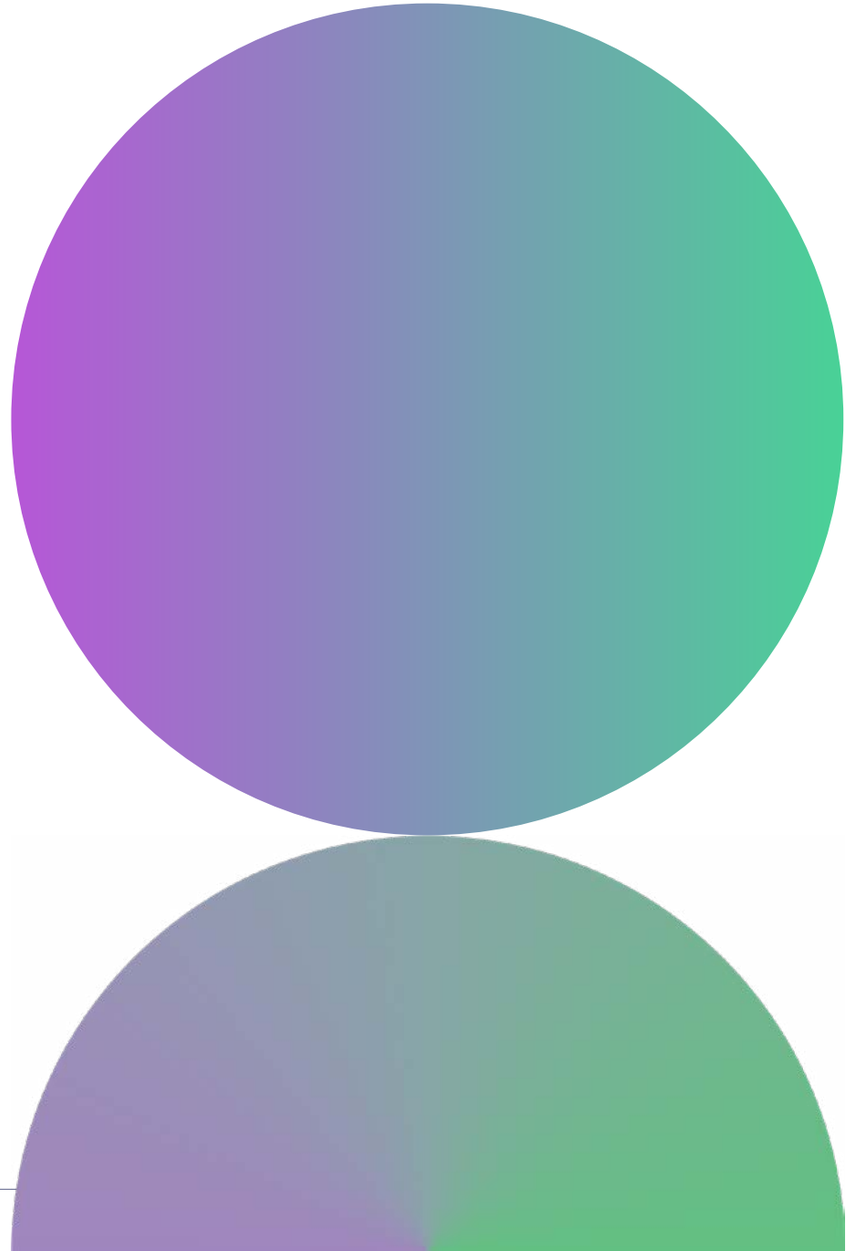

II. Designing and developing AI systems

What it is

Designing and developing AI systems means the set of activities through which an organisation translates an intended use case into an operational AI system, covering problem definition; solution architecture; model selection; data strategy; training, validation and testing. To the extent that AI is a product or is integrated into a product, product manufacturing and integration may also occur at this stage. Though the last step of the AI lifecycle is AI decommissioning, it can also be considered a step in the development stage, when data archiving or deletion is carried out and the decommissioning of physical items of the AI system (waste management processes) is performed.

The main legal touchpoint is the EU AI Act, complemented by product safety, resilience and sustainability laws and data (personal data and non-personal data), intellectual property and competition laws. Other laws, such as those on digital services and accessibility, are also relevant hereto.

In practice, this stage is where risk classification; safety, robustness and resilience; transparency and human oversight design; record-keeping; data governance; quality and risk controls; conformity assessment preparation; and post-market readiness are defined and implemented.

Relevant stakeholders

- AI providers
- Others, such as:
  - AI developers
  - AI manufacturers
  - AI component and service providers, including integrators and model providers
  - Data providers and data analytics providers
  - Infrastructure providers, such as cloud and computing
  - Providers of other inputs, such as hardware, chips and raw materials

Key Topics

The main topics relevant for the AI development stage are:

- Features of the AI system
- Data to train the AI system
- Other resources and partnerships for AI development
- Governance
- Compliance and effects of non-compliance
- IP protection of the AI system
- Particularities for organisations and sectors

Summary of main challenges and how to address them

Challenges	How to address the challenges – summary
Designing and developing AI systems so that they meet legal and ethical expectations Ensuring AI systems – especially high-risk ones – actually meet legal requirements on safety, fundamental rights, resilience, sustainability, privacy, human oversight, auditability, transparency and fairness, and accessibility, rather than treating these as abstract principles	• Assess risk level of the AI system and whether it falls under "high-risk", "prohibited" or other regulated categories • Map all applicable legal frameworks beyond the AI Act (e.g., product safety, resilience and sustainability laws; data protection laws; accessibility laws; competition laws) • Perform structured risk and hazard assessments and define concrete mitigation and elimination measures • Design an integrated compliance plan so that overlapping obligations are met coherently rather than in silos For more details, see FEATURES OF THE AI SYSTEM and PARTICULARITIES FOR ORGANISATIONS AND SECTORS.
Using data that is lawful, high quality and competition compliant Lawfully obtaining and using training, validation and testing data that meets the legally required features (i.e., is relevant, representative, complete, contextualised, accurate and unbiased), while navigating privacy, intellectual property and competition constraints (such as restrictive data licences or data-sharing arrangements by dominant undertakings	• Classify the data you rely on (e.g., personal/non-personal, proprietary/open) and identify lawful bases for use • Assess data sources for lawfulness and quality, and implement processes (e.g., cleaning, labelling, updating, enrichment) to meet required features • Ensure contracts and licences explicitly permit the intended AI uses and reflect any existing mandatory requirements (e.g., on B2B data sharing or on licences for copyrighted works) • Review data-sharing and licensing arrangements (including pricing and exclusivity) for potential competition law issues For more details, see DATA TO TRAIN THE AI SYSTEM.
Structuring sourcing and partnership contracts for AI development Ensuring that contracts for components, services and R&D partnerships allocate rights and obligations clearly (including IP and commercialisation impacts, access to data, SLAs, warranties), while avoiding gaps that block development or clauses that raise competition concerns	• Draft and negotiate contracts to allow for the envisaged AI development • Reflect any mandatory clauses arising from the AI Act and other relevant regimes (e.g., on component supply for high-risk AI systems; on provision of data processing services such as IaaS, PaaS and SaaS) • Capitalise on R&D-friendly provisions (e.g., research exemptions) in collaboration agreements where available • Review the contracts for potential competition law issues (e.g., exclusivities or access restrictions that risk foreclosure of essential inputs such as compute, cloud or foundation models) For more details, see OTHER RESOURCES AND PARTNERSHIPS FOR AI DEVELOPMENT.

Challenges	How to address the challenges – summary
Building effective AI governance without duplication Setting up governance that satisfies AI-specific obligations (such as on quality management, risk management, record keeping, post-market monitoring and data management) and existing frameworks (e.g., Data Protection Impact Assessments (DPIAs), safety and cybersecurity risk assessments), without creating fragmented or duplicated processes	• Design and implement AI governance processes aligned with AI Act requirements for the relevant systems (especially high-risk) • Map all existing governance processes (e.g., AI, safety, resilience, privacy) and ensure coordination among them, also in light of legal requirements on this topic • Define clear accountability structures (roles, committees, escalation paths) and embed a culture of transparency, responsibility and ethical behaviour For more details, see GOVERNANCE.
Demonstrating compliance and dealing with failure Not only being compliant, but being able to demonstrate compliance through conformity assessments, documentation and certifications, and having credible mechanisms to detect, report and remediate incidents or non-compliance	• Determine which conformity assessment procedure(s) must be implemented in light of the applicable laws and assess coordination among them • Identify and leverage relevant standards, certification schemes and guidance to operationalise requirements and streamline evidence gathering • Implement training and literacy initiatives so that teams understand AI risks, obligations and internal procedures • Implement incident detection, reporting, corrective action and liability management processes for malfunctions and breaches For more details, see COMPLIANCE AND EFFECTS OF NON-COMPLIANCE.
Protecting AI-related intellectual assets Securing appropriate IP protection for AI systems, in a context where inventorship, human contribution and protectability criteria can be complex, and where alternatives (trade secrets, technical measures) may be needed	• Design development processes so that protectable subject matter is created (e.g., clear human contribution for copyright and patent purposes) • Develop and implement an effective IP management process to ensure AI system protection is aligned with commercialisation plans • Implement IP management workflows (disclosures, prior checks, filings, registrations, confidentiality controls) from early stages of AI development For more details, see IP PROTECTION OF THE AI SYSTEM.

How to achieve Responsible AI

Responsible AI is achieved when AI systems are developed in ways that are lawful and ethical. To move in that direction, organisations should implement a coherent set of technical, organisational and contractual measures that translate Responsible AI into day-to-day practice, notably by embedding Responsible AI into design and development; ensuring lawful and high-quality data use; building governance and an accountability culture; using contracts, standards, certifications and assessments to operationalise and evidence responsible practices and enable independent scrutiny; monitoring and dealing with real-world impacts, including through reporting and remediation measures; and protecting AI through effective IP management processes that encourage innovation.

Are there any legal requirements regarding the features an AI system should have?

AI systems must meet different legal requirements depending on their functions, integration in (physical) products, processing of personal data, interaction with people, and outputs.

In broad terms, depending on the AI system, it should be designed and developed to be:

- **Safe** – high-risk AI must minimise risks to health, safety and fundamental rights. Systems embedded in certain products (and some standalone safety software) must also meet applicable product safety rules.
- **Resilient** – high-risk AI must achieve appropriate robustness and cybersecurity. Systems must also meet general resilience duties (under horizontal cyber resilience rules).
- **Environmentally sustainable** – where AI is integral to product groups covered by ecodesign rules, it must support compliance with applicable sustainability requirements.
- **Privacy preserving and data accessible** – if processing personal data (see question 11), AI must, where required (see question 5), be designed in line with data protection rules. Connected products and related services (e.g., connected vehicles, medical or fitness devices, smart-home apps) must enable user access to usage data, which integrated AI must support.
- **Overseen by humans** – high-risk AI must allow effective human oversight during use (with similar requirements arising under other regimes).
- **Auditable** – high-risk AI systems should technically allow the automatic recording of events (logs) during their lifetime (with similar requirements arising under other regimes).
- **Transparent and fair** – high-risk AI systems must allow deployers to understand and use outputs appropriately. AI that interacts directly with people must be generally designed and developed so that people are informed that they are dealing with AI, and systems generating synthetic content must ensure that content is marked or detectable as such. Online platform interfaces (which can embed AI or be AI-enabled) must not use dark patterns (see question 8), with forthcoming Digital Fairness Act rules expected to go further with respect to manipulative and addictive designs.
- **Accessible** – certain AI systems (including high-risk AI systems) must comply with EU accessibility requirements.

For more details on each of these requirements, see questions 2 to 10.

The laws provide for certain exclusions. For instance, the AI Act does not apply to AI systems exclusively for military, defence or national security purposes, and those for the sole purpose of scientific research and development. Other laws have their own carve-outs.

To whom the obligations apply

AI Act obligations apply to the **AI provider**; product safety, resilience and environmental rules, to the **product manufacturer**. Generally, they are those that develop (or commission) the system/product and place it on the EU market or put it into service under their own name or trademark. Other actors can become providers or manufacturers by branding or in some cases modifying the system/product.

Data protection obligations apply to **data controllers**. Dark patterns rules apply to **providers of online platforms** and accessibility obligations to **product manufacturers and service providers**. See questions 5, 8 and 10 for an explanation of who these actors are.

AI Checklist

Goal:

Assess why the AI system is being developed and what it is meant to do, and:

- Whether it is **prohibited** or **high-risk**.
- Whether it is a **product** under safety, resilience and sustainability rules and/or a **connected product (or enables a related service)**.
- Whether it is embedded in, or enables, an **online platform**.
- Whether it is a **personal data processing system**.
- Whether it is subject to the **AI Act**, EU **product safety**, **resilience and sustainability laws**, the **Data Act**, the **General Data Protection Regulation** (GDPR), the **Digital Services Act** (DSA) and **accessibility laws**, bearing in mind that some systems (e.g., not made available in the EU or covered by specific exclusions) may fall outside these regimes even if they are high-risk, qualify as products, or meet other criteria.

Questions:

- What are the main objectives of the AI system?
- What are its key features (including autonomy, adaptiveness, predictability)?
- In what context will it be used (simple/predictable or complex/uncertain)?
- Is it embedded in, or does it interact with, a physical product?
- Does it enable the processing of personal data?
- In which countries will it be offered?

Expertise required:

Product & Engineering; Quality; Marketing; Legal.

Prohibited practices

In broad terms, prohibited AI practices include: manipulative, deceptive or vulnerability exploiting systems that distort behaviour and cause significant harm; social scoring and criminal prediction systems based solely on profiling or personal traits; indiscriminate facial image scraping from the internet or CCTV; emotion inferring systems at work or education (save for safety/medical); biometric categorisation systems inferring sensitive traits; most real time biometric ID in publicly accessible spaces by law enforcement; systems that generate or manipulate realistic sexual images of identifiable persons without consent and child pornography material. The Commission's 2025 Guidelines on prohibited AI practices give further explanations and examples.

High-risk systems

High-risk AI systems are those that pose significant risks to health, safety or fundamental rights and are thus subject to stricter requirements. They include, in particular, systems used for biometric identification, critical infrastructure safety, education and vocational training, employment and worker management, access to essential services, law enforcement, migration and border control, and justice and democratic processes (Annex III of AI Act). High-risk AI systems also cover AI that is a safety component of certain regulated products or is itself such a product (for example, machinery), where those products undergo a third-party conformity assessment for health and safety risks (see question 19) (Annex I of AI Act). Under the new updated version of the AI Act, for these latter systems, overlapping AI Act obligations may be limited by the Commission where existing EU product rules already ensure equal or higher protection.

Risk mapping:

Risk of exposure to EU laws	Likelihood	Severity (of non-compliance)	Risk score
List of laws			

See the explanation on the Checklist and Risk Mapping above.

How can AI systems be designed to uphold fundamental rights and maintain safety throughout its lifecycle?

Safety generally refers to products' legal conformity so that they do not cause harm, particularly to people, when used as intended or in a reasonably foreseeable way. AI systems' safety is closely linked to respect for fundamental rights: a "safe" AI system must be designed and developed to prevent or mitigate harm to people's safety, health and rights, in line with the AI Act for high-risk systems and with applicable product safety rules.

In broad terms, depending on the AI system at stake, AI systems should be:

- **Technically safe** – high-risk AI systems must be designed and developed so that, as far as technically feasible, risks to health, safety or fundamental rights are eliminated or reduced. This follows a risk management process (see question 17) and may require updates or redesign. Where risks cannot be eliminated, adequate mitigation and control measures are needed.
- **Tested for safety** – high-risk AI systems must be tested against predefined metrics before market placement, with retesting required only after substantial modifications.
- **Documented and subject to training** – the documentation accompanying high-risk AI systems (see question 20) must cover safety and fundamental rights aspects, including potential discriminatory impacts, and risks to health, safety and fundamental rights. Where appropriate, deployers should receive training.

Regardless of whether the AI Act applies in full, AI systems (high-risk or not) must comply with product safety requirements where applicable. This will typically be the case when the AI is embedded in a physical product, but certain stand alone software (such as safety related AI) can also fall under product safety regimes like the **Machinery Regulation** and the **General Product Safety Regulation** (GPSR) for consumer products.

- The **Machinery Regulation** addresses risks from autonomous and self evolving behaviour.
- The **GPSR** covers safety risks of products (including software) with evolving or predictive functionalities, thereby also capturing low risk AI according to the 2025 Commission guidance.

↪ Under the new updated version of the AI Act, where AI is a safety component of certain regulated products or is itself such a product subject to third-party conformity assessment for health and safety risks, specific AI Act's obligations may be limited by delegated acts of the Commission. This may benefit products such as medical devices, radio equipment or toys. Machinery is excluded from the full high-risk AI regime, with the Machinery Regulation supplemented by delegated acts to incorporate most AI Act health and safety requirements.

In parallel, fundamental rights protections also flow from other EU instruments, such as the General Data Protection Regulation (GDPR), the Digital Services Act (DSA) and non-discrimination laws (on racial equality, employment and occupation, gender equality in access to goods and services, among others). Assessing fundamental rights impacts, even for non-high-risk systems, is therefore a recommended good practice.

To whom the obligations apply

AI Act obligations apply to the **AI provider** for high-risk AI systems, and to the **product manufacturer** under product safety legislation. For an explanation of who these actors are, see question 1.

For other fundamental rights frameworks, GDPR obligations apply to the **data controller** (see question 5 for an explanation of this actor), while DSA obligations apply to **providers of intermediary services**, such as online platforms, offering services in the EU. Non-discrimination obligations typically apply broadly, particularly in contexts as determined by the law at stake, such as in employment, education, access to social security and healthcare, and provision of goods and services.

AI Checklist — Safety features

Goal:

Assess compliance with the safety requirements of the applicable laws, specifically:

- Whether the AI system is designed and developed to reduce or eliminate risks to health, safety or fundamental rights – applicable only to high-risk AI systems under the AI Act.
- Whether the AI system is designed and developed in line with applicable safety requirements – applicable only to AI systems that classify as products under product safety legislation.

Questions:

- Which fundamental rights, and individuals or groups, could be affected by the AI system?
- What safety hazards could the system create (e.g., physical or emotional harm, critical decision errors, disruption, misuse)?
- Which safety features and mitigation and control measures will be in place?
- Which safety tests, metrics and thresholds are planned?

Expertise required:

Product & Engineering; Legal; Compliance; Ethics.

The AI Act as a product safety regulation and its focus on fundamental rights

The AI Act treats many AI systems as products subject to strict safety requirements, aligning them with the EU's broader product safety framework. High-risk AI must therefore meet essential requirements (often via harmonised standards – see question 22), undergo conformity assessment (see question 26), be supported by technical documentation and an EU declaration of conformity, and bear the CE marking (see question 27).

What makes the AI Act distinctive is its explicit focus on fundamental rights, not just safety – risk to fundamental rights is a core organising principle. Actors must thus consider the full spectrum of rights under the EU Charter of Fundamental Rights that can be potentially affected by an AI system, going well beyond the rights typically assessed, such as data protection and non-discrimination, to include, depending on context, rights such as education, work, freedom to conduct a business and freedom of expression.

Risk mapping:

Fundamental rights affected by the AI system	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

Safety risks of the AI system	Description of the impact	Likelihood	Severity	Risk score	Safety features/Mitigation/ elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can the AI system be designed to be resilient throughout its lifecycle?

Resilience refers to the ability of an organisation or system to prevent, protect against, respond to, withstand, mitigate, absorb, accommodate, and recover from an incident. For an AI system to be resilient, it must be designed and developed so that an appropriate level of robustness and cybersecurity are achieved, in line with the AI Act for high-risk systems and the **Cyber Resilience Act** (CRA).

In broad terms, depending on the AI system at stake, AI systems should be:

- **Technically robust and secure** – high-risk AI systems must ensure appropriate robustness and cybersecurity, remaining resilient to errors, faults or inconsistencies, and unauthorised attempts to alter their use, outputs, or performance.
  - ↪ Examples of measures include: backups or fail-safe plans, measures to prevent, detect, respond to, resolve and control for attacks trying to manipulate the training data set or pre-trained components used in training (data and model poisoning), inputs designed to cause the AI model to make a mistake (model evasion), and confidentiality attacks or model flaws.

  AI systems (high-risk or not) must further meet the requirements of the CRA, to ensure that vulnerabilities are minimised and managed across the system's lifecycle, embedding security-by-design and security-by-default principles. These requirements include, among others, ensuring the AI systems have no known exploitable vulnerabilities, have a secure by default configuration and are designed to limit attack surfaces. It also includes vulnerability handling requirements (see question 44). For this purpose, an assessment of the cyber risks must be performed (see question 17).
  - ↪ To the extent an AI system complies with the CRA requirements, it is considered to also meet the AI Act requirements.
- **Tested for security** – under the CRA, effective and regular tests and reviews of the security of the AI system must be performed as part of the vulnerability handling processes that must be implemented (see question 17).
- **Supported by a resilient supply chain** – under the CRA, resilience obligations extend to the entire production and supply ecosystem. Due diligence must be exercised when integrating components sourced from third parties, so that those components do not compromise the cybersecurity of the AI system (see also question 15).
- **Documented** – documentation describing the system's design choices, vulnerability handling processes and cyber risks must be kept (see question 27).
  - ↪ Cyber resilience aspects may also be established in dedicated product safety legislation (see question 2). Under the new updated version of the AI Act, where AI is a safety component of certain regulated products or is itself such a product subject to third-party conformity assessment for health and safety risks, specific AI Act's obligations may be limited by delegated acts of the Commission. This may benefit products such as medical devices, radio equipment or toys, all of which contain cybersecurity requirements. Machinery is excluded from the full high-risk AI regime, with the Machinery Regulation supplemented by delegated acts to incorporate most AI Act health and safety requirements.

To whom the obligations apply

The obligations apply to the **AI provider** for high-risk AI systems under the AI Act, and the **product manufacturer** under the CRA (and dedicated product safety laws). For an explanation of who these actors are, see question 1.

AI Checklist — Resilience features

Goal:

Assess compliance with the resilience and cybersecurity requirements of the applicable laws, specifically:

- Whether the AI system is designed and developed to be robust and secure under the AI Act – applicable only to high-risk AI systems.
- Whether the AI system is designed and developed in line with applicable essential cybersecurity requirements under the CRA and dedicated product safety laws – applicable to all AI systems.

Questions:

- Which cybersecurity vulnerabilities and risks are anticipated (e.g., access control gaps, data poisoning, integrity issues, prompt/command injection, model theft or leakage, denial of service, supply chain risks)?
- Which benchmarks and measurement methodologies are envisaged to assess robustness?
- Which security tests are planned?
- Which resilience and cybersecurity features will the system include?
- What measures are going to be implemented to ensure that third-party sourced components do not compromise the security of the AI system?

Expertise required:

Product & Engineering; Risk & Security; Legal; Compliance.

Risk mapping:

Cyber vulnerabilities and risks of the AI system	Description of the impact	Likelihood	Severity	Risk score	Cyber features | Mitigation/ elimination measures

See the explanation on the Checklist and Risk Mapping above.

The European Cybersecurity Certification Scheme

Cybersecurity certification means formally assessing ICT products, services or processes against predefined criteria and issuing either a certificate or a self-declaration confirming they meet an EU cybersecurity scheme. The EU Cybersecurity Act created this EU-wide framework under which ENISA (the EU Cybersecurity Agency) prepares candidate schemes and, once adopted, manufacturers and providers can seek certification.

Under the AI Act, a high-risk AI system that is certified (or covered by a statement of conformity) under an EU cybersecurity scheme enjoys a presumption of compliance with the AI Act's cybersecurity requirements, to the extent covered. So far, only the EU Common Criteria (EUCC) scheme for ICT products has been adopted, with other schemes under development. In 2026, the Commission proposed a revised "Cybersecurity Act 2" to broaden this framework to include managed security services and an organisation's overall cybersecurity posture.

How can the AI system be designed to be environmentally sustainable throughout its lifecycle?

Environmental sustainability generally refers to requirements established by law to ensure the minimisation of the environmental impacts of products with a view to ensuring a low carbon/environmental footprint, while allowing end-users to choose which products they prefer considering their sustainability. While the AI Act recognises environmental protection as a public interest that must be respected in the development, use and uptake of AI, it does not establish binding environmental requirements.

↳ Despite the above, high-risk AI systems must be adequately designed to mitigate risks to health, safety or fundamental rights (as seen in question 2), and to allow for the logging of events presenting a risk to these rights (see question 7). These rights can be potentially impacted by environmental issues.

↳ The AI Act explicitly encourages codes of conduct covering the assessment and minimisation of the impact of AI systems on environmental sustainability, including as regards energy-efficient programming and techniques for the efficient design, training and use of AI.

Instead, sustainability obligations primarily arise from other EU legislation, most notably the **Ecodesign for Sustainable Products Regulation** (ESPR) and the delegated acts, which are approved for product groups. In broad terms, to the extent the AI system is an integral part of a physical product (e.g., AI-enabled robotics, consumer electronics or connected devices) subject to a delegated act of the ESPR, the AI system, or, in other words, the product including the AI system, should be:

- **Compliant with ecodesign requirements** – products subject to specific ecodesign requirements in delegated acts (e.g., potentially AI-relevant categories such as household dishwashers, fridges, mobile phones, tablets, computers, servers and data storage equipment) must comply with those rules. These cover aspects such as durability and reliability, reusability, upgradability and reparability, maintenance and refurbishment, energy and water use and efficiency, resource efficiency and recycled content, recyclability, environmental impacts (including carbon and environmental footprints) and expected waste generation. Ecodesign requirements may also require that products are not made prematurely obsolete, for instance because software no longer functions after an operating system update or updates are withheld.

  ↳ Software and firmware updates must not degrade product performance beyond legally acceptable margins or reduce functional performance from the user's perspective, and in no case may updates result in the product becoming non-compliant with legal requirements.

- **Tested for sustainability** – tests, measurements and calculations must be performed for the purposes of compliance and verification of compliance with ecodesign requirements.

  ↳ Products must not be placed on the market or put into service if they are designed to alter their behaviour or properties when tested in order to reach a more favourable result.

- **Documented** – products must, in principle, be accompanied by a Digital Product Passport containing the information detailed in the delegated act for each product group (see question 27).

For an analysis of more general environmental concerns that may arise from AI development, see question 32, which may also be relevant hereto.

To whom the obligations apply

Ecodesign obligations apply primarily to manufacturers of AI-enabled physical products. For an explanation of who these actors are, see question 1.

AI Checklist — Sustainability features

Goal:

Assess compliance with the sustainability requirements of the applicable laws, specifically:

- Whether the AI system is designed and developed to reduce or eliminate risks to health, safety or fundamental rights to the extent such risks have environmental impacts – applicable only to high-risk AI systems under the AI Act.
- Whether the AI-enabled product is designed and manufactured in line with applicable ecodesign requirements – applicable to AI-enabled physical products that are subject to EU delegated acts of the ESPR.

Questions:

- Which environmental risks are expected (e.g., energy use and emissions, water-cooling needs, hardware and e-waste)?
- Which ecodesign tests are planned, when will they be run, and what will they assess?
- Which ecodesign features will the AI-enabled product include, and what additional measures will be implemented?

Expertise required:

Product & Engineering; Legal; Compliance.

Risk mapping:

Environmental risks of the AI system	Description of the impact	Likelihood	Severity	Risk score	Ecodesign features | Mitigation/ elimination measures

See the explanation on the Checklist and Risk Mapping above.

Green certification and labels

In the EU, green certifications and labels help promote transparency, sustainability and accountability in the market. The EU Ecolabel is a voluntary label for goods and services that meet high environmental standards across their lifecycle, while the Energy Labelling rules require clear informa3on on the energy efficiency and resource consumption of energy-related (physical) products.

These instruments increasingly intersect with AI. AI supports green labels by enabling data analytics, automated monitoring and lifecycle assessments, and AI components of products are themselves falling within the scope of ecodesign objectives. Recent EU policies increasingly link AI to "green IT" objectives and point towards developing standards, reporting frameworks and possible specific certification or labelling mechanisms, so that digital technologies align more clearly with the EU's environmental goals.

How can the AI system be designed to protect data, notably personal data, throughout its lifecycle?

Data protection generally refers to requirements established by law to ensure that data is protected from unlawful uses and is used only in accordance with the applicable legal requirements. These legal requirements mostly arise from the **General Data Protection Regulation** (GDPR) regarding personal data (see question 11). For an AI system to be privacy preserving, it must, in certain cases, be designed and developed so that the system is aligned with the GDPR from the beginning.

In broad terms, to the extent the AI system processes personal data, it should be:

- **Technically privacy preserving** – in line with the GDPR, AI systems processing personal data (high-risk or not) must embed data protection principles into their design. This means incorporating safeguards and privacy-friendly settings from the outset (e.g., pseudonymisation, encryption, etc.) to ensure protection by design and by default.
  - ↳ Some key data-protection principles to build into AI design include:
    - **Purpose limitation** – data must be collected for specified, clear and lawful purposes, and not re-used in ways that are incompatible with those purposes. Further use is possible in certain cases but under strict safeguards: for example, for scientific research or in AI sandboxes (see question 50) to develop and test AI systems safeguarding a substantial public interest in areas such as public safety, public health or environmental protection.
    - **Data minimisation** – only data that is adequate, relevant and limited to what is strictly necessary for the intended purpose must be collected and processed.
    - **Data accuracy** – only accurate data must be processed, and the deletion or amendment of incorrect data must be ensured.
    - **Integrity and confidentiality** – appropriate technical and organisational measures must be adopted, such as encryption, pseudonymisation or access control (see also question 3).
    - **Storage limitation** – personal data must be kept only for as long as necessary to fulfil the purpose for which it was collected. After this period, the data should be deleted, anonymised, or otherwise securely destroyed, unless continued retention is justified by law.
- **Documented** – high-risk AI systems should be supported by technical documentation that, although seemingly not having to address privacy-preserving features of the system, must describe, among other aspects, its design specifications and key design choices, risk management, as well as foreseeable unintended outcomes and sources of risks to fundamental rights and discrimination in view of the intended purpose (see question 27).

To whom the obligations apply

The obligations apply to the **data controller** — the one determining the purposes and means of processing — regardless of actual data access. Under the GDPR, this covers controllers established in the EU or, if outside, those processing data of individuals in the EU (in connection with offering goods or services), or monitoring them. Under the AI Act, the provider (see question 1 for an explanation of this actor) is a controller only for training, validation, and testing data, while the deployer (see question 35 for an explanation of this actor) is generally the controller for deployment data. Only where the same entity is both the provider and deployer do full privacy-by-design obligations seem to apply with relation to the deployment data.

AI Checklist — Data protection features

Goal:

Assess compliance with the GDPR requirements, specifically:

- Whether the AI system is designed and developed to be privacy-preserving – applicable to all AI systems that enable the processing of personal data when the organisation that is the AI provider is also the deployer.

The point being assessed here is whether the GDPR applies in order to impose privacy features on the AI system itself.

Questions:

- Is the AI system provider also the AI deployer?
- Is the organisation the data controller or processor?
- Which data protection risks are anticipated for the AI system when deployed (e.g., ingestion of excessive data; repurposing data for unrelated purposes; lack of provision of clear information to individuals about AI-driven decisions; fully automated decisions with inadequate human review; inability to identify or remove data from datasets; logs exposing data; outputs leaking personal data)?
- What privacy features will the AI system have to mitigate or eliminate those risks?

Expertise required:

Product & Engineering; Legal; Compliance; Data including Data Protection Officer (see question 18).

Risk mapping:

Data protection risks of the AI system	Description of the impact	Likelihood	Severity	Risk score	Privacy features | Mitigation/ elimination measures

See the explanation on the Checklist and Risk Mapping above.

Privacy certification, seals and marks

Approved certification mechanisms can be used to demonstrate compliance with the requirement of data protection by design and by default, as well as with other GDPR requirements.

Approved certification mechanisms are those endorsed by the competent supervisory authorities (and, where relevant, by the European Data Protection Board – EDPB) and listed in the official register of GDPR certification mechanisms and data protection seals.

Examples of EU-level mechanisms include **Europrivacy**, a GDPR certification scheme approved by the EDPB that certifies the conformity of data processing, including products and services, and **EuroPriSe**, an EU-based privacy seal for processing operations carried out by data processors.

How can the AI system be designed so that it can be overseen by humans during use?

Human oversight refers to the possibility of humans retaining control and intervening in the operation of AI systems. The goal is to prevent or minimise risks to health, safety and fundamental rights that may arise from the operation of an AI system.

For an AI system to be able to be overseen by humans, it must be designed and developed to that the AI system can be effectively overseen by natural persons whilst in use. This requirement stems from the **AI Act** for high-risk systems, though **product safety laws** may also address such requirement.

In broad terms, depending on the AI system at stake, AI systems should be:

- **Technically able to be overseen by humans** – under the AI Act, high-risk AI systems must be designed and developed in such a way as to be effectively overseen by natural persons during their use.
  The human oversight mechanism must allow humans using the AI system to:
  - **Understand and monitor** the operation of the AI system, including in view of detecting and addressing anomalies, dysfunctions, and unexpected performance;
  - **Be aware of automation bias**, i.e., the possible tendency to automatically rely or over-rely on the output produced by a high-risk AI system (automation bias);
  - **Correctly interpret** the high-risk AI system's output, taking into account, for example, the interpretation tools and methods available;
  - **Have authority over decision-making**, including by deciding not to use the high-risk AI system or to otherwise disregard, override, or reverse the output of the high-risk AI system; and
  - **Intervene** in the operation of the high-risk AI system or interrupt the system through a "stop" button or a similar procedure that allows the system to come to a halt in a safe state.
- **Documented** – human oversight can be achieved not only through measures implemented in the AI system, but also measures identified as being appropriate to be implemented by the AI deployer. Documentation should describe human oversight measures (see question 27).

Regardless of whether the AI Act applies in full, AI systems (high-risk or not) may be subject to human oversight requirements under other laws.

- For instance, under the **Machinery Regulation**, machinery with self-evolving behaviour or logic and autonomous mobile machinery must be human supervisable and responsive to people, primarily to protect operators in hazardous environments.

↪ Under the new updated version of the AI Act, where AI is a safety component of certain regulated products or is itself such a product subject to third-party conformity assessment for health and safety risks, specific AI Act's obligations may be limited by delegated acts of the Commission. Machinery is excluded from the full high-risk AI regime, with the Machinery Regulation supplemented by delegated acts to incorporate most AI Act health and safety requirements.

To whom the obligations apply

The obligations apply to the **AI provider** for high-risk AI systems under the AI Act, and the **product manufacturer** under product safety legislation. For an explanation of who these actors are, see question 1.

AI Checklist — Human oversight features

Goal:

- Whether the AI system is designed and developed to be able to be effectively overseen by natural persons during its use – applicable only to high-risk AI systems under the AI Act.
- Whether the AI system is designed and developed in line with applicable safety requirements relating to human oversight – applicable only to AI systems that classify as products under product safety legislation establishing such requirement.

Questions:

- Which human oversight risks are anticipated for the AI system in light of its features (e.g., inadequate ability to monitor; inadequate ability to interpret outputs; inadequate ability to intervene, override or stop the system; over-reliance risks such as automation bias, reduced vigilance, deskilling, loss of accountability)?
- Which decisions or types of decisions will be based on the AI system?
- What types of tools and mechanisms will be made available for human oversight (e.g., "stop" functions, "override" functions, "disregard" functions, escalation paths, alerting mechanisms so that human overseers can detect anomalies or unexpected performance in time to act)?
- Which features will those tools and mechanisms have when it comes to the individuals that can exercise oversight (category, authority, function), the moment/points where oversight can be exercised (e.g., frequency, triggers, thresholders) and the information/decision support tools to assist oversight (e.g., explanations, confidence scores, uncertainty indicators, so that humans can understand when to trust or override the system)? Are they configurable?
- Are there mechanisms to ensure that deployers do not configure or use the AI system in a way that removes or merely formalises human oversight (e.g., fully automated decision chains with no meaningful human involvement)?
- Does the AI system allow the recording of human intervention, including AI decisions disregarded or overridden by humans, and respective reasons?

Expertise required:

Product & Engineering; Legal; Compliance; Ethics.

Risk mapping:

Oversight risks of the AI system	Description of the impact	Likelihood	Severity	Risk score	Oversight features | Mitigation/ elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can the AI system be designed so that it is able to record events/logs?

The auditability of AI systems relates to their ability to record events (logs). The goal is to mitigate opacity and the "black-box" problem, by making the system's decision-making trail traceable and understandable. For an AI system to be auditable, it must be designed and developed to allow for the recording of events. This requirement stems from the **AI Act** for high-risk systems, though other laws may establish similar or specific requirements, as is the case of **safety, resilience and sustainability laws**, or the **General Data Protection Regulation** (GDPR).

In broad terms, depending on the AI system at stake, AI systems should be:

- **Technically auditable** – high-risk AI systems should be designed and developed in a way that permits them to automatically record events over their lifetime in a reliable, tamper evident and time synchronised manner. The scope of logging should support risk identification, post market monitoring and operational monitoring against the instructions for use.
  - ↳ For certain use cases, such as remote biometric identification systems, the AI Act prescribes the logging of additional items.
- **Allow for the retention of logs** – logs automatically generated by high-risk AI systems and under the control of the provider should be kept for a specific period (at least 6 months). Product safety rules may require different or longer retention for specific safety-critical data. Additionally, logs should be protected with robust technical and organisational measures.
- **Documented** – high-risk AI systems should be supported by technical documentation that must describe, among other aspects, design specifications and key design choices (see question 27).

Other laws applying to AI systems may establish equivalent requirements, notably for the purposes of safety, resilience and environmental sustainability. For instance, the **Machinery Regulation**, the **Cyber Resilience Act** (CRA) and the **Ecodesign for Sustainable Products Regulation** (ESPR) all require that products be designed and developed with logging capabilities.

- The **Machinery Regulation** requires, for control systems, the recording of data relating to interventions and uploaded software (which must be kept for five years) and, for autonomous self-evolving machinery, the logging of safety-decision making data (which must be kept for one year).
- The **CRA** requires products to provide security related information by recording and monitoring relevant internal activity, including access to or modification of data, services or functions, with an opt-out mechanism for the user.
- The **ESPR** establishes that products with ecodesign requirements relating to energy-efficiency/energy consumption measurement must, where appropriate, record the in-use data and make such data visible to the end user.
  - ↳ Under the new updated version of the AI Act, where AI is a safety component of certain regulated products or is itself such a product subject to third-party conformity assessment for health and safety risks, specific AI Act's obligations may be limited by delegated acts of the Commission. Machinery is excluded from the full high-risk AI regime, with the Machinery Regulation supplemented by delegated acts to incorporate most AI Act health and safety requirements

When the AI system processes personal data, a record of the processing activities should also be kept (see question 5).

To whom the obligations apply

The obligations apply to the **AI provider** for high-risk AI systems under the AI Act, the **product manufacturer** under the Machinery Regulation, CRA and ESPR, and the **data controller** under the GDPR. For an explanation of who these actors are, see questions 1 and 5.

AI Checklist — Logging Features

Goal:

Assess compliance with the applicable requirements, specifically:

- Whether the AI system is designed and developed to allow the recording of events – applicable only to high-risk AI systems under the AI Act.
- Whether the AI system is designed and developed in line with applicable safety requirements relating to logging – applicable only to AI systems that classify as products under product safety, resilience and sustainability legislation establishing such requirement.
- Whether the AI system is designed and developed to allow logging of the processing of personal data – applicable to all AI systems that enable the processing of personal when the organisation that is the AI provider is also the deployer (see question 5).

Questions:

- Which logging risks are anticipated for the AI system in light of its features (e.g., insufficient logging for traceability; excessive logging creating other risks, such as GDPR risks; logging integrity and availability failures; logging interpretation, review and escalation difficulties)?
- Which events does the AI system allow the logging of?
- Which information is provided in each log? Can it contain personal data? What is its structure and frequency/triggers?
- How are logs made available by the AI system (e.g., APIs, dashboards, export features) and to whom (category, authority, function)?
- Which security features will logs have?
- How long will the AI system allow log retention?
- Are there mechanisms to ensure that deployers do not disable or under-configure critical logging in ways that undermine compliance or traceability?

Expertise required:

Product & Engineering; Legal; Compliance.

Risk mapping:

Logging risks of the AI system	Description of the impact	Likelihood	Severity	Risk score	Logging features | Mitigation/ elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can the AI system be designed when it is intended to interact with individuals?

AI systems intended to interact with natural persons are subject to specific legal requirements, particularly when such interaction can affect individuals' rights, informed decision-making, or safety. These requirements stem not only from the **AI Act**, but also from other EU instruments, in particular **digital services laws**.

In broad terms, AI systems that interact directly with natural persons should be:

- **Transparently designed** – these AI systems (e.g., in digital interfaces, recommender systems, chatbots) must be designed and developed so that individuals are clearly informed that they are interacting with AI, unless this is obvious from the context.
  - ↪ This obligation does not apply if the AI system is authorised by law to detect, prevent, investigate or prosecute criminal offences, subject to appropriate safeguards for the rights and freedoms of third parties, unless those systems are available for the public to report a criminal offence.
- **Able to ensure fairness in interaction with users** – online platform interfaces (which may embed or be enabled by AI systems) must not be designed or organised in a way that deceives or manipulates users, or that otherwise materially distorts or impairs their ability to make free and informed decisions (dark patterns) – for example, visually favouring one option, using misleading nudges, or relying on hard-to-change default settings. This is directly relevant for AI, since AI systems may enable or be integrated into those platforms and may themselves be offered as intermediary services (such as online platforms) within the scope of the Digital Services Act (DSA).
  - ↪ Certain AI systems that may interact with humans are prohibited. This is the case of AI systems deploying subliminal, purposefully manipulative or deceptive techniques to distort behaviour and impair informed decision-making, leading to significant harm, and AI systems exploiting individuals' vulnerabilities due to age, disability, or socio-economic situation to distort behaviour and cause significant harm (see 🔍).
- **Documented** – under the AI Act, high-risk AI systems should be supported by technical documentation that, although not necessarily having to contain separate user-facing information on interactions with individuals, must contain information on its intended purpose, user interface, and design choices, including with regard to persons or groups of persons in respect of who the system is intended to be used. What is more, instructions for use must contain information on characteristics, capabilities and limitations of performance, including information to enable deployers to interpret outputs and use the system appropriately (see question 27).

To whom the obligations apply

The primary obligations under the AI Act lie with the AI provider – for an explanation of this actor, see question 1. In turn, the obligation relating to the prohibition of dark patterns under the DSA applies to providers of online platforms (that design or have the platform designed) offering services in the EU.

AI Checklist — Transparency And Fairness Features | Interaction With Individuals

Goal:

Assess compliance with the applicable requirements, specifically:

- Whether the AI system is designed and developed so that individuals are clearly informed that they are interacting with an AI system– applicable to AI systems intended to interact directly with natural persons, under the AI Act.
- Whether the AI system is designed and developed in a way that the online platform's interface does not deceive or manipulate users– indirectly applicable to AI systems that are integrated in or enable an online platform, under the DSA.

Questions:

- Who are the individuals the AI system is intended to interact with?
- In which contexts (e.g., education, employment, health) and concrete situations will people interact with the AI system, and do these contexts create power imbalances?
- Which transparency and fairness risks are anticipated for the AI system (e.g., no or inadequate AI disclosure, fragmented or inconsistent information, misleading or confusing presentation such as a highly anthropomorphic avatar, override by deployers)?
- Which features and safeguards will the AI system include to address these risks?
- Are there mechanisms to ensure that deployers do not disable or under-configure the AI system in a manner that undermines transparency and fairness?

Expertise required:

Product & Engineering; Legal; Compliance; Data including Data Protection Officer (see question 18).

AI companions

AI companions are digital agents that simulate human-like conversations and relationships, often with persistent personalities that adapt over time to user interactions. They typically combine dialogue models with personalisation and memory (storing preferences, history and patterns), and may include avatars, voices or virtual environments to make interactions feel more immersive and personal.

They intensify existing legal and ethical AI issues around transparency, fairness, privacy (see question 1) and consumer protection (see question 48), by enabling "data extraction through intimacy" and subtle behavioural influence. They also raise wider societal concerns, such as social exclusion, emotional deskilling and "emotional echo chambers", where users primarily engage with responsive digital entities instead of reciprocal human relationships. As a result, AI companions may qualify as a prohibited practice (see 🔍 Prohibited Practices) where they deploy manipulative techniques and exploit vulnerabilities.

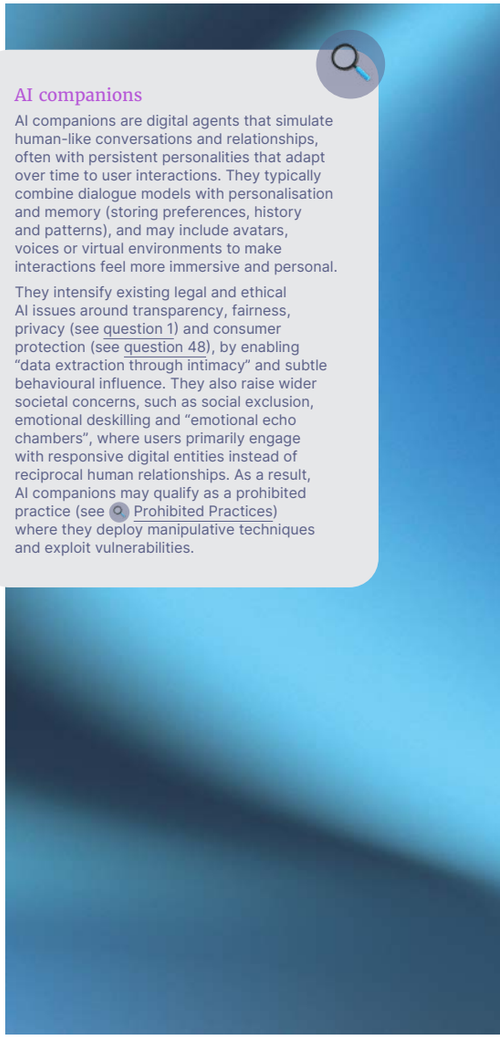

Risk mapping:

Transparency and fairness risks of the AI system	Description of the impact	Likelihood	Severity	Risk score	Features | Mitigation/ elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can the AI system be designed when it is intended to create AI-Generated content?

AI systems that generate content, such as text, audio, images or video, are subject to specific legal requirements to ensure transparency, accountability and the protection of fundamental rights, given the increasing difficulty in distinguishing AI-generated content from human-authored content. These obligations derive primarily from the **AI Act**.

In broad terms, AI systems generating synthetic content should be:

- **Transparently designed** – under the AI Act, technical solutions that are effective, interoperable, robust and reliable, as far as technically feasible, should be implemented to ensure that the outputs of such AI systems are marked in a machine-readable format and detectable as artificially generated or manipulated. The specificities and limitations of various types of content, the costs of implementation and the generally acknowledged state of the art (as may be reflected in relevant technical standards) should be considered when determining the technical solutions. This transparency obligation applies to: (i) autonomously generated content (e.g., AI-written articles, AI art, code generation); and (ii) AI-altered content (e.g., face swapping, voice cloning, content transformation, deep fakes – see Deep Fakes).
  - ↳ The obligation does not apply if the AI system performs an assistive function for standard editing or does not substantially alter the input data provided by the deployer or the semantics thereof; or if the AI system is authorised by law to detect, prevent, investigate or prosecute criminal offences.
- **Able to ensure fairness in relation to users** – certain AI systems that generate content should not be designed and developed in such a way that their AI-generated content can be used to mislead or manipulate users – in these cases, they are prohibited.
  - ↳ This is the case of AI systems deploying subliminal, purposefully manipulative or deceptive techniques to distort behaviour and impair informed decision-making, leading to significant harm, and AI systems exploiting individuals' vulnerabilities due to age, disability, or socio-economic situation to distort behaviour and cause significant harm (see Prohibited Practices).
- **Documented** – under the AI Act, high-risk AI systems should be supported by technical documentation that, although not necessarily having to contain specific information on the marking of AI-generated content as such, must contain information on its intended purpose. What is more, instructions for use must contain information to enable deployers to interpret outputs (see question 27).

To whom the obligations apply

The obligations relating to AI systems generating content apply to **AI providers**. For an explanation of who these actors are, see question 1.

AI Checklist — Transparency And Fairness Features | AI-Generated Content

Goal:

Assess compliance with the applicable requirements, specifically whether the AI system is designed and developed in such a way that the outputs are marked in a machine-readable format and detectable as artificially generated or manipulated – applicable to AI systems generating content, under the AI Act.

Questions:

- What content does the AI system generate and what is the level of AI intervention in the content?
- Could the AI-generated content be mistaken for human-created content?
- What contexts (e.g., education, employment, health) and specific situations will be able to resort to the AI-generated content and do such contexts or situations give rise to power imbalances or undue influence?
- Which transparency and fairness risks are anticipated for the AI system in light of its features (e.g., unlabelled content; weak or removable labelling; ambiguous or misleading labels; misconfiguration or misuse by deployers; inconsistent application across features, channels and content)?
- Which features and safeguards will the AI system include to address these risks, in particular what transparency will be provided (form, timing, content and channel), and can these measures be adapted to different types of content, contexts, situations or users where needed?
- Are there mechanisms to ensure that deployers do not disable or under-configure the AI system in a manner that undermines transparency and fairness?

Expertise required:

Product & Engineering; Legal; Compliance; Ethics; Data including Data Protection Officer (see question 18).

Risk mapping:

Transparency and fairness risks of the AI system	Description of the impact	Likelihood	Severity	Risk score	Features | Mitigation/ elimination measures

See the explanation on the Checklist and Risk Mapping above.

The Code of Practice on Transparency Of AI-Generated Content

The Code of Practice on Transparency of AI Generated Content, issued June 2026, sets out voluntary commitments to help identify AI generated and manipulated content. It rejects single technique solutions and promotes a two layer marking approach to be applied across the value chain. Providers of generative AI are expected to use secured, tamper evident metadata and imperceptible watermarking as primary machine readable markers, with options such as fingerprinting and logging as optional complementary tools.

The Code also addresses measures to reduce the risk of markings removal, improve provenance transparency and allow placement of perceptible markings. It further provides for mechanisms for others to verify whether content is AI generated. It stresses that marking and detection should be effective, reliable, robust and interoperable, and acknowledges proportionate compliance paths that take account of organisations' size and resources, in particular in case of SMEs and small mid-cap enterprises (SMCs) (see question 51).

The European Commission also issued draft guidance on this topic in May 2026.

How can the AI system be designed so that it is accessible?

AI systems are subject to specific legal requirements on accessibility.

Accessibility obligations arise primarily from the **European Accessibility Act** (EAA) and the Web Accessibility Directive (WAD), and thus apply to AI systems that are, or are integrated, in products or services covered by such laws. These laws are complemented by the **AI Act** for high-risk systems.

In broad terms, and depending on the AI system at stake, AI systems should be:

- **Accessible to persons with disabilities** – under the AI Act, high-risk AI systems must be designed and developed to comply with accessibility requirements (of the EAA and WAD) so that their use does not exclude or discriminate against persons with disabilities. The EAA establishes accessibility obligations for a range of products and services (e.g., computers, smartphones, banking applications, etc.), requiring that they be designed by default for accessibility (design for all). If these incorporate AI functionalities, then such AI systems (high-risk or not) must be designed to ensure that persons with disabilities can perceive, understand, navigate and interact with them effectively. Likewise, the WAD focuses on ensuring the accessibility of public sector websites and mobile apps across the EU, requiring public sector bodies to make them accessible.
- **Documented** – for systems covered by the EAA or the WAD, an explanation on how the systems meet accessibility requirements, or an accessibility statement, respectively, must be provided (see also question 27).

↳ The EAA's accessibility requirements apply only insofar as they do not require changes that would fundamentally alter the basic nature of a product or service or impose a disproportionate burden on the relevant economic operators. Similarly, under the WAD, public sector websites and apps may be exempt where meeting accessibility rules would entail a disproportionate burden.

↳ The WAD's accessibility requirements do not apply to certain websites and mobile apps, such as those of public service broadcasters and of NGOs that do not provide services essential to the public or services for persons with disabilities.

To whom the obligations apply

The obligations apply to the **AI provider** for high-risk AI systems under the AI Act (see question 1 for an explanation of this actor) and to the **product manufacturer** (same explanation as indicated in question 1) or **service provider** (providing services on the EU market) under accessibility legislation.

AI Checklist — Accessibility

Goal:

Assess compliance with the applicable requirements, specifically whether the AI system is designed and developed in such a way that the outputs are marked in a machine-readable format and detectable as artificially generated or manipulated – applicable to AI systems generating content, under the AI Act.

Questions:

- Which user interfaces and interaction channels does the AI system have (e.g., web, mobile app, kiosk, API-based UI embedded by clients) and in which of these are persons with disabilities intended to interact with the AI system?
- What accessibility criteria are relevant for the AI system (e.g., perceivable, operable, understandable)?
- Which accessibility risks are anticipated for the AI system in light of its features (e.g., non-accessible user interfaces, non-accessible AI outputs, dependency on third-party components)?
- Which features and mitigation measures will be implemented in the AI system to address the risks, and how will they be tested and when?
- Are there mechanisms to ensure that deployers do not disable or under-configure the AI system in a manner that undermines accessibility?

Expertise required:

Product & Engineering; Legal; Compliance.

Risk mapping:

Accessibility risks of the AI system	Description of the impact	Likelihood	Severity	Risk score	Features | Mitigation/ elimination measures

See the explanation on the Checklist and Risk Mapping above.

What types of data can be used to train, validate and test an AI system?

In broad terms, data may be:

- **Personal vs. non-personal data** – personal data are any information relating to an identified or identifiable natural person (the "data subject"). In turn, non-personal data are any information that are not related to an identified or identifiable natural person.
- **Open vs. protected data** – open data is freely accessible, modifiable, and distributable, while protected data is data safeguarded due to legal, contractual or business requirements. Protected data may include, for instance, personal data, data protected by intellectual property rights, trade secrets and confidential data.
- **Factual vs. non-factual data** – factual data represents acts or facts, having a purely functional character. Non-factual data, by contrast, encompasses more than simple and technical representations of reality – thus, non-factual data may amount to data created, processed or analysed.
- **Raw vs. processed/analysed data** – raw data is automatically generated information that provides a purely technical reflection of reality. Processed or analysed data, in turn, is data that has been handled, treated and interpreted, respectively, to ensure its usability and adaptation to specific use cases.
- **Synthetic vs. real-world data** – synthetic data is data generated by algorithms, whilst real-world data is collected from actual events, interactions and observations of the real world.

Data may be protected by different routes:

  - **Data protection** – applies to personal data, which benefits from protection under the General Data Protection Regulation (GDPR).
  - **Copyright** – protects works or, in other words, human creations, i.e., the exteriorisation of the creative thought of its intellectual creator. Copyrighted data benefits from the protection of copyright laws.
  - **Database *sui generis* right** – protects the qualitative and/or quantitative substantial investment in either the obtaining, verification or presentation of the contents of a database. Data protected by the *sui generis* right benefits from the protection of the Database Directive.
  - **Neighbouring or related rights** – protect "subject matters", i.e., a performance (e.g., of a music), phonogram, videogram, broadcast or press publication.
  - **Trade secret** – protects information that is a secret because it is not generally known or accessible, has commercial value due to its secrecy, and is subject to reasonable steps to keep it confidential. Trade secrets are protected under the Trade Secrets Directive.
  - **Confidentiality** – protects information that is not publicly available, typically through contractual and technical safeguards. Unauthorised access to, or interference with, such information or the systems that store it can amount to criminal conduct under the Cybercrime Directive. Sector-specific secrecy regimes (for example, banking or professional secrecy) may also apply.
  - **Unfair competition** – protects information from being used with an unfair result or purpose as prohibited by law, notably national law.

To whom this is relevant

The AI system is trained, validated and tested by the **AI provider**. Hence, the types of data that can be used for this purpose (and how they are protected, which determines the extent to which they can be lawfully used for training, validation and testing) are especially relevant for the AI provider. See, however, question 1 for an explanation of who the AI provider is.

AI Checklist — Types of Data

Goal:

Assess:

- What type of data is being used – relevant for all AI systems.
- How is the data protected – relevant for all AI systems.

Questions:

- What is the typology of the data being used?
- Is the data protected? If yes, how?
- Is there a clear inventory and data map linking each set of data to routes of protection?
- Which data risks are anticipated (e.g., misclassification of data, poor data inventory)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Legal; Data including Data Scientists, Analytics and Data Protection Officer (see question 18).

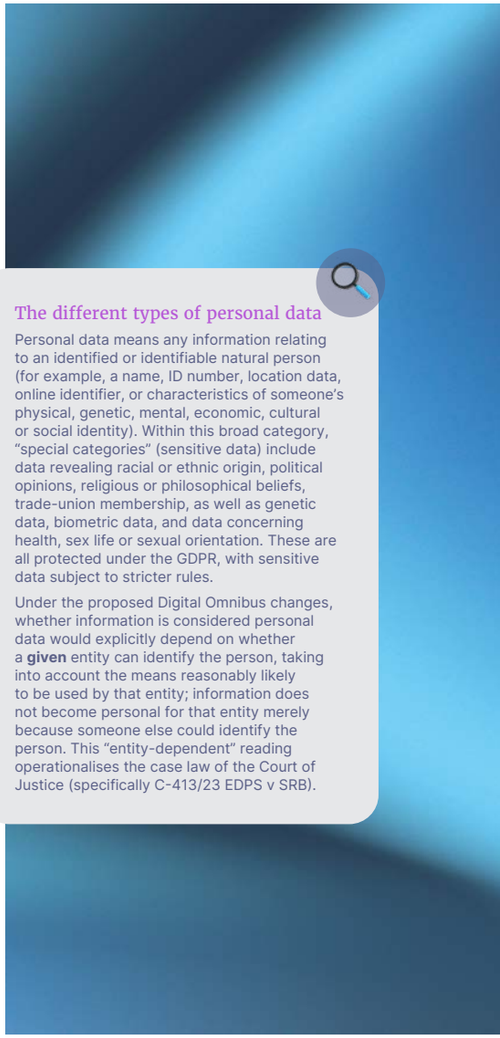

The different types of personal data

Personal data means any information relating to an identified or identifiable natural person (for example, a name, ID number, location data, online identifier, or characteristics of someone's physical, genetic, mental, economic, cultural or social identity). Within this broad category, "special categories" (sensitive data) include data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, trade-union membership, as well as genetic data, biometric data, and data concerning health, sex life or sexual orientation. These are all protected under the GDPR, with sensitive data subject to stricter rules.

Under the proposed Digital Omnibus changes, whether information is considered personal data would explicitly depend on whether a **given** entity can identify the person, taking into account the means reasonably likely to be used by that entity; information does not become personal for that entity merely because someone else could identify the person. This "entity-dependent" reading operationalises the case law of the Court of Justice (specifically C-413/23 EDPS v SRB).

Risk mapping:

Types of data risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can data be obtained to train, validate and test an AI system?

Where data for training, validation and testing of an AI system is not already in the organisation's possession, it must be obtained from third parties, either directly from data holders or data subjects (for personal data), or via intermediaries such as data marketplaces or brokers.

Authorisation

Authorisation is in principle needed to use both personal and non-personal data, though specificities may apply.

For **personal data**, because consent can be difficult to obtain in AI due to scale and reuse, organisations typically rely on other lawful grounds under the **General Data Protection Regulation** (GDPR).

- **Legitimate interests** are often resorted to in the AI lifecycle, provided that a Legitimate Interests Assessment (LIA) confirms the interest is lawful, precise, real and present; the processing is necessary; and the interest is not overridden by data subject rights. This ground does not apply to public authorities in the performance of their tasks.
- Another ground for processing is the processing of special categories of personal data (see 🔍 The Different Types of Personal Data) for **bias detection and correction**. This is allowed under the AI Act for the development of high-risk AI systems, as well as, under the new updated version of the AI Act, for the development and deployment of AI more broadly (high-risk or not), both under tight necessity and protection requirements.

For **non-personal data**, use typically requires a contract or licence where the data is protected (see questions 11 and 14).

↪ A contract is also needed in other situations, such as, in most cases, for using readily available non-personal data generated by users in the Union through connected products and related services (e.g., connected vehicles, medical or fitness devices, smart-home apps) placed on the EU market (Data Act).

A case-by-case assessment is essential to determine whether a contract or licence is required. For instance, whether a licence for copyrighted content or databases is needed may depend on whether acts such as reproduction, transformation, making available, extraction or reutilisation occur in light of the relevant IP regime.

↪ Recent EU and national case law shows divergent views on whether AI "contains" protected works, and key questions on reproduction and communication to the public in AI training and outputs are now before the CJEU in Case C250/25 (Like Company v. Google).

Legal exceptions and limitations

Legal **exceptions and limitations** allow for the use of data even when a licence is lacking. The most relevant of these include data extraction and text and data mining (TDM).

- **Data extraction** – under the Database Directive, non-substantial, non-repeated and non-systematic extractions of data, as well as the extraction of data from connected-product or public sector datasets made available for reuse, are permitted.
- **Text and data mining** (TDM) – under the Digital Single Market (DSM) Directive, TDM of copyrighted content is permitted without authorisation for scientific research uses by research organisations and cultural heritage institutions with lawful access, and more broadly by any user (including commercial) with lawful access, unless rightsholders have clearly opted out by reserving their TDM rights via machine-readable means or terms.

To whom this is relevant

The AI system is trained, validated and tested by the **AI provider**. Hence, it is the AI provider that must obtain such data through any of the means indicated above. See, however, question 1 for an explanation of who the Ai provider is.

AI Checklist — Acquisition of training, validation and testing

Goal:

Assess:

- What are the sources of the data – relevant for all AI systems.
- Whether the data is lawfully obtained for the intended purposes – applicable to all AI systems.

Questions:

- For each dataset, who is the owner/holder?
- How is the data obtained and from whom?
- Is there a clear inventory and data map linking each dataset to sources, means of data acquisition and concrete uses (training, validation, testing, evaluation, fine-tuning, derivatives, etc.)?
- Can the AI system disclose the data used to train, validate and test it?
- Which data risks are anticipated (e.g., unclear ownership of data, unlawful acquisition and/or use of data, over-reliance on legal exceptions, mismatch between acquisition and purpose, regurgitation or memorisation of protected content)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Procurement; Legal; Compliance; Data including Data Scientists, Analytics and Data Protection Officer (see question 18).

Risk mapping:

Data acquisition risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

The EU approach to facilitate data sharing

The EU has created a legal framework to encourage data sharing, which AI actors can use to obtain data for AI.

Public sector data sharing is mainly governed by the Open Data Directive and the Data Governance Act (DGA), which require many public datasets to be made available as open data, with some (including "high value datasets" such as geospatial, earth observation, meteorological, mobility, statistics and company data) to be made available under specific conditions. Protected data can be re-used subject to safeguards such as confidentiality and no re identification. G2G data sharing for public task purposes is largely governed by national law and supported at EU level by the Interoperable Europe Act, which facilitates structured, cross border public sector data exchange.

The Data Act harmonises B2B sharing, protecting businesses (including AI actors) from unfair contractual terms and imposing obligations linked to mandatory data access. It also creates an "exceptional need" regime for B2G sharing, allowing certain public bodies to request data from private entities in public emergencies or certain public interest tasks. Under the Digital Omnibus proposal, the Open Data Directive and DGA would be folded into the Data Act and the general "exceptional need" regime for B2G data sharing would be significantly narrowed.

Should the data used to train, validate and test the AI system have any specific features?

The data used in training, validation and testing must meet the requirements arising from the **AI Act** for high-risk systems.

In particular, datasets used for the development of high-risk AI systems should be:

- **Relevant** – adequate for the result the data is intended to support, assessed in light of the AI system's purpose.
- **Representative** – able to sufficiently reflect, in an accurate manner, the relevant situation, setting or population in view of how the AI system will be used.
- **Complete** – containing all necessary and appropriate information for the AI system's intended use.
- **Contextualised** – taking into consideration the characteristics or elements that are particular to the specific contextual setting within which the high-risk AI system is intended to be used.
- **Accurate** – to the best extent possible, free of errors, so that the data can be trusted from a correctness standpoint.
- **Unbiased** – designed to minimise bias and discrimination through appropriate detection and mitigation measures. Other requirements impact this feature: notably, that the datasets have appropriate statistical properties and take into consideration the system's relevant geographical, contextual, behavioural and functional setting.

Other important features, not expressly detailed in the AI Act but stemming from it and from other legal frameworks, include:

- **Timeliness** – data should reflect the relevant reality at the required point in time and be subject to governance and management practices such as updating (see question 17).
- **Contextuality/metadata** – the extent to which data is enriched with background information (e.g. metadata, annotations). While not explicitly required, the AI Act mandates data governance practices for high risk systems, including annotation and labelling (see question 17).
- **Interoperability** – the ability of data to be reused across systems or entities. While not mandated by the AI Act, it is recognised as important, and EU data law promotes it for public sector, connected product data, and data space contexts.

More broadly, data should also meet general legal requirements, relevant for both high-risk and non-high-risk AI:

- **Privacy** – compliance with the General Data Protection Regulation (GDPR) principles: lawfulness, fairness and transparency; purpose limitation; data minimisation; data accuracy; storage limitation; data integrity and confidentiality (see also questions 5 and 12).
- **Security and integrity** – protection against compromise during data processing. While the AI Act focuses on safeguarding AI systems (e.g. against data poisoning – see question 3), the GDPR imposes security obligations for personal data (see question 5), and the Cybercrime Directive criminalises unlawful data interference.
- **IP compliance** – datasets must be sourced under appropriate licences or exceptions (see question 12).

To whom the obligations apply

The obligation to ensure the datasets meet the above requirements applies to the **AI provider** for high-risk systems under the AI Act. For an explanation of who this actor is, see question 1.

The obligations arising from the GDPR apply to the **data controller** (see question 5 for an explanation of this actor). Under the AI Act, the organisation qualifying as the AI provider is the data controller for the data used to train, validate and test the AI system.

The obligations arising from intellectual property provisions apply, in this context, to the **user of the protected data**.

AI Checklist — Features of training, validation and testing data

Goal:

Assess compliance with the applicable requirements, specifically:

- Whether the data meets the features required by the AI Act – applicable to high-risk AI systems under the AI Act.
- Whether the data meets other features required by other legal acts, such as when it comes to personal data and IP – applicable to all AI systems.

Questions:

- What features does each dataset have?
- Who ensures that the data have the legally required features?
- Which data risks are anticipated (e.g., inadequate features of data, misalignment of data used with stated purpose and risk profile of the AI system)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Quality; Legal; Compliance; Data including Data Scientists, Analytics and Data Protection Officer (see question 18).

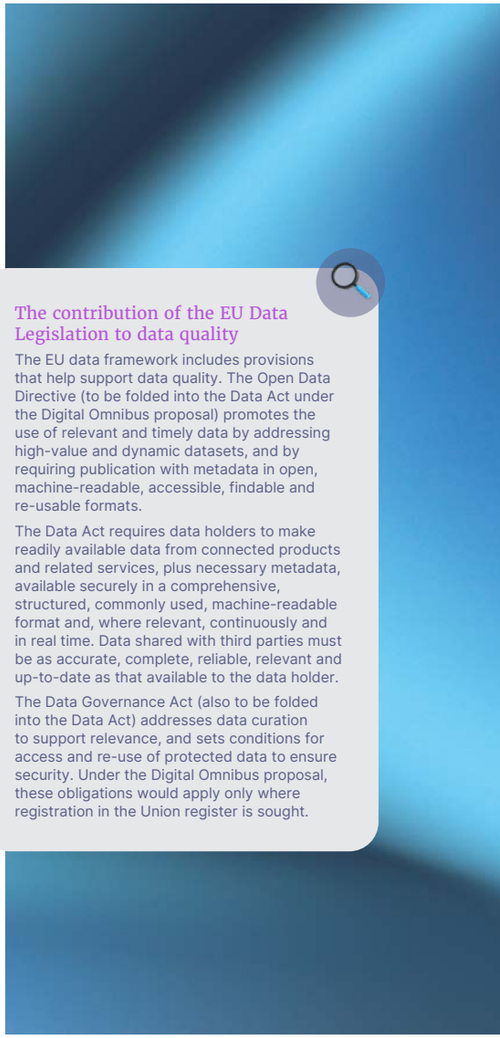

The contribution of the EU Data Legislation to data quality

The EU data framework includes provisions that help support data quality. The Open Data Directive (to be folded into the Data Act under the Digital Omnibus proposal) promotes the use of relevant and timely data by addressing high-value and dynamic datasets, and by requiring publication with metadata in open, machine-readable, accessible, findable and re-usable formats.

The Data Act requires data holders to make readily available data from connected products and related services, plus necessary metadata, available securely in a comprehensive, structured, commonly used, machine-readable format and, where relevant, continuously and in real time. Data shared with third parties must be as accurate, complete, reliable, relevant and up-to-date as that available to the data holder.

The Data Governance Act (also to be folded into the Data Act) addresses data curation to support relevance, and sets conditions for access and re-use of protected data to ensure security. Under the Digital Omnibus proposal, these obligations would apply only where registration in the Union register is sought.

Risk mapping:

Data features risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What main aspects should be included in a contract for data to train, validate and test the AI system?

A contract for data to train, validate and test an AI system is a binding agreement on how data is provided, accessed and used.

In broad terms, it is recommended that such a contract covers the following:

- **The data provided and its features** – precise description of the data, delivery dates and, where relevant, acceptance procedures; indication of quality requirements, especially for high-risk AI, and of the technical means of provision. Regular provision of data may be subject to SLAs.
- **The conditions for use** – clear rules on permitted purposes and fields of use, types of processing (including for data preparation operations required under the AI Act for high-risk systems), storage, internal and external uses, commercial exploitation, access controls, territory, duration and post-termination consequences (e.g., return or deletion).
  - ↳ To the extent the data is to be stored in the AI system and used not only for training, validation and testing, but also as an input for AI generation of results, this should be addressed in the contract.
  - ↳ To the extent the data is provided only for a fixed term, ending access may require retraining, "machine unlearning", or may affect deployed systems and the content they generate.
- **Warranties and legal compliance** – provision of warranties relating to the data (e.g., with respect to quality, accuracy, availability, reliability, suitability), or their exclusion where data is provided "as is"; commitments that data is lawfully provided, does not infringe thirdparty rights, and that both parties comply with personal data and IP rules for protected content; indication of provenance and applied quality controls.
- **Other key provisions** – provisions on reports and audits, changes and additional data, price and payment, security, confidentiality, liability and indemnification, term and termination, assignment, representations and warranties, governing law and dispute resolution, as well as other aspects (such as cooperation and information duties), needed to comply with applicable AIrelated data governance requirements (see question 17).

Certain contracts must also respect mandatory rules. For instance, **contracts providing personal data** must comply with the General Data Protection Regulation (GDPR), **contracts licensing works or performances** (see question 11) must meet requirements on remuneration and transparency under the Digital Single Market (DSM) Directive and any national formalities, **data sharing contracts** under the Data Act must avoid prohibited non-negotiated terms in B2B scenarios (e.g., excluding liability for intent or gross negligence), **data sharing contracts with data intermediation services**, such as data marketplaces, or **copyright management organisations** (see question 12) must respect the specific conditions under the Data Governance Act (DGA) and the Collective Rights Management (CRM) Directive, and **contracts for data processing services** (including database as a service) must include mandatory provisions to facilitate switching between providers under the Data Act.

To whom this is relevant

The organisation that trains, validates and tests the AI system should ensure that the contract contains provisions allowing it to use the data as intended (without prejudice to mandatory requirements that shall be met as indicated above). By default, this is the **AI provide**r. See, however, question 1 for an explanation of who the AI provider is.

AI Checklist — Contracts for training, validation and testing data

Goal:

Assess:

- Whether contracts allow the data to be used as intended and duly protect the parties – relevant for all AI systems.
- Whether contracts comply with applicable obligations – applicable to all AI systems.

Questions:

- What conditions and limits of use does the contract set (types and purposes of use, field, users, territory, term, warranties)?
- Do the authorised uses match the AI system's purposes and risk profile?
- Are there territorial, sectoral, exclusivity or non-compete clauses that constrain use?
- Are the data provider and AI provider actual or potential competitors, and does either hold a dominant position in any relevant market?
- Which data risks are anticipated (e.g., lack of clarity in authorised and prohibited uses, licence does not match intended use, competition risks)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Procurement; Legal; Data including Data Analytics and Data Protection Officer (see question 18).

The risk of anti-competitive effects of data sharing agreements

Data sharing agreements are key for accessing data for AI, but they must not prevent, restrict or distort competition in the EU, whether between competitors or along the value chain. Exchanges of commercially sensitive information, such as competitors pooling data in a common database and accessing each other's data, can be anticompetitive. Data sharing by dominant undertakings must also not foreclose access to essential datasets (for example, via excessive pricing, restrictive licences or exclusivity).

Where data licensing qualifies as a technology transfer, the revised Technology Transfer Block Exemption Regulation (TTBER) offers safe harbours, provided its conditions are met (notably the market-share thresholds, the absence of hardcore restrictions such as price-fixing or output limitations, and the proper treatment of ancillary restraints like exclusivity or grant-backs). More broadly, anticompetitive agreements, decisions or concerted practices may still be lawful if they improve production or distribution or promote technical or economic progress, give consumers a fair share of the benefits, use only indispensable restrictions, and do not eliminate competition.

To unlock pro-competitive benefits, data sharing agreements should be carefully scoped, rely where possible on aggregation and time-lagging, avoid unjustified exclusivity, and include use restrictions to prevent collusive uses of data or AI.

Risk mapping:

Data contract risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What main aspects should be considered when sourcing components and services?

Components and services for AI development are typically procured under binding agreements, which should generally cover the following:

- **The components and services provided** – clear description of items and services, delivery dates and acceptance, as relevant. For physical items, Incoterms can be used to allocate delivery, risk and costs in international trade transactions. Training, knowledge transfer and support/maintenance should also be envisaged where needed.
- **Ownership, risk and conditions for use** – for physical items, transfer of ownership and risk is typically determined in the contract, usually on final acceptance. For intangible items, use may be granted under a licence or the item may be transferred to the client.
- **Warranty for defects** – a warranty period is often established during which the provider corrects non conformities (and, if agreed, optimises performance) at no extra cost.
- **SLAs** – the contract may foresee performance metrics and consequences of breach, such as service credits or penalties (whether liquidated damages or more punitive mechanisms), and, if useful, rewards for exceeding targets.
- **Other key provisions** – provisions on reports and audits, changes and additional components/services, price and payment, security, confidentiality, liability and indemnification, export controls, term and termination, assignment, representations and warranties, governing law and dispute resolution. Specific duties of the provider to supply the information and assistance needed for compliance are especially relevant.

Some types of component sourcing are subject to **mandatory requirements**. For instance:

- Under the **AI Act**, contracts for the supply of AI systems, AI models, tools, services, components or processes to be used or integrated in high-risk AI systems must specify in writing the information, capabilities, technical access and other assistance needed for compliance with AI Act obligations (with possible model terms to be developed by the AI Office). This does not apply to certain free and open-source tools and components made available to the public, though documentation practices (e.g., model cards, data sheets) are encouraged.
- Under the **Cyber Resilience Act**, manufacturers must exercise due diligence when integrating components sourced from third parties (including some free and open-source software) (see also question 3).
- Under the **Ecodesign Regulation**, supply chain actors may be required, in delegated acts, to provide relevant product-related information on request and free of charge to manufacturers and others (see also question 4).
- Under the **Data Act**, contracts for data processing services (such as IaaS, PaaS and SaaS, storage and database as a service) must include mandatory provisions to facilitate switching for customers in the EU.

For an analysis of the anti-competitive risks of contracts for sourcing components and services, see 🔍 The Risk of Anti-Competitive Effects of Sourcing Contracts and Research Partnerships.

To whom this is relevant

The AI system is trained, validated and tested by the **AI provider**. Hence, the AI provider should ensure the contract contains provisions allowing it to use the components and receive the services as intended (without prejudice to mandatory requirements that shall be met as indicated above). See, however, question 1 for an explanation of who the AI provider is.

AI Checklist — Sourcing of components and services

Goal:

Assess:

- Whether contracts for components and services allow for their use as intended and duly protect the parties – relevant for all AI systems.
- Whether the contracts comply with applicable obligations – applicable to high-risk AI systems under the AI Act, and to other AI systems depending on the law at stake.

Questions:

- What services and components are provided, by whom and under which conditions?
- Do they match the AI system's purposes and risk profile?
- Is there a clear inventory linking each set of components/deliverables to concrete uses (training, validation, testing, evaluation, etc.)?
- Are there territorial, sectoral, exclusivity or non-compete clauses that constrain use?
- Are the data provider and AI provider actual or potential competitors, and does either hold a dominant position in any relevant market?
- Which risks are anticipated (e.g., inadequacy of components and deliveries, mismatch between acquisition and purpose, competition risks)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Procurement; Legal.

Risk mapping:

Sourcing risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

The integration of AI models in AI systems

The AI Act defines general-purpose AI (GPAI) models as models with significant generality that can perform many distinct tasks and be integrated into diverse downstream systems, excluding models used only for R&D or prototyping before being placed on the market (see also question 16). Providers of GPAI models have specific responsibilities along the value chain because their models may form the basis of many downstream systems, whose providers need sufficient information to integrate them lawfully and safely.

GPAI providers must therefore give downstream integrators minimum model information, including on capabilities, limitations, acceptable use, integration, and (where applicable) on data used for training, testing and validation, while protecting IP and trade secrets. This obligation does not apply to models released under a free and open-source licence, unless they are classified as having systemic risk. Where models are open source without systemic risk, parameters, architecture and usage information must be public and downstream providers are expected to rely on this public information rather than additional contractual disclosures.

What main aspects should be considered in research and development partnerships for AI?

AI development in R&D partnerships is typically governed by binding agreements, which should generally cover the following:

- **Scope and lifecycle coverage** – definition of what is being researched and developed, intended purpose and lifecycle phases in scope. It is also important to clarify from the outset how any resulting deliverables may later be used or exploited by each party (e.g., internal use, further R&D, commercialisation together or separately), bearing in mind the IP framework that will apply.
- **Roles and responsibilities** – allocation of roles and responsibilities between the parties, ideally mapped in line with relevant legal classifications (e.g., AI provider, manufacturer, controller or processor – see questions 1 and 5 for an explanation of who these actors are).
  - ↳ The determination of each party's role should consider how the future exploitation of results may trigger legal obligations for each party, especially where joint exploitation makes it important to be able to hold the other party liable for non-compliance.
- **Intellectual property** – the determination of who owns the deliverables and how they can be protected. A robust IP management framework is essential.
  - ↳ In collaborative projects, it is essential to clarify when results are jointly owned (and require joint decisions) and when distinct components remain individually owned and separable, to avoid deadlock on exploitation.
- **Governance** – processes for day-to-day management, record-keeping and reporting, IP decision-making, and escalation/dispute resolution.
- **Other key provisions** – provisions on changes, security, confidentiality, liability and indemnification, term and termination, assignment, representations and warranties, governing law and dispute resolution. The use of regulatory sandboxes and real-world testing (see question 50) may be foreseen where relevant.

EU law also provides specific flexibilities for research:

- The **AI Act** does not apply to product-oriented research, testing and development activity prior to the AI system being put into service or placed on the market (except as relates to testing in real world conditions and to requirements for testing in regulatory sandboxes – see question 50).
- Under the **Digital Single Market** (DSM) **Directive**, research organisations and cultural heritage institutions with lawful access to works or other subject matter may carry out text and data mining (TDM) for scientific research without the right holder's authorisation (see question 12).
- Under the **General Data Protection Regulation** (GDPR), personal data (including special categories, subject to extra safeguards) may be further processed and stored for longer periods for scientific research purposes, provided appropriate technical and organisational measures are in place to protect data subjects' rights and freedoms.

The AI Act also expressly allows the processing of personal data collected for other purposes in a regulatory sandbox for the purpose of developing, training and testing certain AI systems in the public interest (see question 50).

To whom this is relevant

In general, all organisations developing AI under a partnership should ensure that the contract contains the provisions detailed above. From the point of view of the AI Act, this would be the **AI provider** developing the AI system (relevant where the AI system researched and developed is put into service or placed on the EU market). See, in any case, question 1 for an explanation of who the AI provider is.

AI Checklist — Research and development partnerships

Goal:

Assess:

- Whether research and development partnerships duly reflect the parties' goals and duly protect them – relevant for all AI systems.
- Whether research and development partnerships take advantage of conditions provided by law for research – relevant for all AI systems.

Questions:

- Who owns pre-existing elements/items and foreground IP?
- How is use of the foreground IP addressed (e.g., use for internal purposes, use for further research or development, independent or joint exploitation)?
- Are there territorial, sectoral, exclusivity or non-compete clauses that constrain use?
- Are the partners actual or potential competitors, and does either hold a dominant position in any relevant market?
- Which risks are anticipated (e.g., unclear ownership of pre-existing or foreground IP, inadequacy or limitations on use, competition risks)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Legal; Data including Data Protection Officer (see question 18).

The risk of anti-competitive effects of sourcing contracts and research partnerships

Component and services contracts and partnerships are important for AI development, but they must not prevent, restrict or distort competition in the EU, whether between competitors or along the value chain. They must also not foreclose access to essential inputs held by dominant undertakings (e.g., cloud and compute, specialised chips, key models), for instance, through refusals to supply, discriminatory terms, restrictions on interoperability, margin squeezes or exclusive partnerships that raise entry barriers or mimic de facto mergers.

The EU R&D Block Exemption Regulation (R&D BER) offers safe harbours for certain R&D agreements, provided its conditions are met (notably the market-share thresholds – the parties' combined market share must not exceed 25% –, access to results – all parties have full access to results and indispensable pre-existing know-how –, and the absence of hardcore restrictions such as price-fixing, output limitations or restrictions on R&D). More broadly, anticompetitive agreements, decisions or concerted practices may still be lawful if they improve production or distribution or promote technical or economic progress, give consumers a fair share of the benefits, use only indispensable restrictions, and do not eliminate competition.

To unlock pro-competitive benefits while limiting harms, contracts should avoid anticompetitive clauses, such as on price or market/customer allocation, and be prepared for closer scrutiny of structured AI partnerships.

Risk mapping:

Partnership risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Are there any governance process that should be implemented while developing the AI system?

AI development requires lifecycle governance, particularly for high risk systems, and in coordination with product safety, cybersecurity, sustainability, and data protection rules (see questions 2, 3, 4 and 5).

In broad terms, depending on the AI system at stake, key governance procedures are:

- **Quality management system**
  The implementation of a documented quality management system (QMS) is required under the AI Act for high-risk systems (to cover, among others, the topics below).
  Where safety, resilience and sustainability laws already require QMS or equivalent, AI Act requirements can be integrated to avoid duplication.
  Where safety, resilience and sustainability laws already require QMS or equivalent, AI Act requirements can be integrated to avoid duplication.
- **Risk management system/assessment**
  High risk AI must be subject to a continuous and iterative risk management process that identifies, analyses, estimates, evaluates and manages risks to health, safety and fundamental rights (see question 2), with systematic review and updating.
  Where product safety and cybersecurity legislation require risk assessments, AI Act risk management steps may be integrated into or combined with those existing processes to avoid duplication.
- **Record-keeping systems**
  Record-keeping systems and procedures are required under the AI Act for high-risk systems to ensure that all relevant documentation and information is recorded and kept (see also questions 26 and 27).
  Similar long term documentation duties exist under product safety and cybersecurity legislation, even for non high risk AI embedded in products.
- **Post-market monitoring system**
  The AI Act requires the implementation of a post-market monitoring system, based on a post-market monitoring plan, to verify ongoing compliance.
  Where a post market monitoring system already exists under product safety laws, the AI Act system may be integrated into systems and plans already existing under that legislation.
- **Data management system**
  High risk AI must be supported by documented data management systems and processes so that datasets meet AI Act quality requirements (see question 13).
  Where personal data is processed, GDPR-compliant measures must be implemented regardless of risk level.
- **Regulatory compliance strategy | Conformity assessment procedures, incident reporting procedures** and **accountability framework** (see questions 19, 21 and 18 respectively).

↳ Under the new updated version of the AI Act, where AI is a safety component of certain regulated products or is itself such a product subject to third-party conformity assessment for health and safety risks, specific AI Act's obligations may be limited by delegated acts of the Commission. This may benefit products such as medical devices and radio equipment. Machinery is excluded from the full high-risk AI regime, with the Machinery Regulation supplemented by delegated acts to incorporate most AI Act health and safety requirements.

To whom the obligations apply

The obligations apply to the **AI provider** for high-risk AI systems under the AI Act, the **product manufacturer** under product safety, resilience and sustainability legislation, and the **data controller** under the GDPR. Under the AI Act, the organisation qualifying as the AI provider is the data controller for the data used to train, validate and test the AI system. For an explanation of who these actors are, see questions 1 and 5.

AI Checklist — Governance Processes

Goal:

Assess:

- Whether all required governance processes have been implemented – applicable to high-risk AI systems under the AI Act, and to all AI systems subject to safety, resilience and sustainability laws and to the GDPR.
- Whether the governance processes have been implemented in a manner avoiding duplication – relevant for high-risk AI systems.

Questions:

- Are there formal, repeatable processes for AI governance?
- Are those processes laid down in writing in internal policies or otherwise?
- Are those processes integrated with existing compliance frameworks?
- Which risks are anticipated (e.g., formal governance processes but weak governance in practice, fragmented governance processes, governance blind spots such as lack of governance processes for lower-risk systems, lack of mechanisms for updating governance processes, inconsistent application across AI systems, inability to demonstrate effective implementation)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Quality; Legal; Compliance; Data including Data Protection Officer (see question 18).

Risk mapping:

Governance processes risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Are there any governance structures that should be implemented for developing AI systems?

The development of AI systems must consider not only the implementation of a set of governance procedures (as seen in question 17), but also of governance structures effectively able to ensure legal compliance.

Although governance structures are not specifically mandated in most cases, two main aspects are worth mentioning:

- **Accountability framework** – one of the elements of the quality management system for high-risk AI systems is an accountability framework. This framework must set out the responsibilities of the management and other staff with respect to all aspects covered by the quality management system.
  - ↳ When personal data is processed, the General Data Protection Regulation (GDPR) also enshrines the principle of accountability, to reflect the responsibility for compliance with the GDPR principles (see questions 5 and 12).
- **Data Protection Officers (DPOs)** – data protection officers are required to be designated in certain cases under the GDPR, such as when the core activities at stake consist of processing operations which, by virtue of their nature, scope and/or purposes, require regular and systematic monitoring of data subjects on a large scale; the core activities consist of large-scale processing of special categories of data or personal data relating to criminal convictions and offences; or the processing of personal data is carried out by a public authority or body (except for courts acting in their judicial capacity). EU or national law may require the designation of DPOs in other situations as well.
  - ↳ AI development (regardless of the AI system being high-risk or not) requires the processing of large amounts of data. However, such processing will not meet the referred requirements in all cases. Therefore, the determination of the need to designate a DPO in AI development depends on a case-by-case assessment. Nevertheless, it could be a best practice to have a person with DPO capacities and functions to ensure adequate levels of privacy and data protection within AI development.

To whom this is relevant

The obligations apply to the **AI provider** for high-risk AI systems under the AI Act and the **data controller** under the GDPR. For an explanation of who these actors are, see questions 1 and 5.

Note, in any case, that the **data processor** (i.e., the one processing personal data on behalf of the controller) is also under the obligation to designate a DPO in the above indicated situations.

AI Checklist — Governance structures

Goal:

Assess:

- Whether the required governance structures have been implemented – applicable to high-risk AI systems under the AI Act and to all AI systems under the GDPR.
- Whether governance processes sufficient to ensure legal compliance and promote a culture of transparency, accountability, and ethical behaviour have been implemented – relevant for all AI systems.

Questions:

- Are there clear governance structures for AI?
- Are tasks and responsibilities clearly allocated to each governance structure, in line with the applicable legal obligations?
- Are AI structures duly coordinated and compatible with other internal structures (e.g., on data protection, cybersecurity, among others)? How?
- Are there channels to escalate concerns?
- Which risks are anticipated (e.g., formal governance structures but weak governance in practice, governance blind spots such as lack of governance structures for lower-risk systems, lack of mechanisms for internal coordination, poor coordination with competent entities)?

Expertise required:

Product & Engineering; Quality; Legal; Compliance; Data Protection Officer.

Ethical leadership

Adopting a dedicated AI leadership structure should be a priority for AI providers, as it supports legal compliance and fosters a culture of transparency, accountability and ethical behaviour, while also strengthening trust and offering strategic market advantages.

Key roles include a Chief AI Officer for strategy and value alignment, an AI Officer for legal compliance (similar to a DPO), and an AI Ethics Board or Committee providing independent oversight and guidance. Such a structure should monitor compliance, align practices with best practice (see question 22), engage with external stakeholders and regulators, and actively promote a culture of responsibility, transparency and respect for fundamental rights across all AI-related activities.

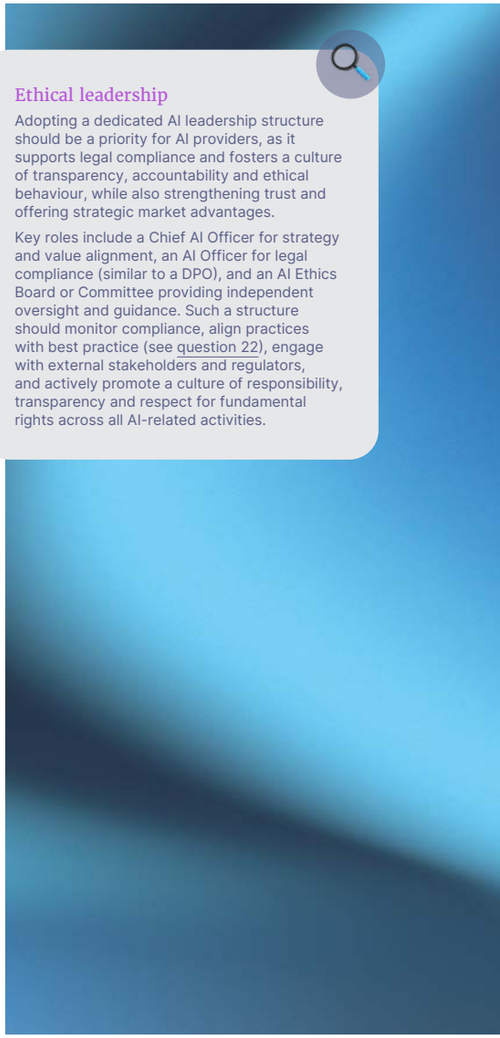

Risk mapping:

Governance structures risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can compliance with legal obligations be shown? Are there any mandatory procedures?

There are mandatory procedures aimed at demonstrating AI compliance, primarily under the AI Act for high risk systems, alongside product safety, resilience, sustainability, and data protection laws (see questions 2, 3, 4 and 5).

In broad terms, such procedures include:

- **Conformity assessment**
  A conformity assessment aims to demonstrate compliance with legal requirements.
  The AI Act determines which conformity assessment procedures to follow depending on the high-risk AI system at stake.
  - ↪ Conformity assessment procedures can be based on internal controls (self-assessment) or be a third-party assessment by a "notified body". Notified bodies can only perform assessments within the scope of their notification, as defined by the list of codes, categories and corresponding types of AI systems established under the AI Act.
  - ↪ Under the new updated version of the AI Act, conformity assessment and tests of certain high-risk systems are the responsibility of the AI Office (which then entrusts the performance of the assessment to notified bodies). These systems include general-purpose AI systems where the model and system share the same developer (save some exceptions) and systems integrated into very large online platforms or search engines (VLOPs and VLOSEs).

  Conformity assessments are also a central element of product safety, resilience and sustainability laws, with legal coordination to avoid duplication for high risk AI indicated by law, including:
  - The AI Act requires high risk AI embedded in or that are regulated products (e.g. medical devices, radio equipment, toys) to follow existing product law conformity assessments, incorporating AI Act requirements, including quality management system requirements – see question 17.
  - The Cyber Resilience Act (CRA) provides that some high risk systems follow AI Act conformity assessments, while others follow CRA procedures.

  Declarations of conformity and CE marking also show compliance – see question 26 for more details.
- **Technical and organisational measures** (TOMs) whenever personal data is processed.
  Technical and organisational measures aim not only to ensure (see question 17), but also demonstrate, that data processing is performed in accordance with the General Data Protection Regulation (GDPR).

In addition, certifications can be applied for with a view to demonstrate compliance with applicable law.

- ↪ **Privacy certifications, seals and marks** can be used to demonstrate compliance with the GDPR (see 🔍 Privacy Certification, Seals and Marks, and question 22).
- ↪ **Cybersecurity certifications** can be used to presume compliance (see 🔍 The European Cybersecurity Certification Scheme and question 22).

To whom the obligations apply

The obligations apply to the **AI provider** for high-risk AI systems under the AI Act, the **product manufacturer** under product safety, resilience and sustainability legislation, and the **data controller** under the GDPR. For an explanation of who these actors are, see questions 1 and 5.

↪ Coordination provisions aim to avoid the duplication of obligations, as seen above. This is especially relevant when the same organisation is the AI provider, product manufacturer and data controller. However, when different organisations/persons are involved, assistance and coordination must be ensured.

AI Checklist — Compliance Demonstration

Goal:

Assess:

- Whether the required conformity assessments have been implemented – applicable to high-risk AI systems under the AI Act and to all AI systems under product safety, resilience and sustainability laws.
- Whether technical and organisational measures (TOMs) have been implemented – applicable to all AI systems when personal data is protected.

Questions:

- Are conformity assessments envisaged? Which ones?
- Are conformity assessments being combined? Which ones?
- Are technical and organisational measures (TOMs) being implemented if personal data is processed? What are those measures?
- Which risks are anticipated (e.g., selection of wrong conformity assessment, incomplete performance of the conformity assessment, poor coordination among applicable conformity assessments, generic non-risk-based TOMs, mismatch between TOMs and the nature, scope, context and purposes of data processing activities)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Legal; Compliance; Data including Data Protection Officer (see question 18).

Risk mapping:

Compliance demonstration risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

The different conformity assessment procedures

Decision No 768/2008/EC sets out the various types of conformity assessment procedures used in EU product legislation. In each specific act, the legislator selects the appropriate procedure for given products and risks, typically requiring (a) an assessment of the design or a specimen against legal requirements (EU type examination) and, where relevant, (b) a subsequent check that products conform to the approved EU type.

Depending on the applicable instruments, the product type and its intended use, different conformity assessment models may be available. These can be carried out internally by the manufacturer (sometimes with the need for external certification of the internal procedure) or may require the direct involvement of a third party. Such third party is a "notified body", i.e., a "conformity assessment body" that submits an application for notification (and meets the established legal requirements) to a competent national notifying authority. Those entities are responsible for conducting assessments and issuing certificates that attest to the conformity of products.

Under the new updated version of the AI Act, manufacturers may use self-assessment for certain high-risk AI products if permitted by the corresponding safety product laws, provided that they apply relevant harmonised standards or specifications under the AI Act (see question 22).

How can it be ensured that the persons dealing with AI meet the applicable obligations?

One of the key factors for ensuring legal compliance is the level of **literacy** among staff and other individuals involved in AI development.

- AI literacy refers to the skills, knowledge and understanding that allow the making of informed decisions about AI systems (regardless of them being high-risk or not) and the gaining of awareness about the opportunities and risks of AI and possible harm it can cause.

Under the AI Act, measures must thus be taken to support the development of AI literacy of staff and other individuals dealing with AI systems, taking into account their technical knowledge, experience, education and training, the context the AI systems are to be used in, and the persons or groups of persons on whom the AI systems are to be used. The new updated version of the AI Act clarifies that this must not be understood as requiring that any specific level of AI literacy of any individual be guaranteed.

- To support the development of AI literacy in a proportionate way, organisations may consider role-specific measures. For instance:
  - Technical staff (e.g., data scientists, engineers) could receive training on AI development, training, validation and testing, data governance, robustness;
  - Legal and compliance staff could receive training on the AI Act, data protection, product safety, resilience and sustainability, liability and sector-specific rules;
  - Management and decision-makers could receive training enabling them to understand AI-related risks, governance obligations and accountability frameworks.

The above is reinforced by other legal frameworks, such as those on data protection.

- Continuous learning and updating are important. Given the rapid evolution of AI technologies and regulatory frameworks, training and awareness activities should ideally not be one-off exercises, but ongoing processes, updated from time-to-time to reflect technological developments, new risks, incident learnings and regulatory changes.
- Literacy measures can be embedded in AI governance and other AI controls (e.g., incident handling), so that feedback from such activities is integrated into literacy efforts and helps keep them up to date.

The Commission and Member States must support and facilitate compliance with AI literacy obligations, in particular by SMEs (including through the facilitation of the drawing up of codes of conduct – see question 22). For this purpose, the Commission must publish practical examples for compliance. In addition, the AI Board will issue recommendations to support the Commission and Member States, including setting common objectives. Likewise, the Cyber Resilience Act (CRA) requires Member States to promote measures and strategies aimed at developing cybersecurity skills.

To whom the obligations apply

The literacy obligation established by the AI Act for all AI systems applies to **AI providers** (for an explanation of who this actor is, see question 1).

AI Checklist — Literacy

Goal:

Assess whether literacy measures have been implemented – applicable to all AI systems under the AI Act.

Questions:

- Who needs AI literacy and for what?
- What level (e.g., basic vs. advanced) and type of literacy (e.g., technical, legal) is needed?
- Which measures are envisaged to deliver and maintain literacy?
- Is there a clear inventory or mapping of literacy measures linking each measure and its recipients within the organisation?
- Are literacy measures documented and recorded? How?
- Are literacy measures embedded in, and coordinated with, AI governance measures?
- Which risks are anticipated (e.g., undefined or narrow scope of literacy recipients, one-size fits all or superficial training, limited literacy measures not covering all required aspects, no literacy documentation or records, no connection between literacy and governance measures impacting literacy updating)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Legal; Compliance; Data including Data Protection Officer (see question 18); Management.

Risk mapping:

Literacy risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What if the AI system is non-compliant or malfunctions?

When an AI system malfunctions or breaches legal requirements, the law foresees reporting, corrective action and liability.

Reporting of incidents and vulnerabilities

Serious incidents of high-risk systems must be reported to market surveillance authorities (or, in some cases, to the AI Office for systems under the oversight of the AI Office, which comprise, in essence, AI systems based on general-purpose AI models where the model and system share the same provider – subject to specified exceptions –, and AI systems that constitute or are integrated into very large online platforms or very large online search engines). This is followed by investigation, risk assessment and corrective action.

In 2025, the European Commission issued draft guidance and a reporting template on serious AI incidents.

Product safety, resilience and data protection laws also require reporting, among other actions depending on the law at stake. For example, actively exploited vulnerabilities and severe incidents must be notified to CSIRTs (Computer Security Incident Response Team)/ENISA (the EU Cybersecurity Agency) and users under the Cyber Resilience Act (CRA), and personal data breaches must be notified to the competent authority (and data subjects) under the General Data Protection Regulation (GDPR).

Corrective measures

Despite some particularities established in the AI Act, nonconformity or risk of high-risk systems triggers corrective measures such as bringing the system into conformity, or withdrawing, disabling or recalling it, among other actions (e.g., informing deployers, investigating the causes of the risk).

Where non-high-risk AI is found non-compliant with the AI Act, market surveillance authorities can still require corrective measures.

Similar duties exist under product safety, resilience and sustainability laws (e.g., bring into conformity, withdrawal/recall dangerous or non compliant products).

↪ Under the new updated version of the AI Act, where AI is a safety component of certain regulated products or is itself such a product subject to third-party conformity assessment for health and safety risks, specific AI Act's obligations may be limited by delegated acts of the Commission. This may benefit products such as medical devices, radio equipment or toys. Machinery is excluded from the full high-risk AI regime, with the Machinery Regulation supplemented by delegated acts to incorporate most AI Act health and safety requirements.

Liability

Breaches may lead to administrative **fines**, civil liability and, in some cases, criminal liability.

When it comes to fines, while some acts leave fines to Member States, many already set penalties for certain breaches.

↪ The AI Act provides for fines up to €35 million or 7% of global turnover for serious breaches, and up to €15 million or 3% for others, with reduced caps for SMEs and small mid cap enterprises (SMCs).

Civil liability arises where breach of legal or contractual duties causes damage, while **criminal liability** is defined mainly at national level (e.g., for IP infringements). Still, criminal law may also apply, in particular under cybercrime rules criminalising illegal access to information systems and data.

To whom the obligations apply

The obligations apply to the **AI provider** under the AI Act, the **product manufacturer** under product safety, resilience and sustainability legislation, and the **data controller** under the GDPR. For an explanation of who these actors are, see questions 1 and 5.

↪ When it comes to reporting and corrective actions, in cases where the AI provider, the product manufacturer and the data controller are different organisations/persons, assistance a nd coordination must be ensured to the extent breaches of one applicable law impact other stakeholders – e.g., non-conformity of the physical product impacts the AI system integrated therein or vice-versa.

AI Checklist — Breach and malfunctions

Goal:

Assess whether there are processes to address incident reporting and corrective actions – applicable to high-risk AI systems under the AI Act, and to all AI systems subject to safety, resilience and sustainability laws and to the GDPR (as well as to all AI systems when it comes to certain corrective actions under the AI Act).

Questions:

- Are there clear processes to handle AI non compliance and malfunctions?
- Are there enough capable human resources to handle these processes?
- Are incident taxonomies and reporting templates in place?
- Do processes include improvement steps after incidents, malfunctions or breaches, with criteria for reporting?
- How are corrective actions decided and implemented, and how are these processes aligned with wider governance arrangements (see question 17)?
- Which risks are anticipated (e.g., formal processes but weak implementation in practice, fragmented processes, misalignment among processes required under several laws, lack of mechanisms for updating processes)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Risk & Security; Legal; Compliance; Data including Data Protection Officer (see question 18).

Civil and criminal liability arising from acts by an AI system

"Algorithm liability" addresses the circumstances in which AI behaviour, stemming from its features, can trigger legal liability. It is often hard to prove causation and fault given AI's black-box nature, with liability generally requiring a person at fault. This raises the question of whether and when the AI provider is liable, especially where the system continues to learn and change after deployment in ways not predetermined by the provider.

The AI Act's human oversight requirements point to a duty to maintain human control. If the AI's design makes oversight impossible, breaches of legal duties and resulting civil or criminal liability may be argued. Where the AI Act does not apply (for example, a privately used autonomous vehicle causing an accident), liability will typically hinge on breach of a duty of care or a duty to prevent that harm. However, AI autonomy affecting the foreseeability of damage or harmful event may exclude fault and thus liability.

Separately, the new Product Liability Directive (PLD II) introduces no-fault liability for defective products, including AI, though it does not apply to non-commercial free and open-source software.

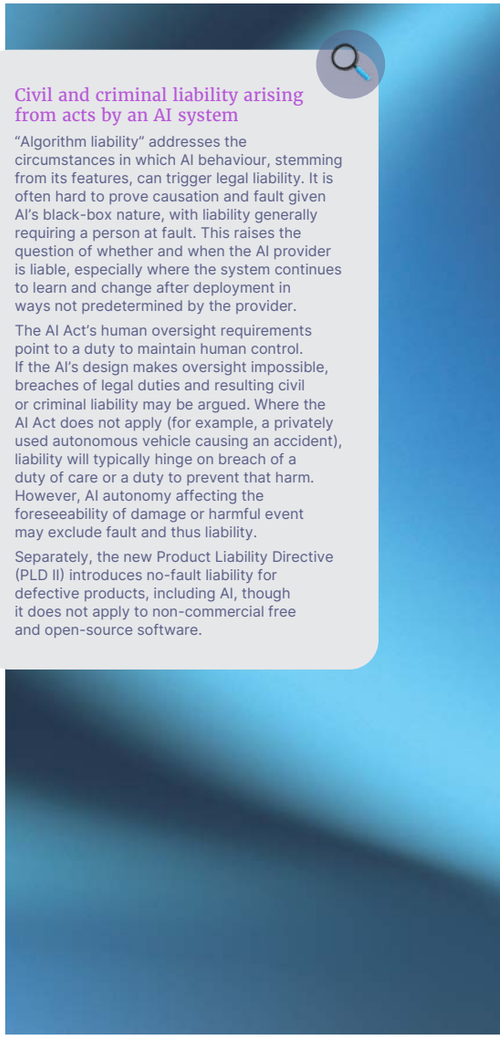

Risk mapping:

Breach and malfunctions risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Are there standards, guidelines or best practices that can be followed?

Standards, certifications, codes of conduct and guidelines are tools that help organisations implement and evidence compliance with AI, product safety, cybersecurity, sustainability and data protection rules.

Standards and common specifications

Harmonised standards (developed by CEN, CENELEC or ETSI and cited in the EU Official Journal) and Commission "common specifications" translate legal requirements into technical rules. Using these gives a presumption of conformity. Dedicated harmonised standards are being developed for AI and cyber resilience, while product safety standards have already been approved. Other standards (e.g., ISO) can support compliance but do not themselves create this presumption unless incorporated by reference in the harmonised standard.

↪ The Commission asked CEN and CENELEC to develop a set of standards in ten areas for high risk systems: risk management, data and data governance, record keeping, transparency, human oversight, accuracy, robustness, cybersecurity, quality management and conformity assessment. Other standards are also being developed, notably on bias management and sustainable AI.

Certifications and labels

Certification formally attests that products, services or organisations meet defined requirements and can create presumptions of compliance. See 🔍 The European Cybersecurity Certification Scheme, 🔍 Green Certification and Labels and 🔍 Privacy Certification, Seals and Marks.

↪ The presumption of compliance created by cybersecurity certifications or statements of conformity issued for an AI system under a cybersecurity scheme presume conformance with the cybersecurity requirements of the AI Act for high-risk systems, and of the Machinery Regulation and Cyber Resilience Act for any AI system.

Codes of conduct

Voluntary codes can operationalise legal obligations and, under the AI Act, may foster the application of legal requirements to non-high-risk systems, including on aspects such as sustainability (energy-efficient programming and efficient design, training and use of AI), literacy and trustworthy AI principles. Under the GDPR, adherence to approved codes can be used as evidence of compliance.

Guidelines

The Commission, the AI Office and AI Board, as well as authorities such as ENISA (the EU Cybersecurity Agency) and the EDPB (the European Data Protection Board), issue guidelines and opinions that clarify how to apply the laws in practice.

↪ Guidance on the AI Act includes the Guidelines on the definition of an AI system (2025) and on prohibited AI practices (2025), as well as forthcoming guidelines on high-risk classification (draft guidance issued in May 2026), transparency (draft guidance issued in May 2026), serious incidents (draft guidance issued in 2025), high-risk requirements, obligations for providers and deployers, responsibilities along the AI value chain, substantial modification, post-market monitoring, quality management systems for SMEs and small mid-cap enterprises (SMCs), and interplay of the AI Act with other EU law. Under the new updated version of the AI Act, guidelines will also be adopted to promote a consistent, complementary and proportionate application of AI-specific requirements alongside existing EU product rules.

To whom this is relevant

The **AI provider** (under the AI Act), **product manufacturer** (under product safety, resilience and sustainability legislation) and **data controller** (under the GDPR) may benefit from the above. For an explanation of who these actors are, see questions 1 and 5.

AI Checklist — Standards, certifications and guidance

Goal:

Assess whether there are standards, certifications and guidance that may apply – relevant for high-risk AI and non-high-risk AI systems under the AI Act, depending on the standard and guidance at stake, and to all AI systems subject to safety, resilience and sustainability laws and to the GDPR.

Questions:

- Which standards and common specifications apply, do they create a presumption of conformity, and will they be used?
- Which certifications or labels are available, do they evidence or presume compliance, and are they planned to be used?
- Have relevant codes of conduct or guidelines been identified, and is adherence or use envisaged?
- Are there clear processes to identify, assess, decide on and implement standards, specifications, certifications, codes and guidelines, and to align this with wider governance (see question 17)?
- Which risks are anticipated (e.g., formal processes but weak implementation in practice, under-relying or over-relying on standards and common specifications, misalignment between certification and legal requirements, treating codes as symbolic commitments?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Risk & Security; Ethics; Legal; Compliance; Data including Data Protection Officer (see question 18).

Risk mapping:

Standards, certification and guidance risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

The ethical principles for AI

Ethics is central to AI. It complements regulation and helps interpret and apply the law where rules are ambiguous, incomplete or still evolving, while remaining more agile and responsive to new societal concerns.

Multiple EU and international frameworks exist, such as the High-Level Expert Group on Artificial Intelligence (AI HLEG), AI Ethics Guidelines for Trustworthy AI and Assessment List for Trustworthy AI (ALTAI), the OECD AI Principles, the UNESCO Recommendation on the Ethics of AI, and the UN HLAB-AI Governance Report. Despite their number, they converge on core principles: human rights and dignity, diversity and non-discrimination, privacy and data governance, human agency and oversight, fairness, harm prevention, transparency and explainability, accountability, and societal and environmental well-being.

These principles have influenced the AI Act itself and also shape its implementation. They help providers navigate compliance, make judgment calls where the law leaves room for discretion, and strengthen trust and confidence in the AI systems they place on the market or put into service.

How can an AI system be protected?

AI systems can generally be protected by intellectual property, mainly as software (copyright) or computer-implemented inventions (patents).

Software

As software, AI systems are protected by copyright if the code reflects human creative choices. Ideas, principles and algorithms as such are not protected, only their expression in code. Copyright requires a human author, so AI-Generated outputs without meaningful human input are not protected, but AI-assisted creations and AI | Human creations can be.

↪ Copyright requires creativity/originality. In order to assess the originality of the AI software, the absence of a purely functional character and the identifiability of original choices, sequences, and combinations, regardless of skill, labour, efforts or investments, are indications of the creativity of a computer program.

Computer-implemented invention

As computer-implemented inventions (CIIs), AI systems can be patented when they have technical character, are new and involve an inventive technical contribution, for example, by solving a technical problem using a computer-implemented method.

↪ CIIs are inventions whose implementation involves the use of a computer, computer network or other programmable apparatus.

Careful drafting of patent applications is required to highlight the technical features of the invention (otherwise the claims may be considered to concern purely mathematical or abstract methods and be excluded from patentability) and to disclose it sufficiently (despite AI's black-box nature), while still preserving other protections such as trade secret.

↪ Because mathematical methods that serve a technical purpose can be patentable, methods for generating training data or training an AI model can contribute to the technical character of an invention when they help achieve a concrete technical effect. This makes it possible, in principle, to obtain European patent protection for AI or machine-learning training methods, or for methods of generating training data, if they can be credibly linked to a reliable and repeatable technical result.

Patent law also requires a human inventor, so inventions generated entirely by AI are not patentable, but inventions where humans design or train AI can be.

AI can also be potentially protected through other means, such as by **trade secret**, **contractual confidentiality** and **technical measures**, with unlawful circumvention potentially amounting to cybercrime (e.g., illegal access or data interference).

See also question 16 on IP issues to consider in contracts for AI development.

To whom this is relevant

The person (legal or natural) that wishes to protect the AI system must assess the applicable routes of protection.

When it comes to **copyright**, the rightsholder over the AI is in principle the natural persons that developed it, or the legal persons designated as the rightsholder under national law – such as an employer or the customer.

When it comes to protection as a **CII**, the holder of the patent will in principle be the inventor, though specific national rules or legal presumptions may be established, such as determining that the holder of the patent is the employer.

The IP holder may be the AI provider under the AI Act, since the provider is the entity that develops or has AI developed and places it on the market, in which case contracts should secure IP ownership for the provider even where development is outsourced. However, the holder may also be a different actor, as entities can become providers by branding the AI system or, in some cases, by modifying it. Where an AI system is modified, the resulting system may qualify as a derivative work or new invention; in principle, the modifier is then the IP holder for that derivative or new invention, without prejudice to the original rightsholder's IP in the underlying AI.

AI Checklist — IP Protection

Goal:

Assess whether the AI system can be protected and by which routes of protection – relevant for all AI systems.

Questions:

- Which elements/assets of AI are to be protected?
- Which routes of protection are possible?
- Are there IP management and governance processes (see also question 17)?
- Who is responsible for these processes?
- Are there sufficient and capable human resources to handle these processes?
- How are these processes coordinated with the governance processes (see question 17)?
- Which risks are anticipated (e.g., unclear routes of protection, unclear ownership of IP, unclear IP management processes, robust processes but weak implementation, lack of coordination between different routes of protection, over-disclosure undermining trade secrets, weak technical protections)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Quality; Legal.

Risk mapping:

IP protection risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Are there any particularities for certain organisations or sectors?

The AI Act addresses certain sectors mostly through its risk-based approach whereby certain AI systems are considered prohibited or high-risk (see 🔍 High-Risk Systems). Under this approach, sector-specific provisions are mostly structured through use cases, such as, among others, for Finance, Healthcare, Employment, Education, Public Services, Critical Infrastructure, and Law Enforcement and Justice (see also question 1).

- **Finance** – high-risk AI includes systems used to evaluate the creditworthiness or credit score of natural persons, affecting access to essential private services (such as loans and insurance).
- **Healthcare** – AI in products covered by existing medical devices legislation falls into the high-risk bucket under the AI Act.
- **Employment** – high-risk AI covers recruitment, worker management and access to self-employment. In turn, AI systems that infer people's emotions are prohibited in workplaces (except when specifically intended for medical or safety purposes).
- **Education** – systems used for access to education and vocational training and for assessment of learning outcomes are also high-risk. In turn, AI systems that infer people's emotions are prohibited in educational institutions (except when specifically intended for medical or safety purposes).
- **Public services** – AI systems used to determine access to, or allocation of, essential public services and benefits (e.g., social security, healthcare coverage) are listed as high-risk.
- **Critical infrastructure** – AI systems used as safety components for managing and operating critical infrastructure (including certain energy, transport and digital infrastructure) are expressly listed as high-risk as well.
- **Law enforcement and Justice** – several AI systems for law enforcement and the administration of justice, as well as for migration, asylum and border control management, are high-risk too. In turn, AI systems that predict a person's likelihood of committing a crime based only on profiling or personality assessment are prohibited.

In addition to this approach, the AI Act further expressly addresses the **financial sector** with a view to facilitating coordination between the AI Act and financial legislation, at two levels: in relation to the quality management system and in relation to logs and technical documentation.

- **Quality management system** (see question 17) – for AI providers that are financial institutions subject to requirements regarding their internal governance, arrangements or processes under Union financial services law, the obligation to put in place a quality management system (with the exception of the obligations relating to the risk management system, post-market monitoring system and reporting of serious incidents) is considered to be fulfilled by complying with the rules on internal governance arrangements or processes pursuant to financial services law.
- **Logs and documentation** (see questions 7 and 27) – AI providers that are financial institutions subject to requirements regarding their internal governance, arrangements or processes under Union financial services law must maintain the logs automatically generated by their high-risk AI systems and technical documentation as part of the documentation kept under Union financial services law.

To whom the obligations apply

The above obligations apply to **providers of high-risk AI systems**, with the specific provisions on the financial sector applying to AI providers that are financial institutions. For an explanation of who the AI provider is, see question 1.

AI Checklist — Sector specificities

Goal:

Assess whether the AI system benefits from dedicated provisions of the AI Act – applicable to high-risk AI systems provided by the financial sector, as the determination of whether the AI system is prohibited or high-risk has been made under question 1.

Questions:

- Is the AI provider a financial institution?
- Is the AI provider subject to requirements regarding its internal governance, arrangements or processes under Union financial services law?
- Which risks are anticipated (e.g., treating AI management, logs and documentation separately from internal processes, duplicated processes, inconsistent controls)?
- Which mitigation measures will be implemented to address the risks?

For other relevant questions, see questions 7, 17 and 27.

Expertise required:

Legal; Compliance.

Risk mapping:

Sector-specific risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above. See also questions 7, 17 and 27.

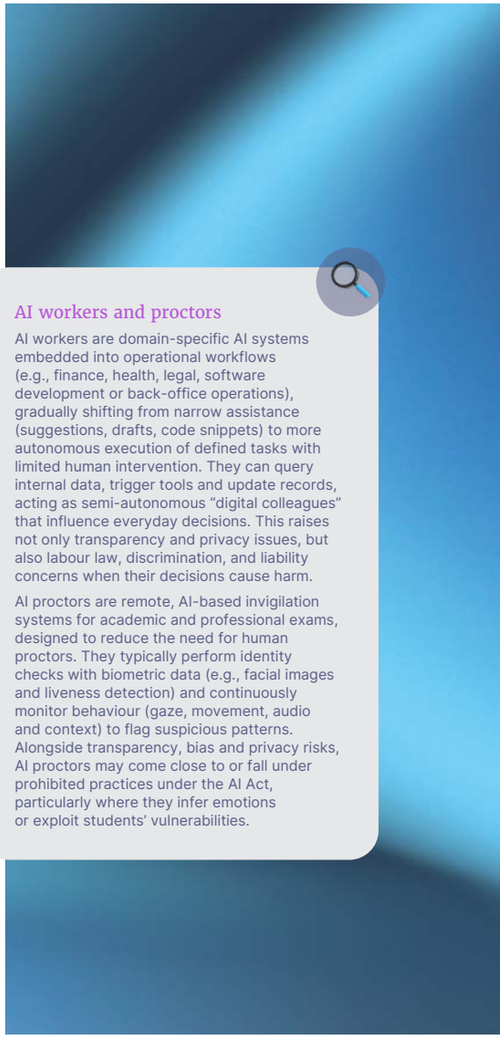

AI workers and proctors

AI workers are domain-specific AI systems embedded into operational workflows (e.g., finance, health, legal, software development or back-office operations), gradually shifting from narrow assistance (suggestions, drafts, code snippets) to more autonomous execution of defined tasks with limited human intervention. They can query internal data, trigger tools and update records, acting as semi-autonomous "digital colleagues" that influence everyday decisions. This raises not only transparency and privacy issues, but also labour law, discrimination, and liability concerns when their decisions cause harm.

AI proctors are remote, AI-based invigilation systems for academic and professional exams, designed to reduce the need for human proctors. They typically perform identity checks with biometric data (e.g., facial images and liveness detection) and continuously monitor behaviour (gaze, movement, audio and context) to flag suspicious patterns. Alongside transparency, bias and privacy risks, AI proctors may come close to or fall under prohibited practices under the AI Act, particularly where they infer emotions or exploit students' vulnerabilities.

III.

Placing AI systems on the market

What it is

Placing an AI system on the market refers, for purposes of this Guide, to the set of activities through which an organisation commercially exploits an AI system. It mainly covers commercialisation – i.e., market deployment –, though the provision of services potentially required to use/deploy an AI system may also be relevant hereto (notably installation, maintenance and support services). Although the last step of the AI lifecycle is AI decommissioning, it can also be considered a step of this stage, under which data archiving or deletion is carried out and the decommissioning of physical items of the AI system (waste management processes) is performed.

The main legal touchpoint is the EU AI Act, complemented by product safety, resilience and sustainability laws. Other laws, such as on data, consumer protection and digital services, are also relevant hereto.

In practice, this stage represents a regulatory gateway in the AI lifecycle. It is the point at which compliance with the obligations detailed in Chapter II becomes legally enforceable and verifiable by the market surveillance authorities. It is also the stage where specific requirements for commercialising the AI system must be met, especially those on documentation, CE marking (where applicable), traceability and registration. Placement on the market (or putting the AI system/product into service) triggers the application of the AI Act and, similarly (despite some particularities), the application of product safety, resilience and sustainability laws.

This Chapter thus focuses on the market placement of the AI system. Other legal requirements that are conditions for, and arise from, such market placement are addressed in Chapter II.

Relevant stakeholders

- AI providers
- Others, notably:
  - AI manufacturers
  - AI vendors, including sellers, licensors, importers, distributors and resellers
  - AI marketplaces
  - AI service suppliers
  - AI user/deployer
  - Infrastructure providers, such as cloud and computing

Key Topics

The main topics relevant for the AI commercialisation stage are:

- Business models
- Conformity features of the AI system (notably EU DoC and CE marking)
- Documentation
- Provision of the AI system – contracts and consumer protection

Summary of main challenges and how to address them

Challenges	How to address the challenges – summary
Determining an effective business model Ensuring that the AI system is placed on the market under a business model, including revenue and pricing structures, that is viable, compliant and aligned with how the system will actually be distributed	• Assess the features and purposes of the AI system, including whether it is a stand-alone software or integrated into another product (tangible or intangible) • Assess existing and anticipated costs (including compliance and ongoing governance costs) and identify viable margins • Determine target customers/segments • Assess market demand, customer preferences and competitor strategies • Determine and document the core revenue and pricing model and how this interacts with your role in the AI value chain For more details, see BUSINESS MODELS.
Placing AI systems on the market in a compliant manner Ensuring AI systems – especially high-risk ones – meet, in addition to the requirements detailed in Chapter II, legal requirements on conformity (EU declaration of conformity and CE marking), documentation (technical documentation and instructions for use), traceability and registration	• Assess the risk level of the AI system and map all applicable legal frameworks beyond the AI Act (e.g., product safety, resilience and sustainability laws; data laws; digital laws) • Prepare, keep and make available all required documentation and information • Ensure that the CE marking is correctly affixed or made accessible, and that it reflects compliance with all applicable EU laws • Register the AI system, where required, in the EU database for high-risk systems • Design an integrated compliance plan so that overlap-ping obligations are met coherently rather than in silos For more details, see MARKING, TRACEABILITY AND REGISTRATION, as well as DOCUMENTATION.
Structuring contracts for AI use Ensuring that contracts for the use of AI allocate rights and obligations clearly (including usage conditions, warranties, rights over input data and AI outputs, liability) and that they do not inadvertently shift roles under the AI Act (for example, by turning a deployer into a provider through modifications)	• Draft and negotiate contracts to allow the use of AI as envisaged, taking into account the features and purpose of the AI system and the parties' respective roles • Reflect any mandatory clauses arising from relevant legal regimes (e.g., on data use, on data processing services such as SaaS) • Define how responsibilities related to AI-specific obligations are allocated between the parties • Assess whether the user qualifies as a consumer and, if so, ensure that consumer protection requirements (including transparency, fairness and remedies) are met in contract terms and system information For more details, see CONTRACTS FOR AI USE.

How to achieve Responsible AI	
	Responsible AI in the commercialisation stage is achieved when AI systems are placed on the market, monetised and licensed in ways that are lawful, ethical and aligned with their risk profile and intended uses. To move in that direction, organisations should ensure that Responsible AI informs how business models are designed, how systems are brought to market and how contracts are structured, notably by choosing revenue and pricing models that do not incentivise unsafe or unlawful practices; ensuring that AI systems placed on the market meet all applicable requirements on conformity, documentation, traceability and registration; and using clear, balanced contracts that allocate rights, responsibilities and liabilities in a way that supports legal compliance, accountability and effective oversight throughout real-world use.

What business models can be implemented to place the AI system on the market?

Different **business models** may be implemented to place AI systems on the market. The aim here, however, is only to show that AI may be placed on the market integrated into another item or **product, notably a physical product** (e.g., smart machinery, autonomous robots, connected vehicles, medical devices) or offered as a **stand-alone digital solution** (e.g., SaaS analytics tools, AI assistants).

- When an AI system is **embedded in a (physical) product**, the system itself is typically made available under a licence granted to the deployer/user. The licensor of the licence may vary: it can be the manufacturer of the product, the developer of the AI system, or any other entity with rights to place the AI system on the market. These roles may be assumed by the same organisation or by different entities.
- When an AI system is offered as a **stand-alone software**, or as a **module of another software/system**, it is likewise typically licensed. Again, the licensor of the licence may vary: e.g., it may be the licensor of the overall software/system, the developer of the AI system, or any other entity with rights to place the AI system on the market.
  AI systems may be provided for local deployment or as software as a service (SaaS/AIaaS), and also via APIs for integration or through marketplaces

Various **revenue and pricing models** can be used and often depend on the delivery model. For on-premise installation, a one-off payment (perpetual licence) is common, sometimes combined with time-limited, renewable licences. In SaaS models, pricing often takes the form of subscription, transaction-based or pay-per-use fees. Economic value, however, can go beyond classic licence or subscription structures.

- In **freemium models**, basic services are provided for free while advanced functionalities are monetised. This is common in AI-powered productivity tools and generative AI applications.
- In **tiered models**, users may choose between basic, pro and enterprise features, with different pricing levels that increase in capability and cost.
- In **data monetisation models**, the AI system is provided at low or no direct monetary cost, but value is extracted through access to user data, which is used for AI training or analytics.
- In **outcome-based AI contracting**, also known as performance-based contracting, payment is linked to measurable results, such as reduced energy consumption or predictive maintenance savings.
- In **dynamic pricing**, prices adjust in real time based on demand, capacity or usage conditions (see also questions 48 and 49).

To whom this is relevant

The choice of business model is particularly important for the AI provider, who develops or has the AI system developed and places it on the EU market or puts it into service under its own name or trademark. However, an organisation may also become a provider simply by branding or, in some cases, modifying an AI system, meaning the AI provider under the AI Act may not be the entity commercialising it. Other actors, such as product manufacturers embedding the AI system and selling to deployers or users, may handle commercialisation and must likewise determine their appropriate business model.

AI Checklist — Business models

Goal:

Assess which business models are most appropriate for the commercialisation of the AI system – relevant for all AI systems.

Questions:

- What is being sold and how?
- What core revenue and pricing model is envisaged, and what are the existing costs, including with compliance, and viable margins over time?
- What are the target customers/customer segments (e.g., enterprises, SMEs, public sector)?
- What are the market demand, customer preferences and competitor strategies?
- Are there third-party components or services in the AI system?
- Which potential customers has feedback been obtained from?
- Which commercialisation risks are anticipated (e.g., underestimation of costs, erosion of margins, vendor lock-in, loss of control over pricing, inadequate revenue model, inadequate pricing model)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Marketing.

Risk mapping:

Commercialisation risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Are there any additional features the AI system should have so that it can be placed on the market?

The placement of a high-risk AI system on the market requires, in addition to compliance with the mandatory features outlined in Chapter II, that an EU declaration of conformity (EU DoC) be drawn up and the CE marking be affixed to the system.

- **EU DoC** – the EU DoC is drawn up following a successful conformity assessment (see question 19) to confirm that the AI system complies with the AI Act. It must be kept at the disposal of the national competent authorities for 10 years after the AI system has been placed on the market or put into service.
  The EU DoC is also required under product safety laws (such as the Machinery Regulation, but not the General Product Safety Regulation), as well as under resilience and sustainability laws.
  - ↪ For high-risk AI systems covered by the AI Act and product safety, resilience and sustainability laws that also require an EU DoC, a single EU DoC must be drawn up containing all the information required under the AI Act and other legal acts.
- **CE marking** – the CE marking must be affixed to the high-risk AI system, thus indicating that the system is in conformity with the AI Act.
  The CE marking is also required under product safety laws (such as the Machinery Regulation, but not the General Product Safety Regulation (GPSR)), as well as under resilience and sustainability laws (as this requirement may be mandated by the delegated acts of the Ecodesign for Sustainable Products Regulation (ESPR)) – see questions 2, 3 and 4.
  - ↪ For high-risk AI systems covered by the AI Act and other laws that also require CE marking, the CE marking must indicate that the high-risk AI system also fulfils the requirements of those other laws.

Additionally, **traceability information** is required:

- The AI system name/trade name and type, and any additional unambiguous reference allowing the traceability of the AI system, is one of the elements required in the EU DoC and in the information to be submitted for the system's registration in the EU database (see 🔍 Registration of AI systems and products). Traceability is also established under safety, resilience and sustainability laws – notably, products must bear a type, batch or serial number for identification. Specifically, when it comes to sustainability laws, the ESPR requires products to have a Digital Product Passport (DPP), which should be linked to a unique product identifier through a data carrier and, where appropriate, to a unique operator identifier and a unique facility identifier.
- The name, registered trade name or registered trademark, and contacts, must also be indicated (for high-risk AI systems, they should appear therein or, where that is not possible, on their packaging or accompanying documentation).

To whom the obligations apply

These obligations apply primarily to the **AI provider** under the AI Act and the **product manufacturer** under product safety, resilience and sustainability legislation. For an explanation of who these actors are, see question 1.

AI Checklist — Marking, traceability and registration

Goal:

Assess:

- Whether the required EU DoC and CE marking have been drawn up and affixed – applicable to high-risk AI systems under the AI Act, and to all AI systems under product safety, resilience and sustainability laws.
- Whether registration has been carried out and traceability information (on the product and on the AI provider/manufacturer) is made available – applicable to high-risk AI systems under the AI Act, and to all AI systems under product safety, resilience and sustainability laws.

Questions:

- Is the AI system subject to an EU DoC and CE marking?
- Are there other applicable laws requiring an EU DoC and CE marking?
- Is a single EU DoC and CE marking envisaged?
- Is the AI system accompanied by the required traceability information?
- Is the AI system registered in the EU database?
- Which risks are anticipated (e.g., incomplete or inconsistent EU DoCs, misuse or omission of CE marking, poor implementation of CE marking, weak traceability of the AI system, inaccurate or outdated information in the EU database)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Marketing; Compliance.

Registration of AI systems and products

The AI Act requires providers of certain high-risk AI systems to register before placing the system on the market. Providers (or their authorised representatives) must enter themselves and their systems in the EU high-risk AI database set up by the Commission.

There are specificities: high-risk systems for biometrics, law enforcement, and migration, asylum and border control are registered in the secure, non-public section of the EU database, while high-risk systems that are safety components of critical infrastructure are registered only at the national level.

Other EU laws also foresee product registration tools: for instance, the GPSR addresses the registration of products in the traceability system the European Commission may set up for products posing serious consumer risks, and the ESPR addresses the set-up by the Commission of the Digital Product Passport (DPP) registry. These systems are still being designed and are not yet fully operational.

Risk mapping:

Marking, traceability and registration risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What documentation and information should accompany the AI system?

AI systems must be accompanied by specific documentation and information, notably technical documentation and instructions for use. These requirements stem from the AI Act for high-risk systems, though product safety, resilience and sustainability legislation (see questions 2, 3 and 4) also address this requirement.

- **Technical documentation** – technical documentation must be drawn up to demonstrate that the high-risk AI system complies with the legal requirements and allow authorities to assess compliance. It must be kept at the disposal of the national competent authorities for 10 years after the AI system has been placed on the market or put into service.
  - ↳ The technical documentation must contain a general description of the AI system, including intended purpose; information on design and development, including architecture, data governance, human-oversight measures and pre-determined changes; a description of how the system functions and foreseeable risks; and a description of the risk management system and post-market arrangements, among other elements.

  Technical documentation is also required under product safety, resilience and sustainability laws.
  - ↳ For high-risk AI systems covered by the AI Act and certain product safety laws, a single set of technical documentation must be drawn up containing all the information required under the AI Act and other legal acts.
- **Instructions for use** – the AI system must be accompanied by instructions for use that include concise, complete, correct and clear information that is relevant, accessible and comprehensible to deployers.
  - ↳ The instructions for use must at least identify the provider and contact details; explain the system's capabilities, limits and any pre-determined changes; describe human oversight and output interpretation measures; specify required computational/hardware resources, expected lifetime, maintenance and update needs; and, where relevant, describe how deployers can collect, store and interpret logs.

  Instructions for use are also generally required under product safety, resilience and sustainability laws.
  - ↳ When a product includes an AI system that is subject both to the AI Act and to other specific product safety laws, the information and documentation required under the AI Act may be integrated into the product documentation prepared under those product safety laws.
- ↳ Under the new updated version of the AI Act, where AI is a safety component of certain regulated products or is itself such a product subject to third-party conformity assessment for health and safety risks, specific AI Act's obligations may be limited by delegated acts of the Commission. This may benefit products such as medical devices, radio equipment or toys. Machinery is excluded from the full high-risk AI regime, with the Machinery Regulation supplemented by delegated acts to incorporate most AI Act health and safety requirements.

To whom the obligations apply

These obligations apply primarily to the **AI provider** under the AI Act and the **product manufacturer** under product safety, resilience and sustainability legislation. For an explanation of who these actors are, see question 1.

AI Checklist — Documentation

Goal:

Assess whether the required technical documentation and instructions for use have been drafted and accompany the AI system – applicable to high-risk AI systems under the AI Act, and to all AI systems under product safety, resilience and sustainability laws.

Note: other documentation may be required under applicable law (e.g., documentation concerning the quality management system – see question 17).

Questions:

- Are technical documentation and instructions for use envisaged to be prepared before placing the AI system on the market?
- Will the technical documentation and instructions for use be combined where possible?
- Does the technical documentation clearly show how the AI system complies with the legal requirements?
- Are the instructions for use accessible and comprehensible to deployers?
- Which risks are anticipated (e.g., non-compliant documentation/instructions, outdated documentation/instructions, poor readability or comprehensibility, instructions for use not tailored to deployers' capabilities, misalignment between technical documentation, instructions and actual behaviour)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Marketing; Compliance.

Cooperation with competent authorities and access to source code

Under the AI Act, providers of high-risk AI systems must cooperate with competent authorities and supply all information and documentation needed to show compliance (including, for example, training, validation and testing data).

Upon a reasoned request, they must also grant access to the system's source code where this is necessary to assess conformity and where analysis of the data and documentation provided has been exhausted or is insufficient.

This approach is in line with other EU regulatory frameworks, notably product safety, resilience and sustainability, which likewise require cooperation with authorities. In addition, the Machinery Regulation also allows national authorities, on a reasoned request, to obtain source code or programming logic from manufacturers when necessary to verify compliance with essential health and safety requirements.

Risk mapping:

Documentation risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What main aspects should be included in a contract for use of an AI system?

Contracts for using AI systems should generally cover the following:

- **The AI system provided and its features** – precise description of the AI system and its intended purpose as declared by the AI provider, plus related services. See also question 27 on documentation.
- **The conditions for use** – clear rules on permitted purposes and fields of use, types of use (including modification), internal and external uses, commercial exploitation, access controls, territory, duration and post-termination consequences (e.g., return or deletion).
  - ↪ Certain modifications by the deployer may lead to it being considered the AI provider, so roles under the AI Act should be clearly allocated.
- **Allocation of duties** – determination of each parties' obligations under applicable law, notably the AI Act, including regarding human oversight, logging, incident reporting and post-market monitoring.
  - ↪ The contract should provide for cooperation in regulatory audits, data breaches and cybersecurity incidents. It should also deal with obligations of the AI provider that affect the deployer (e.g., retraining or removing data in ways that change operation or outputs) and require the provider to supply the information and assistance needed for the deployer to meet its obligations.
- **Warranty period, maintenance and SLAs** – a warranty period may be established during which the provider corrects nonconformities at no extra cost. Maintenance services (preventive, corrective and updates/upgrades – see also question 3) may be foreseen.
  - ↪ SLAs may be established. Breach of the agreed SLAs may lead to penalties (whether liquidated damages or more punitive mechanisms) and, if useful, rewards for exceeding targets may be foreseen.
- **Warranties** – provision of warranties relating to the system and outputs (e.g., with respect to their quality, accuracy, availability, reliability, suitability), or their exclusion where they are provided "as is"; commitments that the system is lawfully provided and does not infringe third-party rights.
- **Other key provisions** – provisions on reports and audits, changes, price and payment, intellectual property, security, confidentiality, personal data, liability and indemnification, export controls, term and termination, assignment, representations and warranties, governing law and dispute resolution.
  - ↪ Provisions on rights to input data and outputs are particularly important (especially if the provider wishes to use them for further training or analytics) and must take into account the applicable rules (e.g., business users' access rights under the Digital Markets Act (DMA), users' data access rights under the Data Act, and General Data Protection Regulation (GDPR) constraints on personal data).
  - ↪ Termination provisions must reflect any mandatory rules, such as the Data Act's requirements on switching for data processing services (such as SaaS).

To whom is this relevant

The licensor and licensee should ensure that the contract contains all provisions needed to enable the lawful use of the AI system. In principle, under the AI Act, these roles correspond to the **AI provider** and the **AI deployer** (see questions 1 and 29 for an explanation of who these actors are).

An organisation can nonetheless become a **provider** simply by branding or, in some cases, modifying a system, so the legal provider may differ from the commercialising entity. In turn, a deployer under the AI Act is a user acting in a professional context, meaning purely personal, non-professional users are not deployers even if they are licensees.

AI Checklist — Contracts for AI use

Goal:

Assess:

- Whether contracts for AI use allow its use as intended and duly protect the parties – relevant for all AI systems.
- Whether the contracts comply with applicable obligations – applicable to AI systems depending on the law at stake.

Questions:

- Who supplies the AI system, who receives it, and is the recipient a consumer?
- What contractual conditions govern the use of the AI system?
- What information or assistance is provided to facilitate the use of the AI system?
- What information and assistance is each party required to provide, and is this sufficient to enable legal compliance?
- How does the contract govern input and output data, and allocate related rights?
- Which risks are anticipated (e.g., lack of clarity in authorised and prohibited uses, granting of rights to deployer leading to it being considered an AI provider, unclear allocation of AI Act obligations and resulting liability, insufficient contractual support for each party's legal obligations, unbalanced liability clauses vis-à-vis the risk)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Product & Engineering; Marketing; Legal; Data including Data Protection Officer (see question 18).

The provision of AI systems to consumers

Supplying AI systems to consumers triggers specific EU consumer law rules, which differ depending on whether AI is provided as a digital service/content or as a good with digital elements. The applicable framework includes, among others, the Consumer Rights, Unfair Commercial Practices, Unfair Contract Terms, Digital Content and Sale of Goods Directives.

Key obligations concern pre-contractual information (e.g., functionality, compatibility, interoperability), conformity with the contract and reasonable consumer expectations, provision of updates to maintain conformity, conditions for modifying the system, fair contract terms, and consumer remedies where there is a lack of conformity. Objective conformity requirements stem from the AI Act and from product safety, resilience and sustainability rules (see questions 2, 3 and 4), so breaches can trigger consumer remedies.

Although the AI Act does not establish stand-alone consumer protection rules, its prohibitions and transparency duties protecting natural persons (see questions 1, 8 and 9) complement horizontal consumer law. Traders must comply with both sets of obligations.

Risk mapping:

Contract risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

IV. Deploying AI systems

What it is

Deploying AI systems refers to the use of AI systems. It covers a set of activities required for such use, including AI acquisition/licence (as well as services acquisition, as necessary), collection and processing of data (input data) to produce an output, and the generation of the output. AI decommissioning can also be considered a step of the deployment stage, under which data archiving or deletion is carried out and the decommissioning of physical items of the AI system (waste management processes) may be performed.

The main legal touchpoint is the EU AI Act, complemented by cybersecurity and resilience laws, sustainability laws (especially on energy consumption and efficiency), and data (personal data and non-personal data), intellectual property and competition laws. Other laws, such as on consumer protection and digital services, are also relevant hereto.

In practice, this stage is where an AI system is configured for a specific purpose or case, input data is fed for AI operation and AI outputs are relied upon to inform or make decisions. In this stage, specific requirements for AI deployment must be met, including when it comes to fundamental rights, human oversight, transparency and governance processes, and covering also cybersecurity, resilience, sustainability and data requirements.

Relevant stakeholders

- AI providers
- Others, such as:
  - AI service suppliers
  - Data providers and data analytics providers
  - Infrastructure providers, such as cloud and computing
  - Organisations and individuals (including consumers) to whom AI-based decisions relate or interacting with the AI system

Key Topics

The main topics relevant for the AI deployment stage are:

- Deployment of the AI system – use conditions
- Data to deploy the AI system
- Governance
- Compliance and effects of non-compliance
- IP protection of the AI output
- Particularities for organisations, sectors and consumers
- Algorithmic collusion

Summary of main challenges and how to address them

Challenges	How to address the challenges – summary
Deploying AI systems so that they meet legal requirements and ethical expectations Ensuring the deployment of AI systems – including but not limited to high-risk ones – actually meets legal requirements on fundamental rights, cybersecurity and resilience, sustainability, privacy, human oversight, auditability and transparency, rather than treating them as abstract principles, whilst also ensuring consumer protection and competition-compliant deployment	• Assess risk level of the AI system and whether it falls under "high-risk", "prohibited" or other regulated categories • Map all applicable legal frameworks beyond the AI Act (e.g., cybersecurity and resilience laws; sustainability laws; data protection laws; consumer protection laws; competition law) • Perform structured risk and impact assessments (including, where relevant, on fundamental rights and data protection) and define concrete mitigation and elimination measures • Design an integrated compliance plan so that overlapping or complementary obligations are met coherently rather than in silos For more details, see DEPLOYMENT OF THE AI SYSTEM – USE CONDITIONS; PARTICULARITIES FOR ORGANISATIONS, SECTORS AND CONSUMERS; and ALGORITHMIC COLLUSION.
Using data that is lawful, high-quality and competition-compliant Using input data that is lawfully obtained and meets the legally required features (is relevant and representative), while navigating privacy, intellectual property and competition constraints (such as restrictive data licences or data-sharing arrangements by dominant undertakings)	• Classify the data you rely on (e.g., personal/non-personal, proprietary/open) and identify lawful bases for use • Assess data sources for lawfulness and quality, and implement processes to meet required features • Ensure contracts and licences explicitly permit the in-tended AI uses and reflect any existing mandatory requirements (e.g., on B2B data sharing or on licences for copyrighted works) • Review data-sharing and licensing arrangements (in-cluding pricing and exclusivity) for potential competi-tion law issues For more details, see DATA TO DEPLOY THE AI SYSTEM.
Building effective AI governance without duplication Setting up governance that satisfies AI-specific obligations and other existing legal frameworks (e.g., Data Protection Impact Assessments (DPIAs), cybersecurity measures), without creating fragmented or duplicative processes	• Design and implement AI governance processes aligned with AI Act requirements for the relevant systems (high-risk) • Map all existing governance processes (e.g., AI, cybersecurity, privacy, environmental) and ensure coordination among them as relevant • Define clear accountability structures (roles, committees, escalation paths) and embed a culture of transparency, responsibility and ethical behaviour For more details, see GOVERNANCE.

Challenges	How to address the challenges – summary
Ensuring compliance and dealing with risk Not only being compliant, but being able to ensure compliance through training and literacy, and having credible mechanisms to detect and report incidents, threats and non-compliance	• Implement training and literacy initiatives so that teams understand AI risks, obligations and internal procedures • Put in place incident and threat reporting, non-conformity and risk suspension and information, and liability management processes • Identify and leverage relevant standards, certification schemes and guidance to operationalise requirements and streamline evidence gathering For more details, see COMPLIANCE AND EFFECTS OF NON-COMPLIANCE.
Protecting AI-related outputs Securing appropriate IP protection for AI outputs, in a context where human contribution and protectability criteria can be complex, and where alternatives (trade secrets, technical measures) may be need-ed	• Identify the nature of AI-related outputs and classify them accordingly • Design output development so that protectable subject-matter is created (e.g., clear human contribution for copyright and patent purposes) • Develop and implement an effective IP management process to ensure AI output is protected and, as relevant, can be exploited as intended • Implement IP management workflows (disclosures, prior checks, filings, registrations, confidentiality controls) as relevant For more details, see IP PROTECTION OF THE AI OUTPUT.

How to achieve Responsible AI	
	Responsible AI deployment is achieved when AI systems are used in ways that are lawful, ethical and aligned with their risk profile and real-world impacts. To move in that direction, organisations should ensure that deployment decisions are grounded in a clear risk classification and mapping of all applicable legal frameworks; that lawful, high-quality and competition-compliant data is used; that AI governance, accountability structures and AI-literacy measures support effective human oversight; that systems are monitored in practice, with logs, incident reporting and remediation processes to deal with risks and non-compliance; and that AI-related outputs are protected and managed through appropriate IP, confidentiality and technical safeguards so they can be used and innovated on responsibly over time.

Are there any legal requirements regarding the deployment of an AI system?

Deployment of an AI system is subject to legal requirements that vary depending on its functions, whether it processes personal data, whether it affects natural persons, and the type of outputs it generates.

In broad terms, depending on the AI system, its deployment should be:

- **Respectful of fundamental rights** – certain deployments of high-risk AI systems (see 🔍 High-Risk Systems) require a fundamental rights impact assessment (FRIA) to identify risks towards affected individuals or groups and to define mitigation measures.
- **Resilient** – certain deployments of AI systems must include security and risk management measures (under horizontal resilience rules), and high-risk AI must be deployed in line with its instructions for use (see question 27), including robustness and cybersecurity information.
- **Environmentally sustainable** – incidents or malfunctioning of the high-risk AI system leading to serious harm to the environment must be reported, and other EU environmental and energy efficiency rules may apply.
- **Privacy preserving** – if the AI system processes personal data (see 🔍 The Different Types of Personal Data), data protection requirements should be complied with.
- **Overseen by humans** – high-risk AI may only be deployed where effective human oversight mechanisms are in place.
- **Auditable** – logs automatically generated by high-risk AI must be kept.
- **Transparent** – certain high-risk systems that make or support decisions about natural persons require that those persons be informed of the use of such systems and, in certain cases, provided with an explanation of decisions taken based on the system's outputs. Additional information duties apply to AI that generates/manipulates content, performs emotion recognition, or conducts biometric categorisation.

Further transparency obligations apply under otherlegislation (including bans on dark patterns) (see 🔍 Use of AI in online platforms).

For more details on each of these requirements, see questions 30 to 37.

The AI Act excludes specific contexts, such as exclusive military/defence/national security uses and purely personal, non-professional uses by natural persons.

In addition to the above requirements, the deployment of certain AI systems is **prohibited** under the AI Act. For more details, see 🔍 Prohibited Practices.

To whom the obligations apply

AI Act obligations apply to **AI deployers**, i.e., those using the AI system under their authority (other than for personal non-professional purposes) when established or located in the EU, or when the AI output is used in the EU.

Resilience obligations apply to AI deployers/users that are **essential**, **important and/or critical entities** (see question 31 for an explanation of these actors), whilst environmental obligations (namely on energy efficiency) apply where certain energy consuming thresholds are met (see question 32 on this). Data protection duties fall upon the data controller and **processor** (see question 33 for an explanation of these actors) and dark patterns rules apply to **providers of online platforms** (see question 36 for an explanation of who these actors are).

AI Deployment Checklist

Goal:

Assess the reason for the deployment of the AI system and its objective/intended effects, and:

- Whether the deployed AI system is **prohibited** or **high-risk**.
- Whether the deployment of the AI system is subject to **cybersecurity and resilience obligations**.
- Whether the deployment of the AI system has **environmental impacts**.
- Whether the AI system **processes personal data** during deployment.
- Whether the deployment of the AI system is subject to the **AI Act, resilience laws, sustainability/efficiency laws**, the **General Data Protection Regulation** (GDPR) and the **Digital Services Act** (DSA), bearing in mind that some systems/deployments (e.g., systems not made available in the EU or covered by specific exclusions) may fall outside these regimes even if they are high-risk, process personal data, or meet other criteria.

Questions:

- What are the main objectives for the use of the AI system?
- What are the key features of the system (including autonomy, adaptiveness, predictability)?
- In which context will it be used (simple/predictable or complex/uncertain)?
- Which are the organisations deploying the AI system and where are they established?
- Does it enable the processing of personal data?
- In which countries will it be deployed?

Expertise required:

Technology & Systems; Procurement; Quality; Legal; Management.

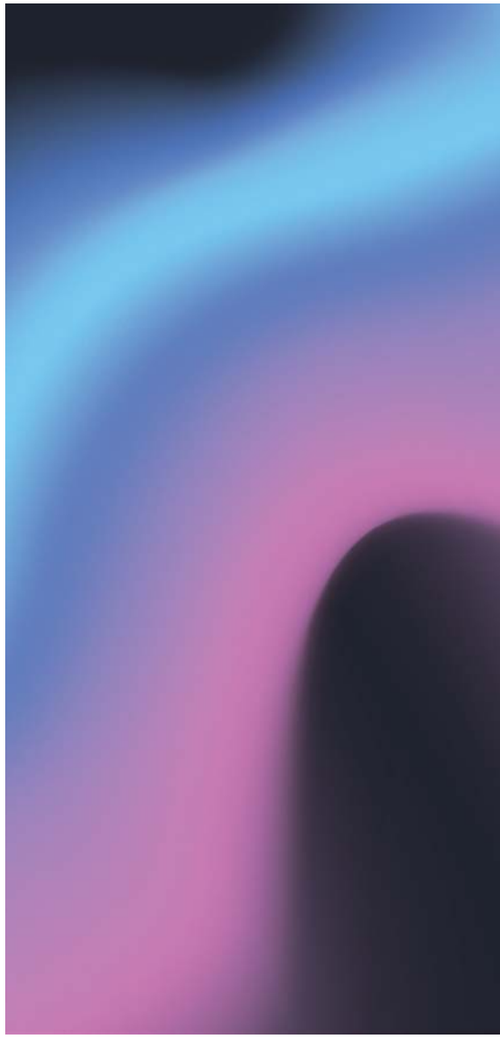

Risk mapping:

Risk of exposure to EU laws	Likelihood	Severity (of non-compliance)	Risk score
List of laws			

See the explanation on the Checklist and Risk Mapping above.

How can a fundamental rights impact assessment – FRIA – be performed?

Under the **AI Act**, a **fundamental rights impact assessment** (FRIA) is a systematic evaluation of the risks posed by the use of high-risk AI systems to individuals' fundamental rights.

This impact assessment must be performed prior to deploying the AI system and applies to its first use.

The assessment must include:

- **Description of processes** – outlining the deployer's processes in which the high-risk AI system will be used, ensuring alignment with its intended purpose.
- **Usage timeline and frequency** – providing a description on the duration and frequency of the high-risk AI system's intended use.
- **Persons and groups affected** – identifying the categories of individuals and groups likely to be affected by the AI system in the specific context.
- **Risk of harm** – describing the specific risks of harm that could affect the identified categories of individuals or groups, considering the information provided by the provider.
- **Human oversight implementation** – describing the implementation of human oversight measures in accordance with the instructions for use of the AI system; and
- **Risk mitigation measures** – i.e., the measures to be taken if the identified risks materialise, including internal governance arrangements and complaint mechanisms.

↳ Several methodologies support fundamental rights assessments, including the Council of Europe's 2024 non binding HUDERIA methodology for evaluating AI impacts on human rights, democracy, and the rule of law.

While performing the assessment, relevant information may be considered, notably, in similar cases, a previously drafted FRIA or existing impact assessments carried out by the provider.

If, during the use of the high-risk AI system, any of the elements listed has changed or is no longer up to date, the necessary steps shall be taken to update the information.

The FRIA may include or cross-reference the relevant parts of the data protection impact assessment (DPIA), when drafted for the same AI system, thus avoiding duplications.

After performing the assessment, the relevant market surveillance authority should be notified.

The AI Office is set to publish a template for a questionnaire, including through an automated tool, to assist in drafting FRIAs.

In parallel, fundamental rights protections also flow from other EU instruments, such as the General Data Protection Regulation (GDPR), the Digital Services Act (DSA) and non-discrimination laws (on racial equality, employment and occupation, gender equality in access to goods and services, among others). Assessing fundamental rights impacts, even in the deployment of non-high-risk systems, is therefore a recommended good practice.

To whom the obligations apply

The FRIA must be performed only by certain of high-risk AI systems (see question 29 for an explanation of AI deployers) of high-risk AI systems: **public bodies** using systems for biometric identification, education and vocational training, employment and worker management, access to essential services, law enforcement, migration, asylum and border control, and justice and democratic processes; **private entities providing public services** using the same systems; any **deployer (public or private) using AI systems** to evaluate the creditworthiness of natural persons or establish their credit score (except for detecting financial fraud) or to set life or health insurance risk and pricing.

AI Deployment Checklist — FRIA

Goal:

Assess whether the deployment of the AI system is subject to a FRIA – applicable only to the deployment of certain high-risk AI systems under the AI Act.

Questions:

- What are the fundamental rights potentially impacted by the deployment of the AI system?
- Which people or groups of people are likely to be affected?
- Is the organisation required to perform a FRIA?
- Is the performance of the FRIA envisaged before deploying the AI system?
- Are there procedures to ensure information updating?
- Are there procedures to coordinate the FRIA with a DPIA, as applicable?
- What risks are anticipated (e.g., incorrectly assuming there is no need for a FRIA, stating only generic risks in the FRIA, no alignment between the FRIA and DPIA, no updates when changes occur, weak documentation and lack of transparency about the FRIA)?
- Which mitigation measures are envisaged?

Expertise required:

Technology & Systems; Legal; Compliance; Ethics; Data including Data Protection Officer (see question 42).

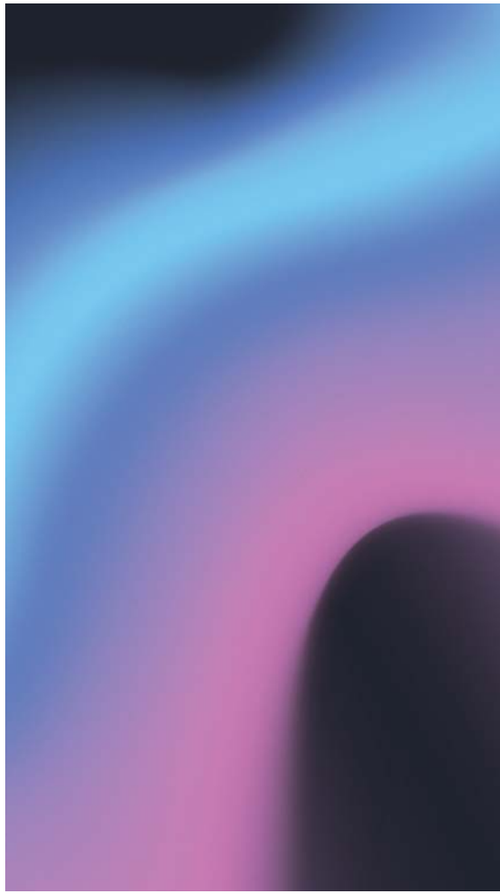

Risk mapping:

Fundamental rights affected by AI deployment	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

Safety risks of AI deployment	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What resilience obligations apply to deployers?

Resilience refers to the ability of an organisation or system to prevent, protect against, respond to, withstand, mitigate, absorb, accommodate, and recover from an incident.

The AI Act does not set explicit resilience obligations for deployment, but high-risk systems must be deployed in line with their instructions for use (see question 27), including robustness and cybersecurity information.

The main resilience requirements arise under cybersecurity and critical infrastructure legislation (the **Network and Information Systems** (NIS 2) and **Critical Entities Resilience** (CER) Directives) and the **General Data Protection Regulation** (GDPR), which require compliance with the following obligations:

- **Implementation of appropriate and proportionate technical, operational and organisational measures (and, under the CER Directive, risk assessments and liaison officers) to manage security risks to network and information systems, and to critical entities/infrastructure.**
  AI will typically form part of such "network or information systems" (under NIS 2) and can constitute or support "critical infrastructure" (under the CER Directive).
  - ↳ Examples of measures include: risk analysis/assessment and security policies, incident handling, business continuity, supply-chain security, secure development and maintenance, cyber hygiene practices and cybersecurity training, human resources security, access controls, asset management, cryptography, multi-factor authentication, physical protection of premises and critical infrastructure, and resilience planning.
- **Implementation of appropriate technical and organisational security measures for the processing of personal data (under the GDPR).**
  - ↳ Examples of measures include: pseudonymisation and encryption, ensuring ongoing confidentiality, integrity, availability and resilience, the ability to restore data after incidents, and regular testing and evaluation of safeguards.

To whom the obligations apply

The **NIS 2 Directive** applies to medium and large enterprises (and, in some cases, smaller but nationally critical ones) in specified sectors, when they provide their services or carry out their activities in the EU. It distinguishes between **essential entities** (subject to full ex ante and ex post supervision) and **important entities** (subject mainly to ex post supervision), based on factors such as size and national designation.

The **CER Directive** applies to **critical entities** designated by Member States because they provide essential services in specified sectors.

- ↳ Sectors covered by both Directives include energy, transport, health, water, space, public administration, whilst banking, financial market infrastructure and digital infrastructure entities must be designated but are not directly bound by CER. The NIS Directive also covers other sectors, such as digital platforms and research, ICT service management, among others.
- ↳ Member States may go beyond these minimum-harmonisation Directives, including by extending their scope to additional entities.

GDPR obligations apply to **controllers and processors**. For an explanation of who these actors are, see question 33.

AI Deployment Checklist — Resilience

Goal:

Assess compliance with the cybersecurity and resilience requirements of the applicable laws, specifically:

- Whether the AI system is deployed in accordance with its instructions for use under the AI Act – applicable to the deployment of high-risk AI systems.
- Whether the AI system is deployed in line with applicable essential cybersecurity requirements under the NIS 2 and CER Directives, and the GDPR – applicable to the deployment of all AI systems by essential, important and critical entities, and by data controllers and processors, respectively.

Questions:

- Is the organisation an essential, important and/or critical entity?
- Does the AI system support or automate functions relating to the essential, important or critical services of the organisation?
- Is the AI system used to process personal data?
- Which security risks are anticipated for the AI system deployment (e.g., treating AI systems outside the scope of NIS 2/CER/GDPR, generic security measures that ignore AI-specific features and threats)?
- Which measures are envisaged to manage the risks?

Expertise required:

Technology & Systems; Risk & Security; Legal; Compliance; Data including Data Protection Officer (see question 42).

Risk mapping:

Cyber risks of AI deployment	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Use of cyber certified AI

Using a cybersecurity-certified AI system means deploying AI that has been formally assessed and certified as meeting predefined security requirements in line with cybersecurity schemes developed under the EU Cybersecurity Act. See **The European Cybersecurity Certification Scheme**.

Under the NIS 2 Directive, Member States – and in some cases the Commission – may require essential and important entities to use, or obtain, EU-certified ICT products, services, processes or certifications. In the absence of such certification schemes, Member States should promote the use of relevant European and international standards.

A 2026 proposal for a revised "Cybersecurity Act 2" would broaden this framework to cover managed security services and an organisation's overall cybersecurity posture, introduce new ICT supply chain risk controls, and more clearly position EU cyber certificates as a key means to evidence compliance and obtain presumptions of conformity under NIS 2. A parallel NIS 2 amendment would clarify that cyber-posture certificates can demonstrate compliance and can be made mandatory for essential and important entities.

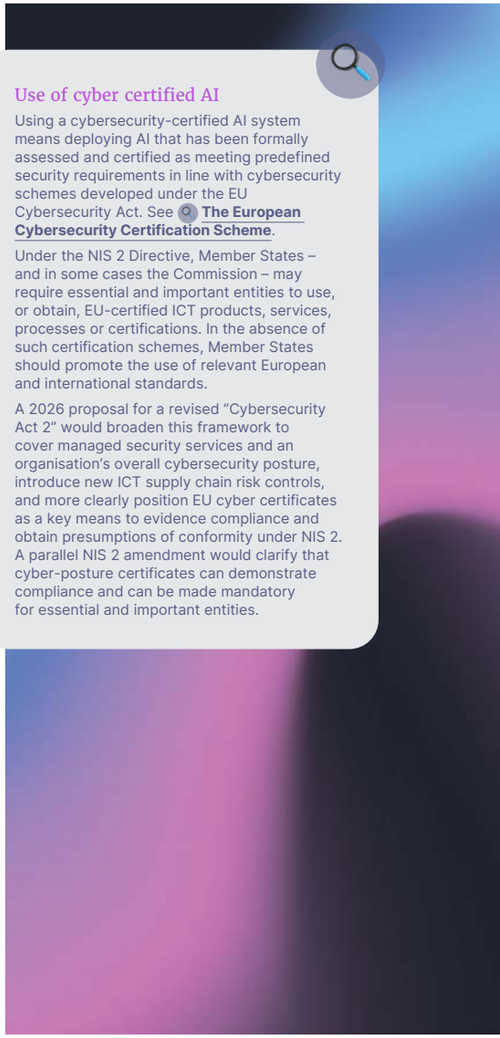

Are there any environmental obligations that apply to deployers?

Environmental sustainability refers to requirements established in law to ensure the minimisation of environmental impacts throughout a business' activity, from the consumption of natural resources (including energy and water), to emissions, pollution and waste. The AI Act does not establish explicit requirements concerning the environmentally sustainable deployment of AI systems, though it requires the reporting of incidents or malfunctioning of the high-risk AI system that leads to serious harm to the environment.

↳ Despite the lack of specific environmental duties, the AI Act recognises environmental protection as relevant to AI use, and, since high-risk AI must be assessed against fundamental rights (see questions 29 and 30), environmental impacts could be considered when evaluating effects on rights such as health and safety.

↳ The AI Act explicitly encourages codes of conduct covering the assessment and minimisation of the impact of AI systems on environmental sustainability, including as regards energy-efficient programming and techniques for the efficient design, training and use of AI.

Instead, sustainability obligations arise primarily from other EU legislation, notably legislation on energy consumption – the **Energy Efficiency Directive** (EED). Energy consumption is a central topic in AI, as AI use can be energy intensive. In broad terms, the EED requires compliance with the following:

- Implementation of an **energy management system**, or the performance of energy audits, and drafting of an action plan based on the recommendations arising from the audits.

  ↳ Enterprises with certified energy performance contracts or environmental management systems that already include compliant energy management or audits are exempt.

  Since the AI Act sets no specific environmental management or reporting duties (neither caps nor incentives for energy consumption), AI-related energy use is only captured in the general monitoring and reporting under the EED.

Other environmental regimes could also be assessed, such as: EU water law (e.g., the Water Framework Directive and Urban Wastewater Treatment Directive) setting the framework for national rules on abstraction and discharge; the emissions framework (e.g., the Industrial Emissions Directive), requiring, among other things, environmental management systems for certain industrial activities; and waste and recycling laws (e.g., the Waste Framework Directive and the electrical and electronic equipment – WEEE – Directive) targeting waste prevention, repair and reuse. Again, the obligations arising from these legal frameworks apply to an organisation's underlying activities (e.g., water use/discharge, waste generation) rather than its use of AI.

↳ Because the AI Act sets no specific environmental rules on water use, emissions or waste (and the EU waste framework does not apply to pure software), there are currently no EU-level requirements on these aspects specifically applicable to AI deployment.

To whom the obligations apply

Environmental obligations apply to **AI deployers**/users only to the extent that they meet the specific conditions or thresholds established under the relevant legislation – and not because they are using AI systems (see question 29 for an explanation of the concept of deployer under the AI Act). Notably, the obligations apply:

- Where deployers' energy consumption exceeds the thresholds established under the EU energy framework – enterprises (within the territory of a Member State) that consume energy exceeding 85 TJ over the previous three years (for the obligation to implement an energy management system) and enterprises with an average annual consumption higher than 10 TJ of energy over the previous three years, which do not implement an energy management system (for the obligation to perform energy audits and an action plan);
- Where deployers use or discharge water under regulated conditions, carry out pollution-generating industrial activities (e.g., battery production), or generate waste covered by applicable waste laws.

AI Deployment Checklist — Sustainability

Goal:

Assess:

- Whether the AI system is deployed in compliance with fundamental rights (which may be impacted by environmental factors) – applicable to the deployment of high-risk AI systems.
- Whether the AI system is deployed in line with applicable environmental requirements, notably on energy – applicable to the deployment of all AI systems by organisations above certain consumption thresholds.

Questions:

- Does the organisation surpass the energy consumption thresholds established in law?
- What is the contribution of AI deployment to energy consumption and generally to the organisation's environmental footprint?
- Is there visibility on AI's energy and resources use?
- What sustainability reporting obligations is the organisation subject to?
- Which environmental risks are anticipated for the deployment of the AI system (e.g., oversized or inefficient AI deployment, lack of internal guidance for environmentally sustainable AI use)?
- Which measures are envisaged to manage the risks?

Expertise required:

Technology & Systems; Legal; Compliance.

Risk mapping:

Environmental risks of AI deployment	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Sustainability reporting

Sustainability reporting is the disclosure of non-financial information on an enterprise's environmental, social and governance (ESG) performance. Certain undertakings (notably large companies and some SMEs) must report sustainability information under the Accounting Directive (as amended by the Corporate Sustainability Reporting Directive (CSRD)), including how and to what extent their activities qualify as environmentally sustainable under the Taxonomy Regulation and its Delegated Acts, using the European Sustainability Reporting Standards (ESRS) as a reference.

This framework aims to improve transparency on ESG matters and progressively guides sustainability reporting. However, simply deploying AI is not itself a taxonomy-eligible activity and does not, by itself, show that a company's activities are environmentally sustainable, even if AI use may have beneficial environmental effects.

Separately, the Corporate Sustainability Due Diligence Directive (CSDDD) will phase in human rights and environmental due diligence duties for some companies, but again, the mere use of an AI system does not in itself trigger its application.

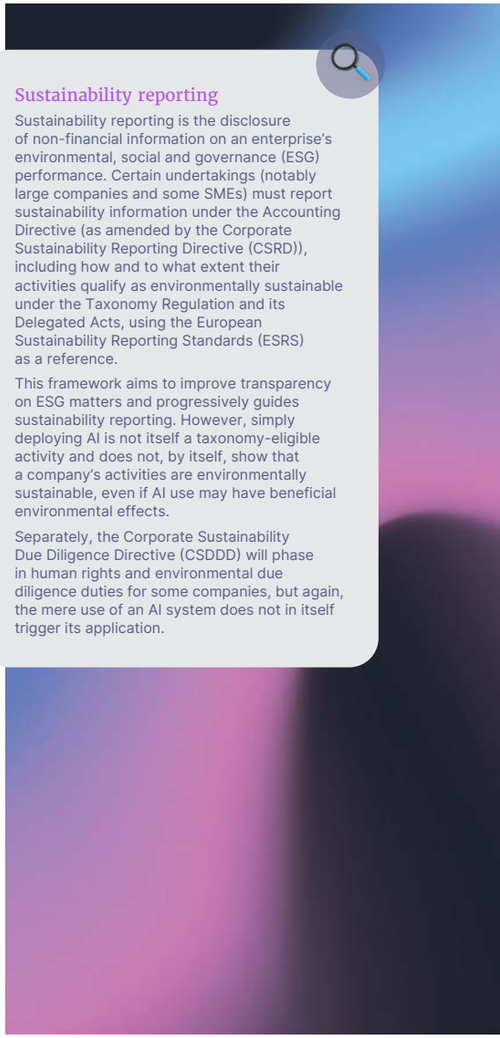

What measures should be implemented to ensure that the use of the AI system protects data, notably personal data?

Data protection generally refers to requirements established in law to ensure that data is protected from unlawful uses and is used only in accordance with the applicable legal requirements. These legal requirements mostly arise from the **General Data Protection Regulation** (GDPR) in relation to personal data (see question 11 and 🔍 The Different Types of Personal Data).

In broad terms, to the extent personal data is processed when deploying an AI system, then the following obligations must be met:

- Processing personal data in line with data protection principles:
  - **Lawfulness, fairness and transparency** – data must be processed lawfully, fairly and in a transparent manner in relation to the data subject.
  - **Purpose limitation** – data must be collected for specified, clear, and lawful purposes, and not re-used in ways that are incompatible with those purposes. Further processing for certain purposes, such as scientific research purposes, is in principle possible, provided that appropriate safeguards (e.g., technical and organisational measures, such as pseudonymisation) are adopted.
  - **Data minimisation** – only data that is adequate, relevant and limited to what is strictly necessary for the intended purpose should be collected and used.
  - **Data accuracy** – only accurate data must be processed, and the deletion or amendment of incorrect data must be ensured.
  - **Integrity and confidentiality** – appropriate technical and organisational measures must be adopted, such as encryption, pseudonymisation or access control (see also question 31).
  - **Storage limitation** – personal data must be kept only for as long as necessary to fulfil the purpose for which it was collected. After this period, the data should be deleted, anonymised, or otherwise securely destroyed, unless continued retention is justified by law.
- To the extent allowed, **configuring, parametrising or customising the AI system so that it remains privacy preserving.**
- Implementing **measures and procedures** to ensure legal compliance – see question 41 for more details. One such measure relates with ensuring the data subject's right not to be subject to solely automated decisions with legal or similarly significant effects, except when permitted under the GDPR – see also question 36.

To whom the obligations apply

The primary obligation to ensure that the use of the AI system protects data applies to the **data controller** — the one determining the purposes and means of processing — regardless of actual data access. In the deployment phase, the data controller is in principle the AI deployer/user (for an explanation of the concept of deployer under the AI Act, see question 29).

The **data processor** – that who processes personal data on behalf of the controller – is also subject to certain obligations, such as those on security of processing.

Under the GDPR, this covers controllers and processors established in the EU or, if outside, those processing data of individuals in the EU (in connection with offering goods or services), or monitoring them.

AI Deployment Checklist — Data protection

Goal:

Assess compliance with the GDPR requirements, specifically:

- Whether the AI system is deployed in compliance with the GDPR – applicable to the deployment of all AI systems that process personal data.

Questions:

- Is the AI system used to process personal data?
- Is the organisation the data controller or processor?
- Which data protection risks are anticipated for the deployment of the AI system (e.g., blurred controller/processor roles in AI use, weak implementation of GDPR principles, customisation of the AI system impacting data protection features)?
- Which measures are envisaged to mitigate or eliminate those risks?

Expertise required:

Technology & Systems; Legal; Compliance; Data including Data Protection Officer (see question 42).

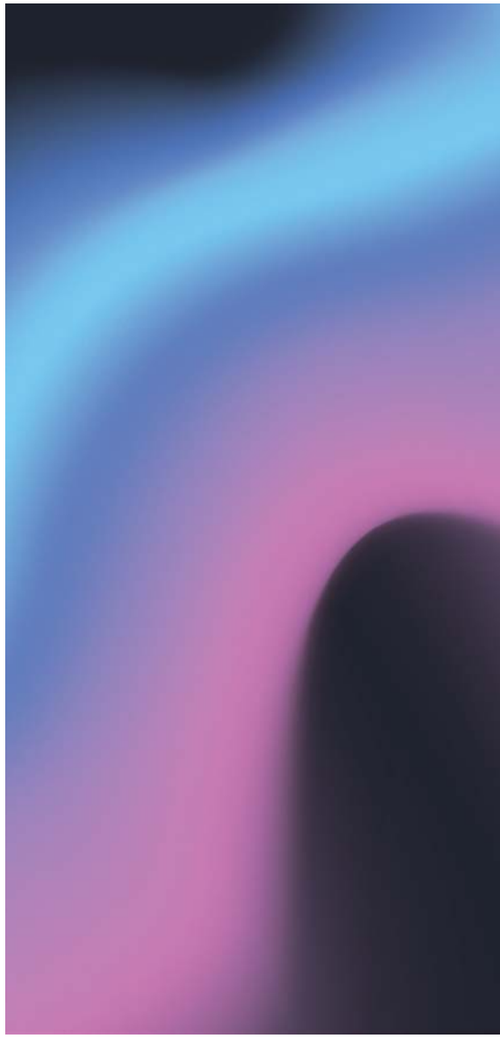

Risk mapping:

Data protection risks of AI deployment	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can human oversight be ensured when deploying an AI system?

Human oversight refers to the possibility of humans retaining control and intervening in the operation of AI systems. The goal is to prevent or minimise risks to health, safety and fundamental rights that may arise from the operation of an AI system. Oversight can take various forms, including direct monitoring of the system's functioning, the possibility of intervening or interrupting its operation, or measures ensuring that the system remains interpretable and that its limitations are understood by the human overseer.

The **AI Act** establishes obligations relating to human oversight in the deployment of high-risk AI systems. Notably:

- Human oversight must be assigned to competent and duly trained people.
- Human oversight must be delegated to people with the necessary authority to carry out this role.
- The above people must be provided with the necessary support for this function.

The deployer is free to organise its own resources and activities for the purpose of implementing human oversight measures.

An AI system may also be subject to human oversight requirements from other applicable laws.

- For instance, under the **General Data Protection Regulation** (GDPR), where automated individual decision-making is allowed (e.g., when necessary for a contract or based on explicit consent), suitable safeguards must be in place, including at least a right to human intervention. In this "human-out-of-the-loop" model, upon the data subject's request, a qualified person reviews the AI-supported decision and can change or override it. This differs from the AI Act's focus on assigning ongoing human oversight to natural persons ("human-in-the-loop").

To whom the obligations apply

The obligations described above apply primarily to **deployers of high-risk AI systems**. For an explanation of who the AI deployer is, see question 29.

The obligation arising from the GDPR applies to **data controllers**. For an explanation of who these actors are, see question 33.

AI Deployment Checklist — Human oversight

Goal:

Assess compliance with the applicable requirements, specifically:

- Whether AI deployment can be effectively overseen by natural persons during its use – applicable only to the deployment of high-risk AI systems under the AI Act.
- Whether AI deployment involves human intervention when processing personal data – applicable to all AI deployments leading to automated individual decision-making.

Questions:

- What decisions or types of decisions will resort to the AI system?
- Which features of the AI system allow human oversight (e.g., "stop" functions, "override" functions, "disregard" functions, escalation paths, alerting mechanisms so that human overseers can detect anomalies or unexpected performance in time to act)? Can they be configured and, if yes, how?
- Which types of oversight are envisaged (preventive, concurrent, *ex post*)?
- What human resources are allocated to human oversight (category, authority, function) and how are they trained?
- What human oversight risks are anticipated for AI deployment (e.g., oversight on paper only; no clear oversight models or measures per different uses; inadequate capacity for oversight and to interpret outputs; inadequate ability to intervene, override or stop the system; no structured escalation paths; over-reliance risks such as automation bias, reduced vigilance, deskilling, loss of accountability)?
- Which measures are envisaged to address the risks?

Expertise required:

Technology & Systems; Legal; Compliance; Ethics.

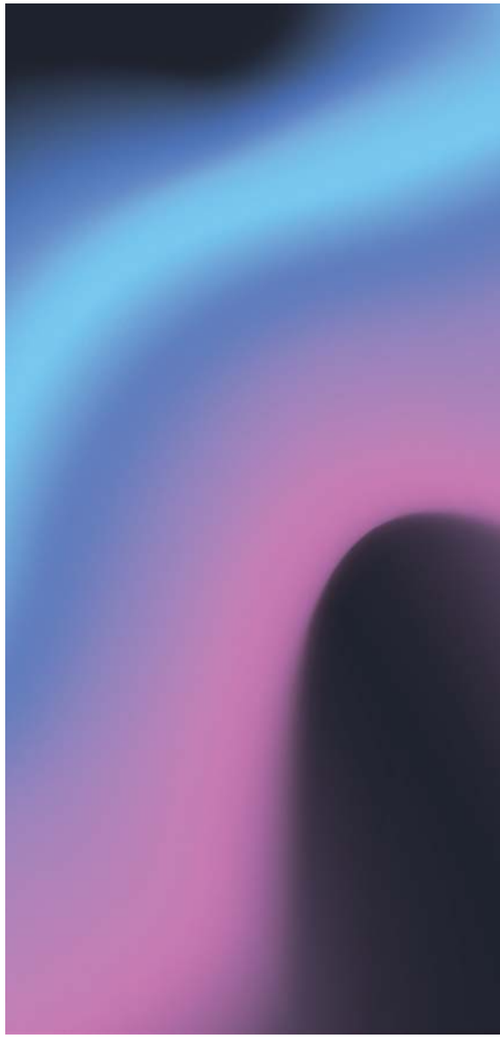

Risk mapping:

Oversight risks of AI deployment	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can the deployer keep the logs generated by the AI system?

Under the **AI Act**, logs automatically generated by the high-risk AI system that are under the deployer's control must be retained for a period appropriate to the intended purpose of the high-risk AI system – and for at least six months. The six-month retention period is, however, subject to exceptions where otherwise provided by Union or national law, particularly under Union law concerning the protection of personal data.

↪ It is important to note that the effective fulfilment of these obligations may depend on the provider's instructions for use, which must include, where relevant, a description of the mechanisms included within the high-risk AI system enabling deployers to properly collect, store and interpret logs (see question 27).

In this regard, it should be noted that:

- Personal data may be stored for no longer than is necessary for the purposes for which the personal data are processed;
- Personal data may be stored for longer periods provided that:
  - The data are processed solely for archiving purposes in the public interest, scientific or historical research purposes, or statistical purposes; and
  - The appropriate technical and organisational measures required by the General Data Protection Regulation (GDPR) are implemented to safeguard the rights and freedoms of data subjects.

↪ When the AI system processes personal data, a record of the processing activities should also be kept (see question 33).

To whom the obligations apply

The obligations described above apply to **deployers of high-risk AI systems**, with personal data requirements being applicable to the **data controller** – for an explanation of who these actors are, see questions 29 and 33.

AI Deployment Checklist — Logging

Goal:

Assess compliance with the applicable requirements, specifically:

- Whether the logs are kept as required by the AI Act – applicable only to the deployment of high-risk AI systems.
- Whether the logs with personal data are kept in line with the GDPR – applicable to all AI deployments.

Questions:

- Which logs are generated by the AI system?
- Do the logs integrate or reveal personal data?
- Are those logs under the organisation's control?
- For how long are the logs going to be kept? Where are they stored? And which access and security controls exist for the logs?
- Which logging risks are anticipated (e.g., insufficient logging; excessive logging creating risks such as GDPR risks; logging integrity and availability failures; logging interpretation, review and escalation difficulties)?

Expertise required:

Technology & Systems; Legal; Compliance;
Data including Data Protection Officer (see question 42).

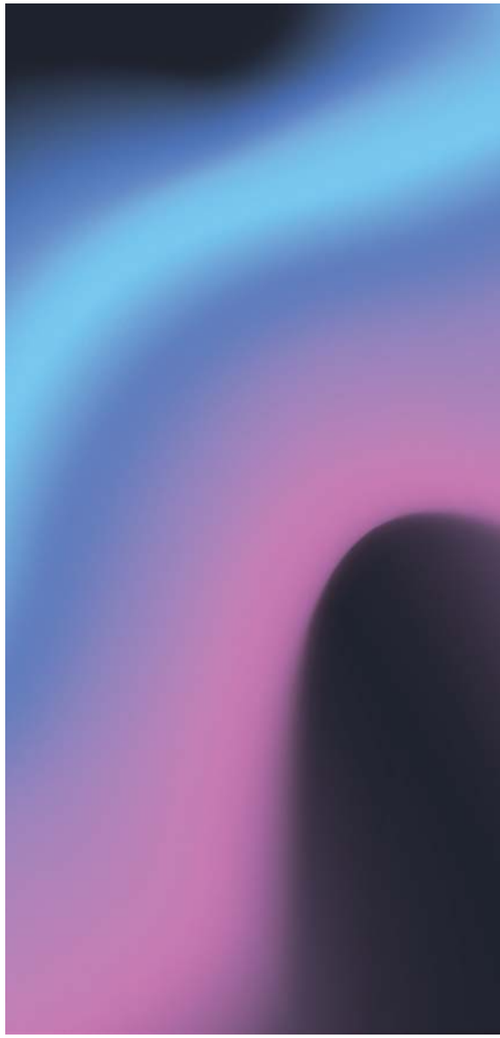

Risk mapping:

Logging risks of AI deployment	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What requirements must be met when the use of an AI system impacts individuals?

When the use of an AI system impacts individuals, specific requirements must be met in its deployment. These requirements arise mainly from the **AI Act**, but also from other EU instruments such as **consumer protection legislation** and the **General Data Protection Regulation** (GDPR).

In broad terms, the following obligations must be met:

- **Provide meaningful explanations** – explanations must be provided to persons affected by decisions (mainly) made on the basis of the output of certain high-risk AI systems, where the decision produces legal effects or similarly significantly affects that person in a way that they consider to have an adverse impact on their health, safety or fundamental rights. The explanations provided must be clear and meaningful and cover the system's role in the decision-making procedure and the main elements of the decision itself.

  ↳ The systems covered by this obligation comprise high-risk systems for biometric identification, education and vocational training, employment and worker management, access to essential services, law enforcement, migration, asylum and border control, and justice and democratic processes.

  This right of the individual applies only to the extent that it is not otherwise provided for under Union law (as in the GDPR).

- **Provide information on the use of AI** – under the AI Act, clear information must be provided to individuals in several situations:
  - When the above high-risk AI systems are used to make, or assist in making, decisions related to individuals, those individuals must be informed that they are subject to the use of said high-risk AI system.
  - When emotion recognition or biometric categorisation systems are deployed, the individuals exposed to them must be informed of the operation of the system (except where such systems are lawfully used for criminal detection, prevention, or investigation).

- **Ensure fairness in interaction with users (including consumers)** – AI systems should be used in a manner that does not result in misleading, deceptive, or coercive practices, namely when dealing with consumers – for more details, see question 48. In addition, online platform interfaces (which may embed or be enabled by AI systems) must not be operated in a way that deceives or manipulates users or that otherwise materially distorts or impairs their ability to make free and informed decisions (dark patterns) within the scope of the Digital Services Act (DSA).

  ↳ The use of certain AI systems that may impact humans are prohibited. See Prohibited Practices.

To whom the obligations apply

The obligations described above primarily apply to **deployers of AI systems**. For an explanation of the concept of AI deployer, see question 29.

Additionally, the obligations arising from the GDPR apply to deployers/users that are **data controllers**. For an explanation of who data controllers are, see question 33.

The prohibition of dark patterns under the DSA applies to **providers of online platforms** (operating their online interfaces) offering services in the EU.

AI Deployment Checklist — Transparency | explanations/information

Goal:

Assess compliance with the applicable requirements, specifically:

- Whether AI deployment is done in such a way that individuals receive information and meaningful explanations when affected by AI decisions – applicable to the deployment of certain high-risk AI systems (and to the deployment of all AI systems under the GDPR in case of automated decision-making).
- Whether AI deployment is done in such a way that individuals receive information on the use of AI – applicable to the deployment of (most) emotion recognition systems and biometric categorisation systems.

For consumer topics, see question 48.

Questions:

- What decisions impacting individuals are made or assisted by the AI system?
- Who are the individuals impacted by such decisions?
- What adverse impacts, if any, are envisaged for the individuals as a result of the AI intervention in decision-making?
- Which risks are anticipated (e.g., under-recognising when decisions require meaningful explanations and information, no or inadequate AI explanation or information, fragmented or inconsistent information, no documentation of decision processes involving AI)?
- Which measures will be implemented to address the risks? Notably, what will be the form, timing, content and channel for the provision of explanations and information?
- Does the AI system have features that enable the provision of information and explanations, or that prevent it?

Expertise required:

Technology & Systems; Legal; Compliance; Ethics; Data including Data Protection Officer (see question 42).

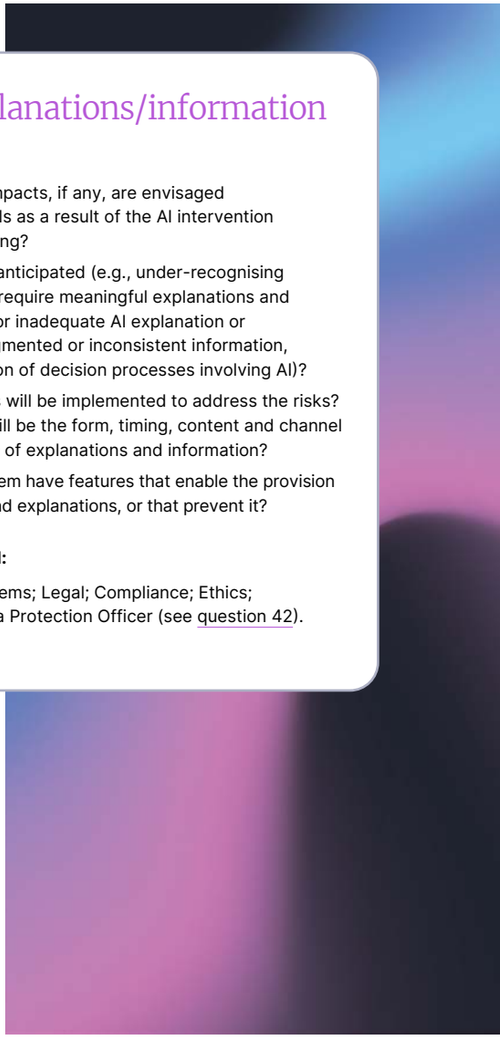

Risk mapping:

Transparency risks of AI deployment	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What requirements must be met when deploying AI systems that create deep fakes?

AI systems that generate or manipulate image, audio or video content that looks deceptively real – "**deep fakes**"–, as well as AI systems that generate or manipulate text informing the public on matters of public interest, are subject to specific transparency and disclosure obligations under the **AI Act**:

↳ Deep fakes are highly realistic, AI-Generated or AI-manipulated images, audio or video that depict events or people that never actually existed. Their main risks include deception and misinformation, reputational and fundamental rights harms (e.g., defamation, privacy breaches, non-consensual intimate content), data protection infringements, impersonation-based fraud and social engineering, and broader compliance issues such as EU transparency duties for synthetic media and unfair practice rules where consumers are misled.

- **Disclosure that the content is artificially generated or manipulated** – AI-generated or AI-manipulated content must be clearly disclosed as such.
  This obligation does not apply:
  - When use of the content is authorised by law to detect, prevent, investigate or prosecute a criminal offence.
  - For certain deep fakes that are evidently artistic, creative, satirical, fictional or analogous works or programmes, where disclosure can be limited to indicating that such AI-generated or manipulated content exists, in a way that does not interfere with the display or enjoyment of the work.
  - For AI-Generated text that has been subject to human review or editorial control and where a natural or legal person holds editorial responsibility for its publication.

In addition, the **Digital Services Act** (DSA) requires that measures be implemented to ensure that an item of information (generated or manipulated image, audio or video) closely resembling real persons, objects, places or events, and which could misleadingly appear authentic or truthful, is distinguishable through prominent markings when presented on online interfaces. An easy-to-use functionality must be provided to enable users to flag this kind of content.

To whom the obligations apply

The transparency and disclosure obligations under the AI Act apply to **deployers** of AI systems. For an explanation of who the AI deployer is, see question 29.

In turn, the obligations of the DSA apply to **VLOPs** and **VLOSEs** (very large online platforms and very large online search engines) offering services in the EU. VLOPs and VLOSEs are platforms and engines with a monthly average number of active recipients of the service in the Union equal to or higher than 45 million (in practice, used by more than 10% of the 450 million people in the EU), and which are designated as such by the European Commission.
See Use of AI in online platforms.

AI Deployment Checklist — Tansparency | Deep fakes

Goal:

Assess compliance with the applicable requirements, specifically whether the deployment of the AI system is done in such a way that deep fakes and AI-generated text are disclosed as such – applicable under the AI Act to most AI deployments generating deep fakes or AI-generated text published with the purpose of informing the public on matters of public interest.

Questions:

- What content does the AI system generate or manipulate?
- Could the AI-generated content be mistaken for human-created content?
- Which risks are anticipated (e.g., unlabelled content; removable labelling; ambiguous or misleading labels; inconsistent application across features, channels and content)?
- Which measures are envisaged to address the risks?
- Does the AI system have features that enable the provision of disclosures, or that prevent it?

Expertise required:

Technology & Systems; Legal; Compliance; Ethics; Data including Data Protection Officer (see question 42).

The Code of Practice on Transparency of AI-Generated Content

The Code of Practice on Transparency of AI Generated Content, published June 2026, set out voluntary commitments to help identify AI generated and manipulated content, especially deep fakes and public interest text.

It lays down design and placement rules requiring that icons, labels or disclaimers appear clearly at first exposure, are consistent and meet accessibility requirements. The Code provides for an EU icon and its further development and usability as a minimum state-of-the-art implementation.

It also sets out specific commitments for artistic and creative works and for public interest text under human review or editorial control, and acknowledges proportionate compliance paths that take account of organisations' size and resources.

The European Commission also issued draft guidance on this topic in May 2026.

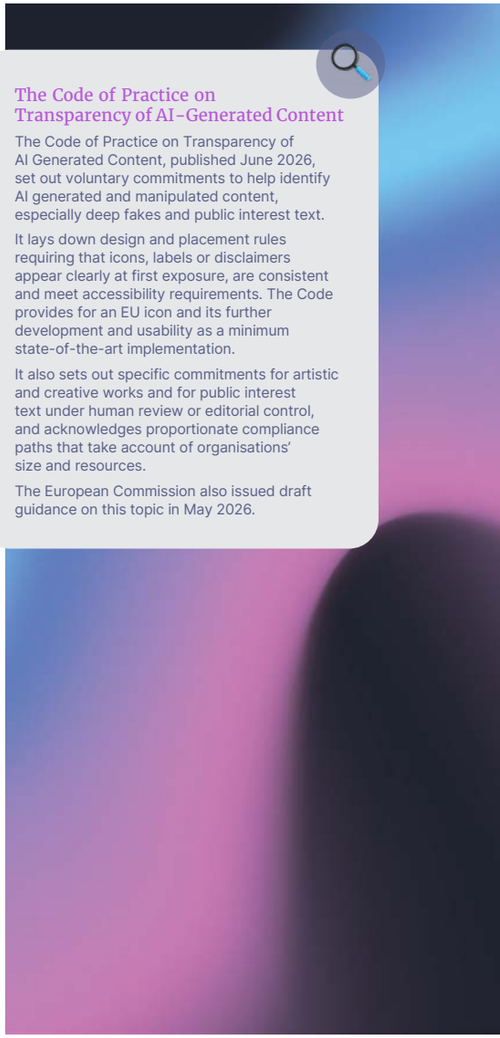

Risk mapping:

Deep fake risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can data be obtained to deploy an AI system?

Where data for AI deployment is not already in the organisation's possession, it must be obtained from third parties, either directly from data holders or data subjects (for personal data), or via intermediaries such as data marketplaces or brokers, which allow bulk access without separate contracts with each holder or subject.

↪ For an assessment of the **types of input data** for the AI system, see question 11, which also applies hereto, and 🔍 The different types of personal data.

↪ For an assessment of the EU legal approach to facilitating data sharing, see 🔍 The EU approach to facilitate data sharing.

Data can be acquired either based on **authorisation** (consent, contract/licence) or by relying on legal exceptions or limitations that allow for the use of data without authorisation. For more details, see question 12, as the same considerations also extend to AI deployment.

Authorisation

Authorisation is in principle needed to use both personal and non-personal data, though specificities may apply.

For **personal data**, because consent can be difficult to obtain in AI due to scale and re use, organisations typically rely on other lawful grounds under the **General Data Protection Regulation** (GDPR).

- **Legitimate interests** are often resorted to in the AI lifecycle, provided that a Legitimate Interests Assessment (LIA) confirms the interest is lawful, precise, real and present; the processing is necessary; and the interest is not overridden by data subject rights. This ground does not apply to public authorities in the performance of their tasks.
- Another ground for processing is the processing of special categories of personal data for **bias detection and correction**. This is allowed under the new updated version of the AI Act for the deployment of AI (high-risk or not) under tight necessity and protection requirements.

For **non-personal data**, use typically requires a contract or licence where the data is protected (see questions 11 and 40).

↪ A contract is also needed in other situations, such as, in most cases, for using readily available non personal data generated by users in the Union through connected products and related services (e.g., connected vehicles, medical or fitness devices, smart home apps) placed on the EU market (Data Act).

A case-by-case assessment is essential to determine whether a contract or licence is required. For instance, whether a licence for copyrighted content or databases is needed may depend on whether acts such as reproduction, transformation, making available, extraction or re utilisation occur in light of the relevant IP regime.

↪ Recent EU and national case law shows divergent views on whether AI, including outputs, "contains" protected works, and key questions on reproduction and communication to the public in AI outputs are now before the CJEU in Case C 250/25 (Like Company v. Google).

Legal exceptions and limitations

Legal exceptions and limitations allow for the use of data even when a licence is lacking. The most relevant of these include data extraction and text and data mining (TDM).

- **Data extraction** – under the Database Directive, non substantial, non repeated and non systematic extractions of data, as well as the extraction of data from connected product or public sector datasets made available for re use, are permitted.
- **Text and data mining** (TDM) – under the Digital Single Market (DSM) Directive, TDM of copyrighted content is permitted without authorisation for scientific research uses by research organisations and cultural heritage institutions with lawful access, and more broadly by any user (including commercial) with lawful access, unless rightsholders have clearly opted out by reserving their TDM rights via machine readable means or terms.

To whom this is relevant

It is the **AI deployer**/user that needs to obtain data for deploying the AI system through any of the means indicated above. For an explanation of who the AI deployer is under the AI Act, see question 29.

AI Deployment Checklist — Acquisition of input data

Goal:

Assess:

- What are the sources of the data – relevant for the deployment of all AI systems.
- Whether the data is lawfully obtained for the intended purposes – applicable to the deployment of all AI systems.

Questions:

- For each dataset, who is the owner/holder?
- How is the data obtained and from whom?
- Is there a clear inventory and data map linking each dataset to sources, means of acquisition of data and concrete uses?
- Do the outputs of the AI system disclose the input data, directly or indirectly?
- Which data risks are anticipated (e.g., unclear ownership of data, unlawful acquisition and/or use of data, over-reliance on legal exceptions, mismatch between acquisition and purpose, inconsistent handling, regurgitation or memorisation of protected content)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Procurement; Legal; Data including Data Scientists, Analytics and Data Protection Officer (see question 42).

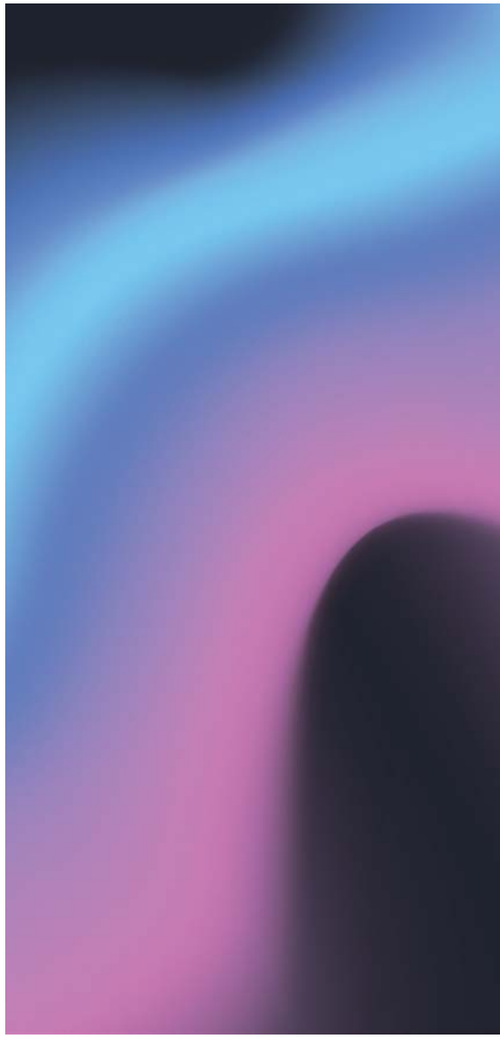

Risk mapping:

Data acquisition risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Should the data used to deploy the AI system have any specific features?

The data used in deployment, i.e., the data on the basis of which the AI system produces an output (the "input data"), must meet the requirements arising from the **AI Act** for high-risk systems.

In particular, input data for deployment of high-risk AI systems should be:

- **Relevant** – i.e., adequate for the result the data is meant to support, assessed in light of the AI system's purpose.
- **Representative** – i.e., able to sufficiently reflect, in an accurate manner, the relevant situation, setting or population in view of how the AI system will be used.

Other important features, not expressly detailed in the AI Act but stemming from it and from other legal frameworks, include:

- **Timeliness** – data should reflect the relevant reality at the required point in time, and be subject to governance and management practices such as updating.
- **Interoperability** – the ability of data to be reused across systems or entities. While not mandated by the AI Act, it is recognised as important, and EU data law promotes it for public sector, connected product data, and data space contexts.

The AI Act establishes that the instructions for use (see question 27) must contain, when appropriate, specifications for the input data, taking into account the intended purpose of the high-risk AI system.

↪ Specifications could include format, modality, metadata, quality, among others, so that the data being fed into the system is known.

More broadly, data should also meet general legal requirements, relevant for both high-risk and non-high-risk AI systems:

- **Privacy** – compliance with the General Data Protection Regulation (GDPR) principles: lawfulness, fairness and transparency; purpose limitation; data minimisation; data accuracy; storage limitation; data integrity and confidentiality (see question 33).
- **Security and integrity** – protection against compromise during processing. While the AI Act does not regulate the security of the data itself, the GDPR and security legislation impose security obligations (see question 31), and the Cy-bercrime Directive criminalises unlawful data interference.
- **IP compliance** – data must be sourced under appropriate licences or exceptions (see question 38).

↪ See 🔍 The contribution of the EU data legislation to data quality.

To whom the obligations apply

The obligation to ensure that data is relevant and sufficiently representative applies to **deployers** of high-risk AI systems with control over the input data. For an explanation of who the AI deployer is, see question 29.

The security obligations apply to **essential and important entities, as well as critical entities**. For an explanation of who these actors are, see question 31.

The obligations arising from the GDPR apply to the **data controller**. For an explanation of this actor, see question 33.

The obligations arising from intellectual property provisions apply, in this context, to the **user of the protected data**.

AI Deployment Checklist — Features of input data

Goal:

Assess compliance with the applicable requirements, specifically:

- Whether the data meets the features required by the AI Act – applicable to the deployment of high-risk AI systems under the AI Act.
- Whether the data meets other features required by other legal acts, such as when it comes to personal data and IP – applicable to the deployment of all AI systems.

Questions:

- What features does each dataset have?
- Who ensures that the data have the legally required features?
- Which data risks are anticipated (e.g., inadequate features of data, misalignment of data used with stated purpose and risk profile of the AI system)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Quality; Legal; Compliance; Data including Data Scientists, Analytics and Data Protection Officer (see question 42).

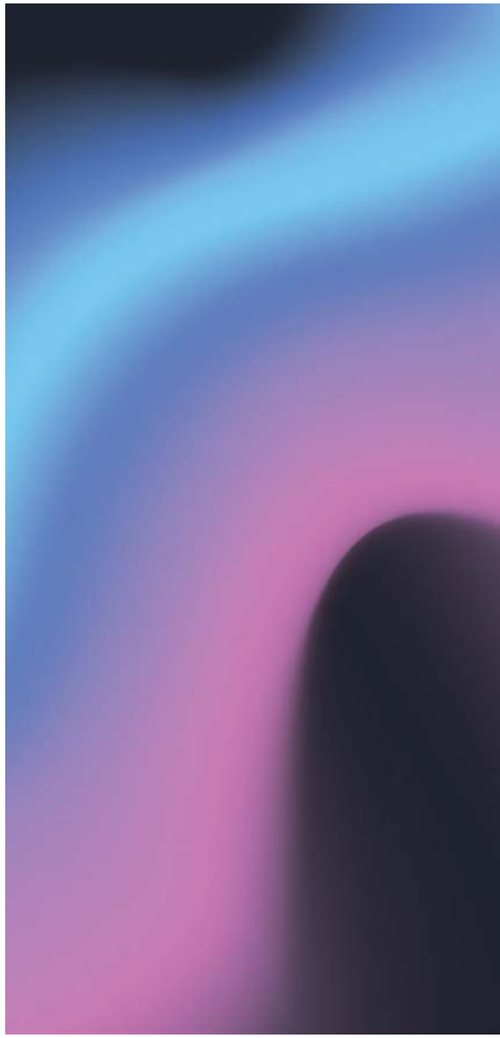

Risk mapping:

Data features risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What main aspects should be included in a contract for data to deploy the AI system?

A contract to supply data for the operation of an AI system (input data) is a binding agreement on how data is provided, accessed and used.
In broad terms, it is recommended that such a contract covers the following:

- **The data provided and its features** – precise description of the data, delivery dates and, where relevant, acceptance procedures; indication of quality requirements, especially for high-risk AI, and of the technical means of provision. Regular provision of data may be subject to SLAs.
- **The conditions for use** – clear rules on permitted purposes and fields of use, types of processing (including any conditions on modification), storage, internal and external uses, commercial exploitation, access controls, territory, duration and post-termination consequences (e.g., return or deletion).
  - ↳ Prohibited uses may also be expressly addressed (e.g., resale, uses for unlawful purposes such as profiling beyond the agreed deployment context).
  - ↳ To the extent the data is provided only for a fixed term, the effects of termination on continued data use should be assessed, including any impact on the AI system's intended operation or content generation.
- **Warranties and legal compliance** – provision of warranties relating to the data (e.g., with respect to their relevance and representativeness), or their exclusion where data is provided "as is"; commitments that data is lawfully provided, does not infringe third-party rights, and that both parties comply with personal data and IP rules for protected content; indication of provenance and applied quality controls.
- **Other key provisions** – provisions on reports and audits, changes and additional data, price and payment, security, confidentiality, liability and indemnification, term and termination, assignment, representations and warranties, governing law and dispute resolution, as well as other aspects (such as cooperation and information duties) needed to comply with applicable AI-related requirements.
  - ↳ **Ownership of AI outputs** should also be addressed, including, as relevant, potential licences to be granted by the rightsholder to the other party over such outputs.

Certain contracts must also respect mandatory rules. For instance, **contracts providing personal data** must comply with the General Data Protection Regulation (GDPR), **contracts licensing works or performances** (see question 38) must meet requirements on remuneration and transparency under the Digital Single Market (DSM) Directive and any national formalities, **data sharing contracts** under the Data Act must avoid prohibited non-negotiated terms in B2B scenarios (e.g., excluding liability for intent or gross negligence), **data sharing contracts with data intermediation services**, such as data marketplaces, or **copyright management organisations** (see question 38) must respect the specific conditions under the Data Governance Act (DGA) and the Collective Rights Management (CRM) Directive, and **contracts for data processing services** (including database as a service) must include mandatory provisions to facilitate switching between providers under the Data Act.

↳ See 🔍 **The risk of anti-competitive effects of data sharing agreements**, which is also relevant for data sharing agreements for AI deployment.

To whom is this relevant

The **AI deployer**/user should ensure the contract contains provisions allowing it to use the data as intended (without prejudice to mandatory requirements that shall be met as indicated above). For an explanation of who the AI deployer is under the AI Act, see question 29.

AI Deployment Checklist — Contracts for input data

Goal:

Assess:

- Whether contracts concluded for the data allow the use of the data as intended and duly protect the parties – relevant for all deployments of AI systems.
- Whether contracts concluded for the data comply with applicable obligations – applicable to all deployments of AI systems.

Questions:

- What conditions and limits of use does the contract set (types and purposes of use, field, users, territory, term, warranties)?
- Do the authorised uses match the AI deployment purposes and risk profile?
- Are there territorial, sectoral, exclusivity or non-compete clauses that constrain use?
- Are the data provider and AI provider actual or potential competitors, and does either hold a dominant position in any relevant market?
- Does the output have to be fed back to the AI provider or vendor?
- Which data risks are anticipated (e.g., lack of clarity in authorised and prohibited uses, licence does not match AI use, over-reliance on data rights obtained by AI vendors, competition risks)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Procurement; Legal; Data including Data Analytics and Data Protection Officer (see question 42).

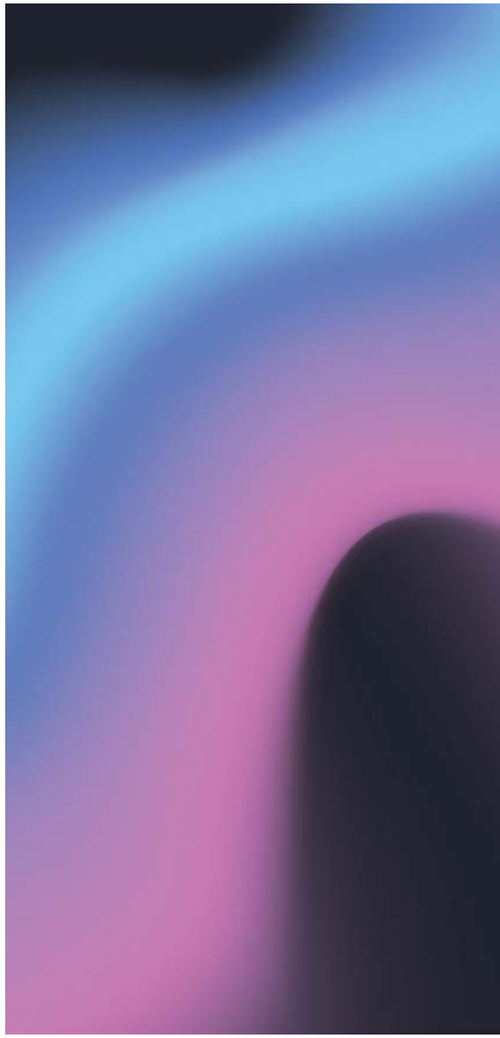

Risk mapping:

Data contract risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Are there any governance processes that should be implemented while deploying the AI system?

At deployment, mechanisms must be in place to monitor the AI system and ensure its compliance with applicable rules. Key measures include:

- Appropriate **technical and organisational measures** to ensure that the high-risk system is used in line with its instructions for use, human oversight measures are implemented, input data is relevant and representative, and logs are kept (see questions 27, 34, 35 and 39). In addition, a fundamental rights impact assessment (FRIA) must be performed for certain uses of high-risk AI systems (see question 30). These obligations arise from the AI Act.
- Appropriate and proportionate **technical, operational and organisational measures** to manage security risks and ensure resilience under the Network and Information Systems (NIS 2) and Critical Entities Resilience (CER) Directives (for more details, see question 31).
- Whenever **personal data** is processed, and regardless of whether the AI system is high-risk or not, a set of procedures must also be implemented, notably General Data Protection Regulation (GDPR)-compliant technical and organisational measures (TOMs), procedures to handle data subject rights and, where processing is likely to be high-risk, data protection impact assessments (DPIAs), which are often triggered in AI contexts due to the large-scale collection and further processing of personal data, and their innovative uses. Because a DPIA is required where processing uses new technologies, it will often be necessary when AI is used, especially where it processes special-category or highly sensitive data (see The Different Types of Personal Data) and/or supports solely automated decisions with legal or similarly significant effects for individuals.
  - ↪ For high-risk AI systems, the provider's information and instructions on how the system functions should be used to conduct the DPIA, and, where the DPIA already covers elements required for the FRIA, the FRIA should build on and complement that DPIA.
- **Energy management systems or energy audits**, required under the Energy Efficiency Directive (see question 32 for more details).

To whom the obligations apply

AI Act obligations apply to **AI deployers** of high-risk AI systems. For an explanation of who they are, see question 29.

Resilience obligations apply to AI deployers/users that are **essential, important and/or critical entities** (for an explanation of who these actors are question 31), whilst environmental obligations (namely on energy efficiency) apply where certain energy consuming thresholds are met (see question 32 on this). Data protection duties fall upon to the **data controller and processor** (see question 33 for an explanation of these actors).

AI Deployment Checklist — Governance processes

Goal:

Assess:

- Whether all required governance processes have been implemented – applicable to the deployment of high-risk AI systems under the AI Act, and to all deployments subject to other laws, notably cybersecurity obligations, the GDPR and energy efficiency legislation.
- Whether the governance processes have been implemented in a complementary manner – applicable to the deployment of high-risk AI systems subject to FRIA and DPIA.

Questions:

- Are there formal, repeatable processes for AI governance?
- Are those processes laid down in writing in internal policies or otherwise?
- Are those processes integrated with existing compliance frameworks?
- Which risks are anticipated (e.g., formal governance processes but weak governance in practice, fragmented governance processes, governance blind spots such as lack of governance processes for deployment of lower-risk systems, lack of mechanisms for updating governance processes, inconsistent application across AI deployments, inability to demonstrate effective implementation)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Quality; Legal; Compliance; Data including Data Protection Officer (see question 42).

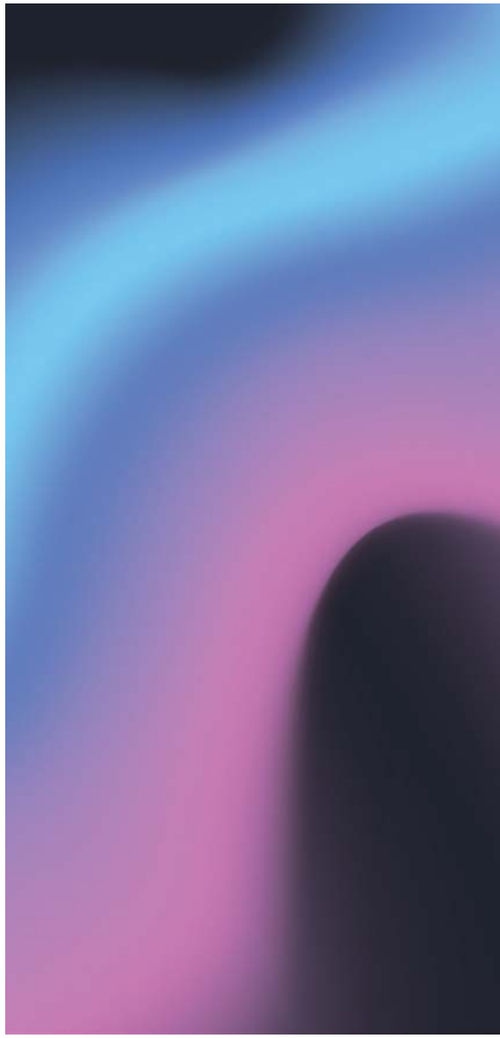

Risk mapping:

Governance processes risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Are there any governance structures that should be implemented for deploying AI systems?

The deployment of AI systems must consider not only the implementation of a set of governance procedures (as seen in question 41), but also of governance structures that are able to effectively ensure legal compliance.

Although governance structures are not specifically mandated in most cases, three main aspects are worth referring:

- **Fundamental rights governance** – under the AI Act, internal governance arrangements must be in place in case risks to fundamental rights materialise from the deployment of high-risk AI systems. This requirement applies when a FRIA must be performed: indeed, under the AI Act, the FRIA must consist of, among other elements, measures to be taken in case of the materialisation of risks, including the arrangements for internal governance.
- **Cybersecurity governance** – under the Network and Information Systems (NIS 2) Directive, management bodies of essential and important entities (see question 31 for an explanation of who these actors are) must approve the cybersecurity risk management measures and oversee their implementation (and are, in addition, liable for infringements thereof). Under the Critical Entities Resilience (CER) Directive, critical entities (see question 31 on what critical entities are) must designate a liaison officer or equivalent as the point of contact with the competent authorities.
- **Data Protection Officers (DPOs)** – data protection officers are required to be designated in certain cases under the General Data Protection Regulation (GDPR), such as when the core activities at stake consist of processing operations which, by virtue of their nature, scope and/or purposes, require regular and systematic monitoring of data subjects on a large scale; the core activities consist of large-scale processing of special categories of data or personal data relating to criminal convictions and offences; or the processing of personal data is carried out by a public authority or body (except for courts acting in their judicial capacity). EU or national law may require the designation of DPOs in other situations as well.

  ↳ The determination of the need to designate a DPO in AI deployment depends on a case-by-case assessment. However, it could be a best practice to have a person with DPO capacities and functions to ensure adequate levels of privacy and data protection within AI deployment, especially in the use of high-risk AI systems.

To whom this is relevant

The obligations apply to the **AI deployer** for high-risk AI systems under the AI Act, to **essential, important and critical entities** under the NIS 2 and CER Directives, and to the **data controller** under the GDPR. For an explanation of who these actors are, see questions 29, 31 and 33. Note, in any case, that the **data processor** is also under the obligation to designate a DPO in the situations indicated above.

AI Deployment Checklist — Governance structures

Goal:

Assess:

- Whether the required governance structures have been implemented – applicable to the deployment of high-risk AI systems subject to FRIA under the AI Act, and to all deployments under cybersecurity obligations and the GDPR.
- Whether governance processes sufficient to ensure legal compliance and promote a culture of transparency, accountability, and ethical behaviour have been implemented – relevant for the deployment of all AI systems.

Questions:

- Are there clear governance structures for AI deployment?
- Are tasks and responsibilities clearly allocated to each governance structure, in line with applicable legal obligations?
- Are AI structures duly coordinated and compatible with other internal structures (e.g., on data protection, cybersecurity, among others)? How?
- Are there channels to escalate concerns?
- Which risks are anticipated (e.g., formal governance structures but weak governance in practice, governance blind spots such as lack of governance structures for deployment of lower-risk systems, lack of mechanisms for internal coordination, poor coordination with competent entities)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Quality; Legal; Compliance; Data Protection Officer (see question 42).

Ethical leadership

Adopting a dedicated AI leadership structure is a keyway for AI deployers to ensure legal compliance and foster a culture of transparency, accountability and ethical behaviour, while also building trust and signalling reliability to the market.

Potential roles may include a Chief AI Officer for strategy and value alignment, an AI Officer for legal compliance (similar to a DPO), and an AI Ethics Board or Committee to oversee AI initiatives and provide guidance. Such a structure would monitor compliance in AI deployment, align practices with best practice (see question 45), engage with external stakeholders and regulators, and actively promote a culture of responsibility, transparency and respect for fundamental rights across all AI-related processes.

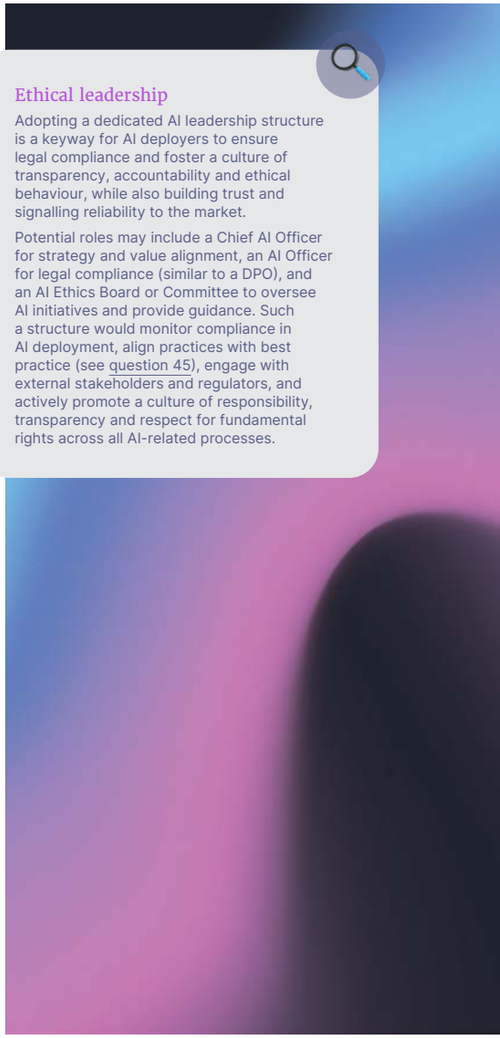

Risk mapping:

Governance structure risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

How can it be ensured that the persons dealing with AI meet the applicable obligations?

One of the key factors for ensuring legal compliance is the level of literacy of staff and other individuals involved in AI development.

↳ AI literacy refers to the skills, knowledge and understanding that allow the making of informed decisions about AI systems (regardless of them being high-risk or not) and the gaining of awareness about the opportunities and risks of AI and possible harm it can cause.

Under the AI Act, literacy measures must thus be taken to support the development of AI literacy of staff and other individuals dealing with AI systems, taking into account their technical knowledge, experience, education and training, the context the AI systems are to be used in, and the persons or groups of persons on whom the AI systems are to be used. The new updated version of the AI Act clarifies that this must not be understood as requiring that any specific level of AI literacy of any individual be guaranteed.

↳ To support the development of AI literacy in a proportionate way, organisations may consider role-specific measures. For instance:

- Technical staff (e.g., data scientists, engineers) could receive training on AI operation and cybersecurity;
- Legal and compliance staff could receive training on the AI Act, data protection, cybersecurity and resilience, consumer protection, liability and sector-specific rules;
- Management and decision-makers could receive training enabling them to understand AI-related risks and accountability frameworks;
- Persons responsible for human oversight could be trained to understand AI outputs, intervene effectively, and escalate risks when needed.

The above is reinforced by other legal frameworks, such as those on data protection.

↳ Continuous learning and updating are important. Given the rapid evolution of AI technologies and regulatory frameworks, training and awareness activities should ideally not be one-off exercises, but ongoing processes, updated from time-to-time to reflect technological developments, new risks, incident learnings and regulatory changes.

↳ Literacy measures can be embedded in AI governance and other AI controls (e.g., incident handling), so that feedback from such activities is integrated into literacy efforts and helps keep them up to date.

The Commission and Member States must support and facilitate compliance with AI literacy obligations, in particular by SMEs (including through the facilitation of the drawing up of codes of conduct – see question 45). For this purpose, the Commission must publish practical examples for compliance. In addition, the AI Board will issue recommendations to support the Commission and Member States, including setting common objectives.

To whom the obligations apply

The literacy obligation established by the AI Act for all AI systems applies to **AI deployers** (for an explanation of who they are, see question 29).

AI Deployment Checklist — Literacy

Goal:

Assess whether literacy measures have been implemented – applicable to the deployment of all AI systems under the AI Act.

Questions:

- Who needs AI literacy and for what?
- What level (e.g., basic vs. advanced) and type of literacy (e.g., technical, legal) is needed?
- Which measures are envisaged to deliver and maintain literacy?
- Is there a clear inventory or mapping of literacy measures linking each measure and its recipients within the organisation?
- Are literacy measures documented and recorded? How?
- Are literacy measures embedded in, and coordinated with, AI governance measures?
- Which risks are anticipated (e.g., undefined or narrow scope of literacy recipients, one-size fits all or superficial training, limited literacy measures not covering all required aspects, no literacy documentation or records, no connection between literacy and governance measures impacting literacy updating)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Legal; Compliance; Data including Data Protection Officer (see question 42); Management.

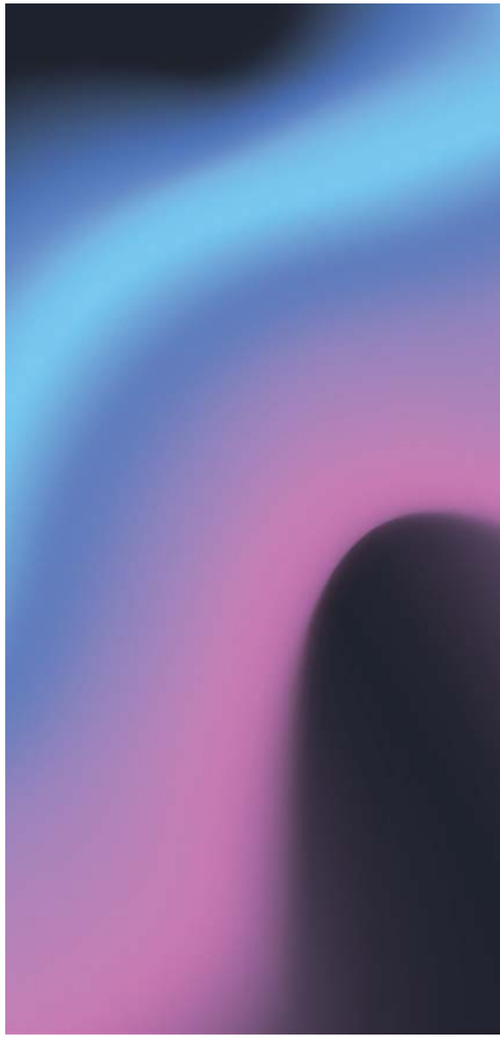

Risk mapping:

Literacy risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

What if the AI system is used in a non-compliant manner or malfunctions?

When an AI system malfunctions or is used in a non-compliant manner, the law foresees reporting, corrective action and liability.

Reporting of incidents

Serious incidents of high-risk systems identified by the deployer must be immediately informed to the AI provider, then to the importer and distributor and the relevant market surveillance authorities.

Cybersecurity laws and the General Data Protection Regulation (GDPR) also address the reporting of incidents. For example, under the Network and Information Systems (NIS 2) Directive, significant incidents must be notified to CSIRTs (Computer Security Incident Response Team) or, where applicable, the competent authority (and affected users). For significant cyber threats, service recipients must be informed of any recommended measures (and the threat itself where appropriate), with public communication on request by the CSIRT or competent entity. Under the Critical Entities Resilience (CER) Directive, incidents that significantly disrupt essential services must be notified to the national authority, while the GDPR requires notification of personal data breaches to the data protection authority, among other actions, such as communication of the breach to affected data subjects.

Suspension and information measures

Despite some particularities established in the AI Act, non-conformity or risk of high-risk systems triggers a set of measures: where use in line with the instructions could result in the AI system presenting a risk, the use of the system should be suspended with notification to the provider or distributor and the competent market surveillance authority; and the system's operation must be monitored against the instructions for use, informing the provider where nonconformities arise.

Corrective measures

Under the new updated version of the AI Act, where non-compliance with AI obligations is found in the deployment of systems under the oversight of the AI Office, measures to ensure compliance must be taken. These systems comprise, in essence, AI systems based on general-purpose AI models where the model and system share the same provider (subject to specified exceptions), and AI systems that constitute or are integrated into very large online platforms or very large online search engines, to the extent the deployer is also the provider or part of the same undertaking.

Liability

Breaches may lead to administrative fines, civil liability and, in some cases, criminal liability.

When it comes to fines, while some acts leave fines to Member States, many already set penalties for certain breaches.

↪ The AI Act provides for fines up to €35 million or 7% of global turnover for serious breaches, and up to €15 million or 3% for others, with reduced caps for SMEs and small mid cap enterprises (SMCs).

Civil liability arises where breach of legal or contractual duties causes damage, while criminal **liability** is defined mainly at national level (e.g., for IP infringements). Still, criminal law may also apply, in particular under cybercrime rules criminalising illegal access to information systems and data.

To whom the obligations apply

The obligations apply to the **AI deployer** under the AI Act, **essential, important or critical entities** under the NIS 2 and CER Directives, and the **data controller** under the GDPR. See questions 29, 31 and 33 for an explanation of who these actors are.

↪ When it comes to reporting and other actions, in cases where the above entities are different organisations/persons, assistance and coordination must be ensured to the extent breaches of one applicable law impact other stakeholders.

AI Deployment Checklist — Breach and malfunctions

Goal:

Assess whether there are processes to address incident reporting, suspension and information measures – applicable to the deployment of high-risk AI systems under the AI Act, and to all deployments subject to cybersecurity rules and to the GDPR.

Questions:

- Are there clear processes to handle AI non compliance and malfunctions?
- Are there enough capable human resources to handle these processes?
- Are incident taxonomies and reporting templates in place?
- Do processes include improvement steps after incidents, malfunctions or breaches, with criteria for reporting?
- How are corrective actions decided and implemented, and how are these processes aligned with wider governance arrangements (see question 41)?
- Which risks are anticipated (e.g., formal processes but weak implementation in practice, fragmented processes, misalignment among processes required under several laws, lack of mechanisms for updating processes, inconsistent application across AI deployments)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Risk & Security; Legal; Compliance; Data including Data Protection Officer (see question 42).

Risk mapping:

Breach and malfunctions risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Civil and criminal liability arising from the use of an AI system

A key question is how far using an AI system can expose a deployer to civil or criminal liability when damage or a crime occurs in connection with AI. The issue is whether the harm stems from the deployer's conduct or from the AI system's own characteristics and functioning.

Because many AI systems are black boxes, it may be hard to demonstrate that the AI caused the damage or offence. Even if causation is proven, liability usually requires a natural or legal person acting with fault. The central question, therefore, is whether the deployer, due to how it used the system, is the legally relevant cause of harm, or whether it would have occurred anyway given the system's features.

Key factors include whether the deployer stayed within the system's intended purpose, respected known limitations, ensured adequate human oversight, and avoided negligent behaviour. In criminal law, it must also be assessed whether the deployer actually committed the offence typified in law by using the AI. Ultimately, liability must be decided case by case.

Are there standards, guidelines or best practice that can be followed?

Standards, certifications, codes of conduct and guidelines are tools that help organisations implement and evidence compliance with AI, cybersecurity, sustainability and data protection rules.

Standards and common specifications

Harmonised standards (developed by CEN, CENELEC or ETSI and cited in the EU Official Journal) and Commission "common specifications" translate legal requirements into technical rules. Using these gives a presumption of conformity. Standards for processes, personnel and data may be especially relevant for AI deployment (e.g., standards for information security management, incident management, data protection measures, environmental and energy management, data quality, among others). Other standards (e.g., ISO) can support compliance but do not themselves create this presumption unless incorporated by reference in the harmonised standard.

↪ Harmonised standards under the AI Act are not yet in place, but the Commission has asked CEN and CENELEC to develop a set of standards in ten areas for high-risk systems: risk management, data and data governance, record-keeping, transparency, human oversight, accuracy, robustness, cybersecurity, quality management and conformity assessment. Other standards are also being developed, notably on bias management and sustainable AI.

Certifications and labels

Certification formally attest that products, services or organisations meet defined requirements and can create presumptions of compliance (see 🔍 Privacy Certification, Seals and Marks). Certifications and labels for products and services do not directly apply to AI deployment. However, using certified products or services can help show that an AI system was chosen with due care, support reputation and, in some cases, satisfy legal obligations where use of certified solutions is mandated.

↪ Under the Network and Information Systems (NIS 2) Directive, Member States and the Commission may require the use of particular ICT products, services and processes that are certified under European cybersecurity certification schemes developed under the Cybersecurity Act.
See 🔍 The European Cybersecurity Certification Scheme.

Codes of conduct

Voluntary codes can operationalise legal obligations and, under the AI Act, may foster the application of legal requirements to non-high-risk systems, including regarding aspects such as sustainability (sustainable use of AI), literacy and trustworthy AI principles. Under the GDPR, adherence to approved codes can be used as evidence of compliance.

Guidelines

The Commission, the AI Office and AI Board, as well as authorities such as ENISA (the EU Cybersecurity Agency) and the EDPB (the European Data Protection Board), issue guidelines and opinions that clarify how to apply the laws in practice.

↪ Guidance on the AI Act includes the Guidelines on the definition of an AI system (2025) and on prohibited AI practices (2025), as well as forthcoming guidelines on high-risk classification (draft guidance issued in May 2026), transparency(draft guidance issued in May 2026), high-risk requirements, obligations for providers and deployers, responsibilities along the AI value chain, substantial modification, and interplay of the AI Act with EU law.

To whom this is relevant

The **AI deployer** (under the AI Act), **essential, important and critical entities** (under the NIS 2 and Critical Entities Resilience (CER) Directives) and the **data controller** (under the GDPR) may benefit from the above (or be required, under the NIS 2 Directive, to use particular certified ICT products, services and processes). For an explanation of who these actors are, see questions 29, 31 and 33.

AI Deployment Checklist — Standards, certifications and guidance

Goal:

Assess whether there are standards, certifications and guidance that may apply – relevant for the deployment of high-risk AI and non-high-risk AI systems under the AI Act, depending on the standard and guidance at stake, and to all AI deployments subject to cybersecurity, data protection and environmental laws.

Questions:

- Which standards and common specifications apply, do they create a presumption of conformity, and will they be used?
- Which certifications or labels are available, do they evidence or presume compliance, and are they planned to be used?
- Are there certified products or services and are they planned to be used?
- Have relevant codes of conduct or guidelines been identified, and is adherence or use envisaged?
- Are there clear processes to identify, assess, decide on and implement standards, specifications, certifications, codes and guidelines, and to align this with wider governance (see question 41)?
- Which risks are anticipated (e.g., formal processes but weak implementation in practice, under-relying or over-relying on standards and common specifications, misalignment between certification and legal requirements, treating codes as symbolic commitments)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Risk & Security; Ethics; Legal; Compliance; Data including Data Protection Officer (see question 42).

Risk mapping:

Standards, certification and guidance risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

The Ethical principles for AI

Ethics is central to AI. It complements regulation and helps interpret and apply the law where rules are ambiguous, incomplete or still evolving, while remaining more agile and responsive to new societal concerns.

At deployment, risks to fundamental rights, safety, fairness and trust crystallise and depend heavily on context, so ethical principles are key for the responsible use, supervision and lifecycle management of AI systems. Core principles repeatedly identified across many initiatives include human rights and dignity, diversity and non-discrimination, privacy and data governance, human agency and oversight, fairness, harm prevention, transparency and explainability, accountability, and societal and environmental well-being.

These principles can guide deployers in complying with AI rules and making judgement calls, especially where the law leaves room for discretion.

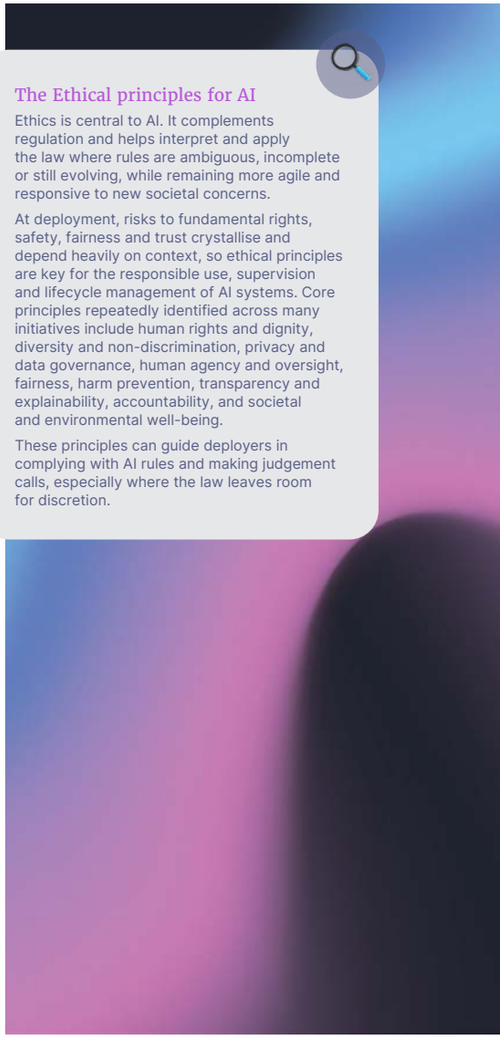

Can the output of an AI system be protected? How?

AI output can generally be protected by intellectual property, either as a copyrighted work, a database, or an invention. Protection by a related right has also been discussed.

Protection as a work

AI output (such as images, text or sound) can be protected by copyright. Copyright requires a human author and, therefore, AI-Generated outputs without meaningful human input are not protected, but AI-assisted creations and AI | Human creations can be. The bar for relevant human intervention is relatively low, so any intellectual input, even if apparently small within the overall creative process, might be relevant to establish copyright protection.

↪ The European Commission developed a four-step test to assess whether AI outputs are copyrightable or not.
Step 1: Production in literary, scientific or artistic domain.
Step 2: Human intellectual effort.
Step 3: Originality/creativity.
Step 4: Expression.

Protection as a database

AI output can be a database and, as such, can also be protected by copyright (when the selection and arrangement of its contents is original/creative) or by the *sui generis* right (which protects substantial investment in obtaining, verifying or presenting the database contents).

↪ The investment may lie in the AI technology and know how used to produce the database, but not in generating new data, as it must relate to pre existing data or its verification or presentation.

Protection as an invention

Inventions can be patented when they have an industrial application, are new and involve an inventive step. Patent law also requires a human inventor, so inventions generated entirely byAI are not patentable, but inventions where humans design or train AI can be.

Related rights

There have been discussions surrounding the protection of AI outputs under a new related/neighbouring right. However, in the EU legal framework, not only is there no related right for this purpose, but also no legislative endeavour or relevant movement towards that direction.

↪ AI outputs can only benefit from existing related rights, such as: phonogram rights for audio, film producer rights for audiovisual content, broadcaster rights for AI-assisted broadcasts, and press publisher rights for the online use of press publications.

AI outputs can also be potentially protected through other means, such as **trade secret**, **unfair competition**, **contractual confidentiality** and **technical measures**, with unlawful circumvention potentially amounting to cybercrime (e.g., illegal access or data interference).

To whom this is relevant

The person (legal or natural) that wishes to protect the AI output must assess the applicable routes of protection.
When it comes to **copyright**, the rightsholder over the output is in principle the natural person(s) that developed it, or the legal person(s) designated as the rightsholder under national law – such as an employer or the customer.

The maker of the database, i.e., the person that takes the initiative and assumes the risk of investing, is, in turn, the rightsholder of the ***sui generis right.***
When it comes to protection as an invention, the holder of the **patent** will in principle be the inventor, though specific national rules or legal presumptions may be established, such as determining that the holder of the patent is the employer.

The rightsholder/holder of the patent can be the AI deployer under the AI Act, since it is the deployer that triggers the acts that give rise to these rights by activating the AI system. Nevertheless, rights over outputs may be allocated to other entities (see, for instance, question 28 on contracts for use of the AI system).

AI Deployment Checklist — IP Protection

Goal:

Assess whether the output of the AI system can be protected and by which routes of protection – relevant for all AI deployments.

Questions:

- Which AI outputs are to be protected?
- How were the AI outputs developed and what was the level of human intervention?
- Who exactly intervened in the development of the AI outputs?
- Are there contractual limitations or requirements on AI ownership (e.g., in contracts with the AI provider, with service providers, with employees)?
- Which routes of protection are possible?
- Are there IP management and governance processes (see also question 41)?
- Who is responsible for these processes?
- Are there sufficient and capable human resources to handle these processes?
- How are these processes coordinated with the governance processes (see question 41)?
- Which risks are anticipated (e.g., unclear routes of protection, unclear ownership of IP, unclear IP management processes, robust processes but weak implementation, lack of coordination between different routes of protection, over-disclosure undermining trade secrets, weak technical protections)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Quality; Legal.

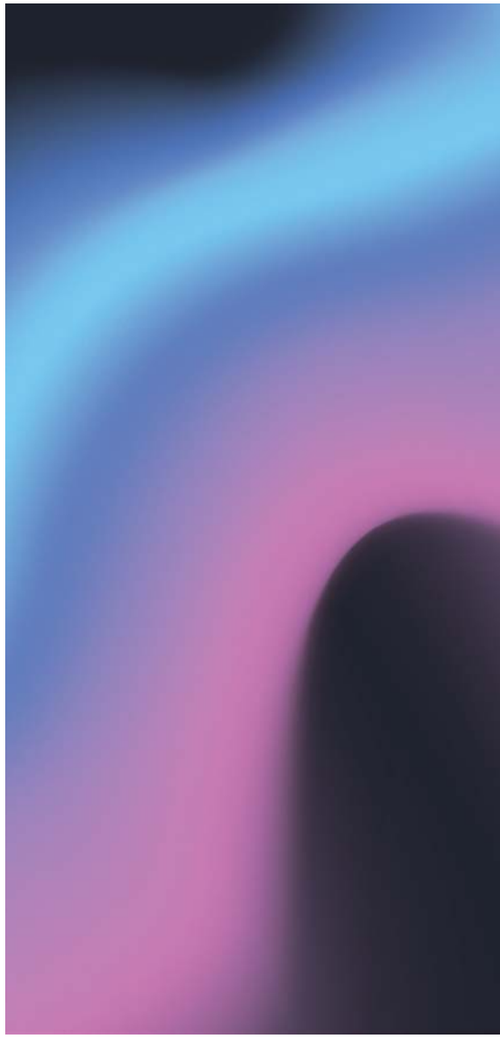

Risk mapping:

IP protection risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Are there any particularities for certain organisations or sectors?

The AI Act addresses the deployment of AI systems by certain sectors and organisations through two main mechanisms:

- Risk-based approach whereby certain AI systems are considered prohibited or high-risk (see question 24).
- Specific provisions for certain organisations, notably:
  - **Finance** – for AI deployers that are financial institutions subject to requirements regarding their internal governance, arrangements or processes under Union financial services law, the obligation to monitor the high-risk AI system on the basis of the instructions for use (see question 27 and 41) is considered to be met by complying with the rules on internal governance arrangements, processes and mechanisms pursuant to the relevant financial service law. In addition, such AI deployers must maintain the logs (see question 35) as part of the documentation kept pursuant to Union financial services law.
    - ↪ Financial institutions are further subject to specific provisions on cybersecurity under the Digital Operational Resilience Act (DORA), notably as regards ICT risk management, incident reporting and oversight of critical ICT thirdparty providers. An integrated or aligned cyber approach that also includes AI is thus recommended.
  - **Employment** – before putting into service or using a high-risk AI system at the workplace, AI deployers who are employers must inform workers' representatives and the affected workers that they will be subject to the use of the high-risk AI system. This information must be provided, where applicable, in accordance with the rules and procedures laid down in Union and national law and practice on information of workers and their representatives.
  - **Public sector** – before putting into service or using certain high-risk AI systems, deployers that are public authorities, Union institutions, bodies, offices or agencies, or persons acting on their behalf, must register themselves, select the system and register its use in the EU database for high-risk systems. If the system is not registered, they shall not use it and shall inform the provider or distributor.
  - **Law enforcement** – deployers using high-risk AI systems for post-remote biometric identification in the framework of an investigation for the targeted search of a person suspected or convicted of having committed a criminal offence, must obtain prior or prompt authorisation from a judicial or specific administrative authority in defined circumstances and report on uses.

We further recall that the obligation to perform a FRIA applies to certain deployers of high-risk AI systems, as seen in question 30.

To whom the obligations apply

The above obligations apply to **deployers** of high-risk AI systems, with the specific provisions on the financial sector, employment and public sector applying to **AI deployers** that are financial institutions, employers or the referred public sector entities. For an explanation of the concept of AI deployer, see question 29.

AI Deployment Checklist — Sector specificities

Goal:

Assess whether the deployment of the AI system is subject to dedicated provisions of the AI Act – applicable to the deployment of high-risk AI systems in finance, employment, public sector and law enforcement.

Questions:

- Is the AI deployer a financial institution, employer, public entity or using post-remote biometric identification systems to search for criminal suspects?
- If the AI deployer is a financial institution, is it subject to internal governance or processes requirements under Union financial services law?
- If the AI deployer is a financial institution, which risks are anticipated (e.g., duplicated processes, inconsistent controls)?
- If the AI deployer is an employer, which risks are anticipated (e.g., lack of information to workers, overly generic or opaque information)?
- If the AI deployer is a public entity, which risks are anticipated (e.g., lack of registration, lack of internal checks for registration)?
- If the AI deployer is using the AI system for law enforcement purposes, which risks are anticipated (e.g., lack of proper or prompt authorisation, mismatch between authorisation and use)?
- Which mitigation measures will be implemented to address the risks?

For other relevant questions, see questions 30, 35 and 41.

Expertise required:

Legal; Compliance.

Use of AI in online platforms

AI can be used in online platforms for many purposes, including in the context of advertising and recommender systems. It may also be used for dark patterns designed to deceive or manipulate users, or otherwise significantly impair their ability to make free and informed decisions (see question 8).

When AI is used for ads or recommender systems, platforms must comply with the Digital Services Act's (DSA) transparency rules, and dark patterns are outright prohibited. These obligations apply to providers of online platforms offering services in the EU, with additional, stricter duties for VLOPs and VLOSEs (very large online platforms and very large online search engines).

The European Commission is preparing a Digital Fairness Act to deal with dark patterns, addictive designs/manipulative online interfaces, and unfair personalisation, with a view to update consumer protection against manipulative digital practices. The proposal is expected in late 2026.

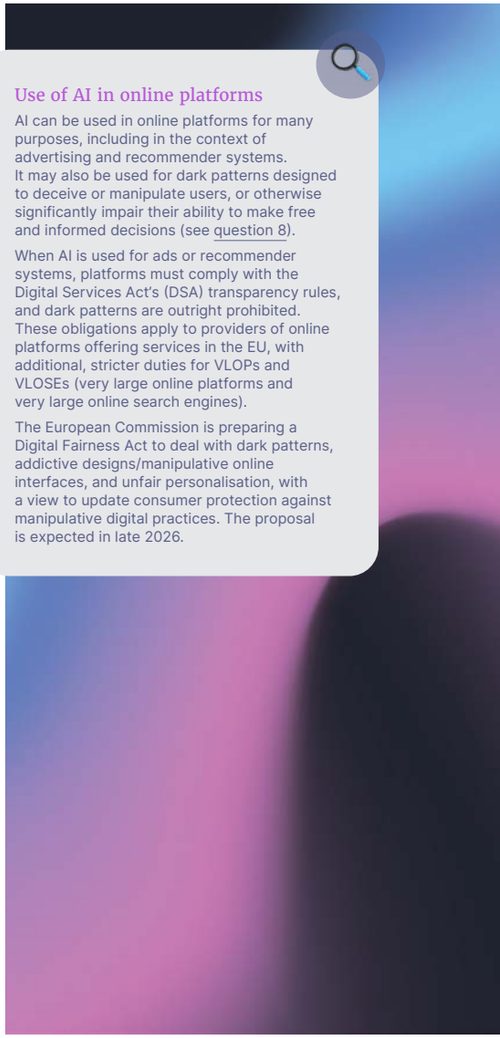

Risk mapping:

Sector-specific risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above. See also questions 30, 35 and 41.

Are there any particularities when the AI system is used in relation to consumers?

Although the AI Act does not explicitly target consumers, AI can play a role at multiple stages of the consumer relationship — from targeted advertising and chatbot interactions to the personalisation or differentiation of consumer offers, as well as the conclusion and performance of contracts.

In all these moments, EU consumer law applies, including notably the Consumer Rights Directive (CRD), the Unfair Commercial Practices Directive (UCPD) and the Unfair Contract Terms Directive (UCTD).

A set of aspects are especially relevant when deploying AI in consumer relations (for an assessment of the provision of AI to consumers, see 🔍 **The Provision of AI Systems to Consumers**). These aspects include:

- **Fair commercial practices**
  Since consumers have less bargaining power, consumer protection law prohibits commercial practices that exploit this imbalance between parties. This includes misleading and aggressive practices such as false and bait advertising; false or deceptive information or omissions; and online search and reviews distortions (e.g., fake testimonials, awards or certifications).
  ↳ AI may be directly employed to carry out all such practices – e.g., use of bots or virtual entities that do not disclose their commercial intent or affiliation, or that influence consumers into taking certain actions; use of AI for personalised or discriminated offers without informing thereof.

  Certain practices are prohibited under the AI Act (see 🔍 **Prohibited Practices**). Two prohibited practices have a unique interplay with EU consumer law: using an AI system that deploys subliminal, purposefully manipulative or deceptive techniques which are significantly harmful and materially influence the behaviour of natural persons or group(s) of persons; and using an AI system that exploits vulnerabilities due to age, disability, or a specific socio-economic situation.

- **Transparent information**
  As parties suffering from information asymmetry, consumers should receive truthful, clear, and comprehensive information about products and services, including their price, quality, and safety.
  ↳ As AI may be used, for instance, to describe products, to draft information required to be made available, to rank products and to personalise prices, it must be deployed in line with consumer transparency requirements.

  The AI Act establishes transparency requirements for AI systems impacting individuals or generating deep fakes (see questions 36 and 37). Notably, use of AI without disclosure is generally prohibited under the AI Act and can also constitute a misleading omission in breach of consumer law.

- **Contract conclusion and execution**
  AI can be used to automate the drafting and conclusion of consumer contracts (including verifying identities, exchange of consent and payment information), execution of the contract, handling of consumer support, and monitoring of contractual compliance.

To whom the obligations apply

The obligations arising from consumer law apply to **traders**, which are natural or legal persons pursuing purposes relating to their trade, business, craft or profession.

In turn, the AI Act obligations apply to **AI deployers**. For an explanation of who they are, see question 29.

A single organisation, notably an AI deployer, fits the concept of trader that triggers most consumer law, in which case all the above obligations would apply.

AI Deployment Checklist — Consumer protection

Goal:

Assess whether AI deployment is done in compliance with consumer protection law, as applicable – relevant for all AI deployments.

Questions:

- How is AI used with consumers (e.g., advertising, information, pricing, eligibility, ranking, contract formation or performance) and for what purposes (e.g., persuasion, personalised offers, ID verification, consent)?
- Which consumers are involved, and do they include vulnerable groups such as children or the elderly?
- Is the use of AI made clear to consumers, and in what way?
- Where AI is used to conclude or perform contracts, can consumers intervene in the automated process and, if so, when and how?
- Which risks are anticipated (e.g., misleading or opaque use of AI, use of AI for manipulating consumers, unfair or discriminatory outcomes for consumers, barriers to exercising consumer rights such as contract termination, failure to recognise when AI deployment is prohibited)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Marketing; Legal; Compliance; Ethics; Data including Data Protection Officer (see question 42).

Digital delegates or assistants

Digital delegates or assistants are AI tools that can take over parts (or all) of the contracting process, from simple assistants that search or order on request, to more advanced agents that autonomously conclude and perform contracts, sometimes with self-learning capabilities. For consumers, delegates allow the outsourcing of decision-making to AI, saving time and effort and potentially improving choices. However, they also raise legal issues.

Key questions include how their actions are attributed to the consumer, compliance with consumer law rules, and protection of consumer rights. In principle, the delegate's acts are attributed to the consumer, who deploys it and intends to be bound by any contracts it concludes. However, this needs qualification to reflect AI autonomy and unforeseen behaviour. Due to this, additional rights and obligations for consumers and providers of AI delegates and assistants have been discussed, such as consumers' right to intervene before contract conclusion, to suspend or deactivate the assistant, or to cancel contracts in certain cases, as well as providers' duties to give clear information on the assistant's functions and limits.

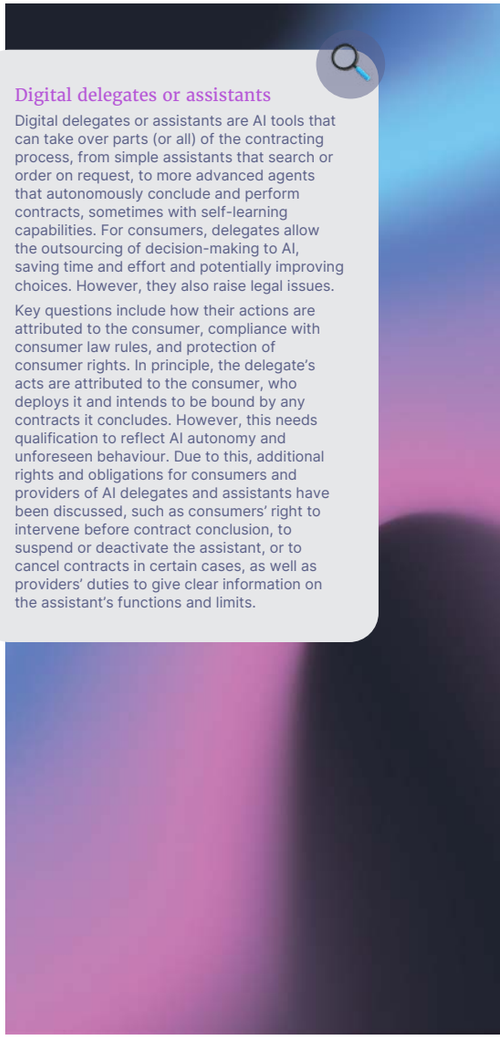

Risk mapping:

Consumer protection risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

Are there limits on organisations' use of AI to coordinate market conduct? – The case of algorithmic collusion

Significant advances in AI have raised concerns about the use of pricing algorithms, particularly those employing AI capabilities. The rising prevalence of such algorithms in monitoring and setting prices has sparked intense debate, as it has been argued that these algorithms foster collusion, which is typically forbidden. This phenomenon has been called "algorithmic collusion", which refers to the use of algorithms to facilitate anti-competitive behaviour among undertakings. There are three types of algorithmic collusion: using algorithms to implement anti-competitive agreements, collusion via hub-and-spoke, and autonomous algorithmic collusion.

- **Using algorithms to implement anti-competitive agreements** - algorithms can monitor real-time market conditions and competitors' prices, recommending prices based on a number of factors and detecting deviations from agreed prices almost instantaneously, thus allowing for more effective monitoring. These characteristics enable undertakings to use algorithms to coordinate their behaviour on essential parameters of competition, such as price, or as a tool to monitor compliance with a previously established agreement.
  - ↳ Using AI to implement anti-competitive agreements may raise challenges to the extent that it allows reducing or eliminating interactions between undertakings. Indeed, competition authorities often rely on interactions between the targeted undertakings to prove the existence and duration of the infringement.
- **Collusion via hub-and-spoke** – hub-and-spoke algorithmic collusion occurs when a third party provides several competing undertakings with the same pricing algorithm, or with algorithms whose pricing strategy principles are somehow coordinated.
  - ↳ The use of AI in this circumstance does not raise particular challenges in the application of competition law: it may infringe competition law, particularly when undertakings intend or accept that their information be shared by the third supplier with their competitors, or when such a scenario is foreseeable.
- **Autonomous algorithmic collusion** – this scenario pertains to cases where algorithms autonomously reach a collusive result, without having been programmed or instructed to coordinate their behaviour and without any form of prior communication or contact between undertakings. Some studies have shown that AI-using algorithms, particularly reinforcement learning algorithms and Q-learning algorithms (a subtype of reinforcement learning algorithms), are capable of learning collusive strategies when programmed to achieve the optimal strategy, particularly for maximising profits.
  - ↳ Since the algorithms coordinate their market behaviour independently, without any agreement or communication, it can be quite difficult to establish an infringement of competition law.

To whom the obligations apply

The rules detailed above apply to **undertakings**, i.e., any entity engaged in an economic activity regardless of its legal status and the way in which it is financed. However, competition rules do not apply to activities which, by their nature, their aim and the rules to which they are subject, do not fall under the scope of economic activity or which are connected with the exercise of public authority powers.

AI deployers under the AI Act using AI for the practices described above would in principle be considered an undertaking and would thus be subject to competition rules. For an explanation of the concept of deployer, see question 29.

AI Deployment Checklist — Algorithmic collusion

Goal:

Assess whether AI deployment is done in a manner that facilitates anti-competitive behaviour – relevant for all AI deployments.

Questions:

- Is AI being deployed to influence prices or other key commercial terms in markets with competitors?
- What markets are at stake?
- Does the AI system have features that create or reinforce coordinated outcomes (e.g., functions, data sources)?
- Does the AI system use market conditions and/or competitors' prices or strategies as inputs?
- Are data or common vendors/tools being shared in a way that could lead to the coordination of parallel strategies?
- Which risks are anticipated (e.g., reliance on dynamic pricing tools without assessing risks of parallel or coordinated behaviour, common vendor or data supplier acting as a coordination hub, no monitoring of pricing/strategy outcomes for collusive patterns)?
- Which mitigation measures will be implemented to address the risks?

Expertise required:

Technology & Systems; Legal; Compliance.

Risk mapping:

Competition risks	Description of the impact	Likelihood	Severity	Risk score	Mitigation/elimination measures

See the explanation on the Checklist and Risk Mapping above.

The use of AI to exclude competitors

AI can enable undertakings with significant market power to exclude rivals, distort competition and entrench their position, especially given network effects and economies of scale in AI markets that can further strengthen the position of early movers and large digital platforms. Dominant platforms may, for example, program algorithms to favour their own products (algorithmic self-preferencing), bundle core services in ways that force users to hand over data as a condition of access (data bundling and lock-in) or exploit extensive consumer data to conduct sophisticated price discrimination.

Such behaviour can constitute abuse of dominance under competition law where it affects trade between Member States, with dominance in AI markets often rooted in the control of essential inputs such as data, compute or proprietary models.

The Digital Markets Act (DMA) adds *ex ante* obligations for designated gatekeepers (large online platforms that provide core platform services, such as Amazon, Apple, Meta or Microsoft) to proactively prevent gatekeepers' anti-competitive behaviour, including a ban on treating their own services more favourably in rankings than comparable third-party offerings.

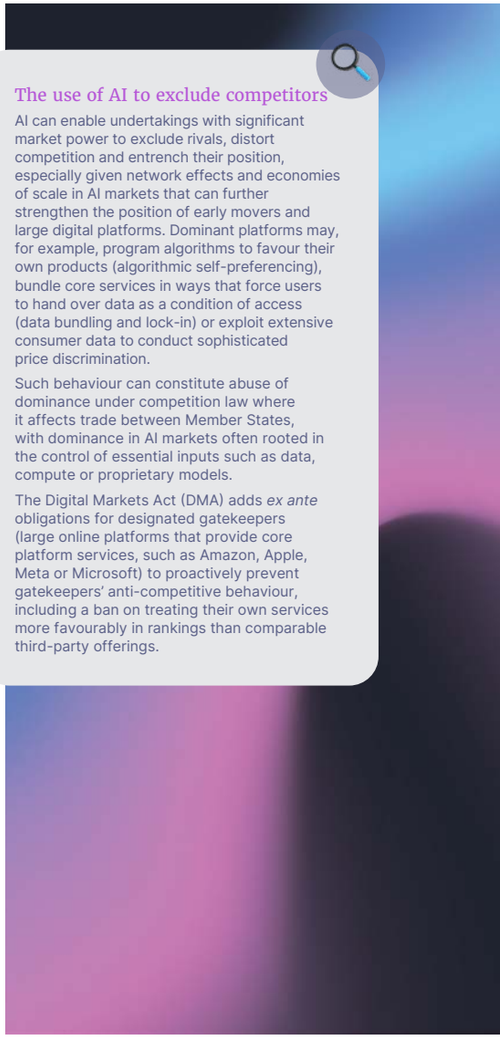

V.

Measures for AI innovation and uptake

Regulatory sandboxes

How can you use regulatory sandboxes to advance your AI activity?

Regulatory sandboxes are controlled frameworks established by competent authorities to allow AI providers or prospective providers to develop, train, test and validate innovative AI systems, where appropriate in real-world conditions, under regulatory supervision. Their primary objectives are to build legal certainty, facilitate early compliance with applicable EU requirements, and accelerate access to the EU market. Each Member State shall ensure that at least one AI national regulatory sandbox is operational by 2nd August 2027.

↪ A European Innovation Act is envisaged to promote access to research and technology infrastructures and to intellectual assets generated by publicly funded R&I. This Act will also contain a common legal definition and basic principles regarding the establishment of regulatory sandboxes, including cross-border or place-based regulatory sandboxes, while ensuring sector-specific needs.

To the extent other regulatory sandboxes exist, Member States must ensure an appropriate level of cooperation between the authorities supervising those other sandboxes and the AI national competent authorities.

↪ For instance, the Cyber Resilience Act (see question 3) addresses the establishment of cyber resilience regulatory sandboxes, clarifying that manufacturers of products with digital elements that are high-risk systems may participate in AI sandboxes.

In Portugal, technological free zones (ZLTs) may be set up to test, in a real or quasi-real environment, innovative technologies, products, services and processes that may span multiple areas or sectors. ZLTs are physical environments, geographically located, used for testing, with the intervention and supervision of competent entities.

Testing in **real-world conditions outside AI sandboxes** is also allowed by the AI Act for most high-risk AI systems and subject to a set of obligations, including a real-world testing plan, free and informed consent of the test subjects, and the possibility to reverse and disregard the output of the AI system.

In addition, under the new updated version of the AI Act, Member States may permit such testing for AI enabled transport products (in aviation, road transport, maritime and agricultural transport) and machinery, through the adoption of frameworks for real-world testing. These frameworks must address, among others, real-world testing plan and free and informed consent of the test subjects, as well as effective governance and accountability arrangements and high level of protection of health, safety and fundamental rights.

A **key point** in this regard is that the AI Act expressly allows the processing of personal data collected for other purposes in a regulatory sandbox for the purpose of developing, training and testing certain AI systems in the public interest.

Who can benefit

Regulatory sandboxes particularly benefit **AI providers**, regardless of the AI system being high-risk or not. For an explanation of the concept of AI provider, see question 1.

Testing outside AI sandboxes is foreseen only for AI providers (established in the EU or with a legal representative in the EU) of high-risk AI systems. See **High-Risk Systems**.

Measures for SMEs

Are there any AI-specific instruments that SMEs can benefit from when developing or deploying AI?

The AI Act requires Member States to undertake a set of actions to facilitate the development and uptake of AI by SMEs. These include providing SMEs, including start-ups, which have a registered office or branch in the Union, with priority (and free of charge) access to AI regulatory sandboxes; organising specific awareness-raising and training activities on the application of the AI Act; using existing dedicated channels and, where appropriate, establishing new ones for communication; and facilitating the participation of SMEs and other relevant stakeholders in the standardisation development process.

In addition, the AI Act also provides for:

- **Simplified technical documentation for SMEs and SMCs**: SMEs and SMCs that are providers of high-risk AI systems may submit technical documentation (see question 27) in a simplified form, using Commission-developed templates.
- **Simplified quality management system for SMEs**: SMEs, including start-ups, may comply with quality management system requirements (see question 17) through simplified procedures, reflecting their organisational constraints. The Commission will issue guidelines to operationalise this assistance and ensure proportional application.

Other laws establish other specific measures for SMEs. For example, the **Cyber Resilience Act** (see question 3) requires Member States to help micro and small enterprises (including startups) comply by offering awarenessraising and training, dedicated communication channels and support for testing and conformity assessment, which can be aligned with AI Act support tools. Similarly, the **Ecodesign for Sustainable Products Regulation** (see question 4) requires Member States to assist SMEs through at least onestop shops or similar mechanisms, and may also offer financial support, organisational and technical assistance, and training.

Who can benefit

The measures detailed above benefit **SMEs** and **SMCs** that are, with relevance for this Guide, AI providers and deployers under the AI Act, as applicable. For an explanation of AI providers and deployers, see questions 1 and 29.

Funding, financing and other measures

What main funding mechanisms are available for AI development and deployment?

Financing options include public, private and hybrid financing, and may cover both initial financing (upfront capital) and ongoing financing needs.

Public financing

Public financing includes grants, procurement-based mechanisms, and public development finance. These instruments are often competitive and administratively intensive but remain a key source of early-stage and higher risk funding.

- EU grants such as **Horizon Europe and Digital Europe Programme** (DIGITAL) are major sources of AI financing, with the AI Innovation Package providing dedicated funds for AI. They form the backbone of EU AI research and innovation investment, supporting the full lifecycle from fundamental research to testing, validation and early market uptake. These programmes enable collaborative R&I, technology transfer and deployment through multi-annual work programmes and co-funded actions.
  In Portugal, national grants (notably from the Agency for Investigation and Innovation (AI2) and the PT2030 programme) can also be considered for AI-related projects.
- Ongoing public funding (for example, performance-based contracts, service contracts, or long-term public procurement of innovation) may also be relevant, especially for AI with public or societal benefits.

↪ Public funding must be compatible with State aid rules. In many cases, Member States must notify planned aid to the European Commission, which assesses its compatibility, but several State aid "safe harbours" exist.
Under the **General Block Exemption Regulation** (GBER), certain categories of aid for research, development and innovation – including many AI-related projects – are deemed compatible with the internal market if they meet defined conditions (e.g., type of R&D, eligible costs, aid intensities) and comply with the relevant formalities.
Aid granted under **Important Projects of Common European Interest** (IPCEI) may also be found compatible, including where AI projects contribute significantly to EU competitiveness, resilience or strategic autonomy.

Private financing

Private financing covers market-based capital such as venture capital, angel investment, venture debt, bank loans, private equity, and crowdfunding, among others. These instruments typically back organisations with a credible business model, a clear path to revenues, or the potential for ownership of critical technology. Private capital often comes with expectations of equity, repayment or commercial returns, but offers flexibility, speed and scalability, helping bridge gaps not covered by public funding.

For ongoing financing, commercialisation and revenue models become the primary source of funding.

- Different revenue and pricing models may be adopted, such as transaction fees, subscriptions, freemium, tiered models, data monetisation models, among others (see question 25 relating to AI commercialisation).

Hybrid financing

Hybrid financing refers to blended arrangements combining public and private capital, such as blended finance instruments, co-investment funds, guarantees or subordinated debt structures. Hybrid approaches leverage the strengths of both sectors: public institutions de-risk early phases and crowd in private capital, while private investors contribute additional funding, expertise and market discipline. EU-level blended instruments and public-private partnerships (including those linked to AI factories and gigafactories under InvestAI) are examples of this approach.

Public measures to facilitate access to resources (such as compute, data and digital infrastructures) and support are equally relevant. These include:

- **AI factories and gigafactories –** large-scale compute infrastructures and ecosystems that provide access to supercomputing power, data resources and expertise, with gigafactories hosting very large GPU clusters to train and deploy advanced AI models under the InvestAI framework.
- **European cloud and edge capacity marketplaces** – cloud and central marketplace for AI tools, datasets and services, which is expected to be addressed in the forthcoming Cloud and AI Development Act (CADA).
- **AI Testing and Experimentation Facilities** (TEFs) – large-scale, specialised sites where technology providers can test AI-based solutions in real-world environments before bringing them to market.
- **European Digital Innovation Hubs** (EDIHs) – one-stop shops aimed at providing experimentation facilities, training, access to finance and support on innovation networking. They now have a reinforced AI focus and act as first-line helpdesks to guide SMEs, small mid-cap enterprises (SMCs) and public organisations on testing, deploying AI and navigating AI Act-related questions.
- **Regulatory sandboxes** – instruments enabling supervised experimentation, where appropriate in real-world conditions, with legal certainty and regulatory guidance (see question 50).

Who can benefit

Beneficiaries of financing mechanisms depend on the financing mechanism at stake.

Public programmes such as Horizon Europe and Digital Europe typically target consortia, bringing together companies (including start-ups and SMEs), universities, research organisations and sometimes public bodies, with eligibility shaped by each call's scope and priorities. National funding in Portugal can support a broad range of Portuguese entities (from SMEs and larger companies to research centres, start-ups and municipalities), again depending on the call and policy focus.

Private instruments differ: venture capital and venture debt usually target higher growth companies (early-stage to scale-ups) with scalable business models and significant upside potential, while bank loans and private equity are more suited to later-stage entities with stable cash-flows, assets or proven market traction. Hybrid or blended finance schemes may combine these, often prioritising projects with strong innovation potential but higher risk profiles, where public participation helps crowd in private investors.

Annex 1 – Timelines

AI Act			Providers	Deployers
February 2, 2025	Prohibitions on most unacceptable AI practices		Prohibited Practices	
	Obligations on literacy		Question 20	Question 43
	Processing of special categories of personal data for bias detection and correction		Question 12	Question 38 (by reference to question 12)
August 2, 2025	Requirements for GPAI models		The Integration of AI Models in AI Systems	
	Penalties/fines		Question 21	Question 44
August 2, 2026	Transparency obligations		Question 8 Question 9	Question 36 Question 37
	Standards and presumption of conformity	Although this applies in law from 2 August 2026, their practical application depends on the entry into force of the high-risk AI system requirements, which has been deferred to 2027 and 2028.	Question 22	Question 45
	Conformity assessment		Question 19	
	EU DoC and CE marking		Question 26	
	Post-market monitoring		Question 17	
	Reporting of serious incidents (by providers)		Question 21	
	Corrective actions requested by market surveillance authorities	Although this applies in law from 2 August 2026, practical application due to non-compliance for high-risk systems depends on the entry into force of the high-risk AI system requirements, which has been deferred to 2027 and 2028	Question 21	Question 44
	Registration		Registration of AI Systems and Products	
	Regulatory sandboxes		Question 50	

AI Act			Providers	Deployers
August 2, 2026	Measures for SMEs	In practice, simplified technical documentation and quality management system depends on the entry into force of the high risk AI system requirements, which has been deferred to 2027 and 2028.	Question 51	
December 2, 2026	Prohibitions on the new unacceptable AI practices – generation or manipulation of realistic sexual images of identifiable persons without consent and child pornography material		Prohibited Practices	
December 2, 2027	High-risk system obligations (Annex III of AI Act)		High-Risk Systems	
			Question 2	Question 30
August 2, 2028	High-risk system obligations of regulated products (Annex I of AI Act) • **Note:** excluding machinery and transport products (in aviation, road transport, maritime and agricultural transport), as in this case only limited AI Act provisions apply directly • **Note:** for machinery, the new version of the Machinery Regulation now requires delegated acts to be approved (to apply by August 2, 2028, as well) to add health and safety requirements for high-risk AI systems which shall ensure that (most) AI Act requirements are reflected. These requirements are those in Questions 2, 3, 6, 7, 13, 17, 27 and 21 (reporting of serious incidents)		Question 3 Question 6 Question 7 Question 10 Question 13 Question 15 Question 17 Question 18 Question 21 (corrective measures) Question 24 Question 26 (obligations on name and trademark) Question 27	Question 31 Question 34 Question 35 Question 39 Question 36 (provision of information on the use of AI make decisions related to individuals) Question 41 Question 42 Question 44 (reporting, suspension and information measures) Question 47

Other Laws		Providers	Deployers
Machinery Regulation	• **January 20, 2027** – for most provisions • **August 2, 2028** – for the delegated act adding health and safety requirements for high-risk AI systems (see table above on this)	Question 2 Question 6 Question 7 Question 19 Question 21 Question 26 Question 27	
General Product Safety Regulation (GPSR)	• **December 13, 2024**	Question 2 Question 21 Question 26 Question 27	
Cyber Resilience Act (CRA)	• **December 11, 2027** – for most provisions	Question 3 Question 7 Question 15 Question 19 Question 21 Question 22 Question 26 Question 27	
	• **September 11, 2026 –** for reporting of vulnerabilities and incidents	Question 21	
Ecodesign for Sustainable Products Regulation (ESPR)	Concrete ecodesign requirements will apply only once the Commission adopts product-specific delegated act and their application dates kicks in. Delegated acts are still under development at the time of writing	Question 4 Question 7 Question 15 Question 19 Question 21 Question 22 Question 26 Question 27	

Other Laws		Providers	Deployers
General Data Protection Regulation (GDPR)	• **May 25, 2018** Amendments from the Digital Omnibus proposal will apply once it is published and its application date kicks in.	Question 5 Question 11 Question 12 Question 13 Question 14 Question 16 Question 17 Question 18 Question 19 Question 21 Question 22 Question 28	Question 31 Question 33 Question 34 Question 35 Question 36 Question 38 Question 39 Question 40 Question 41 Question 42 Question 44 Question 45
Digital Services Act (DSA)	• **February 17, 2024** – for most provisions	Question 8	Question 36 Question 37
European Accessibility Act (EEA)	• **June 28, 2025** – it must be fully applied by Member States (States had until June 28, 2022, to transpose it)	Question 10	
Web Accessibility Directive (WAD)	• **June 23, 2021** – from this date, all public-sector websites and mobile apps had to comply (States had until September 23, 2018) to transpose it	Question 10	
IP framework	• The IP framework comprises several sets of acts (e.g., Database Directive, Digital Single Market (DSM) Directive, Collective Rights Management Directive, European Patent Convention), all of which are already in full application	Question 11 Question 12 Question 13 Question 14 Question 16 Question 23	Question 38 Question 39 Question 40 Question 46

Other Laws		Providers	Deployers
Data Act	• **September 12, 2025** – for most provisions • **September 12, 2026** (with respect to the design of connected products) • **September 12, 2027** – for terms of B2B contracts concluded on or before September 12, 2025, that are of indefinite duration or due to expire at least 10 years from January 11, 2024 Amendments from the Digital Omnibus proposal will apply once it is published and its application date kicks in	Question 1 (with respect to the connected products) Question 12 Question 14 (for terms of B2B contracts) Question 15 Question 28 The EU Approach to Facilitate Data Sharing The Contribution of the EU Data Legislation to Data Quality	Question 38 Question 40
Data Governance Act (DGA)	• **September 24, 2025** Repeal by the Digital Omnibus proposal will apply once it is published and its application date kicks in	Question 14 The EU Approach to Facilitate Data Sharing	Question 40
Open Data Directive	• **September 24, 2023** – it must be fully applied by Member States (States had until July 17, 2021, to transpose it). Repeal by the Digital Omnibus proposal will apply once it is published and its application date kicks in	The Contribution of the EU Data Legislation to Data Quality	
Digital Markets Act (DMA)	• **May 2, 2023** – for most provisions	Question 28	
NIS 2 Directive	• **October 18, 2024** – it must be fully applied by Member States (States had until October 17, 2024, to transpose it) Amendments from the Digital Omnibus proposal will apply once it is published and its application date kicks in		Question 31 Question 41 Question 42 Question 44 Question 45

Other Laws		Providers	Deployers
CER Directive	• **October 18, 2024** – it must be fully applied by Member States (States had until October 17, 2024, to transpose it; and must identify critical entities by July 17, 2026) Amendments from the Digital Omnibus proposal will apply once it is published and its application date kicks in		Question 31 Question 41 Question 42 Question 44
Energy Efficiency Directive (EED) (recast)	• **October 11, 2025** (for most provisions) – it must be fully applied by Member States (States had until this data to transpose it)		Question 32 Question 41
Consumer framework	• The consumer framework comprises several sets of acts (e.g., the Consumer Rights Directive (CRD), the Unfair Commercial Practices Directive (UCPD) and the Unfair Contract Terms Directive (UCTD)), all of which are already in full application (subject to ongoing targeted amendments)	The Provision of AI Systems to Consumers	Question 48
Competition framework	• The competition framework comprises several sets of acts (notably, as relevant herein, the rules under the Treaty on the Functioning of the European Union (TFEU), the R&D Block Exemption Regulation (R&D BER) and the Technology Transfer Block Exemption Regulation (TTBER)), all of which are already in full application (subject to limited transitional arrangements for certain existing agreements)	The Risk of Anti-Competitive Effects of Sourcing Contracts and Research Partnerships The Risk of Anti-Competitive Effects of Data Sharing Agreements	Question 49

Annex 2 – AI Act Governance structure

EU Level					National Level	
					National competent authorities	
					National competent authorities shall be designated by Member States, at least one of each of the following:	
European Commission | AI Office	**AI Board**	**Advisory Forum**	**Scientific Panel**	**Union AI testing support structures**	**Notifying Authority**	**Market surveillance authority**
The AI Office is an administrative structure within the Commission with competences on the implementation and supervision of the AI Act.	The AI Board is an advisory body composed of representatives of Member States, advising the Commission and Member States on the AI Act.	The Advisory Forum is composed of stakeholders from industry, academia, think thanks and civil society, providing technical expertise to the Commission and the AI Board.	The Scientific Panel is a panel of experts providing advice and support to the Commission | AI Office, and national market surveillance authorities, on the enforcement of the AI Act.	These are testing facilities that carry out testing of AI systems at the request of, notably, market surveillance authorities or the Commission, and also provide ad-vice to them and the Board.	National authority responsible for setting up and carrying out the necessary procedures for the assessment, designation and notification of conformity assessment bodies.	It is a national authority carrying out market surveillance activities and measures.

Supervision	
GPAI models	• Commission | AI Office
AI systems	• National market surveillance authorities – in most cases • AI Office – AI systems based on general purpose AI models where the model and system share the same provider (subject to specified exceptions), and AI systems that constitute or are integrated into very large online platforms or very large online search engines
Regulatory sandboxes	• National competent authorities – in principle, those establishing the regulatory sandbox, but also others that may be relevant in light of the activities in the sandbox • AI Office – for sandbox established by it
Testing in real world conditions	• Market surveillance authorities

Contacts

Magda Cocco

Group Executive Partner,
Regulatory & Digital, **VdA**
mpc@vda.pt

Helena Correia Mendonça

Principal Consultant,
Information, Communication & Technology, **VdA**
hcm@vda.pt

Iakovina Kindylidi

Senior International Adviser,
Information, Communication & Technology, **VdA**
imk@vda.pt

Center for Responsible AI

VIEIRA DE ALMEIDA

PRR Plano de Recuperação e Resiliência

REPÚBLICA PORTUGUESA

Financiado pela União Europeia NextGenerationEU