

# Artificial intelligence

Field 1 · EPSO/AD/430/26 · grade AD 8

**240**

successful candidates sought

**AD 8**

entry grade

**30 / 40 min**

field-related test

**13 Oct 2026**

application deadline

[eu-now.com](https://eu-now.com)

Preparation dossier · 15 study blocks with their official sources

Based on notice of competition OJ C/2026/4668 and the official texts that cover its Annex II

[EU-now.com](https://eu-now.com)

## 1 · The competition at a glance

EPSO/AD/430/26 is an open competition based on qualifications and tests for **administrators at grade AD 8** in the institutions, bodies, offices and agencies of the Union. Two fields, and **you may apply for only one**: artificial intelligence (240 successful candidates sought) and cybersecurity (254). You are going for artificial intelligence: 240 places on the reserve list.



**What actually decides:** the three reasoning tests and the EUFTE essay are pass-or-out gates. **Only the field-related MCQ score ranks candidates**, and only the highest scorers have their EUFTE marked and their eligibility checked. Two thirds of your study time belong to the field.



The five phases. Only the field MCQ score ranks candidates; everything else is a pass-or-out gate.

## 2 · The five tests


| Test                     | Language | Questions | Time          | Points      | Pass mark                                  |
|--------------------------|----------|-----------|---------------|-------------|--------------------------------------------|
| Verbal reasoning         | 1        | 20        | 35 min        | 0-20        | <b>10/20</b>                               |
| Numerical reasoning      | 1        | 10        | 20 min        | 0-10        | 10/20 combined with abstract               |
| Abstract reasoning       | 1        | 10        | 10 min        | 0-10        | 10/20 combined with numerical              |
| <b>Field-related MCQ</b> | <b>2</b> | <b>30</b> | <b>40 min</b> | <b>0-30</b> | <b>15/30 and among the highest scorers</b> |
| EUFTE                    | 2        | —         | 40 min        | 0-10        | 5/10                                       |

**Language 1** is any of the 24 official languages. **Language 2** is the language of the field MCQ and of the EUFTE. You need C1 in one and B2 in the other, and you choose them in the application form. The tests are taken remotely with remote proctoring, and the order of the sittings may differ from the order in the notice.

The EUFTE is neither a language test nor a knowledge test: EPSO publishes documentation on EU matters before the test — the same you receive during it — and you answer an assignment based on it. It is marked against published anchors.

### The exam platform

Tests run on TAO, the platform EPSO uses for all its exams. Arrive knowing where everything is: you lose less time and nothing surprises you on the day.



Highlighter and line reader for long stems.

Calculator and scratchpad where the test allows them.

Screenshots of the TAO platform used in EPSO tests.

The reserve list is published in alphabetical order, not by score; everyone tying for the last place goes on it; and being on it guarantees no recruitment.

Sources: notice of competition OJ C/2026/4668, sections 1, 4.2, 4.3.1, 4.3.2 (tables 3 to 5), 4.3.3 and 4.3.4.

### 3 · What the 30 questions look like

The notice publishes no syllabus and no reading list. The contract is **Annex II, «typical duties»**: everything examinable comes from there, which is why this dossier breaks it down in full. Thirty questions in 40 minutes, in language 2: **80 seconds per question**. There is no time to derive anything from scratch — you answer with what you recognise.

A · Notify the CSIRT within 72 hours

B · Early warning to the CSIRT within 24 hours

C · Final report to the Commission within a month

D · Notify the Cooperation Group

## What decides the answer

- 1 **WHO is competent**  
CSIRT, authority, Commission...
- 2 **WHICH exact duty**  
early warning, notification, report
- 3 **THE NUMBER**  
24 h · 72 h · one month

Distractors attack one of the three:  
the neighbouring deadline, the adjacent  
the duty from another phase.

Anatomy of a field question: 80 seconds to recognise three facts.

## The five traps of the format

- **The longest option.** In official papers the key is the longest option only about 4 times out of 10: picking by length does not pass.
- **The absolute.** «Always», «never», «in all cases» usually mark the false option: EU law lives on exceptions.
- **The swapped actor.** The duty is described correctly but hung on the wrong body. Read the subject before the verb.
- **The neighbouring instrument.** The statement is true, but it belongs to another law that does not govern this case.
- **Superseded law.** Wording from before an amendment. Always study the consolidated text.

## The fact sheet that works

| Who it applies to   | What it requires  | Deadlines and figures   | Who supervises    |
|---------------------|-------------------|-------------------------|-------------------|
| entities · products | the exact conduct | 24 h · 72 h · threshold | and who sanctions |

One fact sheet per instrument: four columns, nothing more. That is what answers a question in 80 seconds.

## 4 · The syllabus, block by block

Annex II describes duties, not subjects. Here they are turned into study blocks: what is in, what gets confused and where it is written. What is not in an official text is not examined.



Field 1 examines the whole cycle: designing, deploying and governing AI systems inside an institution.

## 1 · Design, building and operation of AI solutions

Scope and definition

High-risk requirements

General-purpose models

Lifecycle and MLOps

Generative AI: RAG and agents

AI security

## 2 · Compliant, defensible and scalable use

Prohibited practices and risk

Transparency and literacy

Governance and timeline

Trustworthy AI and ALTAI

Risk management (AI RMF)

Personal data

Data governance

## 3 · Policymaking and policy enforcement

EU AI policy

Product liability

Annex II turned into study blocks.

### 1 Scope of the AI Act and what counts as an AI system

Before any duty there is a prior question: is this an AI system, and does the Regulation apply to it? Several questions per exam start here.

#### What is in

- The Article 3(1) definition read with Guidelines C(2025) 5053: machine-based system, autonomy, adaptiveness, explicit or implicit objectives, inference, and outputs that influence physical or virtual environments.
- What Article 2 leaves out: national security and military use, scientific research and development, pre-market testing, personal non-professional use, and free and open-source software except for prohibitions and specific duties.
- The roles and their definitions: provider, deployer, importer, distributor and authorised representative.
- Article 25 and when a deployer becomes a provider: own branding, substantial modification, or a change of intended purpose.
- The Article 113 timeline as amended by Regulation (EU) 2026/1744.

#### What gets confused

- Systems that only optimise mathematically or apply fixed rules do not infer in the sense of Article 3(1): the Guidelines say so expressly.
- Open source does not mean exempt: the Article 5 prohibitions still apply.

**Sources:** Regulation (EU) 2024/1689, Articles 2, 3 and 25 · Guidelines C(2025) 5053 · Regulation (EU) 2026/1744

## 2 Prohibited practices and risk classification

The exam does not ask for the list: it asks whether a case falls inside it, and an element is almost always missing.

### What is in

- The Article 5 practices with their cumulative elements: manipulation with material distortion of behaviour and significant harm, exploitation of vulnerabilities, social scoring, crime prediction based on profiling, untargeted facial scraping, emotion recognition at work and in education, biometric categorisation of sensitive data, and real-time remote biometric identification for law enforcement.
- The exceptions for real-time biometric identification and the prior judicial or administrative authorisation.
- The two routes to high risk: product or safety component covered by Annex I (Article 6(1)) and Annex III use cases (Article 6(2)).
- The Article 6(3) derogation: preparatory task, improvement of a previous human activity, pattern or deviation detection without replacing human assessment — plus the duty to document the assessment.
- Guidelines C(2025) 5052 and the 2026 draft guidelines on high-risk classification.

### What gets confused

- Emotion recognition is prohibited only at work and in education, with medical and safety exceptions.
- Being in Annex III is not enough: Article 6(3) can take the system out of the high-risk tier.

**Sources:** Regulation (EU) 2024/1689, Articles 5 and 6, Annexes I and III · Guidelines C(2025) 5052

## 3 High-risk requirements and deployer duties

The densest block and the one where who-does-what is most often mixed up.

### What is in

- Provider requirements, article by article: risk management (9), data governance (10), technical documentation (11 and Annex IV), automatic logging (12), transparency and instructions for use (13), human oversight (14), accuracy, robustness and cybersecurity (15).
- Quality management system (Article 17) and documentation keeping (Article 18).
- Deployer duties (Article 26): use according to the instructions, assign human oversight to people with competence and authority, keep the logs, monitor operation and suspend use where there is risk.
- Fundamental rights impact assessment (Article 27): who owes it, when and what it contains.
- Post-market monitoring (Article 72) and serious-incident reporting (Article 73).

### What gets confused

- The fundamental rights impact assessment belongs to the deployer, not the provider.
- Article 12 logs are generated by design; keeping them is the deployer's duty under Article 19.

**Sources:** Regulation (EU) 2024/1689, Chapter III, Sections 2 and 3; Articles 72 and 73; Annexes IV and VII

#### 4 General-purpose AI models and the Code of Practice

The newest material and, for that reason, the most heavily examined in recent competitions.

##### What is in

- Duties of every GPAI provider (Article 53): technical documentation, information to downstream providers, copyright compliance policy and a sufficiently detailed public summary of training content.
- Additional duties for models with systemic risk (Article 55): evaluation with standardised protocols, adversarial testing, risk mitigation, serious-incident reporting and cybersecurity protection.
- The systemic-risk presumption above  $10^{25}$  cumulative FLOP and notification to the Commission within two weeks (Articles 51 and 52).
- The 2025 Code of Practice and its transparency, copyright and safety chapters: what signing it means and what it does not.
- The free and open-source exemption and how far it goes.

##### What gets confused

- Signing the Code does not replace the legal duty: it is a way to demonstrate compliance.
- Systemic risk adds duties; it never removes the Article 53 ones.

**Sources:** Regulation (EU) 2024/1689, Chapter V · Guidelines C(2025) 7719 · GPAI Code of Practice 2025

#### 5 Article 50 transparency and AI literacy

Four duties with different addressees: the classic error is hanging them on the wrong one.

##### What is in

- Provider: systems interacting with people must say so (50(1)); synthetic output is marked in machine-readable form (50(2)).
- Deployer: discloses emotion recognition and biometric categorisation (50(3)); reveals deepfakes and AI-generated text of public interest (50(4)).
- The exceptions read in Guidelines C(2026) 5054 and the Code of Practice on marking generated content.
- Article 4 after the Omnibus: providers and deployers take measures that support AI literacy of their staff; the Commission and Member States support them.

##### What gets confused

- The old wording of Article 4 («ensure a sufficient level») is no longer in force: today it is a distractor.
- Marking synthetic output is the provider's job; revealing the deepfake is the deployer's.

**Sources:** Regulation (EU) 2024/1689, Articles 4 and 50 · Guidelines C(2026) 5054 · Regulation (EU) 2026/1744

## 6 Governance, timeline and the Digital Omnibus

Who supervises and since when. This is where the actor questions are decided.

### What is in

- AI Office: supervises general-purpose AI models and enforces Chapter V.
- National market surveillance authorities: high-risk systems; the EDPS for Union institutions.
- AI Board, Scientific Panel and Advisory Forum: coordination and advice, with no supervisory powers.
- The application calendar with the dates moved by Regulation (EU) 2026/1744 and the new Articles 75a to 75d on operators of the systems listed in Article 75(1).
- Penalties: the amounts of Article 99 and the separate regime for GPAI providers (Article 101).

### What gets confused

- Articles 75a to 75d do not apply to general-purpose AI model providers: Articles 88 to 94 and 101 do.
- The AI Board supervises nobody: it coordinates.

**Sources:** Regulation (EU) 2024/1689, Chapters VII and XII · Regulation (EU) 2026/1744

## 7 Trustworthy AI — Ethics Guidelines and ALTAI

Closed lists, easy to ask about and easy to confuse with each other.

### What is in

- Three components of trustworthy AI: lawful, ethical and robust.
- Four ethical principles: respect for human autonomy, prevention of harm, fairness and explicability.
- Seven key requirements, from human agency and oversight to accountability.
- ALTAI as a voluntary self-assessment list and its traceability questions.
- Oversight modes: human-in-the-loop, on-the-loop and in-command.

### What gets confused

- The four principles are not the seven requirements: distractors mix both lists.
- Robustness is a component, not one of the four principles.

**Sources:** Ethics Guidelines for Trustworthy AI (2019) · ALTAI (2020)

## 8 AI risk management frameworks — AI RMF, generative profile and bias

Voluntary guidance that the exam uses as a contrast with legal obligation.

### What is in

- The four AI RMF 1.0 functions: Govern, Map, Measure and Manage, and what each produces.
- The seven trustworthiness characteristics and the lifecycle stages: plan and design, collect and process data, build and use the model, verify and validate, deploy and use, operate and monitor.
- Generative-AI risks in AI 600-1: confabulation, information integrity, information security, value chain, human-AI configuration, data privacy and harmful bias.
- The three types of bias in SP 1270: systemic, statistical-computational and human, and why more data alone does not fix them.

### What gets confused

- AI RMF functions are not the cybersecurity CSF functions.
- AI 600-1 actions are suggested: presenting them as mandatory is the typical mistake.

Sources: NIST AI 100-1 · NIST AI 600-1 · NIST SP 1270

## 9 AI security — adversarial attacks and LLM application risks

The block that connects with cybersecurity, and the one growing fastest.

### What is in

- The NIST AI 100-2 (2025) taxonomy: for predictive AI, evasion, poisoning and privacy attacks; for generative AI, supply chain, direct and indirect prompt injection and misuse.
- Classification by attacker goal: availability, integrity, privacy and abuse.
- The OWASP Top 10 for LLM Applications 2025 with a matching mitigation for each entry.
- The three layers of the ENISA framework where the controls sit.

### What gets confused

- Direct injection is not indirect: what changes is who introduces the malicious text.
- Poisoning is not evasion: one acts at training time, the other at inference.

Sources: NIST AI 100-2 · OWASP Top 10 for LLM Applications 2025 · ENISA multilayer framework

## 10 Generative AI in the organisation — RAG, agents and guardrails

Architecture decisions with control consequences.

### What is in

- When prompting is enough, when retrieval augmentation is needed, when fine-tuning, and when agents with tools.
- The control each pattern demands: grounding and provenance, output handling, tool permissions and approval gates, pre-deployment evaluation and testing, and human oversight.
- What the EDPS requires of an EU institution before it uses generative AI.

- Decision logging and traceability of outputs used as evidence.

### What gets confused

- Retrieval augmentation does not fix bias: it reduces invention over your own data.
- An agent without an authorisation boundary is a security risk, not a productivity gain.

**Sources:** NIST AI 600-1 · OWASP LLM Top 10 2025 · EDPS orientations (2025)

## 11 Lifecycle, MLOps and post-market monitoring

The engineering part of Annex II, with the artefacts the law requires.

### What is in

- The lifecycle sequence and the evidence each stage leaves behind.
- Artefacts that keep a high-risk system auditable: automatic logs (Article 12) with the Article 19 retention, quality management system (17), post-market monitoring plan (72) and serious-incident reports (73).
- Secure development practices for generative AI (SP 800-218A) and how they fit the SSDF.
- Cloud service and deployment models (SP 800-145) and shared responsibility (SP 800-144).

### What gets confused

- Deployment does not close the cycle: post-market monitoring is a continuing provider duty.
- Outsourcing to the cloud does not outsource the deployer's responsibility.

**Sources:** Regulation (EU) 2024/1689, Articles 12, 17, 19, 72 and 73 · NIST SP 800-218A · SP 800-144 and 800-145

## 12 Personal data in AI

Where the AI Act meets data protection — with a different law for the EU institutions.

### What is in

- EDPB Opinion 28/2024: when a model may be considered anonymous and what has to be demonstrated.
- The three-step legitimate-interest test: legitimate interest, necessity and balancing.
- Consequences of unlawful model development for its later use.
- EU institutions are governed by Regulation (EU) 2018/1725, which has no legitimate-interest basis; the role of the data protection officer and of the impact assessment.
- Special categories for bias detection: after the Omnibus the gateway is the new Article 4a, not Article 10(5).

### What gets confused

- Citing GDPR articles for a Union institution is a mistake about which law applies, not a detail.
- A model's anonymity is not presumed: it is justified.

**Sources:** EDPB Opinion 28/2024 · Regulation (EU) 2018/1725 · EDPS orientations 2025 · Regulation (EU) 2026/1744

### 13 Data governance for AI

Annex II calls it data readiness; three instruments meet here and must not be mixed.

#### What is in

- Article 10 of the AI Act: governance practices for training, validation and testing data; design choices, collection, examination of bias and detection of gaps.
- Data Act 2023/2854: connected-product data, business-to-government access and switching between processing services.
- Data Governance Act 2022/868: re-use of protected public-sector data, data intermediation services and data altruism.
- What to require contractually from a training-data supplier.

#### What gets confused

- The Data Act is not a data protection law: it governs access and use.
- Article 10 does not require perfect data: it requires relevance, representativeness and, as far as possible, freedom from errors.

**Sources:** Regulation (EU) 2024/1689, Article 10 · Regulation (EU) 2023/2854 · Regulation (EU) 2022/868

### 14 Product safety and liability

A short block with its own date and practical consequences.

#### What is in

- Directive (EU) 2024/2853 treats software and AI systems as products (Article 4) and excludes free and open-source software supplied outside a commercial activity (Article 2(2)).
- Who is liable along the chain (Article 8) and what happens on substantial modification outside the manufacturer's control.
- The defectiveness test, including learning after the product is placed on the market (Article 7).
- Disclosure of evidence and presumptions (Articles 9 and 10); the development-risk defence and its derogation (Articles 11 and 18); limitation and expiry periods (16 and 17).
- Application date: 9 December 2026.

#### What gets confused

- Complying with the AI Act does not exclude civil liability for a defective product.
- The open-source exclusion does not cover software supplied in the course of a commercial activity.

**Sources:** Directive (EU) 2024/2853

## 15 EU AI policy

Attribution questions: which initiative belongs to which instrument.

### What is in

- AI Factories under the amended EuroHPC Regulation (2024/1732).
- AI Gigafactories and InvestAI in the AI Continent Action Plan (COM(2025) 165) and its five areas.
- The 'AI first' policy and sectoral flagships of the Apply AI Strategy (COM(2025) 723).
- RAISE in the AI-in-Science strategy (COM(2025) 724).
- The 2020 White Paper (ecosystems of excellence and of trust) and the 2021 Coordinated Plan review.
- The Interoperable Europe Act (2024/903): interoperability assessment and the Interoperable Europe Board.

### What gets confused

- AI Factories do not come from the AI Act: they come from EuroHPC.
- The AI Continent Action Plan and the Apply AI Strategy are different documents with different aims.

**Sources:** COM(2020) 65 · COM(2021) 205 · COM(2025) 165, 723 and 724 · Regulations (EU) 2024/1732 and 2024/903

## 5 · How to study it in eight weeks



Eight weeks allow two passes over the syllabus if you answer questions from day one.

| week                    | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|-------------------------|---|---|---|---|---|---|---|---|
| Map and fact sheets     |   |   |   |   |   |   |   |   |
| Block-by-block study    |   |   |   |   |   |   |   |   |
| Weak blocks and figures |   |   |   |   |   |   |   |   |
| Timed mock tests        |   |   |   |   |   |   |   |   |
| Review and reasoning    |   |   |   |   |   |   |   |   |

**Split:** two thirds on the field, one third on reasoning. The field ranks you; reasoning only eliminates. And from week 3, answer questions the same day you study the block: reading without answering fixes nothing.

## 6 · Official library

Everything below is public and citable. No serious preparation needs more, and no commercial manual replaces the consolidated text.

| Document                                                             | What for                                                                             |
|----------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Regulation (EU) 2024/1689 — AI Act, consolidated text                | The core law; read it with the Omnibus amendments already in.                        |
| Regulation (EU) 2026/1744 — Digital Omnibus                          | Which articles changed and since when, including the new 4a and 75a-75d.             |
| Guidelines C(2025) 5053 · C(2025) 5052 · C(2026) 5054 · C(2025) 7719 | AI system definition, prohibited practices, Article 50 transparency and GPAI duties. |
| GPAI Code of Practice 2025                                           | Transparency, copyright and safety chapters.                                         |
| Ethics Guidelines for Trustworthy AI (2019) and ALTAI (2020)         | Components, principles, seven requirements and self-assessment.                      |
| NIST AI 100-1 · AI 600-1 · AI 100-2 · SP 1270                        | Risk management, generative AI, adversarial attacks and bias.                        |
| OWASP Top 10 for LLM Applications 2025                               | Risks of LLM applications, RAG and agents.                                           |
| ENISA — multilayer framework for the cybersecurity of AI (2023)      | How to order the security controls of an AI system.                                  |
| NIST SP 800-218A · SP 800-144 · SP 800-145 · SP 800-146              | Secure development with generative AI and cloud computing.                           |
| EDPB Opinion 28/2024 and EDPS orientations (2025)                    | Personal data in models and generative AI in EU institutions.                        |
| Regulations 2023/2854, 2022/868, 2024/903 and 2024/1732              | Data, data governance, interoperability and EuroHPC.                                 |
| Directive (EU) 2024/2853                                             | Liability for defective products, software included.                                 |

## 7 · Before you call yourself ready

- Given any scenario: is it an AI system, at which risk tier, and who answers for it?
- Can you separate provider duties from deployer duties without hesitating?
- Do you know the Article 113 calendar after the Omnibus, not the original one?
- Can you tell voluntary guidance (NIST, ALTAI) from a legal obligation?
- Can you place the AI Office, national authorities and the EDPS without mixing competences?
- Do you know what the Omnibus changed in Articles 4, 10(5) and 75a-75d?
- Have you done three mock tests of 30 questions in 40 minutes, in your language 2?

**Practise with questions from this field, in the five languages of the exam.**

EPSO/AD/430/26 artificial intelligence bank, with the explanation and the article behind every answer: [eu-now.com/specialist/artificial-intelligence](https://eu-now.com/specialist/artificial-intelligence)



**And when something does not click, ask.** In EU-now, EUgenio answers with the law in front of it and cites the article the answer comes from: the difference between knowing you got it wrong and knowing why.

Figures and deadlines come from the notice published in the Official Journal (C/2026/4668) and from the official texts cited. If EPSO publishes a corrigendum, the Official Journal prevails.