

Fifteen questions a regulator can ask.

Since 27 July 2026 the dates differ from the ones in most checklists: the high-risk regime has been postponed, but part of the regulation applies **today**. Every question states which of the two, and always the same follow-up: can you **demonstrate** it?

THE FIFTEEN CONTROLS

6 BLOCKS · STATUS 28 AUGUST 2026

SIGHT		LINE		
01	02	03	04	05
DISCLOSE			PEOPLE	
06	07	08	09	10
TRAIL			OVERSIGHT	
11	12	13	14	15

The clock was moved, not stopped.

NU			
2 FEB 2025	2 AUG 2026	2 DEC 2026	2 DEC 2027
Prohibited practices	Transparency & enforcement	Marking & wider bans	High risk, Annex III
The prohibited uses in Article 5 and the AI literacy duty apply from this date. The highest maximum penalty in the regulation belongs here too.	The transparency obligations in Article 50 apply and the enforcement framework of the regulation has taken effect. Note: national implementing law is not yet in force in every member state.	End of the grace period for machine-readable marking of AI-generated output, plus an extension of the prohibited practices.	The full high-risk regime, including the obligations for deployers and the fundamental rights impact assessment. For AI in medical devices: 2 August 2028.
The Digital Omnibus (Regulation (EU) 2026/1744, in force 27 July 2026) postponed the high-risk regime to fixed dates , with no conditions attached and no further extension in prospect. What already applies is narrower than most people assume. But what December 2027 asks for, namely being able to show per interaction what happened and keeping those logs for six months, plus documented provenance of your data, is not something you build in a quarter .			

One question per point: evidence, or policy?

- **Demonstrable.** There is evidence an outsider can verify without you standing next to it explaining how it works.
- **On paper.** The policy exists, but nothing enforces it. This is where an inspection stalls, and where most organisations are.
- **Unknown.** Not yet established. That is a result too, as long as it is written down somewhere with a name under it.

A	Sight of what is running	01 · 02
B	The line and the risk class	03 · 04 · 05
C	Saying that it is AI	06 · 07 · 08
D	People who know	09 · 10
E	A trail that holds up	11 · 12 · 13
F	Oversight with authority	14 · 15

Every question states when the obligation applies: **now**, or from a later date.
Rule of thumb for your answer: whatever you cannot back with a record counts as “on paper”.

01

Do you know which AI applications are in use across your organisation?

Not only the systems you bought, but also the AI features that arrived inside software you already had: the electronic health record, the office suite, the rostering tool, the recruitment system.

APPLIES NOW PRECONDITION FOR EVERYTHING BELOW

02

Do you know, per application, whether you are a deployer or a provider?

That role decides which obligations rest on you. Anyone who offers a system under their own name, substantially modifies it, or puts it to a different purpose becomes a provider, with a heavier set of duties.

APPLIES NOW ART. 3 · ART. 25

IN PRACTICE

An inventory in a spreadsheet is out of date on the day you fill it in. Route AI traffic through a single point of access and the list becomes a by-product of use instead of a project.

03

Can you rule out that emotions of staff or applicants are being inferred?

Emotion recognition in the workplace is prohibited, including home working and recruitment. The exception for medical and safety reasons is narrow: wellbeing or stress monitoring of staff does not fall under it.

APPLIES NOW ART. 5 · FINE UP TO 7% OF TURNOVER

04

Are you aware that your HR processes can fall under high risk?

Screening applicants, performance monitoring and behaviour-based task allocation are on the list. This is the most underestimated route: every institution runs personnel policy, including those without any clinical AI.

FROM 2 DEC 2027 ANNEX III, POINT 4

05

Is it recorded per application which class it falls into, and why?

If you rely on the exemption for narrow or preparatory tasks, that assessment has to be on paper before you put the system into use. Where profiling of individuals takes place, the exemption never applies.

FROM 2 DEC 2027 ART. 6

IN PRACTICE

Put the line in the technology, not in a document. As long as a prohibition exists only on paper, practice can drift from it without anyone noticing. That gap is exactly what a regulator looks for.

06

Does everyone talking to your chat assistant know they are talking to AI?

Anyone interacting with an AI system has to know it, unless that is obvious to everybody. If the assistant runs under your name, this is your duty and not your supplier's.

APPLIES NOW

ART. 50(1)

07

Is AI-generated public information marked, or covered editorially?

Patient information or health guidance produced with AI has to be recognisable as such, unless there is demonstrable human editorial review carrying final responsibility.

APPLIES NOW

ART. 50(4)

08

Has your supplier arranged machine-readable marking of AI output?

Synthetic output has to be technically detectable. For systems already on the market before August 2026 the grace period ends on 2 December 2026. Ask for it in writing.

FROM 2 DEC 2026

ART. 50(2)

IN PRACTICE

Arrange the disclosure where the conversation starts, not per tool. One central setting that applies everywhere is the only version you can still guarantee a year later, and the only one you can change in a single move when requirements shift.

09

Can you demonstrate per role what you have done about AI literacy?

The duty covers **all** AI, not only high risk, so it also covers the nurse using a chatbot to draft a report. Since July 2026 you no longer have to guarantee a level, only show which measures you took. There is no separate fine attached, but after an incident this is the first thing a regulator asks for.

APPLIES NOW

ART. 4 · NO DIRECT FINE

10

Have the works council and the staff involved been informed in advance?

Where high-risk AI is used in the workplace, worker representatives and the staff concerned have to be informed before it is put into use. That runs alongside national works council law.

FROM 2 DEC 2027

ART. 26(7) · WORKS COUNCIL LAW

IN PRACTICE

One plenary session without a record did not happen as far as a regulator is concerned.

Record who received which instruction and when. Beyond that, let the tool state the limits itself, because that is the instruction that lands at the moment it matters.

11

Can you see per interaction which model received which data, and what came out?

User, timestamp, model, version, data sent along, output and who acted on it. The GDPR asks for this accountability today; from December 2027 the regulation adds a retention duty.

FROM 2 DEC 2027 ART. 26(6) · GDPR ART. 5

12

Do you keep those logs long enough, and can you get them out?

Six months is the floor in the regulation, not the target: retention periods in healthcare run up to twenty years. Also check whether your supplier actually lets you export your logs.

FROM 2 DEC 2027 ART. 26(6) · SECTOR RETENTION RULES

13

Can you explain to someone what role AI played in a decision about them?

For decisions involving high-risk AI, the person concerned has a right to a clear explanation of the role of the system and the main elements of the decision. That explanation has to come from you, not from the supplier.

FROM 2 DEC 2027 ART. 26(11) · ART. 86

IN PRACTICE

Logging cannot be added retroactively: what was not recorded is gone. Make sure the trail comes into being at the moment of the interaction, and out of reach of whoever performs it.

14

Is someone appointed who can actually stop the system?

Human oversight calls for a person with the knowledge, the time, the authority and the assurance that departing from the AI recommendation will not be held against them. A job description without a mandate is not oversight.

FROM 2 DEC 2027 ART. 26(1) AND (2)

15

Is your DPIA current, and do you know the fundamental rights assessment applies to you too?

Large-scale processing of health data is on the mandatory DPIA list. The DPIA applies now. The fundamental rights impact assessment only follows on 2 December 2027, but healthcare providers count as “private parties providing public services”, so it is coming your way the moment one application is high risk.

APPLIES NOW GDPR ART. 35 · ART. 27 FROM DEC 2027

IN PRACTICE

Build the DPIA and the fundamental rights assessment as one file with cross-references, not as two separate documents. And make sure deviations, blocks and usage spikes surface somewhere, so that intervening does not depend on who happens to notice.

Three assumptions that do not hold.

■ “Everything is postponed, we have until December 2027.”

Half of that is right, and that is what makes it dangerous. Not postponed: the **prohibited practices**, **AI literacy**, and the **transparency obligations** since August 2026. The postponement covers the high-risk regime only.

■ “We do not build AI ourselves, so this does not apply to us.”

As a **deployer** you have obligations of your own: disclosing that it is AI, AI literacy, and later human oversight and log retention. In practice it is the deployer whose daily work is affected, not the builder.

■ “Our AI supplier is compliant, so we are too.”

Compliance is not transferable. A provider can have their part in order; your duties remain: your own data, your own staff, your own log. And a contract is not a record: the evidence has to come from your environment.

Where you can start this month.

01

Finish the inventory.

Which applications are running, who is responsible, and are you a deployer or a provider? Without this the rest stays an assumption.

02

Do what already applies first.

The prohibitions, disclosing that something is AI, and demonstrable AI literacy. This is the part that can be asked of you today, and it is the smallest part of the work.

03

Build the trail before December 2027.

Recording, masking and limiting at the point of access, not in work instructions. Wait on this and you can no longer demonstrate it retroactively.

Governance you show, not promise.

The questions in this self-check can be answered with a record, or with an assumption. That difference is the whole job.



Proven where data is most sensitive.

Numaga is in use at Ontdek Zorg, a home-care provider in the Netherlands, where special categories of personal data under the GDPR are processed: policy and masking on every prompt.

numaga.ai/eu-ai-act

Numaga is a managed, sovereign AI platform by Replikate. Single-tenant, hosted in the Netherlands, ISO 27001-certified supplier.

A practical aid, not legal advice. Status 28 August 2026, after the entry into force of Regulation (EU) 2026/1744, which moved the high-risk dates. Article references to Regulation (EU) 2024/1689 are indicative; check the consolidated text for your own situation. Published by Numaga (Replikate).