

The human in the loop is a job

Role profiles for human oversight of artificial intelligence (AI) systems, for deployers preparing for Article 26(2) of the AI Act

AT A GLANCE

For	deployers of high-risk AI systems
Contains	3 role profiles, a 7-point checklist, a job description template
Applies from	2 December 2027 for high-risk uses such as hiring, credit scoring and public services; 2 August 2028 for AI inside regulated products such as medical devices
Format	Word and Portable Document Format (PDF)
Licence	Creative Commons Attribution 4.0 (CC BY 4.0), free to adapt

2 December 2027

Article 26(2) applies to high-risk AI used in hiring, credit scoring, education, access to public services and the other areas listed in Annex III of the AI Act.

Why this exists

The European Union (EU) AI Act, Regulation (EU) 2024/1689, asks two things of human oversight, and they sit in two different articles.

- **Article 14** is about the system. A high-risk AI system must be designed, and provided to the deployer, in such a way that the people who oversee it are enabled, as appropriate and proportionate, to do five things: understand its capacities and limitations, and monitor it; stay aware of automation bias; interpret its output correctly; decide, in any particular situation, not to use it, or to disregard, override or reverse its output; and intervene or stop it (Article 14(4)(a) to (e)).
- **Article 26(2)** is about the organization that uses it. "Deployers shall assign human oversight to natural persons who have the necessary competence, training and authority, as well as the necessary support."

Article 14 describes what a person must be able to do. Article 26(2) says someone must actually be given that job, with the means to do it. Both apply from 2 December 2027 to high-risk AI used in the areas listed in Annex III of the AI Act, such as recruitment and managing workers, credit scoring, education, and access to essential public services and benefits. For AI in products covered by the EU product laws in Section A of Annex I, such as toys or medical devices, where the AI is the product itself or a part that protects health and safety, and the product needs a check by an independent third party for health and safety reasons, they apply from 2 August 2028. Both dates were set by Regulation (EU) 2026/1744, which amended the AI Act.

Four words carry the obligation: **competence, training, authority, support**. This document turns them into three role profiles that a deployer can adapt, with a map from each capability in Article 14(4) to what the person needs, and a checklist to run before anyone is assigned.

It is practical guidance, not legal advice.

Three roles

Oversight usually needs more than one person, even in a small team. These are roles, not job titles: one person can hold two of them, as long as nobody is asked to oversee a queue that they also designed and are measured on.

Role	In one line	Article 26(2) word it answers most
Oversight operator	Reviews the system's outputs and decides what happens to each one	Competence, training
Oversight lead	Owns the oversight of one system: staffing, time, escalation, stopping	Authority, support
Oversight designer	Designs the review screen and the workflow, and tests whether oversight works	Support (the tools)

Role 1: Oversight operator

Purpose. To check the outputs of an AI system before they take effect, and to change, reject, escalate or stop them when they are wrong.

Responsibilities

- Review each output assigned to them against the information the case provides, not against the system's confidence.
- Reject, correct or escalate outputs that are wrong or that they cannot confirm.
- Report patterns: repeated errors, drift, cases the system handles badly, or conditions that make careful review impossible.
- Stop or pause the system, or trigger the procedure to do so, when the instructions for use or the lead say they should.

Competence (what they must be able to do), mapped to Article 14(4):

Article 14(4)	What the operator needs	How to build it
(a) understand the system's capacities and limitations, and monitor its operation	Know what the system is for, where it is known to fail, and what normal output looks like	Briefing from the instructions for use; examples of known failure cases
(b) stay aware of automation bias	Know that agreeing with a confident system is the default error, and recognize it in their own work	Exercises with outputs that are confidently wrong
(c) interpret the output correctly	Read the output, its explanation and any confidence indicator for what they actually mean	Worked examples; a one-page guide to each indicator on the screen
(d) decide, in any particular situation, not to use the system, or to disregard, override or reverse its output	Know they are allowed to disagree, and how	Written statement of their authority (see below); practice on the real screen
(e) intervene or stop the system	Know when and how to stop it, and whom to tell	The stop procedure, rehearsed

Authority (what they must be allowed to do)

- Reject or override any output, without a second approval.
- Escalate a case without justifying the escalation in advance.
- Trigger the stop procedure when the conditions in the instructions for use are met.

Support (what they are owed)

- Time: enough time per case for a careful check of a typical case and of a hard one. The lead sets the queue so that this time exists.
- A review screen on which disagreeing costs no more effort than agreeing.
- A named person to escalate to, reachable while they work.
- A written commitment that disagreeing with the system will never count against them, and that their rejection rate is not used to assess them.

Role 2: Oversight lead

Purpose. To own the human oversight of one AI system, so that the people assigned to it have the competence, training, authority and support the law requires.

Responsibilities

- Decide who is assigned to oversight, and confirm that each person has been trained before they start.
- Size the queue: work out the time a careful check needs, and staff the review so that this time is available.
- Keep the escalation route working, and answer escalations.
- Decide when the system is paused or stopped, following the instructions for use.

- Set, before any test, the threshold that oversight must meet. For example: the share of known errors that reviewers must catch in a test.
- Act on findings: change the workload, the screen, the training or the system, and record what changed.
- Own the oversight section of any fundamental rights impact assessment. Where Article 27(1) requires a deployer to carry one out, it must include a description of how the human oversight measures are implemented, according to the instructions for use. Giving that section to the lead is our recommendation, not a requirement of the Regulation.

Competence

- Understand the system, its intended purpose and its limitations at least as well as the operators do.
- Read test results as statements about the system, the screen and the working conditions, not about individual people.

Authority

- Change staffing, pacing and targets for the review team.
- Pause the system or take it out of use.
- Require changes from the provider through the organization's contract.

Support

- A budget for review time, not only for the system.
- Access to the provider for questions and fixes.

Role 3: Oversight designer

Purpose. To design the review screen and the review workflow so that operators are able to do what Article 14(4) describes, and to test whether they actually do it.

Responsibilities

- Design the review screen so that rejecting or correcting an output costs no more steps than approving it.
- Put on the screen what an operator needs to disagree: the inputs the output was based on, and what any confidence indicator actually means.
- Remove shortcuts that approve many outputs at once, where decisions affect people.
- Test the oversight: place known errors among real cases, measure how many are caught, and report the result against the threshold the lead set in advance.
- Report findings at the level of the interface and the conditions, never at the level of individual operators.

Competence

- User experience research and interaction design for review and decision-support interfaces.
- Test design: known-error test cases, pass thresholds set in advance, and honest reporting of small samples.
- Knowledge of Article 14 and of the organization's instructions for use.

Authority

- Change the review screen and the workflow, or require the provider to change them.
- Run tests in the live workflow, with the lead's written approval and with operators told in advance that such tests happen.

Support

- Access to the system's logs for test cases, without operator names.
- Time with operators to watch the work being done.

Before you assign anyone: a checklist for deployers

1. **Named people.** Is there a named oversight lead for each high-risk system, and a list of operators?
2. **Training done.** Has each operator been briefed on the system's limitations, on automation bias, on reading its output, and on how to override and stop it?
3. **Authority in writing.** Does each operator have, in writing, the right to reject, escalate and stop without prior approval?
4. **Time exists.** Does the number of cases per operator per hour leave time for a careful check of a hard case?
5. **No penalty for disagreeing.** Is it written down that rejection rates are not used to assess people?
6. **Escalation works.** Can an operator reach the lead, or a deputy, while they work?
7. **A test is planned.** Is there a date for the first test of whether oversight catches known errors, with a threshold set before it starts?

If any answer is no, oversight exists on paper only. Article 26(2) requires that the people assigned have the necessary competence, training, authority and support; the checklist is recommended practice for showing that they do.

Job description template: oversight operator

Adapt to your organization. Keep the authority and support sections: they are how this template gives effect to the authority and support that Article 26(2) refers to. The specific commitments in them are recommended practice, not wording the Regulation prescribes.

Role: Oversight operator, [system name]

Reports to: [oversight lead]

Purpose: To check the outputs of [system name] before they take effect, and to change, reject, escalate or stop them when they are wrong.

You will

- review [number] cases per [day/shift], each with the time a careful check needs;
- reject, correct or escalate any output you cannot confirm;
- report patterns of error and conditions that make careful review impossible;
- follow the stop procedure when [conditions from the instructions for use].

You may, without asking first

- reject or override any output;

- escalate any case to [name/role];
- trigger the stop procedure.

You will be given

- training on [system name]: what it is for, where it fails, how to read its output, how to override and stop it, before your first shift;
- a review screen on which rejecting takes no more steps than approving;
- [minutes] per case on average, and more for cases you flag as hard;
- a written commitment that disagreeing with the system is never held against you, and that rejection rates are not used to assess you.

You will be tested as part of a team, not as an individual: the organization checks from time to time whether oversight catches known errors placed among real cases. You will be told that such checks happen, not when. Results describe the screen and the working conditions, never individuals.

Sources

- Regulation (EU) 2024/1689 (AI Act), Articles 14(1), 14(4), 26(2), 26(3) and 27(1). EUR-Lex: [eur-lex.europa.eu](https://eur-lex.europa.eu/eli/reg/2024/1689/oj) (<https://eur-lex.europa.eu/eli/reg/2024/1689/oj>)
- Regulation (EU) 2026/1744, Article 1 point 40 (application dates). EUR-Lex: [eur-lex.europa.eu](https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202601744) (https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202601744)
- Room to Disagree, the framework's five capabilities (Understand, Monitor, Resist, Interpret, Intervene) mapped to Article 14(4). [uxellence.gr](https://www.uxellence.gr/room-to-disagree) (<https://www.uxellence.gr/room-to-disagree>)

Citation and licence

Mantzios, M. (2026). The human in the loop is a job: role profiles for human oversight of artificial intelligence (AI) systems, for deployers preparing for Article 26(2) of the AI Act (Version 1.0). UXellence. <https://doi.org/10.5281/zenodo.23061032>

© 2026 UXellence. Licensed under [Creative Commons Attribution 4.0 International \(CC BY 4.0\)](https://creativecommons.org/licenses/by/4.0/) (<https://creativecommons.org/licenses/by/4.0/>). Adapt the profiles and templates for your organization, with attribution.