REVYCORE

WHITE PAPER

AI Act & Human Oversight.

Proving that a person remains in charge.

REVYCORE RESEARCH

DECEMBER 2026

EXECUTIVE SUMMARY

Human oversight is a requirement. Proving it is a discipline.

Regulation (EU) 2024/1689, the AI Act, makes human oversight a central requirement for high-risk AI systems: they must be designed so that natural persons can effectively oversee them, and deployers must assign that oversight to people with the competence and authority to exercise it.

As AI agents gain identities, tools and access to data, oversight is no longer only a question of model design. It becomes a governance question: who answers for each agent, what it is allowed to do, and whether a person actually decided.

This paper summarises the oversight principles of the regulation, explains why they meet access governance, and describes the evidence that shows a person remains in charge.

01 **Oversight by people**
Effective oversight by natural persons, with competence and authority.

02 **Agents act**
Identities, tools and data access turn oversight into governance.

03 **Decisions, not policies**
A named person, on a date, with a reason.

04 **Evidence that lasts**
Attestations sealed and verifiable, agent by agent.

CONTENTS

01 What the regulation asks
02 Why agents make it a governance question
03 Oversight in six dimensions
04 Authority that is structural
05 Evidence of oversight
06 Where to start

01

What the regulation asks

The AI Act follows a risk-based approach. For systems classified as high-risk, it sets requirements on data, documentation, transparency, robustness and human oversight. Article 14 requires that such systems be designed and developed so that they can be effectively overseen by natural persons during the period in which they are in use.

Deployers carry their own obligations. Among them, they must assign human oversight to natural persons who have the necessary competence, training and authority, and give them the support they need.

The regulation's obligations apply on a phased timetable, and their scope depends on the classification of each system. Each organisation should confirm which of its systems are concerned and from when.

Human oversight. the ability of designated natural persons to understand, monitor, intervene in and, where necessary, stop the operation of an AI system.

> **Oversight must be effective, not decorative.**

02

Why agents make it a governance question

An AI agent is not only a model. It holds or borrows an identity, reads document sources, calls tools and acts on behalf of users. Each of those capabilities is an access, and access is what governance already knows how to certify.

Oversight that ignores what an agent can reach is incomplete. A person cannot oversee an agent whose tools, sources and audience nobody has listed, nor one whose owner left the organisation months ago.

The question an auditor will ask is therefore concrete: for each agent, who answers for it, what is it allowed to do, and when did a person last confirm it?

> **An agent with access is an identity that acts.**

03

Oversight in six dimensions

Effective oversight starts from a description of what each agent is. Six dimensions are enough to make it governable, and each one can be attested separately by its owner.

DIMENSION	WHAT OVERSIGHT NEEDS TO KNOW
Responsibility	Who answers for the agent, and for what purpose it exists.
Identity	Whether it acts under its own identity, a shared one, or the user's.
Access	Which entitlements it holds or exercises.
Knowledge	Which sources and databases it consults.
Tools	Which actions it can trigger.
Audience	Who can call on it, and whether that audience has grown.

> **Describe the agent before you oversee it.**

04

Authority that is structural

Oversight is credible when the system makes it impossible to bypass. Decisions on access and on agents should be taken by named people, never signed by a system account, never delegated to an automation, and never available as an action to an assistant.

Where an assistant is used to explain or summarise, its own activity should be recorded: what it actually did, with which tools, under which session, and under whose authority. Oversight then applies to the tools of oversight as well.

- **No system signature.** Every decision carries the name of a person.
- **No decision by the assistant.** Assistance explains; it does not decide.
- **A register of the assistant's actions.** What it did, not what it was allowed to do.

> **Humans decide, always. By design, not by policy.**

05

Evidence of oversight

Oversight leaves traces only if they are designed. For each agent, the evidence should show who its owner is, what was attested in each dimension, when, and what changed since. When an agent gains a tool or a wider audience, its attestation should be reassessed, and the new decision recorded.

Sealing those records makes them verifiable by a third party, without trusting the platform that produced them. That is what turns a statement of oversight into a demonstration.

01

Owner

A named person, reachable, who accepted the role.

02

Attestation

Each dimension decided, dated and justified.

03

Reassessment

Triggered when the agent changes.

04

Verification

Sealed records, checked offline.

> **Six months later, in front of a regulator.**

06

Where to start

Most organisations do not yet know how many agents they run. Oversight starts with that inventory.

1. **Inventory the agents** across collaboration, cloud, business and development platforms.
2. **Name an owner for each,** with a stated purpose.
3. **Attest the six dimensions** and record each decision.
4. **Reassess on change:** a new tool, a new source, a wider audience.
5. **Seal the evidence** and keep it ready for the oversight question.

> **Inventory first. You cannot oversee what you have not listed.**

REVYCORE

A person in charge, and the proof of it.

Human decisions. Verifiable evidence.

Request a demonstration in your context: your solutions, your scopes, your agents.

revycore.com

This paper summarises principles of Regulation (EU) 2024/1689 for information purposes; it is not legal advice and does not assess whether any system is high-risk. Obligations apply on a phased timetable that each organisation should verify. © 2026 REVYCORE SAS.