



THE AUTONOMY DIAL

Human-in-the-loop isn't a brake.


It's how you design autonomy.

ONE IDEA • 60 SECONDS • EVERY WEEK

Mehmet Erkek

AI Strategy & Engineering


01 / 08



THE FALSE BINARY

Full autonomy or full manual? Neither.

Teams ask: “Can we trust the agent?” Wrong question. Autonomy isn't a switch. It's a dial you set per decision — based on what happens when the agent is wrong.



THREE MODES

HITL — HUMAN IN THE LOOP

A person is an active step inside the decision loop — reviewing, approving, correcting, or rejecting the AI's output before it takes effect. The AI proposes; the human disposes.


HOTL — HUMAN ON THE LOOP

Agent acts. Human monitors and can intervene. For medium-risk work where speed matters and mistakes are reversible.

HOOTL — HUMAN OUT OF THE LOOP

Agent operates autonomously within a bounded, pre-authorized scope. For well-defined, low-risk actions.

Most teams use one mode for everything. That's the mistake.




WHAT BAD HITL LOOKS LIKE

**Approve everything and
you get neither safety
nor speed.**

If an agent pauses every ten seconds for permission, the ROI of automation collapses. And a reviewer clicking “approve” 200 times a day isn't providing meaningful oversight — they're providing a signature.

**Rubber-stamping is a liability
dressed up as process.**



WHAT CHANGED (AND WHAT DIDN'T)

A 16-month extension — not a repeal.

The EU AI Act's rules for standalone high-risk AI systems (Annex III), including Article 14 on human oversight, will apply from 2 December 2027 — postponed from the original August 2026 deadline via the AI Omnibus.

Oversight must be appropriate to the risks, the level of autonomy, and the context of use.

High-risk areas include employment, credit, essential services and other decisions affecting people's rights.



THE REAL STANDARD

Own the scope, not the clicks.

Meaningful oversight isn't about approving every action. It's about defining who can authorize the system, what it is allowed to do, when a human must intervene, and how those decisions are recorded.

The goal isn't to watch every step. It's to make autonomy bounded, accountable and auditable.



FOUR QUESTIONS PER USE CASE

REVERSIBLE?

Can you undo it quickly?

BLAST RADIUS?

One record — or ten thousand customers?

REGULATED?

Does a law govern this decision?

VISIBLE?

Would a mistake be caught immediately — or months later?

The more irreversible, wide-reaching, regulated and invisible the failure, the closer you move toward human approval.



TIME 'S UP

Autonomy without accountability is a liability.

Human oversight isn't what slows your agents down.
It's what lets you turn them up safely.

Save this before your next agent goes live.