


EU AI Act

# Ship AI fast. Stay in control. Be EU AI Act- ready.



## Challenge

Under the EU AI Act, companies have to account for their high-risk AI systems: what it was or is doing, who approved it, and how fast they can fix it when something goes wrong. And it isn't just regulators asking- review boards, security questionnaires, and procurement want similar evidence.

For most teams, prompts are hardcoded, model changes skip formal review, and the only record of what changed is a Git history.


When something goes wrong, teams dig through logs and commits to reconstruct what happened, then wait for the next release to fix it. By then, regulators or compliance teams may already be asking questions.

## Solution

**LaunchDarkly is your centralised governance and control layer that, paired with a solid compliance program, can help you meet regulatory standards, demonstrate compliance, and maintain release velocity.**

Monitor behavior continuously in production and surface drifts immediately; require sign-off before a prompt or model change ships; and keep a timestamped, attributed record of every configuration change- automatically.

When something needs to change, fix it directly in seconds: no redeploy, no waiting on the next release. Paired with the right processes, that's what turns "We'll figure it out if asked" into already having the answer.



# Close the gap between shipping AI and governing it.

With LaunchDarkly, you can move fast while staying in control of the AI you ship. We enable you to detect drift early and correct it immediately when something goes wrong, keep a human in the loop on every approval, and maintain a configuration change history that shows exactly who did what, and when. That's what it means to govern AI and stay in control, even when regulations kick in.

## Ongoing risk management

### What's required for high-risk systems

Continuous risk monitoring, not a one-time check at launch.

## Human oversight

A human to review AI changes before they go live.

## Change history

A complete configuration change record what changed, when, and by whom.

## Corrective action

Immediate correction when an AI system presents a risk.

### What LaunchDarkly runtime control provides

- ✓ Guarded releases pause and roll back automatically when quality degrades.
- ✓ Online evaluations score every live response continuously.
- ✓ Approvals require sign-off before a prompt or model change ships—can't be bypassed by role.
- ✓ Custom roles define who has authority to approve each change.
- ✓ Every AgentControl config change is timestamped, attributed, and retrievable, with a diff.
- ✓ Shows who approved it (not just what text changed, like a commit history would).
- ✓ Change prompts, models, or guardrails directly from the dashboard without redeploying.
- ✓ Config changes propagate in under 200ms, no build pipeline.



See how LaunchDarkly helps your team control code and agents in production.

[Contact sales](#) →

This material is not legal advice and does not confirm that LaunchDarkly satisfies any specific article of the EU AI Act.