LaunchDarkly

Humans in the AI SDLC: in control, not in the way

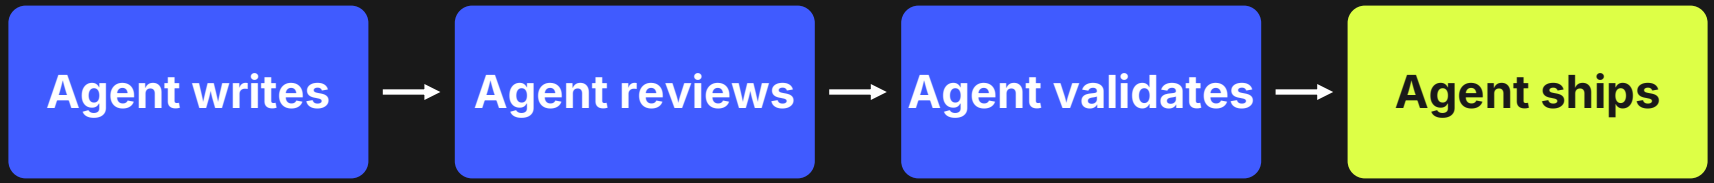


Ciarán Flynn
Senior Solutions Engineer



THE HOOK

The AI software factory is already running



Continuously. Autonomously.

Work that used to take whole teams now takes a pipeline.

THE MISCONCEPTION

Humans move to the control room

1

Set intent

What we build, and why.

2

Set policy

What agents may touch, and what is off limits.

3

Set safeguards

How we know it worked, and how we stop it if it didn't.

Agents do the heavy lifting. Humans run the factory.

THE REGULATORY TURN

Human oversight is in the EU AI Act

Art. 14

Human oversight requirements for high-risk AI systems

- Understand the system's capabilities & limits
- Stay alert to automation bias
- Interpret outputs correctly
- Override, reverse, or stop the system

The stop button is written into the law.

LaunchDarkly

THE REGULATORY TURN

Deferred, not cancelled

27 Jul 2026

Digital Omnibus enters into force

2 Dec 2027

High-risk obligations apply (Annex III, was Aug 2026)

2 Aug 2028

High-risk AI in regulated products (Annex I, was Aug 2027)

Your coding agent isn't automatically high-risk. What you ship might be. Then your pipeline is part of the evidence.

THE PRACTICAL LINE

How LaunchDarkly guardrails the factory

Control

Put agent changes behind flags

Every AI-made change ships behind a flag or AgentControl config, with approvals on the risky steps.

Understand

Watch the guardrail metrics

Guarded rollouts track error rate, latency and evaluation scores while the rollout widens.

Adapt

Roll back automatically

Cross a threshold and LaunchDarkly pauses or reverts. Flip a kill switch, no redeploy.

Run flags and guarded rollouts keep autonomy reversible

Approve approval workflows gate the risky changes

Decide you set the metrics, thresholds and policy

THE CLOSE

What to do after the Conference

Pick one agent workflow. Sort every step into run, approve, or decide.

CTOs and VPs

Deferred isn't cancelled. Build the controls now.

Engineering managers

Review is the core workflow. Staff it and measure it.

Staff+ and principals

Define policies that agents can't edit.

Developers

Your approval is a signature. Know what you're signing.

LaunchDarkly

COME SAY HELLO

See how to safeguard your AI SDLC

BOOTH

201

Come to our booth to see how LaunchDarkly keeps humans in control of the agents in your factory.

Guarded rollouts

Kill switches

AgentControl

LaunchDarkly

THE CLOSE

Humans aren't in the way. They're how the factory stays in control.

By design, and where it counts, by law.



Thank you · Ciaran Flynn
Senior Solutions Engineer