

# From Orchestration to Conversation

Observability as the missing link in agentic FS systems

# **How can we integrate agents safely into Financial Services workflows and systems?**

How can we integrate agents **safely** into  
Financial Services workflows and systems?

■ THE SUB-QUESTIONS

IMPORTANT QUESTION

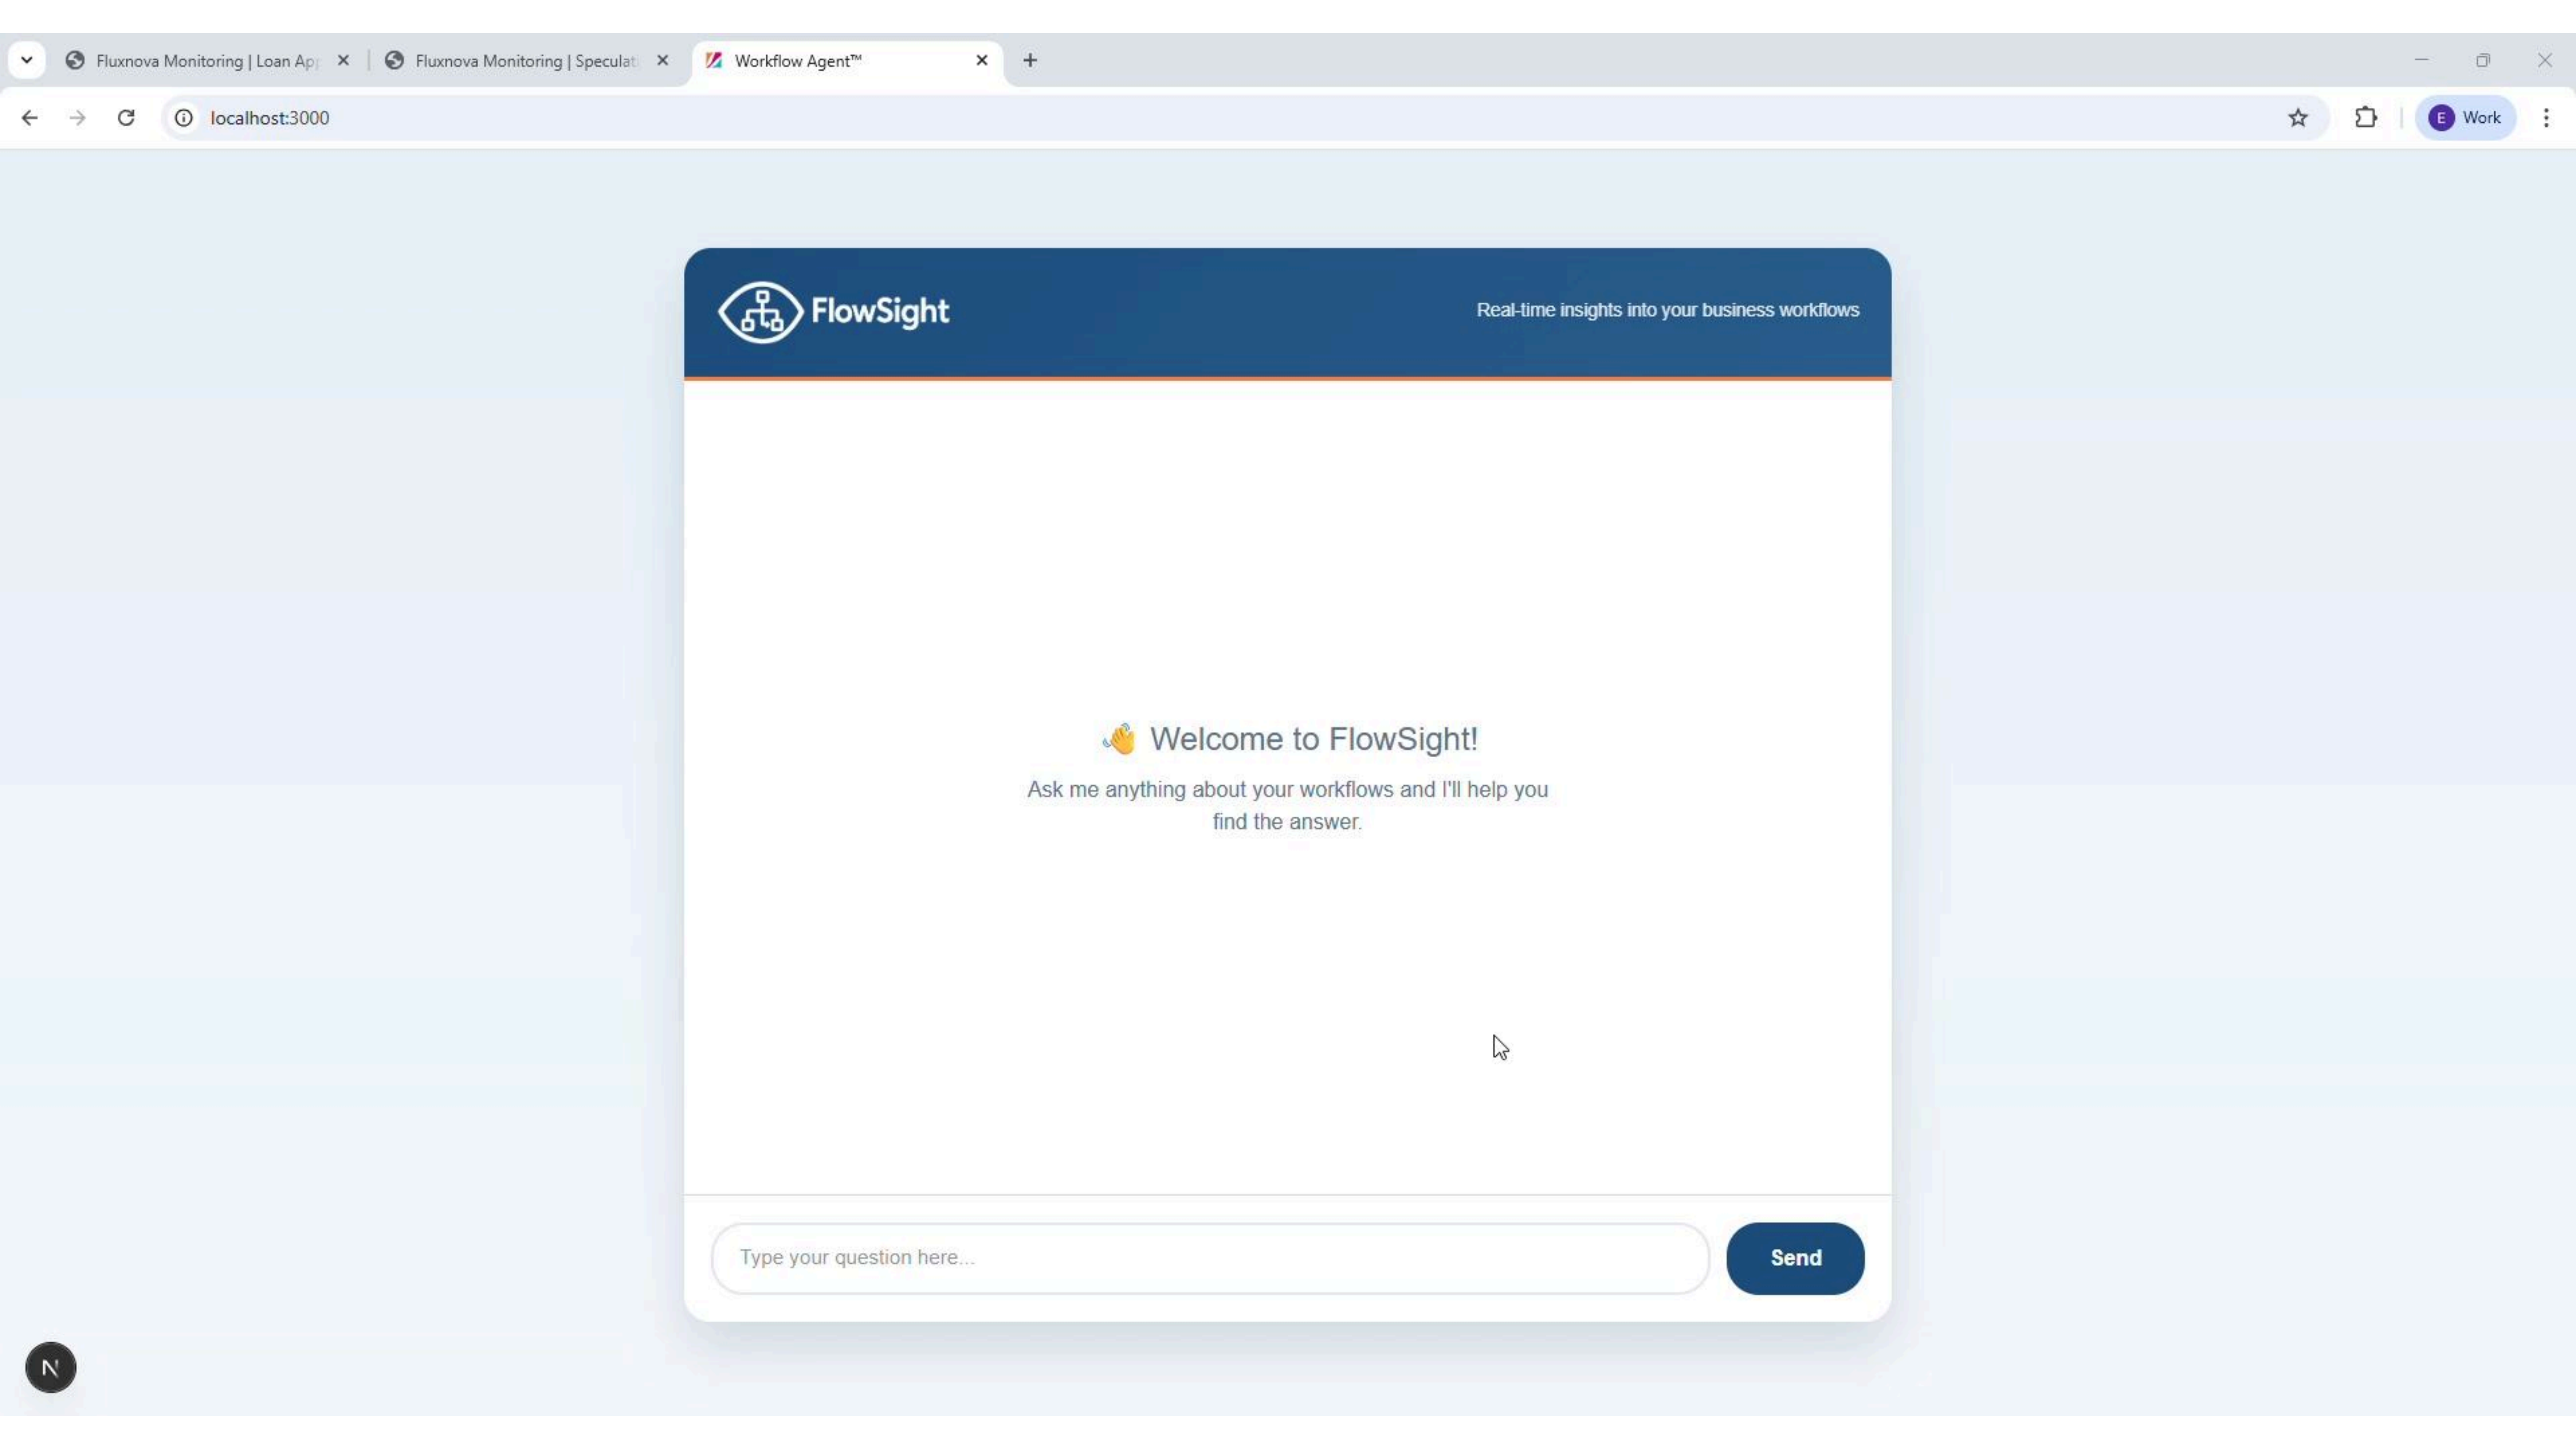
**How can we watch agents?**

MORE INTERESTING QUESTIONS

**What do we let them perceive?**

**Where do we let them act?**





Real-time insights into your business workflows

👋 Welcome to FlowSight!

Ask me anything about your workflows and I'll help you find the answer.

Send



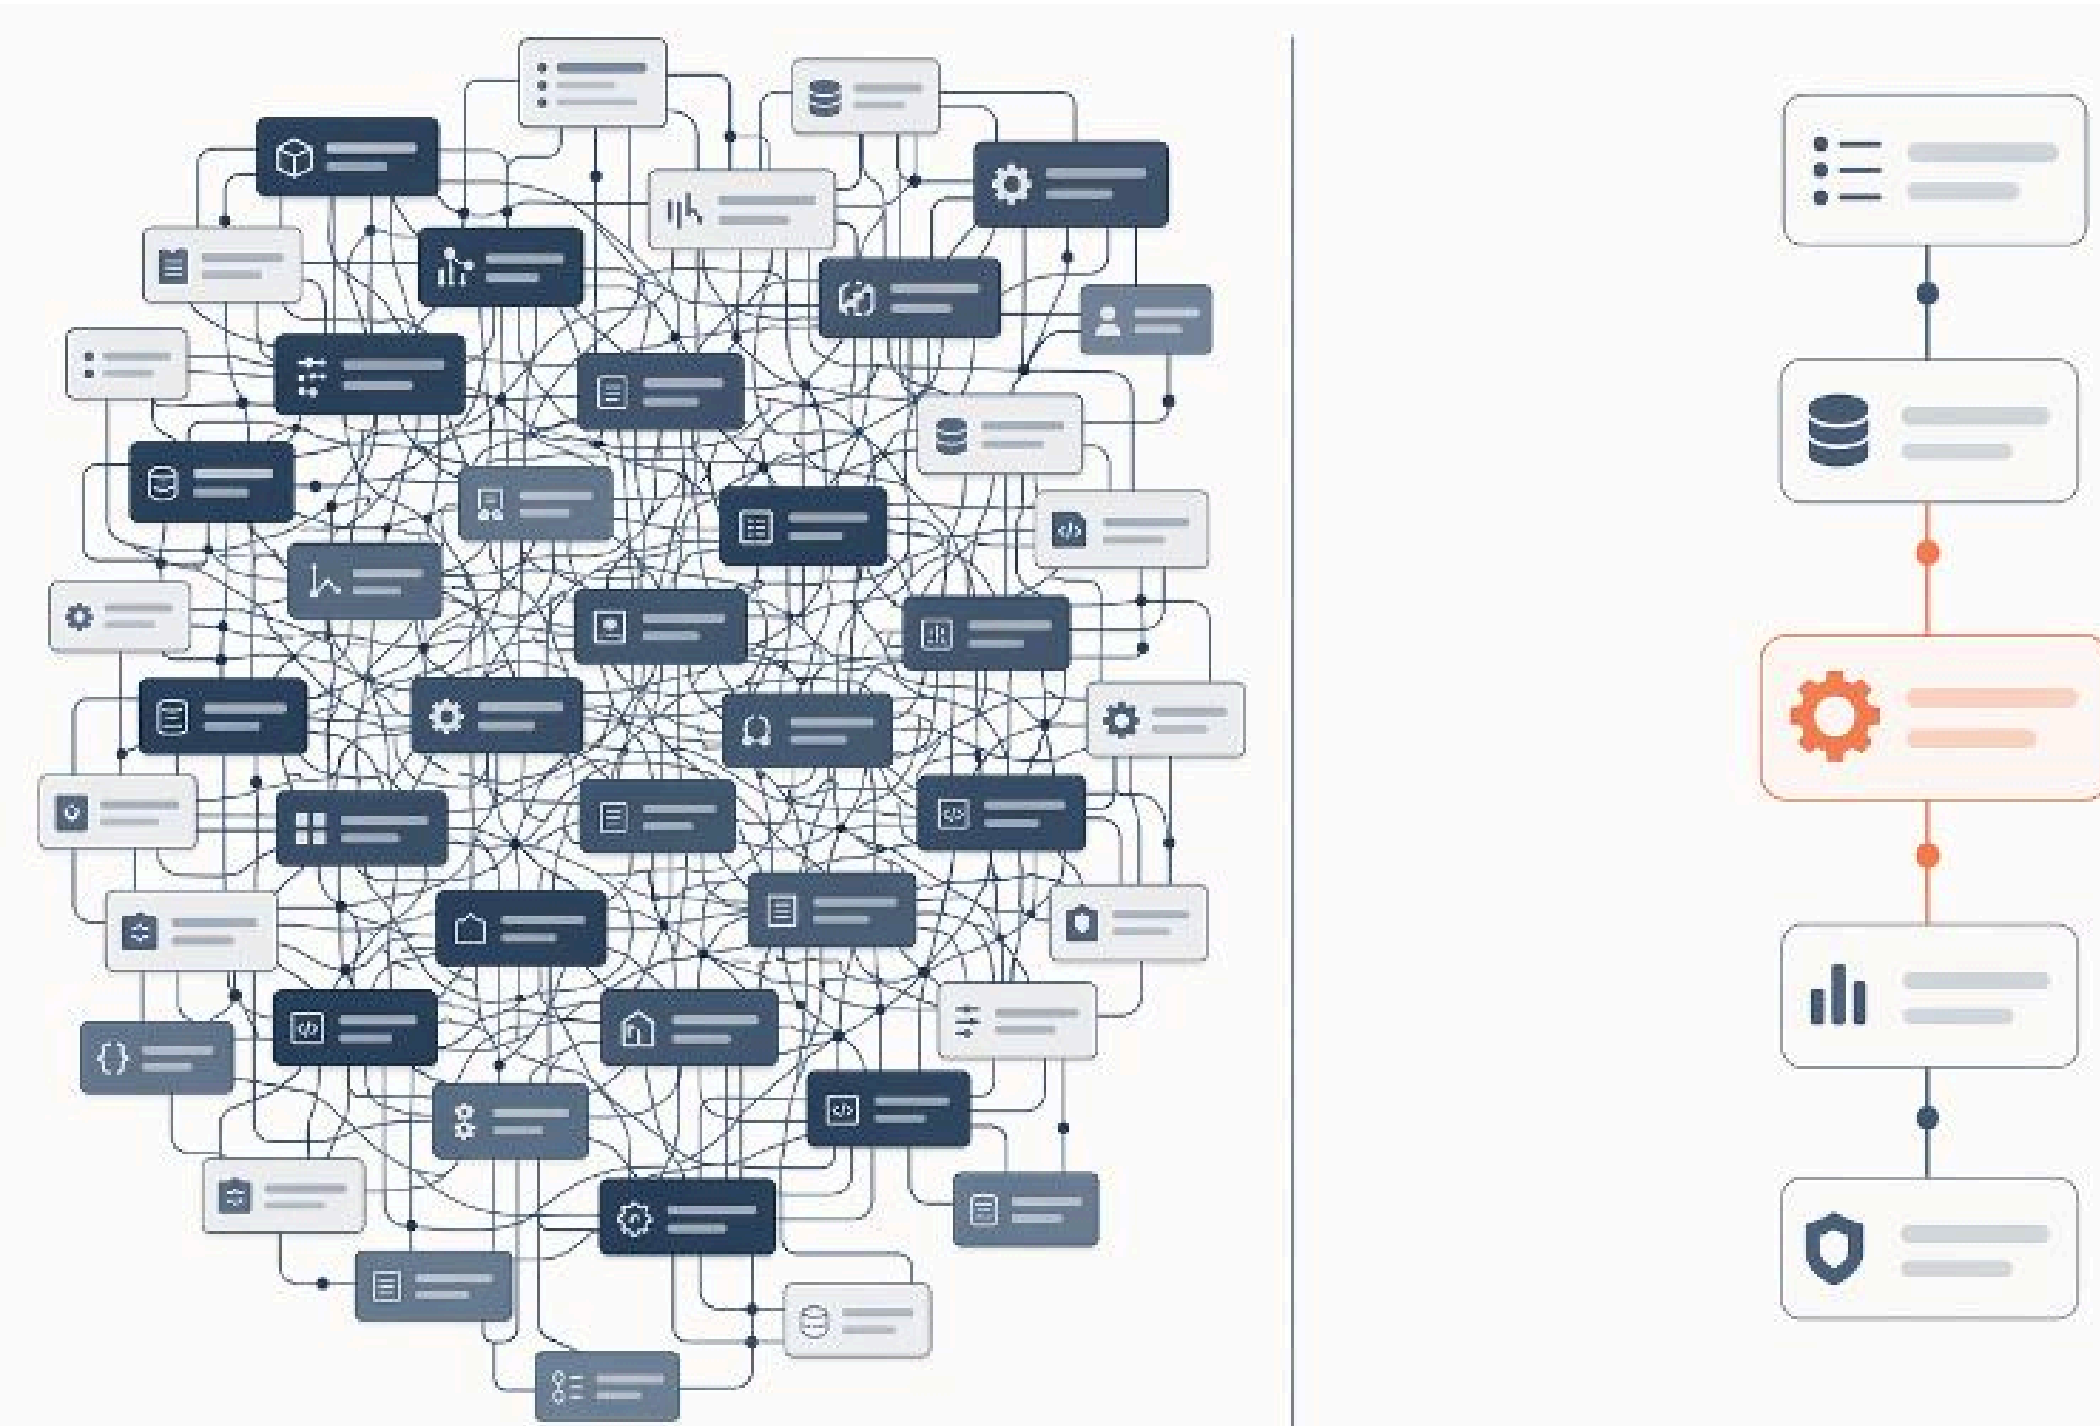
■ WHAT DO WE LET THEM PERCEIVE?

# Human readable ≠ Agent ready

We handed the agent the whole API spec.

It failed twice: too much surface to choose from, and missing the intent a person brings.

Fixing this issue required designing the agent's view intentionally: live process state, shaped for how it reasons, **not** how a person reads docs.



*Raw surface vs. a lean, curated one*

■ BUT WHAT ABOUT AUTONOMY

**Where do we let them act?**

# The constraint you need is a diagnostic.

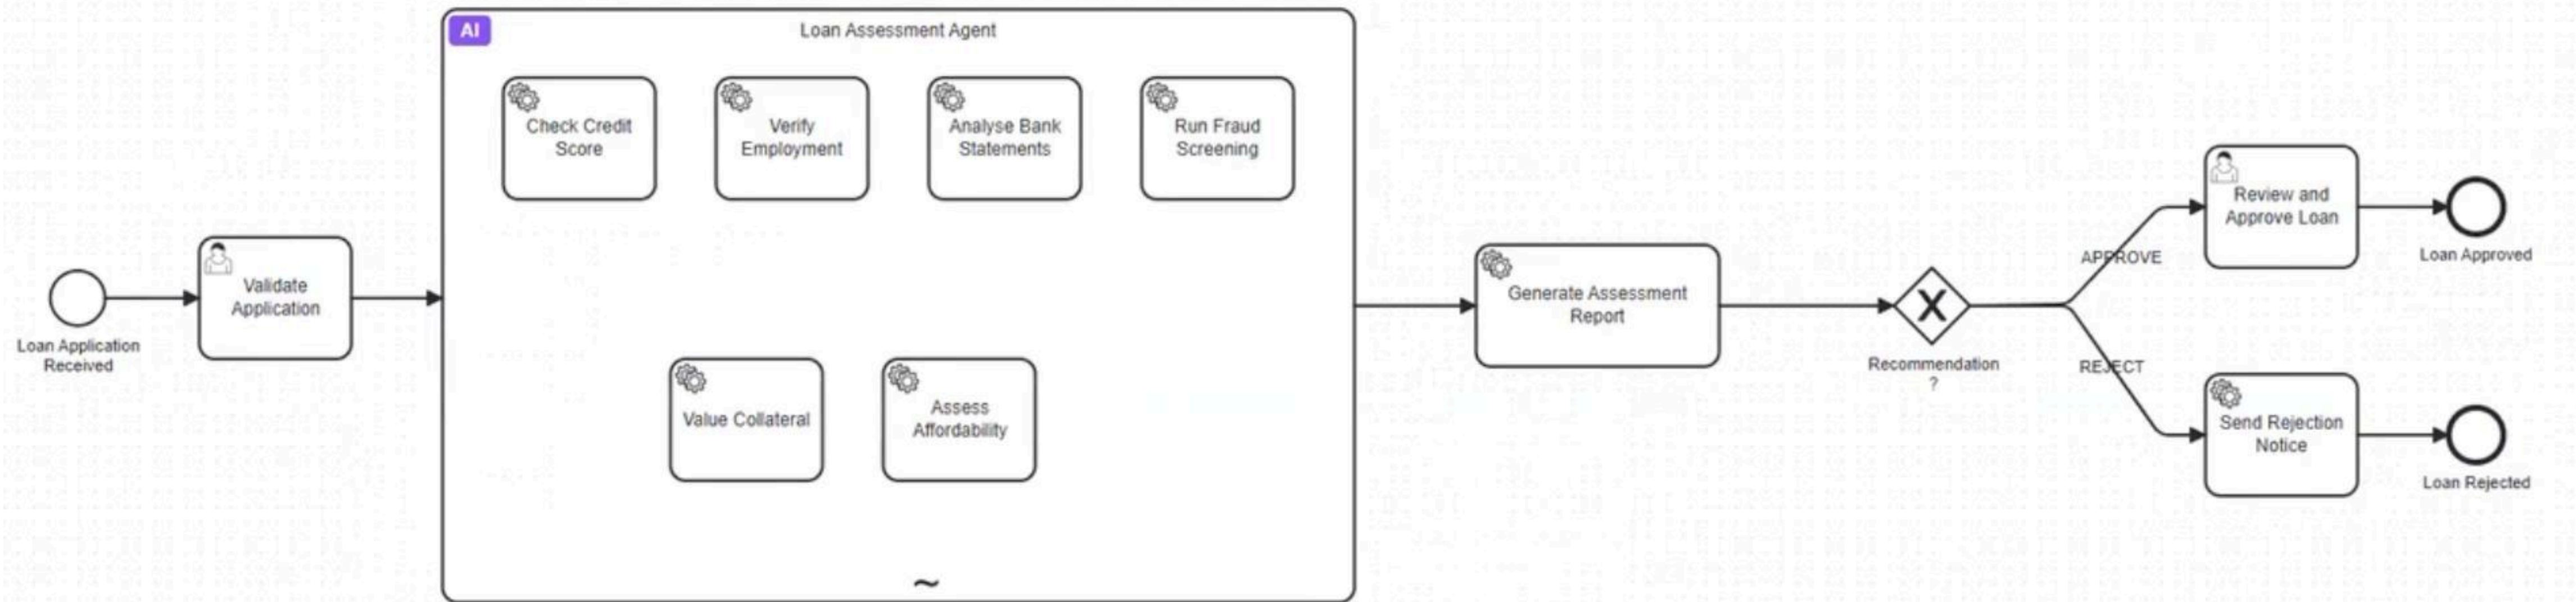
If you're needing to pin the temperature of your agent to zero, or if you run every agent-made decision through a rules engine, you might not need an agent at all; you might need a **deterministic workflow**.

In regulated finance, “**high-stakes and hard to detect**” describes most of the estate, so the bar required to justify reaching for genuine autonomy is justifiably high.

*Most FS work lives on this end of the spectrum*



■ WHERE DO WE LET THEM ACT?



■ AGENTS ARE NOT HUMANS



**The principle hasn't changed.  
The principal has.**

# An agent's intent is whatever's in its context window.

## A Person, a Service Account

Acts on **its own** intent: stable, theirs.

## An Agent

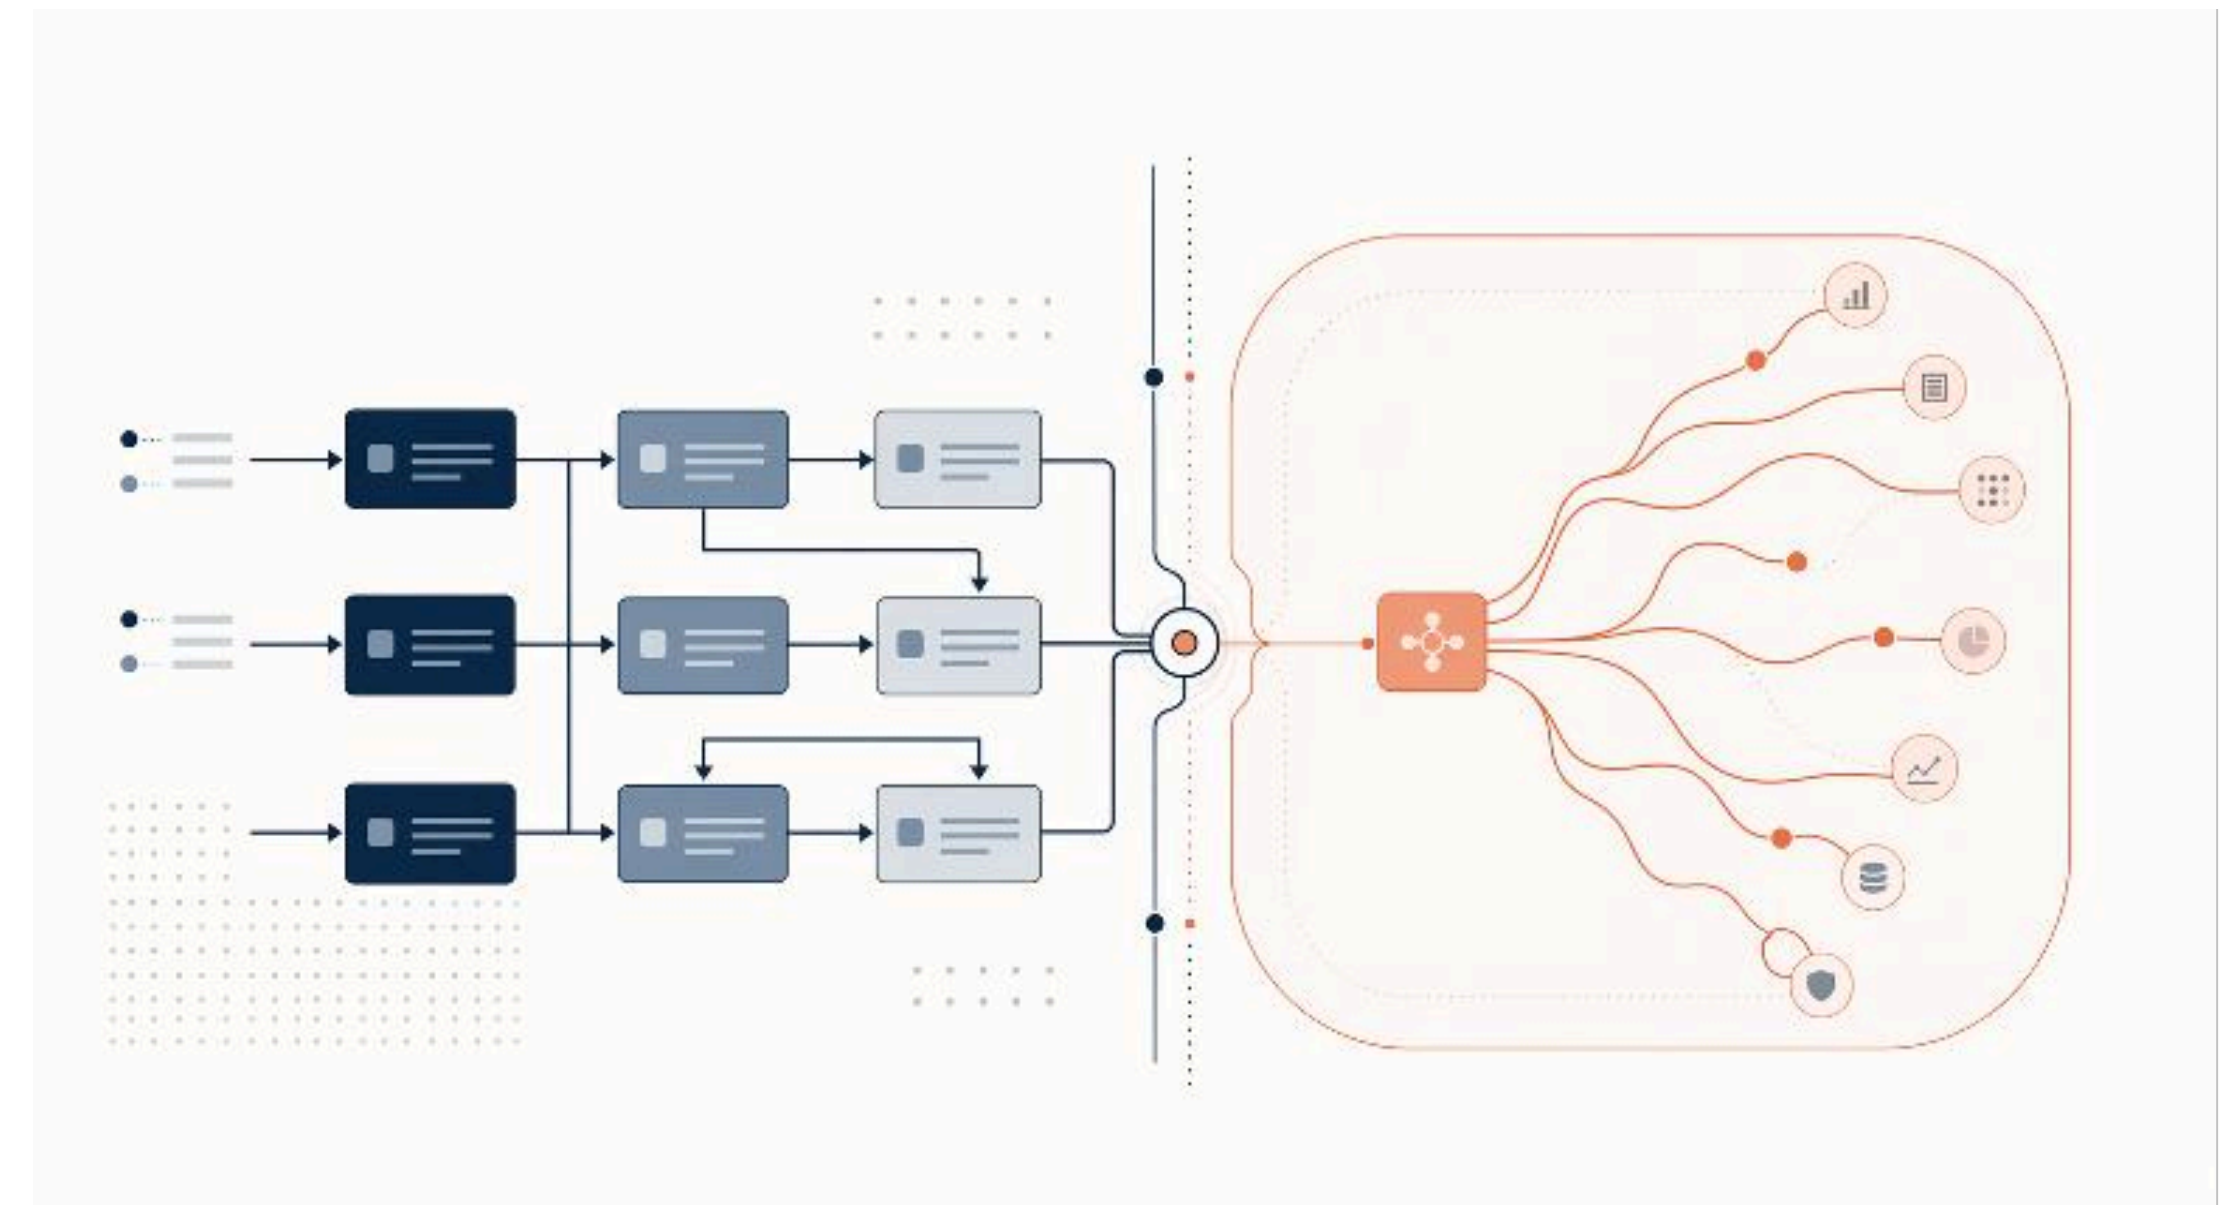
Acts on **borrowed** intent: whoever last filled the window.

If you get text in front of it, **you get a vote on what the agent does.**

# Design the seam: what the agent may perceive, where it may act

observability → context engineering → **design at the seams**

- 1 Design what it perceives, agent-first.**
- 2 Don't trust the agent to police what it does.**
- 3 Blend determinism with agency where it pays.**



*Deterministic structure & bounded agency, joined at the seam*

**Thank-you.**