

14 MAY 2026

LIVE NOW

WEBINAR BEGINS SHORTLY

A TALK ON GOVERNED AUTONOMY

# Scaling AI Agents *Without* Losing Control

*AI proposes. Systems decide.*

PRESENTED BY

## Alan Knox

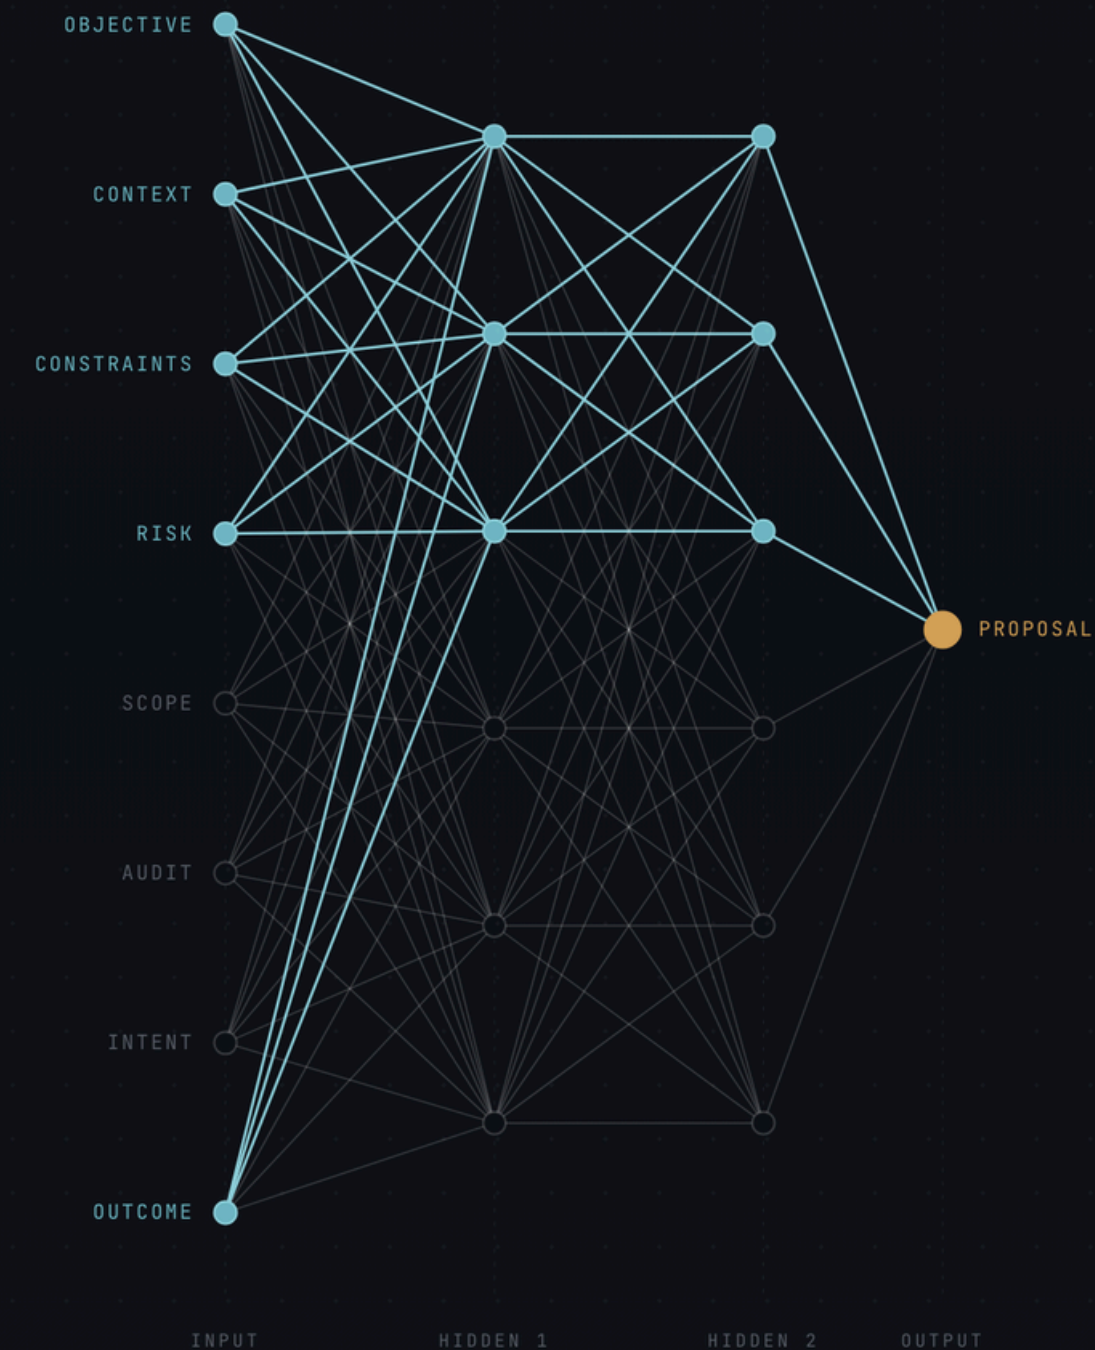
VP Product Development · QuantaViable

ABOUT

*"AI proposes. Systems decide."*

FORWARD PASS · LIVE

INFERENCE



Let me tell you  
about a paramedic.

A paramedic arrives on a scene  
and has to make real decisions.  
*Fast ones. Life and death ones.*

But a paramedic doesn't operate alone.  
They work within a system.



And none of that makes them less  
effective.

*It makes them more effective.*

Now imagine we  
remove all of that.



Confident. Fast. Sometimes right.  
But when they're wrong,  
*there is no system to catch it.*

*Would you trust that  
paramedic with your life?*



And yet, that's how most organizations are deploying *AI agents* right now.



So what does it *actually look like* when you build it?

Something that *plans*. Decomposes a business objective into structured steps, then exits.

AI · PROPOSES

**PLANNER**

DECOMPOSES OBJECTIVE  
→ STRUCTURED PLAN

Something that *executes* each step. Scoped context. One step at a time.



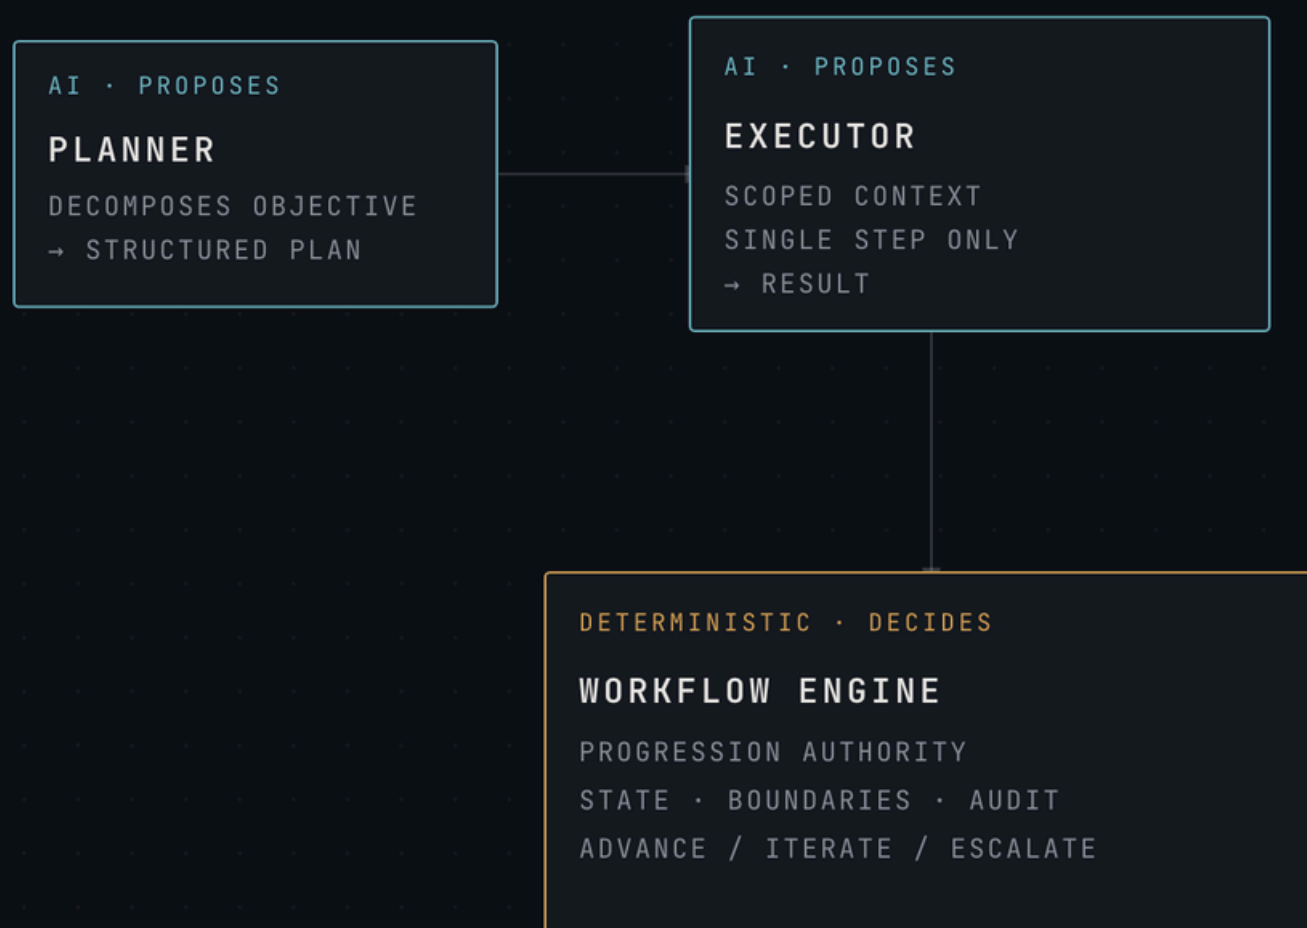
PARAMEDIC PARALLEL

## The paramedic on the call.

The Executor is the paramedic at the patient's side. Highly trained. Making real decisions. But scoped. They handle *this* patient, *this* situation.

They don't decide which call comes next. They don't manage the dispatch queue. They don't audit their own performance. Their authority is bounded, deliberately, and that's what makes them trustworthy at speed.

A deterministic *workflow engine* takes the plan and enforces step boundaries.

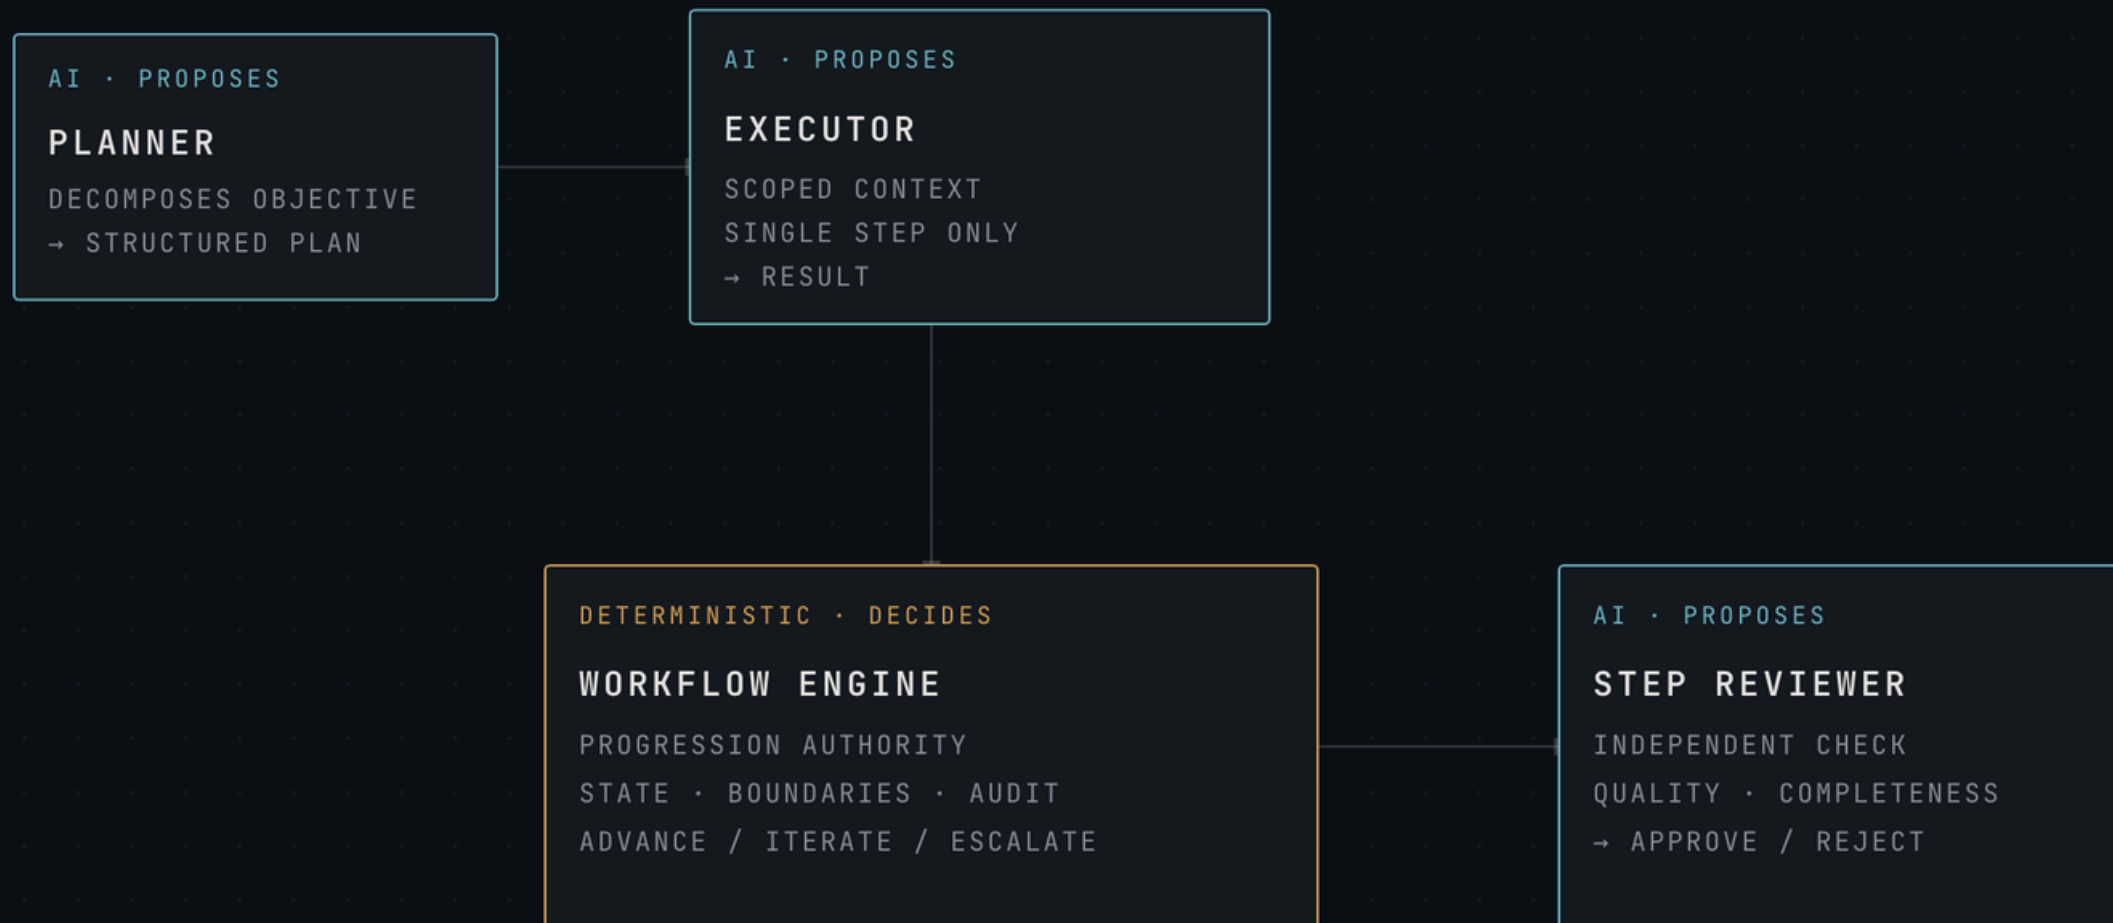


## The protocols themselves.

The Workflow Engine is the protocol book. It's the written, deterministic rules that govern what a paramedic can do, when, and in what order. Protocols don't reason. They don't improvise. They enforce.

Crucially: the paramedic doesn't write the protocols, and the protocols don't make medical decisions. They define the boundaries within which medical decisions get made. That separation, judgment inside the rules and rules outside the judgment, is exactly what the Workflow Engine does for AI.

# A *step reviewer*: independent validation against defined criteria before progression.



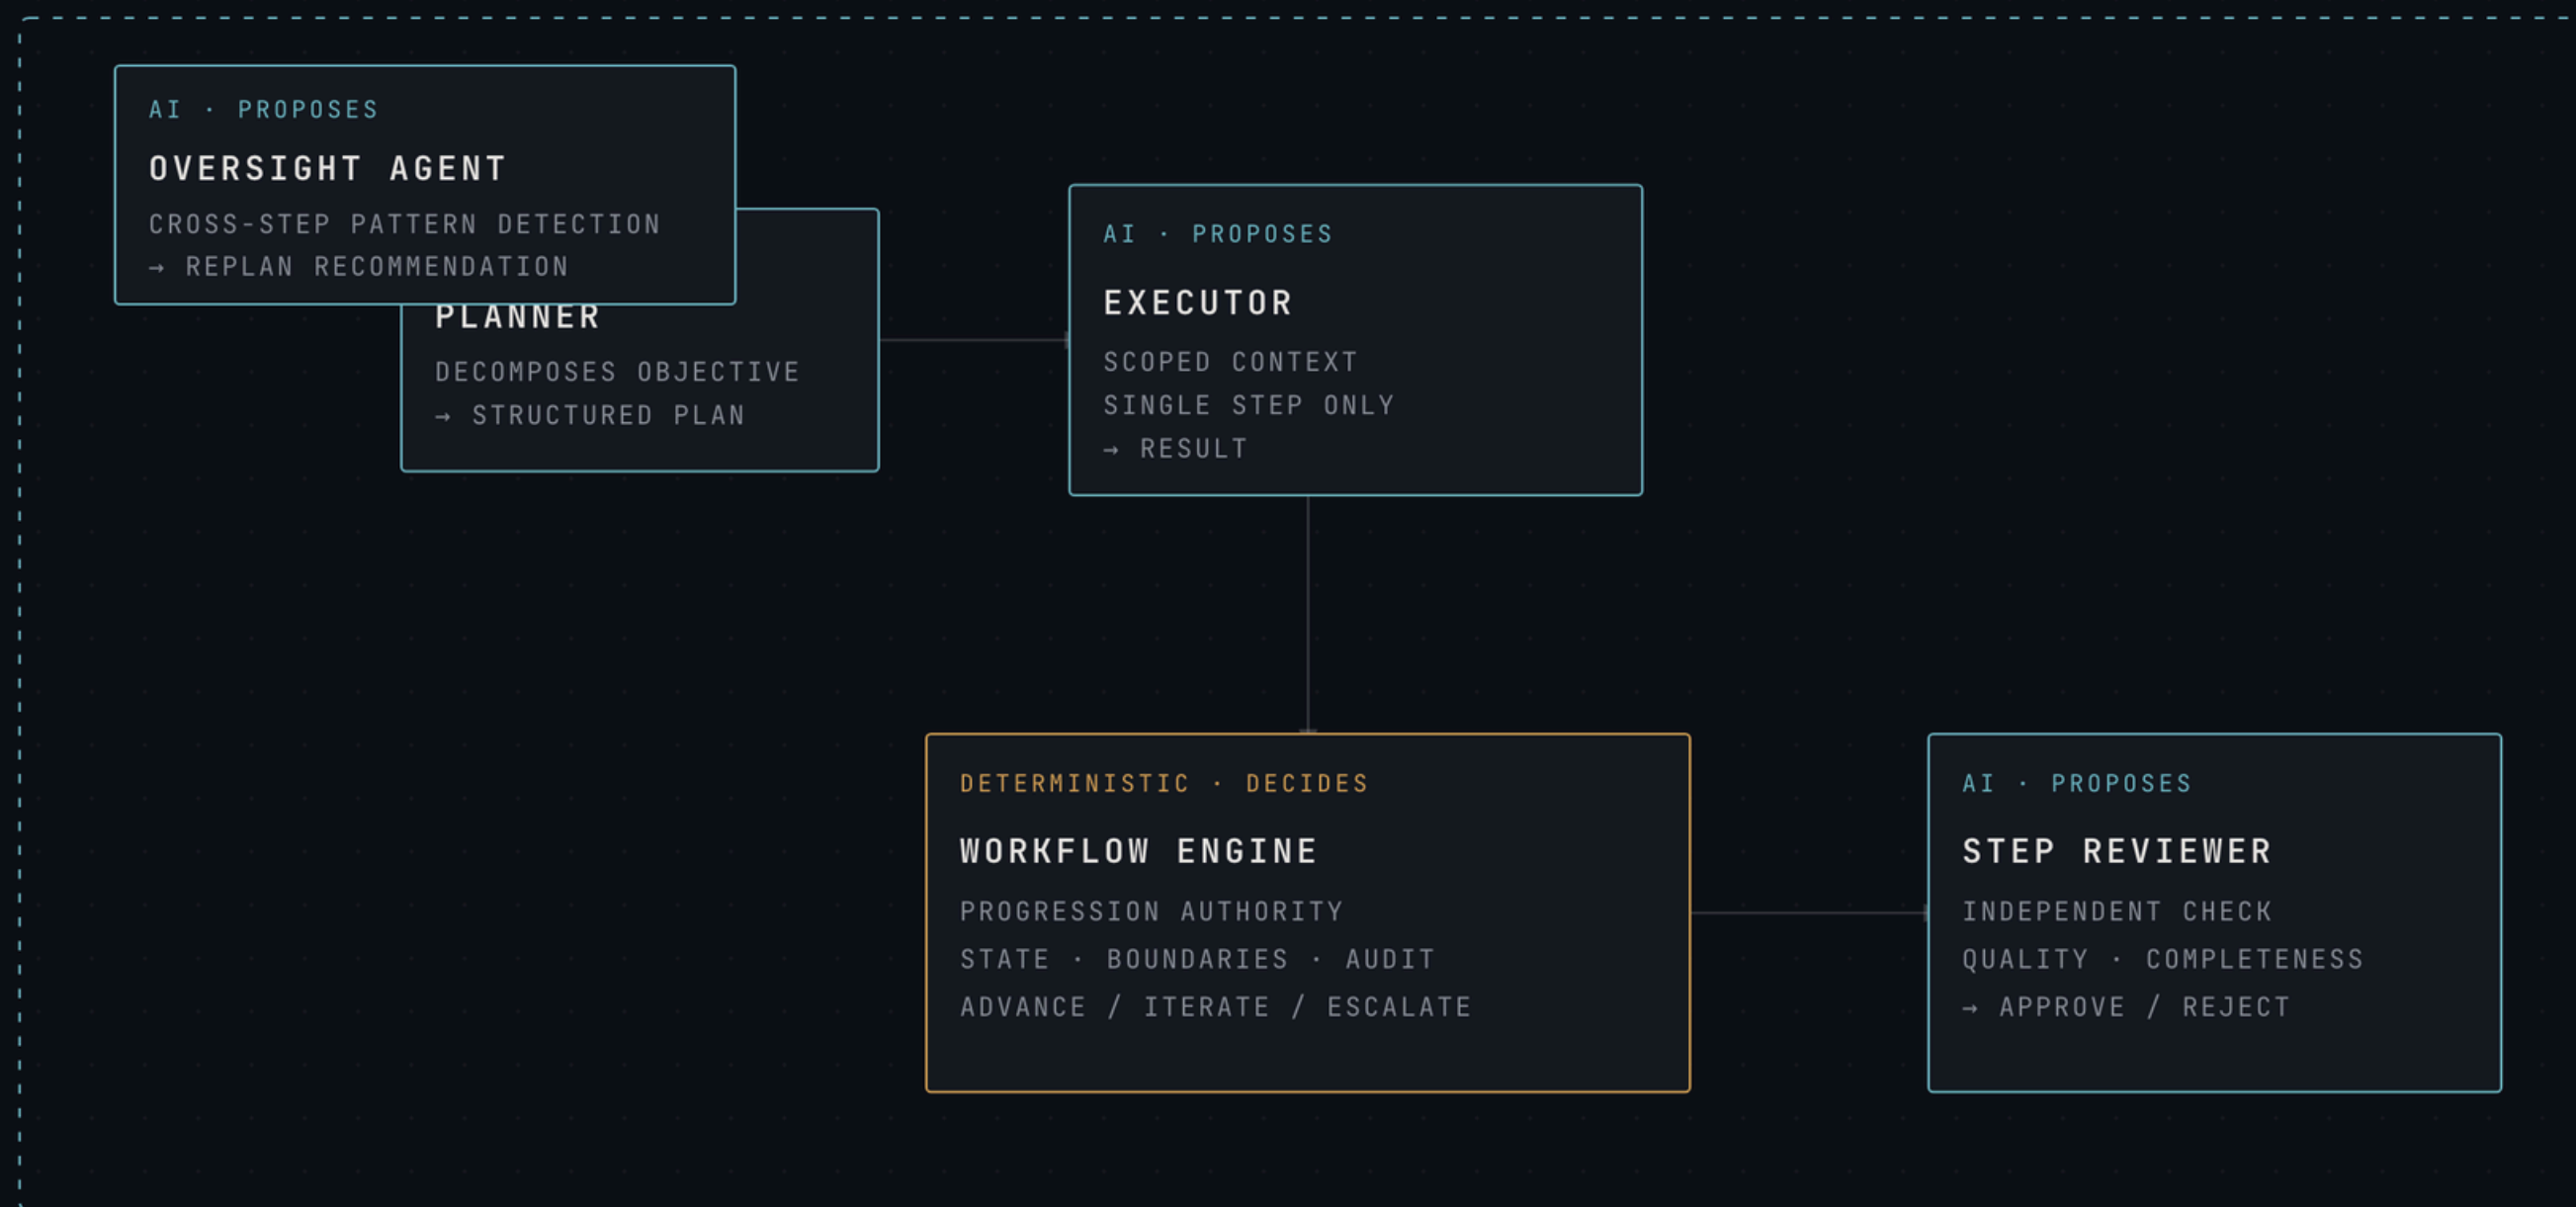
PARAMEDIC PARALLEL

## The Field Training Officer.

My son is an FTO. He doesn't redo other paramedics' work. He reviews it. Validates decisions against protocol before the patient advances. He can approve, flag, or escalate.

The Step Reviewer plays exactly this role. It's not a second executor. It's a different concern: did this step meet the bar? The Engine asks the Reviewer before it agrees to advance.

An *oversight agent* sees across steps. Spots patterns no single executor can.



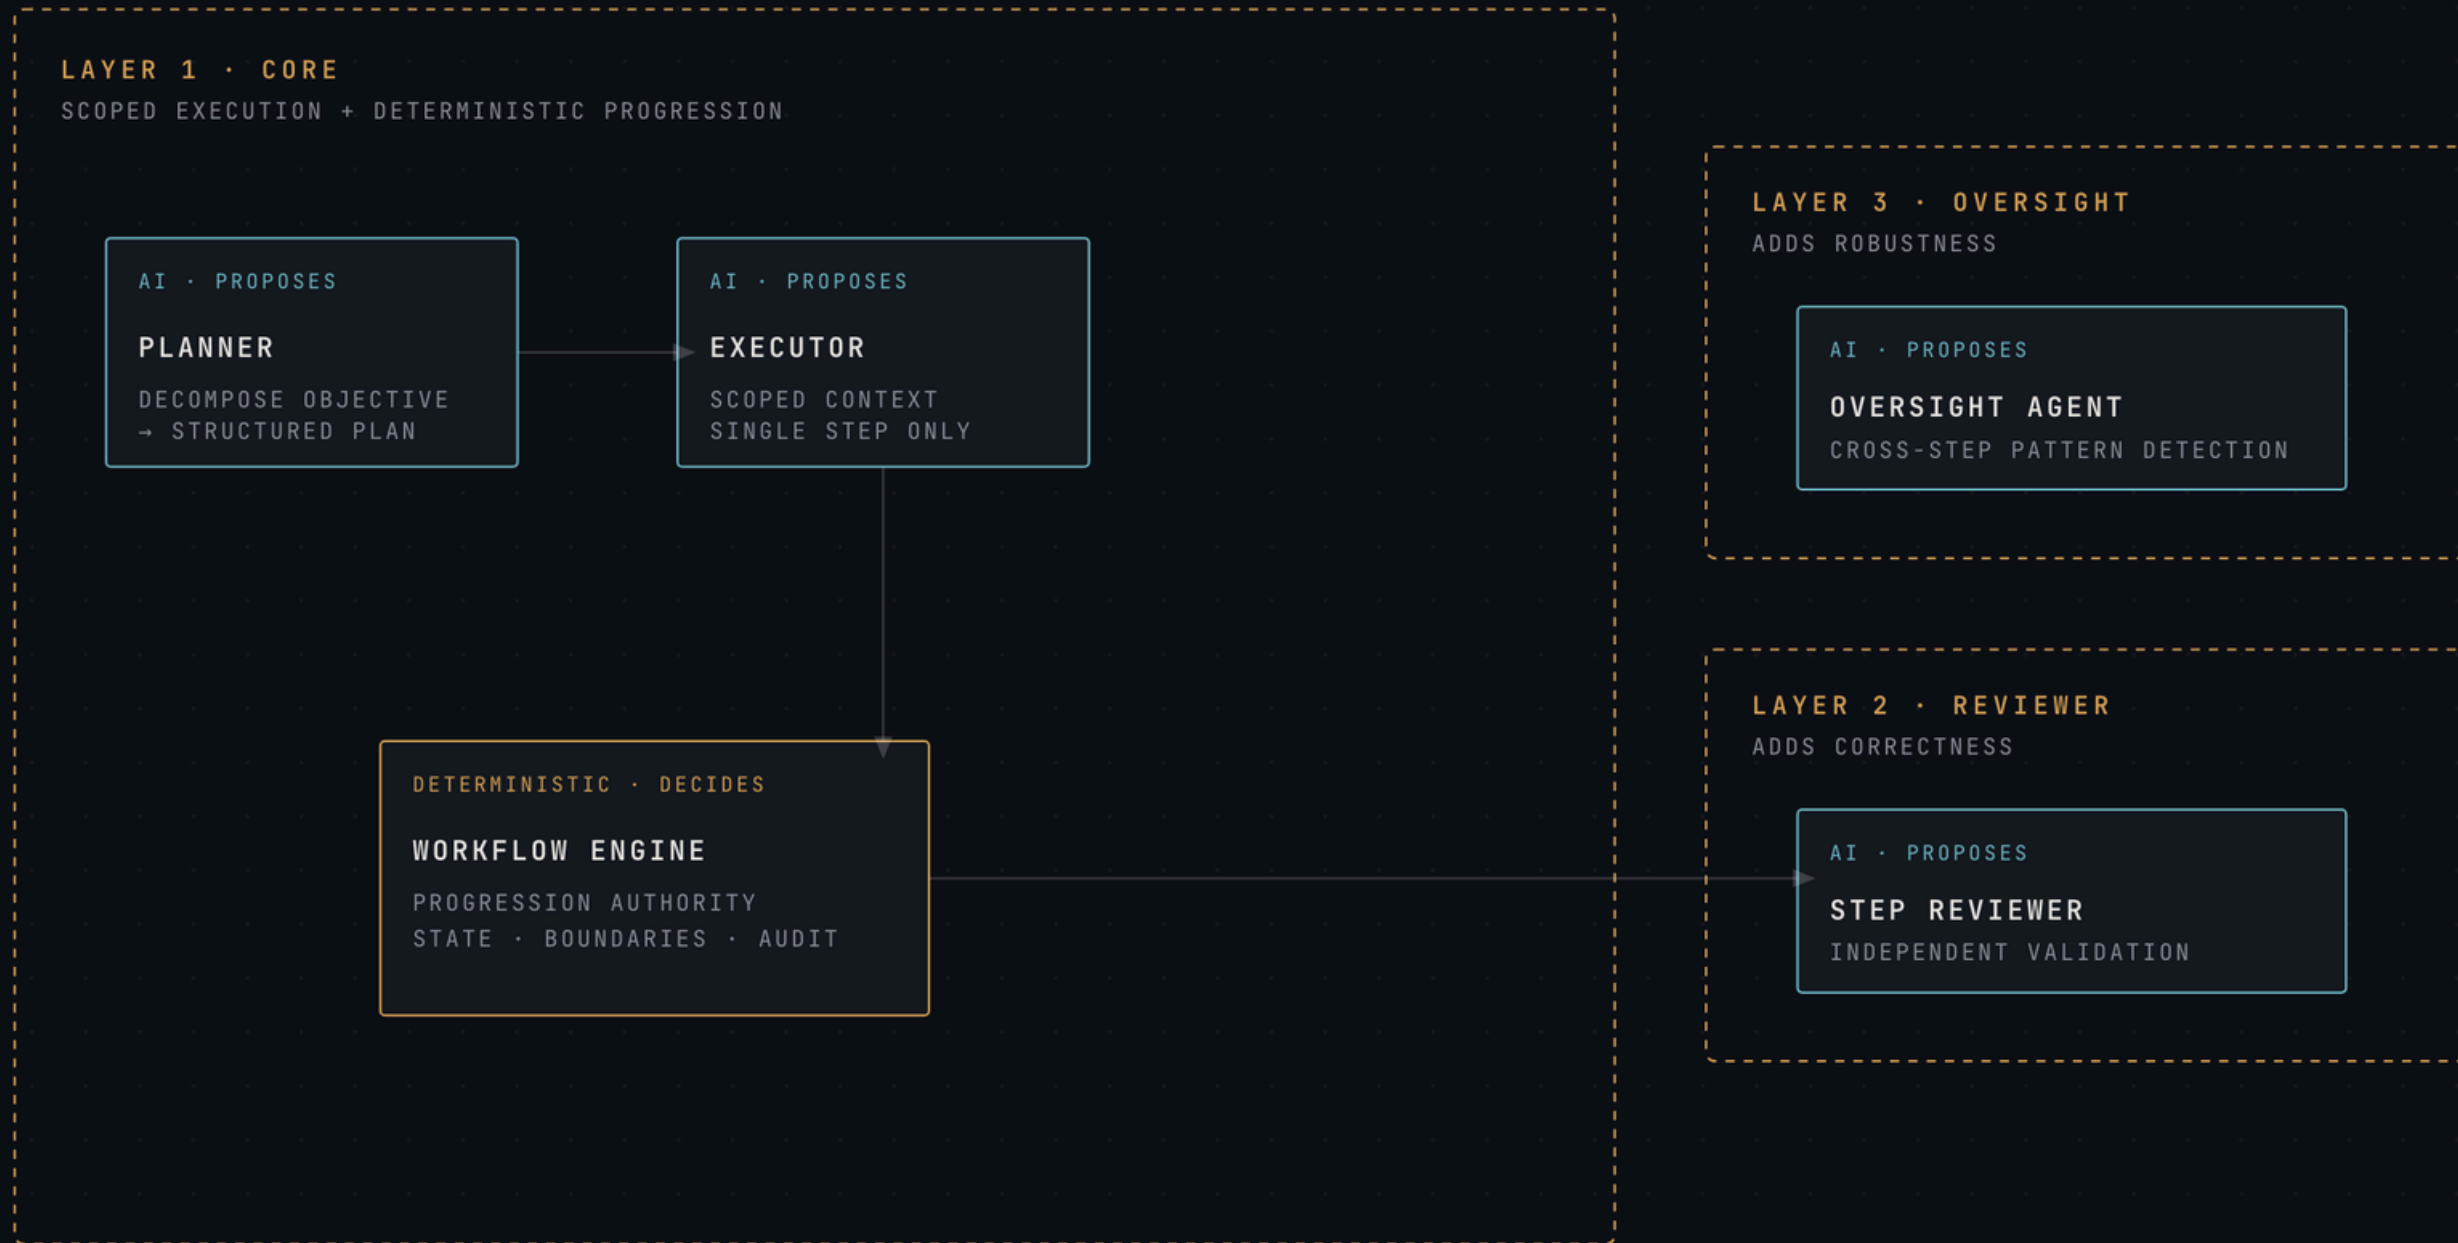
PARAMEDIC PARALLEL

## The Advanced Practice Paramedic.

Some EMS systems have Advanced Practice Paramedics. These clinicians don't do primary work, don't review individual calls, but bring pattern recognition from a high frequency of encounters. They see three patients from the same neighborhood with the same symptoms. They flag systemic issues.

The Oversight Agent is exactly that role: cross-step pattern detection, proactive escalation, replan recommendation when the original plan is failing in ways individual steps can't see.

# The full framework, *one frame.*



# What to take away.

- 01 AI proposes. *Policy-governed systems decide.*
- 02 Authority over time, state, and information flow remains outside AI.
- 03 Every step has explicit contracts defining inputs, outputs, and completion criteria.
- 04 Correctness improves through independent review with tiered criteria.
- 05 Robustness improves through meta-level monitoring and proactive escalation.
- 06 Failures are detected and handled, not hidden or ignored.
- 07 Humans retain control of what matters: policy, escalation, and final authority.
- 08 *Governance is a design principle, not an afterthought.*

Governance is a design  
principle,  
*not an afterthought.*

14 MAY 2026

TALK CONCLUDED

GENAI DAY

THE THESIS

# AI proposes. Systems decide.

SCALING AI AGENTS WITHOUT LOSING CONTROL

PRESENTED BY

**Alan Knox**

VP Product Development · QuantaViable

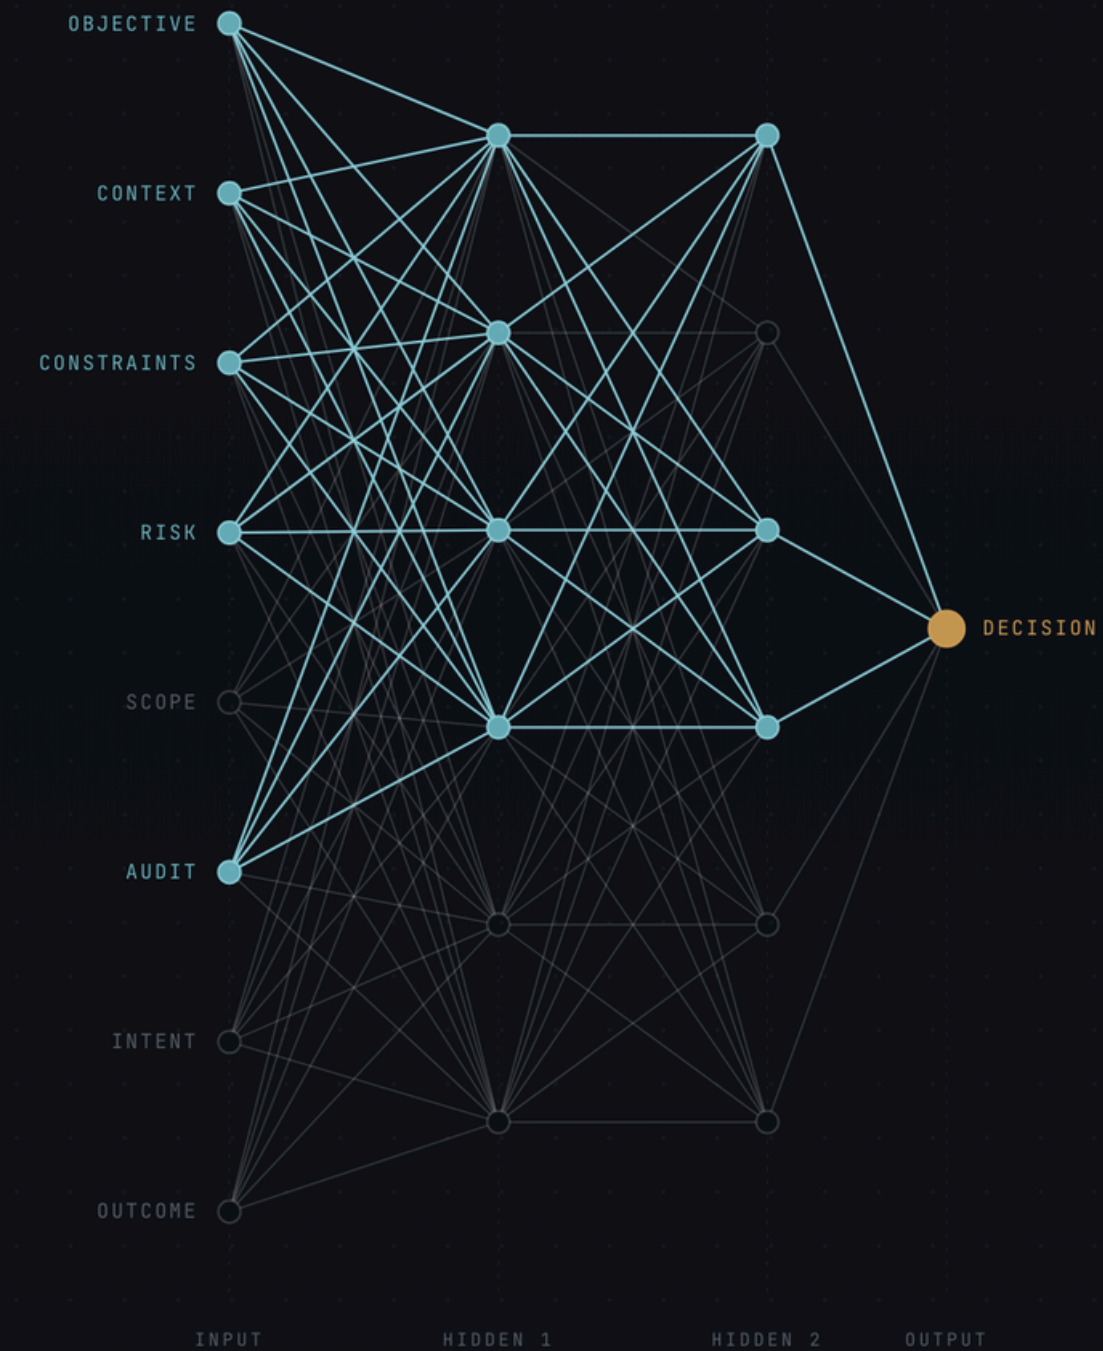
CONTINUE THE CONVERSATION

> [QuantaViable.com](https://QuantaViable.com)

> [BuiltByFathom.ca](https://BuiltByFathom.ca)

FORWARD PASS · ACTIVE

INFERENCE



NETWORK ACTIVE

PASS 120