

FRAMEWORK ON ONE PAGE · VERSION 1.0

Intent Architecture

Authorization architecture · May 2026

TIME TO USE 30 min review / 90 min workshop	AUDIENCE Product, risk, security, and governance leaders	OUTPUT A written intent boundary
--	---	-------------------------------------

?

Before any AI agent is deployed in your environment, does your organization have a formal record of what it was authorized to do, who made that authorization, and under what conditions that authorization must be reviewed?

What it is

Intent Architecture is the organizational design work that must happen before any AI agent goes live. Intent Architecture precedes technical controls. A well-configured agent operating without documented authorization, a named Consequence Owner, and a defined review process is still a governance gap. The configuration is correct.

The three layers

AUTHORIZATION: What the agent is permitted to do	A formal record of the agent's authorized purpose, the specific actions it may take, and the explicit prohibitions that apply regardless of technical capability.
ACCOUNTABILITY: Who owns what it does	A named Consequence Owner, a specific individual who accepted accountability for the agent's behavior before it went live and who is reachable when something goes wrong.
ARCHITECTURE: The design decisions made before deployment	The regulatory environment mapped. The data touchpoints documented. The system integrations scoped. The human review conditions defined.

A worked example

Situation	A generic enterprise is defining the authorized purpose and boundaries of an agent that flags contract clauses.
Where it lands	Intent boundary ready for review
Evidence	The business purpose, prohibited actions, review conditions, and consequence owner are documented before tool access is granted.
Gap found	The review cadence and exception process require final approval.
Next action	Approve the intent record and bind the runtime configuration to its documented boundary.
Output	A written intent boundary ready for authorization.

The one action it produces

A written intent boundary

What good looks like

Any person in the organization, a new compliance officer, an external auditor, or a board member, can be handed the authorization record for any deployed agent and answer the following questions without additional research: what is this agent authorized to do, what is it explicitly prohibited from doing, what data can it access, who approved its deployment, and who is the Consequence Owner accountable for its ongoing behavior.

What it cannot establish

This object can make intent and its boundaries explicit. It cannot prove runtime adherence, assess model quality, or replace technical access controls.

CITE AS