

Nexus Framework: Agent Change Impact Protocol (ACIP)

Purpose: To establish clear boundaries and lifecycle governance for AI-Agent-proposed changes in high-stakes project environments, preventing uncontrolled "Decision Drift".

When AI agents start coordinating models, the danger is not that their modeling is wrong; the danger is that they make *locally optimal* decisions that cause *global drift*. AEC professionals must culturally and operationally shift from "managing clashes" to "governing parameters." ACIP defines the decision lifecycle that agents must pass through, acting as an un-bypassable gate.

The Core Philosophy

ACIP is not another agent; it is the decision gateway agents must pass through. Every machine recommendation must trigger an explicit human-in-the-loop review against the project's risk threshold.

The ACIP Action Lifecycle

A proposed machine change must traverse these five gates in sequence: **Draft** → **Simulate** → **Assess Impact** → **Route for Approval** → **Commit** → **Log**

- 1. Policy Reference (The Guardrails)** Agents can *reference* policy (e.g., safety clearances, codes, contractual constraints) but they **cannot override it**.
- 2. Blast-Radius Simulation** Before an agent proposes a physical move, it must map downstream impacts (e.g., "Moving AHU by 400mm triggers 4 pipe clashes and conflicts with a fire exit route"). The human reviews the consequence, not just the action.
- 3. Structured Authority Routing** Depending on the category of the change, the decision routes to the authorized human counterpart:
 - *Spatial/Clash Impact*: MEP Lead / Discipline Coordinator
 - *Support/Anchor Impact*: Structural Lead
 - *Cost/Schedule Implication*: Project Controls
 - *Life Safety Impact*: Fire & Life Safety Consultant
- 4. Permanent Decision Committing and Logging** The human approves the change. A log is written permanently indicating **who** approved it, **what** downstream impacts were accepted, and **why** the machine recommendation was authorized.

The Human as the Governor

In high-stakes AEC execution, the psychological trap of AI adoption is "operator apathy"—teams blindly trusting the machine's confident output or dodging blame when the machine makes an error.

To counteract this structurally, the human authority must elevate. They are no longer an *operator* of the agent; they are its **Governor**.

- **The Agent** calculates alternative paths, checks the hard limits, and prepares the impact consequence report.
- **The Governor** calibrates the acceptable risk profile, dictates the allowed boundaries, and holds the ultimate execution key.

Automating the execution does not automate the accountability.

© 2026 Nexus Engineering. *Field Notes on BIM & Systems Thinking.*