

AGENTIC AI · SCOPE MANAGEMENT

AI Scope Expansion Does Not Look Like a Decision. It Looks Like a Series of Reasonable Accommodations.

A team finds a new use for a system that is already running. The expansion is small, the system handles it well, and no one flags it for review because no review trigger exists for incremental use case additions. Six months later, the system is doing something materially different from what was authorized. The risk profile has changed. The governance documentation has not.

This is how integration depth accumulates. Not through a single boundary-crossing decision, but through a succession of individually defensible steps that no governance mechanism was designed to catch in aggregate.

The authorization assessed one system. The system currently operating is a different one. Both records show compliance.

That is the structural problem GIAG Stream Two is examining — not whether agencies have governance frameworks, but whether those frameworks have any operational relationship to the systems they are supposed to govern.

Most AI scope expansion does not look like a decision. It looks like a series of reasonable accommodations. None of them individually crosses a governance threshold. Together they represent an unauthorized system operating under an authorized framework. The accountability question — when it arrives — will need to be answered against documentation that describes something else.

Government AI in Practice — thinkcapital.substack.com · thinkcapital.org