

The Executive AI Value Review: seven questions per agent

Portfolio discipline. Run every agent, live or proposed, through the same seven questions at the same table. The output is a verdict per agent: expand, redesign or stop. On operational evidence, not enthusiasm.

Seven questions, one verdict. Any question without a written answer is the finding.

01 What decision? "Which consequential, repeating decision does this agent influence? Named, not a process area."	02 Which outcome? "The enterprise result it must serve, and the approved trade-off when it conflicts with another?"
03 Who owns it? "One accountable human. Not a committee, not the vendor, not "the business"."	04 What may the agent do? "Retrieve, recommend, prepare, execute, escalate, and the threshold on each?"
05 Which evidence governs it? "The inputs, in order, and the rules no input is allowed to override?"	06 When must it stop? "The handoff rule in one sentence: confidence floor, conflicting evidence, consequence above threshold?"
07 What proves value? "The operational measure, baselined before go-live, that shows business performance changed?"	THEN The verdict "Expand, redesign or stop. Recorded with the evidence, reviewed on a date."

USE IT QUARTERLY

Seven questions, every agent, same table. The review is where pilot sprawl becomes a portfolio: fewer overlaps, clearer investment decisions, and a list of what to stop.