

Agent Readiness Scorecard

Twelve questions that tell you whether a decision is ready to be automated, and at what tier of autonomy.

[Download PDF](#)[Print](#)

Score each question 0 (no), 1 (partly) or 2 (yes). Total the column.

#	QUESTION	SCORE
1	Is the decision written down, with its inputs and outputs?	
2	Do we know what a wrong answer costs, in money or trust?	
3	Can the action be undone within a window?	
4	Do we have 50+ real past cases with a known good answer?	
5	Is there a named human who owns this decision today?	
6	Can we enforce hard bounds (caps, floors, step limits) in code?	
7	Can the agent be given only the tools it needs, nothing more?	
8	Can we log inputs, proposal, guardrail results and action for every decision?	
9	Is there a clear rule for when to escalate to a human?	
10	Can we measure accuracy and cost per decision weekly?	
11	Would compliance, finance and customers accept the audit record?	

Reading the score

- **0 to 11:** suggest-only. Let the agent draft; a human acts.
- **12 to 19:** act within bounds. Tight step limits, daily review.
- **20 to 24:** act and report. Sampling audits, weekly review.