

Automation Boundary Worksheet (v0.2)

Before an AI system takes on work, answer these on one page. Record the decision so it can be questioned.

System ID		Record ID	
Component(s) in scope		Design / configuration version	
Consequence of error (low / moderate / high)		Linked records (prior record IDs)	
Date		Owner	
Inherited baseline (later cycles: Decision Record ID)			

1. Eligibility

Automate, assist, defer or do not automate — and why?

2. Boundary

What does the system decide alone, and what stays with a person? What properties of the work (severity, reversibility, ability to oversee) set the line? For each action the system takes alone: can it be reversed directly, or would undoing it require a compensating correction?

3. Decision rights

On each side of the boundary, who decides, who is consulted, who can stop or contain the system, and who is accountable? Who may change the rules, who may change the limits, and where is each authorization recorded?

4. Suppliers

Which parts come from a vendor, and what happens when the vendor changes the model, thresholds or features? Who reviews change notices before activation, which changes need approval, and how is activation held until review and testing are complete?

5. Evidence of influence

Who closest to the work reviewed the proposal? What did they contribute, and what changed as a result — or why not?

6. Proportionality and value

Given the consequences of error, how much of this is a paragraph and how much a full review? What value is expected, and what will the arrangement cost to operate and govern, including oversight time?

What this worksheet is not: it is not a scoring model, a maturity model or a validated assessment. It is a record that makes design decisions visible so they can be questioned. Completing it does not establish that the chosen boundary is safe or appropriate; that takes evidence, including the later Enable and Demonstrate stages.