Human and AI Responsibility Map

Make task boundaries, approval, and escalation explicit.

For: Process owners, domain experts, and delivery teams | Use when: When redesigning a workflow or preparing a handoff | Suggested time: 25 minutes

Version 1.0 | October 5, 2026

Use synthetic or authorized information. Keep sensitive work in approved tools. A named person owns each decision. These tools support discussion and planning; they are not validated assessments.

1. Assign responsibility by step

Record the work at the level where a person can tell whether the system acted within its boundary. Avoid assigning accountability to AI.

AI may

Which bounded action is permitted?

AI may not

Which action is outside its authority?

Human review

Who checks what, using which evidence?

Approval owner

Who authorizes the consequential action?

---

2. Specify exceptions

Write a trigger and a response rather than a general promise of human oversight.

Escalation trigger

What observed condition requires a person?

Escalation destination

Who receives the case, and what context do they need?

Stop and fallback

Who can pause the workflow, and how does work continue?

Review cadence

Who checks whether boundaries still fit the work?

Completed companion example

All organizations, workflows, numbers, and results in these examples are fictional and provided for learning. Targets are proposed decision criteria, not validated standards or promised outcomes.

AI may

Find policy passages and draft a cited explanation.

AI may not

Decide eligibility, access personal files, change a record, or send an answer.

Human review

Policy specialist checks the current policy, citation support, and case category.

Approval owner

Service-desk analyst sends only after specialist approval during the pilot.

Escalation trigger

Missing policy, conflicting sources, individual eligibility, or a request for restricted information.

Escalation destination

Policy specialist receives the question, cited excerpts, and the reason the draft was withheld.

Stop and fallback

Service-desk lead pauses drafting; analysts return to approved manual policy lookup.

Review cadence

Owner reviews exceptions weekly during the pilot and after any source change.

Original guide: https://builtinto.work/resources/human-ai-responsibility-map