The Discovery Sprint Worksheet

Baseline one workflow, rank the opportunities, decide with evidence

Nugawi, an applied-AI consultancy in London, runs every engagement through a Discovery Sprint: two weeks, one high-value workflow, a human-verified baseline and an ROI-ranked roadmap. This worksheet is the same structure, free, no email required. Fill it in with your team and you will know whether AI agents belong in your operation yet, and which workflow earns the first build.

1. PICK THE WORKFLOW

Score each candidate workflow from 1 (weak) to 5 (strong) on the four criteria. The highest total wins the sprint.

Candidate workflow	High value (money or risk)	Measurable (numbers exist)	Consequential (someone cares daily)	Volume (runs often)	Total

Typical winners: order-to-cash, procure-to-pay, shipment-to-settlement, document-to-decision, enquiry-to-resolution.

2. MEASURE THE BASELINE (HUMAN-VERIFIED, NOT SELF-REPORTED)

Sit with the people who run the workflow and measure it as it runs today, for at least one full cycle:

- Volume: transactions per month
- Cycle time: from arrival to done, end to end
- Cost per transaction: loaded cost of the minutes it really takes
- Where the hours go: list every hand-off, rekey, wait and check by hand
- Error and rework rate: how often it comes back

Baseline: _____ transactions/month, _____ per transaction, _____ hours/month of hands-on time.

3. RANK THE OPPORTUNITIES BY ROI

For each AI or automation opportunity in the workflow, estimate the saving against the baseline, then rank. Confidence beats size: a small certain win funds the next sprint better than a large speculative one.

Opportunity	Modelled monthly saving	Confidence (certain / likely / speculative)	Effort to build	Rank

4. GOVERNANCE CHECKLIST (FILL BEFORE ANY BUILD)

- Access matrix: who and what may read, write and approve, per system
- Audit trail: every agent action attributed to a requester
- Data residency: UK or EU region chosen and documented
- Human-in-the-loop points: where a person must approve
- Rollback: the rehearsed way back, tested once before go-live

5. GO / NO-GO

The sprint earns a build only if the answer to all three is yes:

- 1. The baseline was measured by people, not estimated from a survey.
- 2. The top-ranked opportunity has a modelled saving bigger than the cost of building it.
- 3. The governance checklist has no unanswered lines.

Use the worksheet with us in a Discovery Sprint, or hand it to anyone else. No lock-in is the point. If you want the pod that wrote it: nugawi.ai/contact.