First AI workflow card

Fill in one card for the first AI workflow you plan to run. Fill in what you know; every blank line gets an owner and a date. Any blank line means: not ready to run.

QUESTION	ANSWER	OWNER / BY WHEN
Which repeating work goes first: documents, quoting, dispatch support or track and trace?		
Volume: how much of it lands each week, and how many people do it?		
What does an error in that work cost today, and is the worst case a correction?		
Boundary test: does the output move a load, change a rate a shipper sees, watch a driver or cross a border? (yes / no)	☐ Yes ☐ No	
If yes, whose bar applies (load, rate, driver, border), and which named person owns every output before it ships?		
Trained like a new hire and checked like a control, corrected hourly at first, then daily, then weekly? (yes / no)	☐ Yes ☐ No	
What already runs that nobody calls AI (TMS solvers, ETA models, yard systems)?		
What does that count change in the budget line?		
Which vendor claims have run on our freight: our lanes, our quote history, our documents?		
Price quoted, and what it buys?		
The number it moves: what moves, by how much, by when?		
Baseline for that number measured before the pilot starts? (yes / no)	☐ Yes ☐ No	
Who checks the number, and how often?		
Decision and date	**☐ RUN ☐ WAIT FOR THE BAR ☐ RE-TEST ON _____ Date _____ Signed _____**	

First AI workflow card

Example (fictional): status-email drafts at a 140-truck regional carrier

QUESTION	ANSWER	OWNER / BY WHEN
Which repeating work goes first: documents, quoting, dispatch support or track and trace?	Drafting status replies to shipper "where's my load" emails; people review and send	n/a
Volume: how much of it lands each week, and how many people do it?	About 450 status emails a week; three customer service reps handle them	n/a
What does an error in that work cost today, and is the worst case a correction?	A wrong status is fixed in the next reply; no load moves and no rate changes	n/a
Boundary test: does the output move a load, change a rate a shipper sees, watch a driver or cross a border? (yes / no)	☐ Yes ☒ No; the reply reports where a load is, and a person sends it	n/a
If yes, whose bar applies (load, rate, driver, border), and which named person owns every output before it ships?	Not needed at No; the customer service lead still signs every reply before it goes	n/a
Trained like a new hire and checked like a control, corrected hourly at first, then daily, then weekly? (yes / no)	☒ Yes; two leads correct drafts hourly for the first two weeks, then daily	n/a
What already runs that nobody calls AI (TMS solvers, ETA models, yard systems)?	The TMS ETA model and the automated end-of-day status report nobody calls AI	n/a
What does that count change in the budget line?	Feed the ETA model's output to the drafts instead of buying a visibility tool; about $40,000 less in the ask	Ops analyst, by 4 Dec
Which vendor claims have run on our freight: our lanes, our quote history, our documents?	**Fails the bar**: the demo ran on the vendor's own sample emails, not our inbox	Vendor manager, by 11 Dec
Price quoted, and what it buys?	About $60,000 a year for the email module, plus the TMS connection	n/a
The number it moves: what moves, by how much, by when?	Reply time from about 4 hours to under 1 hour, and rep time on status emails from about 30 hours a week to 10, by end of Q1	Customer service lead
Baseline for that number measured before the pilot starts? (yes / no)	**Blank: not ready**	Ops analyst, by 4 Dec
Who checks the number, and how often?	Customer service lead samples 25 sent replies a week and logs wrong statuses	n/a
Decision and date	**☐ RUN ☒ WAIT FOR THE BAR ☐ RE-TEST ON _____ · Date 20 Nov · "No baseline yet, and the vendor has not run on our emails."**	Ops director