

Every supplier bill approved by AI. And by a person.

A self-hosted n8n workflow that reads a supplier invoice, checks it against the approved-supplier record, files what is safe, and puts everything else in front of a human — with a written audit trail for every decision. No cloud AI bill. No invoice paid by a machine.

3.5sPER INVOICE, END TO
END**0**

API COST PER DECISION

100%

DECISIONS AUDITED

8/8

GUARD TESTS PASS

THE PROBLEM

Small businesses approve bills by email and memory. The limit list lives in someone's head or a spreadsheet nobody updates. Bills get paid twice, paid late, or paid to a supplier whose account is on hold. When someone asks “who approved this?”, the answer is a search through a mailbox.

WHAT IT DOES

1 · Capture

Staff paste the invoice into a form, or it arrives by mail.

2 · Check


A deterministic lookup reads the supplier limit and account status from the client's own spreadsheet — the model never fetches or guesses it.

3 · Decide

The local AI agent adjudicates and explains. A second code check refuses any “approved” the record does not support.

4 · Route

Safe bills are filed to the ledger. The rest go to a human by email: APPROVE or REJECT. Every event is written to a Sheet and to an append-only file.



The invoice gate as built: 24 nodes in n8n, three workflows in total (gate, reply handler, error alert). Config and supplier limits are read from the client's own spreadsheet at the start of every run.

THE HUMAN STEP IS THE PRODUCT


The AI prepares the decision and writes the reason. It never pays, never files a bill it was not allowed to file, and never sends the outbound message. A person is the last word.

WHAT IT DOES NOT DO

No payment. No write-back to an ERP. No PDF attachment parsing yet — today the invoice text is typed in or read from the email body. Those are quoted separately, not implied.

EVIDENCE — FROM RUNS, NOT A MOCK-UP

CHECK	RESULT
Invoices approved automatically	2 of 5 cases: known supplier, active account, amount inside the limit → ledger row + local audit line
Invoices stopped for review	3 of 5: over the limit, unknown supplier, supplier account inactive → exception row + approval email, each with the failed rule written down
Reply loop	12 human replies processed, one audit row per reply, one Sheet update per invoice; a second pass adds nothing (idempotent)
Config is live	A limit changed in the Sheet (2,500 → 1,000) rerouted a \$1,840 invoice to review on the next run. No workflow edit, no restart
Approver is live	Approver changed in the Sheet; the next approval email went to the new address
Guard	8/8 tests: model approvals above the limit, for unknown suppliers, for inactive accounts, below the confidence floor, wrapped in markdown, or in prose are all downgraded
Failure visibility	A real trigger-mode failure produced one alert email with workflow, node, execution id and stack



The reply handler: watches the mailbox, updates the Sheet, writes the audit row, labels the email so the same reply can never be processed twice.

WHY IT IS CHEAP TO RUN

AI	Local model (Ollama, qwen2.5:7b) on the same small server. No per-token API bill, no data leaves the machine for the AI step.
Google	Gmail API + Sheets API on the client's own account. No charge at this volume.
n8n	Self-hosted, free. One instance per client.
Hosting	A \$6/month VPS, or the client's own machine
Data	Invoice data stays in the client's Google account. The audit is a Sheet they own.

OPTIONS

Pilot \$750–1,200 One flow, one approver, 2 weeks. Fixed price.	Standard \$2,500–3,500 Full install: config tab, reply loop, error alerts, staff guide.	Support \$150–400/mo Monitoring, changes, monthly report.
--	--	--

Prices are estimates for a single-entity business; scope is fixed in writing before work starts.

HOW TO START

A 20-minute call with one real bill. We map the current approval path, agree the limit table and the approver, and I install a Pilot against your own Google account. You keep the Sheet, the workflow and the audit.