

STANDARD OPERATING PROCEDURE

The Lead Qualification Autopilot

AI reads your inbox. Scores the lead. Drafts the reply. You hit send.

The complete operating manual for an inbound lead system powered by n8n, Gmail, Firecrawl, Claude, and Airtable. Every new email is identified, enriched, scored 0–100, and answered with a personalized draft.

Five minutes of your attention per day. Zero leads lost in the inbox.
Nothing sends without your approval.

SECTION I

What This System Does

Every minute, the workflow checks the connected Gmail inbox for new unread mail. Real inquiries get identified, enriched with a scrape of the sender's company website, scored 0–100 by a hybrid of hard rules and AI judgment, and answered with a personalized 120–180 word draft that references specifics from their email and their business. The draft lands in Gmail Drafts, the lead lands in your Airtable CRM, and everything that came through the door gets logged, including the newsletters and vendor pitches that were never leads at all.

THE FLOW

```
New unread email in Gmail
|
Identify sender: name, email, company domain
|
Firecrawl: scrape company site  <- real context, not guesswork
|
Score 0-100: rules + AI judgment
|
Claude: write personalized reply
|
Gmail draft                      <- the human switch. never auto-sent
|
Airtable CRM upsert  +-> Activity Log (every email, even non-leads)
```

THE RULE THAT MAKES THIS WORK

The draft is the design, not a training wheel. Speed-to-lead is the whole game on inbound, and this makes the default outcome a ready-to-send reply instead of an email sitting unanswered for two days. The Gmail draft keeps your judgment in the loop at a cost of five minutes a day, and the AI explains every score in one sentence so you can trust the queue instead of re-reading every inquiry.

SECTION II

Your Five Minutes a Day

The system runs itself. Your job is the last inch: review, personalize if you want, send. Work the queue top-down by score.

01 Open the queue

Airtable → **Lead Qualification CRM** → **Leads**, sorted by Score. Anything with Status **Draft Ready** is waiting on you.

02 Read the two sentences that matter

The **Score Reason** tells you why the AI rated the lead the way it did. The **Enrichment Summary** tells you what it learned from their website. Ten seconds each.

03 Send the draft

Gmail → **Drafts**, find the reply addressed to that lead. Edit if needed, add your name after “Best,” and hit **Send**. Then flip the Airtable Status to **Contacted**.

04 Close the loop

As conversations progress, update Status to **Qualified** or **Disqualified**. That keeps the CRM honest and gives you real conversion data on what the scorer gets right.

HOT LEADS FIRST

Aim to answer **HOT 70+** leads within a couple of hours. Speed is the single biggest lever on inbound conversion, and it is the one thing the automation cannot do for you. **WARM 40-69** can wait for the daily pass. **COLD <40** is a judgment call: skim, reply if there's a spark, disqualify if not.

SECTION III

The Airtable CRM

Two tables. **Leads** is your working queue, one row per lead, upserted on email address so a repeat inquiry updates the record instead of duplicating it. **Activity Log** is your paper trail, one row per processed email, leads or not.

Leads table

FIELD	WHAT IT TELLS YOU
Score / Tier / Score Reason	The 0–100 score, Hot/Warm/Cold tier, and the AI's one-line explanation for the number.
Status	New → Draft Ready → Contacted → Qualified / Disqualified. Draft Ready is your action signal.
Enrichment Summary	What the AI learned about the company from its website.
Original Email	The inbound message, preserved verbatim. Never rely on a summary when the source is one click away.
Draft Subject / Draft Body	A copy of the outreach draft. The live version sits in Gmail Drafts.

Activity Log table

One row per email with the outcome (**Processed**, **Skipped – Not a Lead**, or **Error**), the score, details, and the n8n execution ID. That ID is your tracer: any row in the log maps straight to a run in n8n → Executions, so you can see exactly what happened at every node.

WHY NON-LEADS GET LOGGED

Newsletters, spam, and vendors selling to you never generate drafts, but they still get a row. **Nothing disappears silently.** When a real lead seems to have vanished, the log tells you in seconds whether it was filtered, errored, or never arrived, instead of you wondering whether the system ate it.

SECTION IV

How Scoring Works

Half the score is deterministic rules you can audit line by line. The other half is AI judgment on the things rules can't see. The split is deliberate: rules keep the floor honest, judgment finds the signal.

RULES FIRST — 0 TO 50 POINTS

Business email domain		+20
Buying-intent keywords (budget, timeline, pricing, hire... +5 each)		+25
Substantive message length		+5

AI JUDGMENT — 0 TO 50 POINTS

Clear project need		+15
Decision-maker signals		+10
Stated budget or timeline		+15
Fit, judged from their website		+10

Tiers: **HOT 70+** **WARM 40-69** **COLD <40**. The AI also filters non-leads out entirely; those are logged as Skipped and no draft is created.

WHY YOU CAN TRUST IT

Every score ships with a one-sentence reason, written by the model that assigned it. When a lead scores oddly, read the reason before you assume the system is broken; most of the time it saw something you skimmed past. When the rubric is genuinely wrong for your business, it's plain English in one node, and you tune it there (Section VI).

SECTION V

Troubleshooting

Every failure mode has a first place to look. Start here before opening a single node.

SYMPTOM	CHECK
No new leads appearing	Is the workflow Active in n8n? Is the email actually unread? Only unread mail triggers a run, and the Promotions and Social tabs are excluded by design.
A lead scored oddly	Read the Score Reason in Airtable; the AI explains itself. If the rubric itself is off, tune it in the “Score Lead & Write Draft” node’s system message.
No website summary	The sender used a personal email (gmail, yahoo, etc.), or the site blocked scraping. Either way the workflow continues; enrichment failure is non-fatal.
Something failed mid-run	n8n → Executions shows every run with the failing step highlighted. Cross-reference the Execution ID column in the Activity Log.
Airtable steps fail with “Forbidden”	The Airtable personal access token doesn’t include this base. At airtable.com/create/tokens , add the “Lead Qualification CRM” base with <code>data.records:read</code> and <code>data.records:write</code> scopes, or grant the whole workspace.

THE DEBUGGING SHORTCUT

The Execution ID in the Activity Log is the bridge between “a lead looks wrong in Airtable” and “here is the exact node that produced it.” One lookup replaces guessing. If you run the n8n connector in Claude, you can also just paste the ID and ask Claude to read the execution and fix whatever broke.

SECTION VI

Tuning & Handoff

Everything an operator would ever want to change lives in five handles, all in the n8n workflow editor. Handing this system to a client means pointing at this page.

HANDLE	WHERE
Which inbox	"New Inbound Email" node → Gmail credential. Swap to the client's account here.
What counts as a lead, rubric, draft tone	"Score Lead & Write Draft" node → System Message. Plain English; edit freely.
Keyword signals	"Normalize Lead" code node → <code>KEYWORDS</code> list.
CRM destination	The three Airtable nodes (base and table IDs).
Polling frequency	Trigger node. Currently every minute.

Safety rails built in

Email never auto-sends; output is always a Gmail draft. **No duplicate CRM records;** upsert matches on email address. **Nothing dropped silently;** skipped and errored emails are logged. **Enrichment failure is non-fatal;** if a website can't be scraped, scoring proceeds without it.

Running costs

n8n cloud on your existing plan. Airtable free tier. Claude API at roughly \$0.01–0.03 per lead on Sonnet. Firecrawl at one credit per business-domain lead, with the free tier covering about 500 a month. Gmail is free. The marginal cost of never losing a lead again rounds to a coffee a month.

Rather have it built for you?

This SOP documents one system from the Soulwire Studio workflow library. We build AI automations for businesses that run on process: intake, scheduling, content, compliance-heavy operations. If you would rather spend your five minutes sending replies than building pipelines, book a workflow audit and we will find the automation hiding in your business.

soulwirestudio.com · [Book a workflow audit](#)