

The Nurture Engine

A starter kit for building your own, with whichever AI you already use. It watches the people you care about, decides when there's a real reason to write, drafts the note, and puts it in front of you. You decide whether to send.

Most independents I talk to have the same problem. I speak at a conference, I get one project or none. The relationships are there. What's missing is knowing what to say, and when. Saying nothing reads as indifference. A generic check-in reads as desperation.

Mine runs every day across 56 people: clients, referral sources, prospects. On October 14 it drafted six notes and put the other fifty on hold, each with a one-line reason.

The one design decision that matters: it drafts everything and sends nothing. It takes away the work of figuring out *what to say and when*. It never takes away the decision of *whether to send*. Keep that gate. The moment it sends on its own, it stops being your voice.

What's in here

1. **Monday-morning version:** no automation, twenty minutes, any AI with web search. Start here.
2. **The full build prompt:** paste it into Claude, ChatGPT or Gemini and it walks you through building the whole engine.
3. **What broke in mine:** five failures from production and the rule each one became. They're already in the prompt.

In September I promised the Umbrex room results that aren't mine. I don't have them yet. If you build this, you could be the first, and I'd like to hear what it finds or where it fights you. — Waz

01 The Monday-morning version

No databases, no automation. One prompt, ten people, twenty minutes. Run it for three Mondays before you build anything bigger. If it isn't earning its keep by hand, a schedule won't save it.

Use an AI that can search the web: Claude with search on, ChatGPT with search, Gemini, or Perplexity. Fill in a block at the bottom for each of the ten people you would most regret losing touch with.

PASTE THIS

I want you to act as my relationship chief of staff. Below are ten people I care about professionally. For each one, do the following, in order:

1. SEARCH the web for anything from the last 14 days about this person. Every search must include their organization or their work email domain alongside their name. Ignore any result that doesn't mention one of those. People share names; I would rather you find nothing than find the wrong person.
2. DROP anything about death, illness, lawsuits, arrests or family matters. Do not mention it to me and never use it as a reason to write.
3. JUDGE. Is there a genuine, specific reason to write to this person this week? A real reason is: they changed roles, published or spoke somewhere, their organization announced something, I owe them something (see my notes), or they wrote to me last and I haven't replied. "It's been a while" is not a reason by itself. Expect the answer to be HOLD for most people. If you find a reason for everyone, you are not judging hard enough.
4. For each person you judge WORTH WRITING, give me:
 - why today, in one line, with the source link
 - a short draft (under 120 words) that names the specific thing, connects it to something I can actually offer them (see my notes), and ends with one real question. Never "just checking in." No pitch.
5. For each person you HOLD, give me one line saying why, and what would change your mind.

Return the WRITE list first, then the HOLD list. If research returns nothing for someone, say so plainly. Do not invent a reason.

MY PEOPLE (repeat this block for each of the ten)

Name:

Organization:

Work email domain:

Who they are to me:

What I could offer them:

Open loops (who owes whom what):

How to know it's working

- Most people come back HOLD. If all ten get a draft, tighten step 3.

- You'd send at least one draft with light edits. If you wouldn't send any, your "what I could offer" column is too thin. That column does most of the work.
- The HOLD reasons sound like something you'd say yourself. That's what earns your trust in the WRITE list.

02 The full engine

Six linked tables and one daily workflow. Paste the prompt below into your AI and it interviews you first, then builds with you step by step.

Your people everyone not marked dormant	Research last 14 days, anchored to their org	Judge worth writing, or hold	One email every draft in full, plus the holds	You decide send, edit, or kill
---	--	--	---	--

Table	What it holds, and why you need it
People	Who they are, organization, email domain, stage, what you could offer, open loops, your next step. You write this table; the machine only reads it.
Signal Log	Every signal found, including the ones it rejected. That's your audit trail.
Content Library	What you can offer: writing, tools, introductions, case studies. Each one is tagged by the problem it speaks to, so drafts pick a real offer and don't invent one.
Outreach Queue	Drafts awaiting review: drafted, approved, sent, or killed.
Engagement History	Every touch, both directions. It stops you repeating yourself and tells you when to stop.
Conversion Tracker	What turned into a conversation, and what came of it. Without this, the whole thing is theatre.

Mine runs on Notion for the tables, n8n for the daily workflow, Tavily for web research, Claude for the judging and drafting, and Gmail for the digest. Each piece has a substitute: Airtable or Google Sheets for the tables, Make or Zapier or a scheduled script for the workflow, any capable model for the judging.

PASTE THIS

Build me a Prospect Nurture Engine: an intelligence-driven system for long-cycle relationships. Every day it researches my people, decides whether outreach is warranted, drafts the message, and sends me one digest email for review. I decide whether to send. It never sends anything on its own.

PHASE 0 – INTAKE. Ask me for all of this before building anything, one or two questions at a time:

- My name, my company, and the email address the daily digest should go to
- Which database: Notion (best, full relations), Airtable, or Google Sheets (simplest, but flat)
- Which automation tool: n8n, Make, Zapier, or a scheduled script
- Which credentials I already have: a model API, a web-research API (for example Tavily or Perplexity), and email sending. If I'm missing one, tell me which one and what it costs before we build around it.
- My pilot person: name, organization, work email domain, who they are to me, what I could offer them, and any open loops between us
- How many days back research should look (default 14)

Do not proceed until I have answered.

PHASE 1 – SIX LINKED TABLES

1. PEOPLE: name, organization, work email domain, role, stage, what I could offer them, open loops (who owes whom what), last contact date, my next step, a PAUSED flag, and a follow-up status (reply owed / due / not due).
2. SIGNAL LOG: source link, date, type, summary, person, and whether it warranted outreach. Keep the rejected signals too.
3. CONTENT LIBRARY: things I can offer, each tagged by the problem it solves.
4. OUTREACH QUEUE: draft, subject, the signal that triggered it, the content offered, and status (drafted / approved / sent / killed).
5. ENGAGEMENT HISTORY: every touch in both directions, with date and channel.
6. CONVERSION TRACKER: what became a conversation, and what came of it.

Link them to each other, not as six flat lists. If the database has no relation fields, use a consistent key and tell me plainly that I lose the linking.

PHASE 2 – THE DAILY WATCHER (one scheduled workflow, every day)

- a. Pull everyone in PEOPLE who is not dormant. SKIP anyone marked PAUSED or whose next step says no outreach, and list the skipped names at the bottom of the digest.
- b. RESEARCH each person over the last N days. The query MUST include their organization or email domain next to their name. Keep only results whose title, link or text mentions that organization or domain.
- c. FILTER OUT anything about death, illness, diagnosis, funerals, lawsuits, arrests, divorce or custody before the model ever sees it.
- d. JUDGE. Give the model the research plus everything in my PEOPLE row (what I could offer, open loops, follow-up status, my next step) and the recent ENGAGEMENT HISTORY. It decides: WRITE or HOLD. A touch is due when a reply is owed, a follow-up is overdue, or there is fresh, verified news. "Just checking in" is never a reason. Most days the answer should be HOLD; a system that finds a reason every day trains me to ignore it.
- e. DRAFT for each WRITE: name the specific signal, connect it to one real item from the CONTENT LIBRARY or my open loops, end with one clear question. Cap the day at 6 drafts, ordered by reply owed first, then overdue, then news.
- f. WRITE signals to SIGNAL LOG and drafts to OUTREACH QUEUE. NEVER write into my own columns in PEOPLE (next step, notes, what I could offer). The machine reads my columns; only I write them.
- g. COUNT CHECK: the number of people judged WRITE must equal the number of drafts produced. If they differ, say so at the top of the digest.
- h. SEND ONE DIGEST EMAIL: the counts (drafts / people reviewed / skipped), each draft in full with its reason and source link, then every HOLD with a one-line reason, then the skipped list. One email a day, never one per signal.

PHASE 3 – THE CLOCK

Alongside the news trigger, track time since last contact and surface anyone going cold even with no news. Use a graceful sequence: an early value-add touch, a direct note with a real question, then a clean close into long-term nurture. Never the same message twice. Tell me plainly when to stop with someone.

NON-NEGOTIABLE RULES

- Everything is drafted, nothing is sent. The approval gate is the product.
- Every draft carries something. If there is nothing to offer, there's no message.

- Show me what it rejected, or I won't trust what it surfaces.
- If research finds nothing for someone, say so. Never invent a reason.
- Never mention anything sensitive, even if it slips through the filter.

BUILD ORDER

Build the tables and seed my pilot person. Build the watcher and run it once against that one person while I watch, including the signals it rejects. Then five people. Then the schedule. At every step tell me what credentials you need and what it will cost per month.

Start with Phase 0. Ask your first question and wait.

03 What broke in mine

All five happened in production. Every rule in the prompt above traces back to one of them.

What happened	Why, and the rule it became
It told me a contact had died.	It hadn't happened. The search ran on her name alone and found an obituary for a stranger with the same name. Rule: every search is anchored to the person's organization or email domain, and sensitive results are dropped before the model sees them.
It found seven people worth writing to and showed me one.	One line of code kept only the first item. For ten days the digest said "0 drafts" while the engine was quietly doing its job. Rule: the count it judged must match the count it drafted, and the digest says so when they differ.
It rewrote my notes, then believed them.	Its draft overwrote my "next step" field. The next day it read its own words back as my instruction. Rule: the machine never writes in my columns.
It matched people to the wrong things.	Someone with a common surname got matched to a public company with the same name. Rule: keep only results that mention the person's organization or domain. Finding nothing beats finding the wrong person.
Its drafts didn't know me.	Early drafts used only the news, not what I'd promised or what the person owed me. They read like a stranger wrote them. Rule: give the model the open loops and what you could offer. That context is what makes the draft sound like you.

The research was never the hard part. The judgment around it is: what it's allowed to see, what it's allowed to write, and how you check it.

Start small, on purpose

1. Run the Monday-morning version by hand for three weeks.
2. Build the full engine and run it on one person while you watch. Read the signals it rejects.

3. Five people for a week. Then everyone, on a schedule.

More free kits, including two n8n workflows you can import directly (no email, no signup):
wazirigaruba.com/kits

Tell me what it finds, or where it fights you: linkedin.com/in/waziri-garuba

Waz Garuba · Harlem Labs · October 2026. Free to use and adapt.