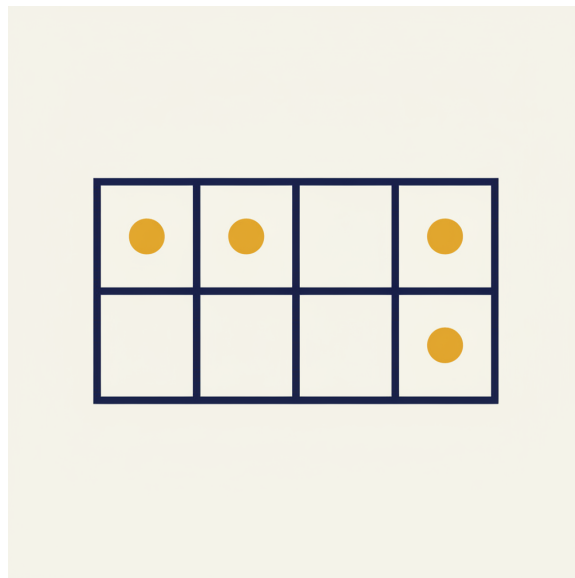




TIER 1 - FOUNDATIONS * V1.0 -- MAY 2026

AI FOR LIFE ADMIN

The boring-but-useful applications that compound into hours-per-week back. Calendars, planning, decisions, household logistics -- the operating procedure for using AI as a personal executive assistant.



BY

Alex Jahn / Agent Logic

v1.0 -- May 2026

Anyone who pays for AI but uses it for fun stuff while real-life logistics still eat their evenings

15-20 minutes

Free. Forever.

EDITION

AUDIENCE

READ TIME

COST

Prepared by Agent Logic / alexanderjahn79@icloud.com / theaiguywi.com

CONTENTS

What's in here

- 1 The 5-hour rebate you're not collecting 3**
Pull up your credit card statement. Find the AI subscription. ChatGPT, Claude, whichever one. Twenty bucks a month. Maybe forty if you...
- 2 What "life admin" actually means 4**
A working catalog. The boring parts of running a life that AI can take a real bite out of:
- 3 The three categories AI handles really well 5**
Out of the catalog above, three meta-categories cover almost everything. Knowing which category you're in tells you what kind of prompt...
- 4 The "executive assistant" prompt pattern 6**
A real personal assistant works for you for months. They learn your kids' names, your spouse's schedule, your business hours, your...
- 5 Five worked sessions 7**
The catalog made abstract. Five concrete prompts you can adapt to your own life today.
- 6 What NOT to delegate 9**
Two boundaries inside life-admin where AI is the wrong tool -- straight from Module 4 (When NOT to Use AI) but worth restating in this...
- 7 The compound effect 10**
Now the math.
- 8 Where to go from here 11**
Tier 1 is now complete. Six modules, in order:

SECTION 1

The 5-hour rebate you're not collecting

The bill says \$20. The rebate says zero.

Pull up your credit card statement. Find the AI subscription. ChatGPT, Claude, whichever one. Twenty bucks a month. Maybe forty if you have two.

Now ask yourself:

If you're an honest answer, it's probably a few minutes of curiosity ("write me a poem about my cat"), maybe one work task, maybe a quick lookup that could have been a Google search. Net rebate against your \$20: somewhere between zero and ninety minutes of life back.

That's not what the tool can do. The tool can hand you back logistics -- the boring stuff that eats your evenings and you don't even notice anymore. Errand planning. Meal planning. Routine emails. Vacation logistics. The decisions that don't deserve 30 minutes of thought but get them anyway.

Most people don't collect this rebate because nobody told them how. The "fun" use cases got all the marketing. The "boring useful" ones are sitting on the table. This module collects them.

What this module is

Twelve pages on the unsexiest, highest-impact application of AI in your daily life:

- A clean catalog of what "life admin" actually means.
- The three categories of life-admin work AI handles really well.
- The single prompt pattern -- context-once, query-many -- that makes AI feel like a real personal assistant.
- Five worked sessions: weekly planning, meal planning, vacation logistics, routine email, decision-narrowing.
- The boundaries: what NOT to delegate to AI even when you can.
- The honest math on what you actually get back.

This is the last module in Tier 1. By the end you'll have the foundation in your hands and a working personal multiplier in your daily life. Then Tier 2 starts: applying this same stack to your actual job.

The flashy AI demos got all the airtime. The biggest gain in your daily life is unphotogenic: it's the 47 minutes of meal planning you used to do every Sunday and don't anymore.

SECTION 2

What "life admin" actually means

A working catalog. The boring parts of running a life that AI can take a real bite out of:

- **Calendars and scheduling.** Resolving conflicts, blocking time, reshuffling around constraints, drafting calendar invites with the right context, finding the realistic slot for a thing.
- **Errand and task planning.** Turning a mess of "things I need to do this week" into a sane multi-day plan that clusters geographically and time-of-day-appropriately.
- **Household logistics.** Meal planning around constraints (allergies, preferences, time, budget). Grocery lists. Routine maintenance reminders (filters, cars, seasonal house tasks). Repair scheduling.
- **Family communication.** Drafting messages to the school. Summarizing kid permission slips. Coordinating logistics with co-parents. Polite but firm replies to the soccer coach.
- **Decision-making at the small-stakes level.** Which insurance plan from the three the broker sent. Which dishwasher in your price range. Which restaurant for the dinner party. Anything that's a 30-minute decision you don't want to spend 30 minutes on.
- **Routine email triage.** "Reply to this with a polite no." "Summarize the actually-important parts of this newsletter." "Draft a follow-up to that contractor."
- **Recurring planning loops.** Weekly menu. Monthly budget review. Seasonal home maintenance checklist. Birthdays / anniversaries / reminders.

That's a lot of boring stuff. It's also where most of your week's friction actually lives. The unsexy truth: getting this back is more life-changing than any "10x productivity" demo.

What it doesn't include

Two bright lines for what's NOT life admin (these get full sections later):

- **Anything privacy-sensitive.** Account numbers, SSNs, custody-level family detail, financial-account-level money management. Module 4 (*When NOT to Use AI*) covers the privacy red line; it applies in full here.

- **Emotional-labor work.** Hard conversations with kids. Conflict with a spouse. Caring for a sick parent. The model can help draft, summarize, and remind, but the human work is still yours.

Inside those lines, you can hand a

huge amo

SECTION 3

The three categories AI handles really well

Out of the catalog above, three meta-categories cover almost everything. Knowing which category you're in tells you what kind of prompt to write.

Three categories of life-admin work AI handles really well:

1. **Sorting** -- Turning a messy list into an organized plan. (errand sequencing, calendar restructuring, task prioritization)
2. **Drafting** -- Writing the message you've been avoiding. (email replies, calendar invites, school-form fill-out, family logistics texts)
3. **Deciding** -- Narrowing options when the decision doesn't deserve full attention. (insurance plan picks, appliance choices, restaurant selection)

If your life-admin task fits one of those categories, AI is reliably the right tool for the first pass. If it doesn't fit -- emotional work, judgment-heavy decisions, anything that needs your specific real-world skin in the game -- that's the category Module 4's "judgment red flag" covers, and AI is the wrong tool.

Why these three work

- **Sorting** plays to AI's strength of pattern-matching to "what does a sane plan for X look like." Combine your messy inputs with the constraints, and the model produces a 3-day plan in 10 seconds. You'd spend 20 minutes doing the same.
- **Drafting** plays to its strength of producing fluent prose conditioned on context. The boring email you've been avoiding is the one where you don't know how to start. The model gives you a draft to react to, which is much easier than starting from scratch.
- **Deciding** plays to its strength of generating "reasonable options across a space" and explaining tradeoffs. It can't make the call for you. But it can compress 30 minutes of research into 3 minutes of focused options.

The pattern:

AI is goo

actual judgment.

SECTION 4

The "executive assistant" prompt pattern

Context-once, query-many

A real personal assistant works for you for months. They learn your kids' names, your spouse's schedule, your business hours, your preferences, your constraints. By month three, you can tell them "book the dinner thing" and they know exactly what to do.

You can give an AI assistant a version of that same context -- once. From then on, every life-admin prompt benefits, because you stop re-explaining your life every single time.

The "context block" template:

Save this somewhere you can paste it. Update it every few months.

``` ABOUT ME (for AI life-admin help):

- Family: [partner name + role / kid names + ages / pet]
- Work: [your job + typical hours + busy/calm seasons]
- Calendar: [what's locked / what's flexible / your "deep work" hours]
- Geography: [your town / nearby errand hubs / commute realities]
- Budget posture: [tight / comfortable / specific limits]
- Preferences: [any consistent things you avoid or prefer -- diet, modes of transport, communication style, working hours]
- Recurring constraints: [school pickup, standing meetings, etc.] ```

Paste that block at the top of any life-admin prompt. The same prompt with the context block produces dramatically better output than without -- because the model isn't guessing at the average user, it's working for

*you.*

## Why this matters

Most life-admin failures come from the model giving you generic advice. "Cluster your errands geographically" is true but useless if it doesn't know your geography. "Plan around your meetings" is true but useless if it doesn't know your meetings.

Once you front-load context, the prompts on the back end can be tiny.

"3-day plan" becomes a one-liner that produces a personalized plan, because the context block already taught the model who

"Sort this  
you are.

## Maintenance -- quarterly refresh

The context block stops being accurate as your life changes. The kids get older. The job changes. The schedule shifts. Re-read it once a quarter, update what's stale, and keep using it. It's the only "personal AI memory" you really need at this level.

## SECTION 5

# Five worked sessions

---

The catalog made abstract. Five concrete prompts you can adapt to your own life today.

## Session 1 -- Weekly task sort

The "mess of errands -> 3-day plan" pattern from Module 2, polished for life-admin use.

*"Want: Sort this week's task list into a 3-day plan (Tue-Wed-Fri).*

*Tuesday morning is locked (job site). Friday afternoon is locked (kid pickup at 3:15). Everything else is flex.*

*tax docs at the accountant, fix the loose handrail at home, draft the proposal for the Schmidt project, schedule the dog's annual checkup, reply to Tom about Saturday softball.*

*errands geographically. Put admin and calls in afternoons (mornings are tools-on time). Mark anything I should push to next week."*

**Context:** [

**Tasks:** ca

**Constrain**

Output: a usable 3-day plan in 10 seconds. You spend the next 10 seconds adjusting the few things that don't match your read on the day.

## Session 2 -- Meal planning around constraints

*"Want: Plan dinners for next week (7 nights). Output as a list with main dish + 2-3 ingredients per night, plus a consolidated grocery list at the end.  
me, partner, kids 6 and 9. Daughter (6) has a tree-nut allergy (hard line). Son (9) won't eat anything green that isn't a pickle. Partner is trying to keep dinners under 600 cal. Wednesday is hectic -- needs to be a 20-minute meal. Saturday we have friends over (4 adults, 2 kids).  
main protein on consecutive nights. No tree nuts in any dish, period. Stay under \$150 total grocery for the week."*

Context: [

Constraint

Output: 7 dinners + grocery list. Maybe two need swapping. Three minutes of prompt + adjustment for ~45 minutes saved.

## Session 3 -- Vacation logistics

*"Want: Help me plan a 4-day road trip from Fond du Lac, WI to Chicago for July 18-21. Output: a day-by-day plan with morning, afternoon, evening blocks; a packing list; and 3 suggested dinners (mid-budget, kid-friendly).  
July. Driving in our car. Booked a downtown hotel near Millennium Park. Daughter gets carsick easily -- needs frequent stops on the drive. Son loves trains and dinosaurs. Partner wants one dinner that's adults-grade (we'll book a sitter through the hotel).  
2-3 hours of unstructured time per day. Suggest only places that make sense given a downtown hotel. Mark anything weather-dependent."*

Context: [

Constraint

Output: a real travel plan you can react to and adjust. Five minutes of prompt time replaces an hour of hotel-research-meets-Reddit-deep-dive.

## Session 4 -- Routine email drafting

*"Want: Draft a polite reply declining a school field-trip volunteering ask.  
block]. The PTA emailed asking me to chaperone the 2nd-grade trip to the science museum on Wednesday March 12. I've already chaperoned twice this year, the date conflicts with a job site, and I'm tapped out. The PTA chair is friendly -- I see her at school events -- and I want to keep the relationship warm.  
over-explain. Mention I appreciate the ask and will sign up next quarter."*

Context: [

Constraint

Output: a draft that takes less editing than starting from scratch. The hardest emails are the polite-decline ones; this is where AI drafting earns its keep.

## Session 5 -- Decision-narrowing on a small stakes choice

*"Want: Help me pick between three appliances. Output a comparison table with the factors that actually matter for my situation, and a recommendation based on my context. context block]. Replacing a 12-year-old dishwasher. Three options: [paste model names + prices + key specs from each]. Family of 4, run dishwasher daily, kitchen is small (so quietness matters), have hard water (need a model that handles it). integration, etc). Focus on durability, water-handling, noise. Recommend one, but tell me when I should override your call."*

**Context:** [

**Constrain**

Output: a focused comparison and a defensible recommendation. The decision still belongs to you (judgment red flag -- this is your kitchen, your money). But you don't spend 30 minutes researching to make a \$700 call.

# 5

## Worked sessions you can run this week.

Pick the two that hit hardest in your actual life. Run them once with your real context. Watch how much time you don't spend.

## SECTION 6

# What NOT to delegate

Two boundaries inside life-admin where AI is the wrong tool -- straight from Module 4 (*When NOT to Use AI*) but worth restating in this context.

## The privacy boundary

Some life-admin tasks involve data that should never enter a third-party log:

- **Money management at the account level.** Bank balances, online-banking credentials, full account numbers, IRS notices with your full SSN, tax-return-level detail.

- **Custody and legal-family material.** Court documents, custody schedules with addresses, anything subject to a custody agreement's confidentiality.
- **Health information.** Diagnostic detail, prescription regimens with names attached, mental-health notes. The kid's allergy ("tree nuts") is fine. The kid's psych eval is not.
- **Anyone else's identifiable data.** Your spouse's salary, your kid's school records with name+ID, your friend's marital problems they told you in confidence.

Substitution rule: redact the load-bearing identifier before pasting.

*allergy*" is fine.

though it's the same fact.

"My 6-yea

"Sarah Ja

## The emotional-labor boundary

This is the subtler one. AI can draft a message. AI cannot  
difference matters more in family life than anywhere else.

be presen

- **Hard conversations.** Talking to your kid about a death in the family. Telling your partner something difficult. Apologizing meaningfully. The model can help you think through what you want to say. Saying it is yours.
- **Comfort and care messages.** When a friend's parent dies, a text from you written by you matters. A polished AI-drafted condolence message is detectable and corrosive to the relationship.
- **Anything where being present is the point.** A handwritten card to your kid's grandma. A note in your partner's lunch. An apology that has to come from you.

The pattern:

**the gift, you handle.**

AI handle

***The point of writing your kid a birthday card by hand isn't the words. It's that your kid knows you spent five minutes thinking about them. AI can give you better words. It can't give your kid those five minutes.***

## SECTION 7

# The compound effect

Now the math.

Take the five categories where AI realistically saves you time on a typical week:

- **Weekly task sort:** ~20 minutes saved
- **Meal planning:** ~45 minutes saved
- **Routine email drafting (3-5 emails/week):** ~30 minutes saved
- **Decision-narrowing (1-2 small decisions/week):** ~25 minutes saved
- **Misc admin (calendar, school forms, etc.):** ~30 minutes saved

That's about 2.5 hours a week, conservatively. Some weeks more (vacation planning, big project decisions, seasonal household tasks). Some weeks less. Average: somewhere around 3 hours.

If you're more aggressive about it -- using AI as a real personal assistant with the context block, hitting all of the categories -- 5 hours a week is a realistic ceiling.

# 250

**Hours per year. 5 hours per week × 50 weeks = 250 hours.**

Six full work-weeks of life back. Not 10x productivity. Not transformative. Just a quarter of one workday a week, every week, forever -- applied to the parts of life that don't deserve to eat your time but currently do.

That's the case for AI literacy that doesn't get made enough. It's not the dramatic productivity story. It's the daily-life-quality story. The hours don't show up on a dashboard -- they show up as you reading to your kid Tuesday night because the meal planning didn't take 45 minutes.

## SECTION 8

# Where to go from here

---

Tier 1 is now complete. Six modules, in order:

1. **What an LLM actually is** -- the mental model.
2. **The 3-question prompt framework** -- the operating procedure.
3. **Reading AI output critically** -- the editorial discipline.

4. **When NOT to use AI** -- the boundaries.
5. **AI as a tutor** -- the personal-learning multiplier.
6. **AI for life admin**

(this one)

You now have the foundation. If you've read all six and applied even half of them, you're already operating ahead of 90% of AI users. Most never get past the prompt-of-the-week stuff.

The next step is

Role-specific templates, multi-step task chains, picking the right model, privacy at work, the team-rollout playbook. If you're a small-business owner or operator, Tier 2 is where the consulting offer starts to make economic sense --

a good place to start.

Tier 2 -- P

Strap In: T

**Get the next module the day it drops: [theaguywi.com/training](https://theaguywi.com/training)**

One email per release. No drip. No spam. Opt out anytime.

If you want this same operating procedure trained into your team -- or installed across your business so the same multiplier applies to everyone, not just you -- that's the consulting offer. We do it the way I run it in my own carpentry business: every employee, every role, with their own context block.

**Reach out: [alexanderjahn79@icloud.com](mailto:alexanderjahn79@icloud.com)**

A short call. Honest scope. We figure out together if it's a fit.

## Closing -- the lock-in line

The last sentence of Tier 1 is the one that matters most:

# 6

## **Six foundations in your hands.**

Mental model, prompt framework, editorial discipline, boundaries, tutoring, life admin. Internalize them and the rest of the curriculum is tactics on top of solid ground. Skip them and every advanced tactic stays wobbly.

Tier 1 is closed. Tier 2 is where it gets practical at work. The same plain English. The same 12-15-page primers. Same hands-on voice. See you there.

### **Agent Logic --**

Fond du Lac, WI. This is module 6 of 6 in Tier 1 (Personal). starts next.

© 2026 Agent Logic. Share freely.

*theaiguyn*

**Tier 1 cor**