

A FIELD GUIDE

The Claude *Code* Field Guide.

Install Your AI Operating System in 30 Days — without burning your weekly Claude.ai quota, hiring an engineer, or learning a framework that'll be obsolete next quarter

AUTHORED BY

Noxlee Automations



NOXLEEAUTOMATIONS.COM

Table of Contents

- **Module 0** — Why install an AIOS at all
 - **Module 1** — The 5-step install playbook
 - **Module 2** — Skills, not mega-skills
 - **Module 3** — Memory, planning, and scheduling
 - **Module 4** — Productize your install (the money module)
 - **Module 5** — What to ignore
 - **Final Activity** — Ship your portfolio
 - **Bonus Track A** — The 6 levels of mastery diagnostic
 - **Bonus Track B** — Use case catalog
 - **Appendix** — Vocabulary, troubleshooting, and what's next
-

Module 0 — Why install an AIOS at all

Tuesday morning, two versions

It's 7:42am. You sit down to the same workflow you've run 200 times — drafting an invoice for a retainer client, transforming a Stripe payout into the format your accountant wants, writing the Tuesday LinkedIn post.

Version A — you open Claude in your browser. You re-explain the project. You re-paste the brand guide. You ask for the invoice format and the model gets the columns wrong because it doesn't remember the conversation from last Tuesday. You spend 25 minutes nudging the output into the right shape. You ship at 8:34am, mildly annoyed.

Version B — you open Claude Code in your project folder and type `draft-invoice 2024-INV-047`. The model loads `CLAUDE.md` and knows your project. It loads the `draft-invoice` skill and knows the column format. It pulls the Stripe payout from the path your scheduled task already wrote at 6:00am. The draft lands at 7:43am. You proofread, ship at 7:46am, and pour a second coffee.

The difference between version A and version B is not about Claude being smarter. It's about whether Claude is **installed**. You either have an AI Operating System running for you, or you have an AI subscription you keep re-explaining context to.

What "AIOS" actually means

The phrase "AI Operating System" caught on in April 2026 when [a wave of creator masterclasses](#) converged on the same architectural picture. An AIOS is **not** a product you buy. It is the combination of four things you assemble in your project, in this order:

1. **An alignment layer** — `CLAUDE.md` and the project context, so every session starts pre-loaded with your standards, vocabulary, and decision rules.
2. **A skill library** — small, modular capabilities the model can invoke by keyword, each ≤200 lines, each doing one thing well.
3. **A memory and persistence layer** — what carries across sessions, scheduled cron tasks that fire when you're not watching, and the artifacts those produce.
4. **A productized surface** — the offer you can sell to a client *because* of the install, with concrete ROI math.

That's the system. Anything less, and you're just using a chatbot.

The diagnostic question

If you're not sure whether you have an AIOS, ask yourself: **am I using Claude, or am I installing Claude?**

Using Claude looks like opening `claude.ai` a few times a day, asking a question, copying the answer somewhere. Hours of "am I using AI?" effort, but no compounding leverage. Every prompt starts cold.

Installing Claude looks like every Claude Code session in your project starting hot — the model already knows your project, the skill you needed is already loaded, the morning artifact is already on disk. The hours go down, the leverage goes up, the system gets better every week without you actively investing in it.

This course is the install path. By Day 30 you'll have the four parts of an AIOS running in one of your real projects, plus a public portfolio URL pointing at it, plus one productized offer scoped that you can sell *because* you've installed it.

Key takeaway, Module 0 — *An AIOS is four parts (alignment + skills + persistence + productized surface) installed in your project, not a product you buy. By the end of this guide, you'll have all four installed in 30 days.*

Who this guide is for

- **Business owners** with 1–10 person teams who run real workflows that touch documents, leads, scheduling, or content

- **Solo operators** who have already paid for a Claude / OpenAI / Anthropic subscription and want it to start paying back
- **Agency owners and consultants** who need to install the same pattern in multiple client environments and want a repeatable playbook

This guide is **not** for: people who have never opened a terminal, people who want a beginner's intro to AI, or people who want a tool tour. It assumes you can run a command, edit a markdown file, and deploy a site to Vercel. If any of those are unfamiliar, do the prerequisites first — the install will frustrate you otherwise.

What you need before you start

ITEM	WHY	COST
Claude Code installed (Pro plan or higher)	The install itself runs in Claude Code	\$20+/mo
One real project folder	The install needs a real subject — not a <code>~/test</code> placeholder	Free
A terminal you're comfortable in	bash / zsh / PowerShell all fine	Free
Vercel account	Final activity deploys a portfolio page	Free tier OK
30 days of weekly checkpoints (~12 min/day)	Self-paced, but the cadence is structural	—

That's it. No paid Figma, no paid CRM, no n8n license required for the core curriculum. Bonus tracks introduce optional tools.

The arc of the course

The five modules are sequential and each one ends with a hands-on activity in the [Practice Workbook](#) that produces a tangible artifact. The workbook activities stack into a complete install — and a complete portfolio page — by the time you reach the Final Ship.

```

Week 1 → Module 1 + Activity 1 → Working CLAUDE.md
Week 2 → Module 2 + Activity 2 → Your first invocable skill
Week 3 → Module 3 + Activity 3 → First scheduled autonomous run
Week 4 → Module 4 + Activity 4 → One productized offer scoped
Final → Module 5 + Activity 5 → Public Vercel portfolio URL

```

Don't skip the activities. They are the course. The reading is just the framing.

→ Start the install: [Workbook Activity 1](#)

Module 1 — The 5-step install playbook

The honeymoon and the wall

Most people open Claude Code, ask it to build something cool, watch it work, and walk away convinced AI is the future. They never come back to that project. They open the next chat, ask the next thing, get the next single-shot output. Six months later they have 200 conversations and zero compounding system.

The wall isn't capability. It's **architecture**. Without an install, every session pays a re-context tax. With an install, every session inherits the previous one's progress.

The 5-step sequence

The install path that converged across the April 2026 creator masterclasses (Nick Puru's 42-business sample size, Jono Catliff's full course, Simon Scrapes' skill systems work) is the same five steps, in the same order. Source: [wiki/04 \(AIOS install playbook\)](#).

Step 1 — Priority matrix. Pick **one** project to install in first. Not three. Not "across my whole stack." One. The right project has all three of: (a) you touch it ≥ 3 days/week, (b) it has at least one repetitive task you'd happily delete, (c) the artifact it produces matters enough that you'd notice if it stopped.

The most common install failure is starting in a project you don't actually live in daily. Without daily contact, the install never gets tuned, never compounds, and dies quietly in three weeks.

Step 2 — First install. Drop a real `CLAUDE.md` in the project root. Not a template. Not a generic "you are a helpful assistant" line. A `CLAUDE.md` that captures *this* project's vocabulary, *this* project's standards, *this* project's "what to ask before doing." (Activity 1 is the interview that produces this file.)

This is the alignment layer. Every Claude Code session in this project from now on starts pre-aligned. The first compounding win.

Step 3 — Second install. Build one skill. One. Not a skill library. Pick one repetitive judgment task you do at least weekly, scope it to ≤ 200 lines of `SKILL.md`, push the long stuff into `references/`, and test that you can invoke it by keyword. (Activity 2 ships this.)

Two installs is enough to feel the compounding. The skill consumes `CLAUDE.md`'s alignment for free. You didn't have to re-explain the project to the skill — it inherits.

Step 4 — Compounding. Add agency. Schedule one task to run when you're not watching. (Activity 3 ships this.) Now your morning starts with an artifact already on disk. The install crosses from "tool I use" to

"system that works for me." This is the moment the metaphor shifts from "AI assistant" to "AI operating system."

Step 5 — Handoff. Productize. Turn the install into a one-page offer you could send to a real prospective client today. (Activity 4 ships this.) The install has now generated a sales asset, which means the install has a path to paying for itself many times over.

That's the full sequence. Five steps, one per week of the 30-day install, four shipped artifacts plus a portfolio page.

Why this order is non-negotiable

Each step depends on the one before it.

- A skill written without a `CLAUDE.md` re-invents project alignment in every skill (bloated, brittle).
- A scheduled task written without a skill is a one-off script masquerading as automation.
- A productized offer written without an install is "I could maybe do this" — useless in a sales conversation.

If you skip a step, the next one wobbles. If you do them in order, they compound.

What "compounding" actually means

Most software gets worse over time — entropy, dependency creep, framework churn. An installed AIOS gets *better* over time, for one specific reason: every artifact you add inherits everything before it.

Your second skill takes 60 minutes instead of 90 because the patterns are familiar. Your third skill reuses the references the second one created. Your fifth skill is doing 80% of its work by composing the previous four. Your tenth skill is the one you didn't realize you needed until you started getting paid to deliver it.

This is the leverage curve every creator masterclass is gesturing at. It's not real on the first day. It's overwhelmingly real by Day 90.

Key takeaway, Module 1 — Five steps in this order: priority matrix → CLAUDE.md → first skill → first scheduled task → first productized offer. Each step depends on the previous one. Don't jump ahead.

→ Run Step 2: [Workbook Activity 1 — Build your CLAUDE.md](#)

Module 2 — Skills, not mega-skills

The anti-pattern that wastes most installs

The first instinct, after writing one skill that works, is to write a bigger skill that does more. "I'll add a second mode here. And a third mode for the edge case. And a fourth for that one client. And..." Six weeks later, the skill is 1,400 lines, takes 8 seconds to load, and answers most queries by ignoring 90% of itself.

This is the **mega-skill** anti-pattern. It is the single most common failure of installs that started well. Source: [wiki/03 \(skill systems\)](#), reinforced by every Anthropic-adjacent voice in the field.

The fix is structural: **one skill, one job**. When you want to add a second job, build a second skill. The cost is a few extra files. The reward is everything stays composable, every skill stays under 200 lines, and the model can reason about which skill to invoke without burning context loading nine variants of the same thing.

The 200-line cap is teachable, not arbitrary

Why 200 lines? Two reasons.

Context tax. Every skill loaded in a session pays a context tax — the tokens spent reading it that aren't spent on your actual problem. The 200-line cap keeps the tax under 5% of a typical session's context budget. Above that, the skill starts crowding out the actual work.

Composability. A 200-line skill is small enough that you can read it end-to-end in 2 minutes when something goes wrong. A 1,400-line skill is small enough that you can't. Debuggability is composability's second-order effect.

Anything longer than 200 lines goes to `references/` and loads on demand. Your skill's `SKILL.md` says "Load `references/linkedin-style-guide.md` before drafting" — and that file gets pulled into context only when needed, not every session.

Progressive disclosure

This is the core architectural pattern. Source: [wiki/03 \(skill systems\)](#) and the [super-skills raw note](#).

```
SKILL.md          (≤200 lines, always loaded when skill fires)
├─ references/
│   style-guide.md      (loaded on demand, only when drafting)
│   csv-format.md      (loaded on demand, only when transforming)
│   client-intake-template.md (loaded on demand, only when scaffolding)
├─ scripts/          (optional: pure mechanical operations)
│   validate-input.py
```

The `SKILL.md` is the front door. It tells Claude what the skill does, what triggers it, what inputs it needs, and which references to load when. Everything heavy lives behind that door, loaded only when the workflow actually needs it.

The result: a skill library can grow indefinitely without ballooning the per-session context cost. Your tenth skill is no slower than your first, because each fires with the same lean front door.

Why Anthropic's own teams do this

The growth-marketing team at Anthropic, when they shipped their internal ad-copy pipeline, broke it into a headline-writer skill, a description-writer skill, and a CTA-writer skill — three small modular skills with a thin orchestrator chaining them together. Source: [wiki/03 \(skill systems\)](#).

This is the canonical move: when one skill is doing two jobs, split it. When two skills need to coordinate, write a thin orchestrator. Don't build one big skill that does both.

This is also why the framework-migration narratives (Hermes, Open Claw, Gravity Claw) miss the point. Source: [hermes raw note](#). The pattern that wins isn't "adopt a framework that wraps everything." It's "compose small skills the framework people would have built individually." If a framework offers a useful prompt, cherry-pick the prompt, don't migrate the project.

Skill discovery mechanics

When you type a message in Claude Code, it scans the `.claude/skills/*/SKILL.md` frontmatter for `description` lines that match keywords from your message. The match is **keyword-rich**, not semantic — your `description` should literally contain the words a future-you would type.

Bad:

```
description: Helps with content tasks.
```

Good:

```
description: Drafts a LinkedIn post from a YouTube URL using Noxlee brand voice. Triggers on "dra
```

Notice the second one explicitly enumerates trigger phrases. This is doing the work for the matcher, not hoping it figures out you meant LinkedIn when you said "social".

One skill that you'll actually use

The best first skill is one you'll invoke at least 3 times in the 7-day test. The second-best first skill is one that automates something so painful you've been delaying it for months. (Activity 2 walks the scope →

scaffold → ship → test loop.)

Frequency-of-use beats importance-of-use for your first skill. A skill you invoke daily teaches you the install. A skill you invoke once a quarter is a file you'll forget exists.

Key takeaway, Module 2 — One skill, one job. ≤200 lines in `SKILL.md`. Heavy stuff to `references/` loaded on demand. Composable beats comprehensive.

→ Run Step 3: [Workbook Activity 2 — Ship your first skill](#)

Module 3 — Memory, planning, and scheduling

The three persistence layers

After an alignment layer (`CLAUDE.md`) and a skill, the next install moves are about **persistence** — what carries across time. There are three layers, ordered by lifespan:

LAYER	LIFESPAN	WHERE IT LIVES
In-conversation memory	One session	The plan/todo list inside an active Claude Code session
Cross-conversation memory	Across sessions	<code>~/.claude/.../memory/</code> files indexed in <code>MEMORY.md</code>
Cross-session agency	Always-on	Cron / Windows Task Scheduler / managed agents firing on schedule

Most users use only the first layer (in-conversation), and it's why their installs feel forgetful. The second and third layers are where the install crosses from "tool I open" to "system that runs."

Sources: [wiki/06 \(operational tactics\)](#), [local-vs-remote raw note](#), [agentic-os 9 pieces raw note](#).

Layer 1 — In-conversation memory

This is the plan/todo state inside one Claude Code session. It's free, automatic, and most users underuse it.

Two tactical tips:

- Use `/context` to see how full your context is. When it's >70%, your session is degrading — the model is losing earlier turns. Time to wrap or `/compact`.

- Use `/compact` to summarize and continue. It hands the model a compressed version of your turn-history, freeing context budget without losing the thread. Underused.

These commands shipped quietly in early 2026 and nobody learned them. They're the lowest-effort, highest-yield operational improvements available. Source: [wiki/06, 32-hacks raw note](#).

Layer 2 — Cross-conversation memory

`MEMORY.md` is the index Claude Code loads on every session start. Underneath it are individual memory files (one fact per file, organized by topic) that get pulled into context when relevant.

The pattern that works:

- **One fact per file.** Easier to find, easier to evict.
- **MEMORY.md as index, not content.** One line per memory file, ≤150 chars. The index gets loaded; the files get loaded on demand.
- **Memory files for `feedback` and `project` decisions, not for code patterns.** The code is the record for code patterns. Memory is for things you'd otherwise have to re-explain — durable preferences, in-flight project context, references to external systems.

The example everyone copies is the operator's own `MEMORY.md` with topic-organized memories like `feedback_chalk_tactics_chip_opacity.md` (one specific decision, with the reason) — small, indexed, evictable, recallable.

Layer 3 — Cross-session agency (the schedule)

This is the moment the install starts producing artifacts when you're not watching.

Three platform options:

Cron (Mac/Linux/WSL). Built-in, free, opaque. Every modern Unix-like has it. The downside: cron's default behavior is to silently swallow errors and email you on stderr — both wrong for AI tasks. Always log explicitly:

```
30 6 * * 1-5 /bin/bash /full/path/to/scripts/task.sh >> /full/path/to/logs/cron.log 2>&1
```

Windows Task Scheduler. Built-in to Windows, GUI-driven, more discoverable than cron. PowerShell command builds a task in 30 seconds.

Claude Code managed agents (`/schedule`). Easier syntax, paid (per Anthropic pricing). Great if you don't want to think about cron and don't mind the cost.

n8n cron node triggering a Claude Code shell call. The right pick if you're already running n8n for client work. The cron node fires on schedule, runs a Code node that shells out to `claude --print < prompt.txt`,

and pipes the output anywhere n8n can post.

For the install, pick whichever platform you'll remember to debug. A scheduled task whose log you can't find is worse than no scheduled task — silent wrong-output is worse than loud failure.

Headless Claude Code (`claude --print`)

The unlock for scheduled tasks is `claude --print` — non-interactive mode. The session loads, runs the prompt, writes the output, and exits without waiting for a TTY. Source: [wiki/06](#), [local-vs-remote raw note](#).

```
claude --print "Summarize today's PRs in /tmp/digest.md"
```

That's the entire scheduling primitive. Wrap it in a shell script with logging and a guardrail (the script checks output size; if empty, it writes to `.failed` instead of overwriting yesterday's good output), and you have a production-grade scheduled task.

The guardrail rule

Every scheduled task needs a guardrail. The pattern:

1. Run the task.
2. Check the output looks valid (non-empty, parses, contains expected markers).
3. If valid → write to the canonical output path (overwriting yesterday's).
4. If invalid → write to a `.failed` file instead, and emit an alert.

This rule alone prevents the failure mode that destroys most automation: the task "ran successfully" but the output was garbage, you didn't notice for two weeks, and downstream consumers (you, your team, your client's Slack channel) were getting wrong data the whole time.

Loud failure beats silent wrong-answer every time.

The kill-switch rule

Always know how to disable a scheduled task in 30 seconds. `crontab -r` deletes everything; that's not a kill-switch, that's a panic button. You want one line you can comment out:

```
# 30 6 * * 1-5 /bin/bash /path/to/task.sh >> /path/to/log 2>&1 # disabled 2026-05-12 - investig
```

When (not if) a task starts misbehaving, you'll be glad you can disable it without losing the schedule definition.

Key takeaway, Module 3 — Three persistence layers (in-conversation, cross-conversation, cross-session). Layer 3 is where the install starts working for you. Always log, always guardrail, always have a kill-switch.

→ Run Step 4: [Workbook Activity 3 — Schedule your first autonomous run](#)

Module 4 — Productize your install (the money module)

Why this is the module that pays for the course

Modules 1-3 are the install. Module 4 is what you sell *because* of the install.

If you stop after Module 3, you have a productivity gain — useful, but invisible to anyone outside your business. If you do Module 4, you have a sales asset that turns the install into a recurring revenue stream you can pitch to real clients.

The three productized offers below are pulled from the May 2026 creator-masterclass synthesis ([wiki/07, five-money-offers raw note](#)). Each one has been tested against real clients, has documented ROI math, and is small enough to scope into a one-page proposal.

Offer 1 — Document Processing

Industries: Insurance, law, accounting, logistics, healthcare admin.

The pain: Every business that handles repetitive documents — invoices, intake forms, contracts, claims, shipping manifests — spends 6–20 hr/week of human time keying data from one document into another system. The error rate is non-trivial; the cost of each error is significant.

The install: A skill that reads the document (PDF, scan, email body), extracts the structured data, validates it against the destination system's schema, and writes the import file (CSV, JSON, API call). A scheduled task that fires when new documents arrive (folder watcher, email rule, Drive trigger) so the entire pipeline runs unattended.

The ROI math template: - Hours/week the client spends keying: 8 - Loaded hourly cost: \$35/hr - Annual labor cost: \$14,560 - Error rework cost: ~\$2K/yr (1 transposition per month × \$200 to chase) - **Total annual pain: ~\$16,560** - One-time install: \$5,000–15,000 - Year 1 net savings: \$1,560–11,560 - Payback: 4–10 months - 3-year ROI: 230%–870%

The closer line: "For [install fee] one-time, you save ~\$16K/yr forever. Payback in [X] months."

This is the offer with the cleanest ROI math because the savings are in measurable hours and dollars, not vague "productivity."

Offer 2 — Database Reactivation

Industries: Gyms, SaaS, coaching, agencies — any business with a CRM full of leads they haven't touched in 6+ months.

The pain: Most CRMs are graveyards. Sales teams pursue the freshest 5%. The other 95% are leads that converted last year, leads that ghosted halfway through the funnel, leads that paid once and never came back. Each one cost real acquisition dollars. Each one is now sitting in a database doing nothing.

The install: A skill that reads the CRM export, segments leads by recency × engagement × deal size, drafts a personalized re-engagement message per segment using the lead's actual history, and queues them for sending through the client's existing email system. A scheduled task that runs weekly so the reactivation campaign keeps refreshing without the client having to think about it.

The ROI math template: - Cold leads in the CRM: 1,200 - Conservative reactivation rate: 3% (industry benchmark) - Reactivated leads: 36 - Average deal size: \$400 - Reactivated revenue: \$14,400 in 60 days - One-time install: \$4,000 - **Year 1 ROI: ~1,200%** (which is the headline; even if reactivation rate is 1%, ROI is still 200%+)

The closer line: *"You already paid to acquire these 1,200 leads. They're sitting unused. We unlock \$14,400 from them in 60 days for a \$4,000 install."*

This is the highest-conversion offer because the install fee is small relative to the revenue unlock, and the leads already exist — there's no acquisition cost to argue about.

Offer 3 — Internal Reporting & Status Notifications

Industries: Every business with ≥3 people and ≥2 systems they need to keep in sync.

The pain: Every business has a moment each morning where someone manually pulls metrics from 4 different systems, formats them in a Slack message or email, and sends them to the team. Or the equivalent: a weekly report no one reads because it's pulled together at 11pm Sunday. The pain isn't dramatic — it's the cumulative drag of 30 min/day × 365 days × multiple people.

The install: A skill that pulls the metrics from each source (API calls, scrapes, CSV reads), formats them per the client's existing template, and posts them to wherever the team consumes them (Slack channel, email, dashboard). A scheduled task that runs daily/weekly without a human ever touching it.

The ROI math template: - People-hours/week spent on this manual work: 5 - Loaded hourly cost: \$50/hr - Annual cost: \$13,000 - Plus the qualitative cost: the report goes out late, gets skipped on busy days, has typos, includes wrong numbers. (Don't quantify this in the proposal — leave it as "and on the days it's missed, decisions get made on stale data.") - One-time install: \$4,000–8,000 - Monthly retainer

(recommended for this offer): \$300–800/mo for ongoing tuning - **Year 1 ROI: 150–300%, plus the qualitative wins**

The closer line: *"Your team's morning starts with the report already in the channel. They never have to remember to send it again."*

This is the **stickiest** offer — once installed, it's load-bearing for the team's daily rhythm. Cancelling it means going back to the manual process, which nobody wants to do. High retention, predictable retainer revenue.

The doctor-not-pharmacist diagnostic

The trap that ruins most AI proposals is leading with the technology. *"I build agentic workflows in Claude Code"* is a pharmacist line — you've named the pill before diagnosing the disease.

The doctor frame: **diagnose first, prescribe second**. Source: [wiki/07 §"Sales / positioning principles"](#), [agentic-workflows-WAT-framework-sales raw note](#).

Discovery call structure:

1. **Listen for the clog.** Where in their pipe is water backing up? "We're growing but onboarding is breaking" → onboarding is the clog. "Sales team is busy but revenue isn't moving" → leads are the clog. "Office manager works Saturdays" → admin is the clog.
2. **Diagnose, not prescribe.** Reflect the clog back: *"It sounds like the bottleneck isn't lack of leads — it's that the leads sit in the CRM untouched. Is that right?"* Confirm before suggesting anything.
3. **Prescribe ONE offer.** From the three above, the one whose pattern matches the diagnosed clog. Not three offers. One. *"For your situation, the offer that fits is Database Reactivation."*
4. **Walk the ROI math live.** Use their actual numbers, even rough. *"You said you have ~1,500 cold leads. At 3% reactivation, that's 45 deals. At your \$400 average, that's \$18K. Our install for this is \$4K."*
5. **One specific yes.** *"Want me to send a one-pager you can show your partner this week?"*

Five steps, ~25 min discovery, one clear yes-or-no for the prospect. This is dramatically higher conversion than "let me show you everything we do."

Value-based pricing, never hourly past the first wins

Once you have one or two paid installs as proof, switch to value-based pricing. Source: [wiki/07 §"Sales / positioning principles"](#).

The line that closes it: *"If we can build something in 30 minutes that saves you 20 hours/week, that's not a 30-minute job. That's tens of thousands of dollars in value to your business — for the rest of time."*

Hourly pricing punishes you for being good at the job. Value pricing rewards it. The ROI math template above is the asset that makes value pricing defensible — clients can argue with hours; they can't argue with their own savings number.

What this offer is NOT

In every proposal, include a section titled "What this is NOT." Counterintuitively, naming what you don't do increases trust because it shows you've thought about scope.

Examples:

- *"This is not an AI chatbot for your customers."* (For Document Processing)
- *"This does not replace your office manager — it removes the worst part of her job so she can do the parts she's actually good at."* (For Internal Reporting)
- *"This is not 'set it and forget it' — we'll spend the first 30 days tuning it together based on edge cases."* (For Database Reactivation)

This is the line that separates serious vendors from hype merchants. Most AI proposals over-promise and under-define scope. Doing the opposite is a competitive advantage.

Key takeaway, Module 4 — *Three productized offers (DocProc, Reactivation, Reporting). Pick the one that matches your install. ROI math template makes value pricing defensible. Doctor-not-pharmacist closes more deals than tool-tour pitches.*

→ Run Step 5: [Workbook Activity 4 — Scope one productized offer](#)

Module 5 — What to ignore

Anti-hype list

The April–May 2026 creator-masterclass wave produced a lot of useful signal. It also produced a lot of noise — pitches, framework migrations, marketplace hype, and obsessions with hardware that have nothing to do with whether your install actually works. This module is the noise filter.

Source: [wiki/07 §"What I'd ignore"](#) plus the cross-cutting skepticism from the masterclass synthesis ingest.

Ignore the framework wars

There's an emerging industry of "agentic frameworks" — Hermes, Open Claw, Gravity Claw, and the next one shipping Tuesday — each one promising a comprehensive structure that wraps everything you do.

Source: [hermes raw note](#).

The right move is **cherry-picking**, not migration. If a framework offers a useful prompt or a clever skill pattern, copy that prompt into your install. Don't migrate the whole project. The cost of migration is real (your install loses six months of accumulated tuning); the benefit is usually one or two prompts you could have just copied.

The frameworks that survive will be the ones that ship a small set of obviously-useful patterns. The frameworks that don't will be the ones that demand you adopt the whole stack. Default skeptical.

Ignore the "agent that self-heals in production" narrative

Many demos show an agent that detects an error, debugs itself, and ships a fix — all autonomously, in production. This is genuinely impressive in the demo. It is **not** how production agents should work.

The WAT framework (Write-Apply-Test) breaks the build-time and run-time loops apart. Source: [wiki/07 §"Cross-cutting principles"](#), [agentic-workflows-WAT-framework-sales raw note](#).

- **Build time:** Yes, agents iterate, debug, retry. This is where self-healing earns its keep.
- **Run time:** No. Production agents do exactly what they were tested doing. They write, they apply, they test — they do **not** invent new behaviors. If they detect an error, they alert and fail safely. They don't fix.

This is the difference between an agent you'd deploy and a demo you'd post on Twitter. The demo is fine; the deployment requires WAT discipline.

Ignore the 800K-skill marketplaces

Several skill marketplaces (skillsmpp.com, agentskills.io, others) have launched promising hundreds of thousands of pre-built skills you can drop into your install. Most of these skills are AI-generated, untested, and shallow.

The pattern that works: write skills for *your* project, against *your* CLAUDE.md vocabulary. The skills that ship by hand from people who've installed AIOSeS in real projects (your own, your client's, the tools you trust) will outperform any marketplace skill 9 times out of 10.

If you must browse the marketplaces, treat them like Stack Overflow — read, understand, copy the *idea*, never the verbatim file. Cherry-picking, not bulk import.

Ignore the Mac Mini obsession

You'll see a lot of content evangelizing dedicated always-on Mac Mini setups for personal AI. The argument: a Mac Mini in the corner running Claude Code 24/7 is the home base for your AIOSeS.

It's not wrong. It's just over-prescribed. If you already have a desktop or laptop you keep awake (or can configure to wake on schedule), you have the equivalent. The hardware doesn't matter — what matters is the install. Spend the \$600 on tokens or training, not on a redundant computer.

Ignore the framework-style 4-hour beginner courses

There's an explosion of multi-hour beginner crash courses on Claude Code. If you're past the absolute beginner tier (you can run a command and edit a file), these are largely a waste of time. They're slow, they teach to the lowest common denominator, and they spend 60% of the runtime on tool tours rather than installs.

If you need a foundational primer: [wiki/01 \(mental models\)](#) is 12 minutes of reading and covers what those courses spend 4 hours on. Use the time you saved to actually do Activities 1-5.

Ignore the brevity-vs-depth false choice in your own work

A March 2026 paper found that **forced-concise responses correlate with correctness**. Source: [wiki/07 §"Cross-cutting principles"](#). The mechanism: when a model is told "be brief," it's also implicitly told "don't add wrong-direction reasoning to pad the answer."

Apply this to your install: when writing a `CLAUDE.md`, a `SKILL.md`, or a prompt template, prefer the shorter version. The longer version is usually padding. The 200-line `SKILL.md` cap is this principle made structural.

It also applies to your sales conversations. The 25-minute discovery call beats the 60-minute one for conversion.

What's actually worth your time

Five things, in priority order:

1. **Finish Activities 1-5.** Everything else is secondary until you've shipped the install.
2. **Add `/statusline` to your terminal.** Continuous visibility into context %, model, and cost. 5-min install. Source: [32-hacks raw note](#).
3. **Use `ultrathink` on hard problems.** Type the literal word `ultrathink` (no slash) before architecture decisions or complex debugging. Allocates 32K thinking tokens. Source: [wiki/07 Tier 1](#).
4. **Install Context 7 MCP.** Pulls live docs from your most-used libraries before Claude writes code. Stops the "function deprecated" errors. Source: [wiki/07 Tier 2](#).
5. **Audit your install monthly.** What's working, what's not, what's worth pruning. The install is a living system.

Key takeaway, Module 5 — Most of what you'll see online is noise. Cherry-pick patterns, never migrate frameworks. Build to your project, never to a marketplace. WAT at runtime, self-healing only at build time.

→ Run the final ship: [Workbook Activity 5 — Assemble portfolio + deploy URL](#)

Bonus Track A — The 6 levels of mastery diagnostic

This is a self-assessment. Source: [wiki/02 \(six levels of mastery\)](#), originally compiled from Nick Puru's "35 concepts, 6 levels" masterclass. The levels are not strictly sequential, but they're a useful diagnostic for "where am I, and what's the next unlock."

Level 0 — Curious user

You've heard of Claude Code. Maybe you've installed it. You've never run it on a real project of your own.

Unlock: Module 0 of this guide. Pick a project. Run Activity 1.

Level 1 — Single-shot user

You open Claude Code, ask one thing, get one answer. No `CLAUDE.md`, no skills, no memory. Every session starts cold. **Unlock:** Activity 1. The `CLAUDE.md` is the entire difference between Level 1 and Level 2.

Level 2 — Aligned user

You have a `CLAUDE.md` that actually changes how the model behaves in your project. You can name one moment in the past week where the model defaulted to the right thing because of `CLAUDE.md`. **Unlock:** Activity 2. The first skill is the unlock.

Level 3 — Skilled user

You have at least one invocable skill that you used at least 3 times in the past week. You understand the 200-line cap. You've broken at least one mega-skill into smaller skills. **Unlock:** Activity 3. The scheduled task is the unlock.

Level 4 — Operating-system user

You have a scheduled task that fires without your input and produces an artifact you check daily. The install is starting to compound — you're shipping faster than you used to, and you can name two specific places where the install bought you back time. **Unlock:** Activity 4. Productizing is the unlock.

Level 5 — Selling-the-install user

You've sent at least one productized offer to a real prospect (paid or pilot). You can speak to the install with concrete ROI math, citing your own portfolio URL. The install isn't just for you anymore — it's a service you can sell. **Unlock:** Sell three more, refine the offer, repeat.

Level 6 — Teaching-the-install user

You've installed an AIOS in someone else's project (a teammate, a client, a partner business) and they're using it without your daily intervention. The install has graduated from personal leverage to delivered product. **Unlock:** This is where you turn the playbook into a productized service business. Source: [wiki/05 \(use case catalog\)](#) for the next-level patterns.

Pick the level that honestly describes you today. The next activity in this guide is your unlock. Don't skip ahead — Level 4 doesn't skip Activity 1, it returns to Activity 1 and tightens its `CLAUDE.md`.

Bonus Track B — Use case catalog

Most of the install patterns from the field collapse to a small number of repeatable shapes. Source: [wiki/05 \(use case catalog\)](#), drawn from 42+ business installs across the masterclass field.

Use this catalog when scoping productized offers (Activity 4) — pick the shape that matches the client's diagnosed clog and adapt the install template.

The Document → System pipeline

Read structured or semi-structured documents (invoices, intake forms, emails, scans) → extract data → write to destination system (CRM, accounting, spreadsheet, API). **Match for:** insurance, law, accounting, logistics, healthcare admin.

The Database → Outreach pipeline

Read CRM or contact database → segment by recency / engagement / value → draft personalized outreach → queue for sending. **Match for:** gyms, SaaS, coaching, agencies, ecommerce.

The Multi-source → Daily Digest pipeline

Pull metrics or status from 3+ systems → format per client template → post to channel (Slack, email, dashboard) on schedule. **Match for:** any business with team-wide morning standups, leadership dashboards, or client status reports.

The Conversation → Followup pipeline

Read meeting transcripts (Fireflies, Otter, Zoom) → extract decisions, action items, dependencies → draft followup emails / Asana tasks / Linear tickets per attendee. **Match for:** consulting, sales-heavy orgs, agencies running discovery calls.

The Catalog → Drafted Content pipeline

Read product/SKU/inventory data → draft product descriptions, marketing copy, social posts → queue for review. **Match for:** ecommerce, content marketing, publishers.

The Inbound → Triaged Routing pipeline

Read inbound messages (email, support tickets, form submissions) → classify by intent, urgency, ICP → route to right human or auto-respond. **Match for:** customer support, sales operations, inbound-heavy businesses.

For each pattern, the install shape is the same: one skill that does the transform + one scheduled task that runs the trigger + one `CLAUDE.md` that aligns the vocabulary + one productized offer with ROI math.

The patterns are the menu. Pick one per client engagement. Resist the urge to bundle three patterns into one offer — that's how scope creep enters the conversation and conversion drops.

Appendix

Vocabulary

- **AIOS** — AI Operating System. The four-part install: alignment + skills + persistence + productized surface.
- **CLAUDE.md** — The project-root file Claude Code loads every session. The alignment layer.
- **Skill** — A modular capability invocable by keyword, defined in `.claude/skills/<name>/SKILL.md`, ≤200 lines.
- **References** — Files in `.claude/skills/<name>/references/` loaded on demand by a skill.
- **Headless mode** — `claude --print`, the non-interactive Claude Code invocation used for scheduled tasks.
- **Scheduled task** — A cron / task scheduler / managed agent firing on a schedule, usually wrapping `claude --print`.
- **Productized offer** — A one-page proposal for a recurring service, with concrete ROI math, scoped from the install.

Troubleshooting

Skill doesn't fire on its trigger phrase. The `description` frontmatter isn't keyword-rich enough. Rewrite to literally enumerate trigger phrases.

Scheduled task ran but the artifact is wrong. Check the log file. Most common cause: the skill or prompt assumed an input that wasn't there at run time. Add a guardrail check before writing the canonical output.

`CLAUDE.md` is over 80 lines and growing. Split rules into `.claude/rules/<topic>.md` and have `CLAUDE.md` reference them. Same progressive-disclosure pattern as skills.

Context fills up too fast. Use `/context` to diagnose what's eating it. Common culprits: long MCP server outputs (consider whether you need the MCP), giant `CLAUDE.md` (split it), unread tool results staying in context. Use `/compact` to recover.

Sales conversations stall after the discovery call. The proposal is probably too long or too generic. Re-read Module 4 — the one-page offer with ROI math beats the multi-page deck every time.

What's next

You finished the install. Three concrete next moves:

1. **Sharpen one offer and pitch it.** The first paid install is the one that pays back the course many times over.
2. **Add a second skill, then a third.** Compounding shows up around skill #5–7 in any given install.
3. **Install in a second project.** The patterns you learned here transfer; the activation energy is dramatically lower the second time.

When you're ready to install AIOs in client environments as a productized service — or want help installing in your own — that's something we do at Noxlee Automations. noxlee.chumperum@gmail.com.

Now ship something else.

— Noxlee