

The model is one component. The system does the work.

How I built an agentic operating system
around my actual work



How many of you have had an excellent AI conversation...

that produced absolutely nothing reusable the next day?

Great conversation.

Gone
tomorrow.

A chatbot answers.
A system remembers.

KNOWS the person and the job

ACTS through tools on real systems

LEARNS through reusable procedures

RUNS on triggers and schedules

PROVES what happened



Seven layers make it dependable.

- 1 Person + purpose
- 2 Context + constraints
- 3 Skills + instructions
- 4 Tools + models
- 5 Schedules + subagents
- 6 State + knowledge
- 7 **Verification + approval**



Local. Routed. Mostly quiet.

49

active scheduled tasks

117

installed skills

Reusable procedures
for research,
publishing,
operations,
documents, and
maintenance.

Mac

local execution

Gateway

always running

Telegram

mobile surface

GPT-5.6


primary conversation

Other models

task-fit routing

The count is not the accomplishment. The filtering is.

Collection is not interruption.



DISCOVER

VERIFY


FILTER

DELIVER

Many sources enter.

At most three signals leave.

A good loop is a small operating policy.



A long wish

Every morning, find AI news and summarize it for me.

No sources. No threshold. No state. No proof.

An operating policy

- **Named source categories**
- Publication dates verified
- Duplicates and stale items rejected
- Direct URLs preserved
- Confidence and implications labeled
- **Silent when nothing crosses the bar**
- **Never publishes or contacts anyone**

We build one live.


A recurring institute research brief, built as a safe temporary workflow.



If the network fails, the artifact should still tell the story.

State stops the agent from rediscovering Tuesday.

- last successful run
- sources reviewed
- **items already surfaced**
- open follow-ups
- **suppressed items + why**
- verification status



Silence is a feature.

[SILENT]

when nothing requires attention, a decision, preparation, or a worthwhile intellectual move

MAXIMUM

3

**items in the Daily
Field Report**

The system earns trust by not demanding it.

Automation watches the automation.



Execution success and message delivery are separate gates.

Approval is architecture.

We do not buy AI. We approve it.

SECRETS

1 Password or Keychain

TOOLS

least privilege

EVIDENCE

facts separate from claims

EXTERNAL ACTION

human approval

**Approve the bounded workflow,
not the impressive demo.**



Design one real loop.

Choose work your team repeats at least monthly.

Name + owner

Trigger

Inputs

Context

Discovery

Decision rule

Action

Verifier

State

Notify when

Never do

Escalate when

What improves? Where could it be wrong? What stays silent?

Context.

Judgment.

Proof.

**The most useful agent knows what matters,
proves what it did, and leaves you alone when
nothing needs you.**

Questions