

# Station: The Active Knowledge Surface

A local-first, agent-integrated second brain that catches half-thoughts, asks the right questions, and turns ideas into reality.

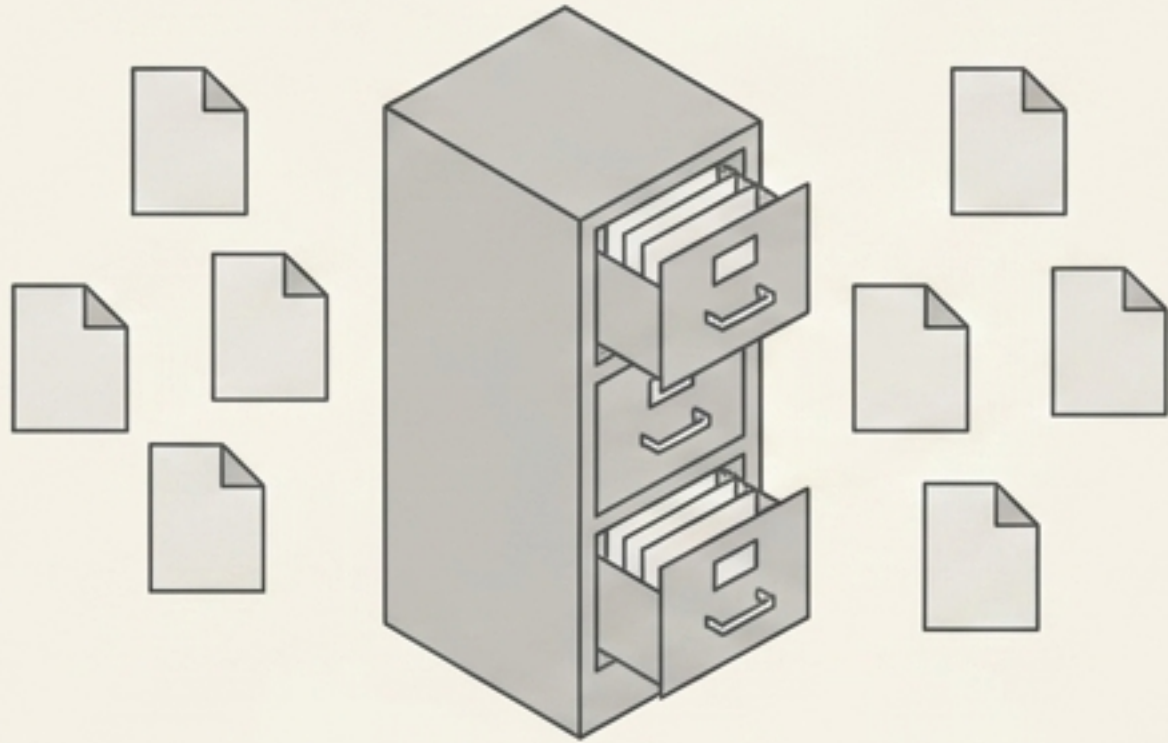
## THE STATUS QUO VS. SOVEREIGN COMPUTING

Capabilities	Dumb Storage (Notes, Markdown)	Siloed SaaS (Notion, Todoist)	Cloud AI (ChatGPT, Claude)	Station
Active Thinking: Does it challenge your ideas?	✗	✗	✓	✓
Cross-Domain Context: Tasks + Docs + Homelab?	✗	✗	✗	✓
Data Sovereignty: Runs on hardware you control?	✓	✗	✗	✓
Agent Interoperability: Native external AI plug-ins?	✗	✗	✗	✓

The status quo forces a choice between intelligence, integration, and privacy. **Station is what happens when you build the tool you actually wanted on hardware you control.**

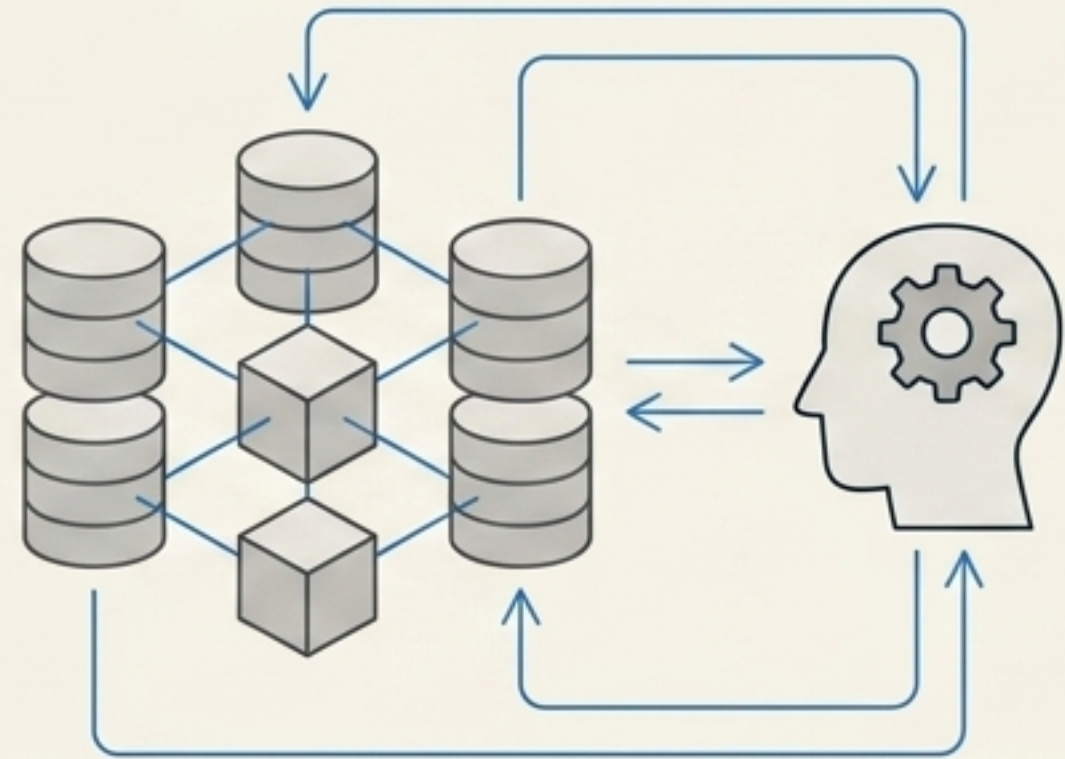
I have a lot of ideas. Most of them are bad.  
The ones that aren't bad usually die because I  
never wrote them down properly.

## The Old Way

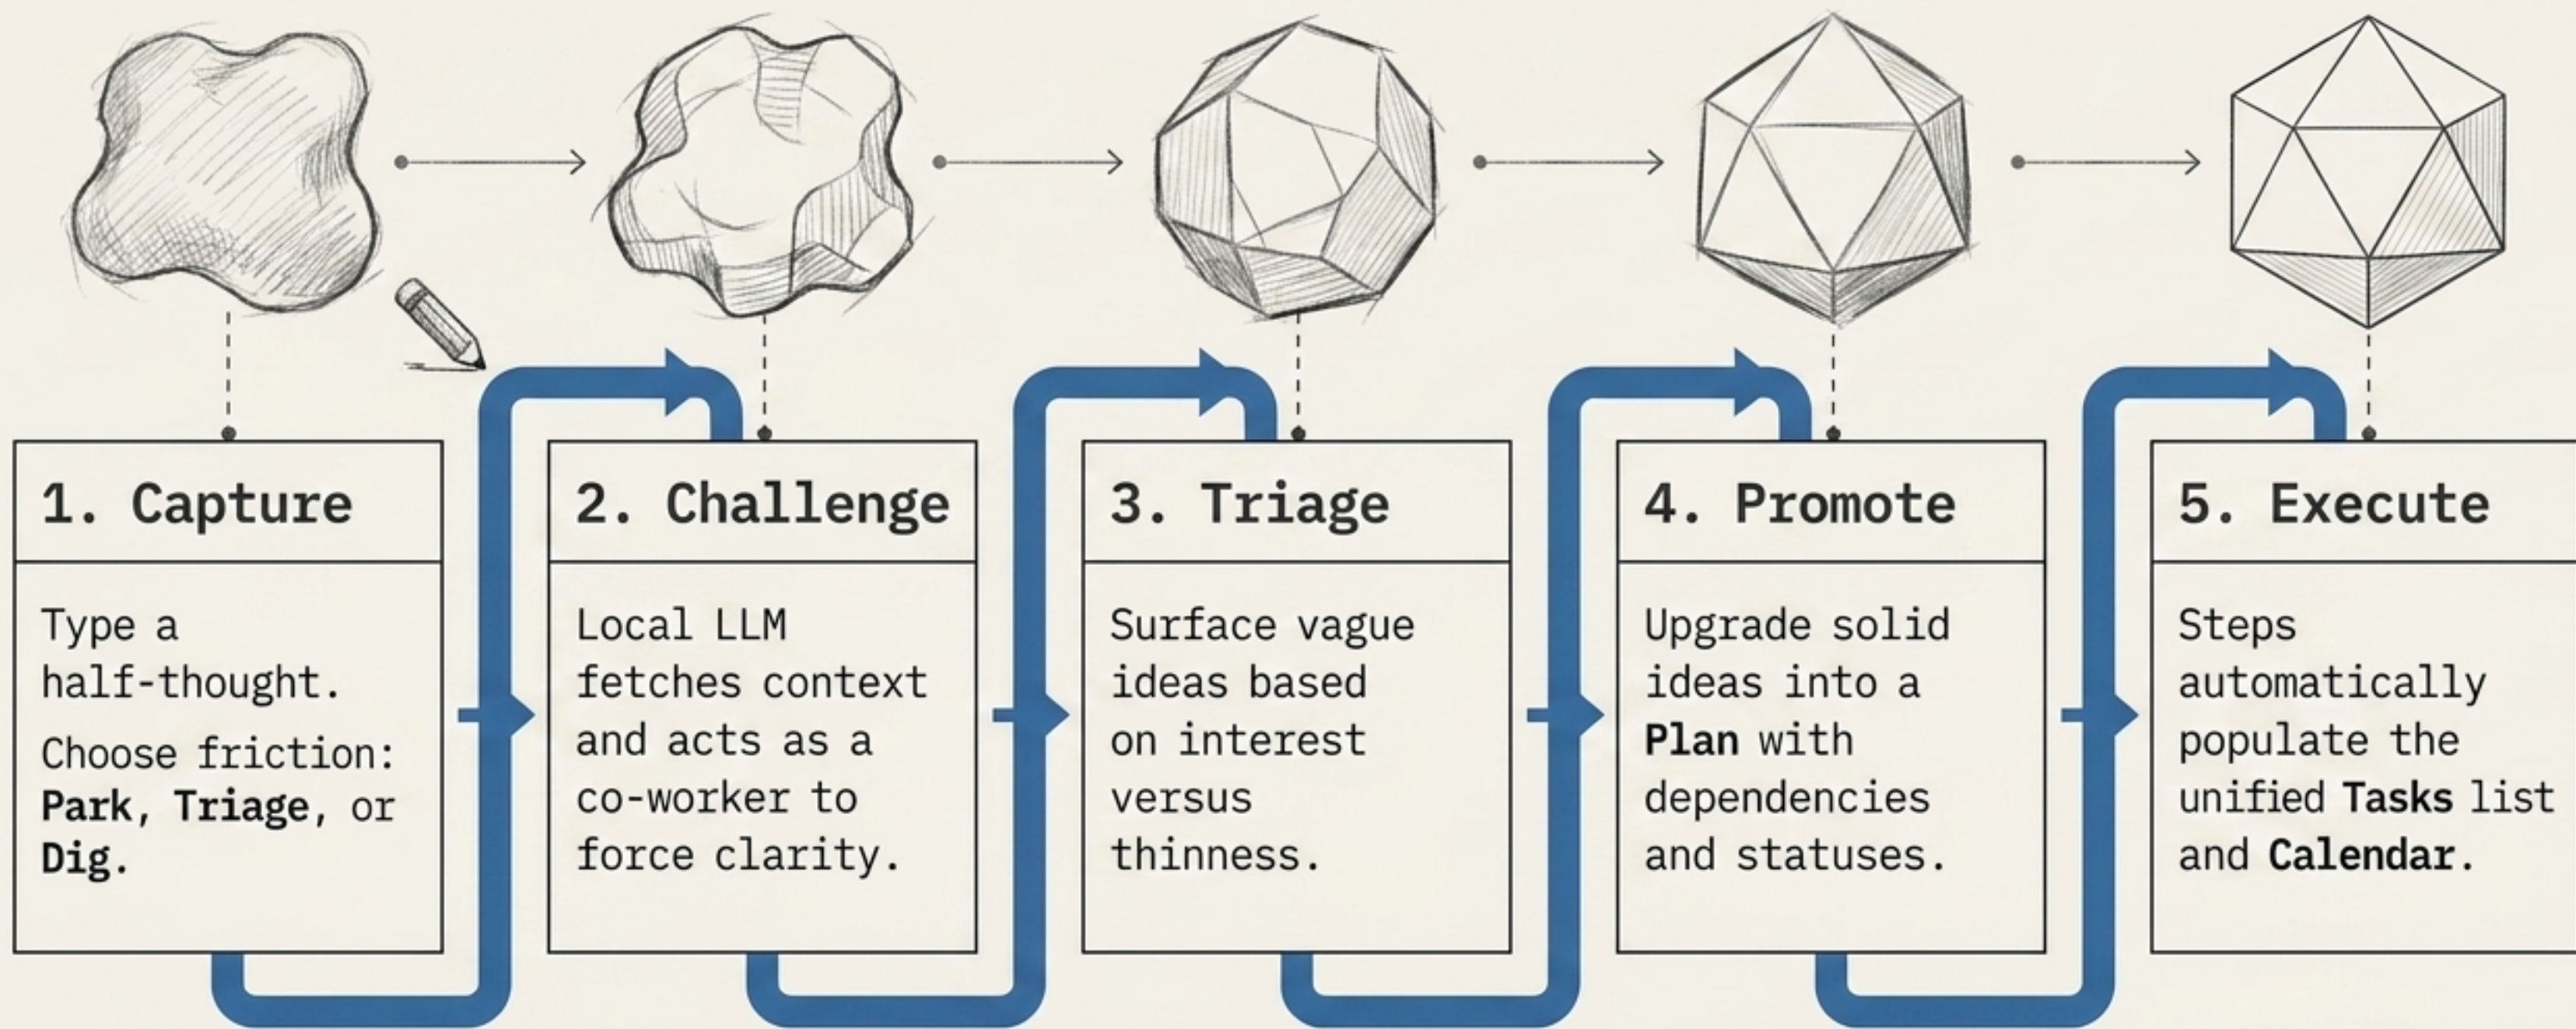


Capture, but don't think. Requires you to  
do all the work of pulling threads.

## The Station Paradigm

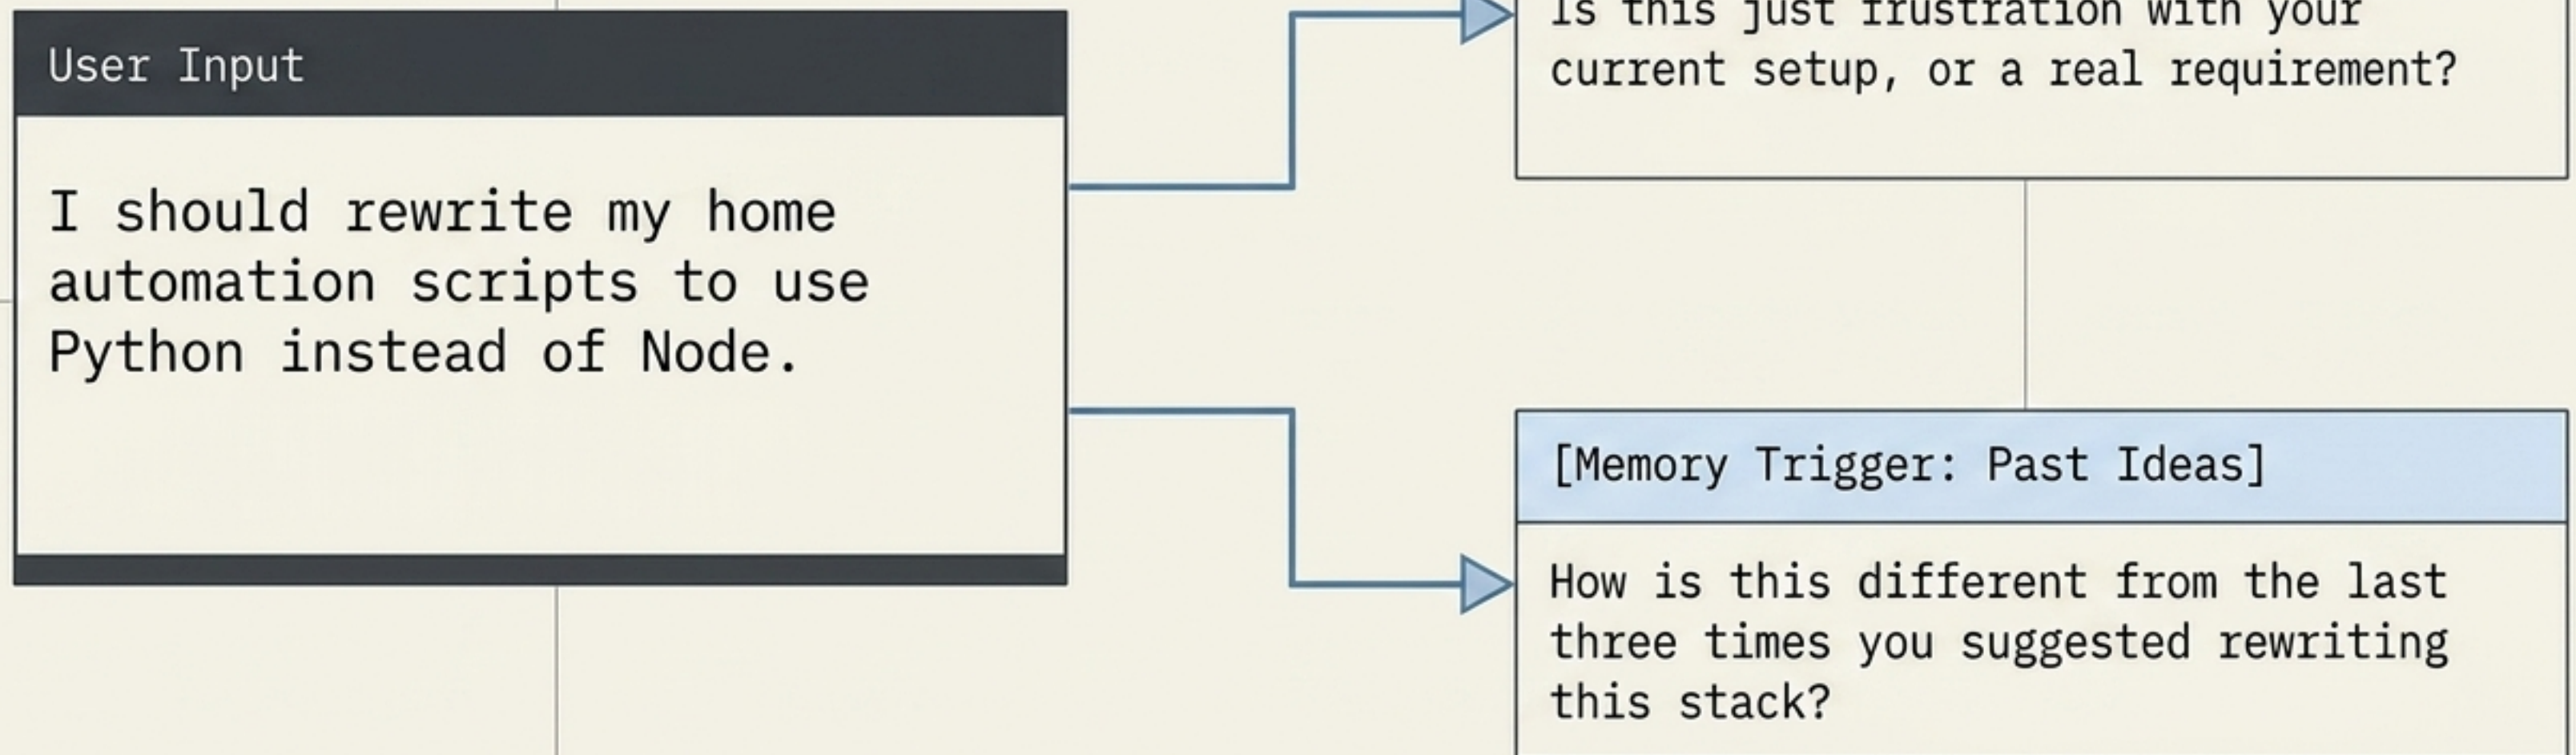


The place my brain offloads to. Asks the  
follow-up questions a good co-worker would.



**Note:** No separate to-do app to sync. Notes and tasks are the exact same database row, just viewed differently.

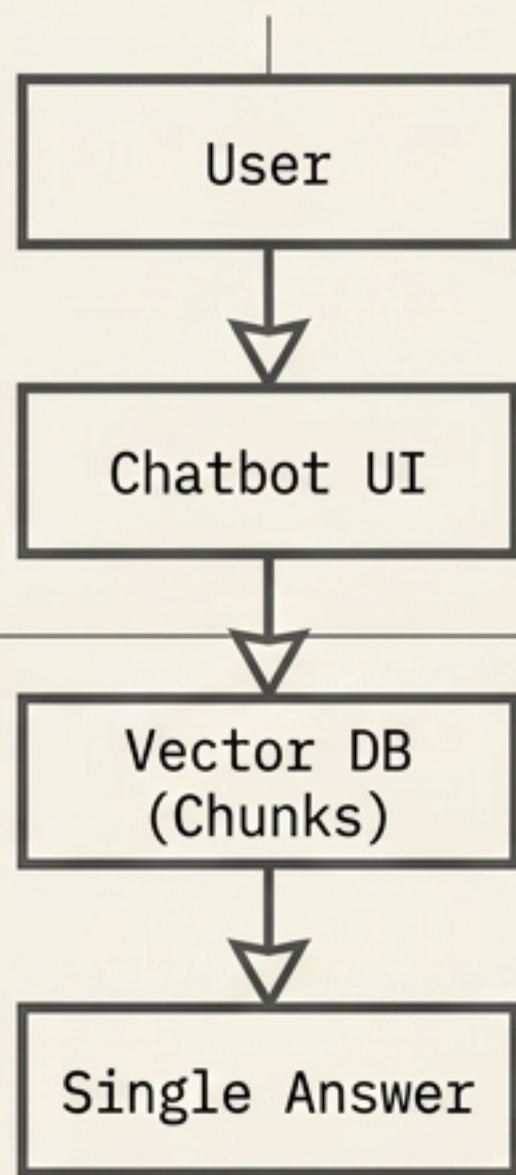
# THE CHALLENGE PHASE: FORCING CLARITY



Mechanism: Answer in plain text. Station updates the idea's completeness score and stops querying once the thought is fully baked.

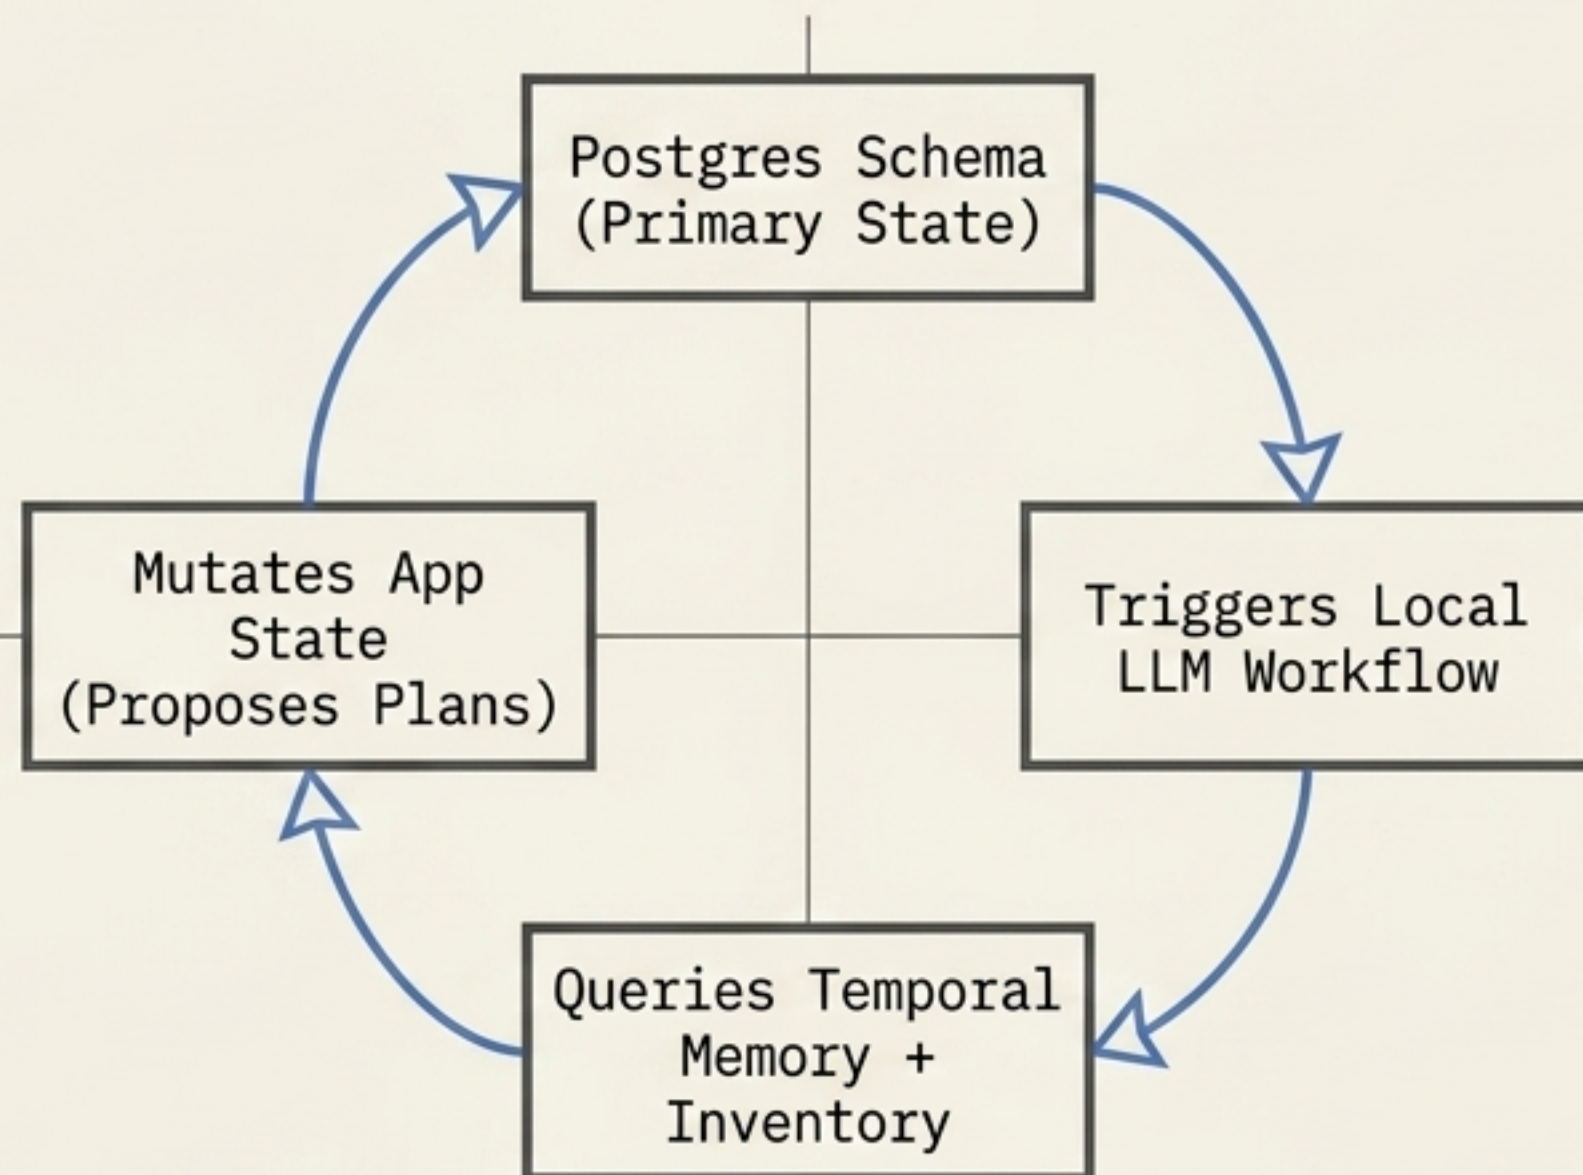
# ARCHITECTURE: BEYOND CONVENTIONAL RAG

## Traditional RAG



[Passive, Chat-Driven, Ephemeral]

## Station Workflow



[Active, Opinionated, Stateful]



Station has RAG-shaped pieces, but RAG is a feature, not its identity.  
The LLM is a participant in a structured workflow, not a chatbot wrapper.

When an idea is captured, Station cross-references external knowledge, personal history, and live infrastructure to ground responses in actual evidence, not vibes.



**Topics:** Lightweight semantic grouping.

**Memory:** Temporal facts about projects and preferences.



**Knowledge:** External docs, chunked and embedded.

**Inventory:** Live state of homelab hosts and containers.




**Research:** Local GPT Researcher deep-dive jobs.

Fact: Prefers Node.js  
for homelab.




[Superseded]

Idea captured:  
Transitioning  
homelab to Python.




Fact: Prefers Python  
for homelab.



[Active]

Every fact has temporal validity. Facts are forgotten or superseded, never silently overwritten. Station maintains a real history of what it once believed, giving crucial context to past ideas.

# INVENTORY: BRIDGING KNOWLEDGE AND HARDWARE



The semantic graph allows the LLM to answer:  
"Which of my ideas relate to the specific service running on this exact host?"

Station SSHs into Docker hosts, parses compose files, snapshots configs, and generates synopses from READMEs.


# THE ENGINE TEARDOWN

## Backend Core

Python 3.12, FastAPI,  
SQLAlchemy 2.x async.



## Unified Storage



PostgreSQL (Source of Truth)  
+ Qdrant (Vectors).  
One DB, no microservices.




## Living Wiki



System docs  
editable in-app,  
syncing to git,  
feeding the  
model's index.

## Frontend Stack

Server-rendered Jinja  
+ HTMX + Tailwind.  
Zero SPA.  
Zero build step.




## Local Intelligence



LocalAI (Quantised 70B  
chat + 2560-dim embedder).

## Deployment

5 Docker Compose containers on  
a single home server.




## Safety Net



30-day  
soft-delete  
recycle bin  
for everything.

# Sovereign Schematics

Station exposes an MCP server. External clients can instantly search your ideas, your memory, inspect your homelab inventory, and write research back to the database.



**Station isn't a chat interface bolted on top of a database. It is a unified, local knowledge surface that any AI agent can use to understand you.**

- [OK] Knows me without a corporation owning that knowledge.
- [OK] Asks good questions instead of waiting passively.
- [OK] Unifies thought instead of fragmenting across a dozen apps.
- [OK] Plays well with agents via open MCP standards.
- [OK] Runs on local hardware with sovereign LLMs.

---

A most bodacious second brain. Be excellent to each other. -> GitHub: loryanstrant/Station