


# THE BUILD STAIRCASE

Hardening probabilistic exploration into deterministic products, step by earned step.



# Undisciplined AI building falls into two distinct traps.

## Stuck at Cool Prompt Swamp

Getting stuck validating cool prompts that never graduate into reliable tools.

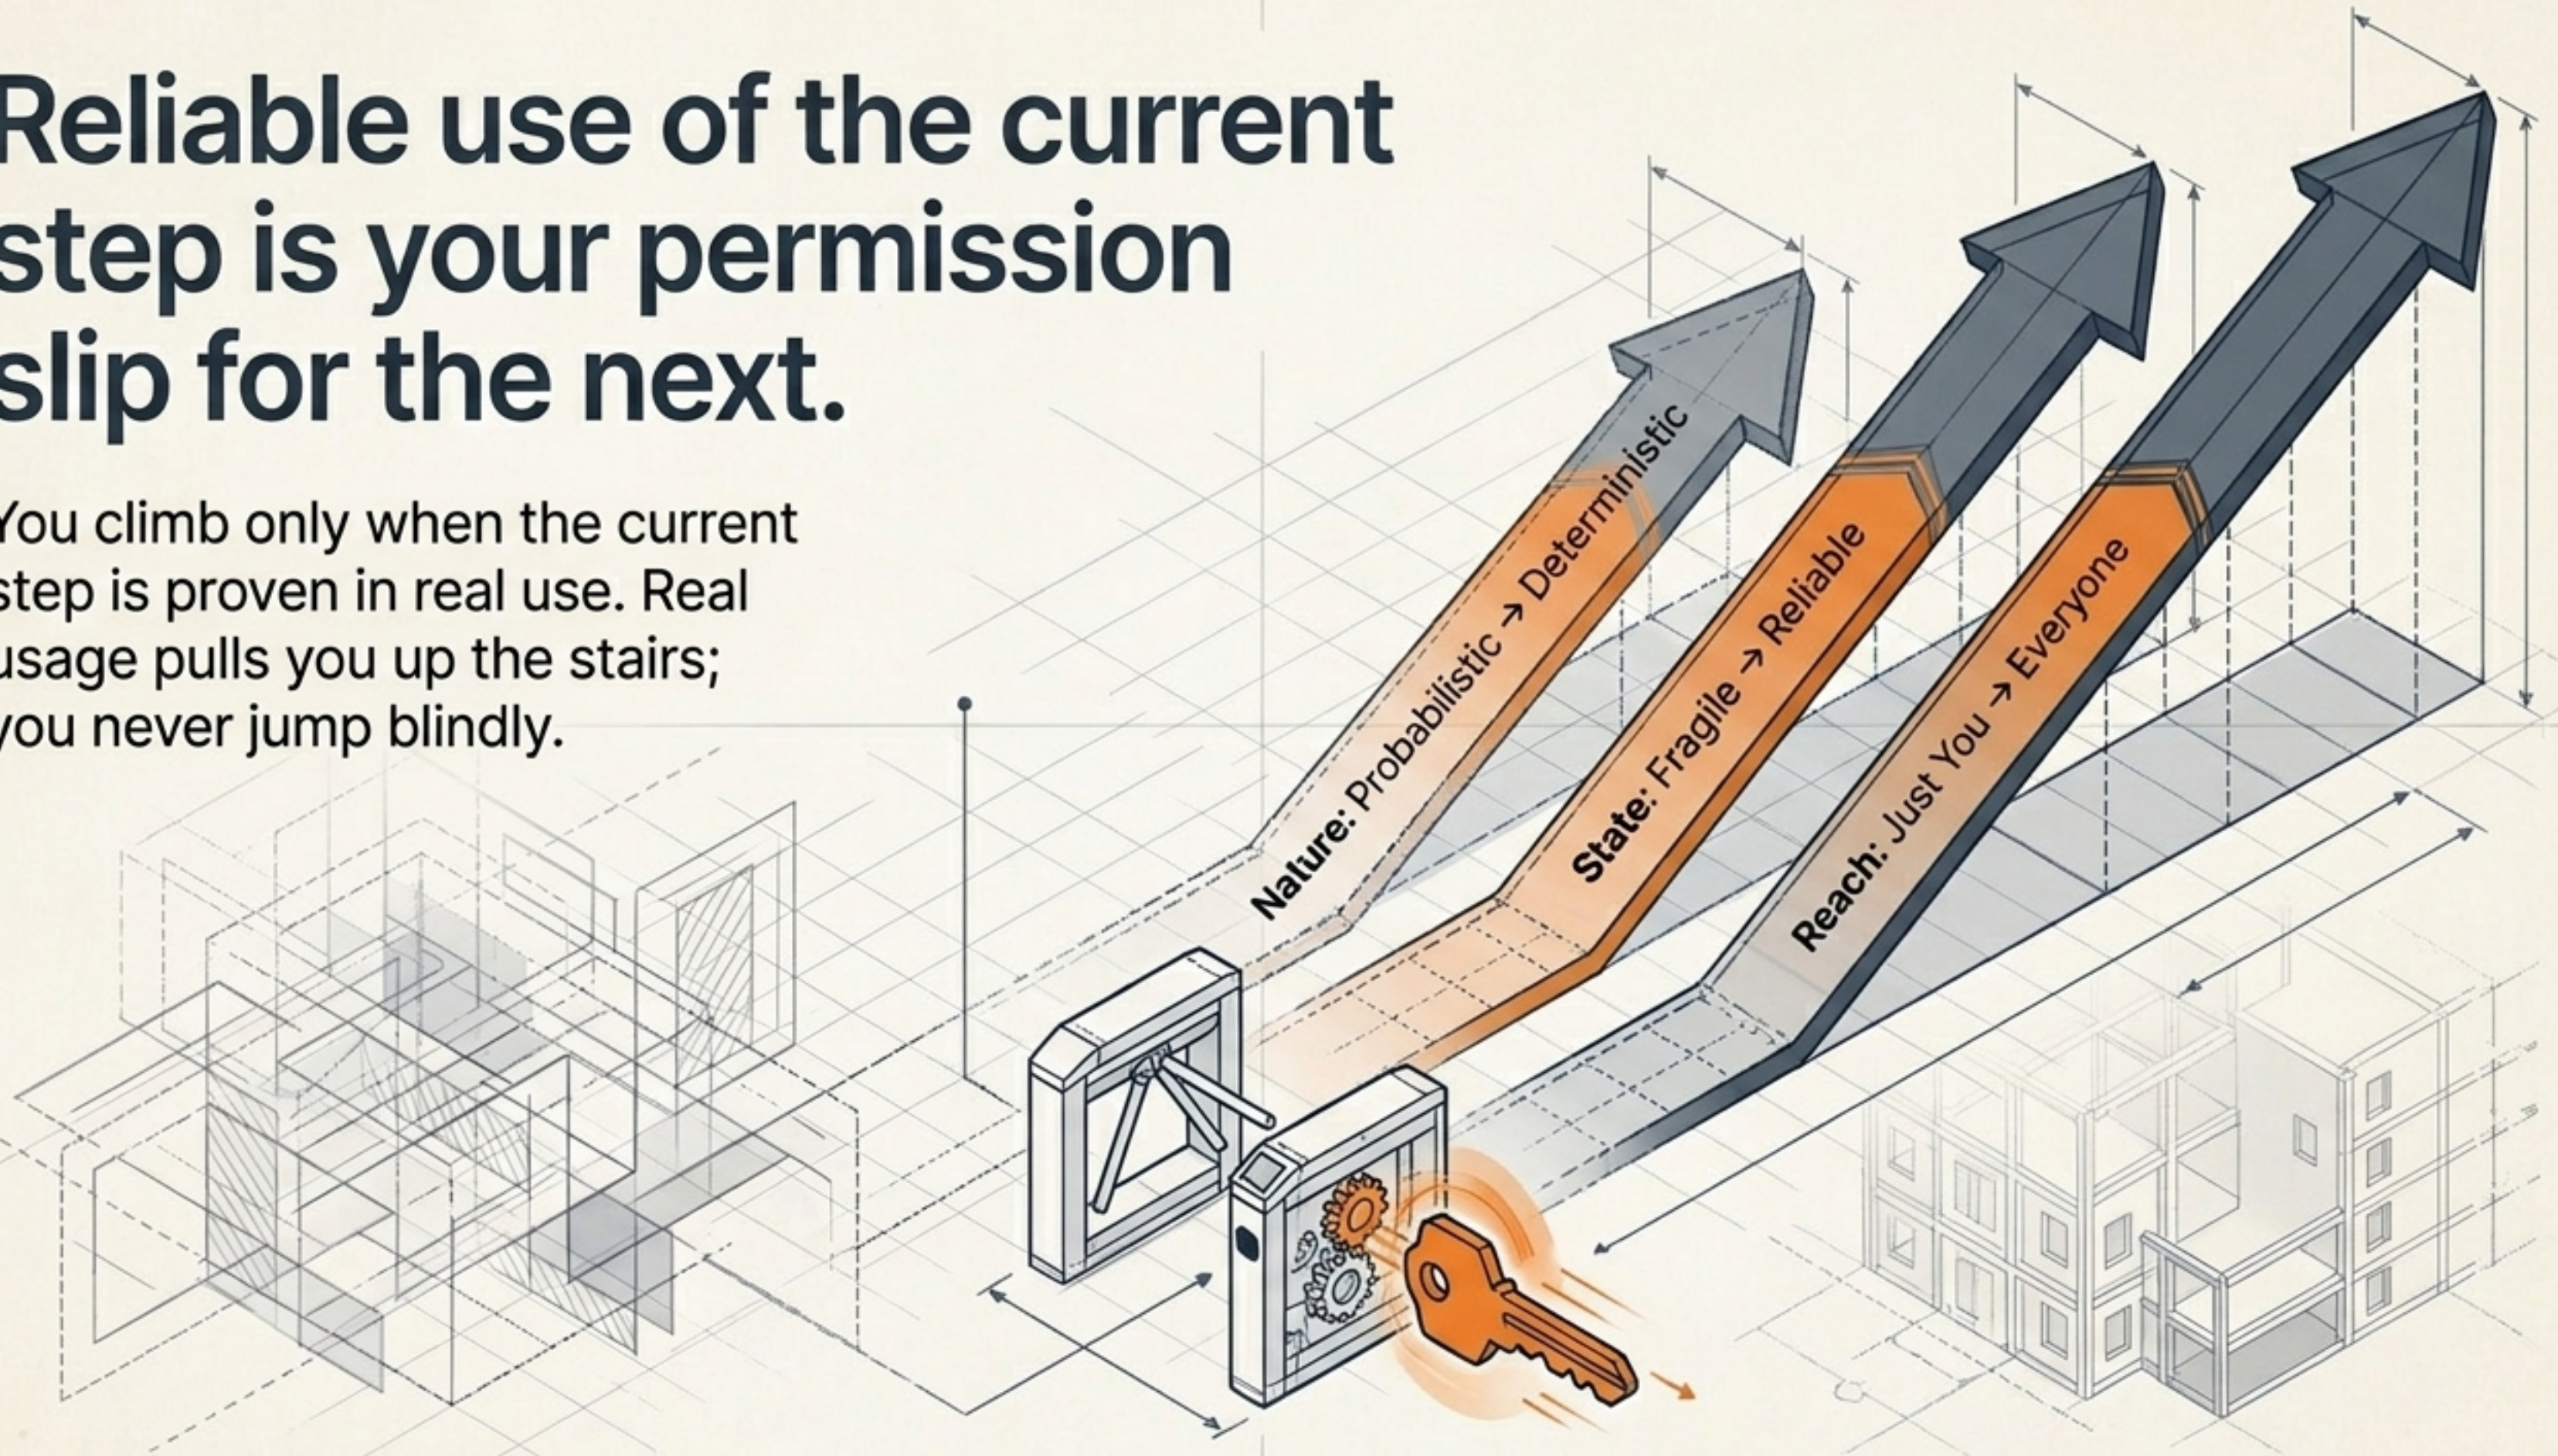
## Overbuilt App with Zero Users

Wasting heavy engineering cycles to build web apps for AI workflows nobody actually uses yet.

We need an evidence-gated bridge between these two extremes.

# Reliable use of the current step is your permission slip for the next.

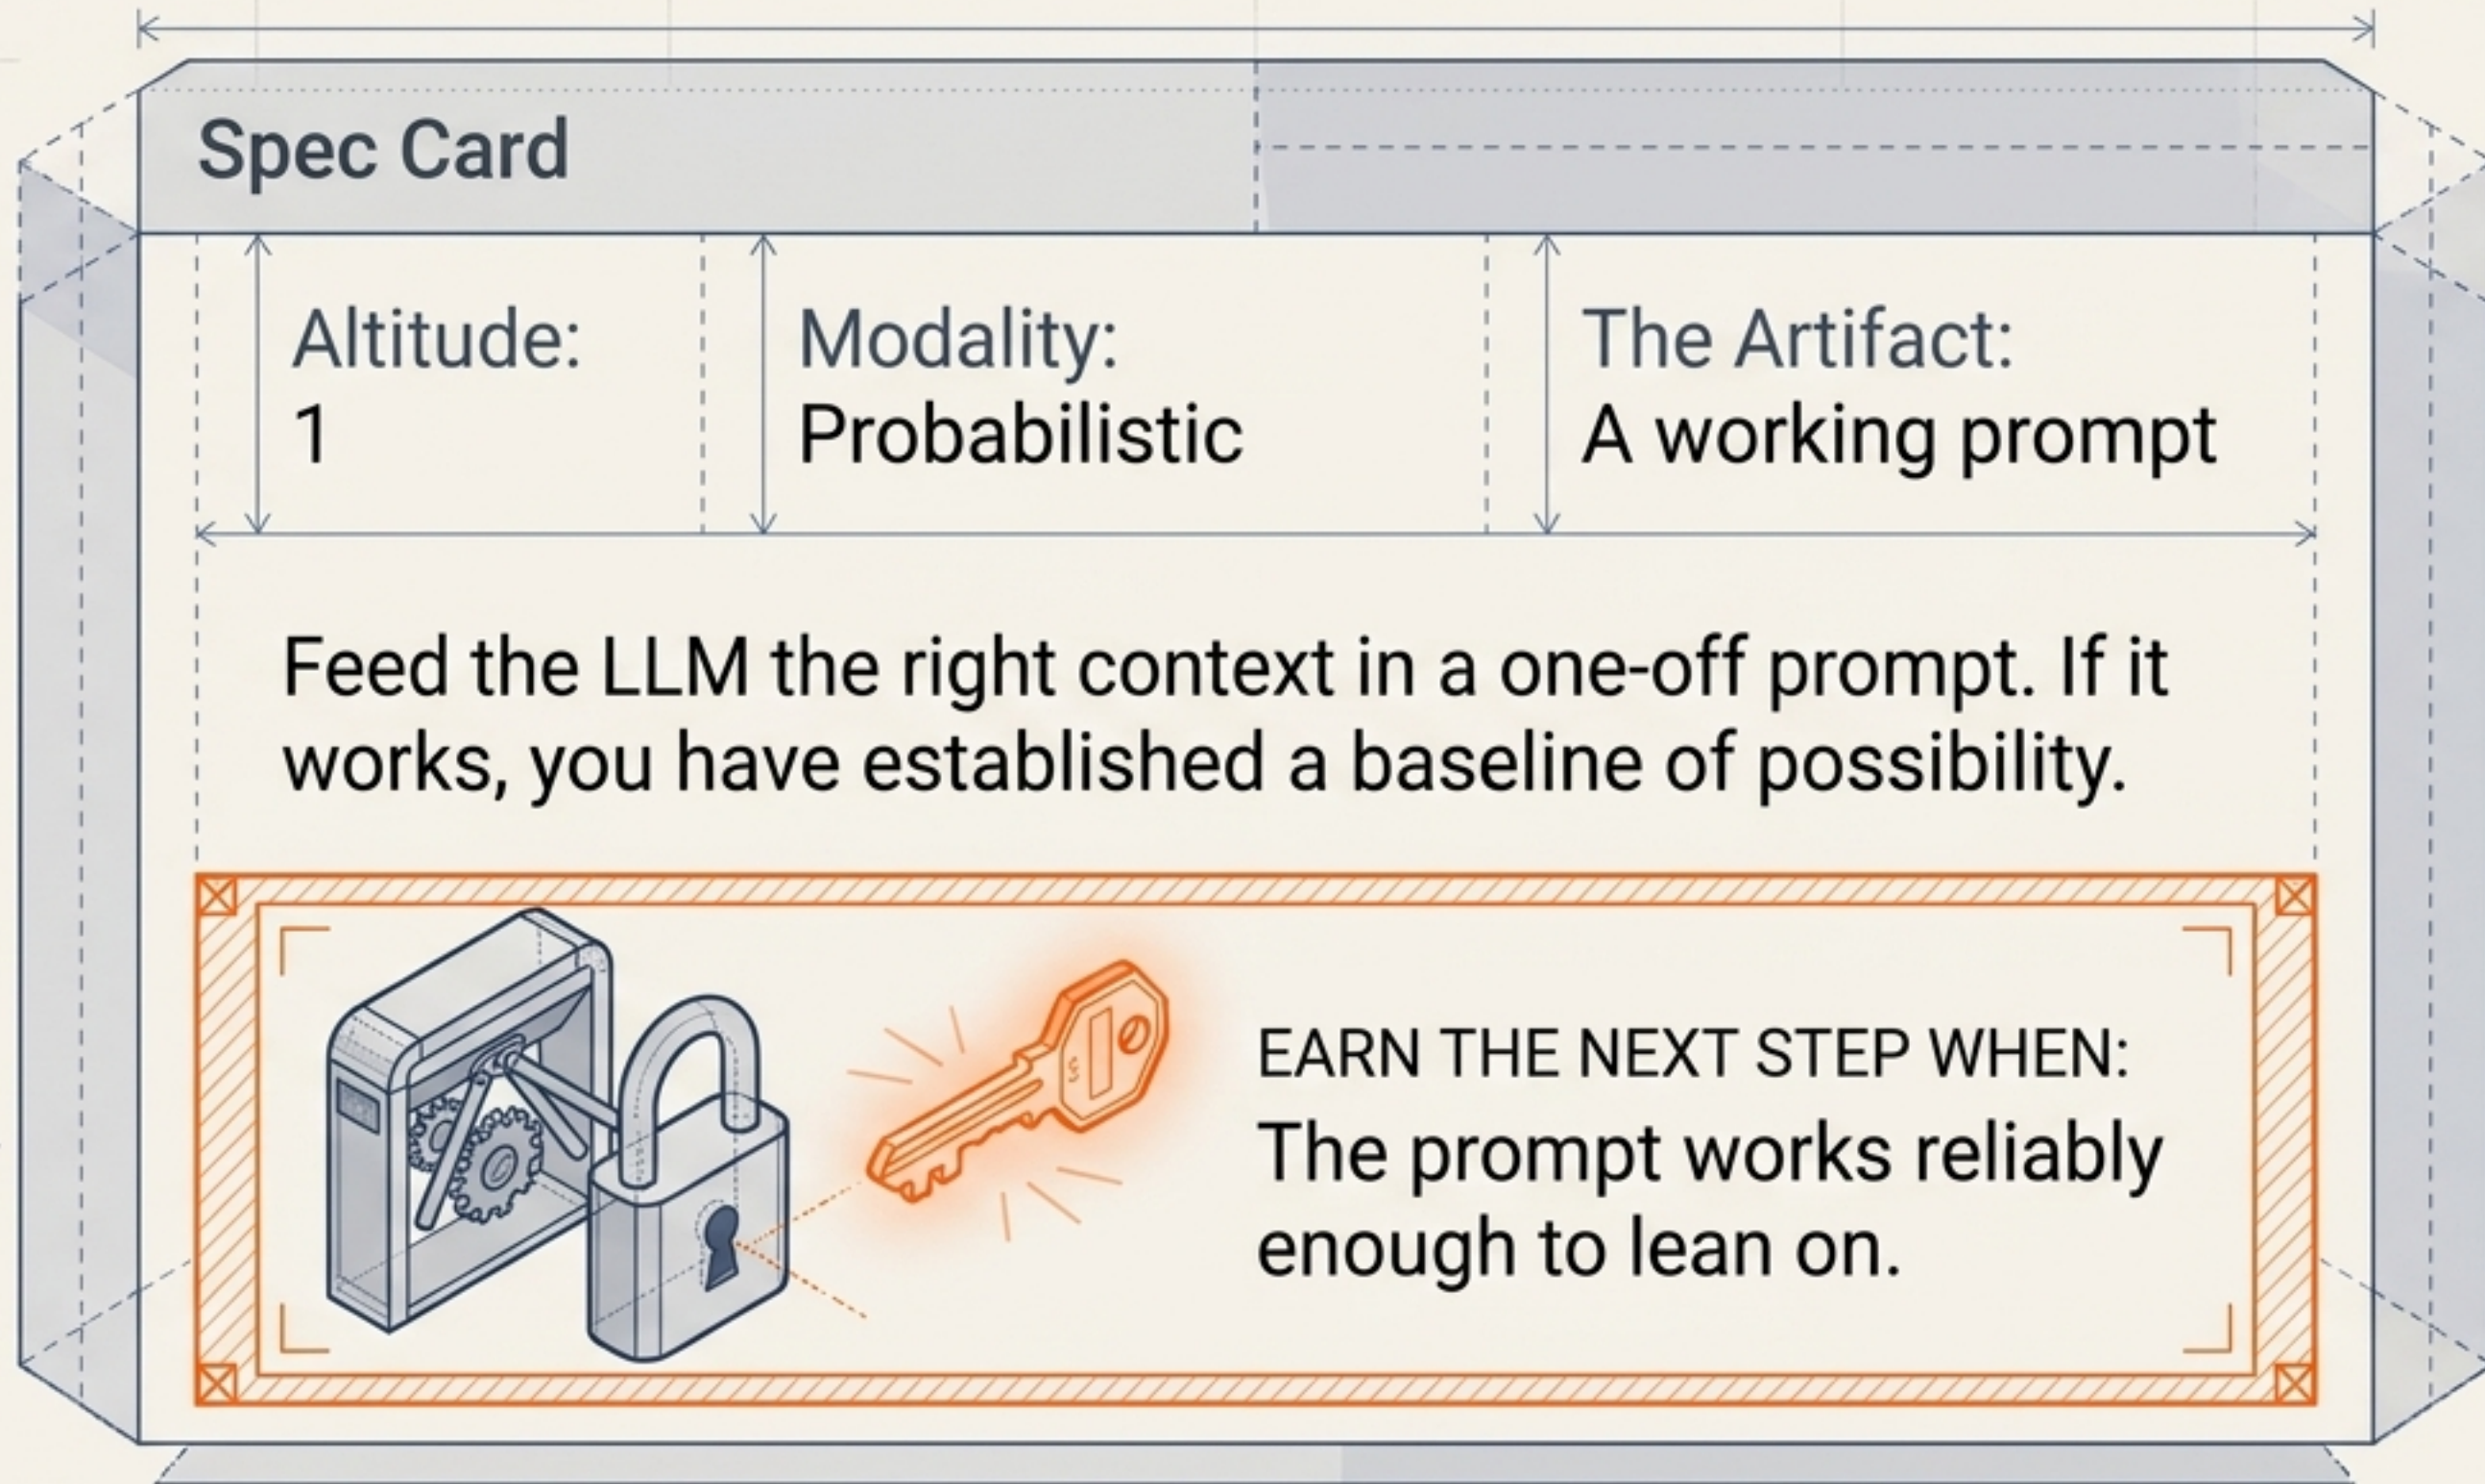
You climb only when the current step is proven in real use. Real usage pulls you up the stairs; you never jump blindly.



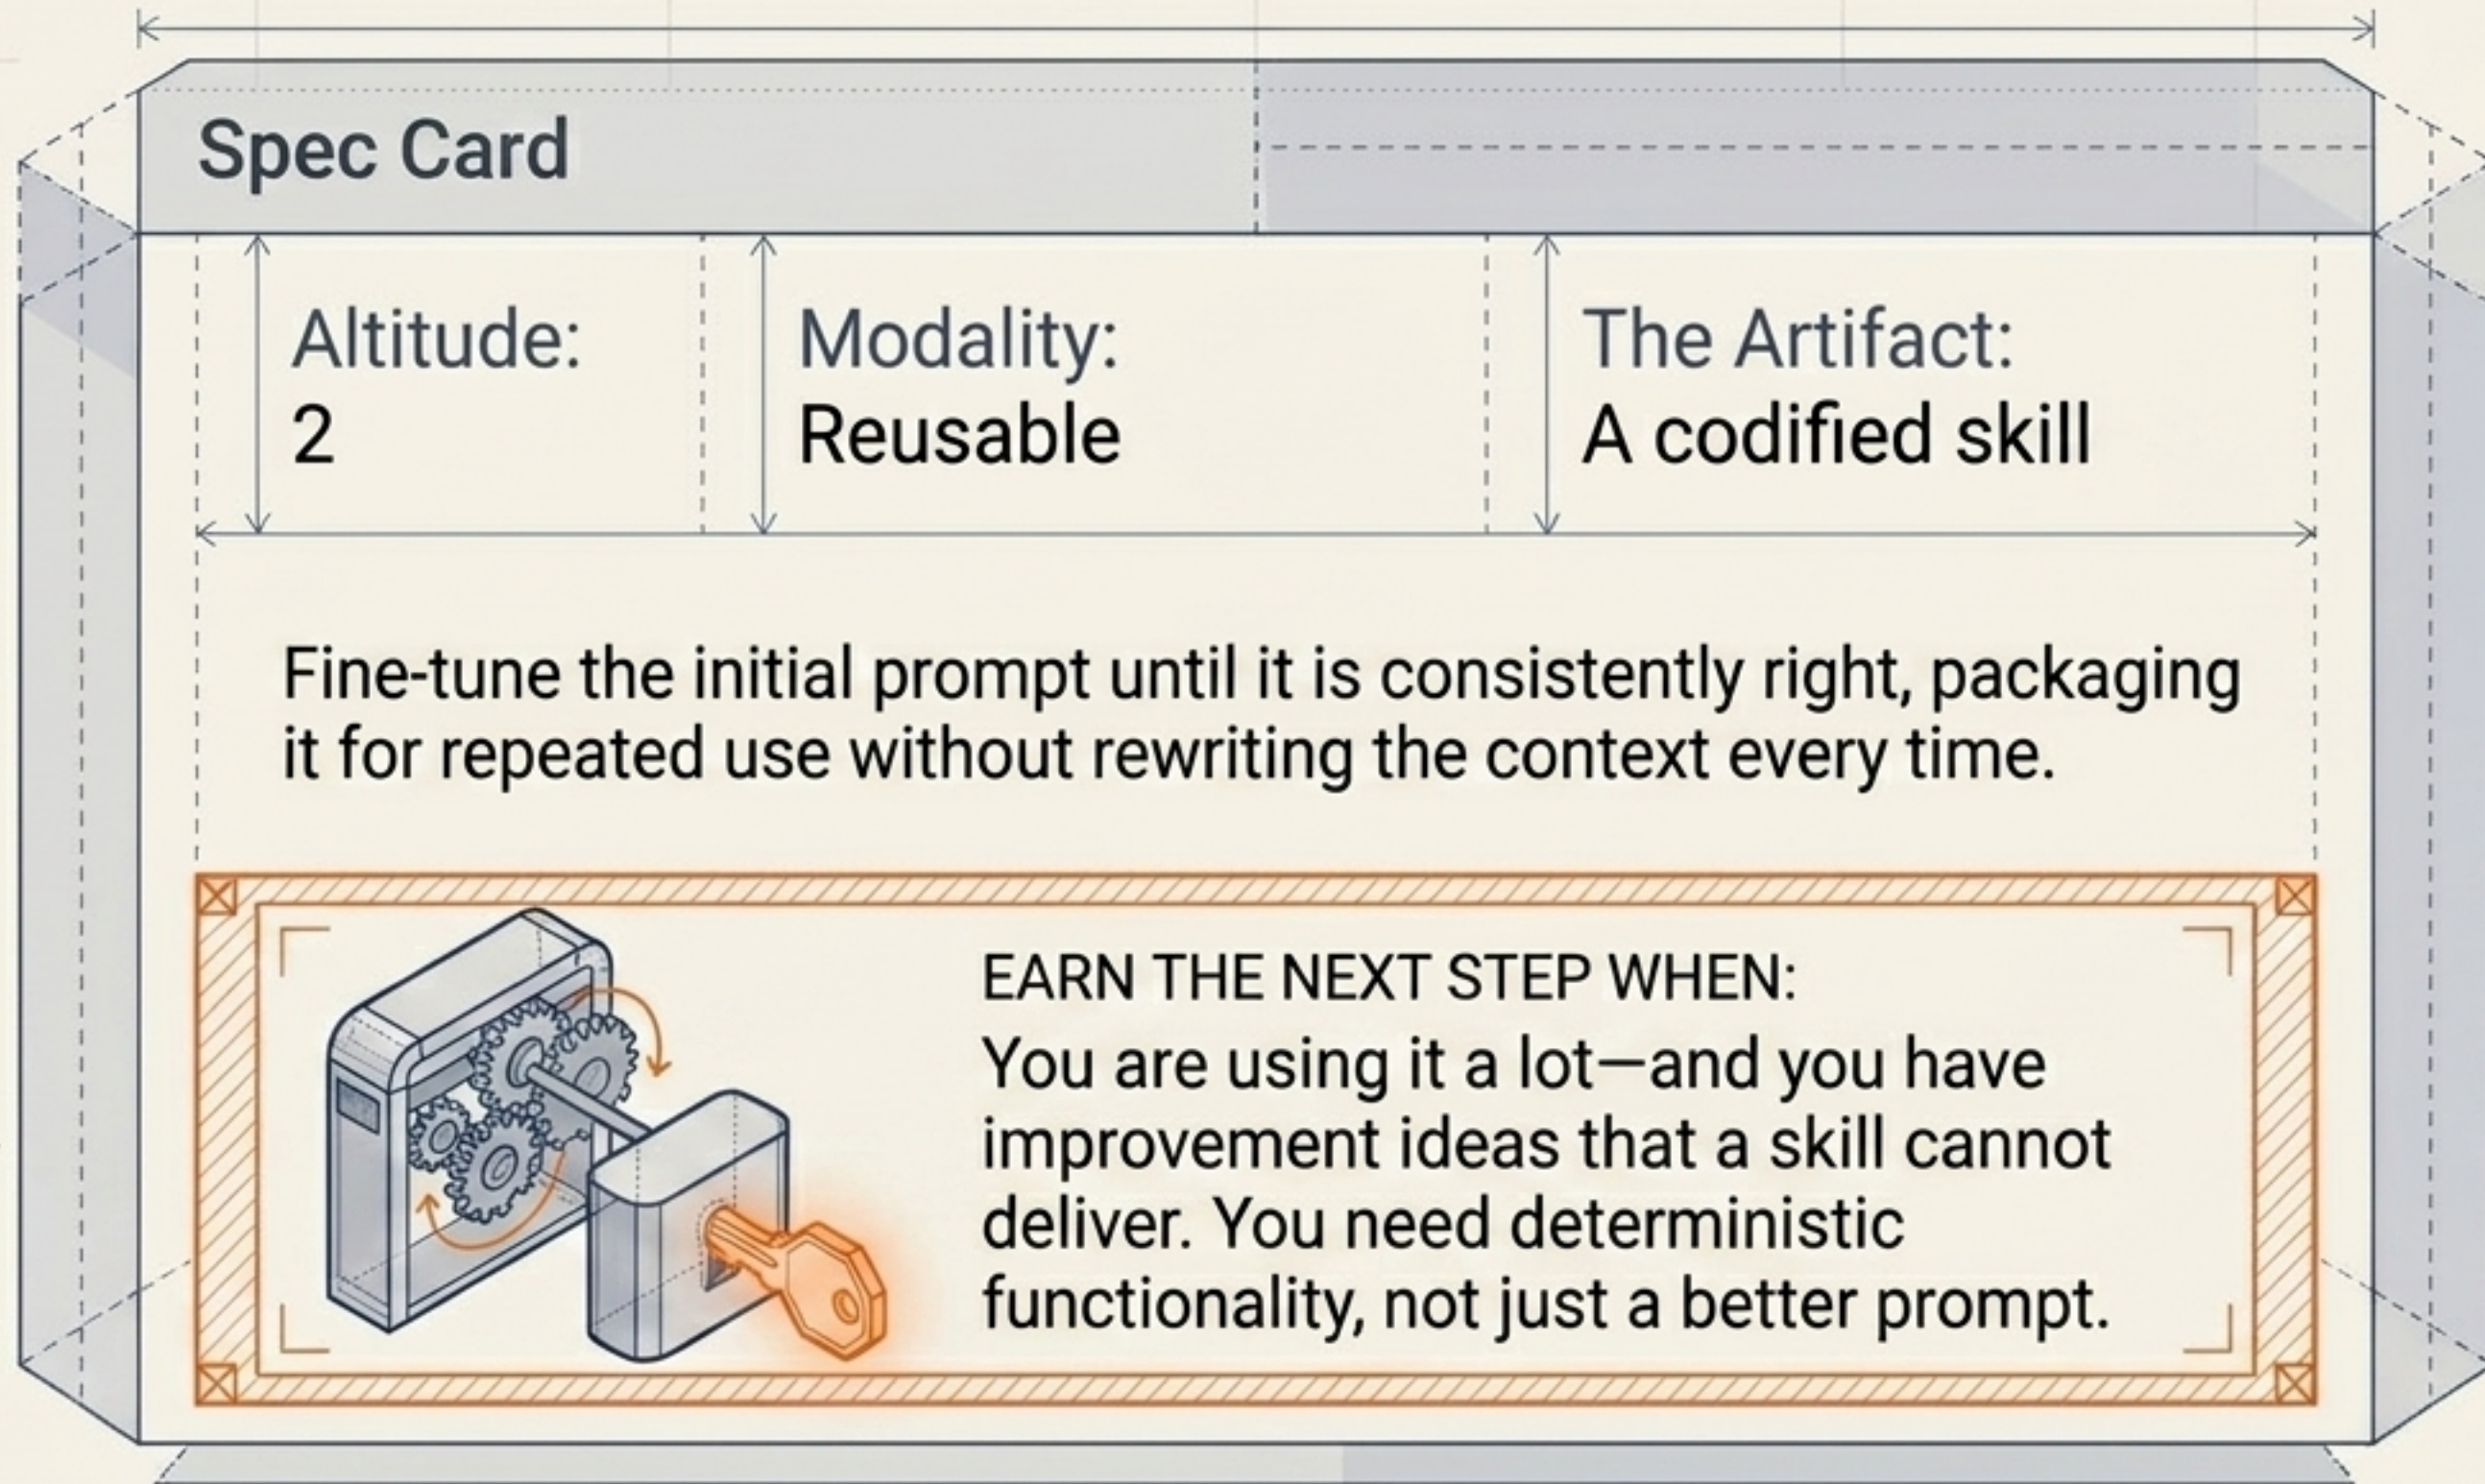
# The Staircase Diagnostic Matrix

Rung	Nature	Artifact	Primary User	The Earn-Up Gate
1. Prompt	Probabilistic	Working Prompt	Solo Explorer	Works reliably enough to lean on
2. Skill	Reusable	Codified Skill	Power User	Used a lot; needs deterministic functionality
3. CLI	Deterministic	CLI Tool	Human/Agent	Commands run reliably, every single time
<b>4. Agent</b>	<b>Orchestrated</b>	<b>Agent ↔ CLI Loop</b>	<b>Agent</b>	<b>Speed and reliability gains proven in use</b>
<b>5. App</b>	<b>Product</b>	<b>Web App</b>	<b>Everyone</b>	<b>The summit is reached</b>

# Spot a capability and prove it is possible.



# Codify the prompt into a reusable skill.



# Harden the judgment into **deterministic scripts**.

## Spec Card

Altitude:  
**3**

Modality:  
**Deterministic**

The Artifact:  
**A CLI Tool**

The probabilistic judgment gets wrapped in hardened, testable code. These command-line scripts can be called reliably by humans or agents.



**EARN THE NEXT STEP WHEN:**  
The commands run reliably,  
every single time.

# Point an agent at the CLI to orchestrate execution.

## Spec Card

Altitude:  
4

Modality:  
**Orchestrated**


The Artifact:  
**An Agent ↔  
CLI Loop**

### Description

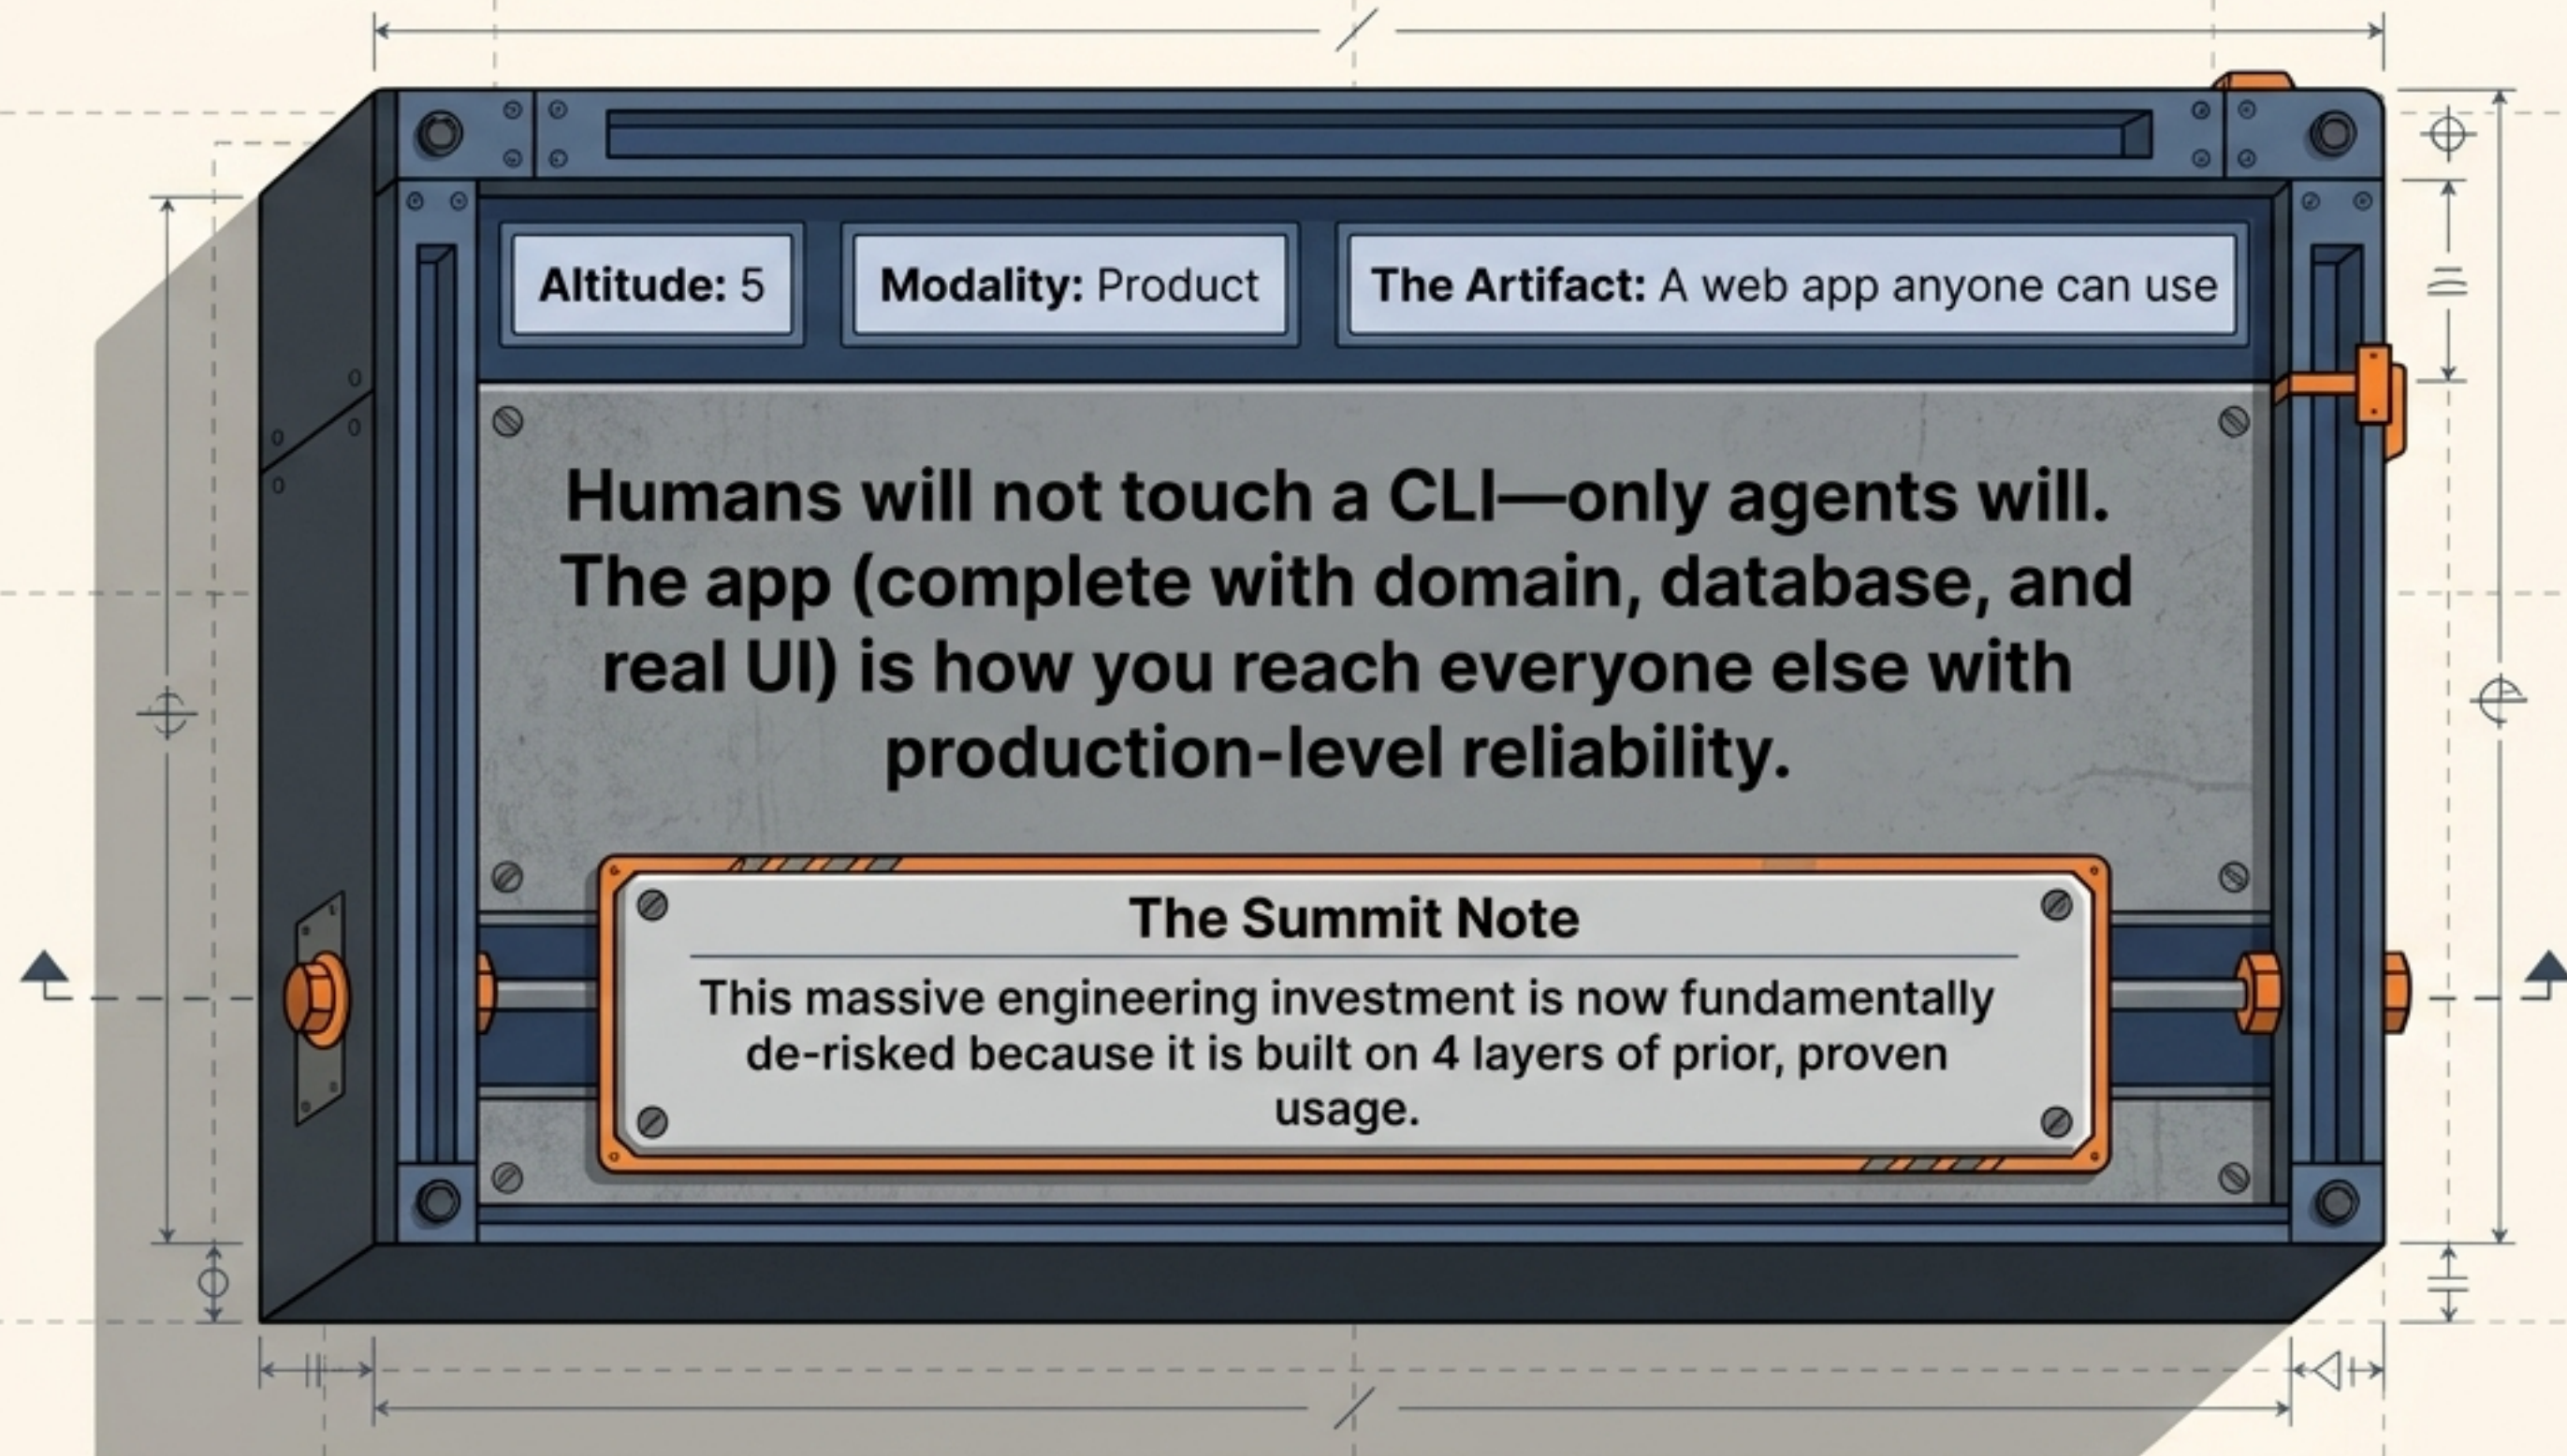
The work stays legible and auditable. You get a massive speed boost, but you can literally see whether the agent ran the correct command.

### EARN THE NEXT STEP WHEN:

The speed and reliability gains are proven in real use.



# Democratize access with a hardened product.



# MCP is a scout, not a structural rung.

## The Drone (MCP):

Runs alongside early rungs to explore what is possible. Cheap discovery, zero build effort.

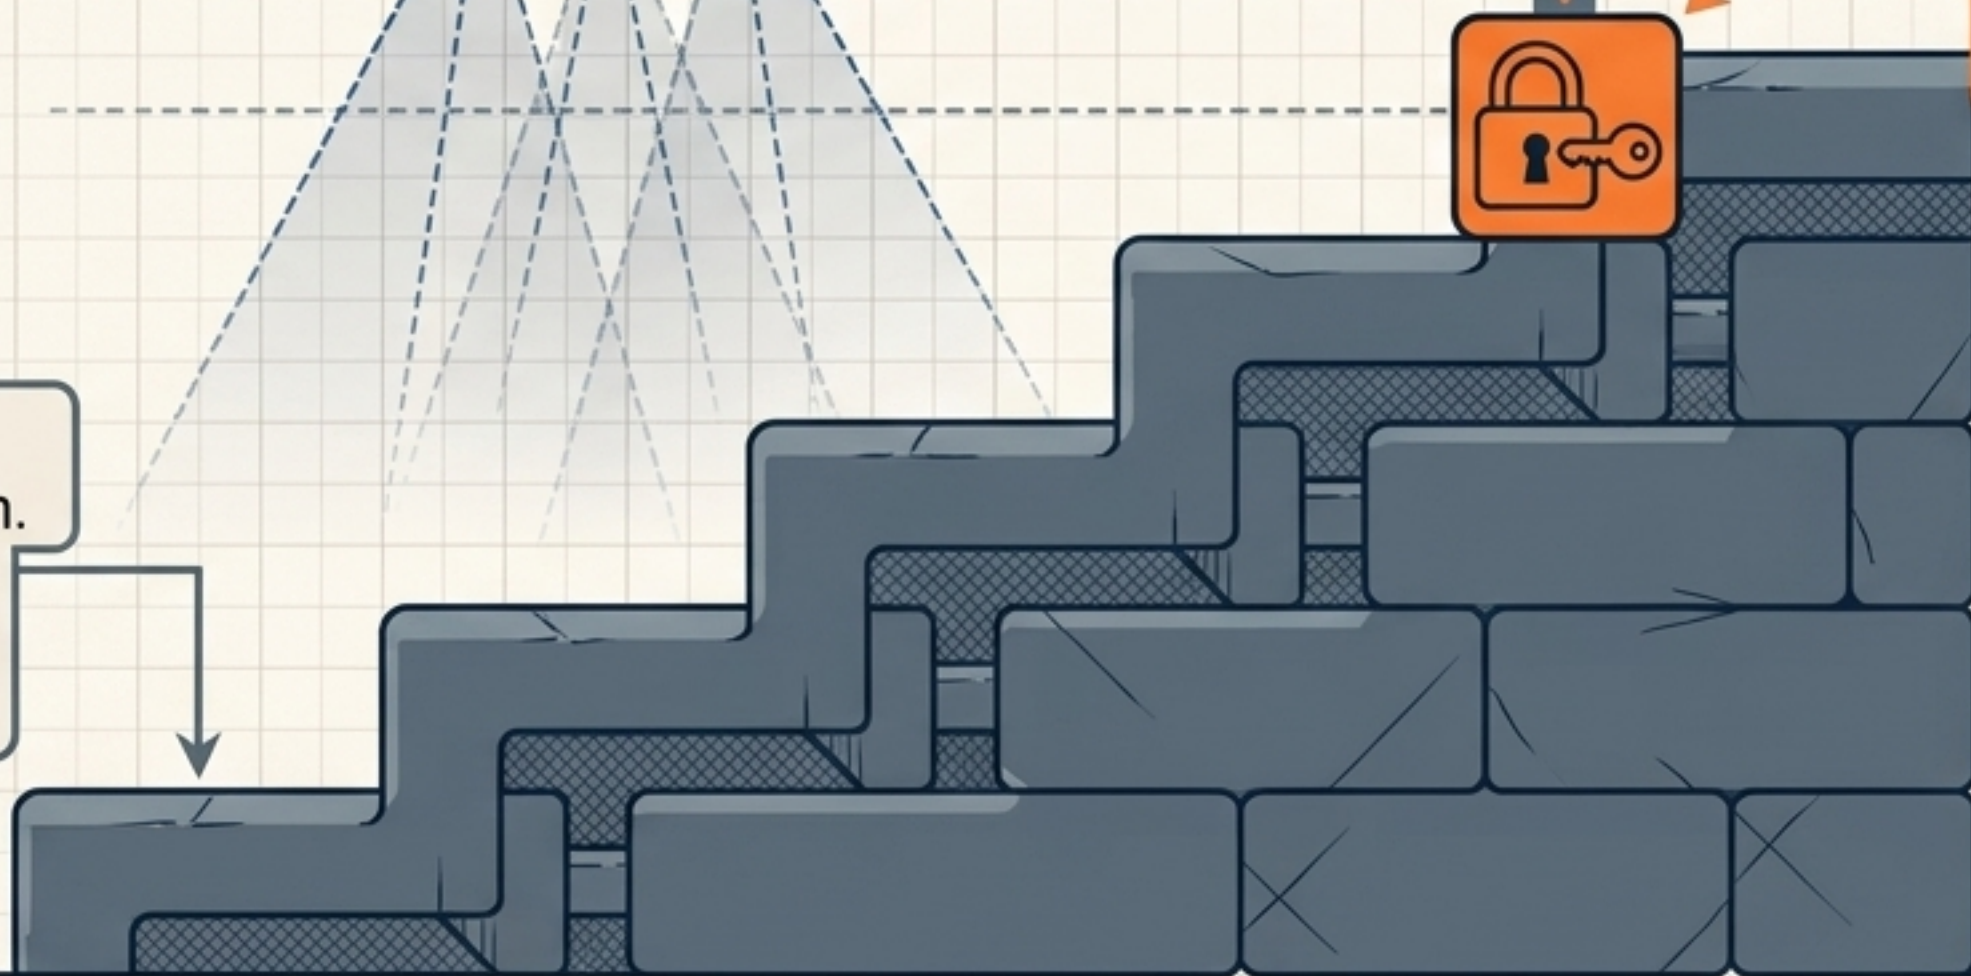


## The Graduation Rule:

When an MCP run is proven and repeated, it must pass the evidence gate and graduate into a deterministic CLI command.

## The Stone (CLI):

The structural foundation. If you are building it yourself, build a CLI, not a bespoke MCP.

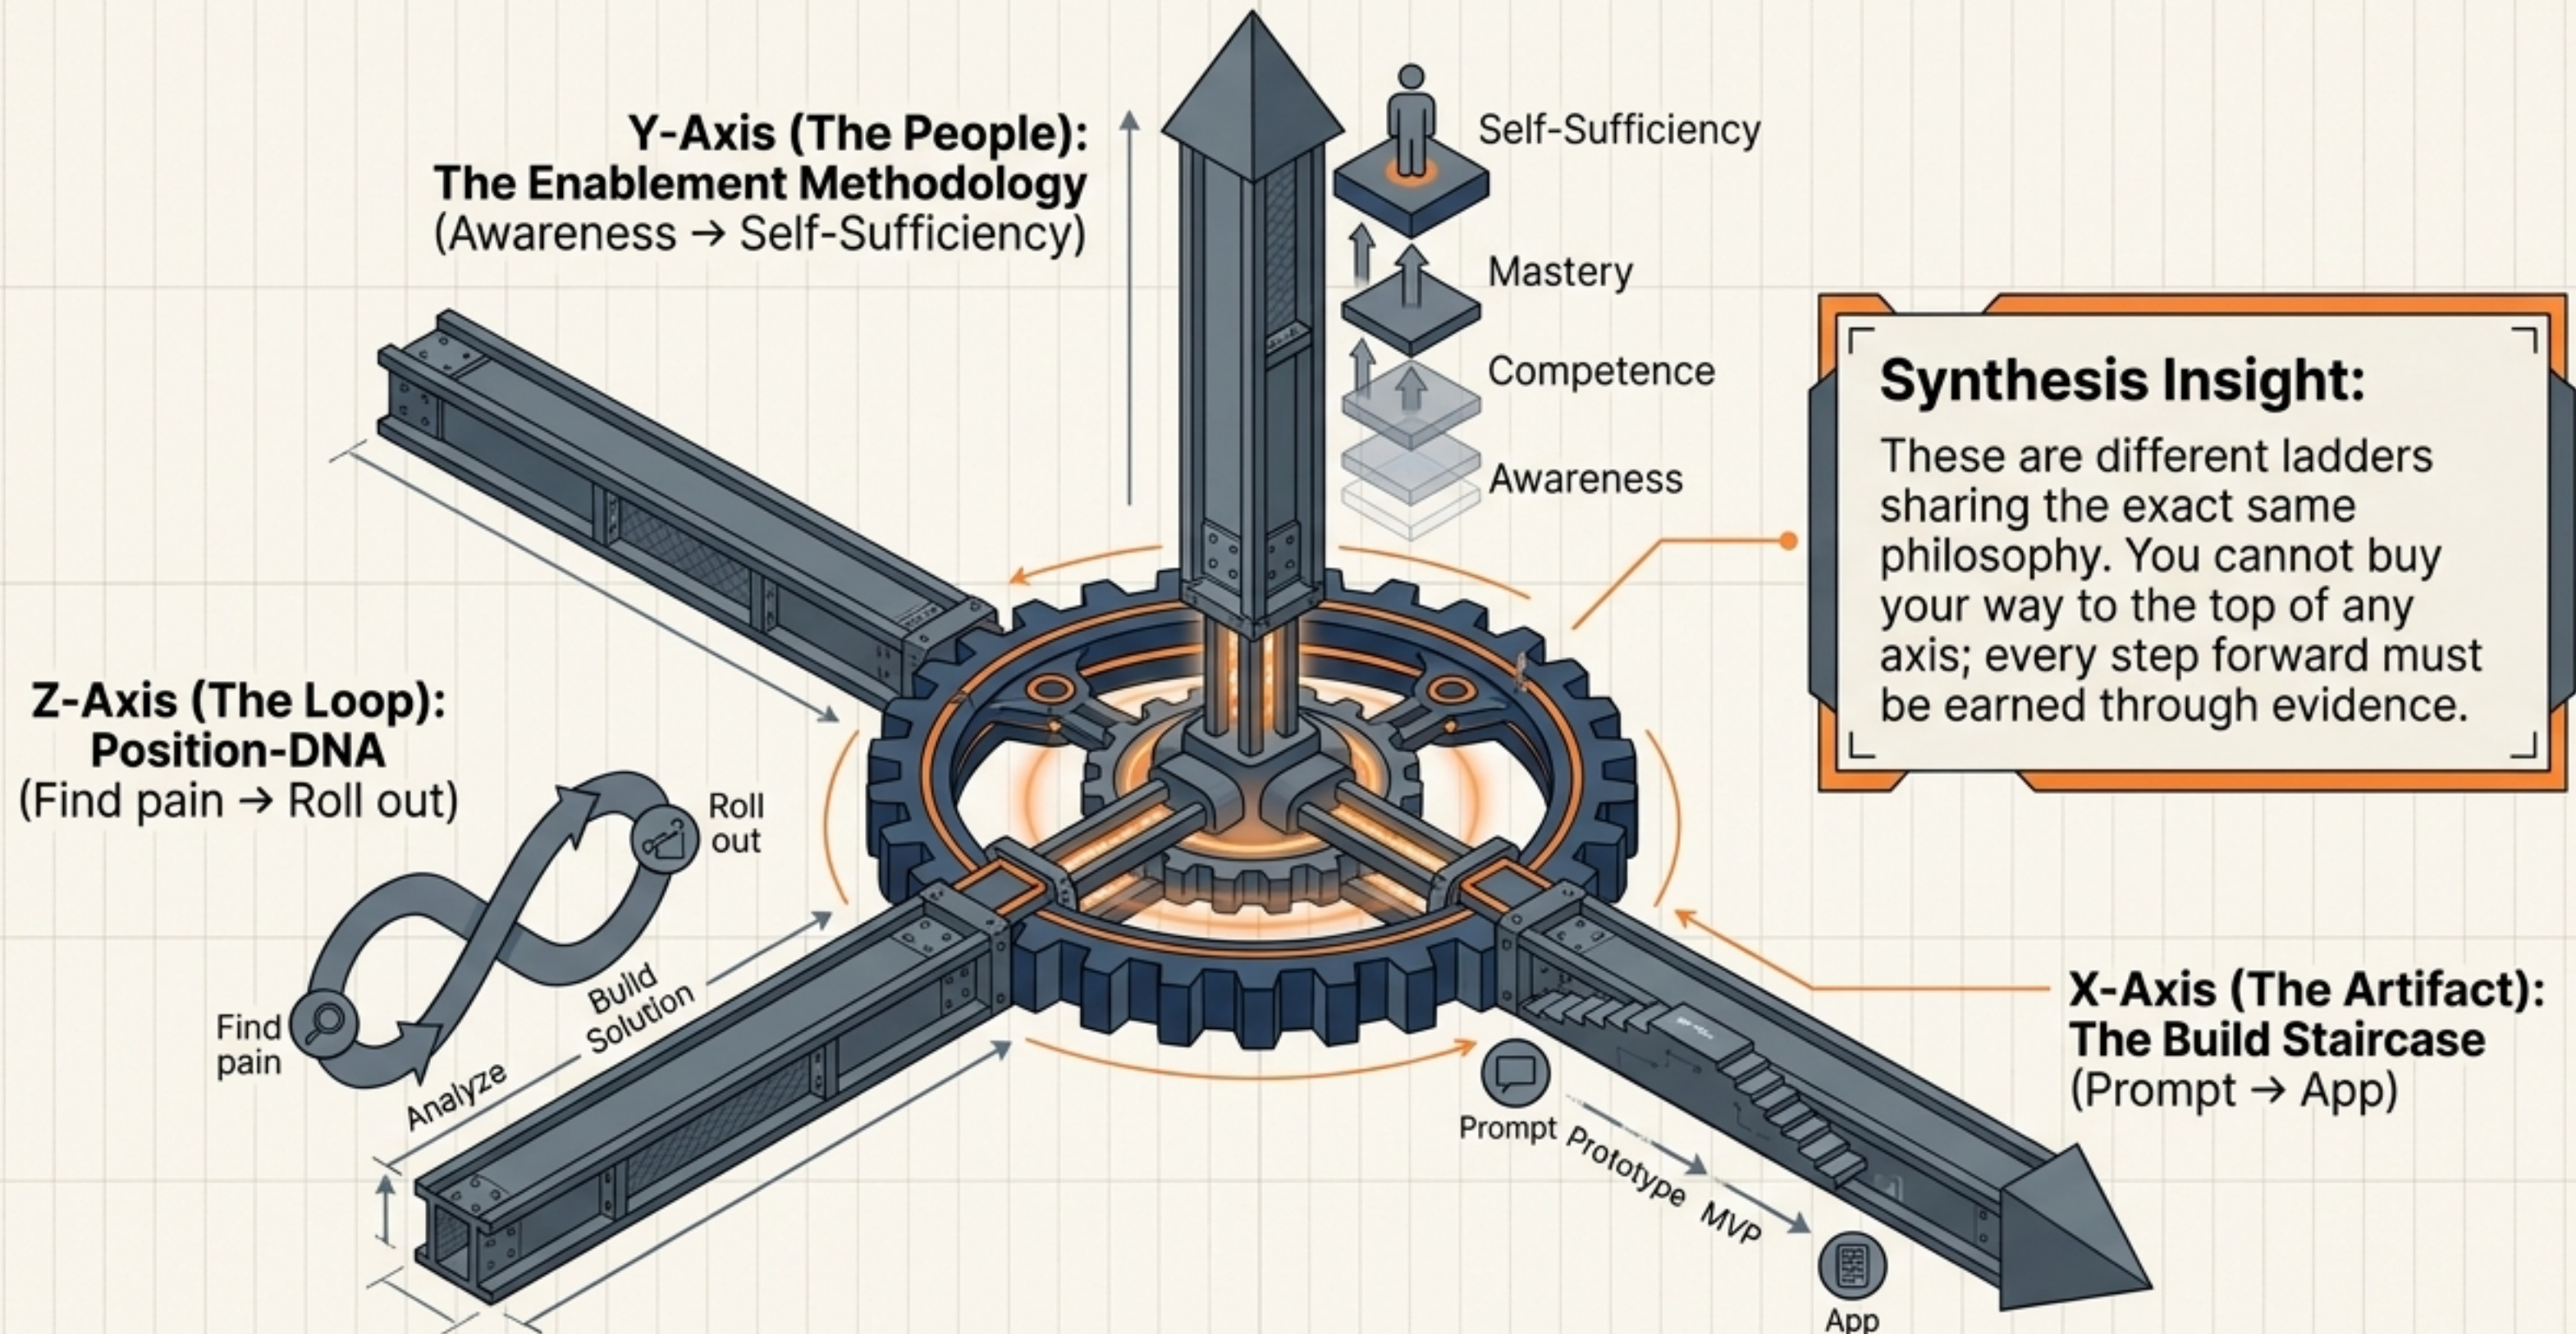


# The Rule: MCP to discover, CLI to drive.

Model Context Protocol (MCP)	Command Line Interface (CLI)
<b>The Scout</b>	<b>The Driver</b>
Exploration of the possible	Deterministic execution
Cheap discovery (if existing)	Hardened investment
<b>Heavy</b> (tool definitions tax agent context)	Light & contained


Takeaway: Do not let heavy context taxes degrade your agent's performance. Harden proven discoveries into CLIs.

# The philosophy of earned steps is fractal.



# Usage is the ultimate signal.

Disciplined AI building means never investing the next tier of engineering until the current tier proves its value.



**This evidence-gated escalation is the only way to guarantee that your AI capabilities remain reliable, reachable, and fundamentally business-owned and measured.**