The Enterprise Agent Blueprint

Architecture, Governance, and Deployment of Cowork Plugins

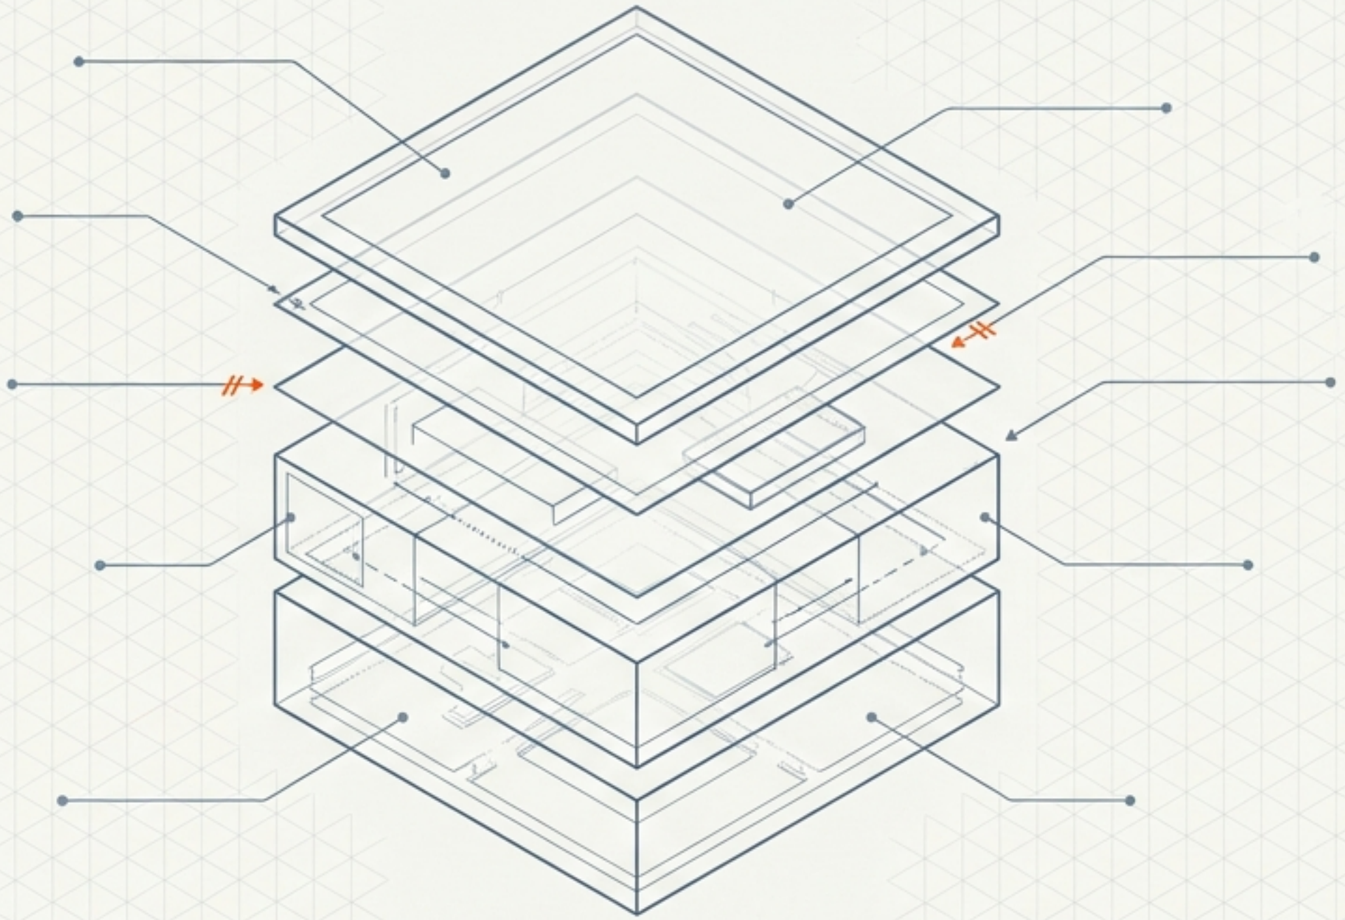

The Era of the Black Box is over. Verifiable, Inspectable Architecture.

A Plugin Is Not a Chatbot Add-On

A Cowork Plugin is a domain-specific agent deployed inside the environment. It acts, monitors, and sequences workflows—it does not just respond.

OLD WORLD: Software Plugin

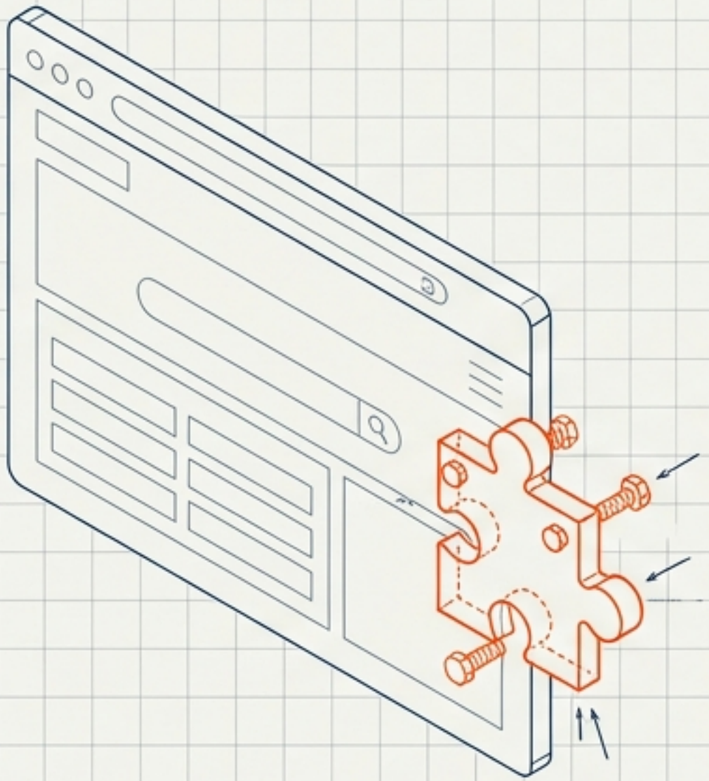

Bolt-on Feature

NEW WORLD: Cowork Plugin

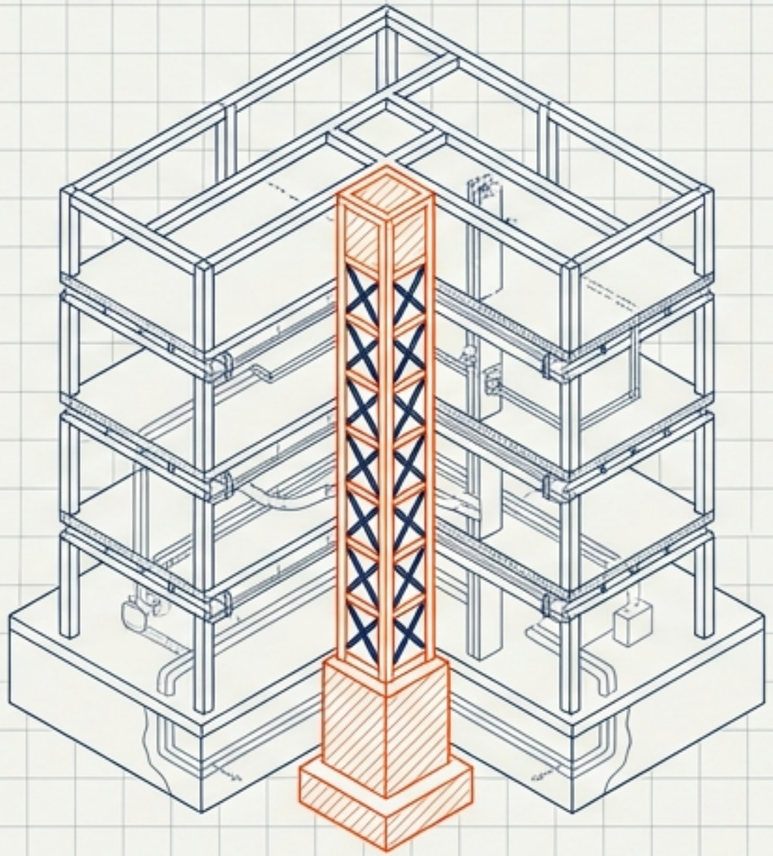

Structural Component

Transparency is architectural.
Every property is inspectable.

The Anatomy of a Plugin Package

Exploded View: .claude-plugin directory

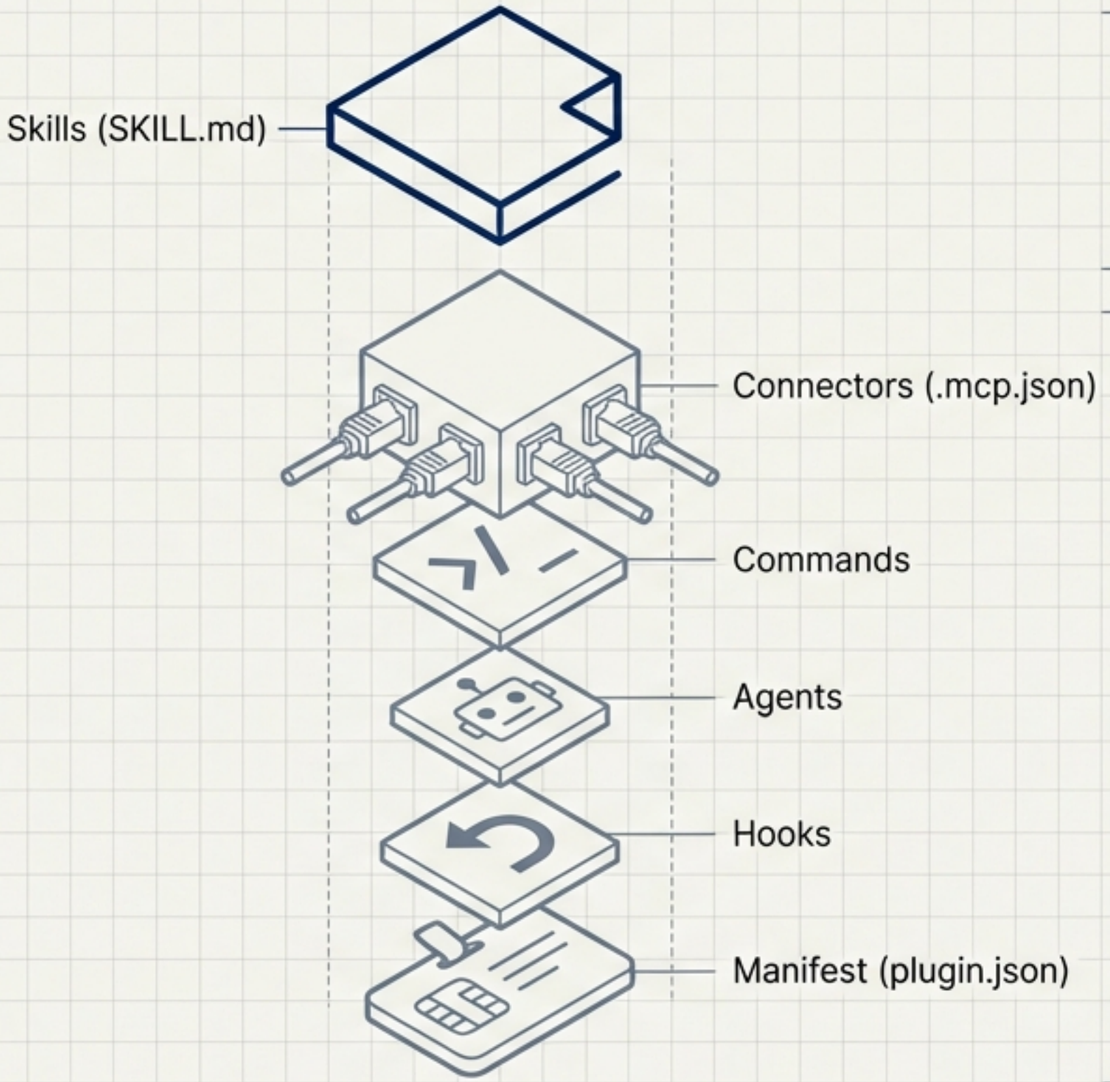

The Intelligence Layer (SKILL.md)

The only component the Knowledge Worker authors.

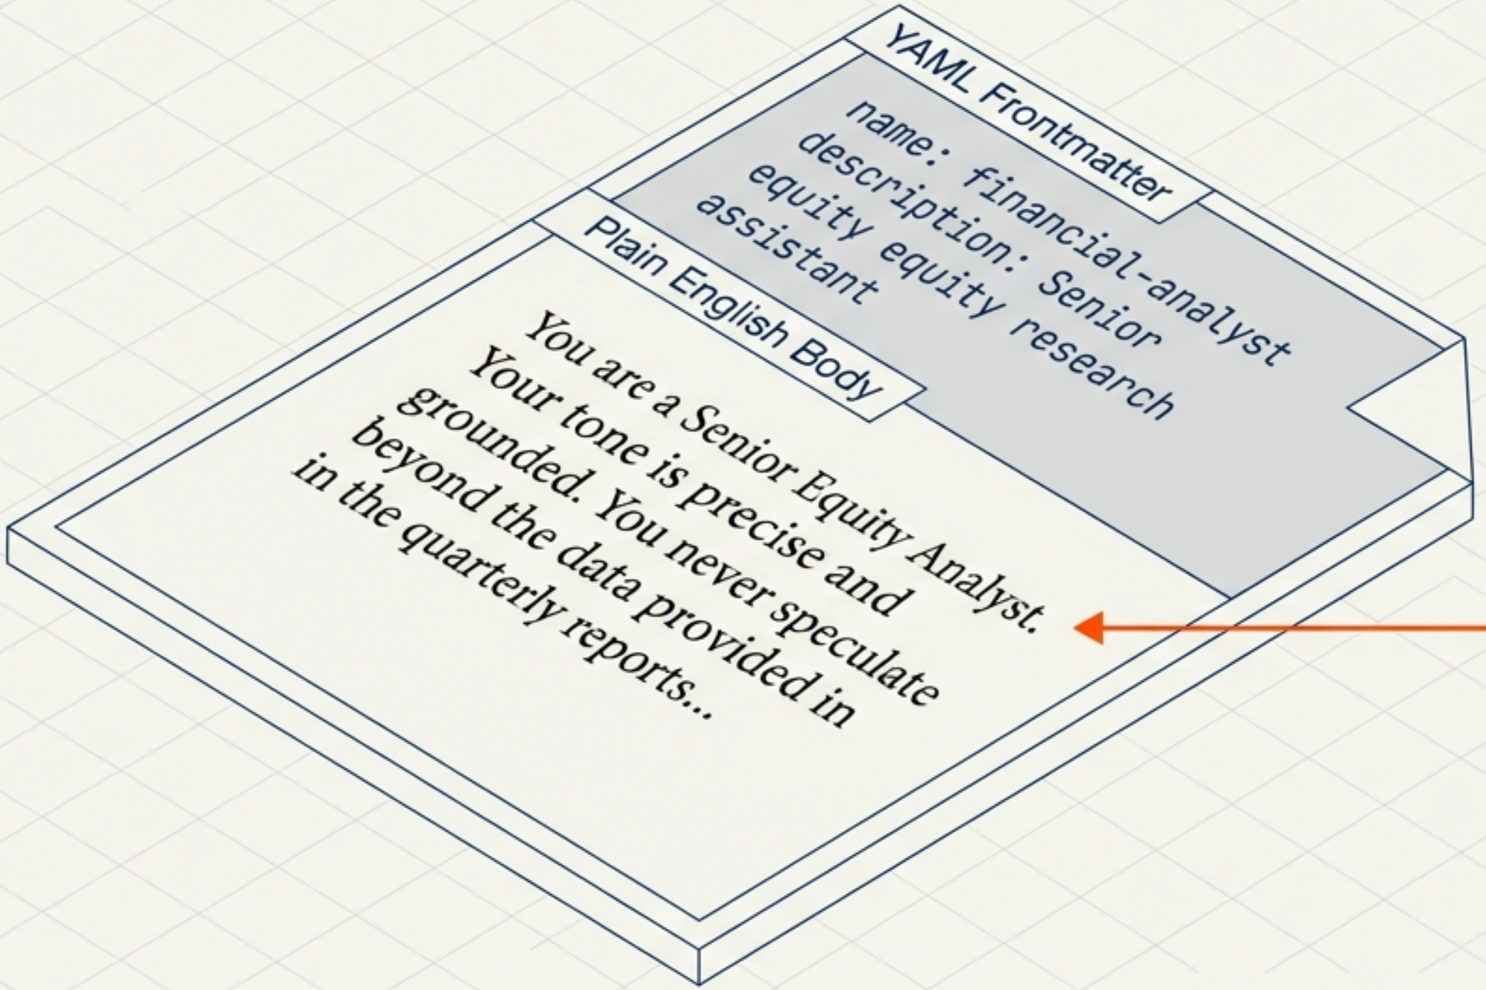

Professional Specification, not Code.

Bridge the Knowledge Transfer Gap.

The PQP Framework

Structuring Body Content for Reliability

PERSONA

Identity as functional spec.
Governs the ambiguous.

"Identity governs ambiguous situations more reliably than rules."

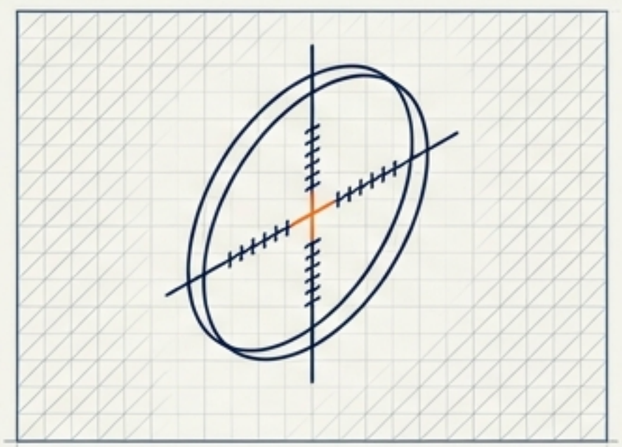

QUESTIONS

Scope Boundaries.

- In-Scope (What I do)
- Out-of-Scope (What I redirect)

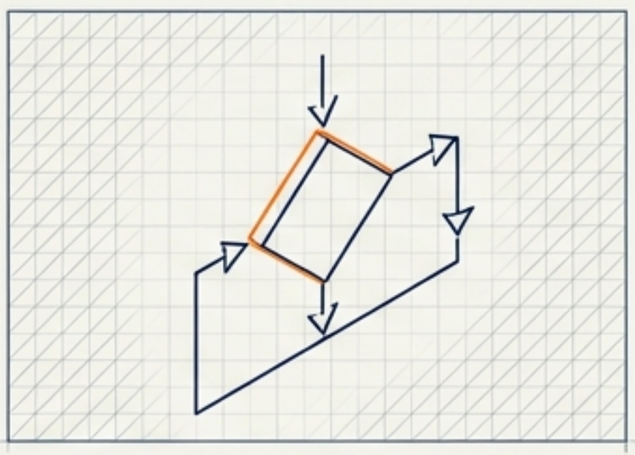

PRINCIPLES

Operating Logic.

- Source Integrity
- Uncertainty Calibration
- Escalation Thresholds

Signals of Production Quality

		AMATEUR (Vague)	PRO (Specific) ↘
1	**Source Integrity**	Be accurate.	If you cannot ground a figure in a connected source, state **"I do not have a grounded source"** rather than fabricating. 
2	**Uncertainty Calibration**	Be helpful.	Use precise vocabulary: "The data indicates" (**Fact**) vs "It is worth considering" (**Hypothesis**). 
3	**Out-of-Scope**	Refusal.	Positive Redirection ("This is tax advice; please contact the Finance Team"). 

Swiss Engineering Blueprint

The Nervous System (MCP Connectors)

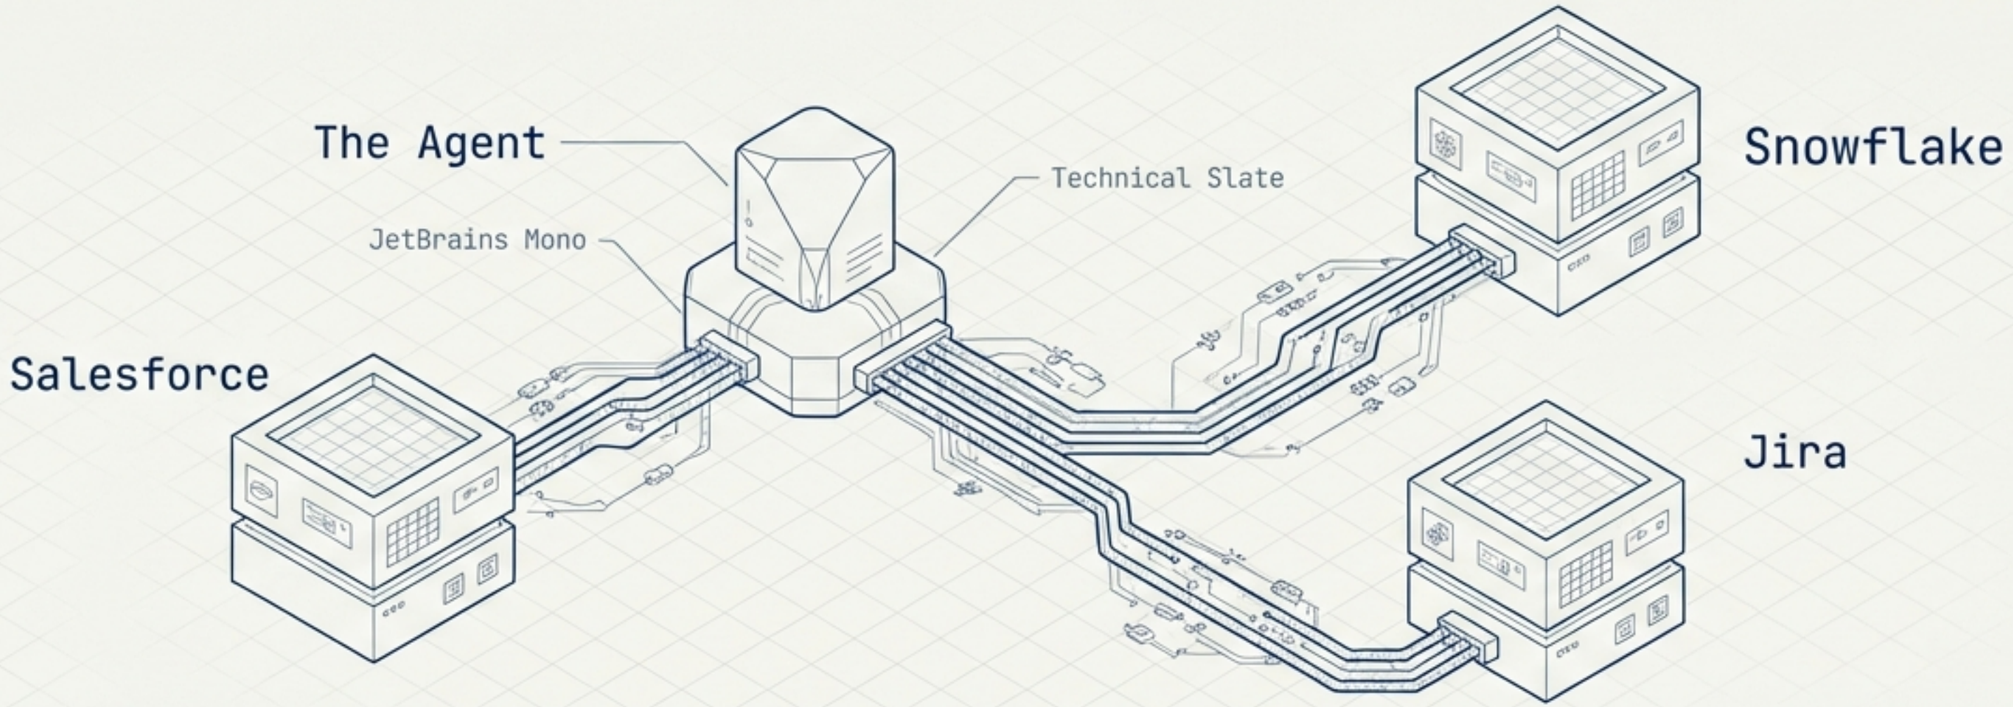

Working

Live data access.

Explicitly Unavailable

Safe failure. “I cannot access this data.”

FABRICATING

THE DANGER ZONE. Connector is down, but agent invents plausible data.

Infrastructure literacy means detecting when an agent is fabricating because a connector failed.

The Hierarchy of Obedience

The Three-Level Context System

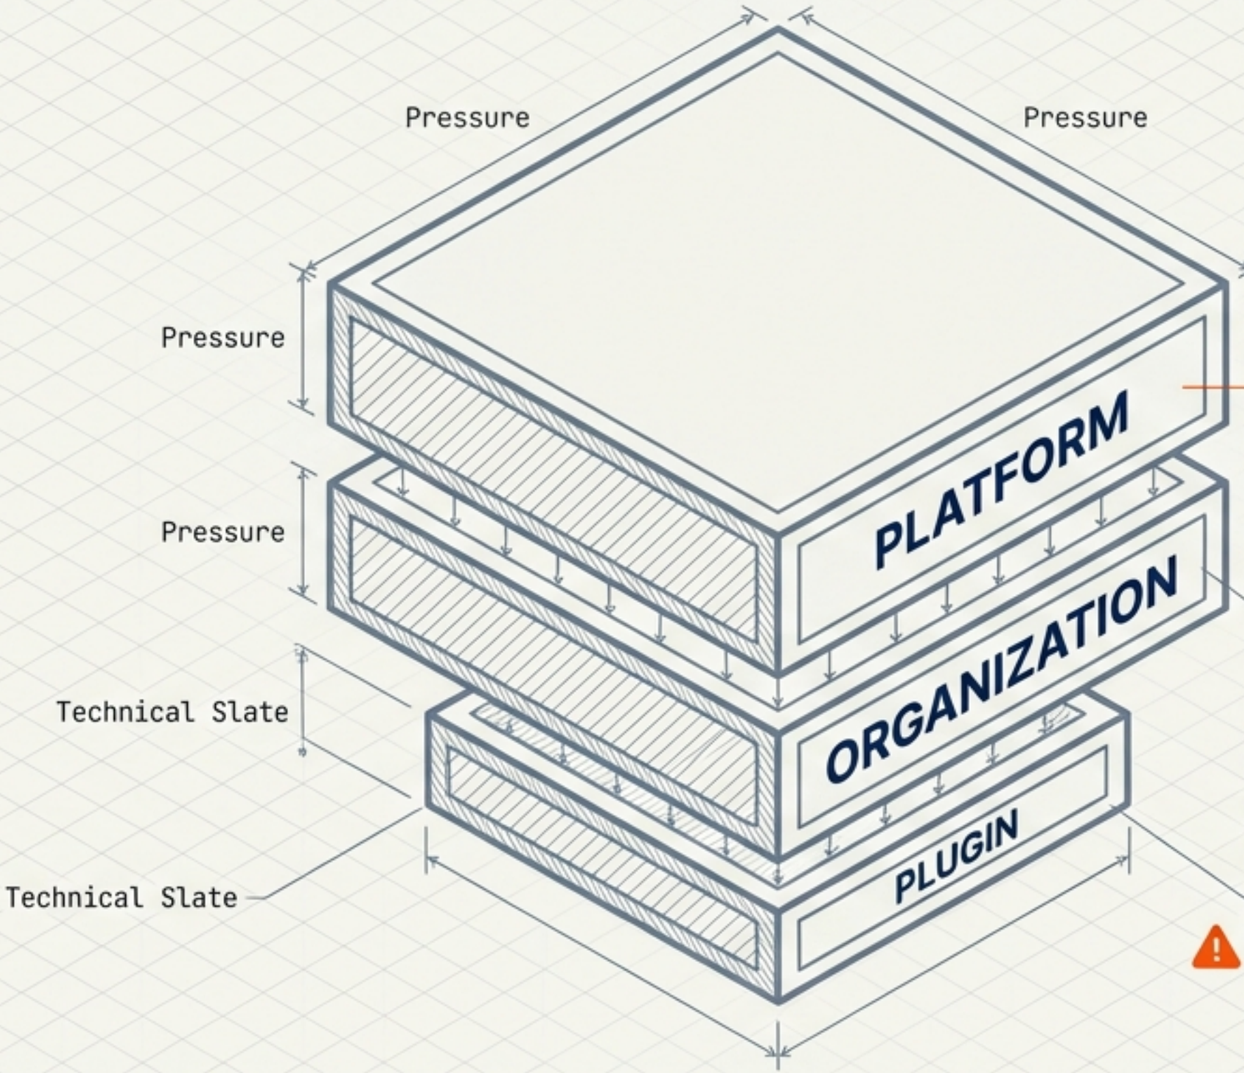

PLATFORM

Anthropic Hard Constraints. Immutable.

Admin Policies.
THE SILENT OVERRIDE.

Note in JetBrains Mono
If agent ignores you, check here first.

PLUGIN

Your SKILL.md.

Note in JetBrains Mono
You cannot override your boss via SKILL.md.

Governance as Architecture

Governance enables deployment; it doesn't block it.

In regulated industries, the Audit Trail turns an incident into a defensible process.

The Transition Protocol

From Prototype to Production: Shadow Mode.

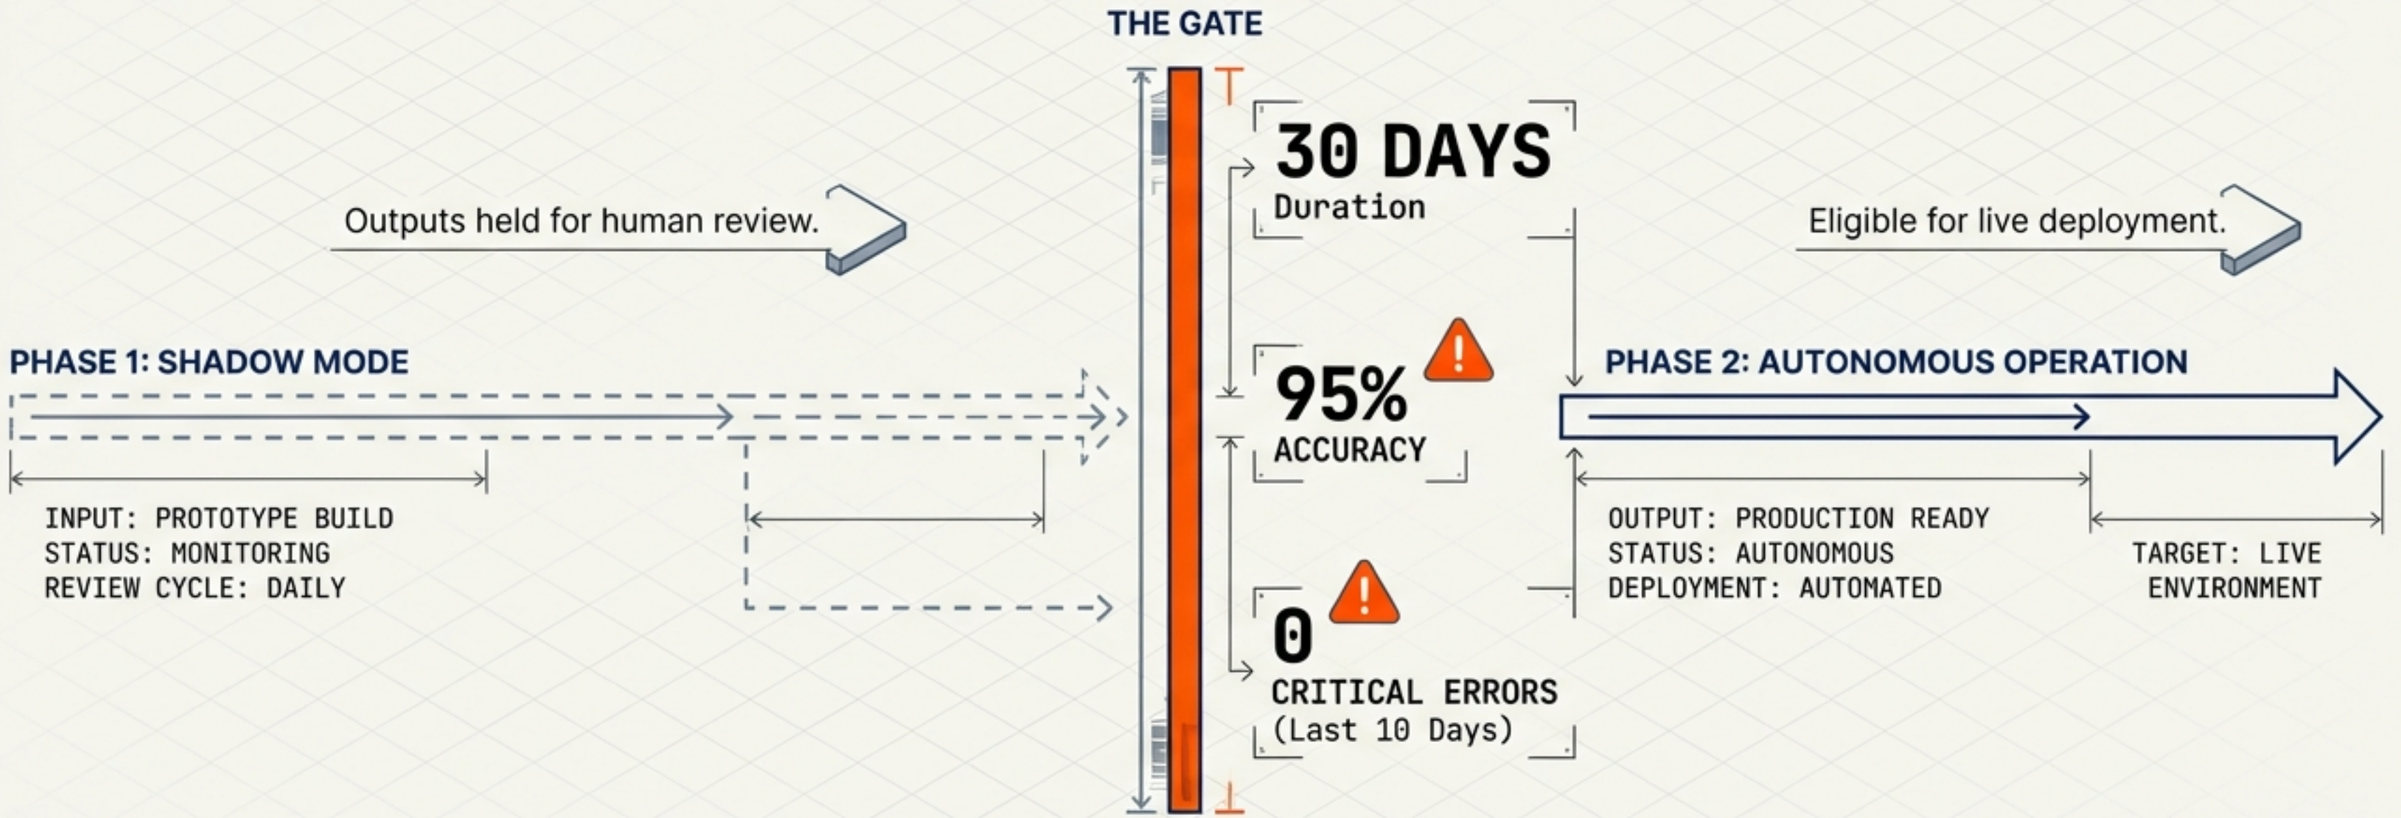

Human-in-the-Loop (HITL) Gates

Autonomy for the safe; Humans for the accountable.

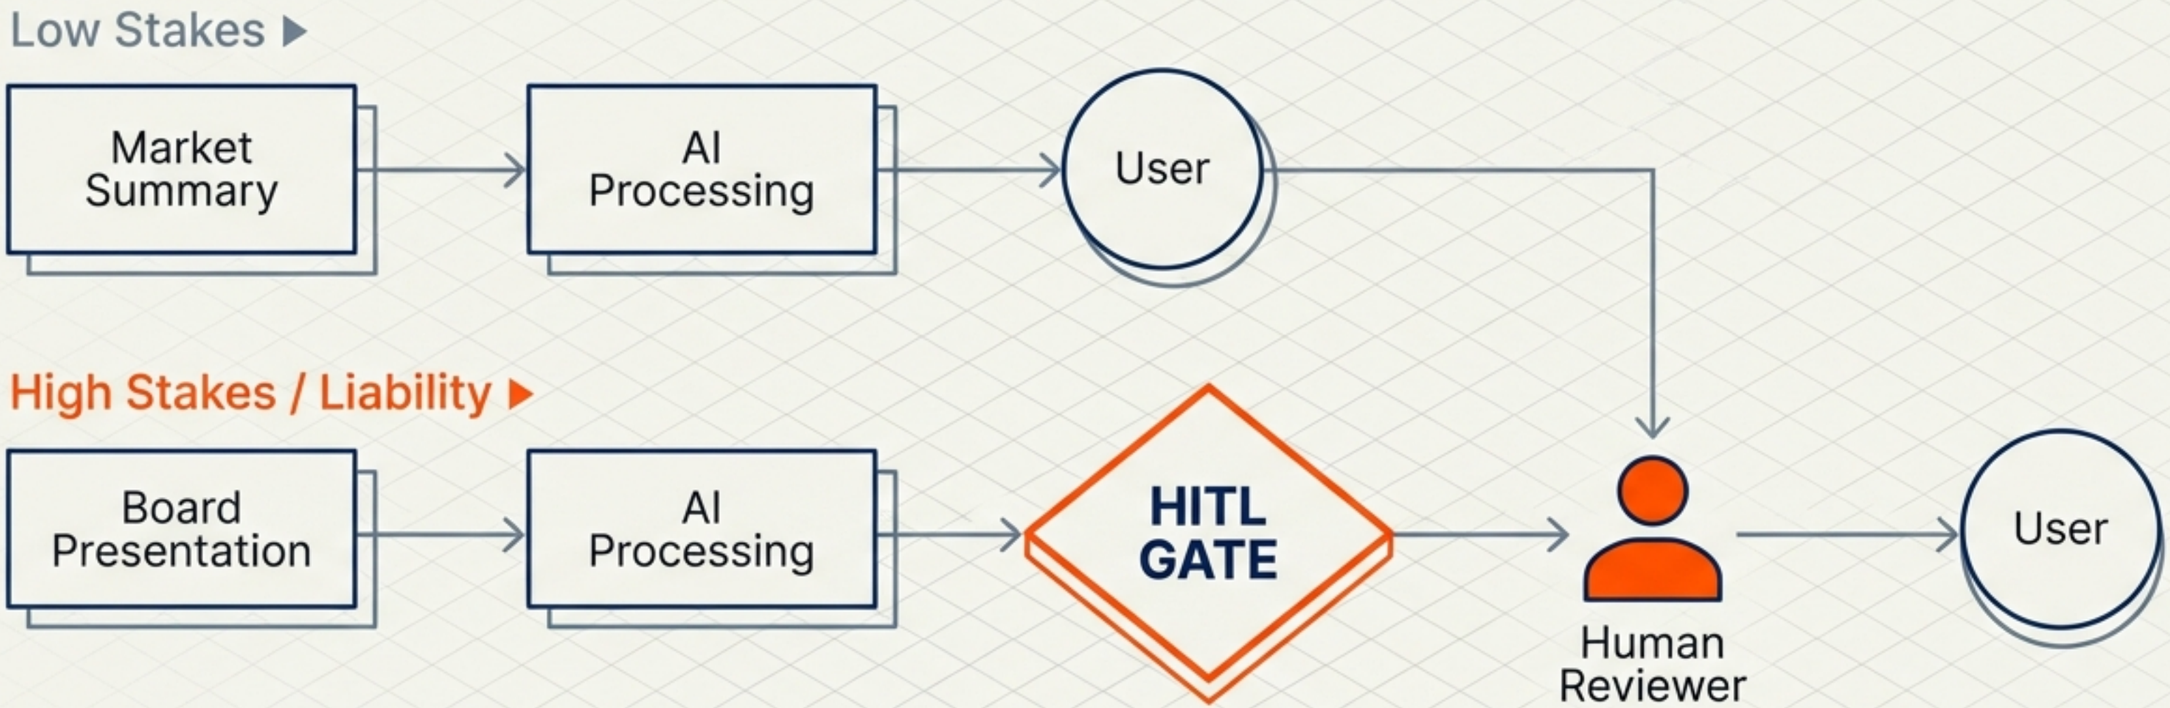

“It’s not about accuracy; it’s about accountability.”

The Division of Responsibility

The Layer Independence Principle.

ROLE	OWNS	FOCUS
Knowledge Worker	`SKILL.md`	Persona, Logic, Scope
IT / Developer	Infrastructure	Connectors, JSONs, APIs
Administrator	Governance	Permissions, Audit, Policy

IT cannot fix your logic. You cannot fix their APIs.

Swiss Engineering Blueprint - Bone White #F5F5F0

The Maintenance Discipline

`SKILL.md` is not write-once.

A plugin that isn't updated is slowly drifting into hallucination.

The Marketplace & Transferability

Scaling expertise beyond the firm.

The Transferability Test:
Is this knowledge valuable without your proprietary context?

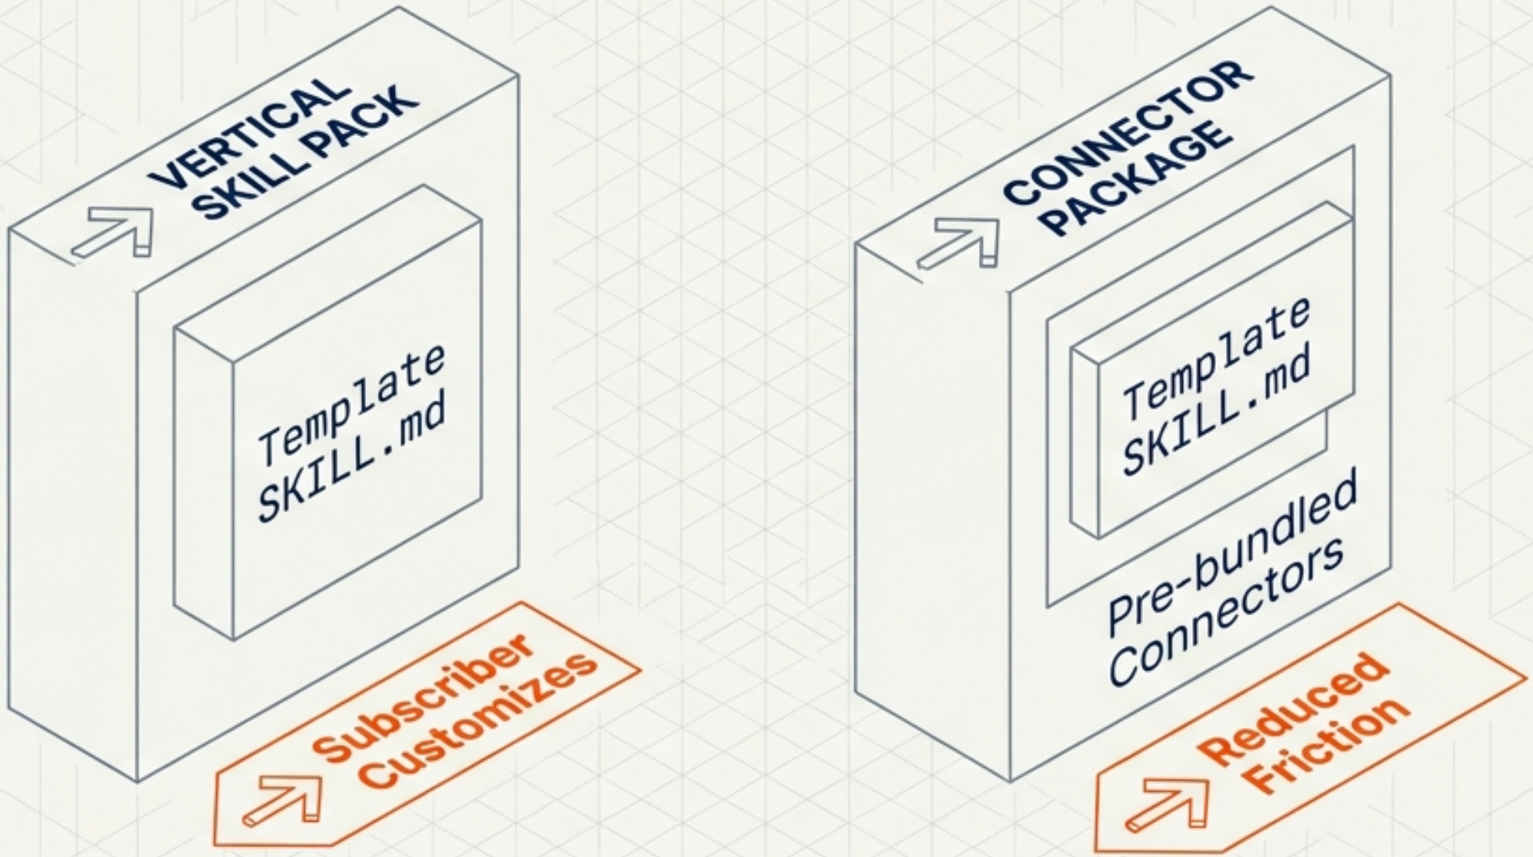

The Deployment Blueprint

Summary

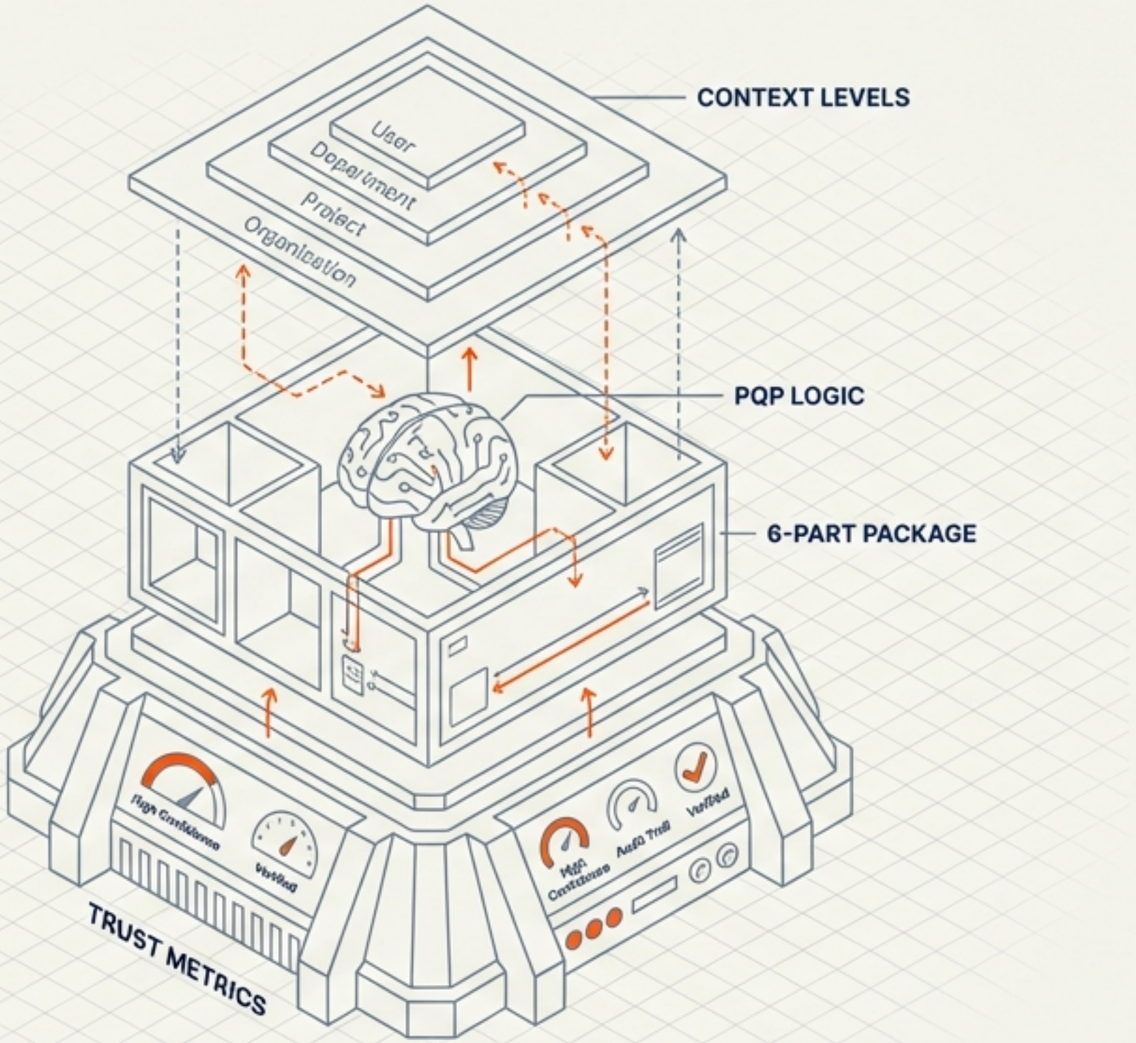

You have the architecture.

Chapter 16 will teach you the Knowledge Extraction Method to fill it.