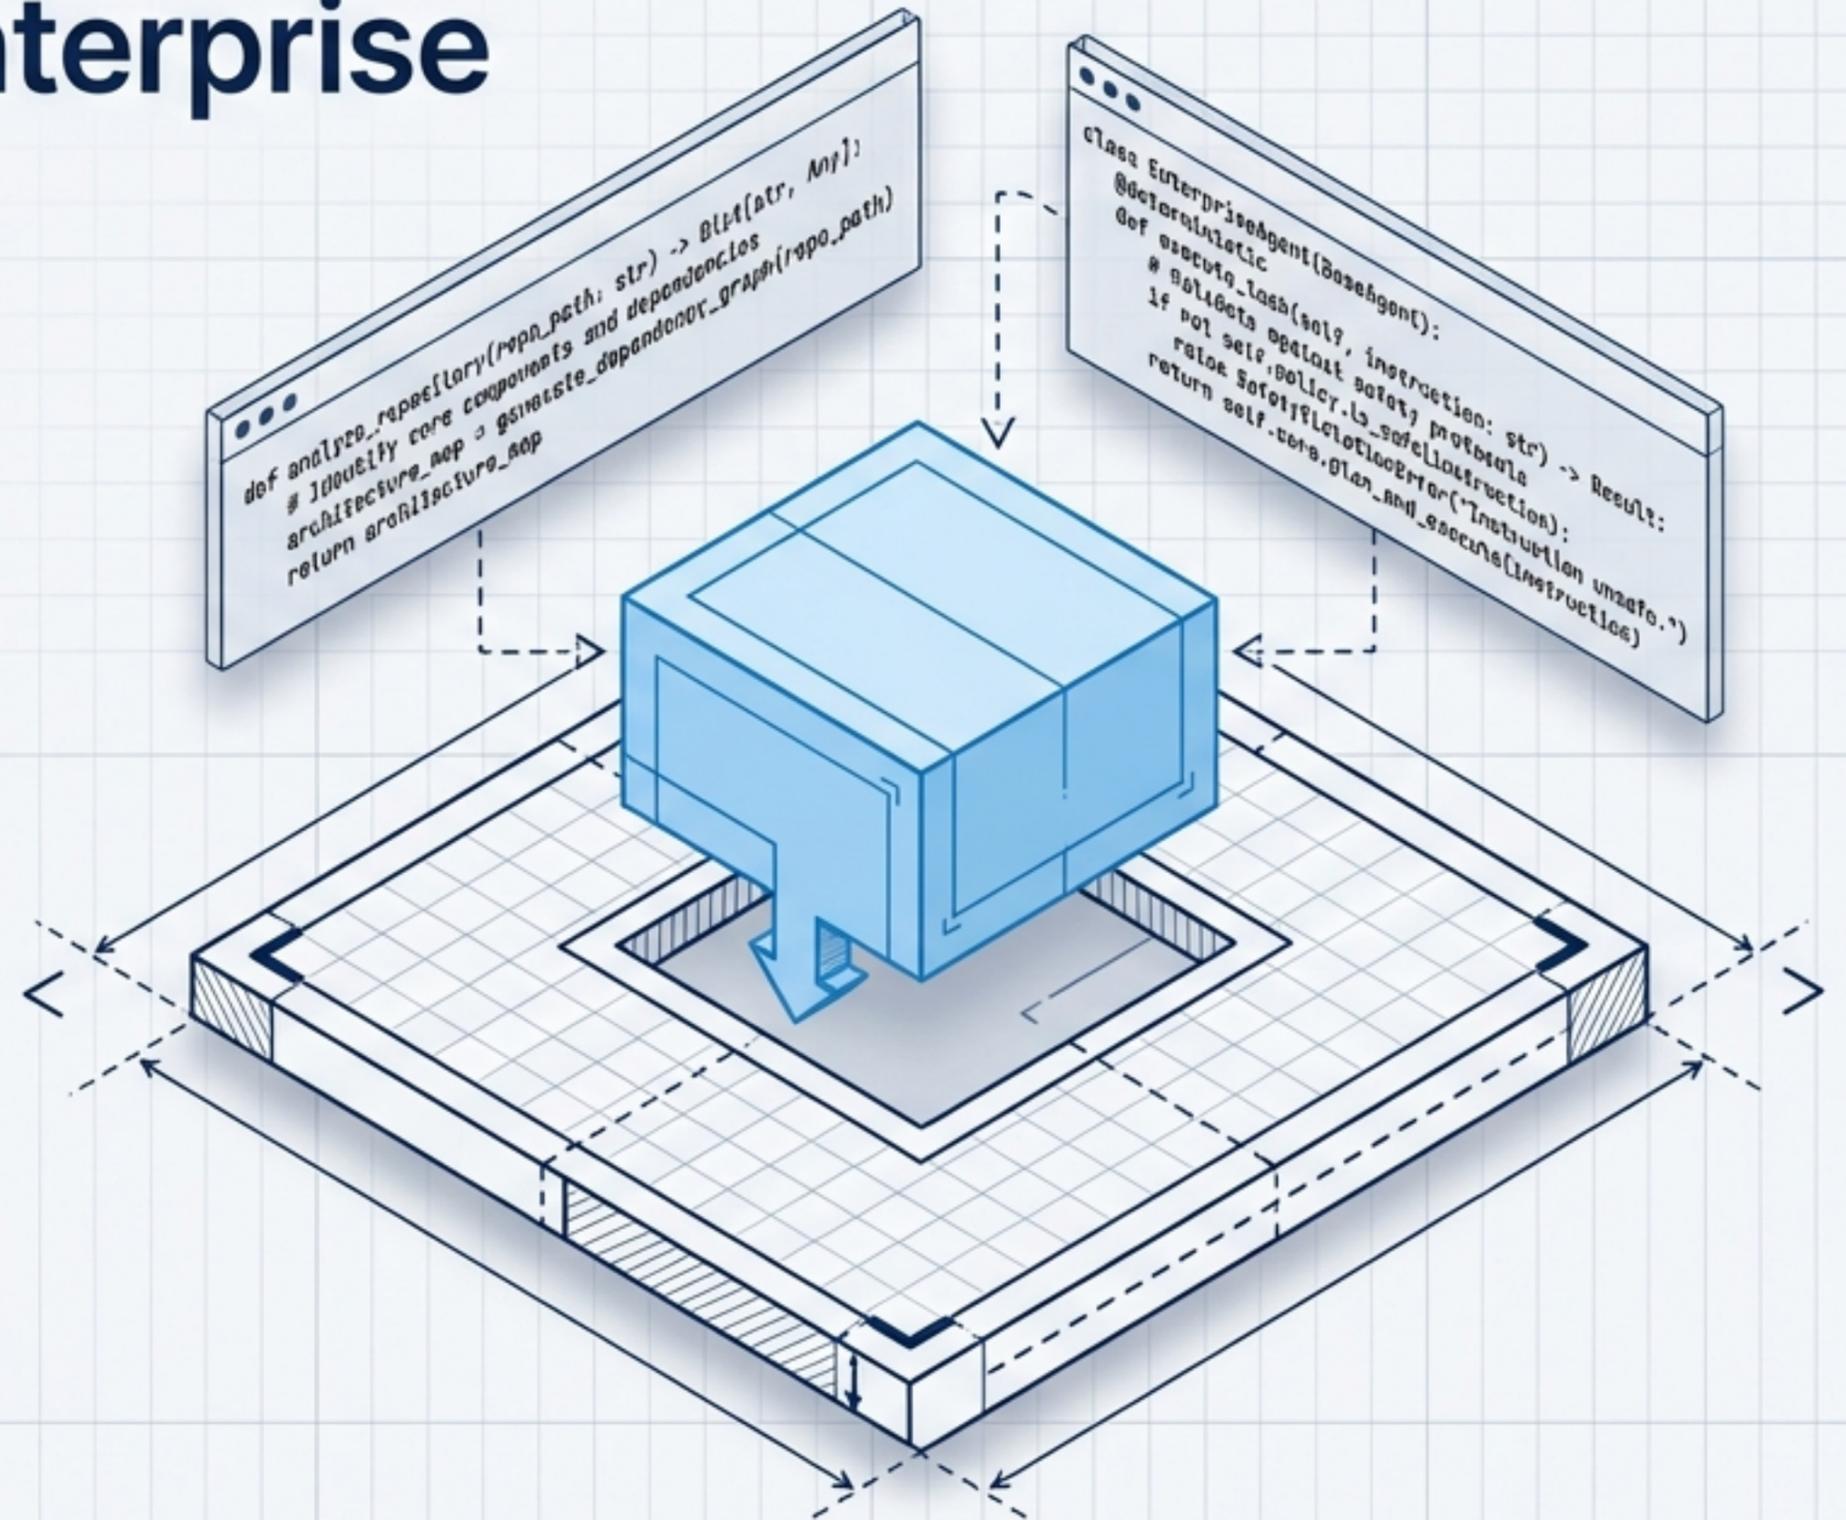
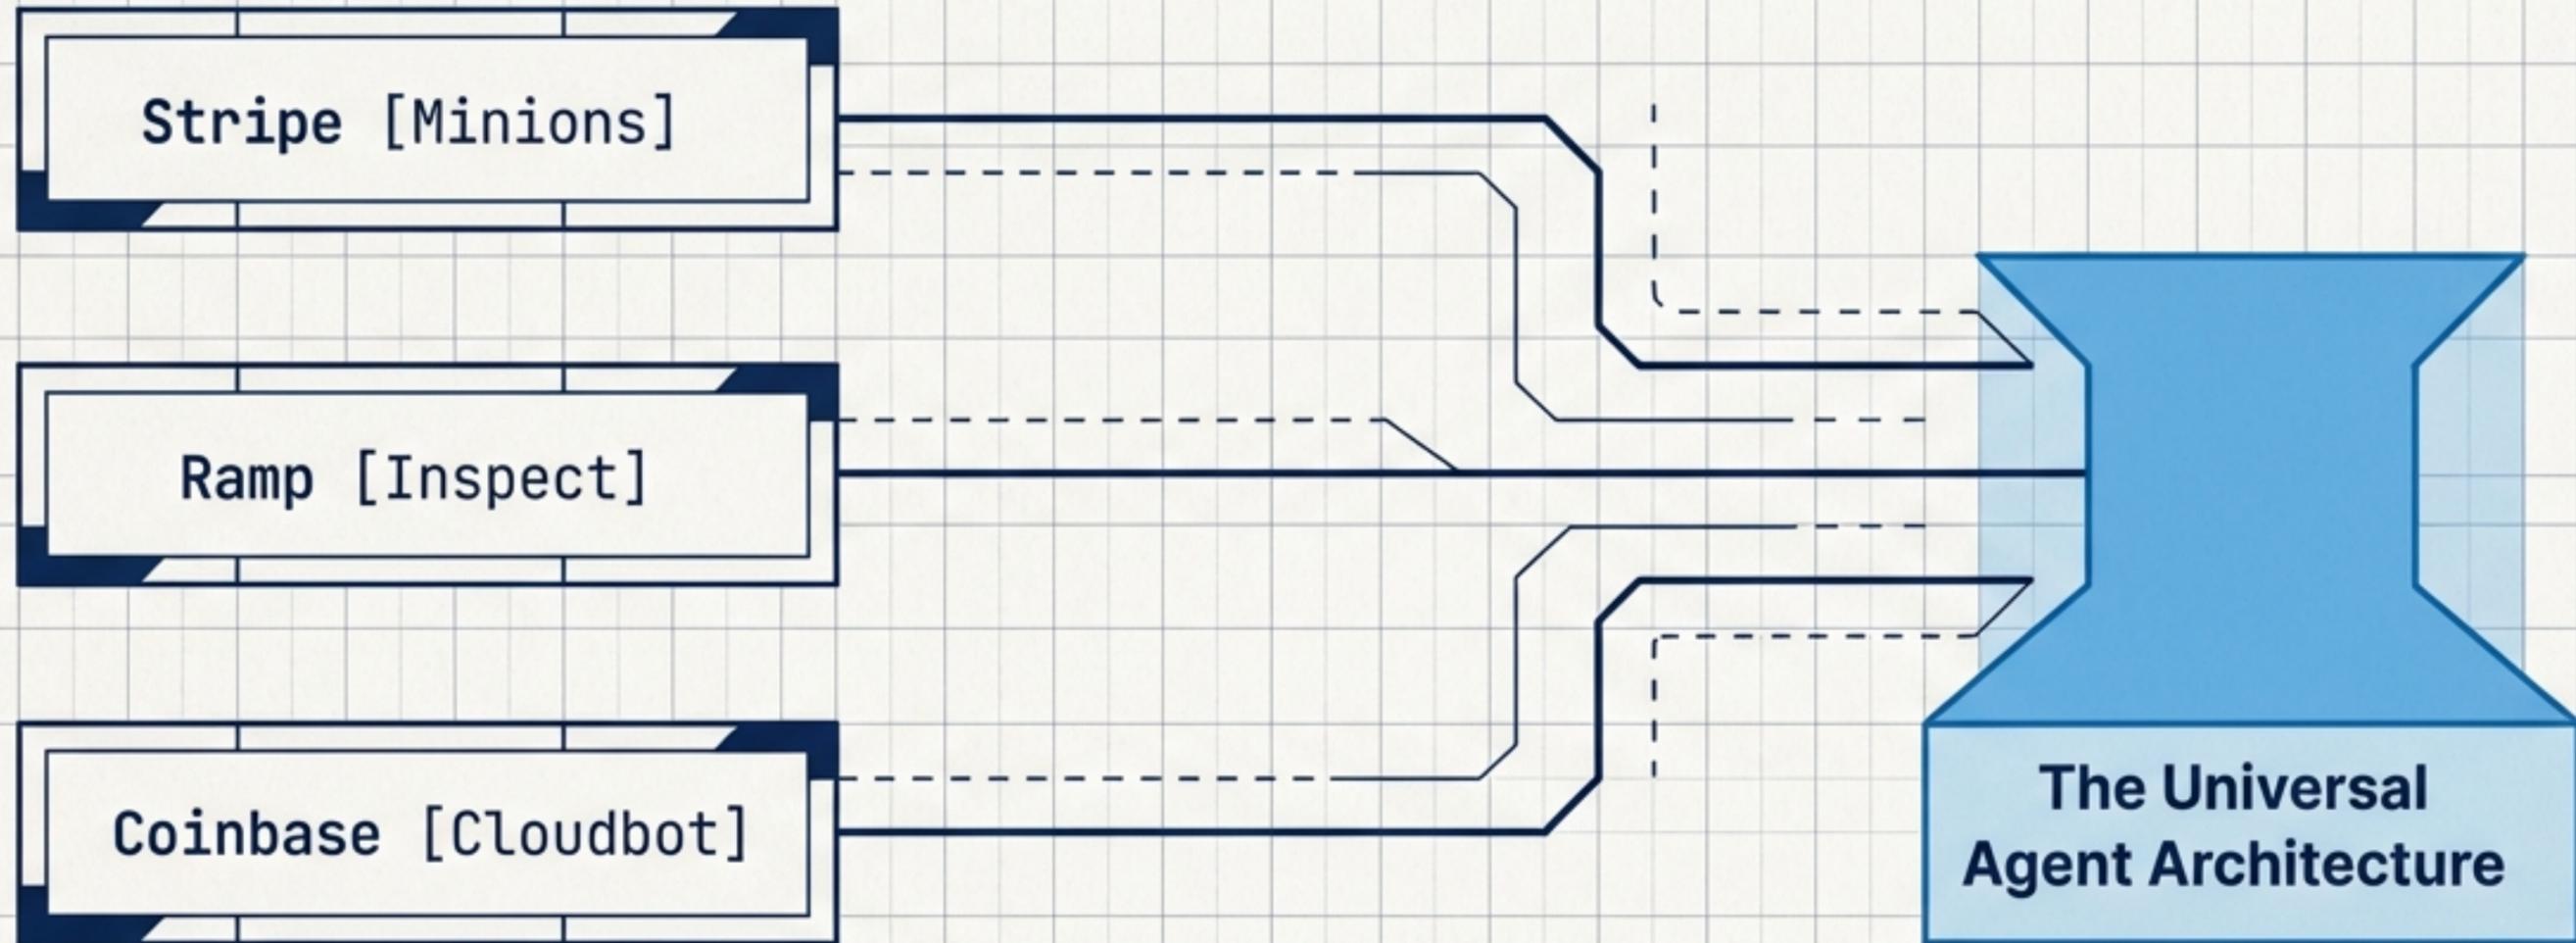


The Blueprint for Enterprise Coding Agents

How the internal architectures of top-tier engineering organizations converged—and the open-source framework that democratizes their DNA.



Independent Teams. Identical Conclusions.



Key Insight: Despite independent development, top-tier engineering orgs have arrived at the exact same architectural decisions. Production reality forced a shared convergence.

The Common DNA of Production Agents



Isolated Environments

Cloud sandboxes with **strict boundaries** and **full, isolated execution rights**.



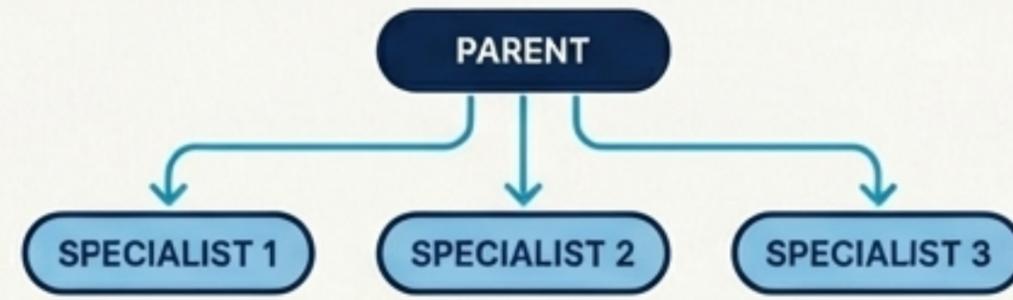
Curated Toolsets

Intent and **quality** over **massive tool bloat**.



Native Workflows

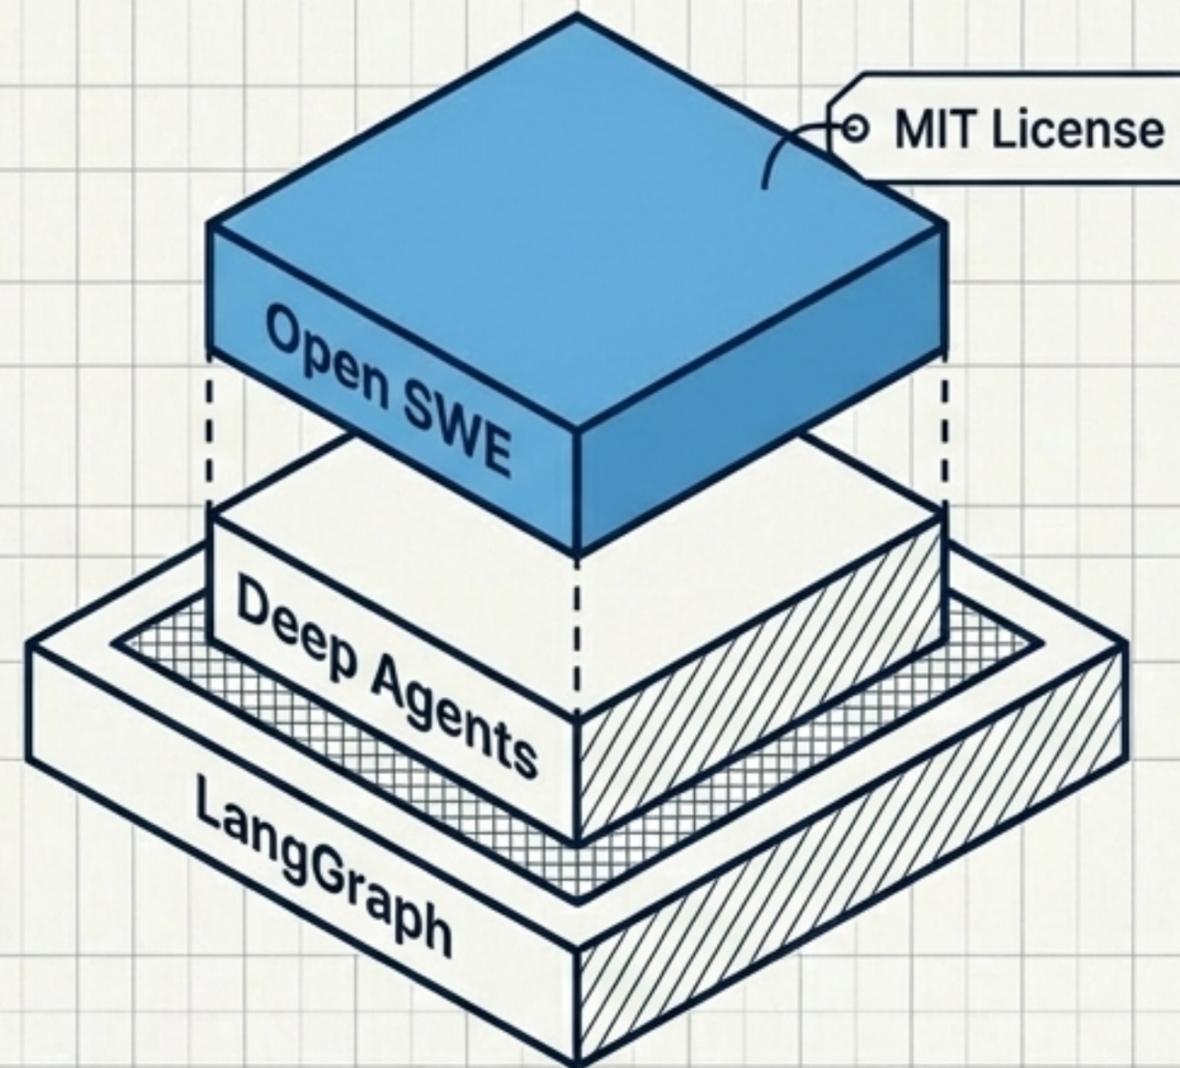
Slack-first, zero context-switching.



Sub-Agent Orchestration

Decomposing complex tasks to **isolated specialists**.

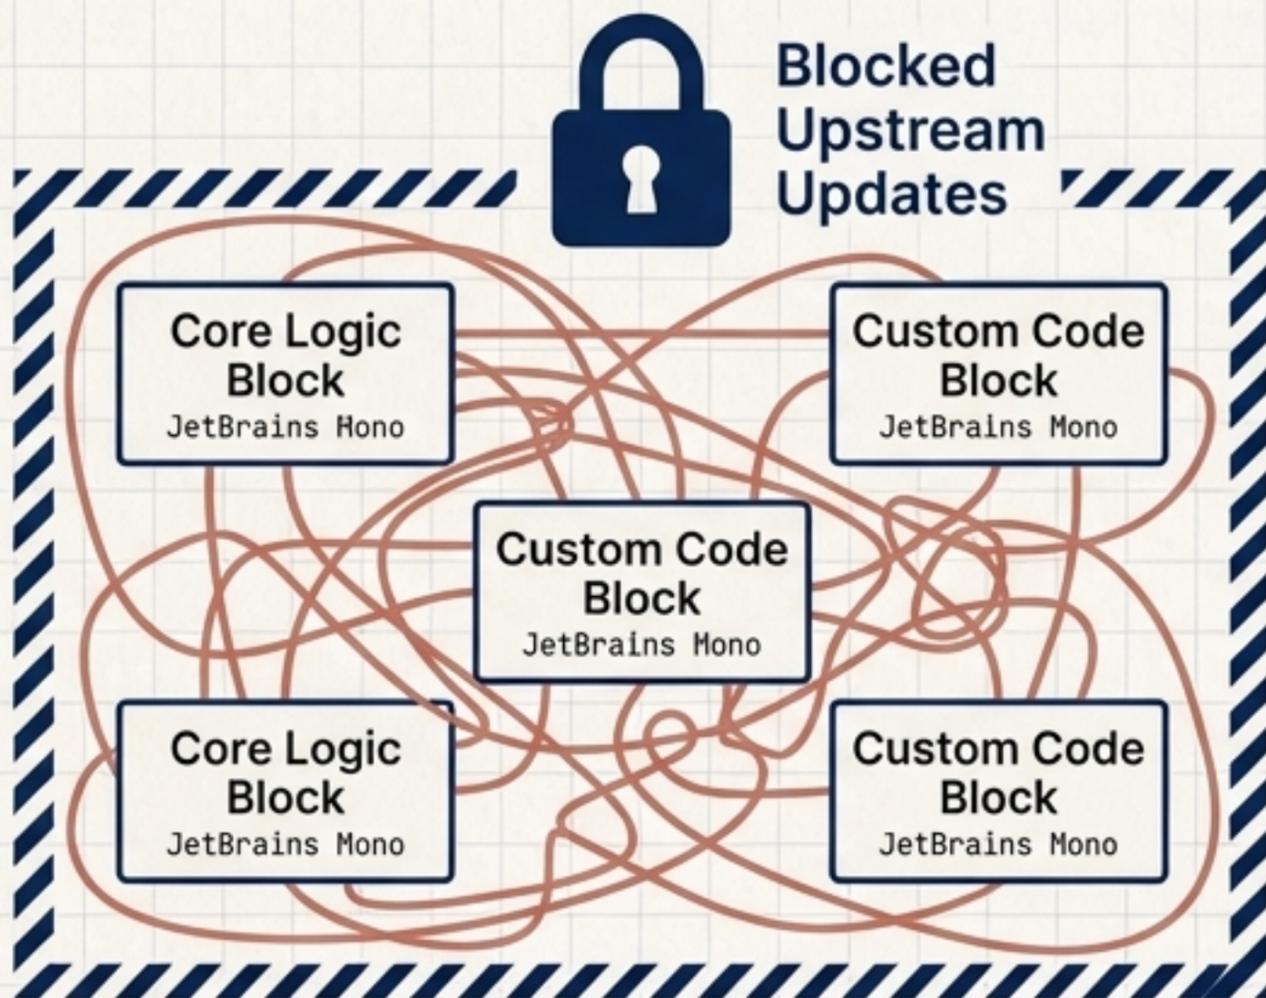
Enter Open SWE: The Open-Source Standard



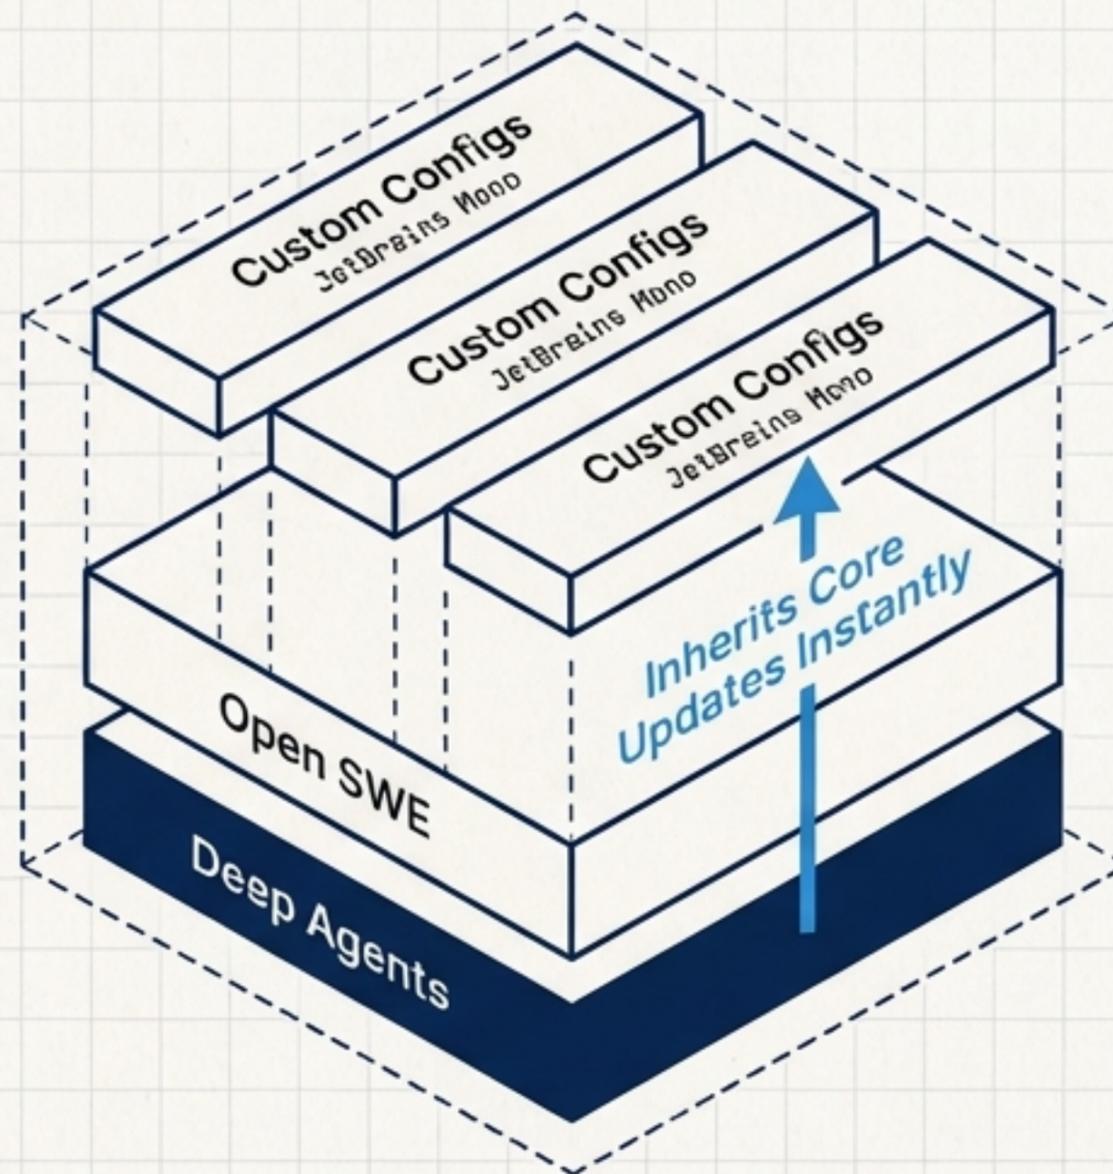
Open SWE provides a production-proven starting point for **internal coding agents**. It is not a rigid product, but an entirely pluggable framework built on composition.

Why Maintain When You Can Compose?

The Forking Trap

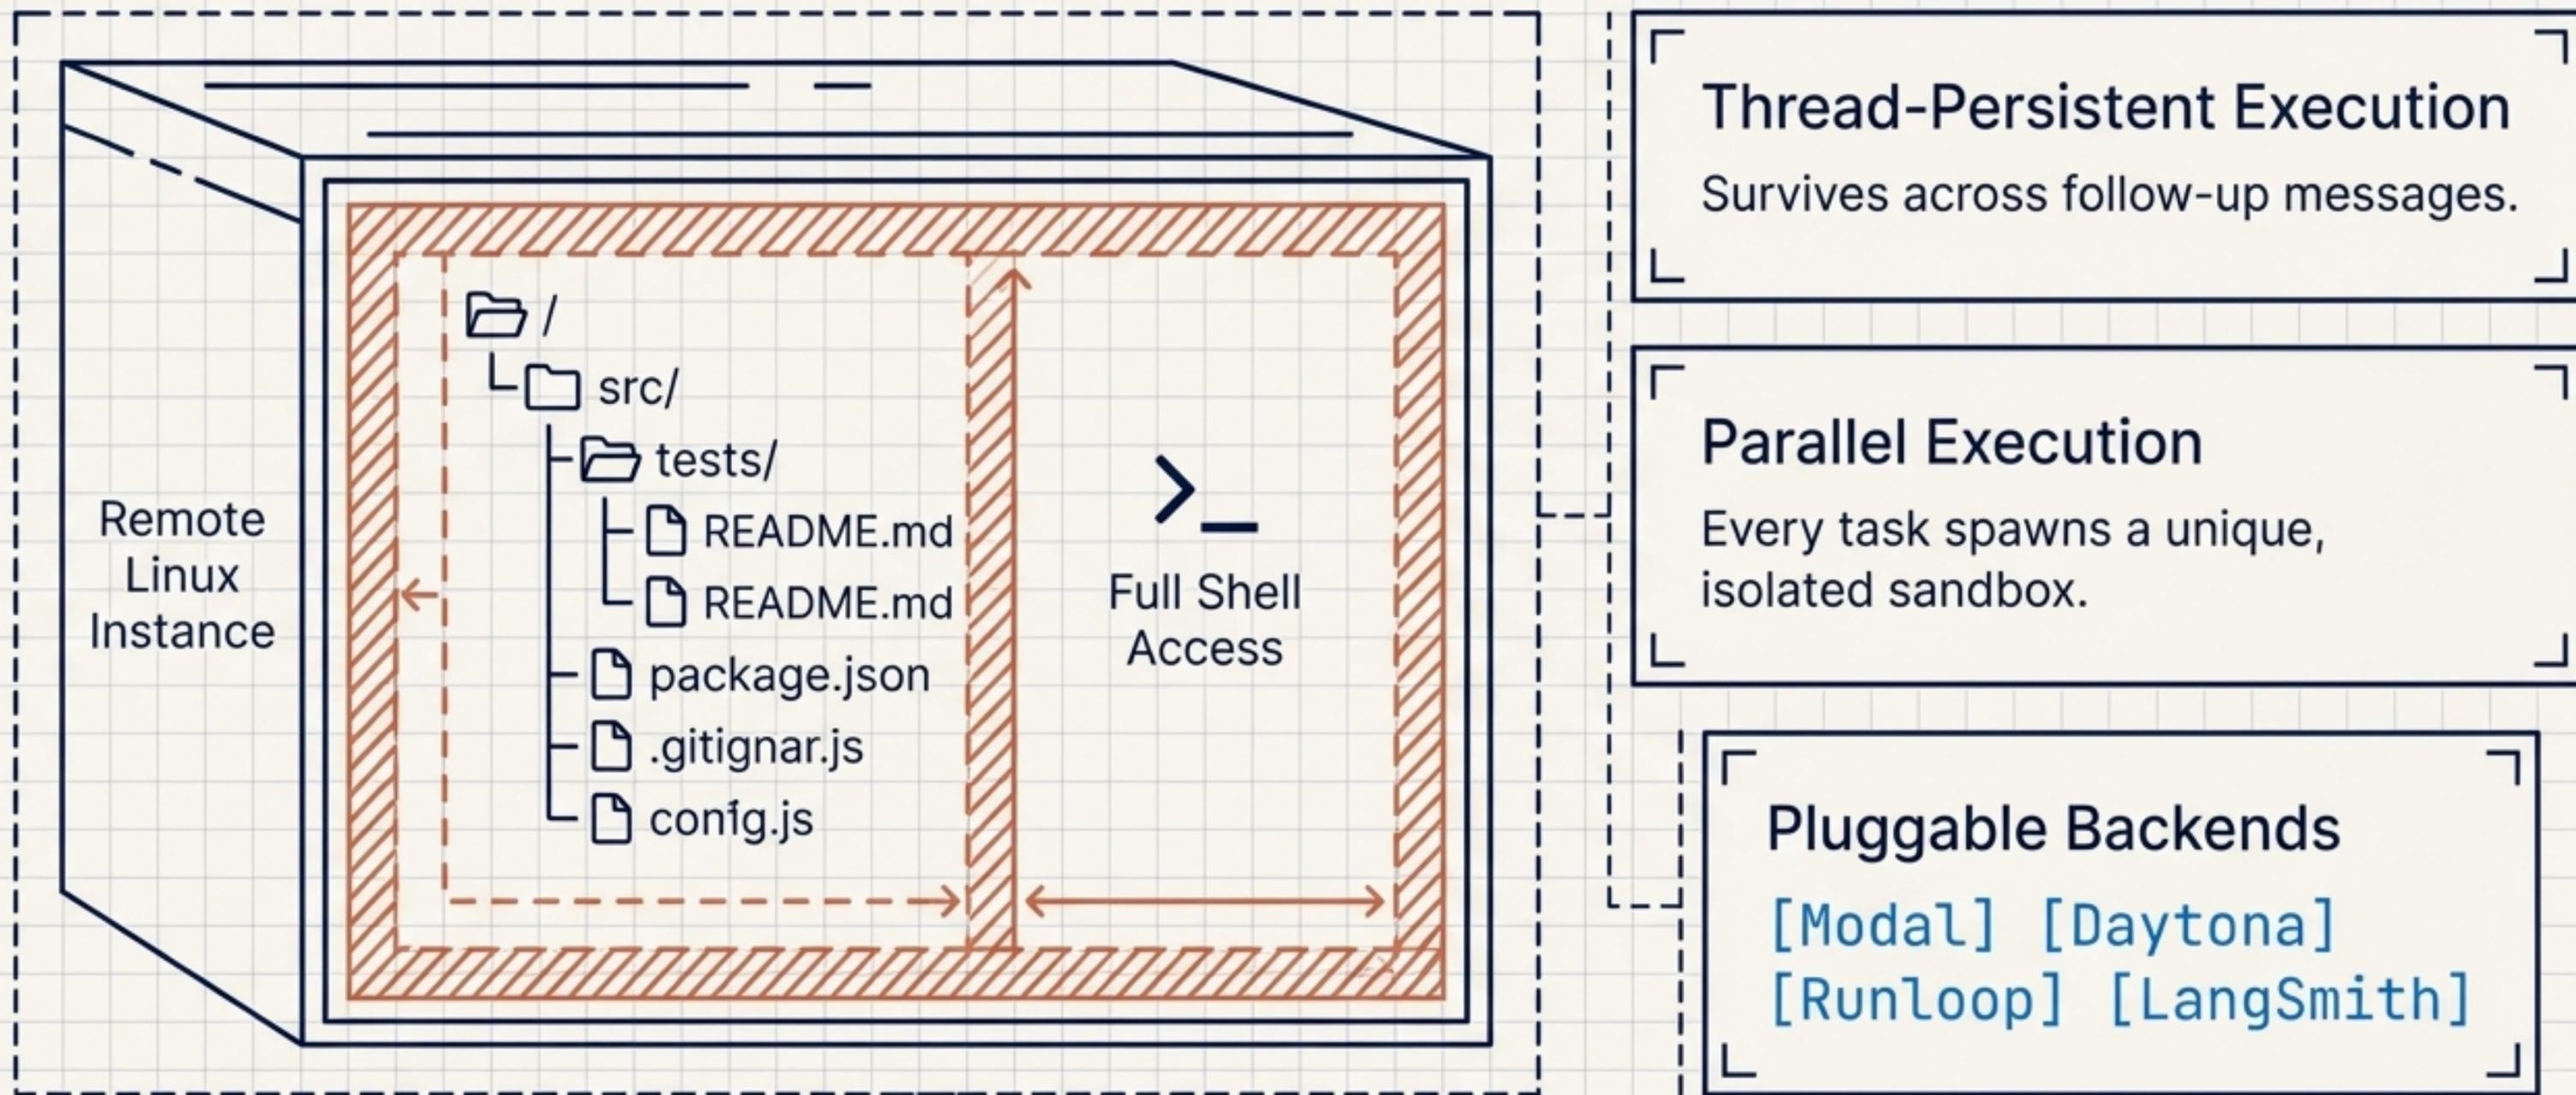


The Composition Model

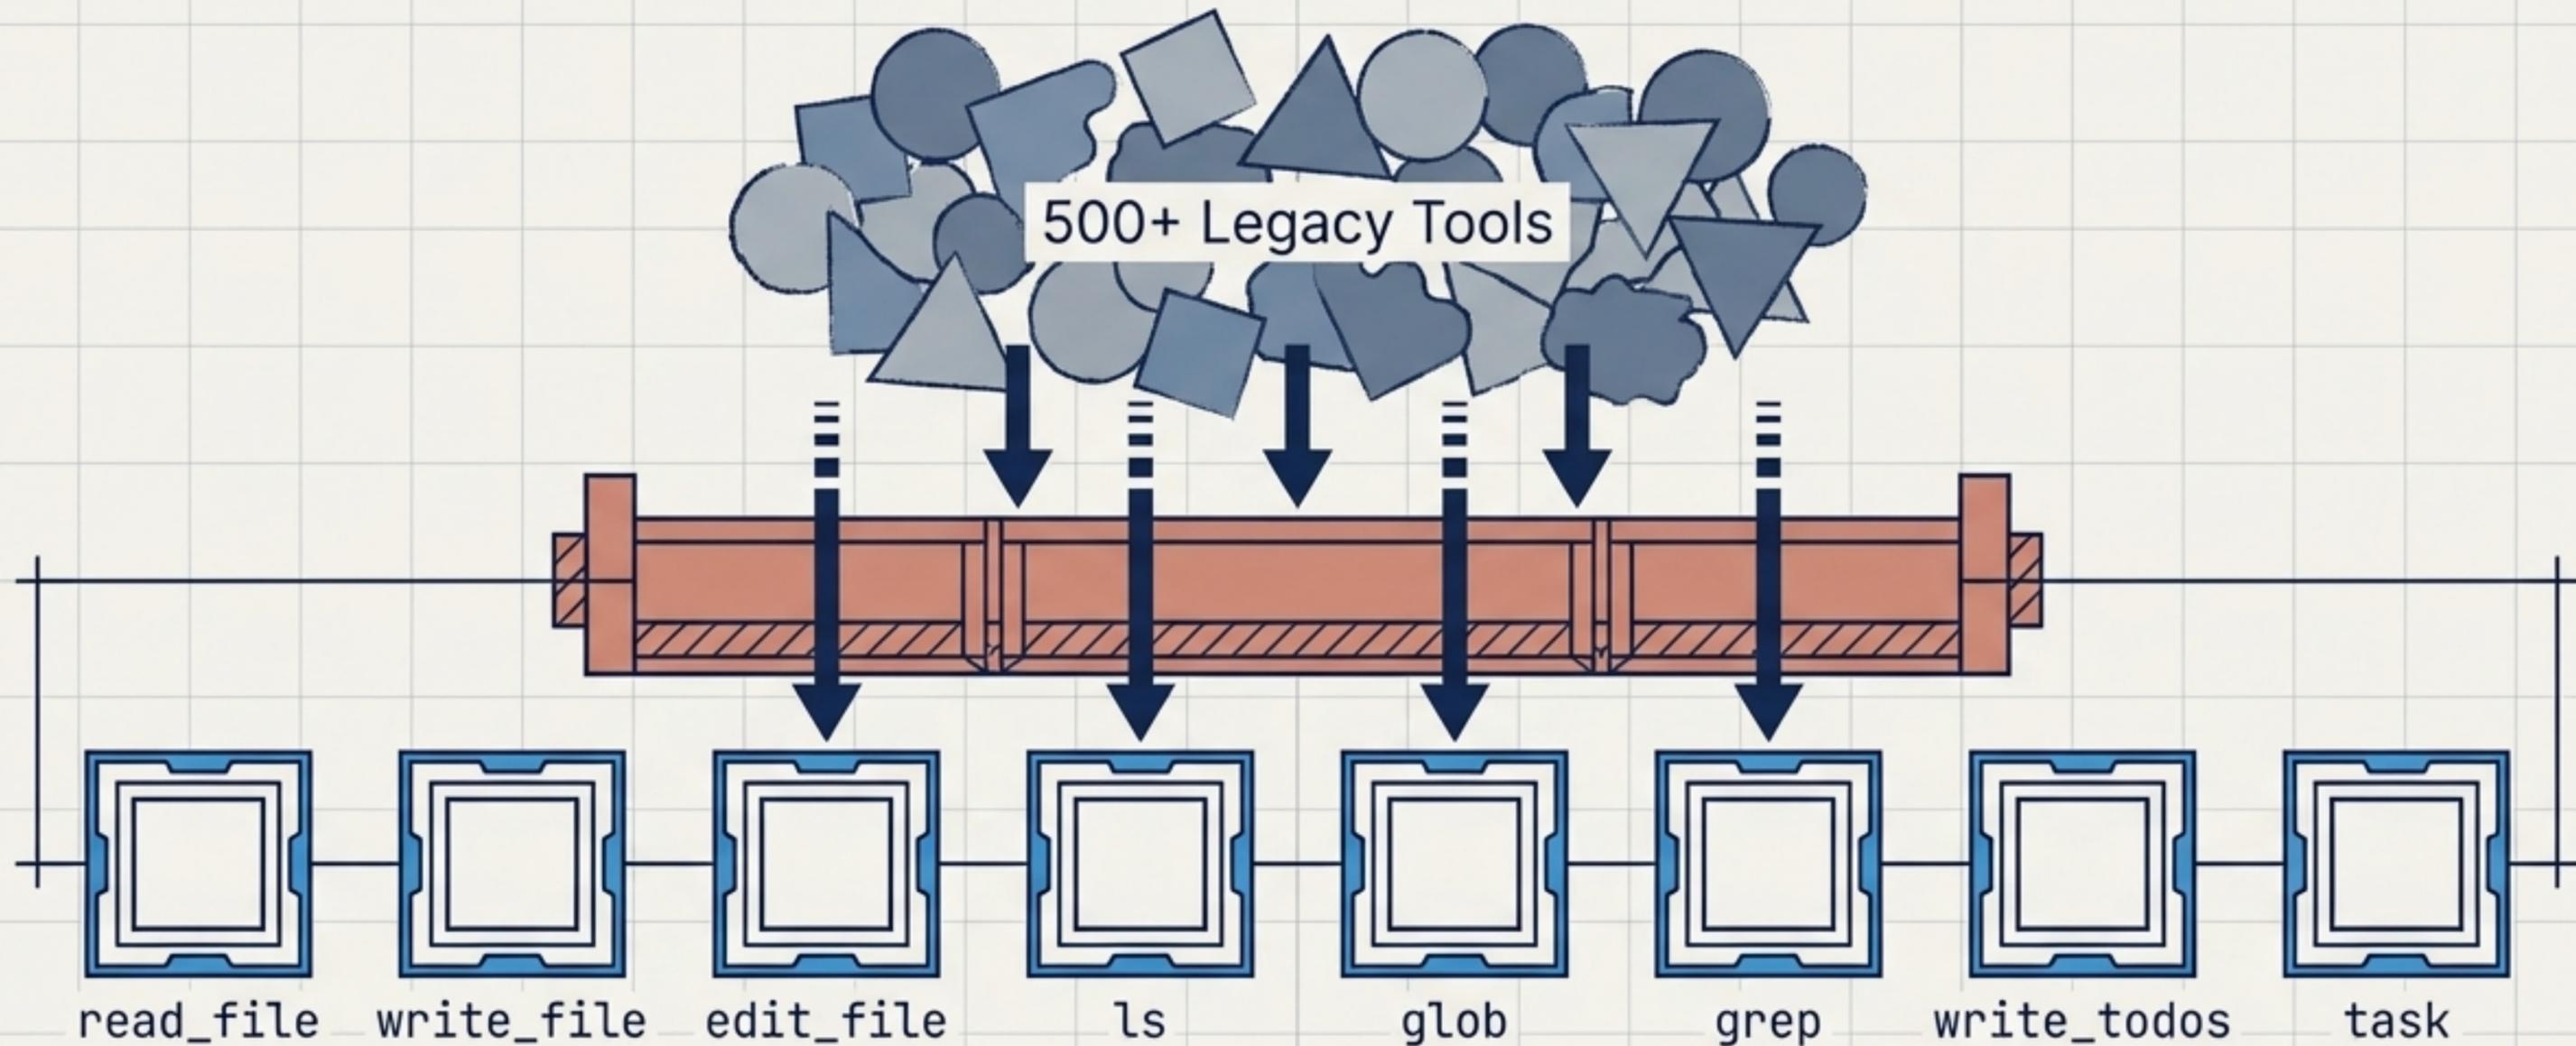


Organization-specific tools and workflows are maintained as configurations, never by modifying core agent logic.

Layer 1: The Persistent Sandbox

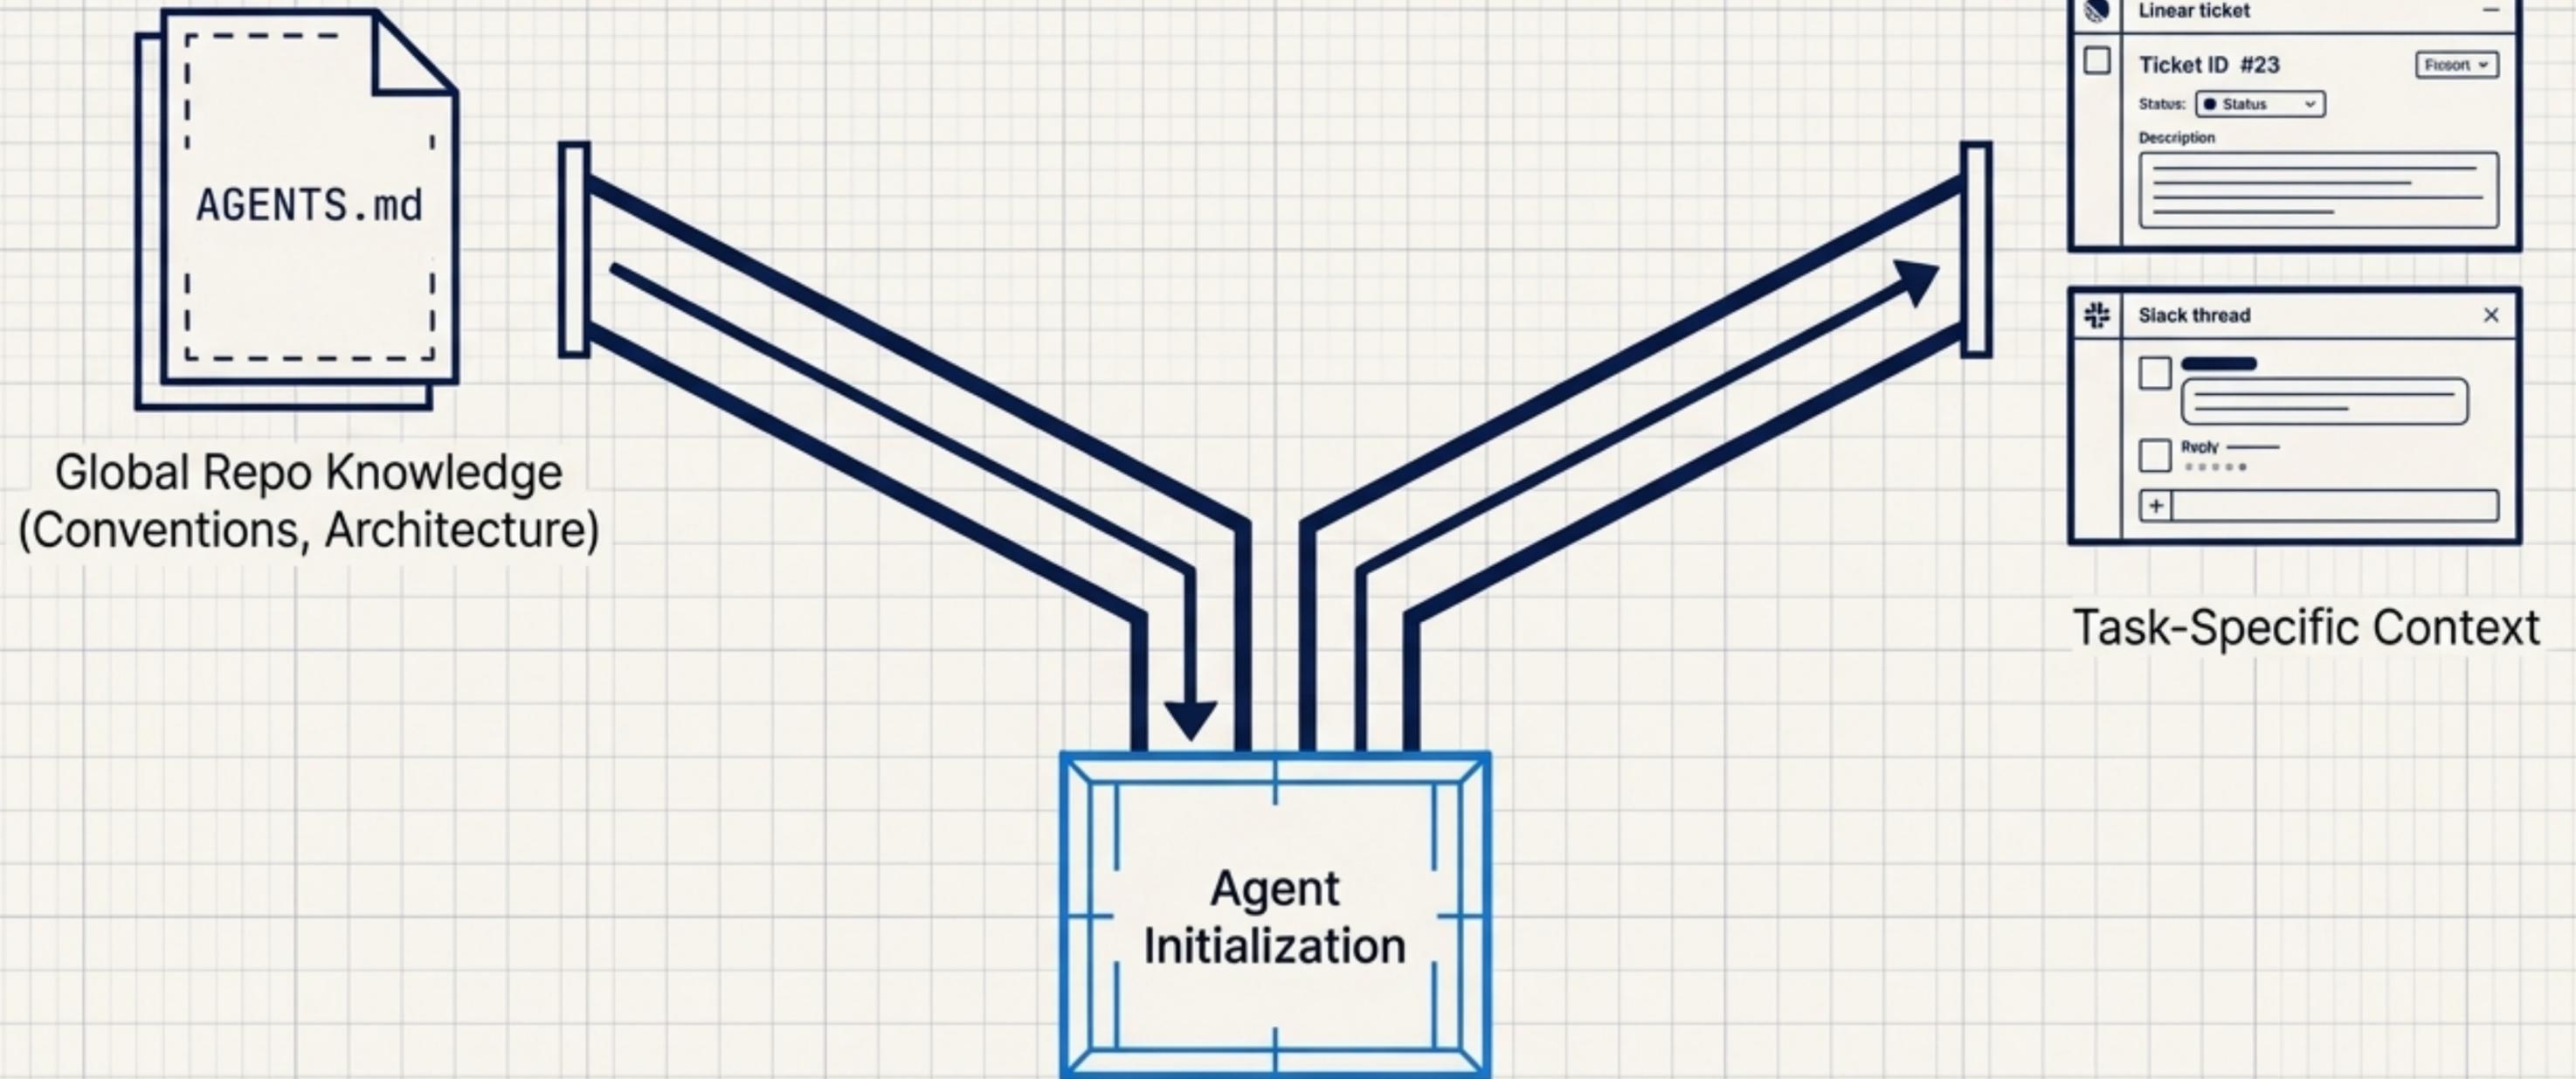


Layer 2: Tool Curation (Precision Over Bloat)



Curated toolsets enable flawless reasoning and easier testing.
Internal APIs and deployment systems snap in only when explicitly required.

Layer 3: Context Engineering via Dual-Intake



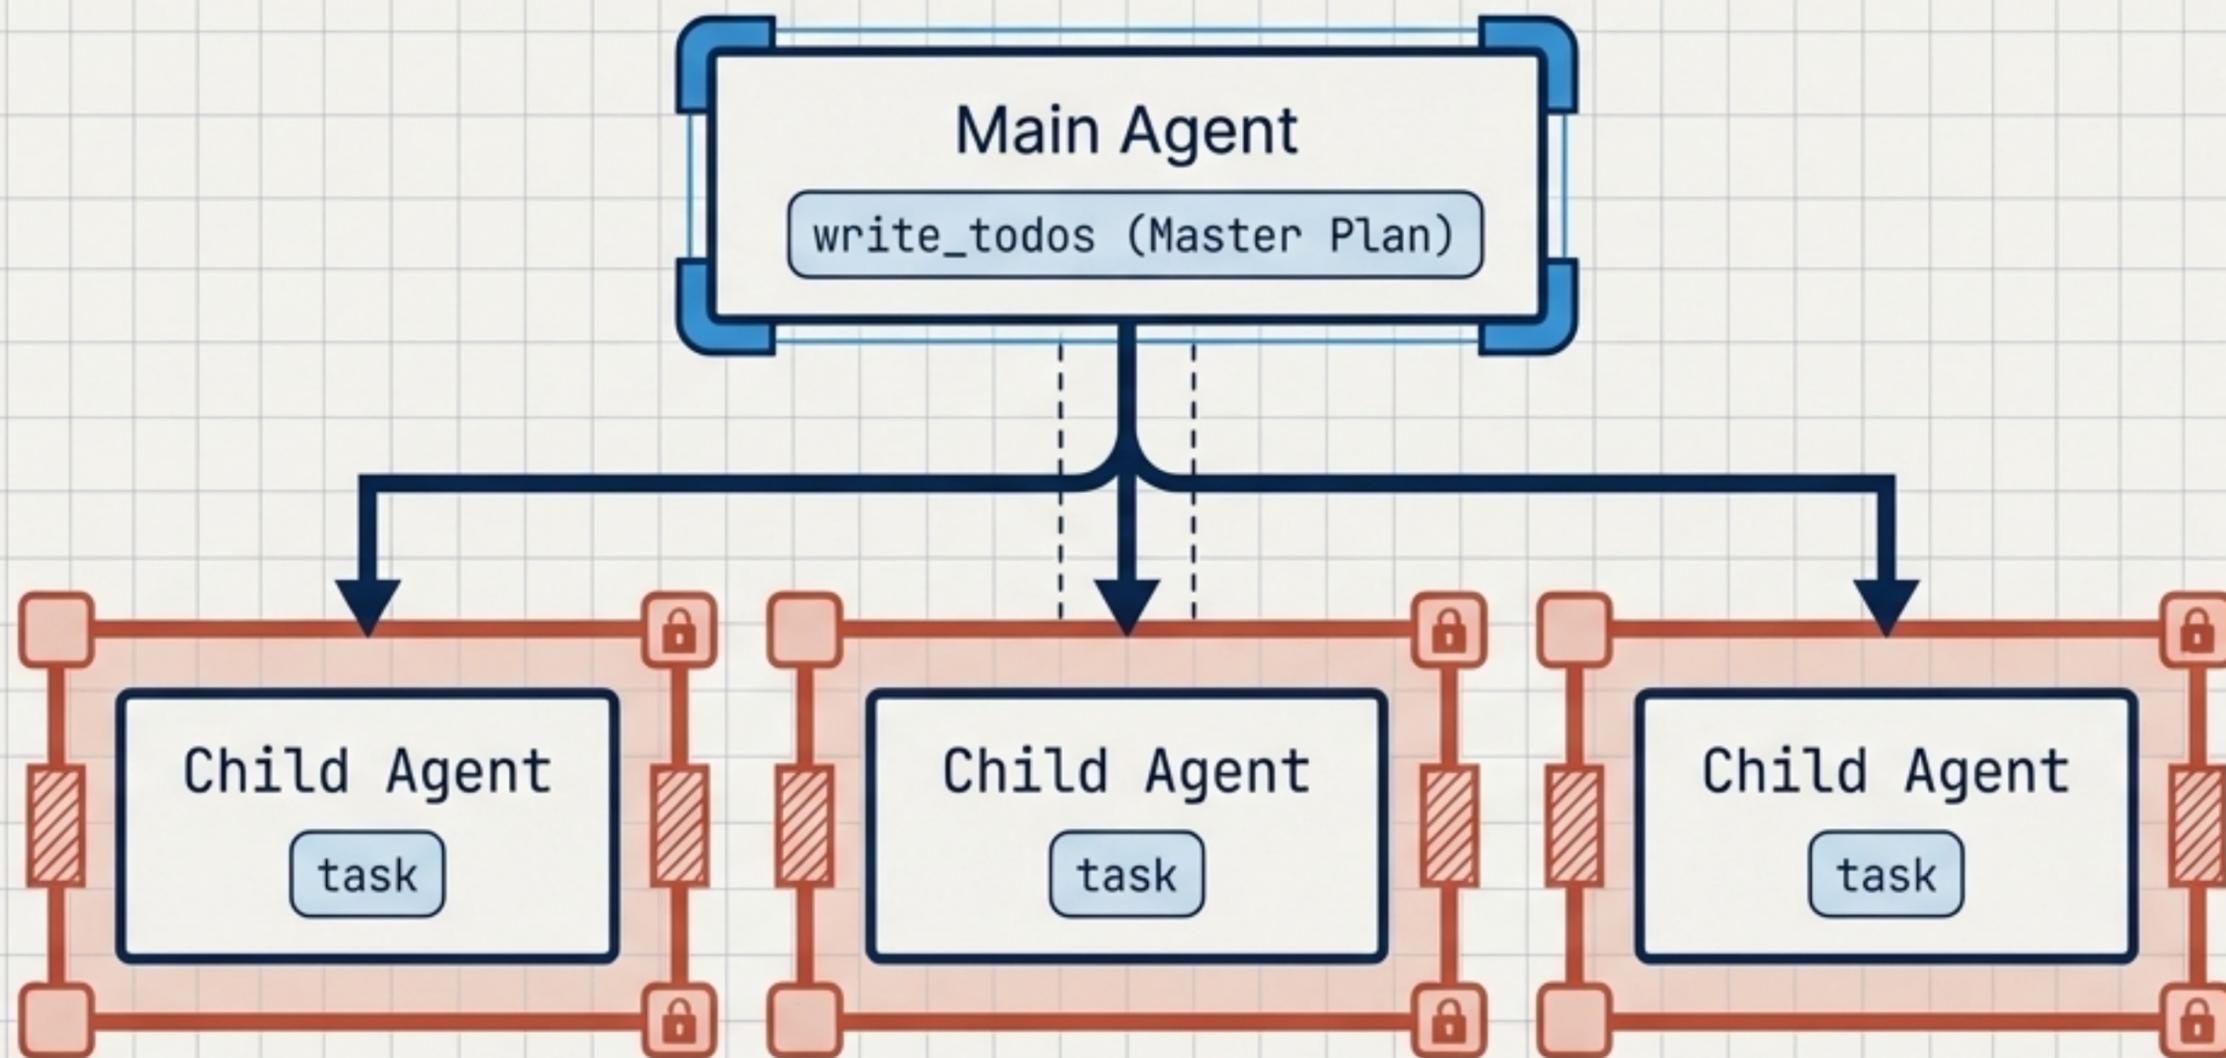
Zero-overhead discovery. Agents are hydrated with full organizational and task context before the first tool is ever called.

Layer 4: Orchestrating Brains and Brawn

Agentic Orchestration (Model-Driven)	Deterministic Orchestration (Middleware)
<p>Focus: Flexible reasoning and adaptive planning.</p> <p>Actions: Spawning isolated sub-agents via the "<code>task</code>" tool, managing complex multi-step plans via "<code>write_todos</code>".</p>	<p>Focus: Rigid safety nets around the agent loop.</p> <p>Actions: "<code>check_message_queue_before_model</code>" (injects mid-task replies), "<code>ToolErrorMiddleware</code>" (graceful failures).</p>

Balancing the brilliant flexibility of LLMs with the cold, hard reliability of deterministic middleware hooks.

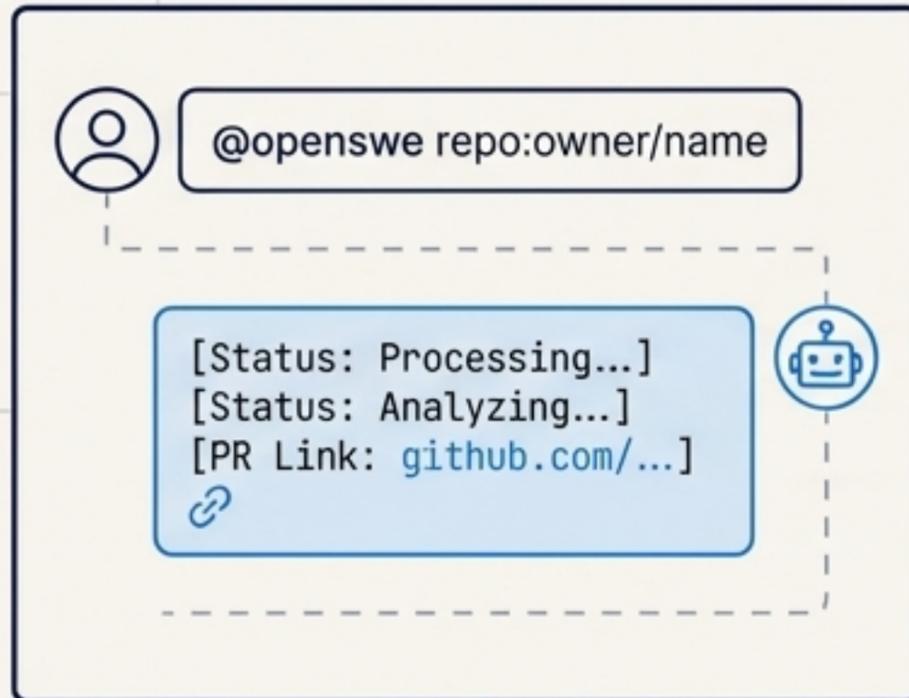
Preventing Context Collapse via Sub-Agent Isolation



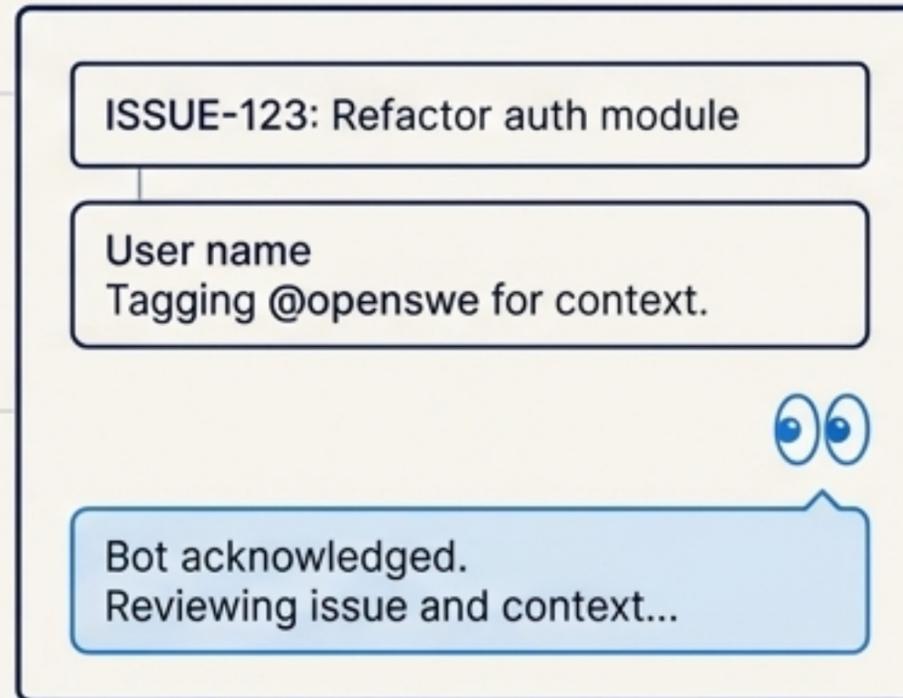
Long-running tasks are offloaded to files. Context windows stay pristine. Reasoning stays sharp without cross-pollution.

Layer 5: Native Invocation (Where Developers Live)

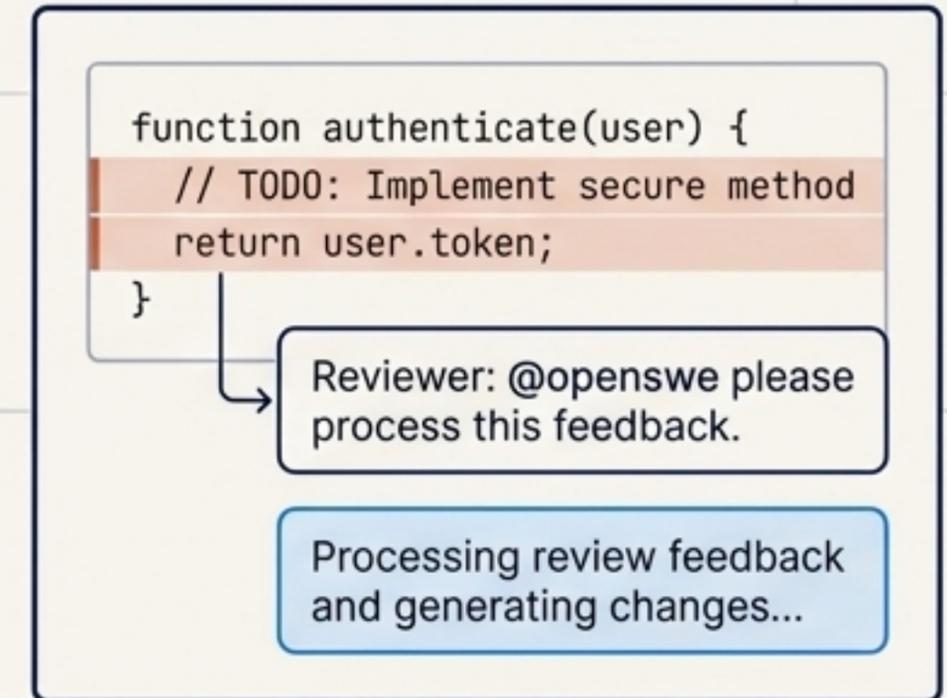
Slack Integration



Linear Integration



GitHub PRs



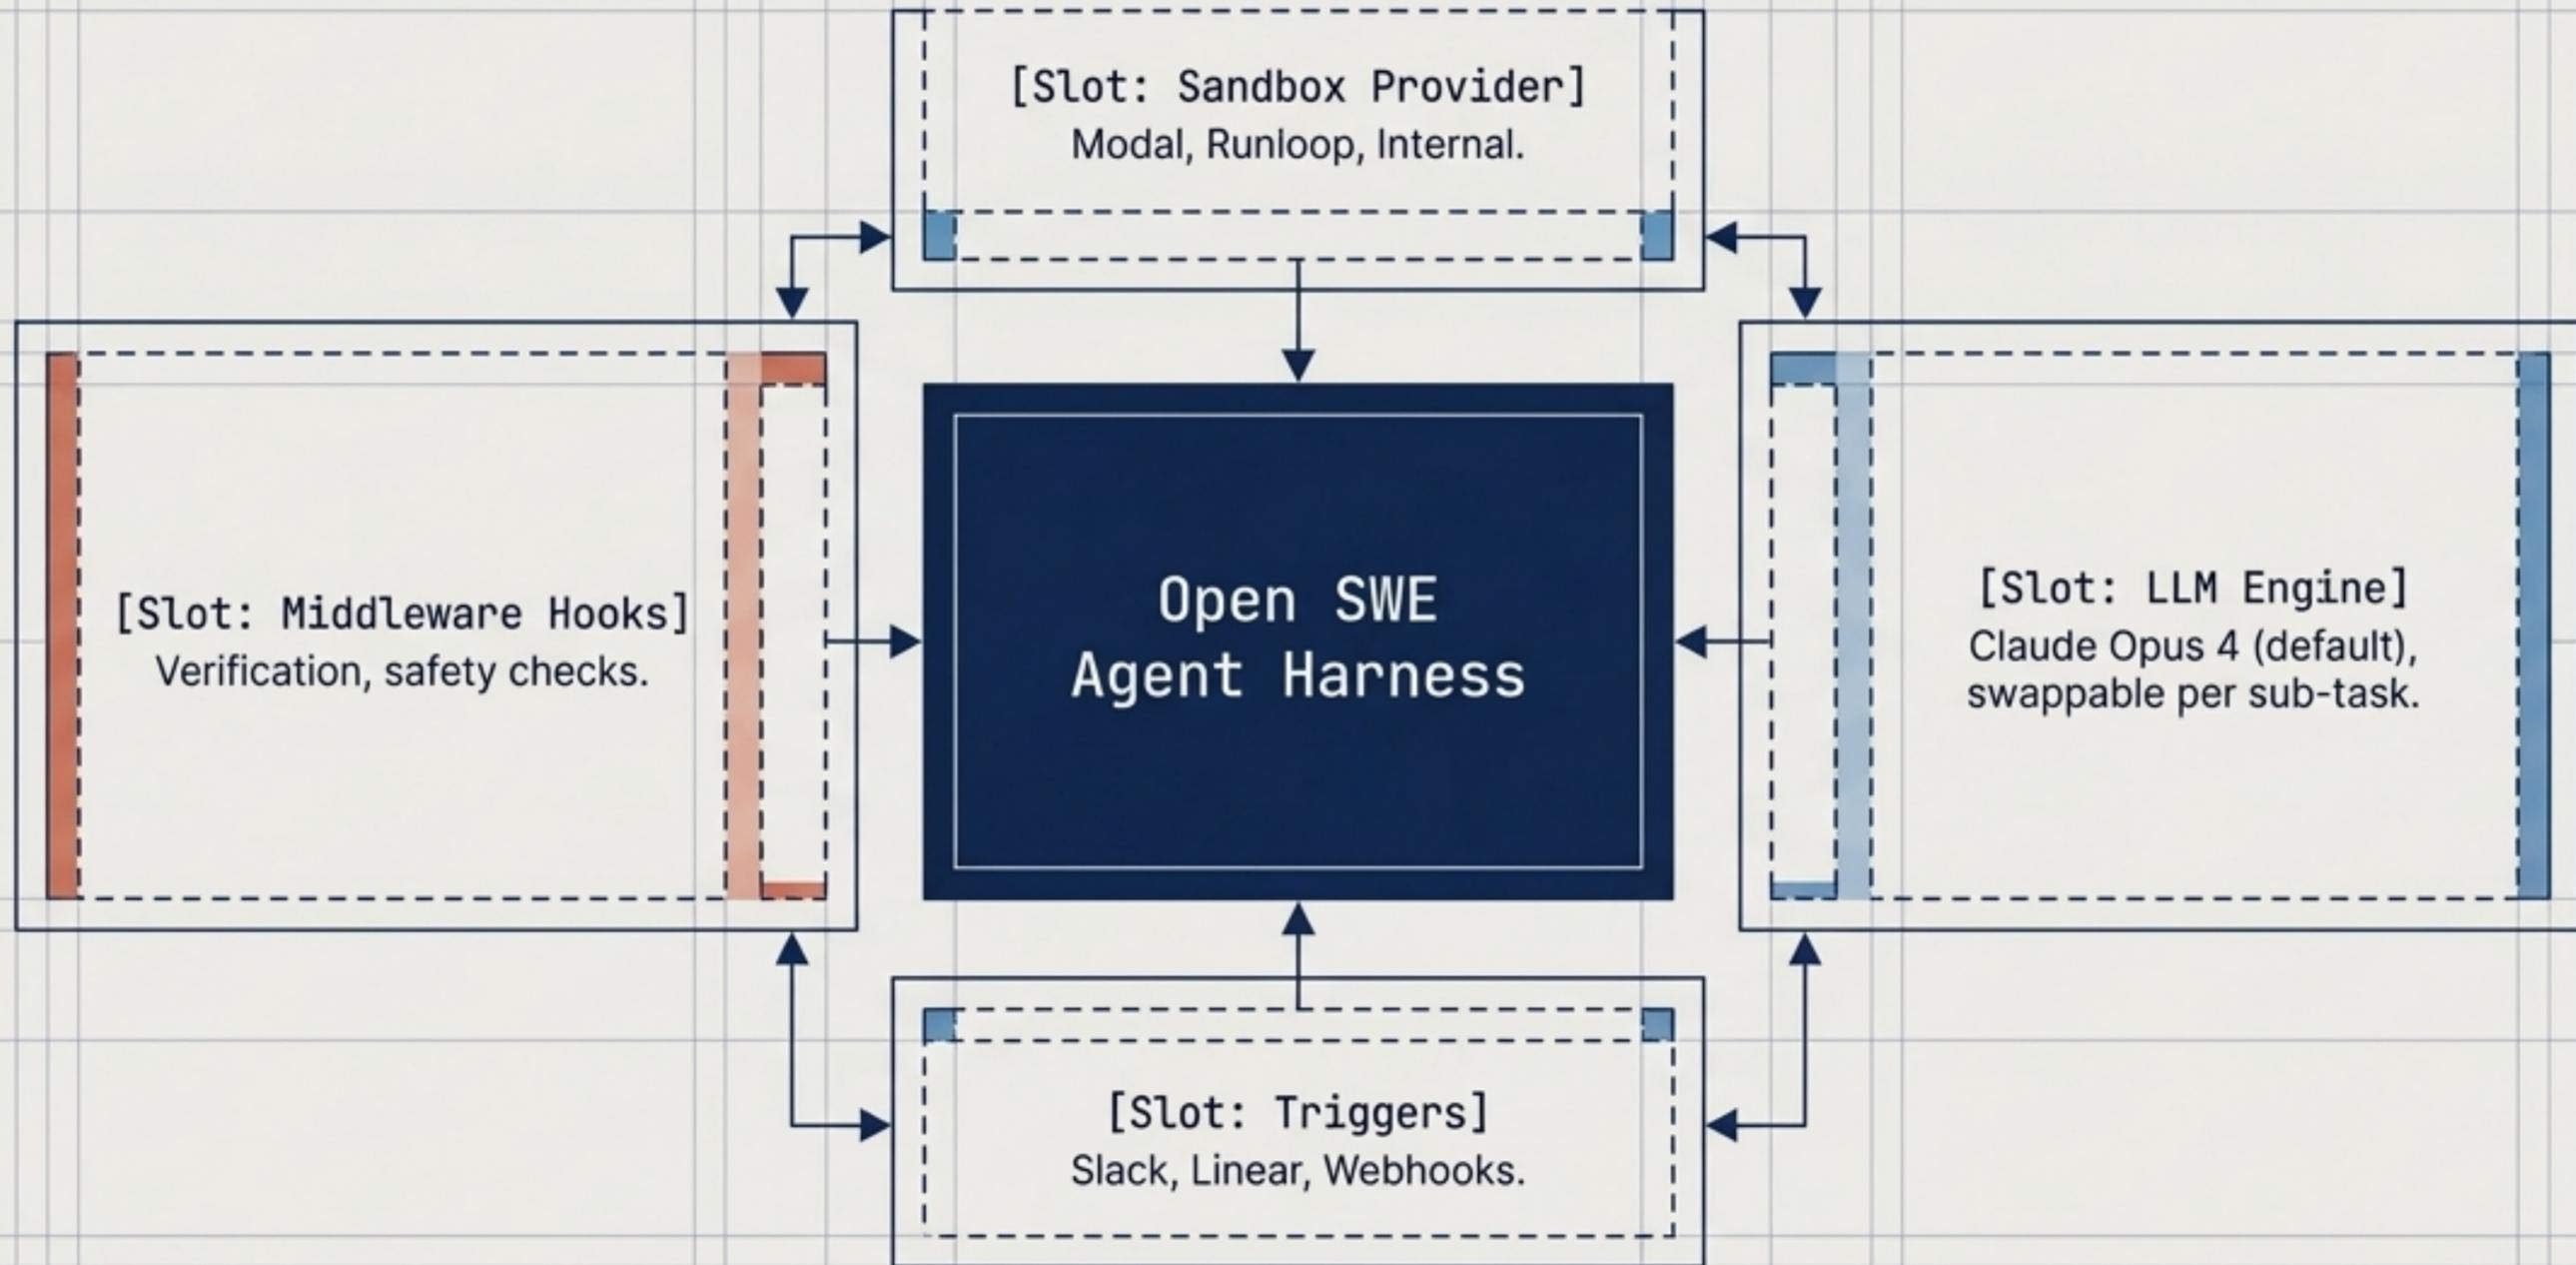
Deterministic Thread IDs ensure follow-up messages automatically route to the exact same running agent. No new dashboards required.

Layers 6 & 7: Verification and Deterministic Fallbacks



Expandable with CI checks and visual validation gates via custom middleware hooks.

Synthesis: The Pluggable Enterprise Framework



Open SWE is not a rigid product. It is a highly customizable foundation designed to mold to your existing infrastructure.

Why Build When You Can Adopt?

	Internal Scratch Build	Open SWE
Context Management	Custom engineering, highly prone to context overflow.	Deep Agents native file-memory & isolated sub-agents .
Architectural Upgrade Path	Stagnant; requires constant manual refactoring.	Inherits open-source improvements automatically .
Time to Value	Months of trial and error.	Day 1 deployment.

Skip the architectural guesswork.
Start with production-validated patterns.

Day 1 Blueprint

github.com/langchain-ai/open-swe

1

Install GitHub App
(Repo access & PR generation)

2

Set up LangSmith
(Tracing & sandbox connection)

3

Connect Triggers
(Linear / Slack integrations)

Available now under the MIT License. Download the framework, plug in your infrastructure, and deploy your enterprise coding agent today.