Patchwork

AI agents you can actually trust

tim@cosmicglue.ai

The Problem

AI agents are **powerful** but **dangerous**.

- LangChain: CVSS 9.3 — prompt injection exfiltrated API keys from env vars
- CrewAI: CVSS 9.2 — leaked a GitHub admin token through error handling
- Zapier/Make: constraints are prompt instructions the LLM might ignore
- Nobody enforces per-operation allow-lists or typed parameter

The Solution

Patchwork — the LLM proposes, the runtime validates.

Capability-scoped agents where constraints are enforced in **compiled Rust**, not natural language. Credentials are encrypted and resolved at runtime — the AI never sees your tokens.

```
User → AI reasons → proposes tool call → Runtime validates → executes
```

The LLM is in **userspace**. The runtime is the **kernel**.

How It Works

Patches

Atomic capabilities. Each wraps a single API operation with a typed contract.

Quilts

Scoped compositions. Which patches, which operations, what constraints, whose credentials.

Runtime

Demo: Calendar Booking

"When am I free this week?"

`calendar_id: { fixed: "primary" }` — the LLM doesn't even know the parameter exists.

Constraints enforced by compiled code, not prompt instructions.

The Architecture

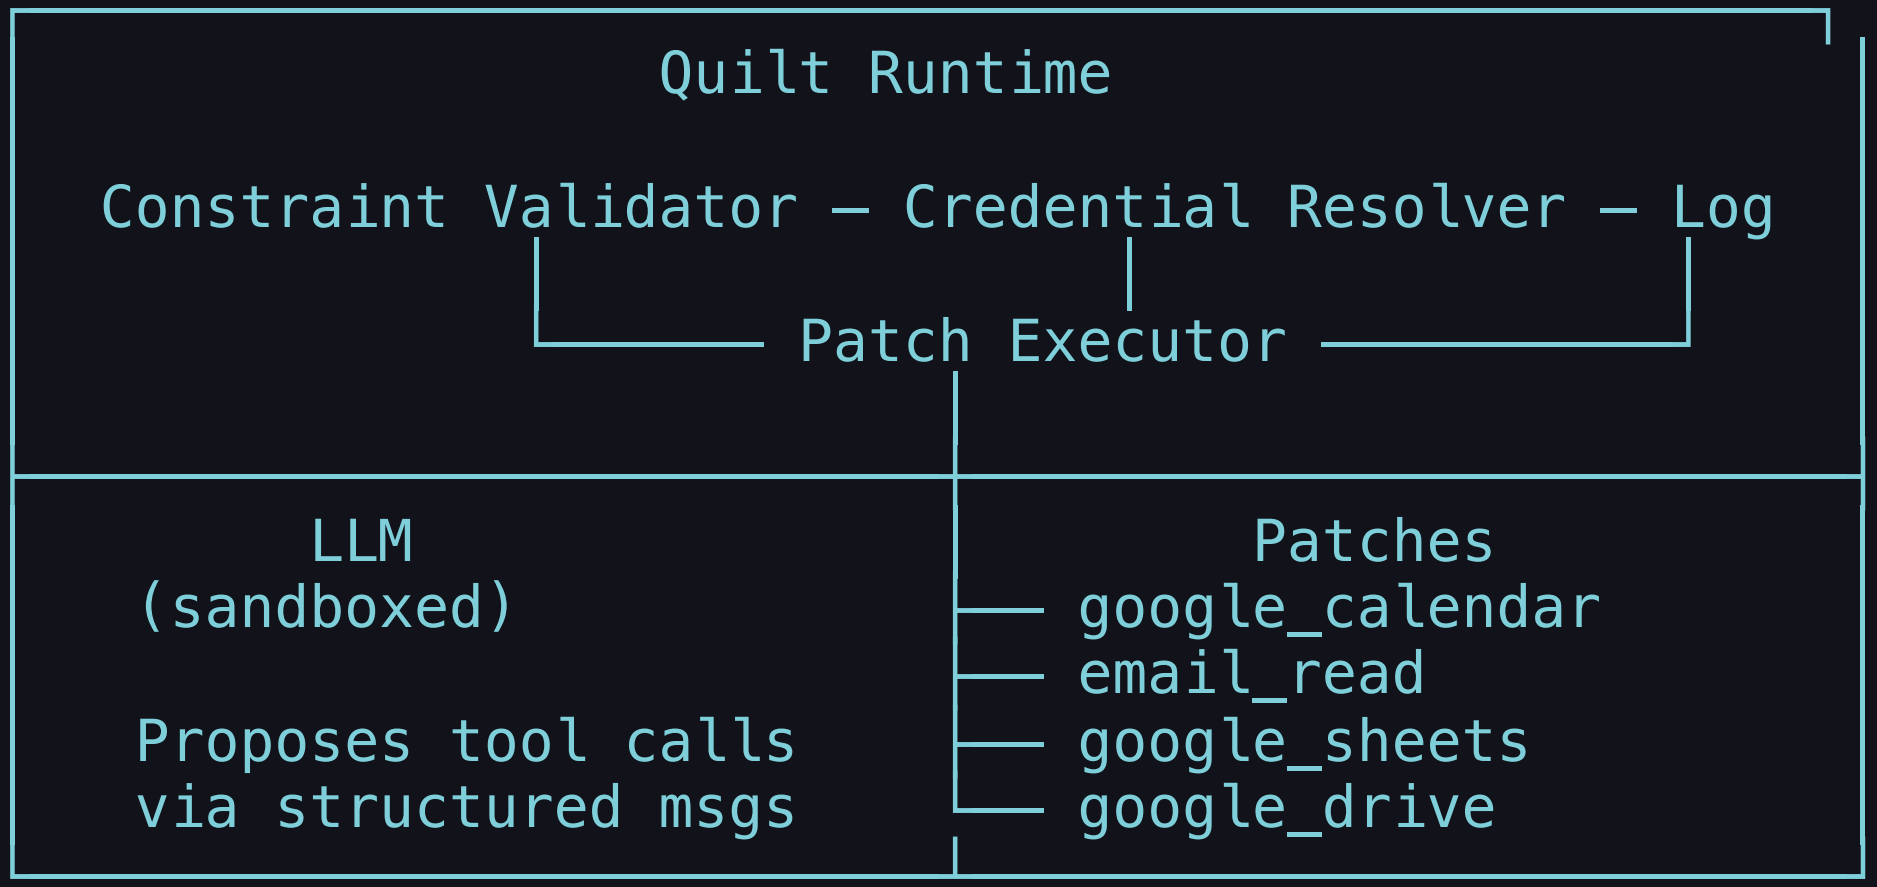

Constraint Model

Four types — all enforced at runtime, invisible to the LLM:

Type	What it does	Example
Fixed	Locked value.	`calendar_id: "primary"`

Violations return structured errors — the LLM retries within a budget.

Credential Isolation

The **#1 security differentiator** vs every competitor.

- LLM sees `$OWNER_OAUTH`, never the actual token
- Tokens encrypted at rest (AES-256), decrypted only at execution time
- OAuth refresh handled by credential resolver, not the LLM
- Sharing a public quilt = sharing a **capability**, not sharing **access**

> *LangChain stored API keys in env vars accessible to the LLM*

Market Opportunity

TAM $50B — AI Automation

SAM $8B — AI-Native Workflow Automation

SOM $500M — Security-First AI Agents

Zapier has no AI reasoning. ChatGPT has no automation. LangChain has no security model.

Patchwork is the only platform where capabilities are scoped, constraints are compiled, and credentials are

Business Model

Pro — $49/mo

Build and run quilts. OAuth connections. Full audit trail.

Team — $149/mo

Shared quilts, team credentials, API mode for agent-to-agent.

Enterprise — Custom

Self-hosted runtime. Custom patches. Compliance features.

All tiers include full observability. We don't gate security

Traction

Built and running — Conversational quilts, 10+ patches, OAuth integrations, SSE streaming, constraint engine, credential isolation, full audit logging

Patches live — Google Calendar, Email, Google Sheets, Google Drive, Domain Checker, Web Scrape, Trademark Search, Crunchbase, Google Search

Partnerships

Team

Tim — Founder. Full-stack engineer. Built AI products. Ships daily.

Hiring — 1 founding engineer

The rest? Claude and Patchwork. Agents all the way down.

Meta — This pitch deck, the marketing site, and the product itself were built with AI assistance. We eat our own cooking.

The Ask

Raising $400K pre-seed — 12 months runway

Use of Funds:

- 70% Engineering (founding team + infrastructure)
- 20% Go-to-market (developer community, content)
- 10% Infrastructure + ops

12-month goals: Public launch, 500 active quilts, marketplace beta, $10K MRR

Why Now?

- AI crossed the capability threshold — agents can reason and use tools
- **But trust hasn't caught up** — every platform trusts the LLM by default
- LangChain/CrewAI CVEs prove the market needs architectural security
- First mover in **capability-scoped agents** — compile-time constraints, not prompts

The LLM proposes. The runtime validates.

Your data stays yours.