

Forum:[Hack Forums](#)

Thread: [Agent v0 a AI terminal CLI opensource](#)

thatcyplex34 - April 16, 2026, 1:34 am

Introducing Agent v0 an open-source multi-agent AI orchestration terminal built for security researchers. It connects multiple AI providers into one unified CLI, runs 10 specialized security agents in sandboxed environments, encrypts everything at rest, and logs every action to a tamper-evident audit trail. Built with Rust, TypeScript, and Python. Here's how it evolved from a rough prototype to where it is today.

v0.1.0 — "Agent Cyplex"

Where it all started. Originally called Agent Cyplex. Basic multi-agent orchestration with local AI support through Ollama and LM Studio. Interactive setup wizard, background daemon, streaming AI responses in the terminal. Raw, experimental, and full of ambition.

v0.2.0 — Package Overhaul & UI Redesign

First major cleanup. Redesigned terminal UI with block-letter banner, braille loading spinner, and boxed Unicode panels. Bumped every major dependency — Anthropic SDK 0.39 to 0.81, OpenAI 4.77 to 6.33, React 18 to 19, Commander 13 to 14. Removed local AI support entirely. Cloud providers only from here on: Anthropic, OpenAI, Google Gemini.

v1.0.0 — Security Audit & SSH Tunnel Removal

The rename. Agent Cyplex became Agent v0. Deleted the entire Go SSH tunnel module after a security audit found 9 vulnerabilities — disabled host key verification enabling MITM attacks, race conditions in reconnection, missing key file permission checks. Rather than patching 9 issues in a module we no longer needed, we deleted the whole thing. Also removed the ssh2 npm dependency and fixed a path traversal vulnerability in skill intake.

v1.2.2 — Web Dashboard & Server Hardening

First web dashboard. HTTPS server with self-signed TLS certs, Socket.IO real-time connection, dark terminal-themed UI. Added security headers across all HTTP servers — CSP, HSTS, X-Frame-Options, X-Content-Type-Options, Referrer-Policy, Permissions-Policy. Rate limiting on auth at 5 attempts per 60 seconds. JSON body limit capped at 10KB.

v1.3.0 — Memory System & Audit Trail

Added persistent memory system for agent state and context. Built a tamper-evident audit trail — every agent action logged to a SHA-256 hash chain where each entry links to the previous one. Web dashboard got the Security Audit tab for real-time audit viewing. Shell command validation layer added.

v1.3.2 — Rust Crypto Hardening (22 fixes)

The big Rust security pass. 22 vulnerabilities patched. Timing attack prevention with constant-time comparison on session tokens and audit hash verification. All crypto random generation moved to OS-seeded CSPRNG. IPC protocol capped at 16 MiB from an uncapped 4 GB which was a DoS risk. HMAC operations return Result instead of panicking. Hard-coded test keys replaced with random values. Wildcard domain matching normalized to case-insensitive per RFC 4343.

v1.4.4 — TypeScript & Python Hardening (21 fixes)

21 more vulnerabilities patched across TypeScript and Python. Every execSync call replaced with

execFileSync to eliminate shell injection across the entire TS codebase. Agent Bash tool now blocks backtick substitution, command substitution, pipe-to-bash, and reverse shells. Secret redactor hardened against prototype pollution. API keys removed from the env template and stored only in the encrypted keystore. All 5 Python forensics modules validate file paths and block directory traversal. OSINT tools validate domain and email format before external requests.

v1.5.0 — macOS Apple Silicon

Native support for M1 through M5. One install command works on both Linux and Mac. Linux gets bubblewrap sandboxing with namespaces and seccomp. macOS gets sandbox-exec with deny-by-default profiles. Platform-specific paths, socket locations, daemon auto-start with systemd on Linux and LaunchAgent on macOS, and Homebrew detection all handled automatically.

v1.6.0 — Claude Code Adapter

New provider option that routes AI requests through Anthropic's Claude Code CLI instead of direct API calls. Spawns claude as a subprocess with no separate API key needed using your existing Claude Code auth. Supports both single-shot and streaming responses. Full tool-use loops and extended thinking behind the standard gateway interface.

v1.7.0 — Neon Architect Web UI

Complete web dashboard redesign with the Neon Architect design system. Tonal sectioning with no hard borders and panel separation through surface color shifts. Glassmorphism with translucent modals and backdrop-filter blur. Typography using Space Grotesk headlines, Manrope body, and JetBrains Mono code. Neon accent gradient from blue to purple. Left icon rail with agent fleet sidebar, tabbed content for Terminal, Tasks, Memory, and Security, plus an AI Chat panel with model selector and streaming responses.

v1.8.0 — Subscription-Based Auth

No more managing API keys. Set auth mode to subscription and authenticate through your existing paid subscriptions. Anthropic routes through Anthropic's Claude Code CLI with zero credentials stored by Agent v0. ChatGPT uses your session token from chatgpt.com encrypted in the keystore. Gemini uses gcloud ADC with OAuth tokens that auto-expire in 60 minutes and are cached for 50. All tokens encrypted at rest with AES-256-GCM and Argon2id. Setup wizard updated with step-by-step guidance for each provider.

v1.9.0 — Web Dashboard Complete

Every broken feature in the web dashboard fixed and shipped. All 7 side panels wired up including Fleet, Search, Tools, Extensions, Settings, Account, and Chat toggle. Memory management with delete on individual memories, clear all with confirmation, and error toasts on failure. AI chat with full markdown rendering including fenced code blocks, inline code, bold, and italic. Socket events for chat complete, memory error, and memories cleared are now properly emitted and handled. Design system compliance with no solid borders and tonal backgrounds only.

By the numbers:

18 releases across 16 days

43 security vulnerabilities patched

4 languages used: Rust, TypeScript, Python, and Go which was removed

AES-256-GCM encryption, Argon2id key derivation, Ed25519 signatures, SHA-256 audit chains

Runs on Linux and macOS Apple Silicon

3 AI providers: Anthropic, OpenAI, Gemini — with Claude Code as an alternative Anthropic route

2 auth modes: API key or subscription

10 built-in security agents

Open source at <https://github.com/centeler34/Agent-v0>

