ALPHA — Project Description

One-liner

ALPHA is a personal agent-assistant development hub: one interface to build, connect, orchestrate, observe, and heal
autonomous AI agents, swarms, bots, MCP servers, APIs, and databases.

Short description (≤ 350 chars, for repo “About” fields)

Autonomous Agent Intelligence Platform — Agentic AI Mesh, Asynchronous DAG workflows, A2A comms, agent swarms, MCP
registry, SSD agent memory, LLMOps and self-healing ops in one zero-trust, post-quantum-ready command center GUI.

Suggested topics

`agentic-ai` `multi-agent` `agent-mesh` `dag` `workflow-orchestration` `llmops` `agentops` `mcp` `a2a` `agent-swarm` `self-healing` `zero-trust` `post-quantum-cryptography` `nextjs` `typescript` `dashboard`

Problem

Teams building with AI agents juggle disconnected tools: one for prompt/LLM telemetry, another for workflow
orchestration, another for tool/MCP servers, and ad-hoc logs for failures. Agent-to-agent traffic, memory, and
remediation are rarely visible in one place, making autonomous systems hard to trust.

Solution

ALPHA unifies ten operational surfaces behind a single navigation shell:

1. **Command Center** — fleet-wide health at a glance
2. **Agent Mesh** — topology of every agent, tool, and data source
3. **DAG Workflow Editor** — design asynchronous, branching, looping workflows
4. **A2A Communications** — inspect and compose agent-to-agent messages
5. **Swarm Manager** — spawn and steer collectives with explicit coordination strategies
6. **MCP Registry** — MCP servers, APIs, and databases as first-class resources
7. **Memory Store (SSD)** — layered, persistent agent memory with vector search
8. **AgenticOps / LLMOps** — cost, tokens, latency, evals, routing
9. **Self-Healing Console** — detect → diagnose → heal → resolve
10. **Assistant Hub** — natural-language control plus an agent-builder wizard

Design principles

- **Controlled autonomy** — every loop has a budget, a stop condition, and an escalation path (Controlled Loop Engineering).
- **Zero-trust by default** — no implicit trust between agents, tools, or users; deny-by-default ACLs.
- **Crypto-agility** — architecture designed for NIST PQC (FIPS 203/204/205) hybrid migration.
- **Observability first** — nothing an agent does should be invisible.
- **Human-in-the-loop for irreversible actions** — payments, production changes, external communications.

Audience

AI/agent engineers, AgenticOps & platform teams, automation consultants, and solo builders designing personal agent assistants.

Status

Preview. The GUI is complete; all data is simulated. See ROADMAP.md.

Maintainer

Adonis Jimenez — @donny-devops