ALPHA Roadmap

Status legend: ✅ done · 🚧 in progress · ⏳ planned · 🔬 research

Phase 0 — GUI Foundation ✅ (current)

- ✅ Cyberpunk-enterprise design system (glass panels, neon palette, three-font system)
- ✅ Ten-module navigation shell with code-split pages
- ✅ Mock data engine: 30+ mesh nodes, swarms, MCP servers, A2A messages, memory entries, models, incidents
- ✅ Interactive mesh graph, DAG canvas, charts, Kanban, chat UI, agent-builder wizard
- ✅ ALPHA branding applied across UI, code, and docs
- ✅ Governance & security documentation, CI and policy workflows

Phase 1 — Persistence & State ⏳

- ⏳ Database schema for agents, workflows, runs, messages, memory, incidents
- ⏳ Replace `lib/mock-data.ts` reads with API routes (keep mock mode behind a flag for demos)
- ⏳ Save / load DAG workflows as versioned JSON; import/export YAML
- ⏳ Authentication with role-based access (owner, operator, viewer)

Phase 2 — Live Agent Runtime ⏳

- ⏳ Assistant Hub wired to a real LLM endpoint with system-state context
- ⏳ DAG executor: topological scheduling, async fan-out/fan-in, retries, timeouts, per-node budgets
- ⏳ Controlled Loop Engineering: max-iteration, token-budget, wall-clock, and convergence stop conditions
- ⏳ A2A message bus with envelope schema, trace IDs, and delivery guarantees
- ⏳ Real MCP client connections (stdio / SSE / streamable HTTP) with per-tool allow-lists

Phase 3 — Observability & Self-Healing ⏳

- ⏳ OpenTelemetry traces for every agent step, tool call, and LLM request
- ⏳ Token / cost / latency metrics from real runs
- ⏳ Healing playbooks executed automatically for low-severity faults (retry, model fallback, reroute, memory flush)
- ⏳ Human-in-the-loop queue for high-severity or irreversible remediation

Phase 4 — Zero-Trust & Post-Quantum Hardening 🔬

- 🔬 Per-agent workload identity and short-lived credentials
- 🔬 Deny-by-default agent × tool ACL matrix enforced server-side
- 🔬 Hybrid key exchange (X25519 + ML-KEM-768/1024, FIPS 203) for A2A channels
- 🔬 Signed audit log entries (ML-DSA FIPS 204 / SLH-DSA FIPS 205)
- 🔬 Tamper-evident, append-only audit trail

Phase 5 — Agentic Commerce & Business Automation 🔬

- 🔬 Spend mandates with per-transaction and per-period limits (docs/AGENTIC_COMMERCE.md)
- 🔬 Payment-provider integration in sandbox mode first; HIL approval above thresholds
- 🔬 Business automation template marketplace (support, ETL, content, DevOps, research)

Non-Goals

- Building a proprietary LLM
- Unbounded autonomous loops without budgets or human escalation
- Claiming certification or compliance that has not been independently assessed

Suggestions welcome via a **Feature request** issue.