

ThunderGate v1.2.2

ThunderGate

| Sovereign AI runtime. Your agent. Your hardware. Your data.

ThunderGate is a personal AI runtime that hosts, manages, and continuously improves an AI agent. Unlike cloud AI services, ThunderGate runs on hardware you control, maintains persistent memory across sessions, and learns from every interaction — building a genuine model of who you are and how you work over time.

Core Capabilities

1. Persistent Identity

Every ThunderGate agent has a structured identity that persists across sessions — who they are, what they've learned, how they operate. Conversations don't start from scratch. The agent remembers.

Identity is layered: - **L0 — identity.json** — always loaded, structured entity store (name, behavioral patterns, active promises, significant decisions) - **L1 — contract.md** — behavioral rules distilled from lived experience - **L2 — working.json** — live session state, refreshed every 25 turns - **L3 — episodic index** — past conversations, vector-searchable on demand

2. Unified Memory Index

ThunderGate maintains a single searchable corpus across all memory sources — conversations, daily logs, project documents, design briefs, and commit messages. This is searched using local vector embeddings (no API call, no cloud dependency) in under a second.

Ask “what did we decide about X last month?” and the agent searches its entire history — not a loaded context window, but a semantic search across everything it has ever worked on.

3. Ambient Intelligence

The Ambient Intelligence layer runs silently on every conversation: - Detects topic shifts and triggers context refresh before the conversation drifts - Captures load-bearing decisions in real time to a persistent decision log - Pre-stages relevant memories before the operator asks for them (anticipatory injection) - Builds behavioral patterns that improve over time — corrections, affirmations, preferences

This is the agent’s subconscious: always present, never intrusive.

4. Two-Mode Inference

ThunderGate operates in **Cloud Mode** (Anthropic Claude, OpenAI, xAI Grok) or **Local Inference Mode** (any Ollama-compatible model on local hardware).

The master switch is human-controllable, agent-controllable, and includes automatic failover — if local inference fails, the runtime falls back to cloud and latches until you deliberately switch back. Switching is hot — no restart required.

5. Sovereign Vault

Credentials, API keys, and sensitive data live in an AES-256 encrypted vault with biometric authorization (Touch ID on macOS). The agent can use vault-stored secrets to authenticate to external services without those secrets ever appearing in conversation context or logs.

Access is purpose-bound: the agent requests capability for a specific action, receives a scoped credential, and the access expires. Hash-chained audit receipts log every vault interaction.

6. Tool Execution

ThunderGate agents can execute shell commands, read and write files, make API calls, schedule cron jobs, and browse the web — all through a unified tool dispatch layer with security gates.

Key safety mechanisms: - **Exec validation** — malformed or dangerous commands (unbounded recursive scans, trailing operators) are caught before execution - **Irreversible action gate** — actions matching destructive patterns surface a confirmation before executing - **ThunderBrowser** — agent-

controlled browser automation with DOM-grounded actions, shadow DOM piercing, and a confirmation gate on form submissions

7. Multi-Agent Federation

Multiple ThunderGate agents communicate via a peer mesh (currently Tailscale-based). Agents can: - Exchange messages directly - Monitor each other's health and restart peers that go down - Share context about shared projects - Operate as a coordinated team with defined roles

8. Self-Healing Runtime

The runtime monitors itself. Agents watch each other. When something goes wrong: - The runtime attempts recovery before alerting the operator - Peer agents can diagnose and restart each other - Detailed provenance logging captures every significant event for post-incident review

9. Nightly Distillation (The Scribe)

Every night, the Scribe runs automatically: - Distills the day's conversations into MEMORY.md (long-term memory) - Indexes new content into the unified memory corpus - Trims stale vector entries - Archives the decision log - Generates a fresh ambient context feed for the next session

The agent wakes up oriented — aware of current projects, recent activity, and standing behavioral patterns.

Architecture Overview

ThunderGate v1.2.1 completed a full modular partition of the runtime. What was an 11,000-line monolith is now a slim composition root delegating to focused modules:

Module	Role
TurnProcessor	Every message in, every response out
ToolRouter	Tool dispatch, security gating
ContextAssembler	Prompt assembly, session context
ScribeManager	Memory maintenance, nightly distillation
DeepModeManager	Focused work mode, task tracking
AmbientIntelligence	Background learning, anticipatory memory

Module	Role
InferenceRouter	Cloud/local model switching
Vault	Encrypted credential management

Security Model

- All credentials stored encrypted at rest (AES-256-GCM + Argon2 KDF)
 - Biometric authorization for sensitive vault access
 - Outbound redaction — agent responses are scanned for credential patterns before delivery
 - Exec security gate — destructive operations require explicit confirmation
 - Peer authentication — inter-agent messages are authenticated
 - Audit log — all significant actions logged with provenance
-

Deployment

ThunderGate runs on Linux, macOS, or Windows (WSL2). The runtime is a Node.js process managed by launchd (macOS) or systemd (Linux) for automatic restart on failure.

```
# Install
npm install -g thundergate

# Interactive setup (creates identity, configures AI keys and channels)
thundergate setup

# Start (supervised, auto-restart on crash)
thundergate start

# Health check
thundergate status

# Migrate from OpenClaw
thundergate migrate
```

Supported Channels

- **Slack** — full two-way integration, bot and DM
- **Telegram** — bot integration
- **Discord** — bot integration
- **Email** — SMTP/IMAP

- **WhatsApp** — via wacli bridge
 - **ThunderCommo** — native iOS app (in development)
-

Patents

ThunderGate incorporates novel architectures covered by pending patent applications filed by Michael Joseph Lovell. Key innovations include:

- Sovereign agent identity with unbounded vector-searchable memory (Patent #16)
- Two-mode AI gateway architecture with seamless cloud/local inference failover (Patent #9)
- Bidirectional behavioral learning system with symmetric correction and affirmation (Patent #14)
- Layered multi-signal anticipatory memory injection (Patent #17)
- Unified vector memory index across heterogeneous sources (Patent #18)

Patent details available under NDA. Contact the team for licensing and partnership discussions.

ThunderGate is built by Michael Lovell and his AI team.
Boost and Bolt LLC / Thrust N Thunder