



cognipeer

Open Source AI Ecosystem

developer.cognipeer.com



Console
Core Platform



Agent Server
Serving



Agent SDK
Runtime



Chat UI
Experience

cognipeer Opensource AI Ecosystem





Console
Core Platform



Agent SDK
Runtime



Agent Server
Serving



Chat UI
Experience



Console

LLM Gateway

Vector Orchestration

Tracing & observability

Files Pipeline

RAG

Alerting

Cost & Budget Control

Multi-tenant isolation

Memory Management

Licence & feature control



Agent SDK

React Agent

Planning (TODO)

Native Guardrail

Human-in-the-loop

Smart Agent

Pause / Continue

Summarisation

Structed Output

Detailed Tracing



Agent Server

Multi-Agent Registry

Publish & Version

REST API

SSE Streaming

Auth (JWT/Token)

Storage (Postgres/Mongo)

File Upload

Swagger / OpenAI



Chat UI

Full Chat

Streaming UI

Tool Cards

File Attachments

Themes (Dark/Light)

Cognipeer Console — AI Gateway & Control Plane

Self-Hosted AI Infrastructure, Fully Controlled

A deterministic TypeScript agent runtime for planning, safety, memory, and multi-agent orchestration.



Console

LLM Gateway

Enables existing clients to operate as a seamless drop-in replacement, supporting chat, embeddings, streaming, and tool calling.

Vector orchestration

Manages multiple vector databases through a unified control layer, including index lifecycle and operational workflows.

Tracing & observability

Tracks agent executions, tool calls, and LLM requests with integrated platform-level visibility.

Files Pipeline

Handles document uploads and optional Markdown conversion for RAG.

RAG

Supports knowledge grounding via document ingestion and vector search.

Alerting

Delivers proactive alerts on spend, performance, and operational anomalies.

Cost & Budget Control

Provides real-time usage visibility, budget enforcement, and cost optimisation.

Multi-tenant isolation

Provides tenant-level separation across quotas, limits, tokens, API keys, and configuration settings.

Memory Management

Maintains contextual continuity through structured memory storage and retrieval.

Licence & feature control

Enables plan- and tenant-based feature flagging and access governance.

Console — AI Gateway & Control Plane Comparison



Product	Primary Role	Self-Hosted	Open Source	LLM Gateway	Routing / Fallback	Rate Limiting / Quotas	Caching	Prompt Management	Tracing / Observability	Multi-Project
Cognipeer Console	AI Gateway + Control Plane	✓	✓ (MIT)	✓	✓	✓	✓	✓	✓	✓
LiteLLM Proxy	AI Gateway	✓	✓	✓	✓	✓	△	×	△	△
Helicone	Gateway + Observability	✓	✓	✓	△	△	△	×	✓	×
Portkey	Managed AI Gateway	△	△	✓	✓	✓	✓	×	✓	△
Cloudflare AI Gateway	Edge AI Gateway	×	×	✓	✓	✓	✓	×	✓	×
Kong AI Gateway	Enterprise API Gateway	✓	△	✓	✓	✓	✓	×	△	△
Envoy AI Gateway	Kubernetes-native Gateway	✓	✓	△	✓	✓	△	×	△	△
Langfuse	LLM Observability Platform	✓	✓	×	×	×	×	✓	✓	△



Cognipeer Agent SDK

Build Production-Grade AI Agents with Full Runtime Control

A deterministic TypeScript agent runtime for planning, safety, memory, and multi-agent orchestration.



Agent SDK

Message-First

Executes agent reasoning through a transparent, message-driven runtime with full conversational traceability.

Planning (TODO)

Structures agent workflows through managed task planning and explicit execution steps.

Summarisation

Compresses long conversational context using token-aware memory summarisation.

Guardrails

Applies runtime safety and policy controls to govern agent behaviour and outputs.

Interrupt

Allows live agent execution to be stopped, redirected, or overridden in real time.

Pause / Continue

Persists agent state to enable workflow pausing and asynchronous continuation.

Human-in-the-loop

Introduces approval and review checkpoints within agent decision processes.

Structured Tracking

Captures detailed execution traces, tool calls, and reasoning events.

Agent SDK — Agent Runtime Framework Comparison



Product	Deterministic Runtime	Planning	Summarisation	Guardrails	Interrupt	Pause / Continue	Human-in-the-Loop	Multi-Agent	Structured Output
Cognipeer Agent SDK	✓	✓	✓	✓	✓	✓	✓	✓	✓
LangChain (JS)	×	△	×	×	×	×	×	△	△
LangGraph	×	△	×	×	×	△	×	✓	△
OpenAI Assistants API	×	×	×	△	×	×	×	×	△
CrewAI	×	△	×	×	×	×	×	✓	×
Microsoft Autogen	×	△	×	×	×	×	×	✓	×

Cognipeer Agent Server — Agent Serving Infrastructure

Serve, Publish, and Scale AI Agents via Unified APIs



A production-ready serving layer that registers, versions, and exposes multiple AI agents as secure, streaming-enabled REST services.

 Agent Server

Multi-Agent Registry

Registers and manages multiple agents within a central serving catalogue.

Auth (JWT / Token)

Secures endpoints using token-based and JWT authentication models.

Storage (Postgres / Mongo)

Persists conversations, messages, and metadata in scalable databases.

File Upload

Supports document ingestion and attachment handling within agent interactions.

Publish & Version

Controls agent lifecycle through versioning, release, and publishing workflows.

REST API

Exposes agents as standardised RESTful endpoints for application integration.

SSE Streaming

Streams real-time agent responses via Server-Sent Events.

Swagger

Auto-generates interactive API documentation and testing interfaces.

Agent Server — Agent Serving & API Infrastructure Comparison



Product	Primary Role	Multi-Agent Serving	Streaming (SSE)	Built-in Auth	Built-in Storage	File Handling	OpenAPI / Swagger	Framework Adapters
Cognipeer Agent Server	Agent → REST API Infrastructure	✓	✓	✓	✓	✓	✓	✓
Custom FastAPI / Express	DIY API Layer	△	△	△	△	△	△	✓
OpenAI Assistants API	Managed Runtime Endpoint	×	✓	✓	✓	△	×	×
LangServe (LangChain)	LangChain App Serving	△	△	△	△	△	△	△

Cognipeer Chat UI — AI Conversational Interface Kit

Production-Ready Chat Experiences for AI Agents

Chat UI is a React component library for building modern, real-time conversational interfaces powered by AI agents. It delivers a ChatGPT-style experience with streaming responses, tool transparency, and full customisation — ready for enterprise applications.



Chat UI



Chat UI

Full Chat

Delivers a complete conversational interface with message history and threading.

Streaming UI

Renders live agent responses through real-time streaming updates.

Tool Cards

Delivers a complete conversational interface with message history and threading.

File Attachments

Visualises agent tool usage, inputs, and outputs within the chat experience.

Themes (Dark / Light)

Provides built-in theming with full support for brand customisation.

Chat UI — AI Chat Front-End Comparison



Product	Primary Role	Streaming (SSE)	Tool Call Visualisation	File Uploads	Theming (Dark / Light)	Hooks API	Agent Server Ready	Open Source
Cognipeer Chat UI	Agent Chat UI Kit	✓	✓	✓	✓	✓	✓	✓
Vercel AI SDK UI	UI Building Blocks	✓	△	△	△	✓	△	✓
Generic Chat Components	Basic Chat UI	△	×	△	△	△	×	△
Fully Custom Build	From Scratch	✓	✓	✓	✓	✓	✓	×

cognipeer Team

10+ Professionals

**20+ years
Experience**

50+ industries

500+ Projects

1000+ Use Cases

Co-Founders & Managing Partners



Anıl Güleroğlu
AI Lead



M. Seçkin Bedük
Product Lead



Yasin Şükrü Gürdemir
Sales & Growth



Thank you!



hello@cognipeer



/company/cognipeer/



@CognipeerAI

Explore cognipeer
Developer tools

