

Start here

Radium documentation, organised by what you're trying to do. Call the API directly, point an existing tool at it, or wire up tracing and evals.

- [Start here](#)

Start here

Radium is an API. There's no chat interface to log into, you connect your existing client, agent framework, or IDE to Radium the same way you'd connect it to OpenAI or Anthropic, and everything downstream keeps working. This page is the fastest way into whichever part of that you need.

Radium API

How to call Radium directly: authentication, models, environments, tool calling, and the errors you'll actually hit.

- [API quickstart](#)
- [Tool calling and MCP](#)

Coding agents & IDEs

Point your editor or CLI at Radium.

- [Run Claude Code on Radium](#)
- [Use Radium with Cursor](#)

Frameworks & SDKs

Radium inside the frameworks your application already runs on.

- [Use the OpenAI Python SDK with Radium](#)
- [Use LangChain with Radium](#)
- [Use LangGraph with Radium](#)
- [Use LlamaIndex with Radium](#)
- [Use CrewAI with Radium](#)
- [Use DSPy with Radium](#)
- [Use AG2 with Radium](#)
- [Use Agno with Radium](#)

- [Use PydanticAI with Radium](#)
- [Use the Vercel AI SDK with Radium](#)

Observability, evals & guardrails

Trace it, test it, and guard it the same way you would any other model in production.

- [Observability and evals with Radium](#)

Everything, by shelf

- [Radium API](#) — the API itself
- [Integrations](#) — coding agents, frameworks and SDKs, observability