

The 6-Framework Agent Decision Guide

The same 3-step agent, built in 6 orchestrators — which to standardize on, and why.

FRAMEWORK	CAMP	BEST FOR
★ LangGraph	Opinionated state (typed graph)	Explicit control + debuggability when runs go sideways
★ CrewAI	Defer state to you	Fastest to a working agent (~30 lines)
Pydantic AI	Opinionated state (typed)	Typed-Python shops
LlamaIndex	Retrieval-native	When the agent is mostly retrieval
AutoGen	Conversational	Microsoft / enterprise stacks
OpenAI Swarm	Minimal (handoffs)	Learning, not production

THE VERDICT

LangGraph and CrewAI separated from the pack and it's close — two opposed architectures that both win. CrewAI gets you to a working agent fastest; LangGraph gives you explicit control and typed state for when things go wrong. The tiebreaker isn't the tool, it's your team's priority.

- The hard problem is identical in every framework: making the agent stop, retry, and hand off reliably. Syntax changes; that problem does not.
- Count transitive dependencies before you commit — one of the six pins 40+ conflicting LangChain versions.