Do we need Python developers to build multi-agent workflows?

No. CrewAI Studio is a visual editor that enables anyone within an organization to build sophisticated multi-agent workflows via drag-and-drop configuration and a helpful AI copilot. However, Python developers can build custom multi-agent workflows via the IDE of choice with intuitive CrewAI APIs.

How can we automatically trigger workflows with external enterprise systems?

CrewAI Enterprise includes dozens of out-of-the-box connectors for enterprise systems such as GitHub, Google Workplace, HubSpot, Salesforce, Zendesk -- with many more currently in development. These connectors will automatically trigger deployed multi-agent workflows once configured.

How can we ensure that multi-agent workflows are acting as expected?

CrewAI Enterprise includes detailed visual tracing and monitoring, enabling anyone to see which tasks, agents and tools were selected to run a multi-agent workflow, it's current progress and the input and output of each step.

Can we monitor multi-agent workflows from external observability services?

Yes. CrewAI Enterprise includes an event bus which supports custom listeners as well as webhooks, enabling external systems, applications and tools to observe CrewAI events in real time and monitor or react to them. In addition, AgentOps, Arize Phoenix, Langfuse, Langtrace, OpenLit, Opik, Patronux, Portkey and Weave all provide native integration with CrewAI.

Can we built agents that access knowledge stored in vector databases?

Yes. CrewAI Enterprise includes native support for vector databases such as Qdrant and Weaviate. In addition, developers can add support any knowledge base by building a custom tool with the CrewAI API and adding it to the organization's private tools repository.

What large language models (LLMs) can we deploy?

CrewAI Enterprise supports the use of any LLM, on premises or in the cloud, including those from inference providers such as AWS, Anthropic, Azure, Cerebras, Cohere, Fireworks, Google, Grok, Hugging Face, IBM Watson, Meta, Mistral, Nvidia, Ollama, Open Router, OpenAI, Perplexity and SambaNova.

Can we use agents and tools available in other agent frameworks and services?

Yes. CrewAI Enterprise enables local agents to access remote Amazon Bedrock agents via a configurable tool. In addition, agents can take advantage of native support for LlamaIndex and LangChain tools.

How can we determine which LLM an agent should use?

CrewAI Enterprise can test multi-agent workflows with multiple LLMs, and produce quality scores for each of them. However, a best-of-breed approach is recommended with different agents using different LLMs based on their tasks (e.g., generation vs. tool use).

How can we ensure high-quality outputs, and improve them as needed?

CrewAI Enterprise includes automated testing and training features. The testing feature can be used to produce multi-agent workflow quality scores while the training feature can be used to improve the outputs of specific agents. Further, quality, hallucination and ROI metrics can be monitored across versions.

How can we keep an eye on LLM token costs?

CrewAI Enterprise includes usage dashboards to monitor token count (both overall and for each multi-agent workflow), performance and ROI. The estimated ROI metrics are calculated based on current LLM provider token costs and the amount of time saved by automating a task or business process.

Ready to automate at scale with AI agents built for the enterprise?

Start free trial