

Agentic AI with LangChain

Design, build and evaluate reliable AI agents with LangChain and LangGraph — from a single tool call to multi-agent workflows.

DURATION

3 days · 24 hours

LEVEL

Advanced

CURRICULUM

10 modules · 6 labs

DELIVERYLive online, Onsite,
Corporate

Course overview

AI agents go beyond answering questions: they plan, call tools, read data and take actions across several steps. Building agents that behave reliably is an engineering discipline, and this three-day course teaches it hands-on using LangChain and LangGraph.

Day one establishes the building blocks: chat models, prompts and structured output with LangChain, tool calling, and the reasoning-and-acting loop. You will then move to LangGraph to model agents as explicit graphs of state, nodes and edges, which makes control flow, retries and branching far easier to reason about.

Day two adds knowledge and memory. You will build agentic retrieval-augmented generation that decides when and how to search, add short- and long-term memory, persist state with checkpoints and include human approval before sensitive actions. You will also connect agents to external tools and data, including the Model Context Protocol.

Day three focuses on reliability: multi-agent patterns such as supervisors and specialists, tracing and debugging, building evaluation datasets, measuring task success and cost, and adding guardrails against prompt injection and unsafe actions. The course ends with a capstone in which you build, evaluate and demonstrate an agent that automates a realistic business workflow.

This is an original Ace Skills course with no vendor exam. On completion you receive an Ace Skills Consulting course completion certificate.

Highlights

- 24 hours of live, trainer-led sessions
- Build agents with LangChain and LangGraph
- Agentic RAG, memory and human-in-the-loop steps
- Evaluation, tracing and guardrails for production
- Capstone: a multi-step agent for a business workflow

What's included

- LangChain and LangGraph from fundamentals to multi-agent systems
- Tool calling, agentic RAG and memory
- Human-in-the-loop approvals and checkpoints
- Model Context Protocol integration
- Tracing, evaluation and cost tracking
- Guardrails and security patterns
- Capstone agent project
- Course completion certificate from Ace Skills Consulting

Who this course is for

Who should attend

- Python developers building LLM applications
- Machine learning and data engineers
- Solution architects designing AI systems
- Generative AI for Professionals graduates with coding experience

Prerequisites

- Solid Python programming skills
- Experience calling LLM APIs and basic RAG concepts
- Familiarity with JSON and REST APIs

Detailed curriculum — 3 days, 24 hours, 10 modules, 6 hands-on labs

Day 1 Agent Foundations with LangChain and LangGraph

8 hours

1.1 From LLM calls to agents

- What makes an agent
- When not to use an agent
- Reasoning-and-acting loops
- Cost, latency and reliability trade-offs

1.2 LangChain building blocks

- Chat models and prompt templates
- Defining tools
- Structured output with schemas
- Tool calling

LAB Build a tool-calling assistant for currency and calendar tasks · 45 min

1.3 LangGraph fundamentals

- State, nodes and edges
- Loops and stopping conditions
- Conditional routing
- Error handling and retries

LAB Rebuild the assistant as a LangGraph agent · 60 min

Day 2 Knowledge, Memory and Integration

8 hours

2.1 Agentic RAG

- Retrieval as a tool
- Self-checking answers against sources
- Query rewriting
- Hybrid search

LAB Agentic RAG over a product knowledge base · 60 min

2.2 Memory and persistence

- Short-term conversation state
- Checkpoints and resuming runs
- Long-term memory stores
- Human-in-the-loop interrupts

LAB Add memory and an approval step before actions · 45 min

2.3 Connecting to the world

- Model Context Protocol servers and clients
- Streaming and async execution
- Calling internal APIs and databases

Day 3 Multi-Agent Systems, Evaluation and Capstone

8 hours

3.1 Multi-agent patterns

- Supervisor and specialist agents
- Planning and plan-execute patterns
- Hand-offs
- Shared state

LAB Build a supervisor with two specialist agents · 45 min

3.2 Evaluation and observability

- Tracing runs
- Task success and trajectory evaluation

3.3 Guardrails and deployment

- Prompt injection defences
- Output validation

3.4 Capstone

- Designing the workflow
- Demonstration and review

- Evaluation datasets
- Tracking cost and latency


- Tool permissions and least privilege
- Deploying agents as services

- Building and evaluating

LAB Build, evaluate and demo a business workflow agent · 120 min

Skills and tools

Skills you will build



Tools used in class



Frequently asked questions

Is there an exam?

No. This is an original Ace Skills course with no vendor exam; you receive an Ace Skills course completion certificate.

How much Python do I need?

You should be comfortable writing functions, classes and working with packages and APIs; this is an advanced course.

Which models do the labs use?

Labs are model-agnostic and work with several leading LLM providers; we discuss trade-offs between them.

Will the code work after the course?

Yes. You keep the notebooks and capstone code, and we note which library versions they were tested against.

Should I take Generative AI for Professionals first?

If you have not built a RAG application or called LLM APIs before, yes — it covers the groundwork this course assumes.

Book a batch or talk to an advisor

Course page: aceskillsconsulting.com/course/agentic-ai-with-langchain · Email: info@aceskillsconsulting.com

Chennai office

Guindy, Chennai, Tamil Nadu

Delhi office

B-95, Old Post Office Street, Shakarpur
New Delhi 110092
Office: 98180 12523

All product names, logos, brands, trademarks and registered trademarks are the property of their respective owners. Ace Skills Consulting is an independent training provider; unless explicitly stated, it is not affiliated with, endorsed by or sponsored by any vendor mentioned on this site. Exam details are provided for guidance and may change — always confirm on the vendor's official website. Curriculum generated on 23 Oct 2026; topics and order may be adjusted by the trainer to suit each batch.