Agentic AI Engineering & Architecture Accelerator. - Mastering Agentic AI using LangGraph, AutoGen, and CrewAI Frameworks

Program Description

This comprehensive curriculum path provides the strategic mechanics required to design, orchestrate, and deploy autonomous, intelligent multi-agent systems powered by enterprise-grade Large Language Models (LLMs). Through rigorous engineering tutorials, participants master industry-leading frameworks like LangGraph, AutoGen, and CrewAI. The Program bridges architectural theory with immediate execution by integrating modern standards like the Model Context Protocol (MCP).

By eliminating technical trial-and-error, this intensive accelerator saves enterprise engineering teams estimated 30 to 40 hours of raw R&D, framework debugging, and infrastructure setup.

Executive Overview & Financial ROI

- **Quantified Engineering Efficiency:** By utilizing pre-built corporate boilerplates and architecture templates, this program slashes estimated **30 to 40 hours of R&D time per engineer**, translating into savings of approximately USD3,000 to USD**4,000 USD** in engineering overhead.
- **Mitigating Integration Costs:** Implementing the Model Context Protocol (MCP) systematically solves the costly integration bottleneck
- **Accelerated Deployment Capital:** Instead of dedicating weeks to debugging raw open-source documentation across three competing frameworks, teams shift from concept to production-ready multi-agent workflows within single-sprint cycles.

Benefits of the Program

- **Master Multi-Framework Orchestration:** Gain engineering mastery across LangGraph, AutoGen, and CrewAI ecosystems to select optimal framework based on specific corporate compliance and business constraints.
- **Solve the Integration Bottleneck:** Future-proof enterprise data infrastructure by leveraging the Model Context Protocol (MCP) to securely connect LLMs with legacy tools and private cloud datasets.
- **Deploy Production-Ready Automation:** Move far beyond basic conversational chat interfaces to build cognitive enterprise systems capable of multi-step reasoning, conditional routing, and automated task delegation.
- **Accelerate Time-to-Market:** Equip technical teams with clean, reusable boilerplate code and scalable architectural blueprints to slash software development lifecycles.

Course Outline

☑ **Foundations of Agentic AI and LLM-Powered Architectures**
☑ **Building Intelligent Workflows with LangGraph Agents**
☑ **Designing Multi-Agent Systems with AutoGen for Advanced Task Automation**
☑ **Orchestrating Agent Teams with CrewAI**
☑ **Model Context Protocol (MCP) for Agentic AI**

Total Duration

9 Hours of premium, online on-demand engineering curriculum

Target Learners

- AI Developers & Engineers wanting to step beyond prompt engineering into autonomous infrastructure design.
- Software Architects & Tech Leads tasked with building scalable, collaborative multi-agent microservices for their organizations.
- Technical Product Managers & Innovation Officers looking to understand the production challenges, enterprise readiness, and strategic value of Agentic AI.

Technical Prerequisites

- **Programming Proficiency:** Strong intermediate-level competency in Python (handling asynchronous functions, dictionaries, and object-oriented architectures).
- **Development Environments:** Local installation of Docker and Devcontainers for sandboxed template deployment.
- **API Frameworks:** Basic familiarity with environment key management and cloud terminal controls (OpenAI API, Anthropic API, or local Ollama configurations).

Total No. Of Certificates

5 Professional Certificates, one issued for each completed module.

Course Content

Foundations of Agentic AI and LLM-Powered Architectures

- Evolution timeline of autonomous AI agents compared to traditional heuristic systems.
- Deconstructing the essential trio: Perception Module, Cognitive Engine, and Action Module.
- Strategic collaboration mechanics and evaluating multi-agent framework readiness for enterprise deployment.

Building Intelligent Workflows with LangGraph Agents

- Graph-based state management: Developing complex architectures with workflows, nodes, and conditional edges.
- Advanced execution algorithms: Implementing prompt chaining, dynamic routing, and parallel processing execution loops.
- **Applied Production Demos:** Automated corporate client inquiry pipelines, automated product description engines, and smart search optimization frameworks.

Designing Multi-Agent Systems with AutoGen for Advanced Task Automation

- Configuring specialized AutoGen agent personalities and managing context-aware persistent memory modules.
- Collaborative systems engineering: Designing multi-step reasoning pipelines and agent-to-agent communication networks.
- **Applied Production Demos:** Automated IT helpdesk troubleshooting assistants, real-time data interfaces, and automated legal contract compliance reviews.

Orchestrating Agent Teams with CrewAI

- Core structural components: Managing the operational execution trio of Agents, Tasks, and Processes.
- Corporate team formations: Tailoring specialized agent attributes and secure tool access for autonomous collaboration.
- **Applied Production Demos:** AI-driven operational trip planner optimization and custom enterprise application development using Streamlit interfaces.

Model Context Protocol (MCP) for Agentic AI

- Mitigating the complex \(N \times M\) software integration challenges currently facing modern enterprise AI implementations.
- Architectural core elements: Deploying and configuring host systems, operational clients, secure servers, and active connection lifecycles.
- SDKs and Security: Implementing the Python SDK for MCP, optimizing reliable function calling, and conducting robust enterprise risk analysis.

Disclaimers

Third-Party Frameworks & System Updates*: This program references open-source and proprietary software frameworks, including LangGraph, AutoGen, CrewAI, and the Model Context Protocol (MCP). All trademarks, logos, and brand architectures belong to their respective owners. Because these technologies evolve rapidly, the curriculum reflects the framework ecosystem available at the time of production.While we endeavour to update course content post purchase during the access period of client, however, learners need to note there would be time gap for update of the most current course content to be effected in our LMS. Learners also have to note future software changes may require modifications to the provided code bases.*

API Tokens & Third-Party Infrastructure Costs*: Course registration fees strictly cover access to the educational materials and pre-configured repository templates. Learners and corporate organizations must secure their own API credentials. All operational costs incurred from consuming commercial Large Language Model (LLM) processing tokens (including but not limited to OpenAI, Anthropic, or cloud compute resources) during runtime executions remain the sole financial responsibility of the participant or corporate client.*

Data Privacy, Local Governance, & System Security*: The applied production blueprints and deployment configurations provided in this curriculum are intended solely for instructional purposes. Corporate clients are strongly advised to subject any code snippets, API routing structures, or multi-agent communication networks to internal cybersecurity compliance reviews and data safety auditing before attempting deployment into a live, production-grade enterprise architecture.*