

Advanced Certification in Agentic AI Engineering

Architecting Autonomous, Self-Correcting Multi-Agent Ecosystems & Industry MCP Architectures

Course Category	Artificial Intelligence & Platform Engineering
Duration	16 Weeks (150 Hours of Immersive Applied Learning)
Delivery Format	Online Instructor-Led & Sandbox Production Labs
Target Audience	Software Engineers, Cloud Architects, Full-Stack Developers, DevOps & Data Engineers

Course Overview

The landscape of Artificial Intelligence has dramatically evolved from simple prompt construction and text wrappers to **Agentic AI Engineering**. Modern enterprise technical operations focus entirely on building goal-driven, autonomous software systems. These systems are uniquely capable of dynamic planning, navigating complex infrastructure securely, executing contextual tool-calling, and collaborating asynchronously inside decentralized multi-agent networks.

This certification program provides an exhaustive, production-grade learning experience structured to train engineers to safely design, implement, test, and deploy resilient cognitive systems at enterprise scale.

Program Key Highlights

- **100% Applied Engineering:** Built entirely on Python, LangGraph, CrewAI, and Microsoft AutoGen framework components.
- **The Open MCP Standard:** Comprehensive modules on designing custom Model Context Protocol clients and servers to safely integrate AI agents with enterprise filesystems and legacy databases.
- **Production-Grade Metrics & Telemetry:** In-depth setups of AgentOps frameworks using LangSmith and LangFuse to accurately evaluate trace hierarchies, latency matrices, and execution spending.

- **Industry Readiness:** Rigorous architectural alignment matching the strict technical proficiencies tested in enterprise AI tracks from Microsoft Learn and NVIDIA.

Detailed Weekly Curriculum Blueprint

Module 1: Foundations of Agentic Architecture & Reasoning Design Patterns

- **The Shift from GenAI to Agentic AI:** Mapping the autonomous loop (Perceive → Plan → Act → Reflect).
- **Cognitive Framing Architectures:** Implementing ReAct (Reasoning and Acting), Chain-of-Thought (CoT), Plan-and-Solve, and self-reflection loops.
- **Deterministic Schema Enforcement:** Leveraging Pydantic contracts to lock down structural JSON inputs/outputs for function calling.

Module Milestone Deliverable: Build an independent single-agent mathematical and research execution node equipped with native tool-routing mechanisms.

Module 2: State Machines & Stateful Orchestration via LangGraph

- **Graph System Topologies:** Designing cyclic workflows, structural state management objects, Nodes, and Conditional Edge Routers.
- **Persistent Execution Layers:** Configuring thread checkpointers and state management across user sessions using localized SQLite and Redis caching layers.
- **Human-in-the-Loop (HITL) Architectures:** Injecting mandatory application manual verification breaks, execution rollbacks, and runtime parameter overrides.

Module Milestone Deliverable: Construct a highly predictable, state-tracked e-commerce backend processing graph with full transaction state recovery thresholds.

Module 3: Enterprise Tool Integration via Model Context Protocol (MCP)

- **The MCP Core Framework:** Understanding the standardized JSON-RPC web transport architecture separating client context from hosted tool environments.
- **Custom MCP Server Implementation:** Building backend server capabilities to safely broadcast infrastructure properties, schemas, and processing endpoints.
- **Sandboxed Resource Isolation:** Binding agents safely to production filesystems, historical relational schemas, and isolated version control repositories.

Module Milestone Deliverable: Code a secure file-and-database transformation MCP server allowing external models to execute safely isolated script migrations.

Module 4: Multi-Agent Collaboration, Swarms, and Orchestration

- **Hierarchical vs. Choreographed Layouts:** Building distinct manager/worker role patterns utilizing CrewAI.
- **Event-Driven Communication:** Setting up asynchronous inter-agent conversation messaging streams using Microsoft AutoGen.
- **Handoff & Resolution Policies:** Handling consensus anomalies, dynamic role handoffs, and token payload preservation across network components.

Module Milestone Deliverable: Deploy a multi-agent scrum swarm (Product Manager, Engineer, and Quality Assurance nodes) to autonomously execute complete software development lifecycles.

Module 5: Adaptive Knowledge Engineering & GraphRAG Architectures

- **Advanced Retrieval Frameworks:** Implementing Self-RAG pipelines and Corrective Retrieval-Augmented Generation (CRAG) checking loops.
- **Multi-Hop Knowledge Engineering:** Structuring deeply nested corporate datasets using Neo4j integrated with optimized semantic vector weights.
- **Document Parsing & Reranking:** Processing heavy structural tables, raw text chunks, and configuring cross-encoder reranking algorithms.

Module Milestone Deliverable: Build an analytical financial intelligence agent optimized to perform cross-document relationship discovery over multi-page PDF records.

Module 6: AgentOps, Security Guardrails, and Production Scalability

- **Deep Telemetry Tracking:** Unpacking complex trace timelines, node latencies, and total token usage profiles inside LangSmith or LangFuse environments.
- **Automated Regression Testing:** Designing standardized evaluation matrix runs utilizing modern LLM-as-a-Judge validation systems.
- **Defensive Engineering & Safety Guardrails:** Creating execution perimeters against prompt injection vectors, target hallucinations, PII data leakage, and tool abuse.
- **Infrastructure Deployment:** Containerizing autonomous agent setups using Docker and configuring scaling strategies across Kubernetes clusters.

Module Milestone Deliverable: Harden, package, and deploy an enterprise-grade agent infrastructure cluster with automated load testing and trace reporting.