


Enterprise Agentic AI Engineering

A comprehensive 3-day, 24-hour advanced technical course taking engineers, architects and advanced learners from LLM-powered applications to secure, orchestrated enterprise AI agents — with hands-on labs at every stage.

3 DAYS · 24 HOURS

HANDS-ON LABS

ENTERPRISE FOCUS



What You Should Know Before This Course

GenAI & LLM Fundamentals

Understand LLM capabilities, prompts, context and basic model/API use

RAG Fundamentals


Understand embeddings, vector databases and the basic RAG pipeline

Programming & APIs

Comfortable with Python and basic API / function-calling concepts

Software & Data Concepts

Basic application architecture, databases and enterprise workflows



Who This Course Is For

Software & AI Engineers

Build, orchestrate and evaluate enterprise AI agents

Solution & Technical Architects


Design agentic architectures, integration patterns and controls

Data / Platform Engineers

Connect agents to enterprise data, APIs and operational systems

Advanced Students & Technical Leads

Apply GenAI + RAG foundations to real-world agentic solutions



The Learning Journey

Three progressive days — from single-agent foundations to enterprise multi-agent systems.

Day 1 – Agent Foundations

From LLM applications to agents that can reason, retrieve and use tools

Agents & tool calling

Memory, state & planning

Agent frameworks

Build a working single agent

1. Enterprise RAG Architecture
2. Data Ingestion & Preparation
3. Chunking & Embeddings
4. Vector Databases
5. Retrieval Engineering
6. Build a Basic Enterprise RAG

Day 2–3 – Orchestration & Enterprise

From workflows to secure, production-oriented agentic AI

LangGraph & CrewAI

MCP & A2A concepts

Multi-agent systems

Evaluation, security, governance & observability

Capstone project

1. Advanced RAG Techniques
2. RAG Evaluation & Optimization
3. Security & Governance
4. Enterprise Architecture Patterns
5. Deployment & Observability
6. Capstone Project

Day 1 – Building the Agentic AI Foundation

Participants progressively build a working agent through guided labs.

01

Agentic AI Architecture

Agent vs chatbot vs workflow; agent lifecycle, planning and action

03

Memory, State & Planning

State, task decomposition, planning patterns and memory concepts

05

Agent Workflows

Routing, branching, retries, human approval and state graphs

02

Agents & Tool Use

Function calling, APIs, tool definitions and controlled execution

04

Agent + RAG


Grounded agents, retrieval, context management and knowledge access

06

Baseline Agent Application


Connect an agent to tools and knowledge; evaluate its task execution

Day 1 – The Agent Architecture You'll Build



By end of Day 1, participants have a functional agent that can retrieve enterprise knowledge, call controlled tools and execute a defined workflow.

Day 2 – Agent Orchestration & Engineering



Agent Orchestration

- State graphs & workflow routing
- Retries, branching & human-in-the-loop
- Task decomposition & delegation
- Reliable tool execution

Frameworks & Connectivity

- LangGraph & CrewAI patterns
- MCP: tools, resources & enterprise context
- A2A concepts & agent communication

Multi-Agent Systems

- Specialized agents & delegation
- Collaboration and communication
- Supervisor / worker patterns

Multi-agent evaluation

Enterprise Agent Security & Governance

Treated as a core engineering module — not an optional discussion.

Agent Access Control

Identity, authentication, authorization, least privilege, tool permissions and human approval

AI & Tool Security

Prompt injection, indirect injection, data exfiltration, unsafe tool calls and secrets management

Governance & Compliance

Audit trails, agent decisions, evaluation records, model governance and accountability

Responsible AI

Privacy, sensitive information controls, transparency, human oversight and safe autonomy

Enterprise Agentic AI Deployment Architecture

Participants understand the evolution from prototype agent to observable, governed enterprise agent platform.



Key areas: APIs, containers, cloud deployment, identity and permissions, tracing, evaluation pipelines, monitoring, high availability, cost management and controlled autonomy.

Enterprise Agentic AI Capstone

What You'll Build

An enterprise-grade multi-agent application covering the full lifecycle:

- Define business task and agent roles
- Connect RAG knowledge and enterprise tools
- Orchestrate single / multi-agent workflows
- Apply guardrails, access control and human approval
- Evaluate task success, safety, latency and cost
- Expose the solution through an API with monitoring/observability

Suggested Tracks

Digital Engineering

Coding, testing, code review,
DevOps or IT operations agents

Data Management

Data quality, metadata, ETL/ELT,
migration or governance agents

Business / Industry

Insurance, banking, healthcare, capital markets, HR, regulatory



Technology Stack

Architecture-first and implementation-practical. Participants work with a modern agent stack while learning transferable patterns.



Core Development

Python · LLM APIs · LangGraph · CrewAI



Agent & Knowledge

RAG · Vector databases · Memory / state · Tool calling



Enterprise Integration

MCP · REST APIs · FastAPI · Docker · Cloud services · Observability



Evaluation & Governance

Agent evaluation · Guardrails · Tracing · Security · Human-in-the-loop

By the End of This Course, You Will Be Able To

Participants complete the program with practical Agentic AI engineering capability. **Certification**

Design Agentic AI Systems

- Architect agents, tools, RAG, memory, orchestration and enterprise controls
- Build & Orchestrate Agents
- Implement single-agent and multi-agent workflows
- Integrate Enterprise Tools
- Connect agents to APIs, data, applications and MCP resources
- Evaluate & Secure Agents

Apply evaluation, guardrails, security, governance and human oversight

Panaashe Certified Agentic AI Engineer

Awarded upon successful completion of hands-on labs, practical assessment and the end-to-end capstone project.

Data Management

Data quality, metadata, ETL/ELT, migration or governance agents

Business / Industry

Insurance, banking, healthcare, capital markets, HR, regulatory