

To discuss this course and customizations:
Call: +1 434-509-6890 or Email: sales@cloudcontraptions.com

Designing Multi-Agent Systems

Class Duration

14 hours of live training delivered over 2-3 days to accommodate your scheduling needs.

Student Prerequisites

- Professional software development experience in Python or TypeScript
- Familiarity with LLM API usage and function/tool calling

Target Audience

Software engineers and architects designing or building systems where multiple AI agents collaborate on complex, long-horizon tasks. Relevant for teams building internal automation platforms, enterprise AI assistants, or agentic pipelines that go beyond single-agent workflows.

Description

Multi-agent systems unlock capabilities beyond single-agent limits - but they introduce new challenges in orchestration, state management, error recovery, and observability. This course covers the architectural patterns and practical techniques for building reliable multi-agent systems: planner/worker decomposition, agent hand-off protocols, shared memory and state passing, failure detection and recovery, evaluation, and the safety considerations unique to autonomous agent collaboration. Labs build progressively more complex multi-agent pipelines using TypeScript and Python with real model backends.

Learning Outcomes

- Describe the key multi-agent architectural patterns: orchestrator, planner/worker, pipeline, and network topologies.
- Implement agent hand-off with clear task scope, context transfer, and completion signaling.
- Design shared memory and state stores for multi-agent coordination.
- Apply failure detection, retry, and escalation patterns to agent pipelines.

To discuss this course and customizations:
Call: +1 434-509-6890 or Email: sales@cloudcontraptions.com

- Build a planner/worker system with dynamic task decomposition.
- Evaluate multi-agent system behavior using trace analysis and task-completion metrics.
- Apply safety boundaries: capability scoping, confirmation gates, and human-in-the-loop escalation.

Training Materials

Comprehensive courseware is distributed online at the start of class. All students receive a downloadable MP4 recording of the training.

Software Requirements

Python 3.12+ or Node.js 22+, API keys for at least one frontier model, and Git.

Training Topics

Why Multi-Agent Systems

- Task classes that benefit from multiple agents
- Limits of single-agent context and capability
- Tradeoffs: complexity, cost, and latency

Architectural Patterns

- Orchestrator/worker pattern
- Planner/executor decomposition
- Pipeline (sequential) agents
- Peer/network agent collaboration
- Choosing the right topology

Agent Hand-Off Protocols

- Task scope and acceptance criteria definition
- Context package design for hand-offs
- Completion signaling and result validation
- Partial completion and resumption

Shared Memory and State

- In-process vs. external state stores
- Shared context formats and schemas
- Concurrent write safety

To discuss this course and customizations:
Call: +1 434-509-6890 or Email: sales@cloudcontraptions.com

- Memory pruning for long-running systems

Failure Detection and Recovery

- Detecting stuck, looping, or incorrect agents
- Retry strategies per agent type
- Escalation to human-in-the-loop
- Graceful degradation when an agent fails

Dynamic Task Decomposition

- Planner agent design: input → task graph
- Dependency resolution and parallel dispatch
- Handling plan revisions mid-execution
- Task graph visualization and debugging

Evaluation and Observability

- Tracing multi-agent execution end-to-end
- Task-completion metrics and success criteria
- Intermediate step quality evaluation
- Cost attribution across agent roles

Safety Boundaries

- Capability scoping per agent role
- Confirmation gates for destructive actions
- Human-in-the-loop escalation triggers
- Audit logs for fully autonomous pipelines