



Agentic AI Systems Design and Governance Training Course

2 – 6 May 2027
Muscat

Agentic AI Systems Design and Governance Training Course

Ref: 190_117383 **Date:** 2 – 6 May 2027 **Location:** Muscat **Fees:** 5700 Euro

Overview

Agentic AI Systems Design and Governance Training Course is a five-day advanced course for AI product owners, automation specialists, solution architects, business analysts, and technical leads, who leave with an Agentic AI Solution Blueprint. Participants define agentic AI use-case boundaries, design tool contracts and memory strategies, select orchestration patterns, establish human approval gates, and apply agent evaluation and observability controls. Agile Leaders Training Center delivers training in governed agentic AI systems design.

Who Should Attend

- AI product personnel responsible for use-case scope, value, and acceptance criteria
- Automation personnel responsible for workflow state, tools, and exception paths
- Solution-design personnel responsible for architecture, integration, and permissions
- Business-analysis personnel responsible for goals, roles, rules, and process evidence
- Technical leaders responsible for evaluation, oversight, security, and operational accountability

The course assumes participants can analyze digital workflows and solution requirements at work, and it leaves out programming bootcamps, vendor configuration, and autonomous deployment without oversight.

Departments and Industries

The course supports functions designing or governing agentic workflows across financial services, healthcare administration, professional services, technology, logistics, and public services.

- Artificial intelligence product and solution design
- Business process automation and operations
- Data, analytics, and digital transformation
- Information security and technology risk
- Service management and quality assurance

Learning Objectives

By the end of this course, participants will be able to:



- Prioritize agentic use cases by value, autonomy, and risk
- Build tool, memory, context, and permission specifications
- Compare workflow and agent orchestration patterns
- Analyze state, failure, retry, and recovery behavior
- Evaluate agents with task sets, traces, and graders
- Apply human oversight, security, and lifecycle controls

Course Agenda

Day 1: Use Cases and Autonomy Boundaries

- Agentic Use-Case Suitability Canvas
- Workflow-versus-Agent Decision Matrix
- Goal, Role, and Success-Criteria Charter
- Autonomy Boundary and Consequence Map
- Human Decision Rights Register

Day 2: Tools Memory and Context

- Agent Tool Contract Specification
- Least-Privilege Permission Matrix
- Context Source and Retrieval Map
- Working and Persistent Memory Strategy
- Prompt-Injection Exposure Checklist

Day 3: Orchestration State and Reliability

- Sequential, Parallel, and Routing Pattern Comparison
- Orchestrator-Worker Responsibility Diagram
- Evaluator-Optimizer Feedback Loop
- Workflow State and Checkpoint Model
- Retry, Timeout, and Compensating-Action Table

Day 4: Evaluation Security and Operations

- Agent Task Set and Acceptance Rubric
- Trajectory Trace and Tool-Call Review
- Operational Observability Signal Map
- Agent Security Threat Model
- Incident, Escalation, and Controlled-Shutdown Route



Day 5: Agentic Design Practice and Capstone

- Suggested Exercise: Challenge an Agentic Use-Case Boundary
- Suggested Exercise: Review a Tool Permission and Approval Gate
- Suggested Exercise: Diagnose a Failed Agent Trajectory
- Suggested Exercise: Test an Incident and Recovery Decision
- Capstone Exercise: Agentic AI Solution Blueprint

Practical Exercises

The course uses suggested activities that turn agentic design choices into reviewable solution evidence.

- Suggested activity: classify tasks as deterministic workflows, agentic workflows, or unsuitable automation
- Suggested activity: specify tools, permissions, context, memory, and mandatory approval points
- Suggested activity: compare orchestration patterns against state, latency, and failure constraints
- Suggested activity: score a multi-step trajectory and document an operational response

FAQs

Who suits the Agentic AI Systems Design and Governance Training Course, and what does it assume?

The course suits experienced product, automation, architecture, analysis, and technical personnel who can examine workflow requirements and need to design governed agents without building an application in code.

How does agentic AI systems design training differ from generative AI fundamentals training?

Agentic AI systems design training focuses on multi-step goals, tool actions, memory, workflow state, orchestration, evaluation, and operational controls rather than prompt-response concepts alone.

When should an agentic AI workflow require human approval?

Human approval should be mandatory when an action has material, irreversible, sensitive, or uncertain consequences, with the workflow preserving state and evidence so review can occur before execution continues.

How are agentic AI systems evaluated?

Agentic AI systems are evaluated with representative task sets, environment state, tool-call traces, outcome checks, behavioral graders, failure cases, and repeated tests that reveal variation across multi-step trajectories.



What security controls matter in agentic AI systems design?

Relevant controls limit tool permissions and data access, separate trusted instructions from untrusted content, validate actions, protect credentials, record traces, require approvals, detect abnormal behavior, and support suspension or controlled shutdown.

Conclusion

Participants take back an Agentic AI Solution Blueprint that connects business scope with architecture, tools, memory, state, permissions, evaluation, and oversight. It makes agentic design decisions reviewable before implementation or procurement. The blueprint supports clearer approvals, safer failure handling, operational monitoring, incident escalation, and controlled lifecycle change.