JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Agentic AI Builder: Designing and Testing AI Agents Course

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

Agentic AI Builder: Designing and Testing AI Agents Course

Agentic AI Builder: Designing and Testing AI Agents is a 1-week training course for business analysts, automation leads, digital transformation and IT professionals, and each participant leaves with an Agent Prototype Pack: a working agent, its test set and its guardrail design. AI agents that plan steps, call business systems and finish tasks unaided are moving from demonstrations into operations. The course is delivered by Jupiter Prime Training.

Why the Agentic AI Builder Matters Now

Two forces shape agent building today. Large language models can now call APIs, query databases and fill forms, so low-code platforms let analysts assemble agents without large engineering teams. At the same time, security and risk functions worry about agents that act on wrong information, leak data or follow injected instructions. Teams that build agents with clear scopes, tested tools and guardrails automate multi-step work such as claims triage or supplier onboarding; teams that skip design and evaluation ship brittle agents that erode trust. The course uses the NIST AI Risk Management Framework AI RMF and ISO/IEC 42001 to frame agent risk controls.

Who Gains Most from the Agentic AI Builder

- Analysts who map business processes and specify where an agent should act
- Automation specialists who already run robotic process automation or workflow tools
- Digital transformation practitioners who turn use cases into working prototypes
- IT professionals who connect agents to APIs, identity and data stores
- Quality and risk practitioners who test agent behaviour before release

An automation centre of excellence gains most when analysts, integrators and testers attend together and build one agent end to end.

Agentic AI Builder Capabilities

Participants practise each stage of the agent build cycle, from scoping a task to proving the agent behaves safely under test.

- Specify an agent's goal, scope, tools and stopping conditions — process owners
- Design a planner-executor or single-loop architecture for a multi-step task — solution architects
- Build tool calls to business systems with clear input and output schemas — integration teams
- Apply memory and retrieval so the agent uses current, approved knowledge — content owners
- Test agent behaviour with an evaluation set and adversarial prompts — risk and quality reviewers
- Communicate prototype results and release conditions — automation steering groups

Key Agentic AI Builder Terms

- AI agent — software that uses a language model to plan steps, call tools and pursue a goal with limited supervision.
- Tool calling — the model's ability to request a defined function, such as a database query or API action, with structured inputs.
- Agent memory — stored context the agent reads between steps or sessions, such as conversation history or retrieved documents.
- Guardrail — a rule or check that blocks, edits or escalates an agent action outside approved limits.
- Prompt injection — hidden instructions in content the agent reads that try to change its behaviour or extract data.
- Evaluation set — a fixed collection of tasks with expected outcomes used to measure agent accuracy and safety before release.

Agentic AI Builder Day by Day

Day 1: From Chatbots to Agents

- Agent Versus Chatbot Versus RPA — Choosing the Right Automation Type
- Task Decomposition Canvas — Breaking a Process Into Agent Steps
- Insurance Claims and Procurement Examples — Where Agents Fit
- Agent Scope and Stopping Conditions — Limiting Autonomy Safely
- Low-Code Agent Platform Tour — Selecting a Build Environment

Day 2: Architectures and Tools

- Single-Loop and Planner-Executor Patterns — Matching Architecture to Task
- Tool Definitions and Function Schemas — Reliable Calls to Systems
- API Connectors and Authentication — Least-Privilege Access for Agents
- Multi-Agent Hand-Offs — When Specialist Agents Should Collaborate
- System Prompt Design for Agents — Role, Rules and Output Format

Day 3: Memory, Knowledge and Guardrails

- Short-Term and Long-Term Agent Memory — What to Store and Forget
- Retrieval From Approved Knowledge Bases — Grounding Agent Answers
- Input and Output Guardrails — Blocking Unsafe or Off-Scope Actions
- Human Approval Steps — Escalation Before Irreversible Actions
- Prompt Injection Defences — Sanitising Untrusted Content

Day 4: Testing and Release Readiness

- Evaluation Set Construction — Normal, Edge and Adversarial Tasks
- Trace Logging and Step Inspection — Debugging Agent Decisions
- Accuracy, Cost and Latency Metrics — Release Thresholds for Agents
- NIST AI RMF Measure Function — Documenting Agent Test Evidence
- Monitoring Plan for Live Agents — Drift, Failures and Rollback

Day 5: Practising the Agent Prototype Pack

- Build Sprint on a Chosen Use Case — Applying Architecture and Tools
- Red-Team Exercise — Testing Guardrails Against Injection Attempts
- Evaluation Run and Score Review — Deciding Readiness
- Peer Demonstration — Explaining Agent Behaviour to Process Owners
- Agent Prototype Pack — Final Agent, Tests and Guardrail Design

Applying the Agentic AI Builder After the Course

- Rebuild the course prototype against one internal process and run its evaluation set with real cases.
- Present trace logs and guardrail results to the automation steering group to agree release conditions.

Agentic AI Builder Questions FAQ

What does an agentic AI builder do and why does it matter now?

An agentic AI builder designs, connects and tests AI agents that plan and complete multi-step tasks. It matters now because low-code platforms make agents easy to assemble, and organisations need people who can make them reliable and safe.

Which roles gain most from agentic AI builder training?

Business analysts, automation leads, digital transformation practitioners and IT integrators gain most. Quality and risk testers also benefit, especially when a centre of excellence sends a mixed group.

How does agentic AI builder training differ from a strategic agentic AI course?

It is hands-on: participants build, connect and test an agent. A strategic agentic AI course looks at operating models and business cases for leaders; this agentic AI builder course focuses on architecture, tool calls, guardrails and evaluation.

What do you bring back from agentic AI builder training?

You bring back a working prototype pack. It holds the agent configuration, tool definitions, guardrail design, evaluation set with scores and a short release recommendation.

Agentic AI Builder: The Way Forward

The Agent Prototype Pack first informs whether one chosen process is ready for an agent pilot in a controlled environment. Analysts, integrators and risk reviewers use it to agree tools, permissions and release thresholds. After the first live run, the evaluation set and guardrail rules deserve another look, since real requests expose edge cases no test anticipated.