JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Agentic AI Governance and Control Training Course

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

Agentic AI Governance and Control Training Course

Agentic AI Governance and Control is a 1-week training course for risk, compliance and IT governance professionals who must keep autonomous AI agents within agreed authority, and who leave with an agent control register. Organisations now deploy agents that send messages, move data and trigger transactions with little human input. Jupiter Prime Training runs this course for control functions setting the rules before those agents scale.

Why Agentic AI Governance and Control Matters Now

Three forces shape this subject today: agent platforms that let business teams launch autonomous workflows, auditors and boards asking who is accountable when software acts on its own, and attackers targeting agents through manipulated inputs. An organisation that governs agents well gains speed without losing control over payments, data and customer commitments; one that does not discovers unauthorised actions only after damage occurs. The course rests on ISO/IEC 42001, the international standard for an artificial intelligence management system, and the NIST AI Risk Management Framework NIST AI RMF for mapping and managing agent risk.

Who Gains Most from Agentic AI Governance and Control

- Risk specialists who assess new technologies and maintain the risk register
- Compliance officers who confirm automated actions follow policy and contractual limits
- IT governance leads who manage access rights, service accounts and change approval
- Internal auditors who need evidence trails for automated decisions
- Information security staff who defend systems against misuse of agent tools

Control functions gain most when risk, compliance and IT governance colleagues attend together, since agent controls span all three.

Agentic AI Governance and Control Capabilities

Participants leave able to classify, limit, watch and audit autonomous agents within approved boundaries.

- Assess agent risk using autonomy, impact and reversibility scores — risk committees approving deployments
- Design delegation of authority and least-privilege permissions for agents — process and system owners
- Specify human oversight checkpoints for high-impact actions — business approvers
- Build monitoring and audit trail requirements that reconstruct agent behaviour — internal and external auditors
- Test incident playbooks including kill switches and containment — security operations

Key Agentic AI Governance and Control Terms

- Agentic AI — AI systems that plan and carry out multi-step tasks, calling tools and systems with limited human direction.
- Delegation of authority — the documented limits on what decisions or actions a person or agent may take without further approval.
- Least privilege — granting an agent only the minimum access and actions needed for its assigned task.
- Human oversight checkpoint — a defined point where a person must review or approve an agent's proposed action before it proceeds.
- Audit trail — a tamper-evident record of an agent's inputs, decisions, tool calls and outcomes for later review.
- Kill switch — a control that immediately halts an agent's activity when behaviour becomes unsafe or unauthorised.

Agentic AI Governance and Control Day by Day

Day 1: How AI Agents Change the Control Environment

- Agent Anatomy — Goals, Planning Loops, Memory and Tool Calls
- Autonomy Levels Scale — From Suggestion to Unattended Execution
- Agent Inventory and Ownership — Knowing Which Agents Act for the Organisation
- Failure Modes of Autonomous Agents — Goal Drift, Looping and Overreach
- Agent Use Cases in Customer Service and Procurement

Day 2: Delegation of Authority and Permissions

- Delegation of Authority Matrices — Spending, Approval and Communication Limits
- Least-Privilege Tool Access — Scoping APIs, Systems and Actions
- Data Permission Boundaries — Classifying What Agents May Read or Write
- Identity and Credentials for Agents — Service Accounts and Secret Handling
- Segregation of Duties Between Agents and Human Approvers

Day 3: Agent Risk Assessment Against Standards

- ISO/IEC 42001 Controls — Placing Agents Inside the AI Management System
- NIST AI RMF Functions — Govern, Map, Measure and Manage for Agentic Systems
- Agent Risk Scoring — Impact, Autonomy and Reversibility of Actions
- Threat Scenarios — Prompt Injection, Tool Misuse and Data Exfiltration
- Third-Party Agent Platforms — Vendor Assurance Questions

Day 4: Oversight, Monitoring and Incident Handling

- Human Oversight Checkpoints — Approval Gates for High-Impact Actions
- Runtime Guardrails — Rate Limits, Budget Caps and Action Allow-Lists
- Behaviour Monitoring Dashboards — Anomalies, Escalations and Override Counts
- Audit Trails and Action Logs — Reconstructing What an Agent Did and Why
- Agent Incident Playbooks — Kill Switches, Containment and Root Cause Review

Day 5: Practising the Agent Control Register

- Autonomy and Risk Scoring Applied — Classifying Agents from Participants' Organisations
- Permission and Delegation Design — Setting Limits for One Selected Agent
- Oversight Checkpoint and Monitoring Plan — Combining Earlier Control Methods
- Audit and Incident Readiness Test — Walking Through a Simulated Agent Failure
- Agent Control Register — Completing and Presenting the Final Document

Applying Agentic AI Governance and Control After the Course

- Build an inventory of agents in use across the organisation and score each one for autonomy and impact
- Introduce approval checkpoints and action logging for the highest-risk agent and review them with its business owner

Agentic AI Governance and Control Questions FAQ

What is agentic AI governance and control, and why does it matter now?

It is the set of policies, permissions, oversight and audit controls that keep autonomous AI agents within approved authority. It matters now because agents increasingly act directly on systems, data and customers with minimal human review.

Which roles gain most from agentic AI governance and control?

Risk, compliance and IT governance professionals gain most. You approve technology risk, confirm policy adherence or manage access, so you need practical ways to classify, limit and audit agent behaviour.

How does agentic AI governance and control training differ from a general AI governance course?

It focuses on systems that act, not only systems that advise. A general course covers policies and model risk broadly, whereas this subject addresses delegation limits, tool permissions, runtime guardrails, audit trails and incident handling for autonomous agents.

What do you bring back from agentic AI governance and control training?

You bring back a control register for agents in your organisation. It records autonomy and risk scores, permissions, oversight checkpoints, monitoring requirements and incident steps for each agent.

Agentic AI Governance and Control: The Way Forward

The agent control register first informs which agents may run unattended and which need approval gates. Risk, compliance and IT governance teams use it to agree permissions, monitoring and audit evidence with business owners. After the first quarter, revisit autonomy scores and override counts, because real usage shows where limits are too tight or too loose.