JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

AI Systems Architecture and Governance Training Course

14 – 18 December 2026
Geneva

AI Systems Architecture and Governance Training Course

Ref: 46_1164 **Date:** 14 – 18 December 2026 **Location:** Geneva **Fees:** 7500 GBP

AI Systems Architecture and Governance is a 1-week training course for enterprise architects, AI platform owners, IT leaders and technology risk managers that ends with an AI reference architecture and control map for one production system. As machine learning moves from notebooks into core platforms, the way pipelines, feature stores and model endpoints are wired decides how safely each model can be changed or audited. Jupiter Prime Training delivers this AI Systems Architecture and Governance course.

Why AI Systems Architecture and Governance Matters Now

Three forces push AI systems architecture and governance up the agenda: model portfolios growing from a handful of experiments to dozens of production endpoints, security teams facing new attack paths such as prompt injection and data poisoning, and auditors asking for evidence of control at every stage of the model lifecycle. An organisation that designs controls into its platform can retrain, roll back and explain models quickly; one that bolts them on later inherits brittle integrations and blind spots. The course draws on ISO/IEC 42001, the AI management system standard, and the NIST AI Risk Management Framework NIST AI RMF.

Who Gains Most from AI Systems Architecture and Governance

- Professionals who own target-state designs for data platforms, integration layers and model hosting
- Leaders accountable for AI platform budgets, vendor choices and service levels
- Engineers who run MLOps pipelines, model registries and deployment automation
- Specialists who assess cyber threats to models, training sets and inference endpoints
- Risk and assurance staff who test whether AI controls operate as designed

A central technology function gains most when its architecture, platform engineering and technology risk members attend together and agree one set of control points.

AI Systems Architecture and Governance Capabilities

Participants learn to read an AI platform as layers of components and to place a control, an owner and a piece of evidence on each layer.

- Design a layered reference architecture for training, serving and monitoring models — architecture review boards
- Specify data lineage, feature store and model registry requirements — platform owners
- Build MLOps release gates with approval, testing and rollback steps — engineering leads
- Assess threats to models and inference endpoints using structured threat modelling — security officers
- Align lifecycle controls with ISO/IEC 42001 clauses and NIST AI RMF functions — internal auditors

- Calculate monitoring thresholds for drift, latency and cost — service owners and finance partners

Key AI Systems Architecture and Governance Terms

- Reference architecture — a reusable blueprint showing the standard components, interfaces and control points for building AI systems across an organisation.
- Model registry — a versioned catalogue of trained models, their lineage, approvals and deployment status.
- MLOps — practices and automation that move models reliably from experiment to production and keep them maintained.
- Feature store — a shared service that stores, documents and serves model input variables consistently for training and inference.
- Model drift — the gradual loss of model accuracy as live input data or real-world relationships change.
- Control map — a matrix linking each lifecycle stage to its controls, owners and evidence for assurance.

AI Systems Architecture and Governance Day by Day

Day 1: Platform Landscape and Lifecycle Stages

- AI Lifecycle Stage Model — Where Architecture Choices Create Risk
- Current-State Platform Inventory — Duplicated Pipelines and Shadow Models
- Build, Buy or Managed Service Comparison — Hosting Strategy for Models
- Architecture Principles for AI Workloads — Portability, Traceability and Cost
- ISO/IEC 42001 Management System Scope — Boundaries of the AI Estate

Day 2: Data Pipelines and Feature Engineering Layers

- Ingestion and Transformation Pipeline Patterns — Batch Versus Streaming Choices
- Data Lineage and Provenance Tracking — Audit Trail for Training Sets
- Feature Store Design — Consistency Between Training and Inference
- Data Quality Rules and Contracts — Accepting or Rejecting Upstream Feeds
- Privacy-Preserving Techniques — Masking and Minimisation in Pipelines

Day 3: Model Serving, MLOps and Security

- Model Registry and Versioning — Approval Before Promotion to Production
- CI/CD Release Gates for Models — Testing, Sign-Off and Rollback
- Serving Patterns: Batch, Real-Time and Edge — Latency and Cost Trade-Offs
- Threat Modelling for AI Endpoints — Poisoning, Evasion and Prompt Injection
- Access Control and Secrets Handling — Protecting Model Artefacts

Day 4: Monitoring, Assurance and Accountability

- Drift, Latency and Cost Dashboards — Thresholds for Retraining
- NIST AI RMF Govern, Map, Measure, Manage — Control Ownership
- Model Cards and System Documentation — Evidence for Auditors
- Incident Response for Model Failures — Escalation and Kill Switches
- Third-Party Model and API Due Diligence — Supplier Assurance

Day 5: Practising the Reference Architecture and Control Map

- Lifecycle Stage Model From Day 1 — Scoping One Production System
- Pipeline and Feature Layers From Day 2 — Lineage Checkpoints
- Release Gates and Threat Model From Day 3 — Security Controls
- Monitoring Thresholds From Day 4 — Assurance Evidence Schedule
- AI Reference Architecture and Control Map — Board-Ready Design Pack

Applying AI Systems Architecture and Governance After the Course

- Map one live model, such as a credit scoring or demand forecasting service, against the lifecycle stage model and record missing lineage or release gates.
- Run a threat modelling session on a customer-facing chatbot or inference endpoint and add the agreed mitigations to the platform backlog.

AI Systems Architecture and Governance Questions FAQ

What is AI systems architecture and governance and why does it matter now?

AI systems architecture and governance is the design of AI platforms so that every pipeline, model and endpoint carries clear controls and owners. It matters now because model numbers are growing fast and auditors and security teams expect traceable evidence across the lifecycle.

Which roles gain most from AI systems architecture and governance training?

Enterprise architects, AI platform owners, MLOps engineers and technology risk managers gain most. Each sees how design choices in pipelines, registries and serving layers either create or remove control gaps that others must later manage.

How does AI systems architecture and governance differ from an AI policy or ethics course?

It works at platform and component level rather than policy level. You design registries, release gates, monitoring thresholds and threat models, and attach governance controls to them, instead of drafting principles or committee charters.

What do you bring back from AI systems architecture and governance training?

You bring back an AI reference architecture and control map for one production system. It shows components, data flows, control points, owners and evidence, ready for an architecture review board or audit committee.

AI Systems Architecture and Governance: The Way Forward

The AI reference architecture and control map informs the first platform decision: which control gaps to close before the next model release. Architects, platform owners and auditors use it as a shared design record. After the first release under the new gates, revisit drift thresholds and third-party model checks, since live traffic and supplier updates often expose weaknesses the design did not foresee.