
Designing Governed AI Platforms: Controls, Risk and Assurance Training Course

19 – 23 October 2026
Geneva





Designing Governed AI Platforms: Controls, Risk and Assurance Training Course

Ref: 249_18955 Date: 19 – 23 October 2026 Location: Geneva Fees: 5500 GBP

Programme Summary

Designing governed AI platforms is the subject of this one-week training course for senior technology and risk leaders at advanced level, culminating in a governed AI platform blueprint with a control and assurance plan. Organisations are putting models into production faster than they can inventory, approve and monitor them. The approach moves from platform architecture and risk, through control design across the data and model lifecycle, to runtime monitoring, assurance and a worked platform design. Mars Training Centre UK delivers this programme in designing governed AI platforms.

The Business Case

AI systems built on separate tools leave leaders unable to say which models are live, who approved them and what evidence shows they behave as intended. A platform with controls designed in lets teams release faster, because approvals, testing and monitoring become routine, and gives boards a basis for challenge.

- Model Release Cycle Time: Shortens as standard approval gates replace project-by-project negotiation.
- Unmanaged AI Exposure: Falls as an inventory and risk tiering expose systems that run outside oversight.
- Assurance Effort: Reduces as control evidence is collected by the platform rather than assembled before each review.

Principles and Evidence Base

- ISO/IEC 42001: Its AI management system requirements structure the roles, policies and controls designed in Sessions 1 and 3.
- ISO/IEC 23894: Its AI risk management process shapes the risk identification, rating and treatment practised in Session 2.
- NIST AI Risk Management Framework: Its Govern, Map, Measure and Manage functions organise the monitoring and assurance design in Session 4.

Intended Participants

- Executives accountable for the technology platforms on which AI systems are built and run.
- Leaders who set enterprise architecture standards for data, models and integration.
- Heads of risk and control who must evaluate and challenge AI deployments.
- Digital and transformation directors who sponsor AI use cases across business units.
- Senior assurance and compliance leaders who report AI risk to boards and committees.

The programme serves banking, energy, healthcare, government and telecommunications organisations.



Learning Outcomes

1. Participants will be able to construct a reference architecture for an AI platform that shows where each governance control sits.
2. Participants will be able to classify AI use cases by risk tier and justify the oversight each tier requires.
3. Participants will be able to diagnose gaps in data, model and release controls using a lifecycle control map.
4. Participants will be able to design runtime guardrails and monitoring thresholds for live AI systems.
5. Participants will be able to evaluate the quality of control evidence presented for assurance reviews.
6. Participants will be able to prioritise platform investments and ownership for a governed AI platform roadmap.

Programme Structure

Day 1: Governed AI Platform Foundations

Which parts of an AI platform need governing, and who decides?

- AI Platform Reference Architecture Layers
- ISO/IEC 42001 Management System Structure
- AI System Inventory and Classification Scheme
- Decision Rights Matrix for AI Approvals
- Platform Governance Maturity Self-Assessment

Day 2: AI Risk Diagnosis and Treatment

How is AI risk identified, rated and treated across a platform?

- ISO/IEC 23894 Risk Management Process
- NIST Govern, Map, Measure and Manage Functions
- Risk Tiering of Candidate Use Cases
- Third-Party and Foundation Model Risk Review
- Risk Register Construction for AI Systems

Day 3: Data and Model Lifecycle Controls

Which controls belong at each stage from data intake to retirement?

- Data Lineage and Quality Gates
- Model Documentation and Model Cards
- Validation, Bias and Robustness Testing
- Release Approval Workflow Design
- Change Control and Model Retirement Rules



Day 4: Runtime Controls, Monitoring and Assurance

How is a live AI system kept inside its approved limits?

- Inference Gateway and Access Controls
- Prompt and Output Guardrails
- Drift, Performance and Incident Monitoring
- Control Testing and Evidence Collection
- Assurance Reporting to Boards and Committees

Day 5: Applying Governed AI Platform Design

Is the platform blueprint ready for sponsor approval?

- Applying the Session 1 Inventory to a Chosen Platform
- Applying Session 2 Risk Tiers to Priority Use Cases
- Mapping Session 3 and 4 Controls to Platform Layers
- Sequencing Investment, Ownership and Delivery Phases
- Completing the Governed AI Platform Blueprint

Methodology

The programme is built around structured design exercises and case analysis. Participants work on a sample enterprise AI estate, drawing architecture views and specifying controls, then defend their design to a simulated risk board. Cases come from financial services and health settings.

- Suggested activity: Building an AI inventory from a set of fictional project descriptions.
- Suggested activity: Tracing a model from data intake to retirement, marking the control at each step.
- Suggested activity: Presenting the platform blueprint to a simulated architecture and risk board.

Capstone Output

Each participant completes a governed AI platform blueprint for one organisation or business unit. It contains an architecture view, a use case risk tiering, a lifecycle control map, monitoring thresholds, an assurance evidence plan and a phased roadmap. It is used to brief the executive sponsor and scope the first build.

Frequently Asked Questions

Is training on governed AI platforms suitable for leaders who are not engineers?

Yes. Governed AI platforms training is set at advanced level for executives and risk leaders. It assumes familiarity with enterprise technology or risk management, not programming skills.



Which standards do governed AI platforms controls in this programme rest on?

Governed AI platforms design rests on ISO/IEC 42001 for management systems, ISO/IEC 23894 for risk management and the NIST AI Risk Management Framework.

How is the capstone from governed AI platforms training used at work?

The governed AI platforms blueprint briefs the executive sponsor and scopes the first build or remediation. Architecture and risk teams can adopt its inventory, risk tiers and control map.