ARTIFICIAL INTELLIGENCE | INTERMEDIATE | 3 DAYS

AI Governance and Assurance: EU AI Act, ISO 42001 and NIST AI RMF

A three-day programme for the people accountable when an auditor asks who approved an AI system. Covers the revised EU AI Act timeline, the transparency and literacy obligations already in force, ISO/IEC 42001 management systems, the NIST AI risk framework, the OWASP LLM and Agentic risk taxonomies, and hands-on red teaming. Produces a control set mapped across all four frameworks rather than four separate exercises.

Why AI Governance and Assurance?

The calendar just changed

July 2026 amendments moved some deadlines and left others in force. Most compliance plans are now wrong.

Assurance, not awareness

Includes adversarial testing so you can evidence controls rather than assert them.

One control set, four frameworks

Map controls once and satisfy the AI Act, ISO 42001, NIST and OWASP together.

Built around audit questions

Structured on what auditors actually ask: who approved it, where are the logs, how fast can it stop.

What you'll be able to do

- ✓ Map which AI regulations and standards apply to your organisation and systems
- ✓ Explain the EU AI Act's risk tiers, actor roles and extraterritorial reach
- ✓ Apply the revised compliance timeline correctly, distinguishing what moved from what did not
- ✓ Meet transparency and AI literacy obligations that are already in force
- ✓ Prepare for high-risk obligations including risk management, data governance and human oversight
- ✓ Build an AI management system aligned to ISO/IEC 42001 and integrate it with ISO 27001
- ✓ Operate the NIST AI risk framework functions as a working governance loop
- ✓ Apply the OWASP LLM and Agentic risk taxonomies as your security vocabulary
- ✓ Establish non-human identity, permissioning and an AI bill of materials for agents
- ✓ Design and run red-team tests for prompt injection, tool abuse and data exfiltration
- ✓ Stand up an AI inventory, risk register, intake workflow and incident process
- ✓ Produce the documentation set an auditor or client security review will accept

Who this is for

- Risk, compliance and internal audit professionals
- Security architects and CISO-office teams
- Legal, privacy and data protection officers
- AI programme owners and enterprise architects
- Consultants advising on AI assurance

Curriculum

01 The AI Regulatory and Standards Landscape

- What actually applies: horizontal regulation, sectoral rules and voluntary standards
- How the AI Act, data protection law, operational resilience rules and security standards interact
- Approaches taken in the UK, US, India, Singapore and the Gulf, and where they diverge
- Extraterritorial reach: when EU obligations bind a non-EU organisation
- **Hands-on:** build an applicability map for your own AI estate

02 EU AI Act: Structure, Scope and the Revised Timeline

- Risk tiers and how a system is classified in practice
- Prohibited practices, including the categories added in the 2026 amendments
- Actor roles: provider, deployer, importer, distributor, and why the distinction matters
- General purpose AI model obligations and where they sit
- The 2026 deferral: what moved to December 2027 and August 2028, and what did not move at all
- **Hands-on:** rebuild a compliance calendar against the current legal position

03 Obligations Already in Force: Transparency and AI Literacy

- Disclosure when a person is interacting with an AI system
- Marking and labelling of synthetic audio, image, video and text
- Deepfake disclosure and public interest content rules
- Notices for emotion recognition and biometric categorisation systems
- The AI literacy duty and how to evidence it across a workforce
- Legacy systems, the December 2026 transition and technical marking solutions

04 High-Risk Obligations and Conformity Assessment

- Annex III standalone systems against Annex I embedded systems
- The required risk management system and its lifecycle
- Data and data governance requirements, including bias testing
- Technical documentation, record keeping and automatic logging
- Human oversight design, accuracy, robustness and cybersecurity requirements
- Conformity assessment routes, notified bodies and the state of harmonised standards

05 ISO/IEC 42001: Building an AI Management System

- Management system structure and the clause-by-clause requirements
- Annex A controls and selecting an applicable control set
- Integrating with an existing ISO 27001 management system rather than duplicating it
- Scoping decisions and defining the boundary of the AI management system
- AI system impact assessment as a repeatable process
- Internal audit, management review and the certification path
- **Hands-on:** draft a scope statement and initial control selection

06 NIST AI Risk Management Framework

- The govern, map, measure and manage functions as an operating loop
- Trustworthiness characteristics and how to make them measurable
- The generative AI profile and what it adds for foundation model use
- Using the framework where no regulation compels you, and why that still helps
- **Hands-on:** run a mapping exercise on a real AI use case

07 OWASP Top 10 for LLM Applications

- The current risk categories and what changed in the latest edition
- Treating a model as a component against treating it as an actor
- Containment-first architecture: harden the surroundings rather than the model
- Mapping to adversarial threat knowledge bases and weakness taxonomies
- Where these risks land inside a management framework's risk register

08 OWASP Top 10 for Agentic Applications

- Why agents need a separate taxonomy: planning, memory, tools, identity, inter-agent calls
- Goal hijack as the agentic counterpart of prompt injection, with action consequences
- Agent identity: distinct scoped workload identity rather than borrowing a user token
- Memory poisoning, supply chain risk and rogue or ownerless agents
- Inter-agent trust and privilege escalation across an agent fleet
- The AI bill of materials as the inventory artefact auditors now expect

09 Red Teaming and Assurance Testing

- Threat modelling an AI system before it ships
- Direct and indirect prompt injection testing through documents, retrieval and tools
- Jailbreak, refusal bypass and harmful output evaluation
- Tool poisoning, excessive permission and exfiltration path testing
- Behavioural baselines and detecting a hijacked agent from a productive one
- Building a repeatable test suite and capturing results as evidence
- **Hands-on:** red team a running system, then document the findings

10 Standing Up the Governance Programme

- AI inventory and use case registration as the foundation for everything else
- Risk register, risk appetite and treatment decisions
- Roles, accountability and the approval workflow for new AI use
- Intake triage: fast lanes for low risk, real scrutiny for high risk
- Third-party and vendor AI assessment, including model and connector supply chain
- Incident response, kill-switch procedures and post-incident review
- Board and regulator reporting that is accurate without being unreadable

11 Evidence, Audit and Certification Readiness

- Mapping one control set across regulation, management standard and risk taxonomy
- The documentation set: policies, assessments, logs, test results, decision records
- Answering the questions auditors ask: who approved permissions, where are the logs, how fast can it stop
- Running a gap assessment against your current state
- Building a prioritised remediation roadmap with owners and dates
- **Hands-on:** produce a gap assessment and roadmap for your organisation