JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

AI Model Cards, Disclosure Notices and Audit Logging Training Course

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

AI Model Cards, Disclosure Notices and Audit Logging Training Course

AI Model Cards, Disclosure Notices and Audit Logging is a 1-week training course for AI oversight leads, risk officers and internal auditors who will build an AI Transparency Evidence Pack. Customers, boards and auditors now ask how an automated decision was reached and who signed it off. Jupiter Prime Training presents AI Model Cards, Disclosure Notices and Audit Logging as the discipline of making every AI system explainable, traceable and answerable to a named owner.

Why AI Model Cards, Disclosure Notices and Audit Logging Matters Now

Three groups are raising the bar for evidence. Users affected by automated credit, hiring or triage outcomes increasingly expect to be told when AI is involved and why. Boards want concise reporting on which models run, who owns them and what changed. Auditors look for logs that reconstruct a decision months later. An organisation with clear model cards and reliable logs answers challenges quickly; one without them faces disputes it cannot settle. The course rests on ISO/IEC 42001, the AI management system standard, and the NIST AI Risk Management Framework AI RMF.

Who Gains Most from AI Model Cards, Disclosure Notices and Audit Logging

- Professionals who assign ownership and sign-off for AI systems across business units
- Specialists who write and maintain model documentation and change records
- Officers who draft notices telling customers and staff when AI shapes an outcome
- Auditors who test whether logs and records support a decision after the event
- Advisers who prepare AI inventories and summaries for board committees

An organisation gains most when its risk, data protection, audit and model engineering staff attend together, since a model card is only trusted when each function agrees what it must contain.

AI Model Cards, Disclosure Notices and Audit Logging Capabilities

Participants turn scattered model knowledge into records that withstand questions from users, boards and auditors.

- Specify a RACI matrix that names an accountable owner for each AI system — executive sponsors
- Build model cards covering purpose, training data, limits and performance by group — model owners
- Design layered disclosure notices and explanation requests for affected users — customer-facing units
- Configure decision logs that capture inputs, model version and human overrides — internal audit
- Test explanation methods such as feature attribution for faithfulness and readability — validation teams
- Communicate an AI inventory and exception summary in board-ready form — audit committees

Key AI Model Cards, Disclosure Notices and Audit Logging Terms

- Model card — a short structured document stating a model's intended use, data, performance, limits and owner.
- Disclosure notice — a plain-language statement telling a person that AI contributed to a decision about them.
- Decision log — a tamper-evident record of inputs, outputs, model version and reviewer for each automated decision.
- RACI matrix — a chart showing who is responsible, accountable, consulted and informed for an activity.
- Feature attribution — an explanation method estimating how much each input pushed a prediction up or down.
- Lineage — the traceable history linking a model to its datasets, code, approvals and releases.

AI Model Cards, Disclosure Notices and Audit Logging Day by Day

Day 1: Ownership and the AI Inventory

- AI System Inventory And Risk Tiering — Scoping What Needs Evidence
- RACI Matrices For Model Ownership — Assigning Accountable Sign-Off
- ISO/IEC 42001 Documented Information Clauses — Record Expectations
- NIST AI RMF Govern Function — Roles And Escalation Paths
- Third-Party Model Supplier Evidence — Contractual Documentation Requests

Day 2: Model Cards and Lineage

- Model Card Templates And Required Fields — Release Approval
- Datasheets For Training Data — Provenance And Consent Checks
- Performance Reporting By Subgroup — Disclosing Known Limits
- Version Control And Lineage Records — Change Traceability
- Generative Model System Cards — Prompt And Output Boundaries

Day 3: Explanations and Disclosure Notices

- Feature Attribution And Counterfactual Explanations — Choosing A Method
- Explanation Faithfulness Testing — Avoiding Misleading Reasons
- Layered Disclosure Notice Design — Customer And Staff Wording
- Explanation Request And Contest Handling — Human Review Routes
- Chatbot And Synthetic Content Labelling — User Awareness

Day 4: Audit Logging and Board Reporting

- Decision Log Schemas And Retention — Reconstructing Outcomes
- Tamper-Evident Logging And Access Controls — Evidence Integrity
- Human Override Records — Measuring Reliance On Models
- Audit Sampling Of Automated Decisions — Assurance Opinions
- Board AI Dashboard And Exception Summary — Oversight Focus

Day 5: Practising the AI Transparency Evidence Pack

- Applying Model Card Drafting To A Retail Lending Score Case — Release File
- Applying Disclosure Notice Design To A Hospital Triage Assistant Case — Patient Wording
- Applying Decision Log Review To A Recruitment Screening Case — Audit Findings
- Peer Challenge Of Ownership Gaps — Accountability Check
- AI Transparency Evidence Pack — Cards, Notices, Logs And Owners

Applying AI Model Cards, Disclosure Notices and Audit Logging After the Course

- Draft model cards for the three highest-tier AI systems in the inventory and obtain owner sign-off.
- Sample twenty logged automated decisions and test whether each can be reconstructed and explained.

AI Model Cards, Disclosure Notices and Audit Logging Questions FAQ

What are AI model cards, disclosure notices and audit logging and why do they matter now?

They are the records that show what an AI system does, who owns it, what users are told and how each decision can be traced. They matter now because users, boards and auditors increasingly demand that evidence.

Which roles gain most from AI model cards, disclosure notices and audit logging?

AI oversight leads, risk officers, data protection specialists and internal auditors gain most.

How do AI model cards, disclosure notices and audit logging differ from a course on AI ethics principles?

An ethics principles course discusses values and policy statements. This subject focuses on the concrete documents, notices and logs that prove those values are applied to each system.

What do you bring back from AI model cards, disclosure notices and audit logging?

You bring back a draft evidence pack with model cards, user notices, log specifications and named owners for your systems.

AI Model Cards, Disclosure Notices and Audit Logging: The Way Forward

The AI Transparency Evidence Pack brings together model cards, disclosure notices, decision log specifications and the owner of each system. Risk and audit leaders first use it to decide which high-tier systems may go live or stay in service. After the first audit cycle, it is worth revisiting which records were actually consulted.