JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

AI Model Operations and Post-Deployment Performance Training Course

21 – 25 June 2027
Munich

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

AI Model Operations and Post-Deployment Performance Training Course

Ref: 1599_58596 **Date:** 21 – 25 June 2027 **Location:** Munich **Fees:** 6200 GBP

AI Model Operations and Post-Deployment Performance is a 1-week training course for AI product owners, data science leads and IT operations managers who keep live models accurate, reliable and worth their running cost long after launch, and it ends with a Model Run Book and Health Scorecard. Many models quietly degrade once real-world inputs drift from their training history. Jupiter Prime Training runs this course on AI Model Operations and Post-Deployment Performance.

Why AI Model Operations and Post-Deployment Performance Matters Now

Three forces raise the stakes after go-live: growing numbers of models running inside customer and operational systems, shifting customer behaviour and supplier data that erode accuracy, and auditors and boards asking who owns each model once the project team disbands. An organisation that runs models as managed services catches decay early, retrains on evidence and retires stale models cleanly; one that treats launch as the finish line meets silent errors, rising cloud bills and incidents nobody owns. Operational accountability in the course follows ISO/IEC 42001 artificial intelligence management system and the NIST AI Risk Management Framework NIST AI RMF.

Who Gains Most from AI Model Operations and Post-Deployment Performance

- Product owners accountable for a live model's business outcome and backlog
- Data science leads who decide when and how models are retrained
- IT operations and platform managers responsible for uptime, cost and release pipelines
- Service and incident managers who handle model failures alongside other outages
- Model risk reviewers who sign off changes to models already in production

A product line sending its owner, data science lead and operations manager together gains most.

AI Model Operations and Post-Deployment Performance Capabilities

Participants turn a deployed model from a finished project into a service with owners, thresholds and a maintenance rhythm.

- Specify service level objectives for accuracy, latency and cost — product and operations agreement
- Build monitoring for data drift, concept drift and prediction quality — early warning
- Design retraining triggers, champion-challenger tests and rollback paths — release approval
- Apply incident response and root cause methods to model failures — service managers
- Calculate run cost and realised benefit per model — keep, rework or retire decisions
- Assign lifecycle ownership and change control — model risk reviewers

Key AI Model Operations and Post-Deployment Performance Terms

- Data drift — a change in the statistical pattern of a model's inputs compared with the data it was trained on.
- Concept drift — a change in the relationship between inputs and the outcome, so past patterns no longer predict well.
- Model registry — a controlled catalogue of model versions, their training data, approvals and deployment status.
- Champion-challenger testing — running a new model version beside the current one on live traffic before switching.
- Service level objective — a measurable target, such as response time or accuracy, agreed for a running model.
- Model decommissioning — the planned retirement of a model, including dependent systems, records and user communication.

AI Model Operations and Post-Deployment Performance Day by Day

Day 1: Life After Go-Live

- Model Inventory and Criticality Tiering — Where Oversight Effort Goes
- Causes of Post-Launch Decay — Input Shifts, Feedback Loops and Upstream Changes
- Handover from Project to Run Team — Ownership Gaps
- Operational Accountability Under ISO/IEC 42001 and NIST AI RMF — Role Clarity
- Service Level Objectives for Models — Accuracy, Latency and Cost Targets

Day 2: Monitoring and Drift Detection

- Population Stability and Distribution Tests — Data Drift Alerts
- Delayed Ground Truth and Proxy Metrics — Concept Drift Signals
- Segment-Level Performance Slicing — Hidden Degradation in Subgroups
- Observability Dashboards and Alert Thresholds — Signal Versus Noise
- Upstream Data Contract Checks — Breaking Changes from Source Systems

Day 3: Retraining and Release Discipline

- Retraining Triggers: Scheduled, Threshold and Event-Based — Cadence Choice
- Model Registry and Version Lineage — Reproducible Releases
- Champion-Challenger and Shadow Deployment — Safe Switching
- Automated Pipeline Tests and Approval Gates — Release Sign-Off
- Rollback and Fallback Rules — Protecting Users When a Release Fails

Day 4: Incidents, Cost and Benefit Tracking

- Model Incident Classification and Response — Severity and Escalation
- Post-Incident Root Cause Review — Data, Code or Environment
- Compute and Inference Cost Tracking — Unit Cost per Prediction
- Realised Benefit Measurement After Launch — Comparing Forecast with Actual
- Keep, Rework or Retire Assessment — Model Portfolio Pruning

Day 5: Practising the Model Run Book and Health Scorecard

- Monitoring Design for a Participant's Live Model — Drift and Quality Alerts
- Retraining and Rollback Plan Applied — Release Rules
- Incident Playbook Walk-Through — Roles and Communication
- Ownership and Change Control Matrix — Sign-Off Responsibilities
- Model Run Book and Health Scorecard — Thresholds, Owners and Review Rhythm

Applying AI Model Operations and Post-Deployment Performance After the Course

- Add segment-level drift and quality checks to one production model and set alert thresholds.
- Run a keep, rework or retire review of the model inventory using unit cost and realised benefit.

AI Model Operations and Post-Deployment Performance Questions FAQ

What is AI Model Operations and Post-Deployment Performance and why does it matter now?

It is the discipline of running live models as services with monitoring, retraining, incident handling and ownership. It matters now because more models run in core systems while shifting data erodes their accuracy without warning.

Which roles gain most from AI Model Operations and Post-Deployment Performance?

Product owners, data science leads, IT operations managers and incident managers gain most. You benefit most if you inherit models from project teams or answer for their results.

How does AI Model Operations and Post-Deployment Performance differ from a pilot-to-production scaling course?

It begins where scaling ends. A scaling course moves a proven pilot into production; this one keeps an already running model accurate, affordable and owned through drift, retraining, incidents and eventual retirement.

What do you bring back from AI Model Operations and Post-Deployment Performance?

You bring back a run book and health scorecard for one of your models. It defines thresholds, owners, retraining and rollback rules, and incident steps.

AI Model Operations and Post-Deployment Performance: The Way Forward

The Model Run Book and Health Scorecard first informs the next retraining or rollback decision, giving product owners and operations staff shared thresholds. Examples from insurance pricing and demand forecasting in retail show different drift patterns. After its first use, revisit the alert thresholds, because early settings often trigger too many or too few warnings.