JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

AI Pilot-to-Production Scaling and Value Tracking Training Course

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

AI Pilot-to-Production Scaling and Value Tracking Training Course

AI Pilot-to-Production Scaling and Value Tracking is a 1-week training course for delivery leads, product owners and operations managers who must turn promising AI experiments into dependable live services, and each participant builds an AI Scale-Up Readiness and Value Ledger. Many organisations now run dozens of proofs of concept yet move few into daily operations, so capital and credibility leak away. The course is delivered by Jupiter Prime Training.

Why AI Pilot-to-Production Scaling and Value Tracking Matters Now

Three forces shape the move from experiment to live service. Executive boards want evidence that AI spending changes cost, revenue or service levels, not more demonstrations. Operations teams inherit models they did not build and must keep them accurate as inputs shift. Risk functions expect each live model to carry clear records and controls. An organisation that scales well compounds gains release after release; one that scales poorly accumulates brittle models and sceptical users. The course draws on the NIST AI Risk Management Framework AI RMF and ISO/IEC 42001, the AI management system standard.

Who Gains Most from AI Pilot-to-Production Scaling and Value Tracking

- Delivery leads accountable for moving proofs of concept through scale gates
- Product owners who decide which model features reach end users
- Operations managers who absorb new predictions into existing workflows
- Finance partners who verify the gains claimed for each deployment
- Data and platform leads responsible for monitoring and retraining live models

Cross-functional squads gain most when a delivery lead, a process owner and a finance partner attend together.

AI Pilot-to-Production Scaling and Value Tracking Capabilities

Participants leave able to judge when an experiment is fit for live traffic and to prove the gains it delivers afterwards.

- Assess source data and feature lineage against production needs — for the platform lead approving go-live
- Design controlled pilots with comparison cohorts and kill criteria — for sponsors funding the next phase
- Specify deployment patterns, monitoring thresholds and rollback steps — for the operations owner on call
- Calculate realised benefits with counterfactual attribution — for finance partners signing off gains
- Communicate adoption signals and override rates — for process owners redesigning workflows

- Build a gated rollout sequence across sites and waves — for the steering group setting priorities

Key AI Pilot-to-Production Scaling and Value Tracking Terms

- Pilot purgatory — a promising AI proof of concept that never gains the funding or ownership to go live.
- Scale gate — a decision point where evidence decides whether a model moves to wider release.
- Model drift — a decline in prediction accuracy caused by changes in input data or in the behaviour the model predicts.
- Canary release — sending a small share of live traffic to a new model version before full rollout.
- Counterfactual baseline — an estimate of results without the model, isolating its genuine contribution.
- Benefits realisation — tracking how model performance turns into measured operational or financial gains.

AI Pilot-to-Production Scaling and Value Tracking Day by Day

Day 1: Why Pilots Stall Before Production

- Pilot Purgatory Patterns — Which Stalled Proofs of Concept to Revive
- Use-Case Intake Scoring Matrix — Ranking Candidates by Feasibility and Payoff
- Baseline Process Metrics — Measuring Today's Cycle Time Before Change
- Sponsor and Process Owner Mapping — Who Signs Off Each Scale Gate
- Risk Tiering With NIST AI RMF — Matching Controls to Model Impact

Day 2: Data Readiness and Pilot Design

- Data Readiness Scorecard — Deciding Whether Source Systems Can Feed Production
- Feature Lineage Mapping — Tracing Inputs From Source to Prediction
- Controlled Pilot Design — Choosing Treatment Groups and Comparison Cohorts
- Success Thresholds and Kill Criteria — When to Stop or Continue
- Pilot Charter Template — Fixing Scope, Duration and Exit Conditions

Day 3: MLOps Foundations for Scale

- Model Registry and Versioning — Knowing Which Model Runs Where
- Deployment Patterns: Shadow, Canary, Blue-Green — Limiting Rollout Exposure
- Drift and Performance Monitoring — Triggering Retraining Before Accuracy Decays
- Inference Cost and Latency Budgets — Sizing Infrastructure for Real Volumes
- Runbooks and Rollback Plans — Restoring Service When a Model Misbehaves

Day 4: Adoption, Benefits and Value Tracking

- Workflow Redesign Canvas — Embedding Predictions Into Daily Decisions
- Frontline Adoption Signals — Reading Override Rates and Usage Logs
- Benefits Realisation Ladder — Linking Model Metrics to Business Outcomes
- Attribution and Counterfactual Methods — Isolating the Model's True Contribution
- Value Dashboard Cadence — Reporting Gains to Finance and Sponsors

Day 5: Practising the Scale-Up Readiness Review

- Readiness Gate Rehearsal — Applying Data and Pilot Criteria to a Case
- Rollout Sequencing Exercise — Planning Sites, Waves and Monitoring Hooks
- ISO/IEC 42001 Control Check — Confirming Records for Each Live Model
- Peer Challenge Panel — Stress-Testing Value Claims and Assumptions
- AI Scale-Up Readiness and Value Ledger — Completing the Ledger

Applying AI Pilot-to-Production Scaling and Value Tracking After the Course

- Run a readiness gate review on one stalled proof of concept using the data readiness scorecard and kill criteria
- Set up a monthly value dashboard for one live model, pairing drift metrics with counterfactual benefit estimates

AI Pilot-to-Production Scaling and Value Tracking Questions FAQ

What is AI pilot-to-production scaling and value tracking, and why does it matter now?

It is the discipline of moving AI experiments into live services and proving their gains. It matters now because boards increasingly fund only models with measured results.

Which roles gain most from AI pilot-to-production scaling and value tracking?

Delivery leads, product owners and operations managers gain most, alongside finance and platform leads who verify or sustain live models.

How does AI pilot-to-production scaling and value tracking differ from an AI operating model course?

It works on individual models rather than organisation design. An operating model course shapes structures and a centre of excellence; this subject covers readiness gates, drift monitoring and benefit attribution, with logistics and healthcare examples.

What do you bring back from AI pilot-to-production scaling and value tracking?

You bring back a completed readiness and value ledger recording data readiness, pilot results, rollback choices and the counterfactual method used to report gains.

AI Pilot-to-Production Scaling and Value Tracking: The Way Forward

The AI Scale-Up Readiness and Value Ledger first informs the go or stop decision at the next scale gate. Delivery leads, platform owners and finance partners use it to align on evidence before wider release. After its first use, revisit the attribution method, because early gains in a retail or claims-handling setting often shift once users adapt their routines.