



Production-Grade MLOps: Building Reliable Machine Learning Systems Using SRE Principles



Production-Grade MLOps: Building Reliable Machine Learning Systems Using SRE Principles

Course Overview:

This advanced training program connects machine learning engineering with site reliability engineering SRE to create reliable, scalable, and production-ready ML systems. The course covers best practices from software engineering and DevOps throughout the ML lifecycle.

Participants will explore key topics such as ML model monitoring, data reliability, model serving strategies, and incident response, aligned with industry standards like MLOps best practices, machine learning system design, and ML deployment strategies.

Target Audience:

- Machine Learning Engineers
- MLOps Engineers
- Site Reliability Engineers
- Data Scientists
- Data Engineers
- Software Developers integrating ML
- AI Product Managers
- DevOps Professionals entering ML environments

Targeted Organisational Departments:

- Data Science & AI Units
- Engineering & DevOps
- IT Operations & Infrastructure
- Quality Assurance and Risk
- Product and Innovation Teams
- ML Governance & Compliance

Targeted Industries:

- Financial Services
- Healthcare
- E-commerce & Retail
- Telecommunications
- Technology & Cloud Services
- Government & Public Sector

Course Offerings:

By the end of this course, participants will be able to:

- Design reliable ML systems using SRE principles
- Build scalable ML production pipelines
- Apply ML observability tools for monitoring and validation
- Define SLOs and SLIs for ML workflows
- Implement robust ML deployment strategies
- Mitigate ML model reproducibility issues and data drift
- Address ML incident response and recovery using structured playbooks
- Apply privacy, fairness, and ethical ML design considerations

Training Methodology:

This program combines instructor-led sessions, peer discussions, case studies, and simulation labs. Participants will work in small groups to design machine learning system architectures, analyse model failures, and establish Service Level Objectives SLOs and Service Level Indicators SLIs.

Course Toolbox:

- Course ebook and system design templates
- Access to monitoring and observability sandbox e.g., Prometheus, Grafana for ML
- Sample datasets for model training and validation
- Checklists for ML reproducibility and ethical AI assessment
- Templates for SLOs and incident response planning

Course Agenda:

Day 1: Foundations of Reliable ML Systems

- **Topic 1:** Understanding the ML Lifecycle and Reliability Challenges
- **Topic 2:** Core Principles of Site Reliability Engineering for ML Systems
- **Topic 3:** Data Collection, Labeling, and Governance Issues
- **Topic 4:** Building Robust ML Training Pipelines
- **Topic 5:** Failure Modes and Production Risks in ML Workflows
- **Topic 6:** Model Development vs. System Design Trade-offs
- **Reflection & Review:** Lessons from the ML Loop and YarnIt Case Study



Day 2: Data Management and Governance in ML

- **Topic 1:** Designing for Data Durability, Versioning, and Access Control
- **Topic 2:** Feature Stores, Metadata, and Labeling Infrastructure
- **Topic 3:** Data Privacy, Security, and Fairness Considerations
- **Topic 4:** Documentation Practices for Human Annotation and Label Quality
- **Topic 5:** Policy and Compliance Impacts on ML Pipelines
- **Topic 6:** Debugging Data-Driven Failures and Edge Cases
- **Reflection & Review:** Review of Governance Failures and Preventive Design

Day 3: Model Validation, Observability, and Monitoring

- **Topic 1:** Defining Quality Metrics for Model Validity and Effectiveness
- **Topic 2:** Offline Evaluation: Metrics, Distributions, and Benchmarks
- **Topic 3:** Online Evaluation: A/B Testing and Shadow Deployment
- **Topic 4:** Building and Using ML Observability Tools
- **Topic 5:** Designing and Measuring ML-specific SLOs and SLIs
- **Topic 6:** Monitoring for Feature Drift, Data Skew, and Model Degradation
- **Reflection & Review:** Observability Strategy and Dashboard Use Cases

Day 4: Scalable Deployment and Incident Response

- **Topic 1:** Model Serving Architectures: Batch, Online, and Edge
- **Topic 2:** Model Deployment Strategies: Blue/Green, Canary, and Rollbacks
- **Topic 3:** Autoscaling, Caching, and Disaster Recovery Patterns
- **Topic 4:** Developing and Executing Incident Response Playbooks
- **Topic 5:** Root Cause Analysis and Postmortems in ML Contexts
- **Topic 6:** Ethical Risks, Bias Failures, and Operational Accountability
- **Reflection & Review:** Simulation of Outage Response and Model Resilience

Day 5: Organizational Integration and MLOps Best Practices

- **Topic 1:** Designing ML Teams and Roles Across the Organization
- **Topic 2:** Organizational Patterns for ML Integration: Centralized vs. Decentralized
- **Topic 3:** Continuous ML Systems and Real-Time Model Updates
- **Topic 4:** Governance, Ethics, and Lifecycle Ownership
- **Topic 5:** Practical Case Studies: NLP Load Testing, Privacy-Aware Pipelines, Ad Click Prediction
- **Topic 6:** Auditing and Compliance in Enterprise MLOps
- **Reflection & Review:** Capstone Presentations and Peer Feedback

FAQ:



What specific qualifications or prerequisites are needed for participants before enrolling in the course?

Basic understanding of ML concepts, familiarity with DevOps or software engineering practices, and some experience with cloud platforms or ML frameworks e.g., TensorFlow, PyTorch are recommended.

How long is each day's session, and is there a total number of hours required for the entire course?

Each day's session is generally structured to last around 4-5 hours, with breaks and interactive activities included. The total course duration spans five days, approximately 20-25 hours of instruction.

What's the difference between monitoring ML models and traditional software systems?

Monitoring ML models goes beyond basic metrics like uptime and latency. It involves tracking model accuracy, feature drift, data skew, and SLO violations. *Reliable Machine Learning* emphasises the need for specialised observability strategies that address ML-specific failure modes.

How This Course is Different from Other Production-Grade MLOps Courses:

Unlike typical MLOps training, this course emphasises operational excellence. It combines *reliable machine learning* principles with software engineering practices and real-world case studies of ML failures, model drift, and incident recovery.

Incorporating Site Reliability Engineering SRE concepts like Service Level Objectives SLOs and observability, participants learn to effectively build, deploy, and manage machine learning models in complex environments. The course also addresses ethical considerations, feature store design, and continuous deployment, making it a modern choice for professionals seeking scalable and high-performing machine learning systems.



Training Course Categories

Agile PM and Project Management Training Courses

Certified Courses By International Bodies

Communication and Public Relations Training Courses

Data Analytics Training and Data Science Courses

Environment & Sustainability Training Courses

Finance and Accounting Training Courses

Governance, Risk and Compliance Training Courses

Human Resources Training and Development Courses

IT Security Training & IT Training Courses

Leadership and Management Training Courses

Legal Training, Procurement and Contracting Courses

Maintenance Training and Engineering Training Courses

Marketing, Customer Relations, and Sales Courses

Occupational Health, Safety and Security Training Courses

Personal & Self-Development Training Courses

Quality and Operations Management Training Courses

Secretarial and Administration Training Courses



Training Cities

Accra Ghana	Al Jubail Saudi Arabia	Amman Jordan
Amsterdam Netherlands	Athens Greece	Baku Azerbaijan
Bali Indonesia	Bangkok Thailand	Barcelona Spain
Berlin Germany	Cairo Egypt	Cape town South Africa
Casablanca Morocco	Chicago USA	Doha Qatar
Dubai UAE	Frankfurt Germany	Geneva Switzerland
Istanbul Turkey	Jakarta Indonesia	Johannesburg South Africa
Kuala Lumpur Malaysia	Kuwait Kuwait	Langkawi Malaysia
Lisbon Portugal	London UK	Madrid Spain
Manama Bahrain	Marbella Spain	Milan Italy
Montreux Switzerland	Munich Germany	Muscat Oman
Nairobi Kenya	New York USA	Nice France
Paris France	Phuket Thailand	Porto Portugal



Training Cities

Prague

Czech Republic

Riyadh

Saudi Arabia

Rome

Italy

San Diego

USA

Seoul

South Korea

Sharm El-Sheikh

Egypt

Singapore

Singapore

Tashkent

Uzbekistan

Tbilisi

Georgia

Tokyo

Japan

Toronto

Canada

Trabzon

Turkey

Vienna

Austria

Zanzibar

Tanzania

Zoom

Online Training

WHO WE ARE

Agile Leaders is a renowned training center with a team of experienced experts in vocational training and development. With 20 years of industry experience, we are committed to helping executives and managers replace traditional practices with more effective and agile approaches.

OUR VISION

We aspire to be the top choice training provider for organizations seeking to embrace agile business practices. As we progress towards our vision, our focus becomes increasingly customer-centric and agile.

OUR MISSION

We are dedicated to developing value-adding, customer-centric agile training courses that deliver a clear return on investment. Guided by our core agile values, we ensure our training is actionable and impactful.

WHAT DO WE OFFER

At Agile Leaders, we offer agile, bite-sized training courses that provide a real-life return on investment. Our courses focus on enhancing knowledge, improving skills, and changing attitudes. We achieve this through engaging and interactive training techniques, including Q&As, live discussions, games, and puzzles.



AGILE LEADERS
Training Center

CONTACT US



UAE, Dubai Investment
Park First



+971585964727
+447700176600



sales@agile4training.com