DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

MLOps Advanced Training

Scale training and inference, apply MLOps to LLM applications, and govern models with lineage, responsible AI controls and cost tracking.

LIVE ONLINE **Max 8-10 hours**	CORPORATE (DAY-WISE) **5 days · Days 1-5**	SELF-PACED VIDEO **1-year LMS access**
BATCH SIZE **Max 10 learners**	NEXT BATCHES **1st of every month · weekends: first Saturday**	CERTIFICATE **Verifiable certificate of completion**

Foundation	Essential	Intermediate	Advanced
1 day · Included	2 days · Included	3 days · Included	5 days · This curriculum

INSIDE THIS CURRICULUM

What you will achieve	Outcomes and level focus
Who should attend	Audience and prerequisites
Agenda	52 topics · Foundation to Advanced
Hands-on labs	4 labs
Formats and fees	Live, 1:1, self-paced, corporate
Trainer and reviews	Rajesh Kumar
Enrolment	Batches, payment, contacts

WHY DEVOPSSCHOOL

- Rated 4.6 on Google
- 1,500+ verifiable certificates issued in the public registry
- Trainer Rajesh Kumar: 20 years in industry
- Small batches: max 10 learners, every session recorded

Course page: www.devopsschool.com/courses/mlops/mlops-advance-training.html
Talk to us: WhatsApp +91 99057 40781 · +1 (469) 756-6329 · contact@DevOpsSchool.com · reply within 1 hour

OUTCOMES

What you will achieve

By the end of the MLOps course path, you will be able to:

1

Productionise ML models

Turn experiments into versioned, containerised and served models with reproducible pipelines.

2

Automate the ML lifecycle

Build CI/CD and continuous training on Kubernetes, Kubeflow and Vertex AI.

3

Operate governed ML platforms

Monitor, scale, secure and govern models, including LLM applications, in production.

WHAT THE ADVANCED LEVEL ADDS

Scale, LLMOps, governance and ML platforms

Hands-on: Capstone: governed enterprise ML platform. Deliver a multi-tenant pipeline platform that trains on GPUs, serves a canary model and a RAG endpoint, and records lineage, model cards and costs.

WHO SHOULD ATTEND

- Experienced MLOps practitioners responsible for quality, security and scale.
- Senior engineers, architects and leads who design and troubleshoot MLOps solutions.
- Teams making MLOps an organisation-wide standard.

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Promote models through a registry, test data and models automatically, serve them on Kubernetes and detect drift to trigger retraining.
- Basic computer skills and a laptop with a stable internet connection for the live sessions and labs.
- This curriculum includes every earlier level, so it is complete on its own.

AGENDA · FOUNDATION TO ADVANCED

MLOps Advanced Training agenda

52 topics. Each level builds on the one before it, so this curriculum lists every earlier level in full. Days are for corporate (day-wise) delivery; live online batches cover the same agenda in max 8-10 hours.

DAY 1 — Foundation: MLOps concepts, ML lifecycle and first served model

Included from Foundation

1. Welcome to MLOps fundamentals: why and when we need MLOps and data scientists' pain points
2. The concept of DevOps in ML: CI, CD and continuous training, and what differs for data and models
3. Machine learning lifecycle: problem framing, data, features, training, evaluation, deployment, monitoring and retraining
4. ML basics for engineers: features and labels, training versus inference, overfitting, precision, recall and RMSE
5. Roles and hand-offs: data scientist, data engineer, ML engineer and platform engineer responsibilities
6. Python project setup for ML: uv or conda environments, pinned dependencies, scripts over notebooks, project layout
7. Git for ML projects: what to commit, ignoring data and model files, clean notebook diffs with nbstripout
8. Data versioning with DVC: dvc init, dvc add, remote storage and reproducing a dataset version
9. Experiment tracking with MLflow: tracking server, logging parameters, metrics and artifacts, comparing runs
10. First model: scikit-learn Pipeline, train and test split, cross-validation and saving the model artifact
11. Serving basics: a FastAPI /predict endpoint with a Pydantic request schema
12. Packaging the model service: Dockerfile, image build and running the prediction API locally
13. Why models decay: data drift, concept drift and feedback loops in production

DAY 2 — Essential: Containers, GKE, Kubeflow Pipelines and ML CI/CD

Included from Essential

14. Introduction to containers: containers and container images, main Kubernetes components as an optional primer
15. Lab intro: getting started with GCP and Qwiklabs, working with Cloud Build and walking through the lab solution
16. Introduction to Kubernetes and Google Kubernetes Engine: compute options in detail, Kubernetes concepts, control plane
17. GKE concepts and lab: Google Kubernetes Engine concepts and deploying a GKE cluster
18. Deployments: ways to create deployments, services and scaling, updating deployments
19. Release strategies: rolling updates, blue-green and canary deployments, lab creating GKE deployments
20. Jobs and CronJobs: parallel jobs and CronJobs for scheduled training and batch scoring
21. Setting up the tools: AI Platform Pipelines, now Vertex AI Pipelines, overview, concepts, when to use, ecosystem, lab
22. Training, tuning and serving on AI Platform, now Vertex AI: system and concepts, reproducible dataset, tunable model
23. Training containers: build and push a training container, train and tune, serve and query, custom containers lab
24. Kubeflow Pipelines: system and concepts, KFP DSL, pre-built, lightweight Python and custom components, compile, upload, run
25. Continuous training lab: a continuous training pipeline with Kubeflow Pipelines and Cloud AI Platform, intro and solution
26. CI/CD for Kubeflow Pipelines: concept overview, Cloud Build builders, configuration and triggers, lab and summary

DAY 3

Included from Intermediate

Intermediate: Registries, feature pipelines, testing, serving and monitoring

27. Model registry: MLflow registry and Vertex AI Model Registry, versions, champion and challenger aliases, approvals
28. Feature pipelines and stores: Feast or Vertex AI Feature Store, offline and online features, training-serving skew
29. Data validation: schema and distribution checks with TensorFlow Data Validation, Great Expectations or pandera
30. Testing ML systems: pytest for feature code, model quality thresholds and behavioural tests on edge cases
31. Pull request checks for ML: GitHub Actions with CML reports comparing metrics against the main branch
32. Orchestrators compared: Kubeflow Pipelines, Airflow, Prefect, ZenML and Metaflow for ML workflows
33. Hyperparameter tuning at scale: Optuna, Katib and Vertex AI Vizier with early stopping
34. Serving on Kubernetes: KServe InferenceService, BentoML and NVIDIA Triton, autoscaling and health probes
35. Batch inference: scheduled scoring jobs, Vertex AI batch prediction and writing results back to the warehouse
36. Production monitoring: Evidently drift reports, Vertex AI Model Monitoring, latency and error metrics in Prometheus
37. Automated retraining: drift, schedule or new-data triggers with safe validation before redeployment
38. Security basics: Workload Identity, least-privilege service accounts, secrets and vulnerability scanning of images
39. Troubleshooting: failed pipeline steps, caching surprises, resource limits and non-reproducible runs

DAYS 4-5

Advanced: Scale, LLMOps, governance and ML platforms

New at this level

40. MLOps maturity: assessing manual, pipeline-automated and CI/CD-automated stages and building a roadmap
41. Distributed training: Ray Train, PyTorch DDP, Kubeflow Trainer and GPU node pools on Kubernetes
42. Scaling inference: dynamic batching, GPU sharing with MIG or time-slicing, scale to zero with Knative
43. Progressive model delivery: shadow traffic, A/B tests, KServe canary traffic splits and automated rollback
44. LLMOps: prompt and model versioning, evaluation datasets, RAG pipelines, vLLM serving and guardrails
45. Lineage and reproducibility: ML Metadata, artifact lineage and audit-ready records for every model version
46. Model governance: model cards, approval workflows, risk tiers and EU AI Act obligations
47. Responsible AI: fairness checks with Fairlearn, explanations with SHAP and human oversight in production
48. ML security: model supply chain, signed model artifacts, unsafe pickle files, adversarial and data-poisoning risks
49. Cost optimisation: spot GPUs, right-sized instances, cost per training run and cost per thousand predictions
50. Multi-tenant ML platforms: Kubeflow profiles or Vertex AI projects, reusable components and self-service onboarding
51. Portability: comparing Vertex AI, Amazon SageMaker and Azure Machine Learning and avoiding lock-in
52. Capstone: enterprise MLOps platform with CI/CD/CT, monitoring, governance and an LLM use case

PRACTICE

Hands-on labs

You build these yourself during the course, guided by the trainer.

FOUNDATION

From notebook to prediction API

Turn a notebook model into a versioned project with DVC and MLflow tracking, then serve it through a FastAPI endpoint in a Docker container.

ESSENTIAL

Continuous training pipeline on GCP

Build and push a training container, run a Kubeflow pipeline that trains, tunes and deploys a model, and trigger it from a Cloud Build CI/CD pipeline.

INTERMEDIATE

Monitored model with automated retraining

Register a model, deploy it with KServe, generate drifted traffic, catch the drift with Evidently and retrain through a validated pipeline run.

ADVANCED

Capstone: governed enterprise ML platform

Deliver a multi-tenant pipeline platform that trains on GPUs, serves a canary model and a RAG endpoint, and records lineage, model cards and costs.

HOW YOU CAN LEARN

Formats and fees

The same agenda in four formats. Fees are in Indian rupees as of November 2026; the course page shows them in your currency.

FORMAT	DURATION	WHAT YOU GET	FEE
Live & interactive (online)	Max 8-10 hours	Instructor-led public batch of max 10 learners with demos, labs, Q&A and assignment review. Weekday batches start on the 1st of every month, weekend batches on the first Saturday of every month.	**INR 34,999**
1:1 live training (online)	Max 8-10 hours, one-to-one	The same agenda, taught one-to-one by the trainer. Starts immediately, without waiting for the next batch; sessions planned around your working hours, at your pace.	**INR 99,999**
Self-paced video	1-year LMS access	All four MLOps levels, plus 20+ more courses and 50+ tools, at your own pace.	**INR 833/month** billed yearly at INR 9,996
Corporate (day-wise)	5 days for this level	A private batch for your team, on demand. Foundation 1 day, Essential 2, Intermediate 3, Advanced 5 (cumulative).	**Quote on request**

ASSESSMENT AND CERTIFICATE

Your progress is checked through hands-on labs and assignments at every level. On completion you receive a DevOpsSchool certificate of completion with a unique credential ID that anyone can verify at devopsschool.com/certificates.

INCLUDED WITH EVERY LIVE BATCH

- Every live session is recorded and stays in your LMS for a year.
- Can't attend your batch? Move to next month's batch.
- Free demo class before you pay.
- Free 10-15 minute call with an advisor or the trainer to check the course fits you.

YOUR TRAINER

Learn from Rajesh Kumar

Principal Architect, trainer and mentor in DevOps, SRE, DevSecOps and platform engineering, with 20 years across 8+ software MNCs, including Staff DevOps/SRE Engineer for a leading Japanese mobile-payment platform and Principal DevOps Architect and Manager in Europe.

Profile: www.rajeshkumar.xyz · Videos: youtube.com/@TheDevOpsSchool

20 years in industry	**200+** corporate trainings
10,000+ engineers mentored	**70+** organisations coached

REVIEWS

What learners say

Rated 4.6 on Google · more reviews at devopsschool.com/testimonials and in our video testimonials.

> “Rajesh is very good trainer. Rajesh was able to resolve our queries and question effectively. We really liked the hands-on examples covered during this training program.”
>
> Indrayani, India

> “The training was very useful and interactive. Rajesh helped develop the confidence of all.”
>
> Abhinav Gupta, Pune

START YOUR MLOPS TRAINING

Join the next batch

Weekday batches start on the 1st of every month and weekend batches on the first Saturday of every month, with max 10 learners each. 1:1 training can start immediately; corporate batches run on demand. We reply within 1 hour.

Pay by UPI, credit and debit cards, net banking or PayPal. Before you pay, read our refund policy. All levels of this course: www.devopsschool.com/courses/mlops/

ENROL OR RESERVE A SEAT
MLOps Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com