DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Dataops Advanced Training

Scale DataOps into a governed, cost-aware self-service platform with measurable data product SLOs and a migration roadmap.

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	50 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/dataops/dataops-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 Dataops course path, you will be able to:

1

Apply DataOps principles

Map data value streams, remove waste and run data as products with clear SLAs.

2

Automate tested data pipelines

Build pipelines with Git, dbt, quality checks, CI/CD and infrastructure as code.

3

Scale governed data platforms

Operate observable, secure and cost-aware data platforms across many domains.

WHAT THE ADVANCED LEVEL ADDS

Platform scale, governance, cost and AI data

Hands-on: Capstone: governed self-service data platform. Design and partly build a multi-domain platform with scaffolded data products, policy-as-code access, SLO dashboards and a cost report per domain.

WHO SHOULD ATTEND

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

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Automate data infrastructure and deployments, stream changes with CDC, trace lineage and handle data incidents with clear runbooks.
- 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

Dataops Advanced Training agenda

50 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: DataOps principles, assembly line and roadmap

Included from Foundation

1. What DataOps is and is not: collaborative, automated data flows for predictable value; not a tool and not only DevOps for data
2. DataOps Manifesto: lean analytics-as-manufacturing, continuous improvement and quality at every station
3. From past to present: batch BI and waterfall analytics to cloud data platforms, streaming plus batch and product thinking
4. Organisational pain: siloed teams, environment drift, schema breaks, rework, governance, audit and platform sprawl
5. Value stream waste: queues, manual gates and duplicated pipelines found on a value stream map
6. Principles to practice: small batches, fast feedback, quality at the source, automation, reproducibility, observability
7. Release processes: mapping DevOps to DataOps with CI/CD for datasets, SQL and ELT code, policies and environment promotion
8. Operating model and roles: data products with SLAs and SLOs for freshness, completeness and accuracy, golden sources
9. Platform enablement: paved paths, review and sign-off of code plus data contracts
10. The data assembly line: ingest, prepare, validate, transform, serve, observe; dev to stage to prod with rollback readiness
11. CI/CD for data concepts: Git workflows, PR checks for SQL and contracts, quality gates, artifacted releases, auditable runs
12. Orchestration patterns: DAGs, retries, SLAs, backfills, event or schedule driven, hybrid; Apache Airflow and managed options
13. Governance and security: access control, tagging, lineage, policy-as-code, PII; Snowflake Horizon lineage and quality monitoring
14. DORA for data and transition playbook: flow and quality metrics, quick wins, 30/60/90-day plan, checklist and glossary

DAY 2 — Essential: Reproducible pipelines with Python, dbt and checks

Included from Essential

15. Reproducible environments: Docker images for pipeline code, pinned Python dependencies with uv or pip-tools, dev containers
16. Ingestion basics: batch extracts with Airbyte or dlt, incremental loads, idempotent writes and landing zones
17. SQL transformations with dbt Core: sources, staging and mart layers, ref, incremental models and seeds
18. Testing data: dbt generic and singular tests, unit tests, Soda or Great Expectations checks for freshness and volume
19. Data contracts in YAML: schema, owners, SLAs and breaking-change rules enforced on producer pull requests
20. Branching for data teams: trunk-based Git, short-lived branches, CODEOWNERS and SQL review checklists
21. Pre-commit quality: SQLFluff linting, Ruff for Python, secret scanning and YAML validation
22. CI pipeline: GitLab CI or GitHub Actions running lint, dbt build and checks against a temporary schema
23. Asset-based orchestration: Dagster software-defined assets, schedules, sensors and partitions
24. Environment configuration: per-developer targets, dev, test and prod profiles, environment variables and secrets

25. Documentation: dbt docs, column descriptions, ownership metadata and a data catalogue entry per dataset
26. Team rituals: data stand-ups, definition of done for data changes and release notes for consumers

DAY 3 **Intermediate: Infrastructure as code, lineage, CDC and incidents** — Included from Intermediate

27. Infrastructure as code: Terraform for warehouses, schemas, roles and storage on Snowflake, BigQuery or Databricks
28. Isolated test environments: zero-copy clones, Delta and Iceberg time travel and lakeFS branches per pull request
29. Deployment strategies: slim CI with state comparison, blue-green schemas and write-audit-publish releases
30. Change data capture: Debezium on Kafka, schema registry, ordering, deletes and late-arriving records
31. Lakehouse table formats: Delta Lake and Apache Iceberg schema evolution, compaction and partition evolution
32. Lineage: OpenLineage events from Airflow, dbt and Spark viewed in Marquez or DataHub for impact analysis
33. Data observability: freshness, volume, schema and distribution monitors with anomaly alerting
34. Pipelines on Kubernetes: KubernetesPodOperator, container images per task, resource requests and retries
35. Security basics: least-privilege roles, column masking, row access policies, PII tags and secrets in a vault
36. Integration with BI and apps: semantic models, reverse ETL and API serving with consumer SLAs
37. Data incident management: severity levels, on-call rota, runbooks, consumer communication and post-incident reviews
38. Troubleshooting: failed runs, safe reruns and backfills, duplicate loads, schema drift and skewed partitions

DAYS 4-5 **Advanced: Platform scale, governance, cost and AI data** — **New at this level**

39. Self-service data platform: golden paths, project scaffolding, platform team charter and internal developer portal
40. Data mesh in practice: domain data products, ownership, interoperability standards and federated governance
41. Unified governance: Unity Catalog, Horizon or Dataplex for catalogue, access, lineage and audit at enterprise scale
42. Policy-as-code: OPA rules, tag-based access control, automated access reviews and approval workflows
43. Reliability engineering for data: SLIs, SLOs and error budgets for freshness and accuracy with burn-rate alerts
44. High availability: multi-region storage, orchestrator HA, disaster recovery for warehouses and metadata
45. Performance at scale: clustering, partition pruning, materialisation strategy and workload isolation
46. FinOps for data: cost per pipeline and data product, warehouse sizing, auto-suspend and budget alerts
47. Privacy compliance: GDPR and India DPDP Act requirements, retention, deletion and consent propagation in pipelines
48. DataOps for AI: training data versioning, feature pipelines and quality gates for RAG document ingestion
49. Migration: legacy ETL and stored procedures to tested ELT with parallel runs and reconciliation reports
50. Capstone: design a governed multi-domain platform with CI/CD, observability, SLOs and a phased rollout

PRACTICE

Hands-on labs

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

FOUNDATION

Value stream map of one pipeline

Map one real pipeline from source to dashboard, measure its lead time and failure points, and choose your next three improvements.

ESSENTIAL

Containerised ELT pipeline with CI checks

Load a public dataset with dlt, model it in dbt with tests and a contract, orchestrate it with Dagster, and gate merges with a CI pipeline.

INTERMEDIATE

CDC to lakehouse with lineage

Stream changes from PostgreSQL through Debezium into Iceberg or Delta tables provisioned by Terraform, and trace lineage and freshness alerts end to end.

ADVANCED

Capstone: governed self-service data platform

Design and partly build a multi-domain platform with scaffolded data products, policy-as-code access, SLO dashboards and a cost report per domain.

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 Dataops 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 DATAOPS 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/dataops/

ENROL OR RESERVE A SEAT
Dataops Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com