DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Kafka Cluster Advanced Training

Design, tune and run production Kafka clusters with rack awareness, tiered storage, DR and Kubernetes operators.

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	49 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/kafka/kafka-cluster-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 Kafka Cluster course path, you will be able to:

1

Run Kafka clusters

Install, configure and operate KRaft-based Kafka clusters from a single node to multiple brokers.

2

Deliver data reliably

Configure producers, consumers, topics and security so events are durable, ordered and protected.

3

Operate Kafka at scale

Monitor, tune, rebalance and recover production Kafka with tiered storage, DR and Kubernetes.

WHAT THE ADVANCED LEVEL ADDS

Scale, disaster recovery, tuning and Kubernetes

Hands-on: Design and fail over a production cluster. Deploy a rack-aware cluster with Strimzi, replicate it to a second cluster with MirrorMaker 2, and run a documented failover and failback.

WHO SHOULD ATTEND

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

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Secure, monitor and operate a Kafka cluster with TLS, SASL, ACLs, Kafka Connect and partition management.
- 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

Kafka Cluster Advanced Training agenda

49 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: Kafka architecture, KRaft and first cluster

Included from Foundation

1. Event streaming concepts: events, producers, consumers and topics, and why Kafka decouples systems
2. Kafka architecture: brokers, controllers and the KRaft metadata quorum, with ZooKeeper removed in Kafka 4.x
3. Topics and partitions: ordering per partition, offsets, log segments and retention by time or size
4. Replication basics: leaders, followers, replication factor and in-sync replicas (ISR)
5. Installing Apache Kafka: Java 17 runtime, downloading the release, directory layout and config files
6. Formatting storage: kafka-storage.sh random-uuid and kafka-storage.sh format for a KRaft node
7. Starting a combined node: process.roles=broker,controller, node.id and listeners in server.properties
8. Topic CLI: kafka-topics.sh --create, --list and --describe with partitions and replication factor
9. Console clients: kafka-console-producer.sh and kafka-console-consumer.sh with --from-beginning
10. Consumer groups introduction: group.id, partition assignment and kafka-consumer-groups.sh --describe
11. Kafka in containers: the apache/kafka image and Docker Compose for local labs

DAY 2 — Essential: Multi-broker cluster, producers and consumers

Included from Essential

12. Multi-node cluster: dedicated controllers, a three-voter quorum and controller quorum configuration
13. Listeners: listeners, advertised.listeners, listener.security.protocol.map and the inter-broker listener
14. Producer configuration: acks, retries, enable.idempotence, linger.ms, batch.size and compression.type
15. Partitioning: message keys, the default partitioner and choosing a partition count
16. Consumer configuration: auto.offset.reset, enable.auto.commit, manual commits and max.poll.records
17. Rebalancing: the KIP-848 consumer group protocol in Kafka 4.x, cooperative assignment and static membership
18. Topic configuration: retention.ms, cleanup.policy=compact, min.insync.replicas and segment settings
19. Durability trade-offs: acks=all with min.insync.replicas=2, unclean leader election and data loss scenarios
20. Offsets and lag: the __consumer_offsets topic and kafka-consumer-groups.sh --reset-offsets
21. Client libraries: Java producer and consumer APIs, and librdkafka-based clients for Python and Go
22. Message formats: JSON, Avro and Protobuf with a schema registry and compatibility modes
23. Admin operations: kafka-configs.sh for dynamic configuration and adding partitions to a topic

DAY 3 — Intermediate: Security, Connect, monitoring and operations

Included from Intermediate

24. Encryption: SSL/TLS listeners, keystores and truststores, and mutual TLS for clients
25. Authentication: SASL/SCRAM-SHA-512 and SASL/OAUTHBEARER, creating SCRAM users with kafka-configs.sh
26. Authorization: StandardAuthorizer ACLs with kafka-acls.sh for topics, groups and cluster resources
27. Kafka Connect: distributed mode, source and sink connectors, the REST API and converters
28. Kafka Streams basics: KStream, KTable, stateful operations and state stores
29. Exactly-once semantics: transactional producers, transactional.id and isolation.level=read_committed
30. Monitoring: JMX metrics such as UnderReplicatedPartitions, ActiveControllerCount and request latency

31. Metrics stack: Prometheus JMX Exporter, Grafana dashboards and consumer lag exporters
32. Partition reassignment: kafka-reassign-partitions.sh, replication throttles and preferred leader election
33. Quorum health: kafka-metadata-quorum.sh describe --status and --replication
34. Disk management: log.dirs, JBOD, retention clean-up and log compaction tuning
35. Rolling restarts and upgrades: controlled shutdown, metadata.version and kafka-features.sh
36. Troubleshooting: under-replicated partitions, ISR shrink, consumer lag and broker log analysis

DAYS 4-5 **Advanced: Scale, disaster recovery, tuning and Kubernetes** **New at this level**

37. Capacity planning: sizing brokers, partitions, disk throughput and network for target throughput
38. Performance tuning: num.network.threads, num.io.threads, socket buffers, page cache and OS settings
39. Load testing: kafka-producer-perf-test.sh and kafka-consumer-perf-test.sh benchmarks
40. Rack awareness: broker.rack, replica placement across availability zones and fetching from followers
41. Cluster balancing: Cruise Control goals, rebalance proposals and self-healing
42. Tiered storage: remote log storage (KIP-405) to move older segments to object storage
43. Disaster recovery: MirrorMaker 2, offset translation, and active-passive vs active-active designs
44. ZooKeeper to KRaft migration: migrating 3.x clusters before upgrading to Kafka 4.x
45. Kafka on Kubernetes: Strimzi operator, Kafka and KafkaNodePool resources, and storage classes
46. Multi-tenancy: client quotas for byte and request rates, topic naming and ownership governance
47. Queues for Kafka (KIP-932): share groups for queue-style consumption and their current maturity
48. Security hardening: least-privilege ACLs, credential rotation, audit logging and network segmentation
49. Capstone: design a multi-AZ production cluster with DR, monitoring, security and an upgrade runbook

PRACTICE

Hands-on labs

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

FOUNDATION **Run a single-node KRaft Kafka**
Format and start a single KRaft node, create a topic with several partitions and exchange messages between console producers and consumers.

ESSENTIAL **Build a three-broker KRaft cluster**
Stand up three controllers and three brokers, create a replicated topic with min.insync.replicas=2, and prove no data loss when a broker stops.

INTERMEDIATE **Secure and monitor the cluster**
Enable TLS and SASL/SCRAM with ACLs, run a Kafka Connect pipeline, and build a Grafana dashboard that alerts on lag and under-replicated partitions.

ADVANCED **Design and fail over a production cluster**
Deploy a rack-aware cluster with Strimzi, replicate it to a second cluster with MirrorMaker 2, and run a documented failover and failback.

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 Kafka Cluster 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 KAFKA CLUSTER 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/kafka/

ENROL OR RESERVE A SEAT
Kafka Cluster Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com