DevOpsSchool

Course curriculum · November 2026

LEVEL 03 OF 04 · INTERMEDIATE

Clickhouse Intermediate Training

Run replicated and sharded ClickHouse clusters with Keeper, stream data from Kafka, secure access and monitor and back up the cluster.

LIVE ONLINE **Max 8-10 hours**	CORPORATE (DAY-WISE) **3 days · Days 1-3**	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 · This curriculum	5 days · Next step

INSIDE THIS CURRICULUM

What you will achieve	Outcomes and level focus
Who should attend	Audience and prerequisites
Agenda	31 topics · Foundation to Intermediate
Hands-on labs	3 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/clickhouse/clickhouse-intermediate-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 Clickhouse course path, you will be able to:

1

Model data for ClickHouse

Design MergeTree tables, sort keys, partitions and views that keep analytical queries fast.

2

Operate ClickHouse clusters

Run replicated, sharded, secured and monitored clusters with reliable ingestion and backups.

3

Scale real-time analytics

Tune, harden and migrate high-volume analytics workloads to ClickHouse with confidence.

WHAT THE INTERMEDIATE LEVEL ADDS

Replication, sharding, ingestion and operations

Hands-on: Run a replicated, sharded cluster. Deploy a two-shard, two-replica cluster with ClickHouse Keeper, stream events from Kafka and set up monitoring, RBAC and S3 backups.

WHO SHOULD ATTEND

- Regular Clickhouse users who want to automate, integrate and manage it at a deeper level.
- Specialists who support Clickhouse for colleagues or customers.
- Leads preparing Clickhouse for wider use across projects.

BEFORE YOU START

- The Essential level, or equivalent hands-on experience.
- You can already: Design sort keys and partitions, use specialized MergeTree engines, materialized views, skipping indexes, codecs and TTL.
- 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 INTERMEDIATE

Clickhouse Intermediate Training agenda

31 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: Columnar concepts, installation and first queries

Included from Foundation

1. OLAP versus OLTP: columnar storage, vectorized execution and when ClickHouse fits
2. Architecture: MergeTree parts, granules, the sparse primary index and background merges
3. Deployment options: self-managed binaries, Docker, Kubernetes operator and ClickHouse Cloud
4. Installation: clickhouse-server, clickhouse-client, HTTP port 8123 and native port 9000
5. Tables: CREATE TABLE with the MergeTree engine, ORDER BY and PARTITION BY
6. Data types: integers, Decimal, String, LowCardinality, Nullable, DateTime64, Array, Map and JSON
7. Loading data: INSERT, formats such as CSV, JSONEachRow and Parquet, and clickhouse-local
8. Querying: SELECT, WHERE, GROUP BY, ORDER BY, LIMIT BY and aggregate functions
9. External data: s3(), url() and file() table functions
10. First analysis: querying a public dataset such as NYC taxi trips

DAY 2 — Essential: Schema design, MergeTree engines and views

Included from Essential

11. Sort key design: choosing ORDER BY and primary key columns for real query patterns
12. Partitioning: partition keys, dropping and detaching partitions, avoiding too many parts
13. MergeTree family: ReplacingMergeTree, SummingMergeTree, AggregatingMergeTree and CollapsingMergeTree
14. Materialized views: incremental aggregation with -State and -Merge combinators
15. Lookups and refreshes: dictionaries for enrichment and refreshable materialized views
16. Data skipping indexes: minmax, set, bloom_filter and ngrambf_v1
17. Compression codecs: LZ4, ZSTD, Delta, DoubleDelta and Gorilla
18. TTL: row and column expiry and moving data to cold storage volumes
19. Updates and deletes: mutations, lightweight DELETE and deduplication with ReplacingMergeTree
20. Joins: join algorithms, dictionaries versus joins and relevant query settings
21. Client integrations: JDBC and ODBC, Python clickhouse-connect, Grafana and BI tools

DAY 3 — Intermediate: Replication, sharding, ingestion and operations

New at this level

22. Replication: ReplicatedMergeTree, ClickHouse Keeper and insert deduplication
23. Sharding: Distributed tables, sharding keys and ON CLUSTER DDL
24. Streaming ingestion: Kafka table engine, async inserts, batching and ClickPipes in ClickHouse Cloud
25. Query analysis: EXPLAIN, system.query_log, query profiling and trace logs
26. System tables: system.parts, system.merges, system.mutations and system.replicas
27. Access control: users, roles, row policies, settings profiles and quotas
28. Backups: BACKUP and RESTORE to S3 or disk and the clickhouse-backup tool
29. Storage policies: tiered storage, S3-backed disks and local caches
30. Monitoring: Prometheus endpoint, Grafana dashboards and key alerts

31. Troubleshooting: too many parts, memory limits, slow merges and replication lag

PRACTICE

Hands-on labs

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

FOUNDATION

Load and analyse a public dataset

Install ClickHouse, load a public dataset from Parquet files on S3 into a MergeTree table and answer analytical questions with SQL.

ESSENTIAL

Build an events analytics schema

Design an events table with the right sort key, partitions, codecs and TTL, and add a materialized view that keeps daily aggregates.

INTERMEDIATE

Run a replicated, sharded cluster

Deploy a two-shard, two-replica cluster with ClickHouse Keeper, stream events from Kafka and set up monitoring, RBAC and S3 backups.

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 Clickhouse levels, plus 20+ more courses and 50+ tools, at your own pace.	**INR 833/month** billed yearly at INR 9,996
Corporate (day-wise)	3 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 CLICKHOUSE 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/clickhouse/

ENROL OR RESERVE A SEAT
Clickhouse Intermediate Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com

Next step after this level: Clickhouse Advanced Training - Tune ClickHouse for high concurrency and data volume, design for high availability, harden security and migrate analytics workloads.