DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Hazelcast Advanced Training

Design Hazelcast for multi-site resilience, persistence and peak performance, and govern it as an enterprise data platform.

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	45 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/hazelcast/hazelcast-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 Hazelcast course path, you will be able to:

1

Build fast data layers

Use Hazelcast distributed data structures and caching patterns to speed up applications.

2

Process data in real time

Create streaming pipelines with windowing, connectors and fault-tolerant jobs.

3

Operate resilient clusters

Secure, monitor, scale and replicate Hazelcast clusters on servers and Kubernetes.

WHAT THE ADVANCED LEVEL ADDS

Multi-site, persistence, performance and architecture

Hands-on: Resilient multi-site data platform. Connect two Hazelcast clusters with WAN replication, enable persistence, tune performance under load and run a failover and rolling upgrade drill.

WHO SHOULD ATTEND

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

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Build streaming pipelines, use strongly consistent primitives, secure the cluster and run Hazelcast on Kubernetes.
- 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

Hazelcast Advanced Training agenda

45 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: In-memory computing concepts, cluster setup and IMap

Included from Foundation

1. In-memory computing concepts: distributed caching, data grids and stream processing
2. Hazelcast Platform overview: community and enterprise editions and Hazelcast Cloud
3. Architecture: members, clients, partitions (271 by default), partition table and backups
4. Deployment topologies: embedded members versus client-server mode
5. Installation: ZIP distribution, Docker image, Homebrew and the hz CLI
6. Configuration basics: hazelcast.yaml and XML, cluster name, network ports and programmatic config
7. Cluster discovery: multicast, TCP/IP member lists and cloud and Kubernetes discovery
8. Starting a cluster: running members, joining nodes and reading the cluster logs
9. IMap fundamentals: put, get, remove, putIfAbsent and TTL per entry
10. Clients: Java, Python, Node.js, Go and .NET clients connecting to the cluster
11. Management Center: installing it and viewing members, maps and partitions

DAY 2 — Essential: Data structures, caching patterns, queries and persistence hooks

Included from Essential

12. Distributed data structures: IQueue, ITopic, MultiMap, ReplicatedMap, ISet and IList
13. Eviction and expiration: max size policies, LRU and LFU eviction, time-to-live and max idle
14. Near cache: client-side caching, invalidation and consistency trade-offs
15. Caching patterns: cache-aside, read-through and write-through with MapLoader and MapStore
16. Write-behind persistence: batching, coalescing and failure handling with MapStore
17. Serialization: Compact serialization, Portable, IdentifiedDataSerializable and performance
18. Querying with predicates: Predicates API, paging predicates and indexes (sorted, hash, bitmap)
19. Hazelcast SQL: mappings, SELECT queries over maps and Kafka and JDBC sources
20. Entry processors: in-place updates without moving data across the network
21. Spring integration: Spring Boot auto-configuration and Spring Cache with Hazelcast
22. Session clustering: web session replication for Java applications

DAY 3 — Intermediate: Stream processing, CP subsystem, security and Kubernetes

Included from Intermediate

23. Stream processing with the Pipeline API: sources, transforms, sinks and job submission
24. Streaming sources and sinks: Kafka, JDBC, files and change data capture connectors
25. Windowing and aggregation: tumbling, sliding and session windows and event-time watermarks
26. Job fault tolerance: snapshots, processing guarantees and job upgrades
27. CP Subsystem: Raft-based FencedLock, IAtomicLong, ISemaphore and CountDownLatch
28. Split-brain protection: quorum configuration and merge policies after network partitions
29. Security: TLS, member and client authentication and role-based permissions
30. Kubernetes deployment: Helm chart, Hazelcast Platform Operator and Kubernetes discovery

31. Monitoring: metrics with JMX and Prometheus, diagnostics logging and Management Center alerts
32. Troubleshooting: member removal, partition migration delays, GC pauses and slow operations
33. Cluster administration: rolling restarts, graceful shutdown and cluster state changes

DAYS 4-5 **Advanced: Multi-site, persistence, performance and architecture** — New at this level

34. WAN replication: active-active and active-passive replication between data centres
35. Persistence: restarting with data intact and backup of persisted data
36. Tiered storage: memory and disk tiers for large datasets
37. High-Density Memory Store: off-heap storage and avoiding garbage collection pauses
38. Performance tuning: partition count, backup count, serialization choice, threading and JVM settings
39. Benchmarking: load testing with realistic access patterns and measuring latency percentiles
40. Rolling upgrades: version compatibility and upgrading clusters without downtime
41. Disaster recovery: failover clients, blue-green cluster switching and recovery drills
42. Capacity planning: sizing memory, network and member counts for data and backups
43. Vector search and AI workloads: storing embeddings and similarity search in Hazelcast
44. Architecture patterns: microservice caching, event-driven processing and real-time analytics
45. Capstone: a resilient multi-site Hazelcast platform serving cached data and streaming analytics

PRACTICE

Hands-on labs

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

FOUNDATION

Start a cluster and cache data

Run a three-member Hazelcast cluster in Docker, connect Java and Python clients and store, read and expire data in a distributed map.

ESSENTIAL

Build a read-through cache

Put Hazelcast in front of a PostgreSQL database with MapLoader and MapStore, add near cache and indexes and query cached data with SQL.

INTERMEDIATE

Stream and secure on Kubernetes

Deploy Hazelcast with the Operator on Kubernetes, run a Kafka windowed aggregation pipeline and enable TLS and Prometheus monitoring.

ADVANCED

Resilient multi-site data platform

Connect two Hazelcast clusters with WAN replication, enable persistence, tune performance under load and run a failover and rolling upgrade drill.

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 Hazelcast 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 HAZELCAST 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/hazelcast/

ENROL OR RESERVE A SEAT
Hazelcast Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com