

## DataStax Apache Cassandra Administrator Associate Training

OEM: Apache • Duration: 3 Days (24 hrs) • Code: DSX-CASS-ADMIN

### COURSE MODULES & TOPICS

#### Module 1: Cassandra Architecture (Operations Focus)

- Peer-to-peer distributed architecture, ring topology, token ranges, virtual nodes
- Gossip protocol, seed nodes, snitch types and rack/DC awareness
- Replication strategies, RF per datacenter, consistency level math
- Write path: CommitLog, MemTable, SSTable, flush triggers
- Read path: Bloom filters, key cache, partition summary, partition index
- Compaction strategies: STCS (insert-heavy), LCS (read-heavy), TWCS (time-series)
- Tombstones, gc\_grace\_seconds, hinted handoff and read repair

#### Module 2: Cluster Installation & Configuration

- Tarball vs. package installation; cassandra.yaml key parameters
- cluster\_name, seeds, listen\_address, native\_transport\_port, storage\_port, num\_tokens
- JVM configuration: heap size (-Xms/-Xmx), G1GC tuning flags
- Kernel tuning: ulimit (open files, max processes), swap settings, transparent huge pages
- Hardware sizing: CPU, RAM, SSD recommendations, network
- Multi-node cluster bootstrap with seed nodes

#### Module 3: Node Management & Cluster Operations

- Adding a node (bootstrap process, token assignment)
- Removing a node: decommission (graceful) vs. removenode (dead node)
- Replacing a dead node; rebuild after node replacement; cleanup
- nodetool commands: status, ring, info, tpstats, cfstats, compactionstats, repair, decommission, removenode, rebuild, flush, drain, netstats
- Running repair: full vs. incremental, subrange repair
- Multi-datacenter operations and consistency during DC failover

#### Module 4: Performance Monitoring & Tuning

- cassandra-stress: write, read, mixed workload benchmarking
- nodetool tpstats (thread pool stats), nodetool cfstats/tablestats
- JMX metrics, GC log analysis, GC pressure diagnosis
- Compaction monitoring: nodetool compactionstats, disableautocompaction

- Hot partition detection, read/write latency, dropped message alerts

## **Module 5: Backup, Restore & Data Management**

- Snapshot-based backups: nodetool snapshot, incremental backups
- Restore from snapshots: manual SSTable copy and sstableloader
- SSTable utilities: sstablescrub, sstableupgrade, sstabledump
- DSBulk and COPY TO/FROM for data operations
- TTL and tombstone management best practices

## **Module 6: Security**

- Authentication: PasswordAuthenticator vs. AllowAllAuthenticator
- Authorization: CassandraAuthorizer with GRANT/REVOKE CQL commands
- Role management: CREATE ROLE, DROP ROLE, LIST ROLES
- TLS/SSL: internode encryption (none/dc/rack/all), client-to-node encryption
- Auditing basics and DSE audit logging