

Q1. What type of database is Apache Cassandra (relational, document, key-value, graph, time-series) and what are its strengths?

Answer:

Q2. Explain how Apache Cassandra handles ACID properties or the consistency model it provides.

Answer:

Q3. How does indexing work in Apache Cassandra? What index types are available and how do you choose one?

Answer:

Q4. What is the data model in Apache Cassandra and how do you design a schema for a typical use case?

Answer:

Q5. How does Apache Cassandra handle scaling: vertical scaling, horizontal sharding, or replication?

Answer:

Q6. Describe the read/write performance characteristics of Apache Cassandra and how to tune them.

Answer:

Q7. How do you handle transactions in Apache Cassandra? What isolation levels or guarantees are supported?

Answer:

Q8. What backup and point-in-time recovery strategies does Apache Cassandra support?

Answer:

Q9. How do you monitor and diagnose slow queries or performance issues in Apache Cassandra?

Answer:

Q10. What are common anti-patterns or pitfalls to avoid when using Apache Cassandra in production?

Answer:
