

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

Answer:

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

Answer:

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

Answer:

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

Answer:

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

Answer:

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

Answer:

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

Answer:

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

Answer:

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

Answer:

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

Answer:
