55 MongoDB interview questions to ask your candidates

Questions

1. What is MongoDB and how does it differ from traditional relational databases?
2. Explain the concept of documents in MongoDB.
3. What is a collection in MongoDB?
4. How does MongoDB handle schema design?
5. What are the key features of MongoDB?
6. Explain the difference between embedding and referencing in MongoDB.
7. What is sharding in MongoDB and why is it important?
8. How does indexing work in MongoDB?
9. What is the purpose of the _id field in MongoDB documents?
10. Explain the concept of BSON in MongoDB.
11. What are replica sets in MongoDB and why are they used?
12. How do you perform CRUD operations in MongoDB?
13. What is the aggregation framework in MongoDB?
14. Explain the concept of capped collections in MongoDB.
15. How does MongoDB ensure data consistency?
16. What are the main use cases for MongoDB?
17. How do you handle transactions in MongoDB?
18. Explain the difference between $lookup and $unwind operators in MongoDB.
19. How would you optimize a slow-performing MongoDB query?
20. Explain the concept of write concerns in MongoDB and when you might use different levels.
21. How does MongoDB handle concurrency control?
22. What are the key considerations when designing a MongoDB schema for a social media application?
23. How would you approach backing up and restoring a large MongoDB database?
24. Describe a situation where you would choose MongoDB over a relational database, and why?
25. How would you monitor and troubleshoot performance issues in a production MongoDB deployment?
26. How does MongoDB's aggregation pipeline work, and can you provide a use case for it?
27. What strategies would you use to secure a MongoDB database, and what are some common vulnerabilities?
28. Can you explain the process of data modeling in MongoDB and how it differs from relational models?
29. What are the differences between the various types of indexes in MongoDB, and when would you use each type?
30. How would you implement a backup strategy for a MongoDB cluster?
31. What techniques can you use to scale a MongoDB application, both vertically and horizontally?
32. Describe how you would handle a data migration from a relational database to MongoDB.
33. What tools or methods do you use for monitoring the health and performance of a MongoDB instance?
34. How do MongoDB's transactions work, and what limitations exist compared to traditional databases?
35. Can you discuss the role of the oplog in MongoDB and how it facilitates replication?
36. What are the different types of indexes available in MongoDB, and when would you use each type?
37. How would you go about identifying and resolving a slow-performing query in MongoDB?
38. What factors should you consider when creating an index in MongoDB?
39. Can you explain the concept of index cardinality and why it is important in MongoDB?
40. How do you monitor and analyze index performance in MongoDB?
41. What is the role of compound indexes in MongoDB, and how do they differ from single field indexes?
42. How do you decide which fields to index in a MongoDB collection?
43. What are the potential downsides of over-indexing in MongoDB?
44. What is the role of the primary node in a MongoDB replica set?
45. How does MongoDB handle automatic failover in a replica set?
46. Can you explain the difference between synchronous and asynchronous replication in MongoDB?
47. What are the considerations for choosing a shard key in MongoDB?
48. How does MongoDB ensure data balance across shards?
49. What happens if a shard in a MongoDB cluster goes down?
50. Explain the concept of 'read preference' in MongoDB and its significance in a sharded environment.
51. How can you add or remove shards from a MongoDB cluster?
52. What are some challenges you might face when implementing sharding in MongoDB?
53. How does the chunk migration process work in a sharded MongoDB cluster?