62 SQL Coding interview questions to ask your applicants

Questions

1. How would you optimize a slow-performing SQL query?

2. Explain the difference between INNER JOIN and LEFT JOIN with an example.

3. Write a query to find the second highest salary in the employees table.

4. How would you handle duplicate records in a table?

5. Describe a situation where you would use a subquery instead of a JOIN.

6. Write a query to pivot data from rows to columns.

7. Explain the concept of indexing and when you would use it.

8. How would you implement a rolling average calculation in SQL?

9. Write a query to find employees who have never been assigned to a project.

10. Describe how you would design a database schema for a simple e-commerce website.

11. What is a primary key, and why is it important in a database?

12. How would you explain the concept of normalization in databases?

13. Can you describe what a foreign key is and how it works?

14. What are the differences between SQL and NoSQL databases?

15. How would you explain the concept of ACID properties in database transactions?

16. What steps would you take to troubleshoot a failing SQL query?

17. How do you handle missing or NULL values in your SQL queries?

18. Can you explain what a view is in SQL and how it might be used?

19. Explain the concept of window functions and provide an example of when you'd use them.

20. How would you design a query to find the top 3 products in each category by sales volume?

21. Describe a scenario where you'd use a recursive CTE and write a sample query.

22. Explain the difference between RANK, DENSE_RANK, and ROW_NUMBER functions with an example.

23. How would you implement a query to detect and resolve circular references in a hierarchical data structure?

24. Write a query to find the median salary for each department in a company.

25. Explain how you would use SQL to implement a simple recommendation system based on user purchase history.

26. Describe how you would design and query a slowly changing dimension (SCD) Type 2 in a data warehouse.

27. Write a query to find customers who have made purchases on consecutive days.

28. How would you implement a custom aggregation function in SQL?

29. Explain the concept of table partitioning and when you would use it to optimize query performance.

30. Write a query to pivot data dynamically based on user input for column names.

31. How would you handle versioning of database objects in a collaborative development environment?

32. Describe a situation where you'd use a materialized view instead of a regular view, and explain the trade-offs.

33. Write a query to identify and resolve data inconsistencies across multiple related tables.

34. How would you explain the concept of database normalization to a non-technical person?

35. What are the trade-offs between using a star schema versus a snowflake schema in data warehouse design?

36. How would you design a database to handle versioning of documents?

37. Explain the concept of database sharding and when you might use it.

38. How would you design a database schema for a social media platform's messaging system?

39. What is eventual consistency in distributed databases, and how does it differ from strong consistency?

40. How would you design a database to handle hierarchical data, such as an organizational structure or a product category tree?

41. Can you explain what query execution plans are and how they can be used to optimize SQL queries?

42. How would you approach optimizing a query that involves multiple joins on large tables?

43. Describe the role of indexing in query optimization and how you would determine which columns to index.

44. Explain the concept of query hints and provide an example of when you might use them.

45. What is a covering index, and how does it help in query optimization?

46. How would you optimize a SELECT query that pulls data from a table with millions of rows?

47. Describe the impact of normalization on query performance and how denormalization might be used for optimization.

48. How do you use the EXPLAIN command to analyze and optimize a query?

49. Can you discuss the importance of statistics in query optimization and how you keep them up to date?

50. Explain the concept of parameter sniffing and how it can affect query performance.

51. What strategies would you employ to optimize a complex query that involves subqueries and derived tables?

52. How would you optimize a query that performs aggregations on large data sets?

53. You've noticed that a nightly batch job is taking longer to complete each week. How would you investigate and address this performance issue?

54. A customer reports that their order history sometimes shows duplicate entries. How would you approach diagnosing and fixing this data inconsistency?

55. Your team needs to implement a loyalty points system that awards points based on customer purchase history. How would you design the database schema and write a query to calculate points?

56. You're tasked with creating a report that shows daily sales trends over the past year, including running totals and week-over-week comparisons. How would you approach this?

57. A large table in your database is causing storage issues. How would you implement archiving for old records while keeping them queryable?

58. Your e-commerce platform needs to handle peak sales during flash sales. How would you optimize the database to handle sudden spikes in concurrent transactions?

59. You need to merge data from two systems with different schemas following a company acquisition. How would you approach this data integration challenge?

60. A critical query is timing out during peak hours. Walk me through your process for identifying the bottleneck and optimizing the query.