66 PostgreSQL interview questions to ask to hire top developers

Questions

1. What is PostgreSQL and how does it differ from other relational databases?

2. Can you explain the concept of ACID properties in PostgreSQL?

3. How would you create a new database and table in PostgreSQL?

4. What is the difference between TRUNCATE and DELETE commands?

5. Explain the purpose of indexes in PostgreSQL and when you would use them.

6. How do you perform a simple SELECT query with conditions in PostgreSQL?

7. What are PostgreSQL schemas and why are they useful?

8. Can you describe the difference between INNER JOIN and LEFT JOIN?

9. How would you backup a PostgreSQL database?

10. What is the purpose of the VACUUM command in PostgreSQL?

11. Explain the concept of transactions in PostgreSQL.

12. How do you grant and revoke user permissions in PostgreSQL?

13. What are PostgreSQL extensions and can you name a few common ones?

14. How would you optimize a slow-running query in PostgreSQL?

15. Can you explain the difference between a view and a materialized view in PostgreSQL?

16. Can you explain the difference between a primary key and a unique constraint in PostgreSQL?

17. How would you approach optimizing a slow query in PostgreSQL?

18. What is the purpose of the SERIAL data type in PostgreSQL?

19. How does PostgreSQL handle concurrent transactions?

20. What is the difference between a view and a materialized view in PostgreSQL?

21. Can you explain what a foreign key is and how it's used in PostgreSQL?

22. What is the purpose of the EXPLAIN command in PostgreSQL?

23. How would you handle database migrations in a PostgreSQL environment?

24. Can you explain how to use the CTE (Common Table Expressions) in PostgreSQL and provide an example?

25. What are window functions, and how are they used in PostgreSQL?

26. How do you handle JSON data in PostgreSQL, and what are the advantages of using JSONB?

27. What is the process of indexing expressions in PostgreSQL, and when would you use it?

28. Explain partitioning in PostgreSQL and provide scenarios where it would be useful.

29. How do you monitor and troubleshoot performance issues in a PostgreSQL database?

30. Can you describe the role of pg_stat_statements and how it can be used for performance tuning?

31. What are the different types of replication available in PostgreSQL, and when would you use each?

32. How do you implement full-text search in PostgreSQL, and what are its benefits?

33. What is the difference between synchronous and asynchronous replication in PostgreSQL?

34. How would you perform a bulk data insertion in PostgreSQL while ensuring minimal impact on performance?

35. Can you explain the concept of table inheritance and its use cases in PostgreSQL?

36. Can you explain the concept of B-tree indexes in PostgreSQL and when they are most effective?

37. How would you decide whether to create a single-column or multi-column index in PostgreSQL?

38. What is the difference between a partial index and a normal index in PostgreSQL?

39. How does PostgreSQL handle index-only scans, and when are they beneficial?

40. What are GIN indexes in PostgreSQL, and for what types of data are they most suitable?

41. How would you approach index maintenance in PostgreSQL to ensure optimal performance?

42. Can you explain the concept of covering indexes in PostgreSQL and how they can improve query performance?

43. How would you identify and resolve a query that's causing high CPU usage in PostgreSQL?

44. Can you explain the concept of query plan caching in PostgreSQL and its impact on performance?

45. What strategies would you employ to optimize a query involving multiple joins on large tables?

46. How does the PostgreSQL query planner work, and how can you influence its decisions?

47. Can you describe the use of parallel query execution in PostgreSQL and when it's most beneficial?

48. What are the pros and cons of using subqueries versus joins for query optimization?

49. How would you optimize a query that involves heavy use of aggregate functions?

50. Can you explain the concept of query rewriting in PostgreSQL and provide an example?

51. What role does statistics collection play in query optimization, and how would you ensure it's up to date?

52. How would you approach optimizing a query that uses a lot of temporary tables or CTEs?

53. Can you describe scenarios where denormalization might improve query performance in PostgreSQL?

54. What are the considerations for optimizing queries involving text search in PostgreSQL?

55. How would you optimize a query that involves complex date/time calculations?

56. Can you explain the concept of query hints in PostgreSQL and when you might use them?

57. How would you handle a situation where a critical query is consistently timing out? What steps would you take to diagnose and resolve the issue?

58. Imagine you are tasked with migrating a large PostgreSQL database to a new server with minimal downtime. How would you approach this task?

59. If a user reports that they cannot access a database they should have permissions for, how would you troubleshoot and resolve the issue?

60. Suppose you discover that a recently added index is not improving query performance as expected. What actions would you take to investigate?

61. You need to implement a new feature that requires changes to multiple database tables. How would you manage these changes while ensuring data integrity?

62. A senior developer asks for a specific database configuration to optimize performance. How would you evaluate and implement their request?

63. Imagine that you identified a significant increase in disk space usage after a batch job. What steps would you take to investigate and mitigate the issue?