

92 DB2 Interview Questions to Hire Top Engineers

Questions

1. Can you describe what DB2 is, in simple terms?
2. What are the different types of DB2 licenses available, and how do they differ?
3. Explain the difference between a database and a database instance in DB2.
4. What is a tablespace in DB2, and why is it important?
5. How does DB2 handle concurrency, and what are some common locking mechanisms?
6. What are the different data types available in DB2, and when would you use each one?
7. Explain the purpose of indexes in DB2 and how they improve query performance.
8. How do you back up and restore a DB2 database?
9. What is the role of the DB2 catalog?
10. Describe the process of creating a table in DB2, including specifying data types and constraints.
11. What are some common DB2 commands you use on a daily basis?
12. How can you monitor the performance of a DB2 database?
13. What are user-defined functions (UDFs) in DB2, and when would you use them?
14. Explain the concept of normalization in database design and its benefits.
15. What is the difference between clustered and non-clustered indexes in DB2?
16. How do you troubleshoot common DB2 errors?
17. Describe the purpose of the DB2 command line processor (CLP).
18. How can you import and export data in DB2?
19. What is the importance of transaction management in DB2?
20. Explain the concept of isolation levels in DB2 and their impact on concurrency.
21. How can you optimize a slow-running query in DB2?
22. What are the different types of joins in SQL, and how do they work in DB2?
23. How do you handle security in DB2, including user authentication and authorization?
24. Describe the process of upgrading a DB2 database to a newer version.
25. What are some of the new features in the latest version of DB2 that you find particularly useful?
26. How would you optimize a slow-running DB2 query, and what tools would you use to identify the bottleneck?
27. Explain the difference between clustered and non-clustered indexes in DB2, and when would you use each?
28. Describe a situation where you would use a materialized query table (MQT) in DB2, and how it improves performance.
29. What are the different isolation levels in DB2, and what are the trade-offs between them in terms of concurrency and data consistency?
30. Explain the concept of 'explain plan' in DB2. How do you read and interpret it to optimize a query?
31. How do you handle deadlocks in DB2, and what strategies can you implement to prevent them?
32. Describe the purpose of DB2's RUNSTATS command, and how frequently should you run it on different types of tables?
33. What is the difference between online and offline backups in DB2, and what are the advantages and disadvantages of each?
34. Explain how you would implement auditing in DB2 to track changes to sensitive data.
35. Describe the process of migrating a DB2 database to a new server, including pre-migration checks and post-migration validation.
36. How can you monitor DB2 database performance in real-time, and what key metrics would you track?
37. What are the different types of locks in DB2, and how do they impact concurrency?
38. Explain how you would implement data partitioning in DB2 to improve query performance and manage large tables.
39. Describe a scenario where you would use stored procedures in DB2, and what are the benefits of using them?
40. How would you troubleshoot a DB2 database that is experiencing high CPU utilization?
41. Explain the purpose of the DB2 catalog tables, and how can you use them to gather information about the database?
42. How can you implement row-level security in DB2 to restrict access to specific rows based on user roles?
43. Describe the process of recovering a DB2 database after a media failure.
44. Explain the purpose of the DB2 governor, and how can you use it to prevent runaway queries from impacting database performance?
45. How would you implement a rolling upgrade of a DB2 database to a newer version, minimizing downtime?
46. How do you optimize a complex DB2 query involving multiple joins and subqueries for performance?
47. Explain the concept of DB2's autonomic computing features and how they contribute to database administration.
48. Describe the different types of locking mechanisms available in DB2 and their impact on concurrency.
49. How would you implement a data partitioning strategy in DB2 to improve query performance and manageability?
50. What are the considerations when designing a high availability and disaster recovery solution for a DB2 database?
51. Explain the role of the DB2 catalog and how it's used by the database manager.
52. How can you monitor and troubleshoot performance bottlenecks in a DB2 environment?
53. Describe the steps involved in upgrading a DB2 database to a newer version, including pre- and post-upgrade tasks.
54. What are the advantages and disadvantages of using DB2 BLU Acceleration for analytical workloads?
55. Explain the different isolation levels in DB2 and how they affect data consistency and concurrency.
56. How do you implement auditing in DB2 to track data access and modifications for security and compliance?
57. Describe how you would handle a situation where a DB2 database is experiencing severe performance degradation due to locking contention.
58. What are the key differences between DB2 on-premises and DB2 on Cloud, and when would you choose one over the other?
59. How can you use DB2's explain facility to analyze query execution plans and identify optimization opportunities?
60. Describe the process of creating and managing user-defined functions (UDFs) in DB2.
61. What are the security considerations when designing and implementing a DB2 database application?
62. How do you back up and restore a DB2 database, and what are the different types of backup options available?
63. Explain how you would use DB2's workload management (WLM) features to prioritize and manage database resources.
64. What are the challenges and best practices for migrating data from other database systems to DB2?
65. How does DB2 handle data compression, and what are the benefits of using compression?
66. Describe the process of tuning DB2 memory parameters to optimize database performance.
67. What are the different types of indexes in DB2, and how do you choose the appropriate index type for a given query?
68. Explain the concept of DB2 pureXML and how it can be used to store and query XML data.
69. How do you monitor and manage DB2 HADR (High Availability Disaster Recovery) environments?
70. Explain the concept of 'dirty reads' in DB2 and how to prevent them. Can you illustrate with an example?
71. What are the different types of locking mechanisms available in DB2, and when would you choose one over the other?
72. Describe the DB2 RUNSTATS command and its importance in query optimization. What are some best practices for using it?
73. Explain the difference between clustered and non-clustered indexes in DB2. How does each impact query performance?
74. What is a DB2 package? How do packages improve performance, and how are they managed?
75. Describe the DB2 autonomic computing features. How do they help in database administration and performance tuning?
76. How can you monitor DB2 performance in real-time? What tools and techniques can you use to identify bottlenecks?
77. What are the different isolation levels in DB2, and how do they affect concurrency and data integrity?
78. Explain the concept of DB2's 'explain plan' and how it is used for query optimization. How do you interpret an explain plan?
79. Describe the DB2 HADR (High Availability Disaster Recovery) feature. How does it work, and what are the different HADR configurations?
80. What are the different types of DB2 utilities? How can they be used for database maintenance and administration?
81. Explain the concept of DB2's data partitioning. How does it improve performance and scalability?
82. How do you handle deadlocks in DB2? What strategies can you use to prevent or resolve deadlocks?
83. Describe the DB2 BLU Acceleration feature and how it improves query performance for analytical workloads.
84. Explain the concept of DB2's 'materialized query tables' (MQTs) and how they can be used to improve query performance. Give an example.
85. What are the different types of DB2 user-defined functions (UDFs)? How can they be used to extend DB2's functionality?
86. How do you implement security in DB2? What are the different authentication and authorization mechanisms available?
87. Describe the DB2 pureScale environment. How does it provide scalability and high availability?
88. Explain how to recover a DB2 database after a failure. What are the different recovery options available?
89. How can you integrate DB2 with other systems or applications? What are the different integration options available?
90. What are the performance considerations when designing a DB2 database schema? How can you optimize the schema for query performance?
91. Explain the purpose of DB2 Governor and how it can be used to manage database resources.