



SNS COLLEGE OF ENGINEERING

Kurumbapalayam (Po), Coimbatore – 641 107

AN AUTONOMOUS INSTITUTION



Accredited by NAAC–UGC with ‘A’ Grade

Approved by AICTE, New Delhi & Affiliated to Anna University, Chennai

23ITB204 – Modern Database Management Systems

Puzzle Unit – IV

Puzzle 1: RAID Levels

Scenario: A company needs to store sensitive customer data and ensure both performance and redundancy.

Question: Given the options of RAID 0, RAID 1, RAID 5, and RAID 10, which RAID configuration would you recommend for this company? Justify your choice based on performance and data safety.

Puzzle 2: File Organization

Scenario: You are tasked with designing a file organization strategy for a library database that manages books, authors, and borrowers.

Question: Describe two different file organization methods (e.g., sequential and indexed) you would consider, and discuss the advantages and disadvantages of each in this context.

Puzzle 3: Data Dictionary

Scenario: A new database is being implemented, and a data dictionary is needed.

Question: What essential components should be included in the data dictionary? Explain how it will help in managing the database effectively.

Puzzle 4: Column-Oriented vs. Row-Oriented Storage

Scenario: An analytics company needs to run complex queries on large datasets frequently.

Question: Would you recommend a column-oriented storage system or a row-oriented storage system? Provide reasons based on performance and query efficiency.

Puzzle 5: Indexing and Hashing

Scenario: You have a large dataset of user records and need to optimize search queries.

Question: Explain how indexing and hashing differ in their approach to speeding up searches. In which scenario would each be most effective?

Puzzle 6: B+ Tree vs. B Tree

Scenario: A database system is using a B tree for indexing.

Question: What are the advantages of switching to a B+ tree? Discuss how this change would impact performance and functionality.

Puzzle 7: Static vs. Dynamic Hashing

Scenario: A web application requires efficient storage for user sessions, with a variable number of sessions expected daily.

Question: Would you implement static hashing or dynamic hashing? Justify your decision based on the expected workload and data behavior.

Puzzle 8: Query Processing Algorithms

Scenario: You need to implement a query that retrieves data from a large relational database.

Question: Discuss the algorithms you would consider for selection, sorting, and join operations. How would you choose the most efficient algorithm for each operation?

Puzzle 9: Query Optimization Heuristics

Scenario: A complex query is taking too long to execute on a production database.

Question: Identify and describe three heuristics that can be applied to optimize the query. How do these heuristics impact query execution plans?

Puzzle 10: Cost Estimation in Query Optimization

Scenario: You are developing a query optimizer for a database management system.

Question: Explain the factors that influence cost estimation for query execution. How would you implement a cost estimation strategy in your optimizer?