

SRM VALLIAMMAI ENGINEERING COLLEGE

(An Autonomous Institution)

SRM Nagar, Kattankulathur – 603 203

FACULTY OF INFORMATION TECHNOLOGY

QUESTION BANK



I SEMESTER

1924104 – NEXT GENERATION DATABASES

Regulation – 2019

Academic Year 2021 – 2022 (ODD SEMESTER)

Prepared by

Dr. S. JEYALAKSHMI, Assistant Professor (Sel. G)



SRM VALLIAMMAI ENGINEERING COLLEGE

(An Autonomous Institution)
SRM Nagar, Kattankulathur – 603 203.



FACULTY OF INFORMATION TECHNOLOGY
QUESTION BANK

SUBJECT : 1924104 – NEXT GENERATION DATABASES

SEM / YEAR : I Sem / I Year

UNIT I - INTRODUCTION

Database Revolutions - System Architecture - Relational Database - Database Design- Data Storage - Transaction Management - Data warehouse and Data Mining - Information Retrieval.

PART – A

1	What are the issues of traditional file-based systems that make DBMS a superior alternative?	BTL 2	Understanding
2	List some examples of open source and commercial Relational DBMSs?	BTL 4	Analyzing
3	What is a database model?	BTL 1	Remembering
4	Give examples for common database models?	BTL 1	Remembering
5	How do you choose a database model?	BTL 2	Understanding
6	What is ER modeling?	BTL 2	Understanding
7	What is ACID property of transactions?	BTL 4	Analyzing
8	What are the different levels of data abstraction?	BTL 2	Understanding
9	What is Data Warehousing?	BTL 1	Remembering
10	What is “deadlock”?	BTL 3	Applying
11	List the isolation levels?	BTL 6	Creating
12	Discuss about relational data model.	BTL 5	Evaluating
13	Define Information Retrieval system. Show how it differs from the database system	BTL 6	Applying
14	Illustrate the situation to roll back a transaction.	BTL 1	Remembering
15	Illustrate the three major eras of Database Technology with example.	BTL 3	Creating
16	Define transaction.	BTL 1	Remembering
17	Compare information retrieval vs. DBMS.	BTL 1	Remembering
18	What is the purpose of normalization in DBMS?	BTL 4	Analyzing
19	What is data warehouse?	BTL 3	Applying
20	Define data mining.	BTL 5	Evaluating

PART – B

1	With the help of the block diagram, describe the basic architecture of a Database Management System.	BTL 5	Evaluating
2	List and explain the categories of data models in detail.	BTL 1	Remembering
3	Describe the extended ER features available in Entity Relationship diagram with diagrammatic notations and example	BTL 1	Remembering
4	What are isolation levels and explain elaborately the various isolation levels?	BTL 1	Remembering
5	Explain: (i) Purpose of database (7) (ii) Data independence (6)	BTL 2	Understanding
6	Write notes on (i) Primary key (7) (ii) Database Schema (6)	BTL2	Understanding
7	(i) Briefly describe about Views of data.(7) (ii) What are the functions of database administrator?(6)	BTL 4	Analyzing
8	Analyze the structure of Relational database and its various relational operations with example (13)	BTL 2	Understanding
9	Examine the following with examples: (i) DDL (7) (ii) DML (6)	BTL 6	Creating
10	What are Normal Forms. Explain the types of Normal form with an example. (13)	BTL 3	Applying
11	(i) Illustrate the architecture of a typical data mining system. (7) (ii) Classify the different types of data on which mining can be performed. (6)	BTL 1	Remembering
12	What is a data warehouse? With the help of a neat sketch, explain the various components in a data warehousing system. (2+11)	BTL 4	Analyzing
13	Summarize about Information Retrieval and Relevance ranking? (13)	BTL 4	Analyzing
14	Describe the important models of information retrieval. (13)	BTL 2	Understanding

PART C

1	A Car rental company maintains a database for all vehicles in its current fleet. For all vehicles it includes the vehicle identification number, license number, manufacturer, model, date of purchase and color. Special data are included for certain types of vehicles. Trucks : Cargo capacity Sports Cars : horsepower, renter age requirement Vans : number of passengers Off-road vehicles : ground clearance, drivetrain(four or two wheel driven) Construct an ER model for the car rental company database	BTL 5	Evaluating
2	(i) Discuss the various disadvantages in the file system and explain how it can be overcome by the database system. (8) (ii) What are the characteristics of database? Explain. (7)	BTL 6	Creating
3	Compare the advantages and disadvantages of: (i) star schema (5) (ii) snowflake schema(5) (iii)fact constellation schema (5)	BTL 6	Creating
4	Discuss about the querying methods used in Information retrieval. (15)	BTL 5	Analyzing

UNIT II - DOCUMENT DATABASES

Big Data Revolution – CAP Theorem – Birth of NoSQL - Document Database – XML Databases - JSON Document Databases - Graph Databases.

PART – A

Q.No	Questions	BT Level	Competence
1	What is NoSQL?	BTL 3	Applying
2	List some examples for Document Database	BTL 4	Analyzing
3	What are the different types of NoSQL databases?	BTL 3	Applying
4	What are the features of NoSQL?	BTL 1	Remembering
5	What are the architectural characteristics that make Dynamo an architectural model for few nonrelational databases.	BTL 1	Remembering
6	What is database sharding?	BTL 2	Understanding
7	List the drawbacks of Sharding strategy.	BTL 3	Applying
8	What are the goals of Amazon's Dynamo key-value store.	BTL 4	Analyzing
9	Explain CAP Theorem.	BTL 2	Understanding
10	Define Map Reduce	BTL 6	Creating
11	What is Partition tolerance in CAP?	BTL 1	Remembering
12	What are A and P in CAP and the difference between them?	BTL 6	Creating
13	Define Document database.	BTL 5	Evaluating
14	List the tools and standards supporting XML.	BTL 1	Remembering
15	Compare JSON and XML	BTL 2	Understanding
16	Explain JSON format	BTL 1	Remembering
17	Define Graph database.	BTL 2	Understanding
18	Differentiate RDBMS and Graph database.	BTL 4	Analyzing
19	What is Graph Compute Engine?	BTL 1	Remembering
20	List the significant Graph Compute Engines available.	BTL 5	Evaluating

PART – B

1	i) What are the advantages of NoSQL over traditional RDBMS? (7) ii) Describe the utilities and applications built on top of core Hadoop (6)	BTL 2	Understanding
2	Explain elaborately about Big data revolution.	BTL 3	Applying
3	Explain in detail about three major software layers developed by Google in Google software stack	BTL 1	Remembering
4	Illustrate with neat diagram explain the Hive architecture	BTL 6	Creating

5	Compare Hive developed by Facebook and Pig Latin developed by Yahoo	BTL 1	Remembering
6	Illustrate the transition from single web server and database server to multiple web servers, Memcached servers and read only database replicas without sharding and with sharding.	BTL 2	Understanding
7	Explain elaborately with neat diagram the consistent hashing characteristics of Dynamo.	BTL 4	Analyzing
8	Describe with neat sketch how Dynamo will trade off consistency, read performance and write performance with Tunable Consistency	BTL 1	Remembering
9	Explain with neat diagram the generic architecture of XML database	BTL 1	Remembering
10	Describe JSON Document database briefly.	BTL 2	Understanding
11	Compare XML and JSON file format with example.	BTL 5	Evaluating
12	Explain Data Models in Document Databases with examples.	BTL 4	Analyzing
13	Write short notes on: (i) RDF and SPARQL (7) (ii) Property Graphs and Cypher (6)	BTL 4	Analyzing
14	Explain (i) Graph Database Internals (7) (ii) Graph Compute Engines (6)	BTL 3	Applying
PART C			
1	Explain how Dynamo and ACID RDBMS are mapped to CAP Theorem with neat diagram	BTL 6	Creating
2	Describe briefly the two major iterations: Hadoop 1.0 and Hadoop 2.0 of Hadoop architecture.	BTL 5	Evaluating
3	Explain elaborately the JSON based database system	BTL 5	Evaluating
4	What Are the Major Advantages of Using a Graph Database?	BTL 6	Creating

UNIT III – COLUMN DATABASES

Column Databases — Data Warehousing Schemas - Columnar Alternative - Sybase IQ - CStore and Vertica - Column Database Architectures - SSD and In-Memory Databases — In-Memory Databases- Berkeley Analytics Data Stack and Spark.

PART – A

Q.No	Questions	BT Level	Competence
1	What is the difference between columnar and row-based databases?	BTL 2	Understanding
2	What are OLTP and OLAP?	BTL 2	Understanding
3	What is the difference between OLTP and OLAP?	BTL 1	Remembering
4	What is a columnar database?	BTL 1	Remembering
5	Compare Columnar database with row-oriented database	BTL 2	Understanding
6	What is CRUD?	BTL 2	Understanding

7	Define SSD.	BTL 1	Remembering
8	Compare single-level cell SSD and multi-level cell SSD.	BTL 5	Evaluating
9	Define Flash SSD	BTL 1	Remembering
10	Explain Aerospike.	BTL 6	Creating
11	Explain how ACID transactional model is supported by VoltDB	BTL 4	Analyzing
12	Write notes on H-store design	BTL 3	Applying
13	What are the circumstances in Redis in-memory database system that the applications have to wait for IO to complete.	BTL 5	Evaluating
14	Define Enhanced Hybrid Columnar Compression(EHCC).	BTL 1	Remembering
15	List the two changes to traditional database architecture that an in-memory system addresses.	BTL 3	Applying
16	What are the techniques used by in-memory databases to ensure that there is no data loss?	BTL 4	Analyzing
17	Explain the two disk files used by Redis for persistence.	BTL 1	Remembering
18	What are the components of BDAS?	BTL 3	Applying
19	Define RDD.	BTL 4	Analyzing
20	Define Spark streaming and Graphx	BTL6	Creating

PART – B

1	i) How does column-oriented NoSQL differ from document-oriented? (7) ii) Write notes on Sybase IQ and C-Store (6)	BTL 2	Understanding
2	Explain the two big advantages to the columnar architecture database	BTL 1	Remembering
3	Explain the various types of Data Warehouse Schema with diagram.	BTL 1	Remembering
4	Describe with neat diagram the generic columnar database architecture.	BTL 6	Creating
5	Explain with diagram the key disadvantage of the columnar architecture.	BTL 1	Remembering
6	Explain with neat diagram the columnar database with multiple projection	BTL 2	Understanding
7	Briefly explain the Columnar Write Penalty with neat sketch.	BTL 4	Analyzing
8	Explain with neat diagram the Timesten in-memory database system architecture.	BTL 5	Evaluating
9	Describe elaborately the architectural components of Redis in-memory database system.	BTL 3	Applying
10	Explain with neat diagram illustrating how partitioning and replication affect concurrency in VoltDB.	BTL 4	Analyzing
11	Illustrate the essential components of Oracle 12c in-memory column store architecture.	BTL 2	Understanding
12	Explain briefly the Berkeley Analytics Stack and its core components.	BTL 4	Analyzing
13	Compare the in-memory databases.	BTL 1	Remembering
14	Briefly describe the Spark Architecture and the essential features of Spark processing.	BTL 3	Applying

PART C

1	(i) Write notes on SSD enabled Databases. (8) (ii) Explain Columnar technology in other Databases. (7)	BTL6	Creating
2	Compare Columnar and row-oriented storages of Columnar architecture with appropriate examples	BTL 5	Evaluating
3	Illustrate the persistent architecture of SAP HANA with diagram	BTL 5	Evaluating
4	Explain how Hadoop elements like YARN and HDFS interact with Spark and Other elements of BDAS with diagram	BTL 6	Creating

UNIT IV - DISTRIBUTED DATABASES

Distributed Database Patterns — Distributed Relational Databases - Non-relational Distributed Databases - MongoDB - Sharding and Replication - HBase- Cassandra - Consistency Models — Types of Consistency - Consistency in MongoDB - HBase Consistency - Cassandra Consistency.

PART – A

Q.No	Questions	BT Level	Competence
1	What is a distributed database?	BTL 2	Understanding
2	What is database partitioning?	BTL 3	Applying
3	Define Region Replicas.	BTL 1	Remembering
4	What is Sharding in MongoDB?	BTL 4	Analyzing
5	What are the types of Distributed Database?	BTL1	Remembering
6	Explain the classifications of database server	BTL1	Remembering
7	What is Cassandra?	BTL 2	Understanding
8	List the benefits of using Cassandra.	BTL 1	Remembering
9	What is Cassandra Data Model?	BTL 2	Understanding
10	How does Cassandra write?	BTL 3	Applying
11	How do the snitches may be configured in Cassandra?	BTL 5	Evaluating
12	What is regionserver?	BTL 1	Remembering
13	What is CQL?	BTL 6	Creating
14	Compare MongoDB with Cassandra.	BTL 5	Evaluating
15	Does Cassandra support ACID transactions?	BTL 1	Remembering
16	List the levels of data locality in HBase.	BTL 2	Understanding
17	List the consistency levels for read operations in Cassandra.	BTL 6	Creating
18	List the consistency levels for write operations in Cassandra.	BTL 4	Analyzing

19	What is replication factor in Cassandra?	BTL 4	Analyzing
20	Compare HBase & Cassandra	BTL 3	Applying

PART – B

1	State the differences between a node, a cluster, and a data center in Cassandra.	BTL 1	Remembering
2	Explain how the three database architectures rely on monolithic database server with neat diagram	BTL 2	Understanding
3	Write short notes on (i) Log-based replication (ii) Distributed database models	BTL 2	Understanding
4	Illustrate the Sharding mechanisms in MongoDB and compare them	BTL 3	Applying
5	Describe elaborately three member replica set in MongoDB and behavior of Read and Write operations in MongoDB applications	BTL 1	Remembering
6	Explain in detail the elements of HBase architecture with neat sketch.	BTL 4	Analyzing
7	Describe the three levels of Data Locality in MongoDB.	BTL 5	Evaluating
8	(i) Write notes on Replication Strategies of Cassandra. (9) (ii) Explain how the snitches are used by Cassandra? (4)	BTL 4	Analyzing
9	Write notes on (i) Types of consistency (ii) Levels of Consistency (6+7)	BTL 6	Creating
10	Explain (i) Consistency in Mongoddb. (5) (ii) Consistency in HBase (8)	BTL 1	Remembering
11	Explain the levels for Write and Read operations and also the interaction between the consistency levels of both Write and read operations	BTL 4	Analyzing
12	Explain Hinted handoff and Read repair in Cassandra with diagram.	BTL 1	Remembering
13	Explain the operation of Vector clock in Cassandra with an example	BTL 3	Applying
14	Describe how the lost updates are solved with lightweight transactions in Cassandra	BTL 2	Understanding

PART C

1	(i) Explain in detail the shared-nothing database architecture with neat diagram (ii) Explain with neat diagram the shared-disk database architecture. (7)	BTL 6	Creating
2	Explain in detail the Consistent Hashing and how the addition of nodes to a cluster is handled in Cassandra.	BTL 5	Evaluating
3	Explain in detail the Multi Version Consistency model.	BTL 6	Creating
4	Explain how the lightweight transaction is implemented using the modified Cassandra Paxos protocol.	BTL6	Creating

UNIT V - DATA MODELS AND STORAGE

Data Models and Storage – SQL - NoSQL APIs - Return of SQL - Advance Databases - Postgre SQL - Riak - CouchDB - NEO4J - Redis - Future Databases - Revolution Revisited – Counter revolutionaries - Oracle HQ - Other Convergent Databases - Disruptive Database Technologies.

PART – A

Q.No	Questions	BT Level	Competence
1	Define memtable.	BTL 2	Understanding

2	Explain the concept of Bloom Filter.	BTL 1	Remembering
3	What is Key-Value pair?	BTL 2	Understanding
4	What is Column Family?	BTL 4	Analyzing
5	Explain Tombstone in Cassandra.	BTL 1	Remembering
6	What is SSTable?	BTL 1	Remembering
7	What is the difference between MongoDB and Redis database?	BTL 1	Remembering
8	Differentiate MongoDB and CouchDB?	BTL 5	Evaluating
9	Define Column family structure.	BTL 6	Creating
10	Explain collection data type in Cassandra.	BTL 4	Analyzing
11	What are the constructs for building JSON document?	BTL 6	Creating
12	List and explain the data types of RIAK.	BTL 1	Remembering
13	List the Convergent Replicated Data Types	BTL 2	Understanding
14	State the characteristics of BigTable tables.	BTL 1	Remembering
15	Compare MongoDB javascript query and SQL	BTL 4	Analyzing
16	Define Cascading.	BTL 3	Applying
17	State the uses of Couchbase UNNEST command.	BTL 3	Applying
18	What are the database-specific proposals of Quantum computing?	BTL 5	Evaluating
19	List the three major classifications of statements of SQL language.	BTL 2	Understanding
20	Explain g-counter and G-set CRDT types.	BTL 3	Applying

PART-B

1	Explain elaborately the Relational database storage architecture with diagram.	BTL 1	Remembering
2	Describe the Log-structured Merge trees and its architecture using Cassandra terminology	BTL 1	Remembering
3	State and explain read pattern for a log-structured merge tree using Cassandra terminology.	BTL 2	Understanding
4	Write notes on (i) Local Secondary indexing in distributed databases. (7) (ii) Implementation of Secondary indexing in NoSQL databases (6)	BTL 2	Understanding
5	Discuss the Cassandra Query Language.	BTL 4	Analyzing
6	Explain MapReduce technology considering WordCount program as example.	BTL 1	Remembering
7	(i) Differentiate Hive Query language and SQL (7) (ii) State the weaknesses for next generation databases. (6)	BTL 3	Applying
8	(i) Explain the decisions involved in choosing the correct database with diagram (5) (ii) Discuss the functions to query JSON document supported by Oracle. (8)	BTL 5	Evaluating

9	Write short notes on (i) Oracle Graph (7) (ii) Convergent Databases (6)	BTL 6	Creating
10	Explain in detail the Oracle support for JSON with example	BTL 4	Analyzing
11	(i) Explain the features that would be supported by an ideal database architecture (7) (ii) Explain the next generation database features offered by Oracle. (6)	BTL 4	Analyzing
12	Describe briefly the Disruptive Database Technologies	BTL 1	Remembering
13	Explain with neat diagram the B-tree index structure of RDBMS architectural plan.	BTL 2	Understanding
14	Explain (i) Cassandra Data model (7) (ii) JSON Data model (6)	BTL 3	Applying

PART C

1	Briefly explain the Convergent Replicated Data type in Cassandra with an example	BTL6	Creating
2	Explain briefly the mechanisms used by the following NoSQL APIs to perform insertion, modification and retrieval of data (5+5+5) (i) Riak (ii) HBase (iii) MongoDB	BTL6	Creating
3	Explain Oracle Sharding and the elements of Oracle Sharding architecture with diagram.	BTL5	Evaluating
4	Discuss how the JSON documents are accessed using the REST-based SODA protocol with example	BTL5	Evaluating