Roll No.												

PRESIDENCY UNIVERSITY

BENGALURU

End - Term Examinations – December 2025

Date: 15-12-2025 **Time:** 1.00pm to 04.00pm

School: SOCSE	**Program**: B. Tech CSE	
Course Code: CSE3513	**Course Name:** No SQL Data Management	
Semester: V	**Max Marks**: 100	**Weightage**: 50%

CO - Levels	CO1	CO2	CO3	CO4	CO5
Marks	24	24	4	24	24

Instructions:

(i) Read all questions carefully and answer accordingly.

(ii) Do not write anything on the question paper other than roll number.

Part A

Answer ALL the Questions. Each question carries 2marks. **10Q x 2M=20M**

1.	Explain the term "Impedance Mismatch" in the context of relational databases and object-oriented applications.	**2 Marks**	**L2**	**CO1**
2.	Differentiate between an Application Database and an Integration Database.	**2 Marks**	**L2**	**CO1**
3.	Compare Sharding and Master-Slave Replication as distribution models.	**2 Marks**	**L3**	**CO2**
4.	What is the significance of a Quorum in a distributed database system?	**2 Marks**	**L3**	**CO2**
5.	Illustrate the 'Map' phase of the Map-Reduce framework with a simple example.	**2 Marks**	**L3**	**CO3**
6.	Why are Key-Value stores not suitable for querying by data values or set operations?	**2 Marks**	**L3**	**CO3**
7.	Analyze a scenario where a Document-Oriented database would be a more suitable choice than a Relational database for a Content Management System.	**2 Marks**	**L4**	**CO4**
8.	What are the potential challenges of using a document database for applications requiring complex transactions across different documents?	**2 Marks**	**L4**	**CO4**
9.	Evaluate why a Graph Database is exceptionally efficient for running recommendation engine algorithms.	**2 Marks**	**L4**	**CO5**
10.	In the context of Graph Databases, why is the "Availability" characteristic often a challenge in cluster configurations?	**2 Marks**	**L4**	**CO5**

Part B

Answer the Questions. **Total Marks 80M**

11.	**a.**	Explain the consequences of using an Aggregate-Oriented data model. How does it affect data relationships and scalability?	**10 Marks**	**L2**	**CO1**
	b.	Differentiate between Key-Value and Document data models based on their structure and querying capabilities.	**5 Marks**	**L2**	**CO1**
	c.	Summarize the main reasons for the emergence of NoSQL databases.	**5 Marks**	**L2**	**CO1**
Or					
12.	**a.**	Compare the three main types of Aggregate-Oriented databases (Key-Value, Document, Column-Family), highlighting their primary use cases.	**10 Marks**	**L2**	**CO1**
	b.	Explain the concept of a "Schema-less" database and one advantage it offers during application development.	**5 Marks**	**L2**	**CO1**
	c.	Illustrate how "Materialized Views" can be used to model data for different access patterns.	**5 Marks**	**L2**	**CO1**

13.	**a.**	Compare the Peer-to-Peer and Master-Slave replication models in terms of fault tolerance, write scalability, and complexity.	**10 Marks**	**L3**	**CO2**
	b.	Explain how Version Stamps help in managing conflicts in an eventually consistent system.	**5 Marks**	**L3**	**CO2**
	c.	Differentiate between Read Consistency and Update Consistency.	**5 Marks**	**L3**	**CO2**
Or					
14.	**a.**	Analyze the CAP theorem. For a distributed banking system, which two characteristics would you prioritize and why?	**10 Marks**	**L3**	**CO2**
	b.	Explain the trade-off involved when "Relaxing Durability" in a distributed data store.	**5 Marks**	**L3**	**CO2**
	c.	Illustrate how sharding can be combined with replication to achieve both scalability and availability.	**5 Marks**	**L3**	**CO2**

15.	**a.**	Analyze the trade-offs between Consistency, Availability, and Query Features in Document-Oriented databases. How do these trade-offs influence database selection?	**10 Marks**	**L3**	**CO4**
	b.	Evaluate the suitability of a document database for a real-time web analytics application. What specific features make it a good fit?	**5 Marks**	**L3**	**CO4**
	c.	Propose a data model for an E-Commerce product catalog in a document database, justifying the structure of your aggregates.	**5 Marks**	**L3**	**CO4**
Or					

16.	**a.**	Critically evaluate a scenario where a Document-Oriented database would be an inappropriate choice. What specific system requirements would lead to this conclusion?	**10 Marks**	**L4**	**CO4**
	b.	Propose an alternative data Modeling strategy within a document database to handle "varying aggregate structures."	**5 Marks**	**L4**	**CO4**
	c.	How does the scaling strategy (e.g., sharding) of a document database impact the way you design your document IDs?	**5 Marks**	**L4**	**CO4**

17.	**a.**	Analyze why Graph Databases are inherently better at handling deeply interconnected data than relational or aggregate-oriented databases. Use the example of a social network.	**10 Marks**	**L4**	**CO5**
	b.	Propose a use case for a Graph Database in "Routing, Dispatch, and Location-Based Services" that would be inefficient in a SQL database.	**5 Marks**	**L4**	**CO5**
	c.	Evaluate the challenges of scaling a Graph Database in a peer-to-peer cluster.	**5 Marks**	**L4**	**CO5**
Or					
18.	**a.**	Critically assess the transaction and consistency model of a typical Graph Database. How does it support ACID properties compared to other NoSQL stores?	**10 Marks**	**L4**	**CO5**
	b.	Design a simple graph model for a recommendation engine that suggests "friends you may know."	**5 Marks**	**L4**	**CO5**
	c.	Justify why a Graph Database is not the best tool for performing large-scale aggregations across all entities (e.g., total sales sum).	**5 Marks**	**L4**	**CO5**