



VIT

Vellore Institute of Technology

Final Assessment Test – November 2025

Course: BCSE406L - NoSQL Databases

Class NBR(s): 1966/1970

Time: Three Hours

Slot: A1+TA1

Max. Marks: 100

- KEEPING MOBILE PHONE/ANY ELECTRONIC GADGETS, EVEN IN 'OFF' POSITION IS TREATED AS EXAM MALPRACTICE
- DON'T WRITE ANYTHING ON THE QUESTION PAPER

COs	CO Statements
CO1	Define, compare, and use the four types of NoSQL Databases.
CO2	Explain key-value databases and apply queries on those databases
CO3	Explain the detailed architecture, define objects, load data, query data, and performance tune Document-oriented NoSQL databases.
CO4	Demonstrate an understanding of the detailed architecture, define objects, load data, query data, and performance tune Column-oriented and Graph-oriented NoSQL databases
CO5	Evaluate NoSQL database development tools and programming languages.

BL – Blooms Taxonomy Level (1 – Remember, 2 – Understand, 3 – Apply, 4 – Analyse, 5 – Evaluate, 6 – Create)

Answer ALL Questions

(10 X 10 = 100 Marks)

1. A global e-commerce company experiences a surge in customers and transactions, leading to vast amounts of data from user activities, product catalogs, and purchase histories. The organization needs a data management solution that can efficiently manage fluctuating workloads, reduce infrastructure costs, accommodate diverse and evolving data formats, and ensure continuous service availability across regions. Explain how implementing a NoSQL database system can help the company overcome challenges related to scalability, cost-effectiveness, flexibility, and high availability in handling large-scale data operations.

CO1 BL3
2. Illustrate and describe six important design patterns to be considered when designing a key-value database. Explain how each pattern helps in improving data organization and retrieval efficiency in a key-value store.

CO2 BL2
3. ShopEase uses a key-value database (Redis) to manage user browsing activity, shopping cart details, session information, and personalized recommendations in real time. Apply your Redis query knowledge to the following tasks:

 - A developer needs to record the last 5 products a customer viewed using a list. Write a command to add product IDs P101, P102, and P103 to the end of the list named customer:9001:recent_products.

CO2 BL3

- b. Write a command to retrieve the most recently viewed product from the list customer:9001:recent_products.
- c. Write a command to store the hash of the customer profile: name: "Arjun Mehta", email: "arjun.mehta@email.com", and loyalty_points: "450".
- d. Assume the system already contains the following set:
SADD customer: 9001: purchased_products P101 P102 P103
From the above set, write a query to retrieve all the members.
- e. Write a query to remove the product P102 from the above set.

CO3 BL3

4. MediCare is a hospital management system that uses a NoSQL database (MongoDB) to manage patient records, doctor profiles, and appointments. Consider the patients, doctors, and appointments collections and write the MongoDB query for the tasks given.

Patients:

```
{
  "patientId": "P201",
  "name": "Rohan Verma",
  "email": "rohan.verma@email.com",
  "city": "Mumbai",
  "loyaltyPoints": 120
}
```

Doctors:

```
{
  "doctorId": "D301",
  "name": "Dr. Anjali Mehta",
  "specialization": "Cardiology",
  "fee": 1500,
  "availableSlots": 20
}
```

Appointments:

```
{
  "appointmentId": "A401",
  "patientId": "P201",
  "doctorIds": ["D301", "D305"],
  "appointmentDate": "2025-10-05",
  "totalFee": 3000,
  "status": "Completed"
}
```

a. Insert a new doctor into the doctors collection with the following details:

Doctor ID: D310

Name: Dr. Vikram Singh

Specialization: Neurology

Fee: 2000

AvailableSlots: 15

b. Retrieve all patients who live in Mumbai and have more than 100 loyalty points, displaying only their name and email.

c. Increase the loyaltyPoints of patient P201 by 50 after completing an appointment.

d. Determine the average fee of all doctors in the "Cardiology" specialization.

e. Calculate the total revenue generated from all appointments with status "Completed".

5.a) The e-commerce platform's NoSQL database is experiencing performance bottlenecks because the Products table contains thousands of records with numerous attributes. As a database designer, apply your knowledge of horizontal and vertical partitioning to restructure this database for improved performance. Explain how implementing both the horizontal and vertical partitioning strategies can enhance query efficiency, data management, and overall scalability of the system.

CO3 BL3

OR

5.b) A logistics company maintains shipment records, driver schedules, and delivery histories in a document-oriented database. Managers frequently query shipment details along with driver assignments, while executives generate reports such as "average delivery time per region."

CO3 BL3

- i. Analyse how denormalization versus joins affects the efficiency of retrieving complete shipment histories along with driver information.
- ii. Explain how replication and consistency levels impact the accuracy of shipment and tracking data during inter-warehouse synchronization.

6. Identify and explain five fundamental features of Google Bigtable that allow it to efficiently manage large-scale, data-intensive applications. Explain any two features applicability with real-time service scenarios.

CO4 BL2

7. A global e-commerce platform handles millions of real-time transactions, product updates, and user interactions across multiple countries. As the system architect, explain how Cassandra's architecture will efficiently support a high-performance retail system. Elaborate on how the different architectural components may contribute to high availability, fault tolerance, and performance of this platform. CO4 BL3
8. FleetTrack is building a real-time vehicle monitoring system to store and analyze vehicle telemetry data collected from GPS and IoT sensors. The system must efficiently store and query data such as speed, engine temperature, and fuel level for thousands of vehicles over time. Each record should include:
vehicle_id, vehicle_number, record_time, speed, engine_temp, fuel_level.
Task CO5 BL3
- Design a Cassandra table named vehicle_telemetry to efficiently support CQL queries.
 - Update the speed of vehicle V205 at record_time = '2025-10-07 09:00:00' — but only if the current speed is 60 km/h.
 - Create a secondary index on engine_temp to allow querying all vehicles with a specific engine temperature.
 - Write a query to fetch all records for vehicles whose fuel_level is below 20% and speed is above 80 km/h, ordered by record_time descending, limited to the last 50 records.
 - Write a query to delete all telemetry records of vehicle V205 before 2025-01-01.
9. AirTrack is developing a real-time analytics system for a major airline with flights operating across multiple cities and countries. The system collects and stores millions of flight operations records per day from check-in counters, in-flight sensors, and booking systems. As the database architect, you are tasked with designing a column family database to store and query this data. Apply any five of Cassandra's table design guidelines to design the table. CO4 BL3

- 10.a) A hospital administration wants to model the patient referral network among various departments and analyze the spread of hospital-acquired infections. Represent both the referral network and infection transmission using a sample graph structure. Clearly define the elements of a graph and specify at least three relevant properties for each.

CO4 BL3

OR

- 10.b) An e-commerce company is designing a Graph-Based Supply Chain Management System to efficiently manage warehouses, deliveries, customer orders, and inventory. The system collects data about warehouses, delivery routes, customers, and product categories. Explain how each of the following types of graphs can be used to represent different components of the system:

CO4 BL3

- i. Directed Graph
- ii. Undirected Graph
- iii. Flow Graph
- iv. Bipartite Graph

Q/K/TY