**BCS THE CHARTERED INSTITUTE FOR IT**

BCS HIGHER EDUCATION QUALIFICATIONS
BCS Level 6 Professional Graduate Diploma in IT

# ADVANCED DATABASE MANAGEMENT SYSTEMS

Tuesday 7th October 2025 – Afternoon

Answer **any** THREE questions out of FIVE. All questions carry equal marks.

Time: THREE hours

**Answer any Section A questions you attempt in Answer Book A**
**Answer any Section B questions you attempt in Answer Book B**

The marks given in brackets are **indicative** of the weight given to each part of the question.

| Calculators are **NOT** allowed in this examination. |
|---|

**This page is left intentionally blank.**

**Section A**
**Answer Section A questions in Answer Book A**

**A1.**

a) While both are storing data, there are significant differences between data warehouses and online transactional processing (OLTP) databases. Identify and describe the differences considering:

i. The type of data and characteristics of data,
ii. The organisation of data,
iii. The types of queries and transactions,
iv. The purpose.

**(12 marks)**

b) Consider the data cube below. Its dimensions are time (in quarters), products and region and the facts are the total sales value (the data values are fictional).

| Time | Products | Region | Sales |
|---|---|---|---|
| 2023 Q1 | Smartphones | Region 1 | 500,000 |
| 2023 Q1 | Smartphones | Region 2 | 700,000 |
| 2023 Q1 | Laptops | Region 1 | 200,000 |
| 2023 Q2 | Smartphones | Region 1 | 20,000 |
| 2023 Q2 | Laptops | Region 2 | 500,000 |

Execute the following operations on the data cube and show the result.

i. Slice the cube for time = 2023 Q1.

**(3 marks)**

ii. Rollup the cube for year (a year is made up of 4 quarters from Q1 to Q4).

**(4 marks)**

c) Slowly Changing Dimensions (SCD) in a data warehouse are a concept that enables dealing with the history of the data. Define SCD Types 1, 2 and 3 and give a reasoned explanation for when each type is used.

**(6 marks)**

**[Turn Over]**

**A2.**

a) Consider the following relational database schema from a zoo database and answer the questions below:

```
Animals (aid, name, species, hid)
Habitats (hid, name, environment)
FeedingSchedule(sid, aid, food, feedingtime)
```

Primary keys are underlined, `FeedingSchedule.aid` is a foreign key referencing the animal table and `Animals.hid` is a foreign key referencing the habitats table.

Assume you have the following SQL query that retrieves the names of animals and the names of their habitats for animals that are fed a specific type of food, such as 'seeds':

```
SELECT
   a.name AS animal_name,
   h.name AS habitat_name
FROM
   Animals a
JOIN
   FeedingSchedule fs ON a.aid = fs.aid
JOIN
   Habitats h ON a.hid = h.hid
WHERE
   fs.food = 'seeds';
```

i. Draw the initial query tree using relational algebra notation for the given SQL query.

**(5 marks)**

ii. Consider whether the tree is optimal or whether there are further improvements that could be made to enhance performance. If improvements can be made, describe **one** of them.

**(3 marks)**

b) Consider the following relational database schema with **two** tables and **four** given assumptions and answer the questions below:

Schema:

```
Room (rid, capacity)
Student (sid, age, sex, dept, gpa, rid)
```

Assumptions:

- 50% of students are female.
- There are very few single occupancy rooms (i.e. those with capacity of 1).
- Primary keys are underlined, `Student.rid` is a foreign key referencing Room.
- Queries are much more frequent than updates.

One of the most common queries to run on the database is to find grade (gpa) / age relations in data with a query like this (values for gpa and age can obviously change):

'list student ids with `gpa` = 3 and `age` < 19'

Which of the following three indexing approaches is best to improve performance of the above query and why? You need to justify why your choice is best, but also why the others are not as suitable.

Index on `age`
Index on `gpa`
composite index on two columns (`age, gpa`)

**(8 marks)**

c) The database is still running very slowly even after you recommended the right indexes in part (b) and using tuple-level (row-level) locks. A workload study reveals that in practice, the following two transactions are extremely frequent:

- update the gpa of students (may be different each time this is run)
- find the average age of all students group by `rid`

Trace analysis has found that the poor performance is due to concurrency bottlenecks.

Consider the following three options that could be used to improve performance:

- Make use of page granularity (block level) locks
- Vertically partition (i.e. decompose) the `Student` table into two separate tables such that (`sid, age, aid`) is in one table, and (`sid, sex, dept, gpa`) is in another.
- Use two-phase locking (2PL)

Decide which one is most likely to help (select only **one**). Justify your answer by explaining why **each** option is suitable (or not).

**(9 marks)**
**[Turn Over]**

**A3.**

a) Explain how the use of Role-Based Access Control (RBAC) affects the way in which a security strategy (the control of user access to data) can be implemented?

**(5 marks)**

b) Consider the following scenario for a university database system:

- Roles:
  - Admin: Can perform all operations on all tables.
  - Professor: Can read and update all student grades.
  - Student: Can view all grades but not modify them.

- Tables:
  - `Students(StudentID, Name, ProgrammeOfStudy)`
  - `Grades(StudentID, CourseID, Grade)`
  - `Courses(CourseID, CourseTitle, tid)`
  - `Professors(tid, name)`

i. Design an RBAC policy to enforce the above access rules. Specify which roles have access to which operations (SELECT, INSERT, UPDATE, DELETE) on each table.

**(3 marks)**

ii. Write SQL statements to implement the RBAC policy using a database system that supports roles. You can assume that the three roles specified above have already been created.

**(5 marks)**

c) The college has identified that its professors have accidentally changed data for courses they do not teach. Additionally, the permissions for students were too wide-ranging, allowing them to see other students' marks. Views (which had not been used) have been identified as the technical approach to further restrict access to data. You can assume that a system variable CURRENT_USER exists which provides a student ID for students and a staff ID for staff.

Answer the following (assume your Relational Database Management System (RDBMS) is setup as per the prior tasks, so this might need to remove any previously granted rights):

i. Professors should only be able to update the grades on courses they teach; they do need to be able to continue to see other grades for student counselling purposes. Write the required CREATE, GRANT, and REVOKE statements to create a suitable view that can then be used to provide appropriate access.

**(7 marks)**

ii. Students should be able to see only their own grades. Write the required CREATE, GRANT, and REVOKE statements to create a suitable view that can then be used to provide appropriate access. Write a SQL view that filters grades to a specific student number and provide the SQL GRANT statement for students accordingly.

**(5 marks)**

## Section B
## Answer Section B questions in Answer Book B

**B4.**

a) Compare NoSQL databases to Relational Database Management Systems (RDBMS). Your answer should explain what a NoSQL database is and consider the key differences and similarities between RDBMS and NoSQL. Aspects to consider could include the relationships between items, transactions, queries and storage. Your answer should conclude with a clear statement on situations in which each is best suited.

**(14 marks)**

b) ACID and BASE both aim to ensure well-behaved database systems. Considering only the aspect of consistency, define strong versus eventual consistency and explore the advantages and disadvantages in the use of each.

**(6 marks)**

c) A bank transfer involves debiting £500 from Account A in a database (db1) and crediting £500 to Account B in a remote database (db2). Explain how 2-Phase Commit (2PC) enforces that both operations complete successfully or neither operation occurs, thereby ensuring atomicity.

**(5 marks)**

**[Turn Over]**

**B5.**

a) You are a Database Administrator (DBA) and your staff have complained about the database being unresponsive. Further investigation shows that the database is not performing well, and the causes are hardware limitations. Discuss how CPU, RAM and Disks can create bottlenecks, how these manifest themselves and how you would resolve them.

**(9 marks)**

b) Describe the concept of data masking used in the context of database development and why it is desirable.

**(4 marks)**

c) A global e-commerce company initially uses a centralised database to store all its data, including customer information, orders, and inventory. As the company expands, users from different regions (North America, Europe, and Asia) experience significant delays when accessing the database. The inventory is held only in Europe and Asia as all orders are dispatched from warehouses in those regions.

Answer the questions below in the context of this scenario:

i. Explain why fragmentation of the centralised database might be necessary in this scenario.

**(2 marks)**

ii. Select a fragmentation strategy (horizontal, vertical, or hybrid) to address the company's needs. Justify your choice.

**(2 marks)**

iii. For your chosen strategy, describe how the following data might be fragmented. Your answer should state which fragments you will create and include the SQL statements for creating the fragments:

- `Customer Information (CustomerID, Name, Region)`
- `Inventory (ProductID, Stock, Region)`

**(8 marks)**

**END OF EXAMINATION**