BCS THE CHARTERED INSTITUTE FOR IT

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

ADVANCED DATABASE MANAGEMENT SYSTEMS

Wednesday 2nd October 2024 – 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 INTENTIONALLY BLANK

B5.

a) Explain the concept of read and write locks and how they ensure correctness of concurrent transactions.

(6 marks)

b) Consider a banking database with two tables:

```
Accounts(AccountID, Balance, AccountHolder)

Transactions(TransactionID, AccountID, TransactionType,
Amount)
```

Note that `TransactionType` can have the value '`Deposit`' or '`Withdrawal`'.

Design a SQL transaction that transfers £500 from Account 101 to Account 102. Provide the SQL code for the transaction.

(6 marks)

c) Explain why it is essential to ensure that the transfer is an atomic operation.

(3 marks)

d) Consider a database for a library management system with the following two tables:

Books table:

BookID	Title	Author	AvailableCopies
1	"Introduction to SQL"	John Doe	5
2	"Database Design"	Jane Smith	3
3	"Data Modeling 101"	Michael Johnson	7

BorrowedBooks table:

BorrowID	BookID	BorrowerName	BorrowDate
101	1	Alice Brown	2023-01-15
102	2	Bob Miller	2023-02-02
103	3	Carol White	2023-02-10

You are tasked with creating a trigger that will automatically update the `AvailableCopies` column in the Books table whenever a book is borrowed or returned. The trigger should decrement the `AvailableCopies` when a book is borrowed and increment it when a book is returned.

(10 marks)

END OF EXAMINATION

Section B
Answer Section B questions in Answer Book B

B4.

a) Compare a Graph Database to an RDMS. Your answer should explain what a Graph Database is and consider key differences and similarities between RDMS and Graph Databases. Aspects to consider could include relationships between items, transactions, queries and storage.

(12 marks)

b) Distributed Document Databases often use BASE instead of ACID as model for ensuring consistency. Explain what BASE is and why it is more suitable for Distributed Document Databases.

(8 marks)

c) Briefly describe the Two-Phase Commit Protocol used in distributed RDBMs. You might wish to use a diagram to support your answer.

(5 marks)

Section A
Answer Section A questions in Answer Book A

A1.

a) Data warehouses can be organised in different ways. Describe the concept of a star schema and contrast to a snowflake schema. Consider the aspects of disk use, normalisation, hierarchy and efficiency of queries in your answer.

(12 marks)

b) Consider a data cube string frequency of grades achieved by students (dimensions along axes are course, grades, academic year).

Explain the following operations and exemplify with an example based on the given example cube.

i. Dicing.
ii. Rollup.

(8 marks)

c) Data lakes are populated by a process called extract and load (EL). Briefly discuss the challenge of deleted data for incremental extraction (E) and pushing changes (L).

Remember that incremental extraction looks for changed data in the sources and only extracts that, whereas pushed changes update the existing data in the data lake.

(5 marks)

[Turn Over]

A2.

Consider the following relational database schema for an e-commerce system and answer the questions below:

```
Customer (customer_id, name, email, address)
Order (order_id, customer_id, order_date, total_amount)
OrderItem (order_id, product_id, quantity, unit_price)
Product (product_id, name, category, supplier_id)
```

Assume you have the following SQL query that retrieves information about customers who made orders in a specific category:

```
SELECT C.name, O.order_id, O.order_date
FROM Customer C
JOIN Order O ON C.customer_id = O.customer_id
JOIN OrderItem I ON O.order_id = I.order_id
JOIN Product P ON I.product_id = P.product_id
WHERE P.category = 'Electronics'
```

a) Draw the initial query tree for the given SQL query.

(6 marks)

b) Identify and explain **two** possible optimisation techniques for improving the performance of the query.

(6 marks)

c) Illustrate the modified query tree after applying one of the optimisation techniques you mentioned in question A2. b) above.

(8 marks)

d) Database caching can improve performance. Briefly explain the **two** approaches of query result caching and object caching.

(5 marks)

A3.

a) You are the database administrator for a financial institution that manages a database containing sensitive financial data, including customer account information and transaction details. The organisation is concerned about the confidentiality and integrity of the data and has tasked you with implementing logical database security measures.

Discuss and provide recommendations for enhancing the logical security of the database. Consider various aspects such as user authentication, authorisation, encryption, and auditing.

(10 marks)

b) Consider a large e-commerce database that stores customer information, order details, and product inventory. The organisation is concerned about security and compliance and has decided to implement a comprehensive auditing strategy.

Explain the importance and rationale behind auditing the e-commerce database. Discuss the potential risks and benefits of implementing a robust auditing strategy. Consider aspects such as security, compliance with regulations, and detection of unauthorised activities.

(7 marks)

c) As a database administrator for an online retail company that values data security, you are tasked with implementing views to ensure that sensitive information is appropriately controlled and accessed by different departments. The company maintains a database containing information about `products`, `customers`, and `orders`.

Task:

Write SQL statements to create secure views based on the following requirements for the Sales and Inventory departments, with a focus on data security.

i. Create a view named `SalesView` that displays product information along with the total sales quantity and revenue for each product. However, ensure that this view only includes products that have been sold and hides any sensitive customer information.
ii. Create a view named `InventoryView` that shows the current inventory status of each product, including the product name, current stock quantity, and reorder status. However, implement data masking to hide specific details, such as exact stock quantities, and only include products where the stock quantity is below the reorder threshold.

(8 marks)

[Turn Over]