

BCS THE CHARTERED INSTITUTE FOR IT

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

ADVANCED DATABASE MANAGEMENT SYSTEMS

Monday 10th May 2021 - 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.

SECTION A
Answer Section A questions in Answer Book A

A1.

a) The terms “*Volume, Velocity and Variety*”, are known as the 3-Vs that characterise Big Data applications. Explain why these terms are associated with Big Data.

(10 marks)

b) Describe a specific Big Data solution used by a company or organisation and explain how the use of Big Data adds value for that company/organisation.

(6 marks)

c) Discuss the changes to relational database technology that accommodate a Big Data solution. Use the following criteria to guide your answer:

- i) Structure and organisation of data;
- ii) Distributed processing and concurrency control;
- iii) Software techniques used by NoSQL, Hadoop and similar technologies.

(9 marks)

A2.

a) Refer to the sample table (*emp*) Fig A2.1 and the trigger code (Fig A2.2) below.

Describe the overall function of the trigger; an example of how it is invoked; and what output would be produced when the trigger is invoked due to SMITH's salary being updated to 1000.

(8 marks)

Fig A2.1 EMP (Employee) table

EMPNO	ENAME	JOB	MGR	SAL
7782	CLARK	MANAGER	7839	2450
7839	KING	PRESIDENT		5000
7369	SMITH	CLERK	7902	800
7876	ADAMS	CLERK	7788	1100
7566	JONES	MANAGER	7839	2975
7788	SCOTT	ANALYST	7566	2900
7902	FORD	ANALYST	7566	2900
7900	JAMES	CLERK	7698	950
7844	TURNER	SALESMAN	7698	1500
7698	BLAKE	MANAGER	7839	2850

B5.

a) Briefly describe in general terms the following approaches for achieving database distribution.

- i) Full database replication;
- ii) Fragmentation.

(6 marks)

b) Describe in detail how the **TWO** approaches from part a) (fragmentation and full database replication) can be applied to the following scenario to achieve a distributed database solution for the current central database (CDB). Your answer should describe the setup and table fragmentation as well as explaining the advantages and disadvantages of each solution.

Scenario

The company 'cheap-stuff.com' has been operating as an online store for several years supported by a central database. The product portfolio is such that customers often make orders to look and try items and then return a large fraction for refunds, meaning a large overhead on managing shipping. In the light of this a management decision was made to open three physical shops which will deal with customers to provide better service and to not interact with customers anymore through the online route. Customers can use any of the shops in relation to any order.

Your answer should describe how the central database could be distributed. This will require a brief explanation text and a description of how the various tables in the database are altered. To allow a basis for your discussion, the **FOUR** crucial tables from the centralised database are shown below. The names of the columns accurately describe the data held in those columns.

```
Customer(CustID, Name, LastName, Address)
Product(ProductID, name, description)
OrderItems(OrderNo, ProductID, quantity)
CustomerOrders(CustomerID, OrderNo, Branch, OrderDate)
```

(14 marks)

c) Describe the concepts of horizontal and vertical fragmentation and evaluate their use in distributed databases.

(5 marks)

b) It has been agreed that the salaries of Analysts are to be increased to £3000. The table emp (Fig A2.1) is updated accordingly.

Fig A2.2 SQL Trigger Code

```
CREATE OR REPLACE TRIGGER output_salary_changes
  BEFORE INSERT OR UPDATE ON emp
  FOR EACH ROW
  DECLARE
    sal_diff NUMBER;
  BEGIN
    sal_diff := :NEW.sal - :OLD.sal;
    output('Old salary: ' || :OLD.sal);
    output('New salary: ' || :NEW.sal);
    output('Difference: ' || sal_diff);
  END;
```

Write an SQL statement that will display, for **EACH** manager, a count of the number of employees whose salary is now greater than their own salary after the above update has taken place.

The following output should be produced when your statement is executed against the updated emp table

```
MANAGER COUNT_OF_EMPS
-----
JONES                2
```

(7 marks)

c) A new business rule is now required such that the salary of an analyst cannot exceed the salary of their manager.

Discuss the advantages **AND** disadvantages of the following approaches of implementing this business rule:

- i) Within the DBMS, using for example a Trigger;
- ii) Outside the DBMS within an application program using an Object Relational Mapping (ORM) framework.

(10 marks)

End of Examination

A3.

a) Give a definition of a “database transaction”.
(2 marks)

b) Describe, using a sequence of sample SQL statements, how the following key operations control the processing of a database transaction:

- i) SAVEPOINT;
- ii) COMMIT;
- iii) ROLLBACK.

(6 marks)

c) Show and explain how the basic functions and operations performed by a typical concurrency control protocol used by DBMSs can prevent the problems known as:

- i) Inconsistent reads;
- ii) Lost updates.

(8 marks)

d) Describe using the following sequence of transactions (T1 and T2) how the Two-Phase Locking protocol works.

T1	T2
Read(A)	
	Read(A)
Read(B)	
Write(B)	
Commit	
	Read(B)
	Write(B)
	Commit

(5 marks)

e) Transactions that apply the Two-Phase Locking (2PL) protocol can be affected by a situation known as deadlock. Using a simple example, explain how deadlock can occur and briefly explain how a deadlock situation can be resolved.

(4 marks)

SECTION B
Answer Section B questions in Answer Book B

B4.

a) Consider the following tables:

```
Product (ProductNo, ItemName, Price)
Shop (shopID, location)
Stocked (ProductNo, shopID, amount)
```

And the following query:

```
SELECT Product.price
       ,Shop.Location
FROM Product JOIN Stocked
ON Stocked.productno = product.productno
JOIN Shop
ON Shop.shopID = Stocked.ShopID
WHERE Product.ItemName = 'Navy Suit'
AND Shop.location = 'London';
```

Draw the query tree that corresponds to the most efficient way of processing this query.
(10 marks)

b) Relational database design typically strives for normalisation of tables to 3rd Normal Form. However, there are reasons for denormalization. Explain under which circumstances you might wish to denormalise a table and why.
(6 marks)

c) The CRUD paradigm has been developed to provide a standard way of bridging between objects and database operations. A database table called Books has been created together with the following object:

```
Book: {id: int; Title: string; author: string; ISBN: string}
```

Java and other programming languages use JDBC or ODBC type libraries to connect to databases. Write the required SQL query strings for e.g. Java or php database connections (or at least the related SQL query) that implement each of the 4 CRUD operations.

(9 marks)

[Turn Over]