Databases

Sample

Examination Paper

Marking Scheme

Answer ALL questions.

Clearly cross out surplus answers.

Time: 3 hours

The maximum mark for this paper is 100.

Any reference material brought into the examination room must be handed to the invigilator before the start of the examination.

Answer ALL questions

Marks

Question 1

a) Define metadata in the context of a Database Management System. Then provide ONE (1) example of metadata and outline THREE (3) kinds of information it might contain. **4**

Answers:

Metadata is data that contains the structure of other data within the database (1 mark)/data about data (1 mark). It describes the structure, organisation, management, and storage of data (1 mark). Award up to a maximum of 1 mark.

An example of metadata is the schema of a database (1 mark) which is stored in a data dictionary (1 mark). It includes the structure of tables, columns (1 mark), data types (1 mark), and constraints (1 mark). Award up to a maximum of 3 marks.

Any suitable examples of a data dictionary should be accepted as an alternative valid answer.

b) What roles do the Data Definition Language (DDL) and Data Manipulation Language (DML) play in a DBMS? Provide TWO (2) examples of operations for each. **6**

Answers:
DDL is used for defining and modifying the database structure or schema (1 mark). This includes creating, altering, and dropping tables and other database objects (1 mark).

Examples of DDL operations include "CREATE TABLE" (1 mark), "ALTER TABLE" (1 mark), and "DROP TABLE" (1 mark).

DML is concerned with the manipulation of the data itself (1 mark). This includes inserting, updating, deleting, and retrieving data from the database (1 mark).

Examples of DML operations include "INSERT" (1 mark), "UPDATE" (1 mark), "DELETE" (1 mark), and "SELECT" (1 mark).

1 mark for correctly identifying each role, up to a maximum of 2 marks.
1 mark for each correctly identified operation, up to a maximum of 4 marks.

Total 10 Marks

Marks

Question 2

(a) ACID is an acronym that stands for four key properties critical to database transactions. Identify these four properties and provide an explanation of each property. **8**

Answers:

Atomicity (1 mark): This is the property that defines a transaction as an indivisible unit in the sense that either the whole transaction must occur or no part of it must occur (1 mark). Although a transaction could be made up of a number of operations for the transaction to be atomic, ALL operations must be carried out or none of them at all (1 mark).

Consistency (1 mark): A transaction must not leave the database in an inconsistent state. Therefore, all constraints should be followed including enterprise constraints that enforce business rules. (1 mark)

Isolation (1 mark): Transactions should not interfere with other transactions. (1 mark)

Durability (1 mark): When a transaction has taken place, then its effects must be lasting and not vulnerable to being lost, because of a subsequent system failure. (1 mark)

Award 1 mark for each property and 1 mark for its explanation, up to a maximum of 8 marks. Any suitable answers should be accepted.

(b) Define and distinguish the terms ***data*** and ***information***. **2**

Answer:

Data is raw data and unprocessed.

Information is processed data and meaningful.

1 mark for describing each correctly, up to a maximum of 2 marks.

Total 10 Marks

Marks

Question 3

Scenario: **10**

A local hobby club organises various workshops throughout the year, covering different activities such as painting, photography, and cooking. Each workshop can be attended by many members of the club, and each member can sign up for multiple workshops based on their interests.

Draw an Entity Relationship Diagram to represent the scenario described above. You need to identify entities, relationships, cardinality, label all the primary keys and foreign keys. Clearly state which entity the keys belong to.

Answers:

1 mark for each entity, up to a maximum of 3 marks.

A maximum of 3 marks for the correct relationships, deduct 0.5 mark for each incorrect numerical cardinality. The notation types should not be the focus of this marking.

1 mark for each correctly labelled PK and FK, up to a maximum of 4 marks.

Total 10 Marks

Marks

Question 4

(a) The following ERD represents a veterinary clinic system where pet owners bring their pets to the clinic. Each pet can have multiple appointments and treatments. Each vet can handle multiple appointments but can only work on one treatment at a time.

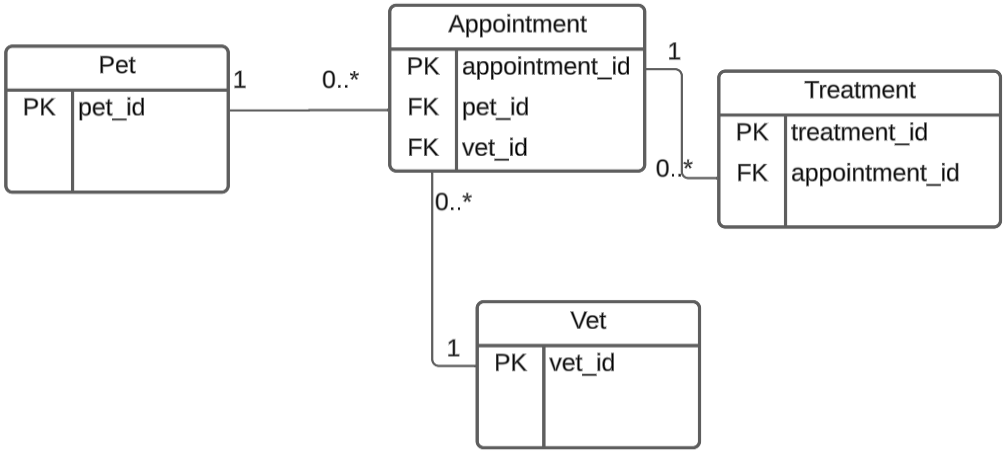

Create a CRUD matrix to show the following transactions:

- Transaction 1 – Register a new Pet
- Transaction 2 – Schedule a new Appointment for a registered pet
- Transaction 3 – Delete a vet and all their treatments.
- Transaction 4 – Update a pet's medical history after a completed treatment

4

Answers:

Award 1 mark for each correct row in a table, up to a maximum of 4 marks. 0.5 mark should be given for a partially correct row.

Transaction	***Pet***	***Appointment***	***Treatment***	***Vet***
T1	***C***			
T2	***R***	***C***		
T3			***D***	***D***
T4	***U***		***U***	

Marks

(b) A number of business rules have been defined for the veterinary clinic database system. Discuss how each might be enforced on the system.

Rule 1: All pets must have a recorded check-in date for their appointments. Pets that have not visited the clinic for over two years should be flagged for a wellness check reminder.

Rule 2: Pets should be categorised based on whether they are regular patients or referred by another clinic (categorised as 'Regular' or 'Referred').

Rule 3: Pets should be allowed to have follow-up appointments for treatments that were not successful, with a record of the initial treatment and the follow-up treatent linked for reference.

6

Answers:

2 marks for each accurately description of the database mechanisms for enforcing the business rules, up to a maximum of 6 marks.

Rule 1: Attributes for check-in date and last visit date should be added to the 'Pet' table. A transaction to flag the pets for a wellness check if the last visit exceeds two years.

Rule 2: A new attribute of PetStatus would need to be added to the 'Pet' table. A check constraint on the PetStatus ensures that a pet is classified as either 'regular' or 'referred'. Or this could be supported by the creation of a separate domain.

Rule 3: The primary key for the 'Treatment' table needs to include an additional attribute to allow multiple treatments for the same issue such as 'TreatmentAttempt', which can increment for each subsequent treatment, ensuring each treatment record remains unique.

Total 10 Marks

Marks

Question 5

(a) A university maintains a spreadsheet to track student enrollments in courses, recording each student's ID and name, along with the name, code, and duration of each course. The current structure of the data does not conform to the First Normal Form. **8**

Convert the provided data into 1NF.

student_id (PK)	student_name	enrollment
35365	John Smith	Computer Science, CS101, 3
27810	Jane Doe	Biology, BIO204, 4
63790	Bob Johnson	Computer Science, CS101, 3

Answers:

student_id(PK)	***first_name***	***last_name***	***course_code(FK)***
35365	***John***	***Smith***	***CS101***
27810	***Jane***	***Doe***	***BIO204***
63790	***Bob***	***Johnson***	***CS101***

course_code(PK)	***course_name***	***duration***
CS101	***Computer Science***	***3***
BIO204	***Biology***	***4***

1 mark for dividing the student_name column into two separate columns: first_name and last_name.

1 mark for each correctly identified table, up to a maximum of 2 marks.

1 mark for including course_code in the student table, and 1 mark if it's identified as a FK.

1 mark for the presence of course_name and duration in the course table.

1 mark for including course_code in the course table, and 1 mark if it's identified as a PK.

Marks

(b) Discuss the importance of relational integrity. **2**

Answers:

It refers to the accuracy and consistency of data within a relational database (1 mark). It ensures that relationships among data in different tables remain valid and that the data is reliable (1 mark).

Any other suitable answers should be accepted.

Total 10 Marks

Question 6

Based on the scenario provided in Question 5. Write the appropriate SQL commands to perform each of the following tasks.

(a) Create the `student` table. Assume the course table has been created. **6**

Answers:

CREATE TABLE student(

student_id, INT PRIMARY KEY,

first_name VARCHAR(20),

last_name VARCHAR(20),

course_code VARCHAR(10),

FOREIGN KEY (course_code) REFERENCES course);

1 mark for each correctly defined row in the SQL command, up to a maximum of 6 marks.

Marks

(b) The `student` table needs to be updated to include a new column named `email` to store student email addresses. **2**

Answers:

ALTER TABLE student

ADD email VARCHAR(255);

1 mark for each correctly defined row in the SQL command, up to 2 marks.

(c) The `duration` column in the `course` table is not needed and should be removed. **2**

Answers:

ALTER TABLE course

DROP COLUMN duration;

1 mark for each correctly defined row in the SQL command, up to 2 marks.

Total 10 Marks

Marks

Question 7

Scenario: Zoo Database System

This zoo database system is designed to manage information about the animals, their habitats, and feeding schedules. Below are the tables with sample data.

animal

animal_id	animal_name	species	habitat_id
1	Leo	Lion	1
2	Stripes	Tiger	2
3	Polly	Parrot	3
4	Slithers	Snake	4

habitat

habitat_id	habitat_name	environment
1	Savannah	Grassland
2	Tiger Trail	Forest
3	Bird Paradise	Tropical
4	Reptile House	Temperate

feeding_schedule

schedule_id	animal_id	food	feeding_time
1	1	Meat	14:00
2	2	Chicken	12:00
3	3	Seeds	08:00
4	3	Fruits	15:00
5	4	Mice	20:00

Marks

(a) Write a SQL command to retrieve a list of all animals along with their habitat names. Include both the animal's name and the habitat name in your results. **3**

Answer:

```
SELECT a.animal_name, h.habitat_name
FROM animal a
JOIN habitat h ON a.habitat_id = h.habitat_id;
```

1 mark for each correctly defined row in the SQL command, up to a maximum of 3 marks.

(b) Write a SQL command to find the total number of feedings scheduled for each type of food. The result should list the food type and the total number of feedings. **3**

Answer:

```
SELECT fs.food, COUNT(fs.food) AS total_feedings
FROM feeding_schedule fs
GROUP BY fs.food;
```

1 mark for each correctly defined row in the SQL command.

(c) Write a SQL command to list the names of the animals that need to be fed more than once per day. Include the names of the animals and the number of times it is fed in your results. **4**

Answer:

SELECT a. animal_name***, COUNT(fs.schedule_id) AS feedings_per_day (1 mark)***

FROM animal a

JOIN feeding_schedule fs ON a.animal_id = fs.animal_id (1 mark)

GROUP BY a. animal_name ***(1 mark)***

HAVING COUNT(fs.schedule_id) > 1; (1 mark)

Total 10 Marks

Marks

Question 8

You are provided with a paper form that every new patient is required to fill out upon registering at a clinic. This form collects various information necessary for the clinic's patient management system. **10**

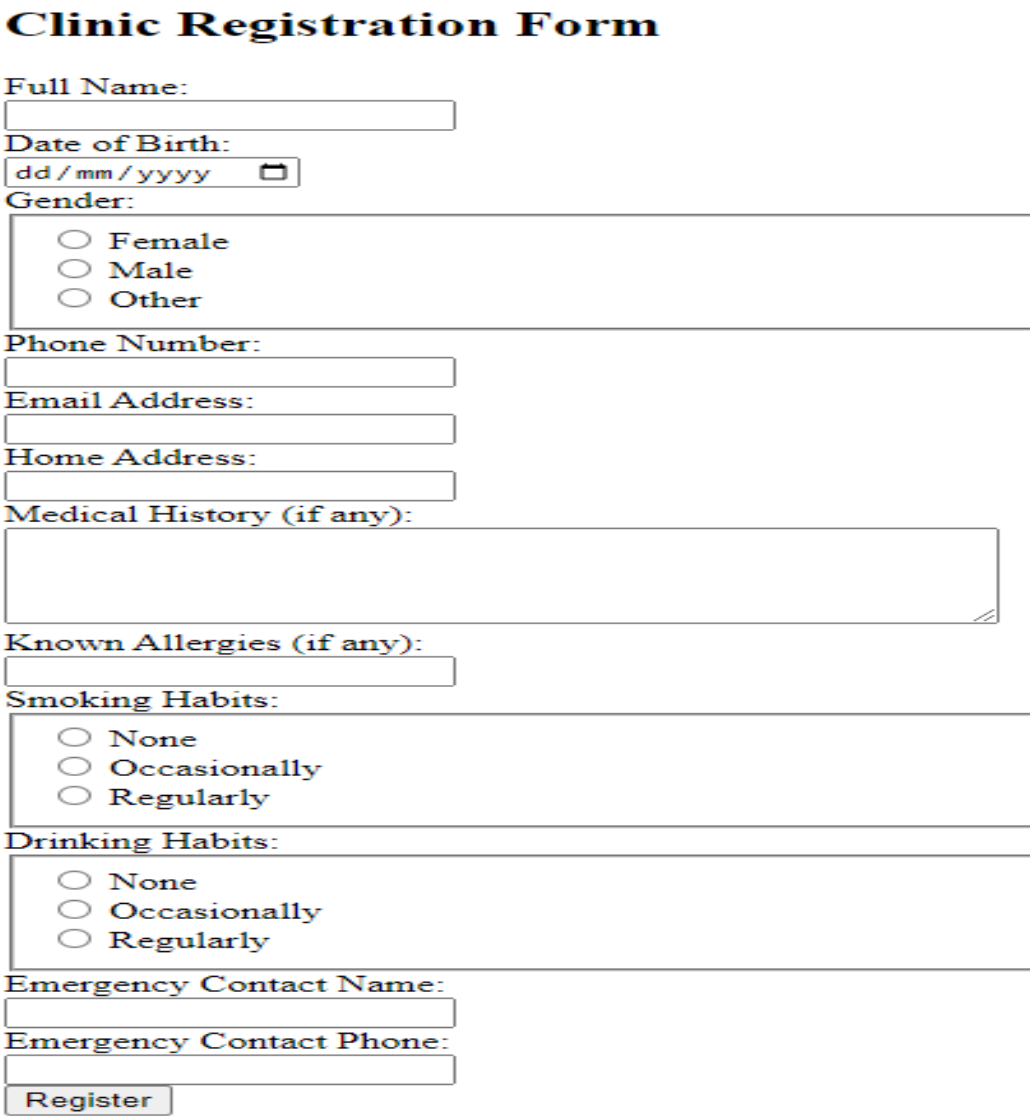

Identify TEN (10) data fields that are collected from the new patient.

Answers:

Examples include first name, last name, data of birth, gender, country code/area code/phone number, email address, home address, medical history, smoking habits, drinking habits, first name and last name of emergency contact, country code/area code/phone number of emergency contact.

1 mark for each correctly identified data field.

Total 10 Marks

Marks

Question 9

(a) It is crucial to protect sensitive information from unauthorised access and potential breaches. Identify THREE (3) strategies that are essential for safeguarding a database. Provide a description of how it contributes to the overall security of the database system. **6**

Answers:

- ***Discretionary Access Control: Access is based on the identity of the user and on discretionary decisions made by the individual responsible for the data.***
- ***Mandatory Access Control: Access decisions are determined by comparing security labels that indicate how sensitive data is with security clearances that indicate system entities' trustworthiness.***
- ***Role-based Access Control: Access decisions are based on the roles that users have within the system and the rules stating what access is allowed for users in those roles.***
- ***Authentication: It prevents unauthorised access and potential data breaches by verifying user identities with strong authentication measures.***
- ***Data masking: Used to protect sensitive information from exposure to unauthorised personnel by obscuring it within a database, using techniques such as character shuffling, substitution with realistic but artificial data, encryption etc.***
- ***Backup and recovery – regular backups involve copying and archiving database data to enable recovery in the event of data loss.***

1 mark for each correctly identified strategy, up to a maximum of 3 marks.

1 mark for each correct description, up to a maximum of 3 marks.

Any other suitable answers should be accepted.

(b) Identify TWO (2) advantages of Cloud databases **and** provide an explanation of each advantage. **4**

Answers:

- ***Scalability: resources can be dynamically adjusted to meet demand.***
- ***Enhanced management: cloud databases often come with a suite of tools for monitoring, performance tuning, and security management.***
- ***Cost-effectiveness: savings from a pay-as-you-go pricing model and the elimination of capital expenditures.***

Marks

- ***Accessibility: Cloud databases can be accessed from anywhere, which facilitates remote work and distributed teams.***

- ***Security: Cloud service providers offer robust security measures that might be challenging to implement on-premises.***

1 mark for each correctly identify advantage, up to a maximum of 2 marks.

1 mark for each correct description, up to a maximum of 2 marks.

Total 10 Marks

Question 10

(a) Big data is often characterised by the 3Vs. Identify the 3Vs and provide a brief explanation of each. **6**

Answer:

- ***Volume: the vast amounts of data generated from various resources.***
- ***Velocity: the speed at which new data is generated and the pace at which data moves through organisations.***
- ***Variety: the different types of data available, ranging from structured, numeric data in traditional databases to unstructured text, video, audio, financial transactions, etc.***

1 mark for each correctly identified characteristic, up to a maximum of 3 marks.

1 mark for each correct description, up to a maximum of 3 marks.

Any other suitable answers should be accepted.

(b) List FOUR (4) types of NoSQL databases. **4**

Answers:

Document databases, key-value databases, wide column, graph databases.

1 mark for each correctly identified database, up to a maximum of 4 marks.

Total 10 Marks

End of Paper

Marks

Learning Outcomes matrix

Question	Learning Outcomes assessed	Marker can differentiate between varying levels of achievement
1	1	Yes
2	4	Yes
3	3	Yes
4	1, 4, 5	Yes
5	2, 4	Yes
6	5	Yes
7	5	Yes
8	1	Yes
9	1, 7	Yes
10	6	Yes

Grade descriptors

Learning Outcome	Fail	Referral	Pass	Merit	Distinction
Understand the concepts associated with database systems	Has little to no awareness of principles and concepts underlying theoretical frameworks and approaches and is unable to identify associated strengths or weaknesses.	Has vague awareness of principles and concepts underlying theoretical frameworks and approaches and is only superficially able to identify associated strengths or weaknesses.	Has satisfactory awareness of principles and concepts underlying theoretical frameworks and approaches and demonstrates sufficient ability to identify associated strengths and weaknesses.	Has very good awareness of principles and concepts underlying theoretical frameworks and approaches and is able in detail to identify associated strengths and weaknesses.	Has an excellent awareness of principles and concepts underlying theoretical frameworks and approaches and is comprehensively able to identify associated strengths and weaknesses.
Understand the concepts associated with the relational model	Has little to no awareness of principles and concepts underlying theoretical frameworks and approaches and is unable to identify associated strengths or weaknesses.	Has vague awareness of principles and concepts underlying theoretical frameworks and approaches and is only superficially able to identify associated strengths or weaknesses.	Has satisfactory awareness of principles and concepts underlying theoretical frameworks and approaches and demonstrates sufficient ability to identify associated strengths and weaknesses.	Has very good awareness of principles and concepts underlying theoretical frameworks and approaches and is able in detail to identify associated strengths and weaknesses.	Has an excellent awareness of principles and concepts underlying theoretical frameworks and approaches and is comprehensively able to identify associated strengths and weaknesses.

					Marks
Understand how to design and develop a database system	Demonstrates little to no ability to adequately use design principles to effectively develop an artefact to solve an identified issue.	Demonstrates an extremely limited ability to adequately use design principles to effectively develop an artefact to solve an identified issue.	Demonstrates a satisfactory ability to use design principles to effectively develop an artefact to solve an identified issue.	Demonstrates a very good ability to use design principles to effectively develop an artefact to solve an identified issue.	Demonstrates an excellent ability to use design principles to effectively develop an artefact to solve an identified issue.
Be able to develop and enhance a logical database design	Demonstrates little to no ability to adequately use design principles to effectively develop an artefact to solve an identified issue.	Demonstrates an extremely limited ability to adequately use design principles to effectively develop an artefact to solve an identified issue.	Demonstrates a satisfactory ability to use design principles to effectively develop an artefact to solve an identified issue.	Demonstrates a very good ability to use design principles to effectively develop an artefact to solve an identified issue.	Demonstrates an excellent ability to use design principles to effectively develop an artefact to solve an identified issue.
Be able to develop and enhance a database system using SQL	Demonstrates little to no ability to adequately develop an artefact to solve an identified issue.	Demonstrates an extremely limited ability to adequately develop an artefact to solve an identified issue.	Demonstrates a satisfactory ability to develop an artefact to solve an identified issue.	Demonstrates a very good ability to develop an artefact to solve an identified issue.	Demonstrates an excellent ability to develop an artefact to solve an identified issue.

Marks

Understand the principles behind Big Data and the NoSQL databases	Displays no understanding of Big Data and NoSQL databases. Unable to recognise the key features and benefits of NoSQL over traditional databases or how Big Data impacts database design and scalability.	Shows a basic awareness of Big Data and NoSQL databases. Recognises some key aspects but fails to understand the broader implications and technical considerations involved.	Demonstrates satisfactory comprehension of Big Data principles and the operational basics of NoSQL databases. Can distinguish NoSQL databases from traditional databases and understands their use cases.	Exhibits a strong grasp of the principles behind Big Data and NoSQL databases. Can articulate the advantages of NoSQL databases in handling large datasets and different data types.	Has an in-depth understanding of Big Data and NoSQL databases. Can critically evaluate and recommend NoSQL database solutions for complex Big Data problems, demonstrating a sophisticated understanding of scalability and performance optimisation.
Understand the concepts and principles of Cloud database	Lacks understanding of the fundamental concepts and principles of Cloud databases. Cannot identify the main advantages or the operational model of Cloud databases.	Has a superficial grasp of Cloud databases. Recognises some concepts but cannot articulate how Cloud databases differ from on-premises solutions or their benefits.	Has a solid understanding of Cloud database concepts and principles. Understands the benefits of scalability, availability, and managed services provided by Cloud databases.	Possesses a very good understanding of Cloud databases. Can discuss the importance of Cloud databases in modern architectures and understands their significance in providing flexible and scalable data storage solutions.	Demonstrates an excellent and comprehensive understanding of the key concepts and principles of Cloud databases. Can describe in detail how Cloud databases offer advantages in terms of accessibility, maintenance, and scalability, and can suggest appropriate use cases for Cloud databases.