



**Pakistan Institute
of Public Finance Accountants**

Model Solutions

**Database Management
System (Application)**

AGP | PG | Public Sector

May Exam-2026

MODEL SOLUTIONS – DISCLAIMER

INTRODUCTION

The Model Solutions are provided to students for clear understanding of relevant subject and it helps them to prepare for their examinations in organized way.

These Model Solutions are prepared only for the guidance of students that how they should attempt the questions. The solutions are not meant for assessment criteria in the same pattern mentioned in the Model Solution. The purpose of Model Solution is only to guide the students in their future studies for appearing in examination.

The students should use these Model Solutions as a study aid. These have been prepared by the professionals on the basis of the International Standards and laws applicable at the relevant time. These solutions will not be updated with changes in laws or Standards, subsequently. The laws, standards and syllabus of the relevant time would be applicable. PIPFA is not supposed to respond to individual queries from students or any other person regarding the Model Solutions.

DISCLAIMER

The Model Solutions have been developed by the professionals, based on standards, laws, rules, regulations, theories and practices as applicable on the date of that particular examination. No subsequent change will be applicable on the past papers solutions.

Further, PIPFA is not liable in any way for an answer being solved in some other way or otherwise of the Model Solution nor would it carry out any correspondence in this regard.

PIPFA does not take responsibility for any deviation of views, opinion or solution suggested by any other person or faculty or stake holders. PIPFA assumes no responsibility for the errors or omissions in the suggested answers. Errors or omissions, if noticed, should be brought to the notice of the Executive Director for information.

If you are not the intended recipient, you are hereby notified that any dissemination, copying, distributing, commenting or printing of these solutions is strictly prohibited.



Q.1. i. Domain

A set of possible values for an attribute.

ii. Entity

An object in the real world that is distinguishable from other objects such as the green dragon toy.

iii. Relationship

An association among two or more entities.

iv. Covering constraint

Within an ISA hierarchy, a covering constraint determines where the entities in the subclasses collectively include all entities in the superclass. For example, with an Employees entity set with subclasses Hourly Employee and Salary Employee, does every Employee entity necessarily have to be within either Hourly Employee or Salary Employee?

v. Weak entity set

An entity that cannot be identified uniquely without considering some primary key attributes of another identifying owner entity. An example is including Dependent information for employees for insurance purposes.

vi. Aggregation

A feature of the entity relationship model that allows a relationship set to participate in another relationship set. This is indicated on an ER diagram by drawing a dashed box around the aggregation.

vii. Role Indicator

If an entity set plays more than one role, role indicators describe the different purpose in the relationship. An example is a single Employee entity set with a relation Reports-To that relates supervisors and subordinates.

Q.2. In the given situation, the following Vs of big data are demonstrated:

a

- Volume:** Demonstrated by “massive amounts of data daily” and “millions of customer reviews”.
- Velocity:** Demonstrated by “thousands of transactions per hour” and the need for “immediate analysis”.
- Variety:** Demonstrated by the combination of “customer purchase records (structured)”, “customer reviews and social media comments (unstructured)”, and “real-time clickstream data (semi-structured)”.

Implications for RiftCart

Implication of Volume: The “massive amounts of data” generated daily by RiftCart put a significant “strain on their existing data infrastructure”. This means they need specialized infrastructure for storage and processing to handle the enormous and continuously growing amount of data.

Implication of Velocity: The need to process “thousands of transactions per hour” and rely on “immediate analysis” means RiftCart requires systems capable of real-time or near-real-time data processing to enable dynamic operations and timely marketing campaign adjustments.

Implication of Variety: The collection of data in “structured”, “unstructured” and “semi-structured” formats lead to challenges in data integration and standardization. This diverse nature can contribute to the “inconsistencies in customer profiles” and make comprehensive analysis difficult, impacting their reporting capabilities.



Q.2. Identification of specific challenges

b

Data Quality and Data Integration: The inconsistency in customer profiles stems from integrating diverse data sources with different structures and formats. Ensuring accuracy and completeness across this varied data landscape is a major challenge.

Storage and Processing Capacity: Delayed reporting is caused by high data volume exceeding the capacity of the current infrastructure to process data efficiently.

Suggestion of one relevant emerging trend or technology and justification

- **Challenge: Data Quality and Data Integration**
Trend / Technology: Artificial Intelligence (AI) and Machine Learning (ML)
Justification: AI-driven tools, particularly those leveraging ML, can “automate data quality checks, detect anomalies” and process diverse data formats. They help standardize and integrate structured, semi-structured and unstructured data. For RiftCart, this would directly assist in resolving “inconsistencies in customer profiles” by standardizing varied data types, thereby contributing to “reliable insights” and improving the accuracy of their reporting.
- **Challenge: Storage and Processing**
Trend / Technology: Edge Computing
Justification: Edge computing brings data processing closer to the source (e.g., at the customer device or local node), reducing load on central servers and minimizing latency. By processing high- volume, real-time data (like clickstreams) at the edge, RiftCart can improve response times and reporting speed, especially during peak seasons.

Q.3.

a

	Relational Database Management Systems (RDBMS)	Object-Oriented Database Management Systems (OODBMS)
Core Characteristics:	In the Relational Model, data is organized into tables, also known as relations, consisting of rows(tuples) and columns (attributes). Each row represents a unique record, and each column represents an attribute of that record. Relationships between different tables are established using primary keys (unique identifiers for each record) and foreign keys (fields in one table referencing a primary key in another). Structured Query Language (SQL) is the standard language for interacting with relational databases.	OODBMS stores data as objects, similar to object-oriented programming. Each object encapsulates both data (attributes) and methods (functions) that operate on that data. This model supports concepts such as class, object, and inheritance, where objects can inherit properties and methods from other objects.
Strengths & Use Cases:	RDBMS offers simplicity due to its table-based structure, which is easy to understand and work with. It provides data independence and flexible querying capabilities through SQL, making it easy to retrieve and manipulate data. RDBMS is optimally used for transactional systems and data analysis where structured data is prevalent, such as customer records and financial transactions.	OODBMS is well-suited for applications that require the management of complex data types, including multimedia, graphics and other real-world objects. The encapsulation of data and operations into a single entity simplifies the modelling of real-world systems.
Contrast (Disadvantages):	Large-scale relational databases may encounter performance issues with complex queries, requiring optimization techniques like indexing. They are also not as scalable as some newer models for handling massive datasets or unstructured data.	Object-oriented databases are generally more complex to design and manage compared to relational databases. They often lack a standard query language like SQL, which can make implementation and maintenance more challenging. The cost and effort involved in developing and maintaining OODBMS and the reluctance of organizations to transition from existing relational systems are also significant drawbacks.



Q.3. Primary Characteristics:

b

NoSQL databases are specifically designed to handle unstructured and semi-structured data. They are characterized by their ability to manage massive volumes of data generated at high velocity from diverse sources.

NoSQL databases are highly scalable, particularly with complex data. Examples include MongoDB and Cassandra.

Typical Applications:

Given their design for unstructured and semi-structured data, NoSQL databases are often leveraged in Big Data environments.

They are suitable for real-time analytics and machine learning applications. This makes them ideal for scenarios involving social media data, IoT device data and web transactions, where data variety and volume are significant.

- (a)** The Network Model organizes data as nodes (entities) connected by links (relationships), allowing a child node to have more than one parent node. This ‘graph-like structure’ provides greater flexibility than the hierarchical model, which is limited to one-to-many relationships, making the Network Model highly suitable for managing complex data relationships that involve many-to-many connections. Its design inherently supports these intricate interconnections, enabling more direct and efficient navigation of linked data.

Examples

- Consider a student’s table and a course’s table. A student can enroll in multiple courses, and each course can have multiple students enrolled. This forms a many-to-many relationship between “Students” and “Courses”. This scenario perfectly illustrates the kind of complex, interconnected data that the Network Model is designed to manage, where multiple instances from one entity type are associated with multiple instances from another.
- Consider a scenario involving “Products” and “Suppliers”. A single product can be supplied by multiple different suppliers, and conversely, a single supplier can provide many different products. This represents a classic many-to-many relationship. In a Network Model, “Product” and “Supplier” could be nodes, and a “Supplies” link could connect them, allowing any product to be linked to multiple suppliers and any supplier to be linked to multiple products, directly representing these complex relationships.

- Q.3.** The Network Model organizes data as nodes (entities) connected by links (relationships), allowing a child node to have more than one parent node. This ‘graph-like structure’ provides greater flexibility than the hierarchical model, which is limited to one-to-many relationships, making the Network Model highly suitable for managing complex data relationships that involve many-to-many connections. Its design inherently supports these intricate interconnections, enabling more direct and efficient navigation of linked data.

c

Examples

- Consider a student’s table and a course’s table. A student can enroll in multiple courses, and each course can have multiple students enrolled. This forms a many-to-many relationship between “Students” and “Courses”. This scenario perfectly illustrates the kind of complex, interconnected data that the Network Model is designed to manage, where multiple instances from one entity type are associated with multiple instances from another.
- Consider a scenario involving “Products” and “Suppliers”. A single product can be supplied by multiple different suppliers, and conversely, a single supplier can provide many different products. This represents a classic many-to-many relationship. In a Network Model, “Product” and “Supplier” could be nodes, and a “Supplies” link could connect them, allowing any product to be linked to multiple suppliers and any supplier to be linked to multiple products, directly representing these complex relationships.



- Q.4.** SQL statements are high-level instructions and each statement is responsible for a specific task. These statements can generally be classified into five categories:

a

Data Definition Language (DDL)

- * This family of SQL commands is used to define database schema.
- * Examples include CREATE, DROP, ALTER

Data Manipulation Language (DML)

- * This family of SQL commands is used to modify the data inside a table.
- * Examples include INSERT, UPDATE, DELETE

Data Query Language (DQL)

- * This family of SQL commands performs query on existing tables.
- * Examples include SELECT

Data Control Language (DCL)

- * This family of SQL commands deals with the rights and permissions.
- * Examples include GRANT, REVOKE

Transaction Control Language (TCL)

- * This family of SQL commands deals with transactions.
- * Examples include COMMIT, ROLLBACK, SAVEPOINT
- * You only need these commands if you have OLTP operation.

- Q.4.** They are closely related. DDL is responsible to define the structure of the table, basically what is allowed to enter the table and what is not allowed. DDL can be regarded as a set of rules that shape the table structure (schema). After DDL defines the schema, then it is DML job to fill the table with the data.

b

- Q.4.** Both of these functions return a single value, the difference is the input, scalar functions operate on a single value while aggregate functions operate on a set of values. I try to clarify the difference with an example. For example, string functions like ISNULL(), ISNUMERIC(), LEN() are scalar functions. They input a single value and return a single value. On the other hand, AVG(), MAX(), SUM() are aggregate functions, they input multiple values and output a single value.

c

Q.5. Possible Inadequacies and Explanations (Any three)

a

• **Data Silos and Fragmentation:**

Sales data is scattered across Excel files from various marketplaces and a separate SQL database. This lack of integration creates data silos, making it difficult to reconcile sales figures with marketing and maintain a consistent view of performance. It violates principles of data integration and data consistency, limiting enterprise-wide reporting.

• **Manual Processing and Error-Prone Data:**

The company relies on manual processes to compile data, which are slow, inconsistent, and prone to human error. This undermines data quality dimensions such as accuracy, timeliness and completeness, impairing reliable decision-making.

• **Weak Authentication and Access Controls:**



Solutions – Database Management System (Application)

Generic login IDs for customer support allow unauthorized or untraceable access to sensitive customer data. This leads to lack of accountability, high risk of costly data breaches, business disruptions and reputational damages.

- **Unmonitored Systems and Poor Security Practices:**

Due to lack of monitoring and poor security practices, a marketing database breach occurred. This points to insufficient IT controls, such as continuous monitoring, logging and threat detection, which are core to any information security framework.

- **Lack of Formal Data Governance Structure:**

With no clear data ownership or standards and limited IT staffing, 'NewTrends' lacks a structured data governance framework. This leads to inconsistent practices in data handling and poor data stewardship, undermining both compliance and strategic analytics.

Q.5. Possible Solutions (Any three corresponding to the inadequacies explained in part (a))
b

- **Implement a Centralized Data Warehouse (CDW) with Extract, Transform, Load (ETL) Processes:**

Implement a CDW, powered by ETL processes.

A single source of truth for data improves consistency across departments (marketing, sales, customer support), fostering better collaboration and reducing time spent reconciling conflicting figures.

- **Automated Data Validation and Cleansing Tools:**

Implement automated tools that validate and cleanse data as it enters the system.

It reduces manual error and ensures only high-quality, consistent data enters reporting systems. This improves trust in analytics, reduces rework and speeds up decision-making cycles.

- **Implement Role-Based Access Control (RBAC) with Unique Login IDs:**

Deploy RBAC to ensure employees only access data relevant to their role, eliminating generic logins. Also, deploy Privileged Access Management (PAM) for sensitive systems like the SQL database.

- **Enhance IT Security Measures and Monitoring:**

Deploy Security Information and Event Management (SIEM) systems for real-time monitoring of all databases and networks to detect suspicious activity and prevent unauthorized access.

It enables proactive identification and mitigation of threats through real-time monitoring and alerts. This reduces downtime, prevents breaches and protects company data assets, preserving continuity and customer confidence.

Robust ITGCs and security measures (access controls, monitoring, IRP) minimize the risk of costly data breaches and operational disruptions, ensuring business continuity and protecting reputation.

- **Establish a Data Governance Framework:**

Introduce formal policies, procedures, standards, and roles (e.g., data ownership and stewardship) for data collection, storage, access and use across the organization.

It establishes consistent policies and accountability for data handling across departments. It improves data quality, promotes efficient collaboration and ensures that strategic initiatives (like predictive analytics) are based on trustworthy data.