70-762^{Q&As}

70-762^{Q&As}

Developing SQL Databases

Pass Microsoft 70-762 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/70-762.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

🔧 **Instant Download** After Purchase

🔧 **100% Money Back** Guarantee

🔧 **365 Days** Free Update

🔧 **800,000+** Satisfied Customers

QUESTION 1

Note: this question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in the series.

Information and details provided in a question apply only to that question.

You are developing an application to track customer sales.

You need to create a database object that meets the following requirements:

-

 Return a value of 0 if data inserted successfully into the Customers table.

-

 Return a value of 1 if data is not inserted successfully into the Customers table.

-Support TRY...CATCH error handling

-Be written by using Transact-SQL statements.

What should you create?

A. extended procedure

B. CLR procedure

C. user-defined procedure

D. DML trigger

E. scalar-valued function

F. table-valued function

Correct Answer: D

DML triggers is a special type of stored procedure that automatically takes effect when a data manipulation language (DML) event takes place that affects the table or view defined in the trigger. DML events include INSERT, UPDATE, or DELETE statements. DML triggers can be usedto enforce business rules and data integrity, query other tables, and include complex Transact-SQL statements.

References:https://msdn.microsoft.com/en-us/library/ms178110.aspx

QUESTION 2

Note: This question is part of a series of questions that use the same answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are developing an application to track customer sales. You create tables to support the application. You need to

create a database object that meets the following data entry requirements: What should you create?

Table name	Data entry requirements
Customers	Inserts must be grouped in a batch of Transact-SQL statements.
CustomerTransactions	Inserts must be grouped in a batch of Transact-SQL statements.
Invoices	Inserts must be grouped in a batch of Transact-SQL statements. All of the data modifications made must become a permanent part of the database if data entry is successful. All data entry modifications must be erased if an error occurs.
Orders	Inserts must be grouped in a batch of Transact-SQL statements. All of the data modifications made must become a permanent part of the database if data entry is successful. All data entry modifications must be erased if an error occurs.

A. extended procedure

B. CLR procedure

C. user-defined procedure

D. DML trigger

E. DDL trigger

F. scalar-valued function

G. table-valued function

Correct Answer: C

References: https://msdn.microsoft.com/en-us/library/ms345075.aspx

QUESTION 3

Note: This question is part of a series of questions that use the same or similar answer choices. An Answer choice may be correct for more than one question in the series. Each question independent of the other questions in this series.

Information and details provided in a question apply only to that question.

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run

during off-peak hours.

You observe that many deadlocks appear to be happening during specific times of the day.

You need to monitor the SQL environment and capture the information about the processes that are causing the deadlocks. Captured information must be viewable as the queries are running.

What should you do?

A. A. Create a sys.dm_os_waiting_tasks query.

B. Create a sys.dm_exec_sessions query.

C. Create a PerformanceMonitor Data Collector Set.

D. Create a sys.dm_os_memory_objects query.

E. Create a sp_configure `max server memory\\' query.

F. Create a SQL Profiler trace.

G. Create a sys.dm_os_wait_stats query.

H. Create an Extended Event.

Correct Answer: F

To view deadlock information, the Database Engine provides monitoring tools in the form of two trace flags, and the deadlock graph event in SQL Server Profiler.

Trace Flag 1204 and Trace Flag 1222 When deadlocks occur, trace flag 1204 and trace flag 1222 return information that is captured in the SQL Server error log. Trace flag 1204 reports deadlock information formatted by each nodeinvolved in the deadlock. Trace flag 1222 formats deadlock information, first by processes and then by resources. It is possible to enable bothtrace flags to obtain two representations of the same deadlock event.

References:https://technet.microsoft.com/en-us/library/ms178104(v=sql.105).aspx

QUESTION 4

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

You have a database named DB1 that contains the following tables: Customer, CustomerToAccountBridge, and CustomerDetails. The three tables are part of the Sales schema. The database also contains a schema named Website. You create the Customer table by running the following Transact-SQL statement:

```
CREATE TABLE Customer
(
    CustomerNumber int NOT NULL,
    CustomerName varchar(50) NOT NULL,
    CreateDate date NOT NULL,
    Gender bit,
    Address varchar(50)
    City varchar(50)
    State char(2),
    CustomerStatus bit NOT NULL,
    MaritalStatus bit,
    Segment varchar(5),
    CountryCode char(2),
    Birthday date,
    PostalCode char(5),
    PhoneNumber varchar(20),
    Account1 char(7),
    Account1Status bit,
    Account2 char(7),
    Account2Status bit,
    CONSTRAINT PK_Customer PRIMARY KEY CLUSTERED (CustomerNumber)
);
```

The value of the CustomerStatus column is equal to one for active customers. The value of the Account1Status and Account2Status columns are equal to one for active accounts. The following table displays selected columns and rows from the Customer table.

Customer ID	CustomerName	Gender	Account1	Account1Status	Account2	Account2Status
101	Name A	0	0001001	0	0001002	1
102	Name B	1	0002001	1	0002002	0
103	Name C	0	0003001	1	0003002	1

You plan to create a view named Website.Customer and a view named Sales.FemaleCustomers. Website.Customer must meet the following requirements:

1.

 Allow users access to the CustomerName and CustomerNumber columns for active customers.

2.

 Allow changes to the columns that the view references. Modified data must be visible through the view.

3.

 Prevent the view from being published as part of Microsoft SQL Server replication. Sales.Female.Customers must meet the following requirements:

1.

Allow users access to the CustomerName, Address, City, State and PostalCode columns.

2.

Prevent changes to the columns that the view references.

3.

Only allow updates through the views that adhere to the view filter.

You have the following stored procedures: spDeleteCustAcctRelationship and spUpdateCustomerSummary. The spUpdateCustomerSummary stored procedure was created by running the following Transacr-SQL statement:

```
CREATE PROCEDURE uspUpdateCustomerSummary
@CustomerId INT
AS
BEGIN
    SET NOCOUNT on;
    UPDATE CustomerDetails SET TotalDepositAccountCOunt - TotalDepositAccountCount + 1 WHERE CustomerID = @CustomerID;
    BEGIN TRAN;
        BEGIN TRY
            UPDATE CustomerDetails SET TotalAccountCount - TotalAccountCount + 1 WHERE CustomerID - @CustomerID;
        END TRY
        BEGIN CATCH
            IF @@TRANCOUNT > 0
                ROLLBACK TRAN;
        END CATCH
        IF @@TRANCOUNT > 0
            COMMIT TRAN;
```

You run the spUpdateCustomerSummary stored procedure to make changes to customer account summaries. Other stored procedures call the spDeleteCustAcctRelationship to delete records from the CustomerToAccountBridge table.

You must update the design of the Customer table to meet the following requirements.

1.

You must be able to store up to 50 accounts for each customer.

2.

Users must be able to retrieve customer information by supplying an account number.

3.

Users must be able to retrieve an account number by supplying customer information.

You need to implement the design changes while minimizing data redundancy.

What should you do?

A. Split the table into three separate tables. Include the AccountNumber and CustomerID columns in the first table. Include the CustomerName and Gender columns in the second table. Include the AccountStatus column in the third table.

B. Split the table into two separate tables. Include AccountNumber, CustomerID, CustomerName and Gender columns in the first table. Include the AccountNumber and AccountStatus columns in the second table.

C. Split the table into two separate tables, Include the CustomerID and AccountNumber columns in the first table. Include the AccountNumber, AccountStatus, CustomerName and Gender columns in the second table.

D. Split the table into two separate tables, Include the CustomerID, CustomerName and Gender columns in the first table. Include AccountNumber, AccountStatus and CustomerID columns in the second table.

Correct Answer: D

Two tables is enough.CustomerID must be in both tables.

QUESTION 5

You are developing a database reporting solution for a table that contains 900 million rows and is 103 GB.

The table is updated thousands of times a day, but data is not deleted.

The SELECT statements vary in the number of columns used and the amount of rows retrieved.

You need to reduce the amount of time it takes to retrieve data from the table. The must prevent data duplication.

Which indexing strategy should you use?

A. a nonclustered index for each column in the table

B. a clustered columnstore index for the table

C. a hash index for the table

D. a clustered index for the table and nonclustered indexes for nonkey columns

Correct Answer: B

Columnstore indexes are the standard for storing and querying large data warehousing fact tables. It uses column-based data storage and query processing to achieve up to 10x query performance gains in your data warehouse over

traditional row-oriented storage.

A clustered columnstore index is the physical storage for the entire table.

Generally, you should define the clustered index key with as few columns as possible.

A nonclustered index contains the index key values and row locators that point to the storage location of the table data. You can create multiple nonclustered indexes on a table or indexed view. Generally, nonclustered indexes should be

designed to improve the performance of frequently used queries that are not covered by the clustered index.

References: https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview?view=sql-server-2017