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

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:

Launch when table data is modified.

Evaluate the state a table before and after a data modification and take action based on the difference.

Prevent malicious or incorrect table data operations.

Prevent changes that violate referential integrity by cancelling the attempted data modification.

Run managed code packaged in an assembly that is created in the Microsoft.NET Framework and located into Microsoft SQL Server.

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: B

You can create a database object inside SQL Server that is programmed in an assembly created in the Microsoft .NET Framework common language runtime (CLR). Database objects that can leverage the rich programmingmodel provided

by the CLR include DML triggers, DDL triggers, stored procedures, functions, aggregate functions, and types.

Creating a CLR trigger (DML or DDL) in SQL Server involves the following steps:

Define the trigger as a class in a .NETFramework-supported language. For more information about how to program triggers in the CLR, see CLR Triggers. Then, compile the class to build an assembly in the .NET Framework using the

appropriate language compiler.

Register the assembly in SQL Server using the CREATE ASSEMBLY statement. For more information about assemblies in SQL Server, see Assemblies (Database Engine).

Create the trigger that references the registered assembly.

References:https://msdn.microsoft.com/en-us/library/ms179562.aspx

QUESTION 2

You are developing an application that connects to a database. The application runs the following jobs:

Job	Transact-SQL statement	Description
JobA	Exec uspDeletePrevRecords	The stored procedure deletes all records from a table named tblBalanceTransactions that were created before the current month by using a single DELETE statement. Approximately 10 million records are deleted each time you run this stored procedure
JobB	Exec uspUpdateCurRecords	This stored procedure updates records in the tblBalanceTransaction table that were created in the current month. Only a few hundred records are updated each time you run this stored procedure.

The READ_COMMITTED_SNAPSHOT database option is set to OFF, and auto-content is set to ON. Within the stored procedures, no explicit transactions are defined.

If JobB starts before JobA, it can finish in seconds. If JobA starts first, JobB takes a long time to complete.

You need to use Microsoft SQL Server Profiler to determine whether the blocking that you observe in JobB is caused by locks acquired by JobA.

Which trace event class in the Locks event category should you use?

A. LockAcquired

B. LockCancel

C. LockDeadlock

D. LockEscalation

Correct Answer: A

The Lock:Acquiredevent class indicates that acquisition of a lock on a resource, such asa data page, has been achieved.

The Lock:Acquired and Lock:Released event classes can be used to monitor when objects are being locked, the type of locks taken, and for how long the locks were retained. Locks retained for long periods of time may cause contention

issues and should be investigated.

QUESTION 3

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 that contains the following tables: BlogCategory, BlogEntry, ProductReview, Product, and SalesPerson. The tables were created using the following Transact SQL statements:

```
CREATE TABLE BlogCategory
(
      CategoryID int NOT NULL PRIMARY KEY,
      CategoryName nvarchar (20)
);

CREATE TABLE BlogEntry
(
      Entry int NOT PRIMARY KEY,
      Entrytitle nvarchar (50),
      Category int NOT NULL FOREIGN KEY REFERENCES BlogCategory
(CategoryID)
);

CREATE TABLE dbo.ProductReview
(
      ProductReviewID IDENTITY(1,1) PRIMARY KEY,
      Product int NOT NULL,
      Review varchar (1000) NOT NULL
);

CREATE TABLE dbo.Product
(
      ProductID int Identity(1,1) PRIMARY KEY,
      Name varchar(1000) NOT NULL
);
CREATE TABLE dbo.SalesPerson
(
      SalesPersonID int IDENTITY(1,1) PRIMARY KEY,
      Name varchar (1000) NOT NULL,
      SalesID Money
)
```


You must modify the ProductReview Table to meet the following requirements:

1.

The table must reference the ProductID column in the Product table

2.

Existing records in the ProductReview table must not be validated with the Product table.

3.

Deleting records in the Product table must not be allowed if records are referenced by the ProductReview table.

4.

Changes to records in the Product table must propagate to the ProductReview table.

You also have the following databse tables: Order, ProductTypes, and SalesHistory, The transact-SQL statements for these tables are not available.

You must modify the Orders table to meet the following requirements:

1.

Create new rows in the table without granting INSERT permissions to the table.

2.

Notify the sales person who places an order whether or not the order was completed.

You must add the following constraints to the SalesHistory table:

-

a constraint on the SaleID column that allows the field to be used as a record identifier

-

a constant that uses the ProductID column to reference the Product column of the ProductTypes table

-

a constraint on the CategoryID column that allows one row with a null value in the column

-

a constraint that limits the SalePrice column to values greater than four

Finance department users must be able to retrieve data from the SalesHistory table for sales persons where the value of the SalesYTD column is above a certain threshold.

You plan to create a memory-optimized table named SalesOrder. The table must meet the following requirments:

-

The table must hold 10 million unique sales orders.

-

The table must use checkpoints to minimize I/O operations and must not use transaction logging.

-Data loss is acceptable.

Performance for queries against the SalesOrder table that use Where clauses with exact equality operations must be optimized.

You need to modify the design of the Orders table.

What should you create?

A. a stored procedure with the RETURN statement

B. a FOR UPDATE trigger

C. an AFTER UPDATE trigger

D. a user defined function

Correct Answer: D

Requirements: You must modify the Orders table to meet the following requirements:

1.

Create new rows in the table without granting INSERT permissions to the table.

2.

Notify the sales person who places an order whether or not the order was completed.

References:https://msdn.microsoft.com/en-us/library/ms186755.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.

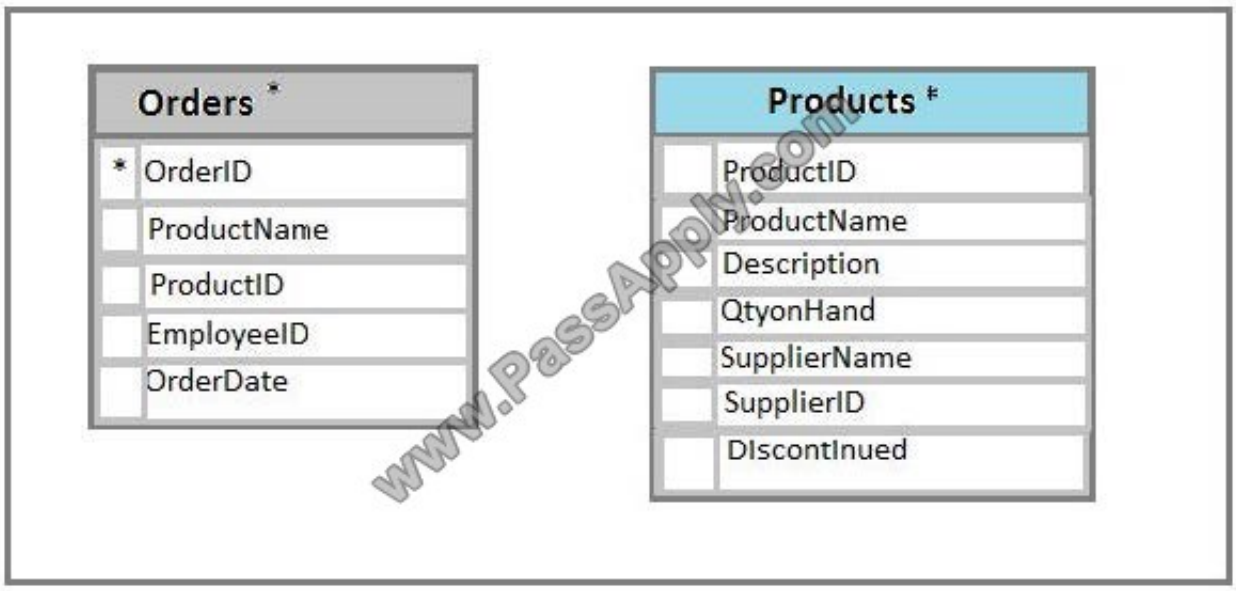

You have a database named Sales that contains the following database tables: Customer, Order, and Products. The Products table and the Order table are shown in the following diagram.

The customer table includes a column that stores the data for the last order that the customer placed.

You plan to create a table named Leads. The Leads table is expected to contain approximately 20,000 records. Storage requirements for the Leads table must be minimized.

You need to modify the database design to meet the following requirements:

In the table below, identify the constraint that must be configured for each table.

NOTE: Make only one selection in each column.

Hot Area:

Correct Answer:

A FOREIGN KEY in one table points to a PRIMARY KEY in another table. Here the foreign key constraint is put on the ProductIDin the Orders, and points to the ProductID of the Products table.

With a check constraint on the ProductID we can ensure that the Products table contains only unique rows.

References:http://www.w3schools.com/sql/sql_foreignkey.asp

QUESTION 5

You have a database that contains three encrypted store procedures named dbo.Proc1, dbo.Proc2 and dbo.Proc3. The stored procedures include INSERT, UPDATE, DELETE and BACKUP DATABASE statements. You have the following requirements:

-You must run all the stored procedures within the same transaction.

-You must automatically start a transaction when stored procedures include DML statements.

-You must not automatically start a transaction when stored procedures include DDL statements.

You need to run all three stored procedures.

Which four Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments to the answer area and arrange then in the correct order.

Select and Place:

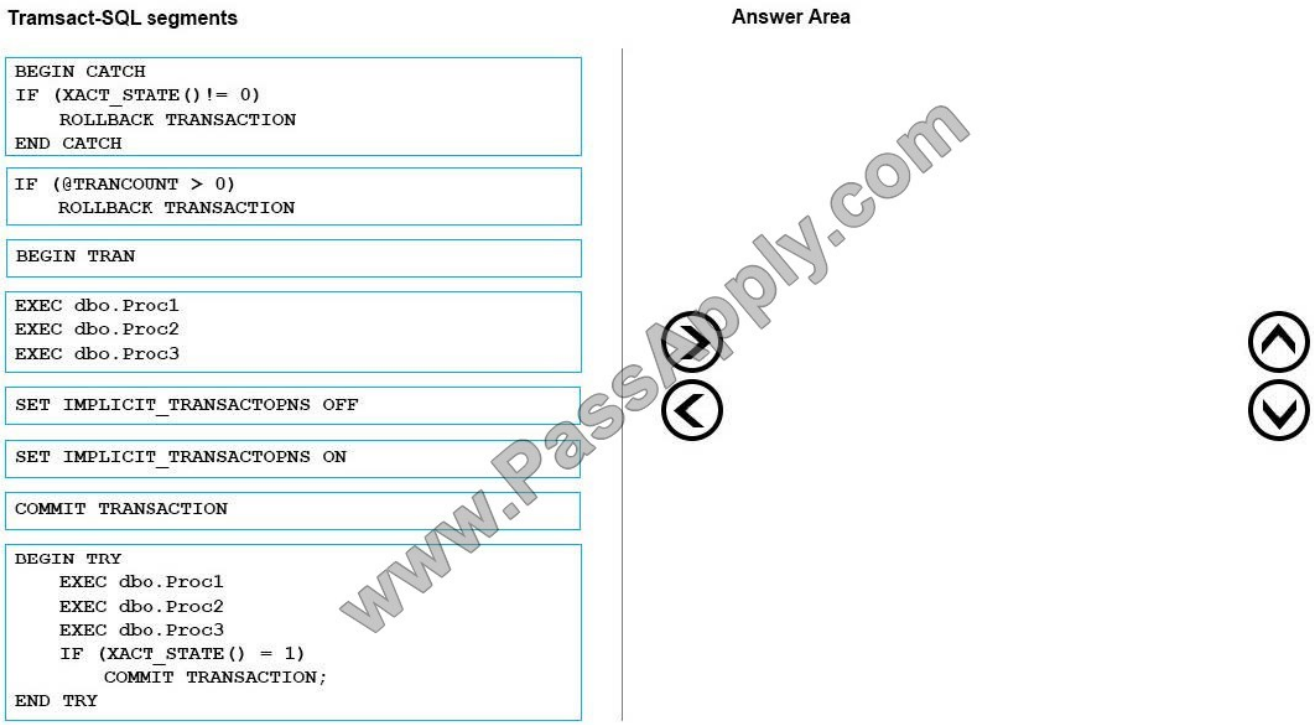

Correct Answer:

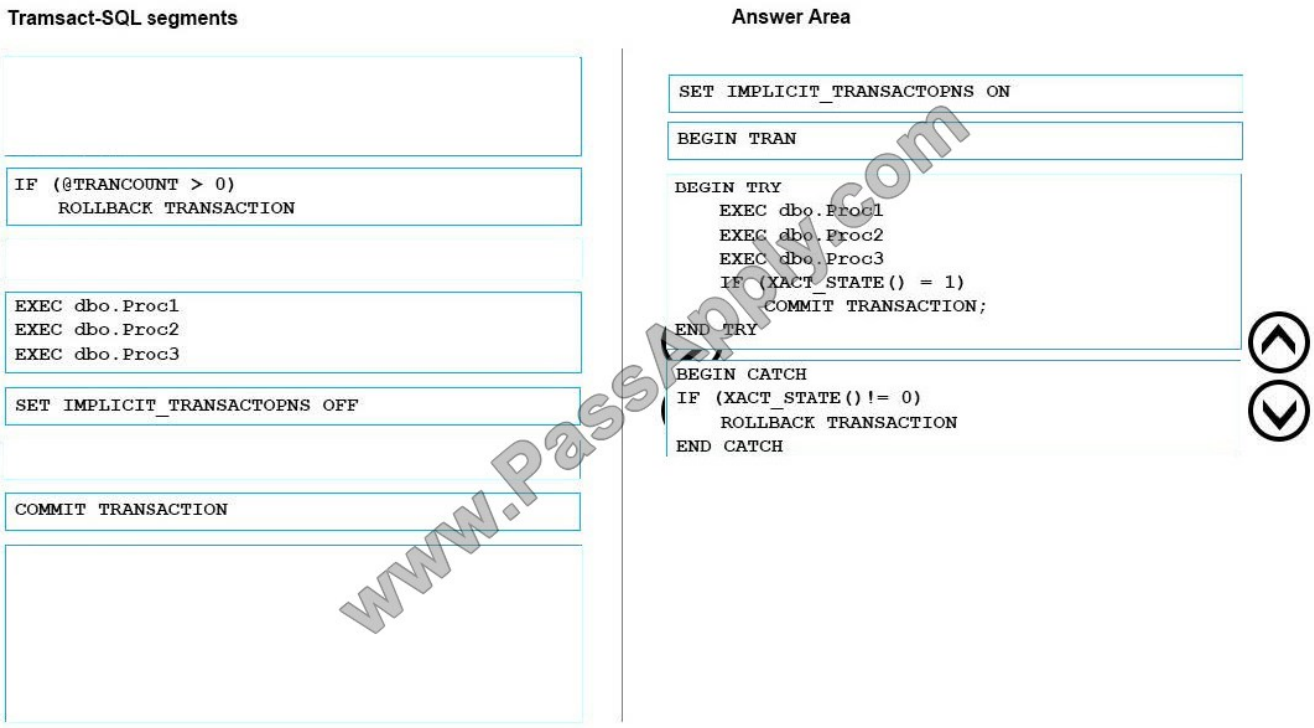

Note:

Implicit transaction mode remains in effect until the connection executes a SET IMPLICIT_TRANSACTIONS OFF statement,which returns the connection to autocommit mode. In autocommit mode, allindividual statements are committed if

they complete successfully.

ALTER TABLE	FETCH	REVOKE
BEGIN TRANSACTION	GRANT	SELECT (See exception below.)
CREATE	INSERT	TRUNCATE TABLE
DELETE	OPEN	UPDATE
DROP	.	.

When a connection is in implicit transaction mode and the connection is not currently in a transaction, executing any of the following statements starts a transaction:

Note 2: XACT_STATE returns the following values.

1 The currentrequest has an active user transaction. The request can perform any actions, including writing data and committing the transaction. The transaction is committable.

-1 The current request has an active user transaction, but an error has occurred that has caused the transaction to be classified as an uncommittable transaction.the transaction is uncommittable and should be rolled back.

0 There is no active user transaction for the current request. Acommit or rollback operation would generate an error.

References:

https://technet.microsoft.com/en-us/library/ms187807(v=sql.105).aspx https://technet.microsoft.com/en-us/library/ms189797(v=sql.110).aspx

QUESTION 6

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 and application to track customer sales.

You need to return the sum of orders that have been finalized, given a specified order identifier. This value will be used in other Transact-SQL statements.

You need to create a database object.

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: F

User-defined scalar functions return a single data value of the type defined in the RETURNS clause.

References:https://technet.microsoft.com/en-us/library/ms177499(v=sql.105).aspx

QUESTION 7

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 need to create Website Customer.

How should you complete the view definition? To answer, drag the appropriate Transact-SQL segments to the correct locations, Each Transact-SQL segment may be used once, more than once or not at all. You may need to drag the split bar

between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

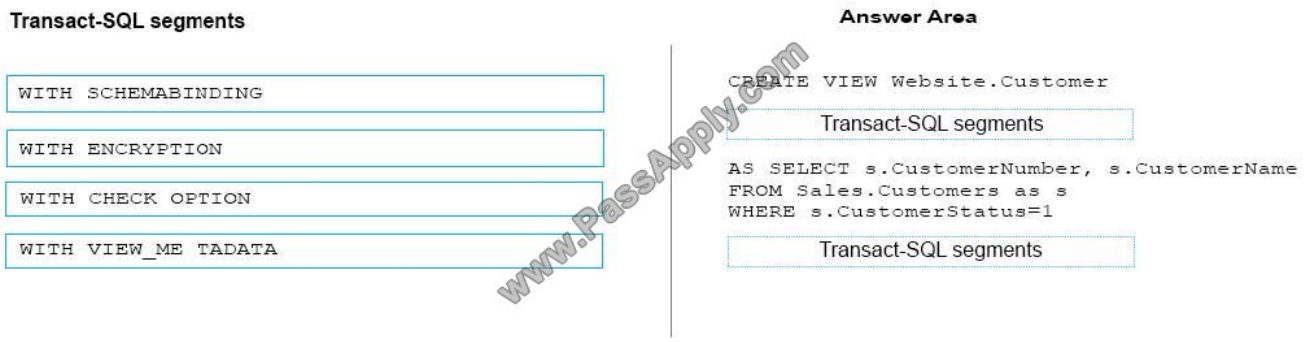

Correct Answer:

Box 1: WITH ENCRYPTION

Using WITH ENCRYPTION prevents the view from being published as part of SQL Server replication.

Box 2: WITH CHECK OPTION

CHECK OPTION forces all data modification statements executed against theview to follow the criteria set within select_statement. When a row is modified through a view, the WITH CHECK OPTION makes sure the data remains visible

through the view after the modification is committed.

Note: Website.Customer must meet the followingrequirements:

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 aspart of Microsoft SQL Server replication.

Incorrect Asnwers:

SCHEMABINDING binds the view to the schema of the underlying table or tables. When SCHEMABINDING is specified, the base table or tables cannot be modified in a way that would affect the view definition.

VIEW_METADATA specifies that theinstance of SQL Server will return to the DB-Library,ODBC, and OLE DB APIs the metadata information about the view, instead of the base table or tables, when browse-mode metadata is being requested

for a query that references the view.

References:https://msdn.microsoft.com/en-us/library/ms187956.aspx

QUESTION 8

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 and application to track customer sales.

You need to create an object that meet the following requirements:

- Run managed code packaged in an assembly that was created in the Microsoft.NET Framework and uploaded in Microsoft SQL Server.

-

Run written a transaction and roll back if a future occurs.

-

Run when a table is created or modified.

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: B

The common language runtime (CLR) is the heart of the Microsoft .NET Framework andprovides the execution environment for all .NET Framework code. Code that runs within the CLR is referred to as managed code.

With the CLR hosted in Microsoft SQL Server (called CLR integration), you can author stored procedures, triggers, user-defined functions, user-defined types, and user-defined aggregates in managed code. Because managed code compiles

to native code prior to execution, you can achieve significant performance increases in some scenarios.

QUESTION 9

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 Sales that contains the following database tables: Customer, Order, and Products. The Products table and the Order table are shown in the following diagram.

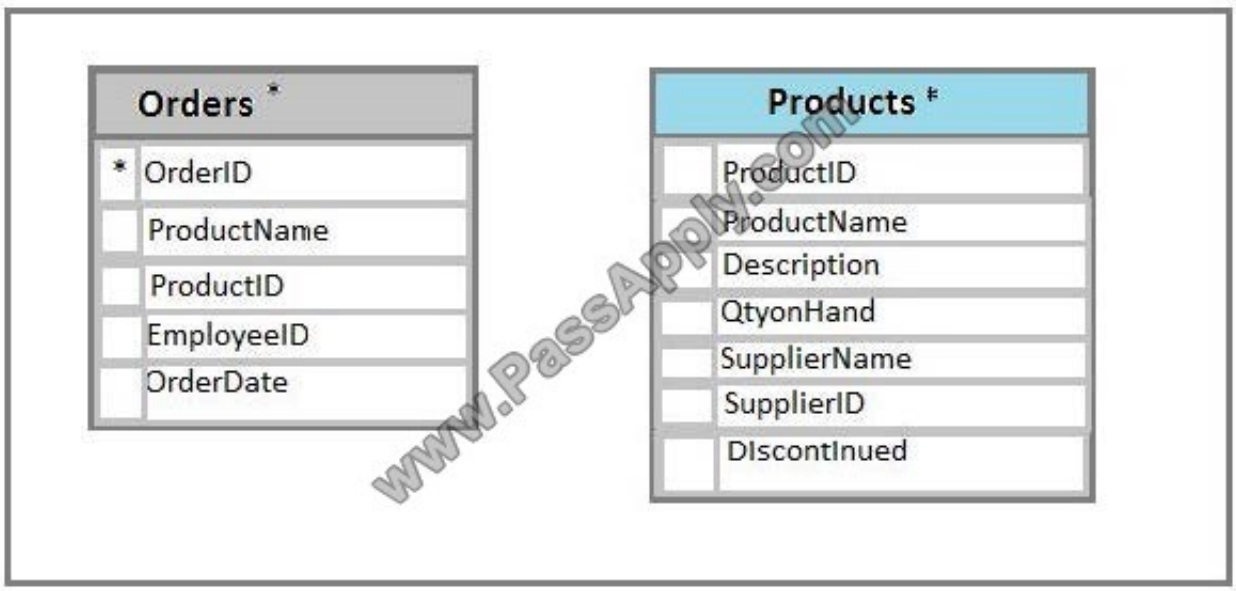

The customer table includes a column that stores the data for the last order that the customer placed.

You plan to create a table named Leads. The Leads table is expected to contain approximately 20,000 records. Storage requirements for the Leads table must be minimized.

You need to create triggers that meet the following requirements:

In the table below, identify the trigger types that meet the requirements.

NOTE: Make only selection in each column. Each correct selection is worth one point.

Hot Area:

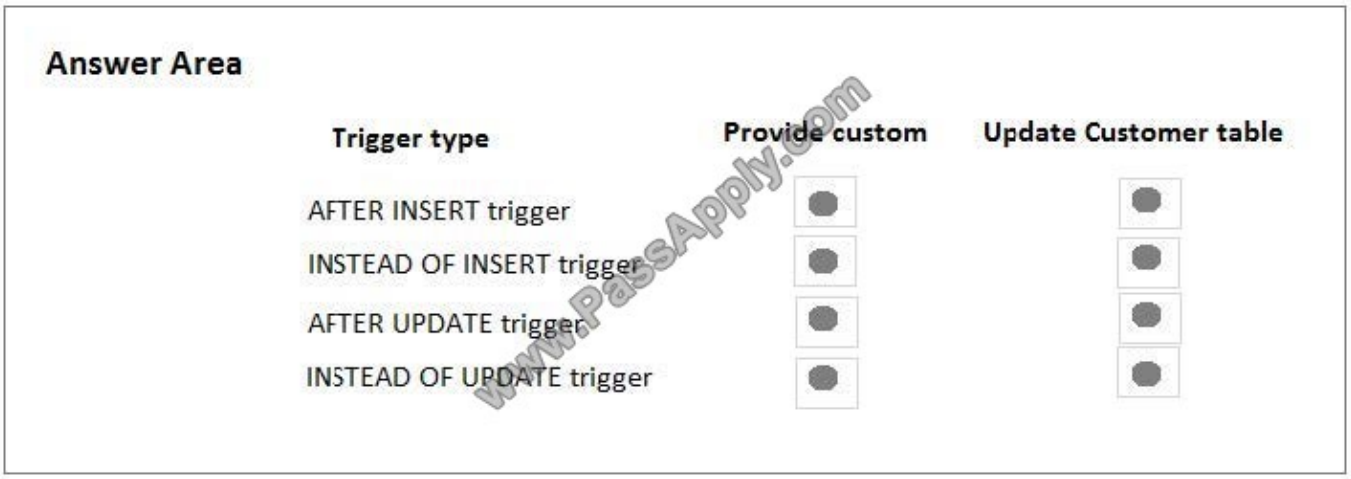

Correct Answer:

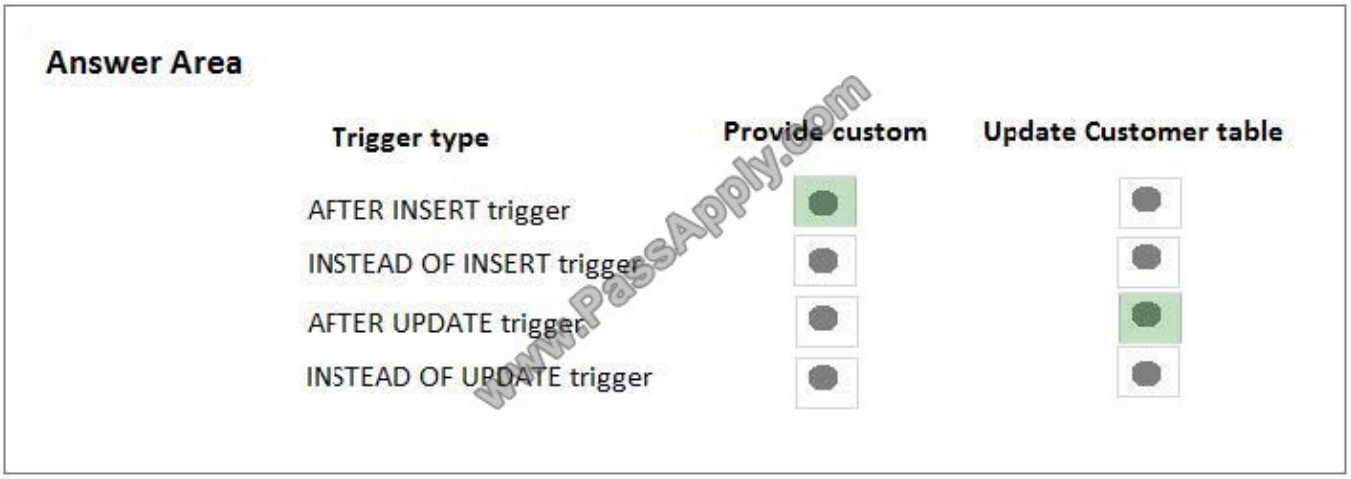

INSTEAD OF INSERT triggers can be defined on a view or table to replace the standard action of the INSERT statement.

AFTER specifies that the DML trigger is fired only when all operationsspecified in the triggering SQL statement have executed successfully.

References:

https://technet.microsoft.com/en-us/library/ms175089(v=sql.105).aspx

QUESTION 10

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 SQL Server instances. There are many SQL jobs that run during off-

peak hours.

You must monitor the SQL Server instances in real time and optimize the server to maximize throughput, response time, and overall SQL performance.

You need to collect query performance data while minimizing the performance impact on the SQL Server.

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 Performance Monitor 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: C

SQL Server Data Collector is a feature for performance monitoring and tuning available in SQL Server Management Studio.

Integration Services packages transform and load the collected data into the Microsoft Data Warehouse database.

Collection sets are defined and deployed on a server instance and can be run independently of each other. Each collection set can be applied to a target that matches the target types of all the collector types that are part of a collection set.

The collection set is run by a SQL Server Agent job or jobs, and data is uploaded to the management data warehouse on a predefined schedule.

Predefined data collection sets include:

The Query Statistics data collection set collects information about query statistics, activity, execution plans and text on the SQL Server instance. It does not store all executed statements, only 10 worst performing ones.

Disk Usage data collection set collects information about disk space used by both data and log files for all databases on the SQL Server instance, growth trends, and average day growth.

Etc.

References:

http://www.sqlshack.com/sql-server-performance-monitoring-data-collector/

QUESTION 11

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a database that contains a table named Employees. The table stored information about the employees of your company. You need to implement the following auditing rules for the Employees table:

-Record any changes that are made to the data in the Employees table.

-

Customize the data recorded by the audit operations.

Solution: You implement a user-defined function on the Employees table.

Does the solution meet the goal?

A.

Yes

B.

No

Correct Answer: A

SQL Server 2016 provides two features that track changes to data in a database: change data capture and change tracking. These features enable applications to determine the DML changes (insert, update, and delete operations) that were

made to user tables in a database.

Change data is made available to change data capture consumers through table-valued functions (TVFs).

References:https://msdn.microsoft.com/en-us/library/cc645858.aspx

QUESTION 12

You have a database that is experiencing deadlock issues when users run queries.

You need to ensure that all deadlocks are recorded in XML format.

What should you do?

A. Create a Microsoft SQL Server Integration Services package that uses sys.dm_tran_locks.

B. Enable trace flag 1224 by using the Database Cpmsistency Checker(BDCC).

C. Enable trace flag 1222 in the startup options for Microsoft SQL Server.

D. Use the Microsoft SQL Server Profiler Lock:Deadlock event class.

Correct Answer: C

When deadlocks occur, trace flag 1204 and trace flag 1222 return information that is capturedin the SQL Server error log. Trace flag 1204 reports deadlock information formatted by each node involved in the deadlock. Trace flag 1222 formats deadlock information, first by processes and then by resources.

The output format for Trace Flag 1222 only returns information in an XML-like format.

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