



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 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 stores information about the employees of your company.

You need to implement and enforce the following business rules:

Limit the values that are accepted by the Salary column.

Prevent salaries less than \$15,000 and greater than \$300,000 from being entered.

Determine valid values by using logical expressions.

Do not validate data integrity when running DELETE statements.

Solution: You implement a foreign key on the table.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

A check constraint is needed.

References: https://en.wikipedia.org/wiki/Check_constraint

QUESTION 2

You run the following Transact-SQL statements:

```
CREATE TABLE Customers (  
    CustomerID INT NOT NULL IDENTITY PRIMARY KEY CLUSTERED,  
    CustomerName NVARCHAR (100) UNIQUE NOT NULL  
)  
  
CREATE TABLE Orders (  
    OrderID INT NOT NULL IDENTITY PRIMARY KEY CLUSTERED,  
    CustomerID INT NOT NULL REFERENCES Customers (CustomerID),  
    OrderDate DATE NOT NULL  
)  
  
CREATE VIEW v_CustomerOrder  
AS SELECT  
    b.CustomerName, a.OrderID, a.OrderDate,  
    (SELECT COUNT(*) FROM Orders c WHERE c.CustomerID = a.CustomerID) AS CustomerOrderCount  
FROM Orders a  
INNER JOIN Customers b ON a.CustomerID = b.CustomerID
```



Records must only be added to the Orders table by using the view. If a customer name does not exist, then a new customer name must be created. You need to ensure that you can insert rows into the Orders table by using the view.

- A. Add the CustomerID column from the Orders table and the WITH CHECK OPTION statement to the view.
- B. Create an INSTEAD OF trigger on the view.
- C. Add the WITH SCHEMABINDING statement to the view statement and create a clustered index on the view.
- D. Remove the subquery from the view, add the WITH SCHEMABINDING statement, and add a trigger to the Orders table to perform the required logic.

Correct Answer: A

The WITH CHECK OPTION clause forces all data-modification statements executed against the view to adhere to the criteria set within the WHERE clause of the SELECT statement defining the view. Rows cannot be modified in a way that causes them to vanish from the view.

References: <http://www.informit.com/articles/article.aspx?p=130855andseqNum=4>

QUESTION 3

You are creating the following two stored procedures: A natively-compiled stored procedure An interpreted stored procedure that accesses both disk-based and memory-optimized tables

Both stored procedures run within transactions.

You need to ensure that cross-container transactions are possible.

Which setting or option should you use?

- A. the SET TRANSACTION_READ_COMMITTED isolation level for the connection
- B. the SERIALIZABLE table hint on disk-based tables
- C. the SET MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT=ON option for the database
- D. the SET MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT=OFF option for the database

Correct Answer: C

Provide a supported isolation level for the memory-optimized table using a table hint, such as WITH (SNAPSHOT). The need for the WITH (SNAPSHOT) hint can be avoided through the use of the database option MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT. When this option is set to ON, access to a memory-optimized table under a lower isolation level is automatically elevated to SNAPSHOT isolation.

Incorrect Answers:

B: Accessing memory optimized tables using the READ COMMITTED isolation level is supported only for autocommit transactions. It is not supported for explicit or implicit transactions. References: <https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/transactions-with-memory-optimized-tables?view=sql-server-2017>

QUESTION 4



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 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 within a transaction and roll back if a failure 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 and provides 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 5

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

You have a Microsoft SQL Server database named DB1 that contains the following tables:



Table Name	Description
TBL1	<ul style="list-style-type: none">The table will contain 10 million records.The frequency of inserting, updating, and deleting records is low.
TBL2	<ul style="list-style-type: none">The table will contain 1 million records.The frequency of inserting, updating, and deleting records is low.

Users frequently run the following query:

```
SELECT TBL1.Column2, TBL2.Column2, SUM(TBL1.Column3), SUM(TBL1.Column4 * TBL1.Column5 * (TBL2.Column4 - TBL2.Column3))  
FROM TBL1  
INNER JOIN TBL2 TBL1.Column1 = TBL2.Column1  
GROUP BY TBL1.Column2, TBL2.Column2
```

Users report that the query takes a long time to return results.

You need to minimize the amount of time requires for the query to return data.

What should you do?

- A. Create clustered indexes on TBL1 and TBL2.
- B. Create a clustered index on TBL1.Create a nonclustered index on TBL2 and add the most frequently queried column as included columns.
- C. Create a nonclustered index on TBL2 only.
- D. Create UNIQUE constraints on both TBL1 and TBL2. Create a partitioned view that combines columns from TBL1 and TBL2.
- E. Drop existing indexes on TBL1 and then create a clustered columnstore index. Create a nonclustered columnstore index on TBL1.Create a nonclustered index on TBL2.
- F. Drop existing indexes on TBL1 and then create a clustered columnstore index. Create a nonclustered columnstore index on TBL1.Make no changes to TBL2.
- G. Create CHECK constraints on both TBL1 and TBL2. Create a partitioned view that combines columns from TBL1 and TBL2.
- H. Create an indexed view that combines columns from TBL1 and TBL2.

Correct Answer: H

[Latest 70-762 Dumps](#)

[70-762 VCE Dumps](#)

[70-762 Study Guide](#)