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:

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. DDL 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 programming model 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

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

You are developing a new application that uses a stored procedure. The stored procedure inserts thousands of records as a single batch into the Employees table.

Users report that the application response time has worsened since the stored procedure was updated. You examine disk-related performance counters for the Microsoft SQL Server instance and observe several high values that include a

disk performance issue. You examine wait statistics and observe an unusually high WRITELOG value.

You need to improve the application response time.

Solution: You replace the stored procedure with a user-defined function.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

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

---

QUESTION 3

You run the following Transact-SQL following statement:

```
CREATE TABLE Customer(
    CustomerId INT IDENITY (1, 1) PRIMARY KEY,
    Code CHAR(S) NOT NULL,
    FirstName VARCHAR (50) NOT NULL,
    LastName VARCHAR (50) NOT NULL
)
```

Customer records may be inserted individually or in bulk from an application.

You observe that the application attempts to insert duplicate records.

You must ensure that duplicate records are not inserted and bulk insert operations continue without notifications.

Which Transact-SQL statement should you run?

A. CREATE UNIQUE NONCLUSTERED INDEX IX_Customer_Code ON Customer (Code) WITH (ONLINE = OFF)

B. CREATE UNIQUE INDEX IX_CUSTOMER_Code O Customer (Code) WITH (IGNORE_DUP_KEY = ON)

C. CREATE UNIQUE INDEX IX Customer Code ON Customer (Code) WITH (IGNORE DUP KEY =OFF)

D. CREATE UNIQUE NONCLUSTERED INDEX IX_Customer_Code ON Customer (Code)

E. CREATE UNIQUE NONCLUSTERED INDEX IX_Customer_Code ON Customer (Code) WITH (ONLINE = ON)

Correct Answer: B

IGNORE_DUP_KEY = { ON | OFF } specifies the error response when an insert operation attempts to insert duplicate key values into a unique index. The IGNORE_DUP_KEY option applies only to insert operations after the index is created

or rebuilt. The option has no effect when executing CREATE INDEX, ALTER INDEX, or UPDATE. The default is OFF.

Incorrect Answers:

ONLINE = { ON | OFF } specifies whether underlying tables and associated indexes are available for queries and data modification during the index operation. The default is OFF.

References: https://docs.microsoft.com/en-us/sql/t-sql/statements/create-index-transact-sql?view=sql-server-2017

QUESTION 4

You need to create a view that can be indexed. You write the following statement.

```
01    create view myView2
02
03    as
04    select col1, col2
05    from dbo.test
06    go
```

What should you add at line 02?

A. with check_option

B. with recompile

C. with view_metadata

D. with schemabinding

Correct Answer: D

The following steps are required to create an indexed view and are critical to the successful implementation of the indexed view:

1.

 Verify the SET options are correct for all existing tables that will be referenced in the view.

2.

Verify that the SET options for the session are set correctly before you create any tables and the view.

3.

Verify that the view definition is deterministic.

4.

Create the view by using the WITH SCHEMABINDING option.

5.

Create the unique clustered index on the view.

References: https://docs.microsoft.com/en-us/sql/relational-databases/views/create-indexed-views?view=sql-server-2017

QUESTION 5

Your company runs end-of-the-month accounting reports. While the reports run, other financial records are updated in the database.

Users report that the reports take longer than expected to run.

You need to reduce the amount of time it takes for the reports to run. The reports must show committed data only.

What should you do?

A. Use the NOLOCK option.

B. Execute the DBCC UPDATEUSAGE statement.

C. Use the max worker threads option.

D. Use a table-valued parameter.

E. Set SET ALLOW_SNAPSHOT_ISOLATION to ON.

F. Set SET XACT_ABORT to ON.

G. Execute the ALTER TABLE T1 SET (LOCK_ESCALATION = AUTO); statement.

H. Use the OUTPUT parameters.

Correct Answer: E

Snapshot isolation enhances concurrency for OLTP applications.

Once snapshot isolation is enabled, updated row versions for each transaction are maintained in tempdb. A unique transaction sequence number identifies each transaction, and these unique numbers are recorded for each row version. The transaction works with the most recent row versions having a sequence number before the sequence number of the transaction. Newer row versions created after the transaction has begun are ignored by the transaction. References: https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/snapshot-isolation-in-sql-server