Leads4Pass

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.leads4pass.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 present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals. You have a table that has a clustered index and a nonclustered index. The indexes use different columns from the table. You have a query named Query1 that uses the nonclustered index. Users report that Query1 takes a long time to report results. You run Query1 and review the following statistics for an index seek operation:

Index Seek (NonClustered)
Scan a particular range of rows from a nonclustered index.

Physical Operation	Index Seek
Logical Operation	Index Seek
Actual Execution Mode	Row
Actual Number of Rows	3571454
Actual Number of Batches	0
Estimated I/O Cost	0.0093577
Estimated Operator Cost	0.0107304 (0%)
Estimated CPU Cost	0.0013727
Estimated Subtree Cost	0.0107304
Estimated Number of Executions	1
Number of Executions	8
Estimated Number of Rows	0
Estimated Row Size	19 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	100

You need to resolve the performance issue. Solution: You defragment both indexes. Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

We see Actual Number of Row is 3571454, while Estimated Number of Rows is 0. This indicates that the statistics are old, and need to be updated.

QUESTION 3

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 table that has a clustered index and a nonclustered index. The indexes use different columns from the table. You have a query named Query1 that uses the nonclustered index.

Users report that Query1 takes a long time to report results. You run Query1 and review the following statistics for an index seek operation:

Index Seek (NonClustered)
Scan a particular range of rows from a nonclustered index.

Physical Operation	Index Seek
Logical Operation	Index Seek
Actual Execution Mode	Row
Actual Number of Rows	3571454
Actual Number of Batches	0
Estimated I/O Cost	0.0093577
Estimated Operator Cost	0.0107304 (0%)
Estimated CPU Cost	0.0013727
Estimated Subtree Cost	0.0107304
Estimated Number of Executions	1
Number of Executions	8
Estimated Number of Rows	0
Estimated Row Size	19 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	100

You need to resolve the performance issue.

Solution: You update statistics for the nonclustered index.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: A

We see Actual Number of Row is 3571454, while Estimated Number of Rows is 0. This indicates that the statistics are old, and need to be updated.

QUESTION 4

You are designing a solution for a company that operates retail stores. Each store has a database that tracks sales transactions. You create a summary table in the database at the corporate office. You plan to use the table to record the

quantity of each product sold at each store on each day. Managers will use this data to identify reorder levels for products.

Every evening stores, must transmit sales data to the corporate office. The data must be inserted into the summary table that includes the StoreID, ProductID, Qtysold, Totprodsales, and Datesold columns.

You need to prevent duplicate rows in the summary table. Each row must uniquely identify the store that sold the product and the total amount sold for that store on a specific date.

What should you include in your solution?

A. Create a unique constraint.

B. Create a foreign key constraint to the StoreID column in each of the store tables.

C. Create a rule and bind it to the StoreID column.

D. Create a check constraint.

E. Create a table-valued user-defined function.

Correct Answer: A

You can use UNIQUE constraints to make sure that no duplicate values are entered in specific columns that do not participate in a primary key. Although both a UNIQUE constraint and a PRIMARY KEY constraint enforce uniqueness, use a UNIQUE constraint instead of a PRIMARY KEY constraint when you want to enforce the uniqueness of a column, or combination of columns, that is not the primary key.

Incorrect Answers:

D: CHECK constraints enforce domain integrity by limiting the values that are accepted by one or more columns.

References: https://docs.microsoft.com/en-us/sql/relational-databases/tables/unique-constraints-and-check-constraints?view=sql-server-2017

QUESTION 5

You have a reporting application that uses a table named Table1. You deploy a new batch update process to perform updates to Table1.

The environment is configured with the following properties:

The database is configured with the default isolation setting.

The application and process use the default transaction handling.

You observe the application cannot access any rows that are in use by the process.

You have the following requirements:

Ensure the application is not blocked by the process.

Ensure the application has a consistent view of the data

Ensure the application does not read dirty data.

You need to resolve the issue and meet the requirements with the least amount of administrative effort.

What should you do?

A. Enable the database for the ALLOW_SNAPSHOT_ISOLATION isolation level. Modify the application for the

SERIALIZABLE isolation level.

B. Enable the database for the READ_COMITTED_SNAPSHOT isolation level.

C. Enable the application for the WITH (NOLOCK) hint.

D. Enable the database for the ALLOW_SNAPSHOT_ISOLATION isolation level. Modify the application and the update process for the SNAPSHOT isolation level.

Correct Answer: B

Snapshot isolation must be enabled by setting the ALLOW_SNAPSHOT_ISOLATION ON database option before it is used in transactions. This activates the mechanism for storing row versions in the temporary database (tempdb).

READ COMMITTED is the default isolation level for SQL Server. It prevents dirty reads by specifying that statements cannot read data values that have been modified but not yet committed by other transactions. Other transactions can still

modify, insert, or delete data between executions of individual statements within the current transaction, resulting in non-repeatable reads, or "phantom" data.

Incorrect Answers:

A: SERIALIZABLE is the most restrictive isolation level, because it locks entire ranges of keys and holds the locks until the transaction is complete. It encompasses REPEATABLE READ and adds the restriction that other transactions cannot insert new rows into ranges that have been read by the transaction until the transaction is complete.

References: https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/snapshot-isolation-in-sql-server

Latest 70-762 Dumps   70-762 PDF Dumps   70-762 Practice Test