

## 70-762<sup>Q&As</sup>

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 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 stored procedure on the Employees table.

Does the solution meet the goal?

A.

Yes

B.

No

Correct Answer: B

We should use table-valued functions, not procedures, to customize the recorded change data. References: <https://msdn.microsoft.com/en-us/library/cc645858.aspx>

---

**QUESTION 2**

Database users report that SELECT statements take a long time to return results. You run the following Transact-SQL statement:

```
SELECT OBJECT_NAME([object_id]) AS [object_name],  
d.equality_columns, d.inequality_columns, d.included_columns  
FROM sys.dm_db_missing_index_details;
```

Object_name	Equality_columns	Inequality_columns	Included_columns
[Users]	[CountryCode]	[UserStatus]	[UserName]

You need to create one nonclustered covering index that contains all of the columns in the above table. You must minimize index key size. Which Transact-SQL statement should you run?

A. CREATE NONCLUSTERED INDEX IX\_User ON Users (CountryCode, UserName);

B. CREATE NONCLUSTERED INDEX IX\_User ON Users (CountryCode, UserStatus) INCLUDE (UserName);

C. CREATE NONCLUSTERED INDEX IX\_User ON Users (CountryCode, UserStatus, UserName);

D. CREATE NONCLUSTERED INDEX IX\_User ON Users (UserStatus, CountryCode) INCLUDE (UserName);

Correct Answer: D

Use the UserStatus as the first column in the index, as it is an in\_equality column. Incorrect Answers:

A: UserStatus is not included.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/indexes/create-indexes-with-included-columns>

### QUESTION 3

Note: This question is part of a series of questions that use the same or similar answer choices. As 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"><li>The table will contain 10 million records.</li><li>The frequency of inserting, updating, and deleting records is low.</li></ul>
TBL2	<ul style="list-style-type: none"><li>The table will contain 1 million records.</li><li>The frequency of inserting, updating, and deleting records is low.</li></ul>

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

---

#### QUESTION 4

You have a view that includes an aggregate.

You must be able to change the values of columns in the view. The changes must be reflected in the tables that the view uses.

You need to ensure that you can update the view.

What should you create?

- A. a nonclustered index
- B. a schema-bound view
- C. a stored procedure
- D. an INSTEAD OF trigger

Correct Answer: B

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. Views or tables that participate in a view created with the SCHEMABINDING clause cannot be dropped unless that view is dropped or changed so that it no longer has schema binding.

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-view-transact-sql>

---

#### QUESTION 5

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 as a 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.

[70-762 VCE Dumps](#)

[70-762 Practice Test](#)

[70-762 Braindumps](#)