



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.pass4itsure.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. 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">• Column1 is configured as the primary key.• The table will contain 20 million records.• The table will contain historical data.• Most queries of TBL1 return a high percentage of rows from the table with aggregates.
TBL2	<ul style="list-style-type: none">• Column1 is configured as the primary key.• The table will contain 25 million records.• The frequency of updates and deleted records in TBL2 is low.• Most queries of TBL2 return a high percentage of rows from the table with aggregates.

There are no foreign key relationships between TBL1 and TBL2.

You need to minimize the amount of time required for queries that use data from TBL1 and TBL2 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: B



References: <http://www.sqlservergeeks.com/sql-server-indexing-for-aggregates-in-sql-server/>

QUESTION 2

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.

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">• The table has 25 columns.• The table will contain 10 million records• Approximately 100,000 records will be inserted monthly.
TBL2	<ul style="list-style-type: none">• The table has 25 columns.• The table will contain 100,000 records.• The frequency of inserting, updating, and deleting records is low

You frequently run the following queries:



```
SELECT *  
FROM TBL1  
WHERE Column1 BETWEEN '01/01/2016' AND '30/04/2016'
```

```
SELECT Column5, Column6  
FROM TBL2  
WHERE Column2 = 'ABC156XYZ'
```

There are no foreign key relationships between TBL1 and TBL2.

You need to minimize the amount of time required for the two queries to return records from the tables.

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: B

QUESTION 4

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 drop the nonclustered index. 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 5

You have a table that stores transactions partitioned by year. Users occasionally experience performance issues when they access the table. The table is locked exclusively when the records are updated. You need to prevent exclusive locks

on the table and maintain data integrity.

What should you do?

A. Set the LOCK_EXCALATION property to DISABLE

B. Set the DATA_COMPRESSION property to ROW at the partition level



- C. Set the LOCK_EXCALATION property to AUTO
- D. Set the LOCK_EXCALATION property to TABLE
- E. Set the DATA_COMPRESSION property to PAGE at the partition level

Correct Answer: C

The default lock escalation mode is called TABLE, it implements SQL Server's lock escalation on all types of tables whether partitioned or not partitioned.

There are two more lock escalation modes: AUTO and DISABLE.

The AUTO mode enables lock escalation for partitioned tables only for the locked partition. For non-partitioned tables it works like TABLE.

The DISABLE mode removes the lock escalation capability for the table and that is important when concurrency issues are more important than memory needs for specific tables.

References:

<https://www.mssqltips.com/sqlservertip/4359/altering-lock-escalation-for-sql-server-tables/>

[Latest 70-762 Dumps](#)

[70-762 PDF Dumps](#)

[70-762 Brindumps](#)