



# 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.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

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 2

You are developing a database reporting solution for a table that contains 900 million rows and is 103 GB.

The table is updated thousands of times a day, but data is not deleted.

The SELECT statements vary in the number of columns used and the amount of rows retrieved.

You need to reduce the amount of time it takes to retrieve data from the table. The must prevent data duplication.

Which indexing strategy should you use?

- A. a nonclustered index for each column in the table
- B. a clustered columnstore index for the table
- C. a hash index for the table
- D. a clustered index for the table and nonclustered indexes for nonkey columns

Correct Answer: B

Columnstore indexes are the standard for storing and querying large data warehousing fact tables. It uses column-based data storage and query processing to achieve up to 10x query performance gains in your data warehouse over



traditional row-oriented storage.

A clustered columnstore index is the physical storage for the entire table.

Generally, you should define the clustered index key with as few columns as possible.

A nonclustered index contains the index key values and row locators that point to the storage location of the table data. You can create multiple nonclustered indexes on a table or indexed view. Generally, nonclustered indexes should be

designed to improve the performance of frequently used queries that are not covered by the clustered index.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview?view=sql-server-2017>

---

### QUESTION 3

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.

Your company has employees in different regions around the world.

You need to create a database table that stores the following employee attendance information:

Employee ID date and time employee checked in to work date and time employee checked out of work

Date and time information must be time zone aware and must not store fractional seconds. Solution: You run the following Transact-SQL statement: Does the solution meet the goal?

```
CREATE TABLE [dbo].[EmployeeAttendance] (  
    EmployeeID int NOT NULL,  
    DateChekedIn datetime2 NOT NULL,  
    DateCheclOut datetime2 NOT NULL)
```

A. Yes

B. No

Correct Answer: B

Datetime2 stores fractional seconds.

Datetime2 defines a date that is combined with a time of day that is based on 24-hour clock. datetime2 can be considered as an extension of the existing datetime type that has a larger date range, a larger default fractional precision, and

optional user-specified precision.

References: <https://docs.microsoft.com/en-us/sql/t-sql/data-types/datetime2-transact-sql?view=sql-server-2017>



<https://msdn.microsoft.com/en-us/library/bb677335.aspx>

#### QUESTION 4

Note: This question is part of a series of questions that use the same 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 are developing an application to track customer sales. You create tables to support the application. You need to create a database object that meets the following data entry requirements: What should you create?

Table name	Data entry requirements
Customers	Inserts must be grouped in a batch of Transact-SQL statements.
CustomerTransactions	Inserts must be grouped in a batch of Transact-SQL statements. Inserts must be grouped in a batch of Transact-SQL statements. All of the data modifications made must become a permanent part of the database if data entry is successful.
Invoices	All data entry modifications must be erased if an error occurs. Inserts must be grouped in a batch of Transact-SQL statements. All of the data modifications made must become a permanent part of the database if data entry is successful.
Orders	All data entry modifications must be erased if an error occurs.

- A. extended procedure
- B. CLR procedure
- C. user-defined procedure
- D. DML trigger
- E. DDL trigger
- F. scalar-valued function
- G. table-valued function

Correct Answer: C

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

#### QUESTION 5

You have several real-time applications that constantly update data in a database. The applications run more than 400 transactions per second that insert and update new metrics from sensors.

A new web dashboard is released to present the data from the sensors. Engineers report that the applications take



longer than expected to commit updates.

You need to change the dashboard queries to improve concurrency and to support reading uncommitted data.

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: A

The NOLOCK hint allows SQL to read data from tables by ignoring any locks and therefore not being blocked by other processes. This can improve query performance, but also introduces the possibility of dirty reads.

Incorrect Answers:

F: When SET XACT\_ABORT is ON, if a Transact-SQL statement raises a run-time error, the entire transaction is terminated and rolled back.

G: DISABLE, not AUTO, would be better.

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.

Note: SQL Server's locking mechanism uses memory resources to maintain locks. In situations where the number of row or page locks increases to a level that decreases the server's memory resources to a minimal level, SQL Server's locking strategy converts these locks to entire table locks, thus freeing memory held by the many single row or page locks to one table lock. This process is called lock escalation, which frees memory, but reduces table concurrency.

References: <https://www.mssqltips.com/sqlservertip/2470/understanding-the-sql-server-nolock-hint/>

[Latest 70-762 Dumps](#)

[70-762 PDF Dumps](#)

[70-762 Exam Questions](#)