



# 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

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

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 independent of the other questions in this series.

Information and details provided in a question apply only to that question.

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run during off-peak hours.

You must monitor and optimize the SQL Server to maximize throughput, response time, and overall SQL performance.

You need to examine delays in executed threads, including errors with queries and batches.

What should you do?

A. Create a sys.dm\_os\_waiting\_tasks query.

B. Create a sys.dm\_exec\_sessions query.



- C. Create a Performance Monitor Data Collector Set.
- D. Create a sys.dm\_os\_memory\_objects query.
- E. Create a sp\_configure `max server memory` query.
- F. Create a SQL Profiler trace.
- G. Create a sys.dm\_os\_wait\_stats query.
- H. Create an Extended Event.

Correct Answer: G

sys.dm\_os\_wait\_stats returns information about all the waits encountered by threads that executed. You can use this aggregated view to diagnose performance issues with SQL Server and also with specific queries and batches. Incorrect Answers:

A: sys.dm\_os\_waiting\_tasks returns information about the wait queue of tasks that are waiting on some resource.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-os-wait-stats-transact-sql>

---

### QUESTION 3

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 an object that meets to following requirements:

Run managed code packaged in an assembly that was created in the Microsoft.NET Framework and uploaded in Microsoft SQL Server.

Run written a transaction and roll back if a failure occurs.

Run when a table is created or modified.

What should you create?

- 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: B

The common language runtime (CLR) is the heart of the Microsoft .NET Framework and provides the execution environment for all .NET Framework code. Code that runs within the CLR is referred to as managed code. With the CLR hosted in Microsoft SQL Server (called CLR integration), you can author stored procedures, triggers, user-defined functions, user-defined types, and user-defined aggregates in managed code. References:

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/introduction-to-sql-server-clr-integration>

---

#### 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 need to create a stored procedure that

updates the Customer, CustomerInfo, OrderHeader, and OrderDetails tables in order.

You need to ensure that the stored procedure:

Runs within a single transaction.

Commits updates to the Customer and CustomerInfo tables regardless of the status of updates to the OrderHeader and OrderDetail tables.

Commits changes to all four tables when updates to all four tables are successful.

Solution: You create a stored procedure that includes the following Transact-SQL segment:

```
BEGIN TRY
    BEGIN TRAN
        UPDATE Customer ...
        UPDATE CustomerInfo ...
        UPDATE OrderHeader ...
        UPDATE OrderDetail ...
    COMMIT TRAN
END TRAN
BEGIN CATCH
    IF XACT STATE() = 1
        ROLLBACK TRAN
END CATCH
```

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

All four tables are updated in a single transaction.



Need to handle the case where the first two updates (OrderHeader, OrderDetail) are successful, but either the 3rd or the 4th (OrderHeader, OrderDetail) fail. Can add a variable in the BEGIN TRY block, and test the variable in the BEGIN CATCH block.

References: <https://docs.microsoft.com/en-us/sql/t-sql/language-elements/begin-transaction-transact-sql>

---

## QUESTION 5

Note: The 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 question in the series.

Information and details provided in a question apply only to that question.

You have a reporting database that includes a non-partitioned fact table named Fact\_Sales. The table is persisted on disk.

Users report that their queries take a long time to complete. The system administrator reports that the table takes too much space in the database. You observe that there are no indexes defined on the table, and many columns have repeating values.

You need to create the most efficient index on the table, minimize disk storage and improve reporting query performance.

What should you do?

- A. Create a clustered index on the table.
- B. Create a nonclustered index on the table.
- C. Create a nonclustered filtered index on the table.
- D. Create a clustered columnstore index on the table.
- E. Create a nonclustered columnstore index on the table.
- F. Create a hash index on the table.

Correct Answer: D

The columnstore index is the standard for storing and querying largedata 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, and up to 10x data compression over the uncompressed data size.

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

---

## QUESTION 6

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questions sets might have more than one correct solution, while others might not have a correct solution.



After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a 3-TB database. The database server has 64 CPU cores.

You plan to migrate the database to Microsoft Azure SQL Database.

You need to select the service tier for the Azure SQL database. The solution must meet or exceed the current processing capacity.

Solution: You select the Standard service tier.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Premium service is required for 3 TB of storage. Single database DTU and storage limits

	Basic	Standard	Premium
Maximum storage size	2 GB	1 TB	4 TB
Maximum DTUs	5	3000	4000

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tiers-dtu>

## QUESTION 7

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 8

### Background

You have a database named HR1 that includes a table named Employee.

You have several read-only, historical reports that contain regularly changing totals. The reports use multiple queries to estimate payroll expenses. The queries run concurrently. Users report that the payroll estimate reports do not always run. You must monitor the database to identify issues that prevent the reports from running.

You plan to deploy the application to a database server that supports other applications. You must minimize the amount of storage that the database requires.

### Employee Table

You use the following Transact-SQL statements to create, configure, and populate the Employee table:

```
CREATE TABLE dbo.Employee
  (EmployeeId INT PRIMARY KEY,
  LastName varchar(50),
  FirstName varchar(50),
  DepartmentId int,
  HireDate datetime,
  TerminationDate datetime,
  SupervisorId int,
  CostCenterNumber int,
  EmployeeStatus int,
  EmployeePayRate int)
) GO
CREATE INDEX IX_1 on dbo.Employee (LastName, FirstName, DepartmentId) INCLUDE (HireDate)
CREATE INDEX IX_2 on dbo.Employee (LastName) INCLUDE (EmployeeId, FirstName, DepartmentId)
CREATE INDEX IX_3 on dbo.Employee (LastName, FirstName) INCLUDE (DepartmentId)
CREATE INDEX IX_4 on dbo.Employee (LastName, FirstName) INCLUDE (HireDate, DepartmentId)
GO

INSERT INTO Employee (EmployeeID, LastName, CostCenterNumber) VALUES(1001, 'Employee A', 3001001)
INSERT INTO Employee (EmployeeID, LastName, CostCenterNumber) VALUES(1002, 'Employee B', 3001001)
GO
```

### Application

You have an application that updates the Employees table. The application calls the following stored procedures simultaneously and asynchronously:

UspA: This stored procedure updates only the EmployeeStatus column.

UspB: This stored procedure updates only the EmployeePayRate column.

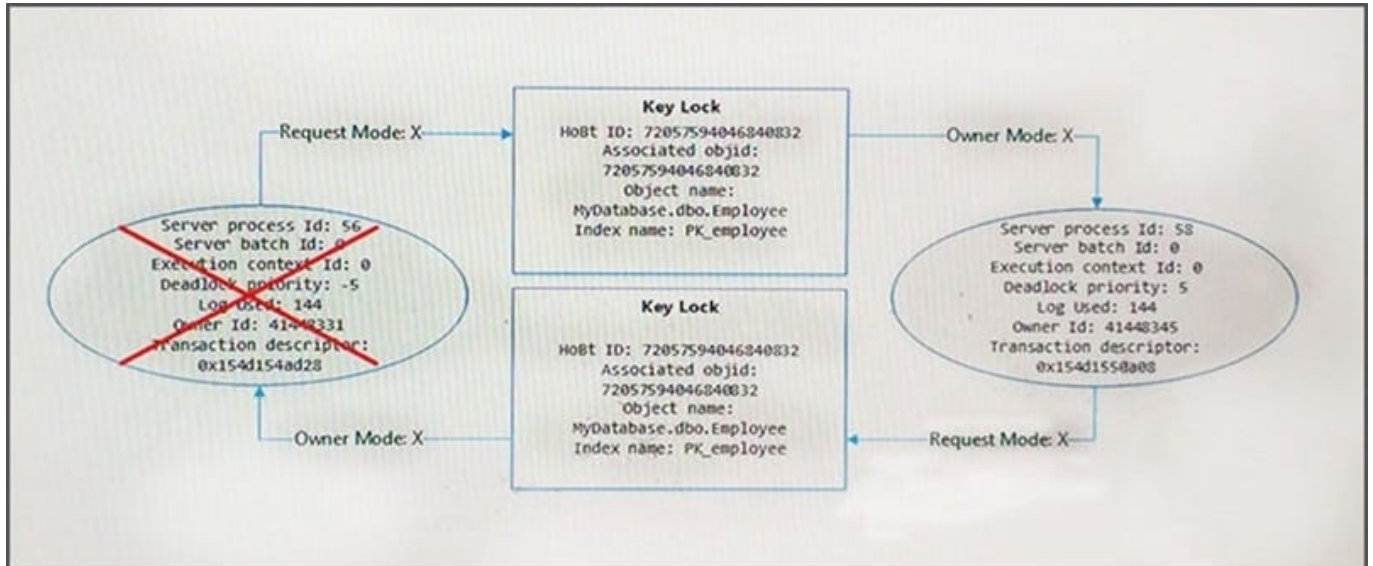
The application uses views to control access to data. Views must meet the following requirements:



Allow user access to all columns in the tables that the view accesses.

Restrict updates to only the rows that the view returns.

Exhibit



You are analyzing the performance of the database environment. You discover that locks that are held for a long period of time as the reports are generated.

You need to generate the reports more quickly. The database must not use additional resources.

What should you do?

- A. Update all FROM clauses of the DML statements to use the IGNORE\_CONSTRAINTS table hint.
- B. Modify the report queries to use the UNION statement to combine the results of two or more queries.
- C. Apply a nonclustered index to all tables used in the report queries.
- D. Update the transaction level of the report query session to READ UNCOMMITTED.

Correct Answer: D

Transactions running at the READ UNCOMMITTED level do not issue shared locks to prevent other transactions from modifying data read by the current transaction. This is the least restrictive of the isolation levels. References: [https://technet.microsoft.com/en-us/library/ms173763\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms173763(v=sql.105).aspx)

## QUESTION 9

You have a Microsoft Azure SQL Database. You enable Query Store for the database and configure the store to use the following settings:

SIZE\_BASED\_CLEANUP\_MODE = OFF

STALE\_QUERY\_THRESHOLD\_DAYS = 60





MAX\_STORAGE\_SIZE\_MB = 100

QUERY\_CAPTURE\_MODE = ALL

You use Azure Query Performance Insight to review queries. You observe that new queries are not displayed after 15 days and that the Query Store is set to read-only mode.

If the Query Store runs low on data space, the store must prioritize queries that run regularly or queries that consume applicant resources.

You must set the Query Store to read\_write mode and determine the performance of queries from the past 60 days.

Which three actions should you perform? Each correct step presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Set the value of the CLEANUP\_POLICY setting to (STALE\_QUERY\_THRESHOLD\_DAYS = 75)
- B. Set the value of the QUERY\_CAPTURE\_MODE setting to AUTO
- C. Increase the value for the MAX\_STORAGE\_SIZE\_MB setting
- D. Set the value of the SIZE\_BASED\_CLEANUP\_MODE setting to AUTO
- E. In the Azure portal, navigate to Query Performance Insight. Use the Custom tab to select a period of 2 months.

Correct Answer: BCD

B: Capture mode:

All ?Captures all queries. This is the default option.

Auto ?Infrequent queries and queries with insignificant cost are ignored. (Ad hoc recommended)

None ?Query Store stops capturing new queries.

C: Max Size (MB): Specifies the limit for the data space that Query Store can consume within the database. This is the most important setting that directly affects operation mode of the Query Store.

While Query Store collects queries, execution plans and statistics, its size in the database grows until this limit is reached. When that happens, Query Store automatically changes the operation mode to read-only and stops collecting new

data. You should monitor this closely to make sure you have sized the store appropriately to contain the full history you'd like to retain.

D: Size Based Cleanup Mode: Specifies whether automatic data cleanup will take place when Query Store data size approaches the limit.

It is strongly recommended to activate size-based cleanup to makes sure that Query Store always runs in read-write mode and collects the latest data.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store>

---



## QUESTION 10

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\_COMMITTED\_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>

---



### QUESTION 11

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>

---

### QUESTION 12

You have a data warehouse fact table that has a clustered columnstore index.

You have multiple CSV files that contain a total of 3 million rows of data.

You need to upload the data to the fact table. The solution must avoid the delta group when you import the data.

Which solution will achieve the goal in the least amount of time?

- A. Load the source data to a staging table. Load the data to the fact table by using the INSERT\_SELECT statement and specify the TABLOCK option on the staging table
- B. Create a Microsoft SQL Server Integration Services (SSIS) package. Use multiple data flow tasks to load the data in parallel.
- C. Load the source data to the fact table by running bcp.exe and specify the ?TABLOCK option



D. Load the source data to the fact table by using the BULK INSERT statement and specify the TABLOCK option

Correct Answer: A

If you are loading data only to stage it before running more transformations, loading the table to heap table will be much faster than loading the data to a clustered columnstore table. In addition, loading data to a [temporary table][Temporary]

will also load much faster than loading a table to permanent storage.

A common pattern for data load is to load the data into a staging table, do some transformation and then load it into the target table using the following command

```
INSERT INTO
```

```
SELECT FROM
```

This command loads the data into the columnstore index in similar ways to BCP or Bulk Insert but in a single batch. If the number of rows in the staging table

loaded into compressed rowgroup. One key limitation was that this INSERT operation was single threaded. To load data in parallel, you could create multiple staging table or issue INSERT/SELECT with non-overlapping ranges of rows from

the staging table. This limitation goes away with SQL Server 2016 (13.x). The command below loads the data from staging table in parallel but you will need to specify TABLOCK.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-data-loading-guidance?view=sql-server-2017#plan-bulk-load-sizes-to-minimize-delta-rowgroups>

---

### QUESTION 13

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. table-valued function
- B. a schema-bound view
- C. a partitioned view
- D. a DML trigger

Correct Answer: D

---

### QUESTION 14



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.

<b>Physical Operation</b>	Index Seek
<b>Logical Operation</b>	Index Seek
<b>Actual Execution Mode</b>	Row
<b>Actual Number of Rows</b>	3571454
<b>Actual Number of Batches</b>	0
<b>Estimated I/O Cost</b>	0.0093577
<b>Estimated Operator Cost</b>	0.0107304 (0%)
<b>Estimated CPU Cost</b>	0.0013727
<b>Estimated Subtree Cost</b>	0.0107304
<b>Estimated Number of Executions</b>	1
<b>Number of Executions</b>	8
<b>Estimated Number of Rows</b>	0
<b>Estimated Row Size</b>	19 B
<b>Actual Rebinds</b>	0
<b>Actual Rewinds</b>	0
<b>Ordered</b>	True
<b>Node ID</b>	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 15

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. DML 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>