

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

You are experiencing performance issues with the database server.

You need to evaluate schema locking issues, plan cache memory pressure points, and backup I/O problems.

What should you create?

- A. a System Monitor report
- B. a `sys.dm_exec_query_stats` dynamic management view query
- C. a `sys.dm_exec_session_wait_stats` dynamic management view query
- D. an Activity Monitor session in Microsoft SQL Management Studio.

Correct Answer: C

`sys.dm_exec_session_wait_stats` returns information about all the waits encountered by threads that executed for each session. You can use this view to diagnose performance issues with the SQL Server session and also with specific queries and batches.

Note: SQL Server wait stats are, at their highest conceptual level, grouped into two broad categories: signal waits and resource waits. A signal wait is accumulated by processes running on SQL Server which are waiting for a CPU to become available (so called because the process has "signaled" that it is ready for processing). A resource wait is accumulated by processes running on SQL Server which are waiting for a specific resource to become available, such as waiting for the release of a lock on a specific record.

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 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. DDL 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>

QUESTION 3

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 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 identify the indexes that are not being used so that you can remove them from the database.

Solution: You run the sys.dm_db_index_physical_stats dynamic management view.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

QUESTION 5

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>

[Latest 70-762 Dumps](#)

[70-762 Practice Test](#)

[70-762 Study Guide](#)