



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.geekcert.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 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 database that is 130 GB and contains 500 million rows of data.

Granular transactions and mass batch data imports change the database frequently throughout the day. Microsoft SQL Server Reporting Services (SSRS) uses the database to generate various reports by using several filters.

You discover that some reports time out before they complete.

You need to reduce the likelihood that the reports will time out.

Solution: You increase the number of log files for the database. You store the log files across multiple disks.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Instead, create a file group for the indexes and a file group for the data files.

QUESTION 2

You are developing an ETL process to cleanse and consolidate incoming data. The ETL process will use a reference table to identify which data must be cleansed in the target table. The server that hosts the table restarts daily.

You need to minimize the amount of time it takes to execute the query and the amount of time it takes to populate the reference table.

What should you do?

A. Convert the target table to a memory-optimized table. Create a natively compiled stored procedure to cleanse and consolidate the data.

B. Convert the reference table to a memory-optimized table. Set the DURABILITY option to SCHEMA_AND_DATA.

C. Create a native compiled stored procedure to implement the ETL process for both tables.

D. Convert the reference table to a memory-optimized table. Set the DURABILITY option to SCHEMA_ONLY.

Correct Answer: D

If you use temporary tables, table variables, or table-valued parameters, consider conversions of them to leverage



memory-optimized tables and table variables to improve performance.

In-Memory OLTP provides the following objects that can be used for memory-optimizing temp tables and table variables: Memory-optimized tables

Durability = SCHEMA_ONLY Memory-optimized table variables

Must be declared in two steps (rather than inline): CREATE TYPE my_type AS TABLE ...; , then DECLARE @mytablevariable my_type;

References: <https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/faster-temp-table-and-table-variable-by-using-memory-optimization?view=sql-server-2017>

QUESTION 3

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 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 DML trigger
- B. a schema-bound view
- C. a stored procedure
- D. a DDL trigger

Correct Answer: A

When you create a view, you must obey the following conditions in order to make the view updateable:

The columns being modified in the view have to directly reference the underlying table columns. As a consequence, the view\\s columns being modified cannot be the result of an aggregate function or a computed column.

The only way to make data changes on a non-updateable view is by using INSTEAD OF DML triggers. This way you can use procedural code to overcome the limitation.

References:

<https://www.mssqltips.com/sqlservertip/5984/sql-server-trigger-on-view-example/>



QUESTION 5

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 the SQL Server instances in real time and optimize the server to maximize throughput, response time, and overall SQL performance.

What should you do?

- A. Create `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: B

`sys.dm_exec_sessions` returns one row per authenticated session on SQL Server. `sys.dm_exec_sessions` is a server-scope view that shows information about all active user connections and internal tasks. This information includes client version, client program name, client login time, login user, current session setting, and more. Use `sys.dm_exec_sessions` to first view the current system load and to identify a session of interest, and then learn more information about that session by using other dynamic management views or dynamic management functions.

Examples of use include finding long-running cursors, and finding idle sessions that have open transactions.

[Latest 70-762 Dumps](#)

[70-762 VCE Dumps](#)

[70-762 Exam Questions](#)