



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.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 use Microsoft SQL Server Profiler to evaluate a query named Query1. The Profiler report indicates the following issues:

At each level of the query plan, a low total number of rows are processed.

The query uses many operations. This results in a high overall cost for the query.

You need to identify the information that will be useful for the optimizer.

What should you do?

- A. Start a SQL Server Profiler trace for the event class Auto Stats in the Performance event category.
- B. Create one Extended Events session with the sqlserver.missing_column_statistics event added.
- C. Start a SQL Server Profiler trace for the event class Soft Warnings in the Errors and Warnings event category.
- D. Create one Extended Events session with the sqlserver.missing_join_predicate event added.

Correct Answer: D

The Missing Join Predicate event class indicates that a query is being executed that has no join predicate. This could result in a long-running query.

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 are creating the following two stored procedures: A natively-compiled stored procedure An interpreted stored procedure that accesses both disk-based and memory-optimized tables

Both stored procedures run within transactions.

You need to ensure that cross-container transactions are possible.

Which setting or option should you use?

- A. the SET TRANSACTION_READ_COMMITTED isolation level for the connection
- B. the SERIALIZABLE table hint on disk-based tables
- C. the SET MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT=ON option for the database
- D. the SET MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT=OFF option for the database

Correct Answer: C

Provide a supported isolation level for the memory-optimized table using a table hint, such as WITH (SNAPSHOT). The need for the WITH (SNAPSHOT) hint can be avoided through the use of the database option MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT. When this option is set to ON, access to a memory-optimized table under a lower isolation level is automatically elevated to SNAPSHOT isolation.

Incorrect Answers:

B: Accessing memory optimized tables using the READ COMMITTED isolation level is supported only for autocommit transactions. It is not supported for explicit or implicit transactions. References: <https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/transactions-with-memory-optimized-tables?view=sql-server-2017>

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_operational_stats dynamic management view.

Does the solution meet the goal?

- A. Yes



B. No

Correct Answer: B

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 database named DB1. The database does not use a memory-optimized filegroup. The database contains a table named Table1. The table must support the following workloads: You need to add the most efficient index to support the new OLTP workload, while not deteriorating the existing Reporting query performance. What should you do?

Workload	Type	Description
Reporting	Existing	The reporting workload must scan most of the records in the table to aggregate on a number of columns. A clustered columnstore index is already created on the table to support this workload.
OLTP	New	The OLTP workload must support 3,000 transactions per second. Rows are identified by using two columns. The filter is variant on one of the two columns while constant on the other. Only a small number of records with a few columns are returned by the query.

- 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: C

A filtered index is an optimized nonclustered index, especially suited to cover queries that select from a well-defined subset of data. It uses a filter predicate to index a portion of rows in the table. A well-designed filtered index can improve query performance, reduce index maintenance costs, and reduce index storage costs compared with full-table indexes.

References:[https://technet.microsoft.com/en-us/library/cc280372\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/cc280372(v=sql.105).aspx)