



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 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 2

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 need to collect data from the following two sources:

1.

The performance counters of the operating system

2.

Microsoft SQL Server events

You must analyze the two datasets side-by side by using a single tool.

Solution: You use dynamic management views and Data Collector Sets (DCs) in Performance Monitor to collect performance data. You use SQL Server Management Studio (SSMS) to analyze the data.

Does this meet the goal?



A. Yes

B. No

Correct Answer: B

Reference: <https://www.sqlshack.com/dba-guide-sql-server-performance-troubleshooting-part-2-monitoring-utilities/>


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

You have a database named DB1. There is no memory-optimized filegroup in the database. You run the following query:

```
DECLARE @LastName varchar(20) = "contoso";  
SELECT P.FirstName, Addr.*  
FROM Person.Person P  
JOIN Person.BusinessEntityAddress ON P.BusinessEntityID =  
BusinessEntityAddress.BusinessEntityID  
JOIN Person.[Address] Addr ON BusinessEntityAddress.AddressID =  
Addr.AddressID  
WHERE P.LastName = @LastName;
```

The following image displays the execution plan the query optimizer generates for this query:



Users frequently run the same query with different values for the local variable @lastName. The table named Person is persisted on disk. You need to create an index on the Person.Person table that meets the following requirements:

1.

All users must be able to benefit from the index.



2.

FirstName must be added to the index as an included column.

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: B

By including nonkey columns, you can create nonclustered indexes that cover more queries. This is because the nonkeycolumns have the following benefits:

They can be data types not allowed as index key columns.

They are not considered by the Database Engine when calculating the number of index key columns or index key size.

QUESTION 4

You have a relational data warehouse that contains 1 TB of data.

You have a stored procedure named usp_sp1 that generated an HTML fragment. The HTML fragment contains color and font style.

You need to return the HTML fragment.

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: G



A SQL Server stored procedure that you can call is one that returns one or more OUT parameters, which are parameters that the stored procedure uses to return data back to the calling application. References: <https://docs.microsoft.com/en-us/sql/connect/jdbc/using-a-stored-procedure-with-output-parameters?view=sql-server-2017>

QUESTION 5

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. The Account table was created using the following Transact-SQL statement:

```
CREATE TABLE Account
(
    AccountNumber int NOT NULL,
    ProductCode char(2) NOT NULL,
    Status tinyint NOT NULL,
    OpenDate date NOT NULL,
    CloseDate date,
    Balance decimal(15,2),
    AvailableBalance decimal(15,2)
);
```

There are more than 1 billion records in the Account table. The Account Number column uniquely identifies each account. The ProductCode column has 100 different values. The values are evenly distributed in the table. Table statistics are refreshed and up to date.

You frequently run the following Transact-SQL SELECT statements:

```
SELECT ProductCode, SUM(Balance) AS TotalSUM FROM Account WHERE ProductCode
<> 'CD' GROUP BY ProductCode;
SELECT AccountNumber, Balance FROM Account WHERE ProductCode = 'CD'
```

You must avoid table scans when you run the queries. You need to create one or more indexes for the table. Solution: You run the following Transact-SQL statement:

```
CREATE CLUSTERED INDEX PK_Account On Account(AccountNumber);
CREATE NONCLUSTERED INDEX IX_Account_ProductCode On Account(ProductCode)
INCLUDE (Balance);
```

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: A

Create a clustered index on the AccountNumber column as it is unique. Create a nonclustered index that includes the ProductCode column.



VCE & PDF

PassApply.com

<https://www.passapply.com/70-762.html>

2024 Latest passapply 70-762 PDF and VCE dumps Download

References: <https://msdn.microsoft.com/en-us/library/ms190457.aspx>

[Latest 70-762 Dumps](#)

[70-762 PDF Dumps](#)

[70-762 Practice Test](#)