



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.pass4itsure.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 manage a database that supports an Internet of Things (IoT) solution. The database records metrics from over 100 million devices every minute. The database requires 99.995% uptime.

The database uses a table named Checkins that is 100 gigabytes (GB) in size. The Checkins table stores metrics from the devices. The database also has a table named Archive that stores four terabytes (TB) of data. You use stored

procedures for all access to the tables.

You observe that the wait type PAGELATCH_IO causes large amounts of blocking.

You need to resolve the blocking issues while minimizing downtime for the database.

Which two actions should you perform? Each correct answer presents part of the solution.

- A. Convert all stored procedures that access the Checkins table to natively compiled procedures.
- B. Convert the Checkins table to an In-Memory OLTP table.
- C. Convert all tables to clustered columnstore indexes.
- D. Convert the Checkins table to a clustered columnstore index.

Correct Answer: AB

Natively compiled stored procedures are Transact-SQL stored procedures compiled to native code that access memory-optimized tables. Natively compiled stored procedures allow for efficient execution of the queries and business logic in the stored procedure.

SQL Server In-Memory OLTP helps improve performance of OLTP applications through efficient, memory-optimized data access, native compilation of business logic, and lock- and latch free algorithms. The In-Memory OLTP feature includes memory-optimized tables and table types, as well as native compilation of Transact-SQL stored procedures for efficient access to these tables.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/natively-compiled-stored-procedures>

<https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/memory-optimized-tables>

QUESTION 2

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 check constraint on the Employees table.

Does the solution meet the goal?

A.

Yes

B.

No

Correct Answer: B

Check constraints cannot be used to track changes in a table. References: <https://msdn.microsoft.com/en-us/library/bb933994.aspx>

QUESTION 3

You have the following stored procedure that is called by other stored procedures and applications:

```
CREATE PROCEDURE UpdateCustomer @CustomerId INT
AS
BEGIN
    EXEC ProcessCustomer_Internal @CustomerId
    DECLARE @Status INT
    SELECT @Status = Status FROM Customer WHERE CustomerId =
@CustomerId
END
```

You need to modify the stored procedure to meet the following requirements:

Always return a value to the caller.

Return 0 if @Status is NULL.

Callers must be able to use @Status as a variable.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Replace NULL values with 0. Add a PRINT statement to return @Status.

B. Add a RETURN statement.

C. Replace NULL values with 0. Add an output parameter to return @Status.

D. Replace NULL values with 0. Add a SELECT statement to return @Status.

E. Add a PRINT statement.



F. Add a SELECT statement to return @Status.

G. Add an output parameter to return @Status.

Correct Answer: BC

There are three ways of returning data from a procedure to a calling program: result sets, output parameters, and return codes.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/stored-procedures/return-data-from-a-stored-procedure?view=sql-server-2017>

QUESTION 4

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 5

You have a database that users query frequently.

The users report that during peak business hours, the queries take longer than expected to execute.

A junior database administrator uses Microsoft SQL Server Profiler on the database server to trace the session activities.

While performing the trace, the performance of the database server worsens, and the server crashes.

You need to recommend a solution to collect the query run times. The solution must minimize the impact on the resources of the database server.

What should you recommend?

- A. Increase the free space on the system drive of the database server, and then use SQL Server Profiler on the server to trace the session activities.
- B. Collect session activity data by using SQL Server Extended Events.
- C. Clean up tempdb, and then use SQL Server Profiler on the database server to trace the session activities.
- D. Collect performance data by using a Data Collector Set (DCS) in Performance Monitor.

Correct Answer: B

SQL Server Extended Events has a highly scalable and highly configurable architecture that allows users to collect as much or as little information as is necessary to troubleshoot or identify a performance problem. Extended Events is a light weight performance monitoring system that uses very few performance resources. Extended Events provides two graphical user interfaces (New Session Wizard and New Session) to create, modify, display, and analyze your session data.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/extended-events/extended-events?view=sql-server-2017>

[70-762 VCE Dumps](#)

[70-762 Exam Questions](#)

[70-762 Braindumps](#)