



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

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 and optimize the SQL Server to maximize throughput, response time, and overall SQL performance.

You need to identify previous situations where a modification has prevented queries from selecting data in tables.

What should you do?

- A. Create a `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: G

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

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 stores information about the employees of your company.

You need to implement and enforce the following business rules:

Limit the values that are accepted by the Salary column.

Prevent salaries less than \$15,000 and greater than \$300,000 from being entered.



Determine valid values by using logical expressions.

Do not validate data integrity when running DELETE statements.

Solution: You implement a FOR UPDATE trigger on the table.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

References: <http://stackoverflow.com/questions/16081582/difference-between-for-update-of-and-for-update>

QUESTION 3

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 user-defined function on the Employees table.

Does the solution meet the goal?

A.

Yes

B.

No

Correct Answer: A

SQL Server 2016 provides two features that track changes to data in a database: change data capture and change tracking. These features enable applications to determine the DML changes (insert, update, and delete operations) that were

made to user tables in a database.

Change data is made available to change data capture consumers through table-valued functions (TVFs).

References: <https://msdn.microsoft.com/en-us/library/cc645858.aspx>



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 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 create a file group for the indexes and a file group for the data files. You store the files for each file group on separate disks.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

Consider creating two additional File Groups: Tables and Indexes. It is best not to put your stuff in PRIMARY as that is where SQL SERVER stores all of its data and meta-data about your objects. You create your Table and Clustered Index (as that is the data for the table) on [Tables] and all Non-Clustered indexes on [Indexes].

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 PDF Dumps](#)

[70-762 Study Guide](#)