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

SATISFACTION GUARANTEED
100%
SATISFACTION GUARANTEED

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 observe that many deadlocks appear to be happening during specific times of the day.

You need to monitor the SQL environment and capture the information about the processes that are causing the deadlocks. Captured information must be viewable as the queries are running.

What should you do?

A. A. Create a sys.dm_os_waiting_tasks query.

B. Create a sys.dm_exec_sessions query.

C. Create a PerformanceMonitor 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: F

To view deadlock information, the Database Engine provides monitoring tools in the form of two trace flags, and the deadlock graph event in SQL Server Profiler.

Trace Flag 1204 and Trace Flag 1222 When deadlocks occur, trace flag 1204 and trace flag 1222 return information that is captured in the SQL Server error log. Trace flag 1204 reports deadlock information formatted by each nodeinvolved in the deadlock. Trace flag 1222 formats deadlock information, first by processes and then by resources. It is possible to enable bothtrace flags to obtain two representations of the same deadlock event.

References:https://technet.microsoft.com/en-us/library/ms178104(v=sql.105).aspx

QUESTION 2

You have the following stored procedure:

```
CREATE PROCEDURE AddNextNumber @Number INT
AS
BEGIN
    SET ANSI_DEFAULTS ON
    INSERT INTO Numbers (Number) VALUES (@Number)
END
```

The Numbers table becomes unavailable when you run the stored procedure. The stored procedure obtains an exclusive lock on the table and does not release the lock.

What are two possible ways to resolve the issue? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Remove the implicit transaction and the SET ANSI_DEFAULTS ON statement.

B. Set the ANSI_DEFAULT statement to OFF and add a COMMIT TRANSACTION statement after the INSERT statement.

C. Add a COMMIT TRANSACTION statement after the INSERT statement.

D. Remove the SET ANSI_DEFAULTS ON statement.

Correct Answer: CD

SET ANSI_DEFAULTS is a server-side setting that the client does not modify. When enabled (ON), this option enables SET IMPLICIT_TRANSACTIONS (and some other options).

The SET IMPLICIT_TRANSACTIONS, when ON, the system is in implicit transaction mode.

This means that if @@TRANCOUNT = 0, any of the following Transact-SQL statements begins a new transaction. It is equivalent to an unseen BEGIN TRANSACTION being executed first: ALTER TABLE, FETCH, REVOKE, BEGIN

TRANSACTION, GRANT, SELECT, CREATE, INSERT, TRUNCATE TABLE, DELETE, OPEN, UPDATE, DROP.

References: https://docs.microsoft.com/en-us/sql/t-sql/statements/set-implicit-transactions-transact-sql?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

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 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.

You need to collect query performance data while minimizing the performance impact on the SQL Server.

What should you do?

A. 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: C

SQL Server Data Collector is a feature for performance monitoring and tuning available in SQL Server Management Studio.

Integration Services packages transform and load the collected data into the Microsoft Data Warehouse database.

Collection sets are defined and deployed on a server instance and can be run independently of each other. Each collection set can be applied to a target that matches the target types of all the collector types that are part of a collection set.

The collection set is run by a SQL Server Agent job or jobs, and data is uploaded to the management data warehouse on a predefined schedule.

Predefined data collection sets include:

The Query Statistics data collection set collects information about query statistics, activity, execution plans and text on the SQL Server instance. It does not store all executed statements, only 10 worst performing ones.

Disk Usage data collection set collects information about disk space used by both data and log files for all databases on the SQL Server instance, growth trends, and average day growth.

Etc.

References:

http://www.sqlshack.com/sql-server-performance-monitoring-data-collector/

QUESTION 5

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/

70-762 Practice Test   70-762 Study Guide   70-762 Braindumps