



Microsoft

Exam Questions 70-765

Provisioning SQL Databases (beta)

NEW QUESTION 1

HOTSPOT - (Topic 1)

You plan to migrate a Microsoft SQL Server workload from an on-premises server to a Microsoft Azure virtual machine (VM). The current server contains 4 cores with an average CPU workload of 6 percent and a peak workload of 10 percent when using 2.4Ghz processors.

You gather the following metrics:

	Minimum IOPS	Average IOPS	Maximum IOPS
Data Drive	100	938	7253
Transaction Log Drive	12	145	350
TempDB Drive	300	900	1900

You need to design a SQL Server VM to support the migration while minimizing costs.

For each setting, which value should you use? To answer, select the appropriate storage option from each list in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

VM setting	Value ▼
Data drive	Local storage Premium storage Standard storage
Transaction log drive	Local storage Premium storage Standard storage
TempDB drive	Local storage Premium storage Standard storage
VM size	A3 D3 DS3

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Data drive: Premium Storage Transaction log drive: Standard Storage TempDB drive: Premium Storage

Note: A standard disk is expected to handle 500 IOPS or 60MB/s. A P10 Premium disk is expected to handle 500 IOPS.

A P20 Premium disk is expected to handle 2300 IOPS. A P30 Premium disk is expected to handle 5000 IOPS.

VM size: A3

Max data disk throughput is 8x500 IOPS

References: <https://docs.microsoft.com/en-us/azure/virtual-machines/virtual-machines-windows-sizes>

NEW QUESTION 2

HOTSPOT - (Topic 1)

You use Resource Manager to deploy a new Microsoft SQL Server instance in a Microsoft Azure virtual machine (VM) that uses Premium storage. The combined initial size of the SQL Server user database files is expected to be over 200 gigabytes (GB). You must maximize performance for the database files and the log file.

You add the following additional drive volumes to the VM:

Drive volume	Storage	Host caching
E:	Premium storage	ReadOnly
F:	Premium storage	None

You have the following requirements:

You need to deploy the SQL instance.

In the table below, identify the drive where you must store each SQL Server file type. NOTE: Make only one selection in each column. Each correct selection is worth one point.

Answer area

Drive	Data files	Log files
C:	<input type="radio"/>	<input type="radio"/>
D:	<input type="radio"/>	<input type="radio"/>
E:	<input type="radio"/>	<input type="radio"/>
F:	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Enable read caching on the disk(s) hosting the data files and TempDB.
 Do not enable caching on disk(s) hosting the log file. Host caching is not used for log files.

NEW QUESTION 3

- (Topic 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 stated goals.

You manage a Microsoft SQL Server environment with several databases.

You need to ensure that queries use statistical data and do not initialize values for local variables.

Solution: You enable the PARAMETER_SNIFFING option for the databases. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

PARAMETER_SNIFFING = { ON | OFF | PRIMARY} enables or disables parameter sniffing. This is equivalent to Trace Flag 4136.

SQL server uses a process called parameter sniffing when executing queries or stored procedures that use parameters. During compilation, the value passed into the parameter is evaluated and used to create an execution plan. That value is also stored with the execution plan in the plan cache. Future executions of the plan will re-use the plan that was compiled with that reference value.

References:<https://msdn.microsoft.com/en-us/library/mt629158.aspx>

NEW QUESTION 4

HOTSPOT - (Topic 2)

You need to ensure that a user named Admin2 can manage logins.

How should you complete the Transact-SQL statements? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

▼	Admin2 WITH password = 'Pa\$\$w0rd';	
CREATE USER		
ALTER SERVER ROLE		
CREATE LOGIN		

▼	Admin2User FROM	▼	Admin2
CREATE USER		WINDOWS	
ALTER SERVER ROLE		EXTERNAL PROVIDER	
CREATE LOGIN		LOGIN	

ALTER ROLE ' loginmanager	▼
dbmanager	
bd_ddladmin	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: CREATE LOGIN

First you need to create a login for SQL Azure, it's syntax is as follows: CREATE LOGIN username WITH password='password';

Step 2, CREATE USER Step 3: LOGIN

Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a user:

CREATE USER readonlyuser FROM LOGIN readonlylogin; Step 4: loginmanager

Members of the loginmanager role can create new logins in the master database.

References:

<https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/> <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-manage-logins>

NEW QUESTION 5

- (Topic 3)

A company has an on-premises Microsoft SQL Server 2014 environment. The company has a main office in Seattle, and remote offices in Amsterdam and Tokyo. You plan to deploy a Microsoft Azure SQL Database instance to support a new application. You expect to have 100 users from each office.

In the past, users at remote sites reported issues when they used applications hosted at the Seattle office.

You need to optimize performance for users running reports while minimizing costs. What should you do?

- A. Implement an elastic pool.
- B. Implement a standard database with readable secondaries in Asia and Europe, and then migrate the application.
- C. Implement replication from an on-premises SQL Server database to the Azure SQL Database instance.
- D. Deploy a database from the Premium service tier.

Answer: B

Explanation:

References:<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-geo-replication-transact-sql#add-secondary-database>

NEW QUESTION 6

- (Topic 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 stated goals.

You have a mission-critical application that stores data in a Microsoft SQL Server instance. The application runs several financial reports. The reports use a SQL Server-authenticated login named Reporting_User. All queries that write data to the database use Windows authentication.

Users report that the queries used to provide data for the financial reports take a long time to complete. The queries consume the majority of CPU and memory resources on the database server. As a result, read-write queries for the application also take a long time to complete.

You need to improve performance of the application while still allowing the report queries to finish.

Solution: You create a snapshot of the database. You configure all report queries to use the database snapshot.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Use a Resource Governor instead.

References: <https://msdn.microsoft.com/en-us/library/bb933866.aspx>

NEW QUESTION 7

HOTSPOT - (Topic 4)

You need to optimize SRV1.

What configuration changes should you implement? To answer, select the appropriate option from each list in the answer area.

Answer Area

How should you modify the tempdb configuration?

- Change the recovery model of tempdb.
- Change the number of tempdb files.
- Change the size of the tempdb log file.
- Change the MAXDOP property.

How should you reconfigure the tempdb database?

- Add additional tempdb files.
- Remove tempdb files.
- Add tempdb log files.
- Remove tempdb log files.
- Set MAXDOP to 8.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

From the scenario: SRV1 has 16 logical cores and hosts a SQL Server instance that supports a mission-critical application. The application has approximately 30,000 concurrent users and relies heavily on the use of temporary tables.

Box 1: Change the size of the tempdb log file.

The size and physical placement of the tempdb database can affect the performance of a system. For example, if the size that is defined for tempdb is too small, part of the system- processing load may be taken up with autogrowing tempdb to the size required to support the workload every time you restart the instance of SQL Server. You can avoid this overhead by increasing the sizes of the tempdb data and log file.

Box 2: Add additional tempdb files.

Create as many files as needed to maximize disk bandwidth. Using multiple files reduces tempdb storage contention and yields significantly better scalability. However, do not create too many files because this can reduce performance and increase management overhead. As a general guideline, create one data file for each CPU on the server (accounting for any affinity mask settings) and then adjust the number of files up or down as necessary.

NEW QUESTION 8

HOTSPOT - (Topic 5)

You need to configure the data entry and business intelligence databases. In the table below, identify the option that you must use for each database. NOTE: Make only one selection in each column.

Answer Area

Option	Data entry	Business intelligence
Elastic database pools only	<input type="radio"/>	<input type="radio"/>
Geo-replicated database only	<input type="radio"/>	<input type="radio"/>
Elastic database pools and geo-replicated databases	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Data Entry: Geo-replicated database only

From Contoso scenario: Each location database for the data entry application may have an unpredictable amount of activity. Data must be replicated to secondary databases in Azure datacenters in different regions.

Business intelligence: Elastic database pools only

From Contoso scenario: For the business intelligence application, corporate executives must be able to view all data in near real-time with low network latency. SQL DB elastic pools provide a simple cost effective solution to manage the performance goals for multiple databases that have widely varying and unpredictable usage patterns.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool>

NEW QUESTION 9

DRAG DROP - (Topic 6)

You create a login named BIAppUser. The login must be able to access the Reporting database.

You need to grant access to the BIAppUser login in the database.

How should you complete the Transact-SQL statements? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Code segments	Answer area
Reporting	USE [Code segment]
master	GO
CREATE USER	Code segment [BIAppUser] Code segment
ALTER LOGIN	GO
ALTER USER	
FOR LOGIN [BIAppUser]	
FOR USER [BIAppUser]	
WITH LOGIN = [BIAppUser]	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Reporting

The user is to be created in the Reporting database.

Box 2: CREATE USER

Box 3: FOR LOGIN [BIAppUser]

Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. Here is some sample Transact-SQL that creates a user:

CREATE USER readonlyuser FROM LOGIN readonlylogin;

References: <https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/>

NEW QUESTION 10

- (Exam Topic 7)

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 question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are tuning the performance of a virtual machines that hosts a Microsoft SQL Server instance. The virtual machine originally had four CPU cores and now has 32 CPU cores.

The SQL Server instance uses the default settings and has an OLTP database named db1. The largest table in db1 is a key value store table named table1.

Several reports use the PIVOT statement and access more than 100 million rows in table1.

You discover that when the reports run, there are PAGELATCH_IO waits on PFS pages 2:1:1, 2:2:1, 2:3:1, and 2:4:1 within the tempdb database.

You need to prevent the PAGELATCH_IO waits from occurring. Solution: You add more tempdb databases.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

From SQL Server's perspective, you can measure the I/O latency from sys.dm_os_wait_stats. If you consistently see high waiting for PAGELATCH_IO, you can benefit from a faster I/O subsystem for SQL Server. A cause can be poor design of your database - you may wish to split out data located on 'hot pages', which are accessed frequently and which you might identify as the causes of your latch contention. For example, if you have a currency table with a data page containing 100 rows, of which 1 is updated per transaction and you have a transaction rate of 200/sec, you could see page latch queues of 100 or more. If each page latch wait costs just 5ms before clearing, this represents a full half-second delay for each update. In this case, splitting out the currency rows into different tables might prove more performant (if less normalized and logically structured).

References: <https://www.mssqltips.com/sqlservertip/3088/Explanation:-of-sql-server-io-and-latches/>

NEW QUESTION 10

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database.

You need to ensure that the size of the transaction log file does not exceed 2 GB. What should you do?

- A. Execute sp_configure 'max log size', 2G.
- B. use the ALTER DATABASE...SET LOGFILE command along with the maxsize parameter.
- C. In SQL Server Management Studio, right-click the instance and select Database Setting
- D. Set the maximum size of the file for the transaction log.
- E. in SQL Server Management Studio, right-click the database, select Properties, and then click Files. Open the Transaction log Autogrowth window and set the maximum size of the file.

Answer: B

Explanation:

You can use the ALTER DATABASE (Transact-SQL) statement to manage the growth of a transaction log file

To control the maximum the size of a log file in KB, MB, GB, and TB units or to set growth to UNLIMITED, use the MAXSIZE option. However, there is no SET LOGFILE subcommand.

References: [https://technet.microsoft.com/en-us/library/ms365418\(v=sql.110\).aspx#ControlGrowth](https://technet.microsoft.com/en-us/library/ms365418(v=sql.110).aspx#ControlGrowth)

NEW QUESTION 14

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 failover cluster that contains two nodes named Node A and Node B. A single instance of SQL Server is installed on the cluster.

An additional node named Node C has been added to the existing cluster.

You need to ensure that the SQL Server instance can use all nodes of the cluster. What should you do?

- A. Run the New SQL Server stand-alone installation Wizard on Node C.
- B. Run the Add Node to SQL Server Failover Cluster Wizard on Node C.
- C. Use Node B to install SQL Server on Node C.
- D. Use Node A to install SQL Server on Node C.

Answer: B

Explanation:

To add a node to an existing SQL Server failover cluster, you must run SQL Server Setup on the node that is to be added to the SQL Server failover cluster instance. Do not run Setup on the active node.

The Installation Wizard will launch the SQL Server Installation Center. To add a node to an existing failover cluster instance, click Installation in the left-hand pane. Then, select Add node to a SQL Server failover cluster.

References:

<http://technet.microsoft.com/en-us/library/ms191545.aspx>

NEW QUESTION 17

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database.

You have a SQL Server Agent job instance that runs using the service account. You have a job step within the job that requires elevated privileges.

You need to ensure that the job step can run using a different user account. What should you use?

- A. a schedule
- B. an alert
- C. an operator
- D. a proxy

Answer: D

Explanation:

A SQL Server Agent proxy defines the security context for a job step. A proxy provides SQL Server Agent with access to the security credentials for a Microsoft Windows user. Each proxy can be associated with one or more subsystems. A job step that uses the proxy can access the specified subsystems by using the security context of the Windows user. Before SQL Server Agent runs a job step that uses a proxy, SQL Server Agent impersonates the credentials defined in the proxy, and then runs the job step by using that security context.

References: [https://technet.microsoft.com/en-us/library/ms189064\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms189064(v=sql.105).aspx)

NEW QUESTION 20

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 server. One of the databases on the server supports a highly active OLTP application.

Users report abnormally long wait times when they submit data into the application.

You need to identify which queries are taking longer than 1 second to run over an extended period of time. What should you do?

- A. use SQL Profiler to trace all queries that are processing on the server
- B. Filter queries that have a Duration value of more than 1,000.
- C. Use sp_configure to set a value for blocked process threshold
- D. Create an extended event session.
- E. Use the Job Activity monitor to review all processes that are actively running
- F. Review the Job History to find out the duration of each step.
- G. Run the sp_who command from a query window.
- H. Run the DBCC TRACEON 1222 command from a query window and review the SQL Server event log.

Answer: A

NEW QUESTION 21

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database.

The database contains a Product table created by using the following definition:

```
CREATE TABLE dbo.Product
(ProductID INT PRIMARY KEY,
Name VARCHAR(50) NOT NULL,
Color VARCHAR(15) NOT NULL,
Size VARCHAR(5) NOT NULL,
Style CHAR(2) NULL,
Weight DECIMAL(8,2) NULL);
```

You need to ensure that the minimum amount of disk space is used to store the data in the Product table. What should you do?

- A. Convert all indexes to Column Store indexes.
- B. Implement Unicode Compression.
- C. Implement row-level compression.
- D. Implement page-level compression.

Answer: D

NEW QUESTION 22

- (Exam Topic 7)

You are the administrator for a SQL Server 2016 instance that stores the data for an online transaction processing sales system. The company takes full backups every week; differential backups on the days with no full backups; and hourly transaction backups.

These backups are stored on a backup server in the company's data center.

Every week, the company places the full backup on a tape and sends it to a third-party backup storage system. The company is worried that a disaster might occur that could destroy their computer center and cause them to lose orders.

You need to determine the best method for providing the smallest amount of data loss and downtime without leasing or purchasing additional physical locations.

What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Set up SQL Server Always On with a SQL Azure database as a replica.
- B. Set up SQL Server Always On by using a SQL Server on a Windows Azure Virtual Machine.
- C. Put the differential backup on tape and send it to the third-party backup storage system.
- D. Use the Microsoft SQL Server Backup to Microsoft Windows Azure Tool to direct all backups to a different geographical location.

Answer: D

Explanation:

Microsoft SQL Server Backup to Microsoft Azure Tool enables backup to Azure Blob Storage and encrypts and compresses SQL Server backups stored locally or in the cloud.

References: <https://www.microsoft.com/en-us/download/details.aspx?id=40740>

NEW QUESTION 23

- (Exam Topic 7)

You have an on-premises Microsoft SQL server that has a database named DB1. DB1 contains several tables that are stretched to Microsoft Azure.

From SQL Server Management Studio (SSMS), a junior database administrator accidentally deletes several rows from the Azure SQL database and breaks the connection to Azure.

You need to resume Stretch Database operations.

Which two stored procedures should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. sys.sp_rda_reconcile_batch
- B. sys.sp_rda_reconcile_indexes
- C. sys.sp_rda_reauthorize_db
- D. sys.sp_rda_reconcile_columns
- E. sys.sp_rda_set_rpo_duration

Answer: CD

Explanation:

sys.sp_rda_reauthorize_db restores the authenticated connection between a local database enabled for Stretch and the remote database.

If you have accidentally deleted columns from the remote table, run sp_rda_reconcile_columns to add columns to the remote table that exist in the Stretch-enabled SQL Server table but not in the remote table.

NEW QUESTION 27

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 Enterprise Edition server that uses 64 cores.

You discover performance issues when large amounts of data are written to tables under heavy system load. You need to limit the number of cores that handle I/O.

What should you configure?

- A. Processor affinity
- B. Lightweight pooling
- C. Max worker threads
- D. I/O affinity

Answer: D

Explanation:

The affinity Input-Output (I/O) mask Server Configuration Option.

To carry out multitasking, Microsoft Windows 2000 and Windows Server 2003 sometimes move process threads among different processors. Although efficient from an operating system point of view, this activity can reduce Microsoft SQL Server performance under heavy system loads, as each processor cache is repeatedly reloaded with data. Assigning processors to specific threads can improve performance under these conditions by eliminating processor reloads; such an association between a thread and a processor is called processor affinity.

References:

<http://msdn.microsoft.com/en-us/library/ms189629.aspx>

NEW QUESTION 32

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft azure virtual machine that has 12 databases. All database files are in the same Azure Blob storage account.

You need to receive an email notification if I/O operations to the database files exceed 800 MB/s for more than five minutes.

Solution: You run the Get-Counter cmdlet and specify the -counter '\physicaldisk:disk Transfers/sec' parameter.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 37

- (Exam Topic 7)

You have an on-premises database that runs several maintenance jobs. You move the database to a Microsoft Azure SQL database.

You need to ensure that the maintenance jobs on indexes continue to run after the move.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions, Select from these

Answer Area, Place here

Create a runbook

1.

Create an Automation Account

2.

Configure a schedule

3.

Create a credential

4.

Publish a runbook

5.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

General steps for a solution to automate the maintenance you Azure SQL DB statistics: References:

<https://blogs.msdn.microsoft.com/azuresqldb/2018/01/15/automating-azure-sql-db-index-and-statistic-m>

NEW QUESTION 40

- (Exam Topic 7)

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 question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have Microsoft SQL Server on a Microsoft Azure virtual machine that has a database named DB1. You discover that DB1 experiences WRITE_LOG waits that are longer than 50 ms.

You need to reduce the WRITE_LOG wait time. Solution: Add additional data files to DB1. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

In SQL Server, if we have a transactional based system and find a high WRITELOG wait type this is a performance bottleneck and can cause the transaction log file to grow rapidly and frequently.

It is being recommended to SQL server users that they must archive the log files on a separate disk for getting better performance.

References: <https://atdhebuja.wordpress.com/2016/06/20/resolving-sql-server-transaction-log-waits/>

NEW QUESTION 44

- (Exam Topic 7)

You have a Microsoft Azure SQL database in the US West region. You need to create a replica in the US East region.

Which cmdlet should you run first?

- A. New-AzureRmAvailabilitySet
- B. New-AzureRmLoadBalancer

- C. New-AzureRmSqlDatabaseSecondary
- D. New-AzureRmSqlElasticPool
- E. New-AzureRmVM
- F. New-AzureRmSqlServer
- G. New-AzureRmSqlDatabaseCopy
- H. New-AzureRmSqlServerCommunicationLink

Answer: G

Explanation:

The New-AzureRmSqlDatabaseCopy command creates a copy of a SQL Database that uses the snapshot at the current time.

References:

<https://docs.microsoft.com/en-us/powershell/module/azurearm.sql/new-azurermsqldatabasecopy?view=azurermp>

NEW QUESTION 49

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database that includes a table named Application.Events. Application.Events contains millions of records about user activity in an application.

Records in Application.Events that are more than 90 days old are purged nightly. When records are purged, table locks are causing contention with inserts.

You need to be able to modify Application.Events without requiring any changes to the applications that utilize Application.Events.

Which type of solution should you use?

- A. Partitioned tables
- B. Online index rebuild
- C. Change data capture
- D. Change tracking

Answer: A

Explanation:

Partitioning large tables or indexes can have manageability and performance benefits including:

You can perform maintenance operations on one or more partitions more quickly. The operations are more efficient because they target only these data subsets, instead of the whole table.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/partitions/partitioned-tables-and-indexes>

NEW QUESTION 52

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance that contains a financial database hosted on a storage area network (SAN).

The financial database has the following characteristics:

The database is continually modified by users during business hours from Monday through Friday between 09:00 hours and 17:00 hours. Five percent of the existing data is modified each day.

The Finance department loads large CSV files into a number of tables each business day at 11:15 hours and 15:15 hours by using the BCP or BULK INSERT commands. Each data load adds 3 GB of data to the database.

These data load operations must occur in the minimum amount of time.

A full database backup is performed every Sunday at 10:00 hours. Backup operations will be performed every two hours (11:00, 13:00, 15:00, and 17:00) during business hours.

You need to ensure that your backup will continue if any invalid checksum is encountered. Which backup option should you use?

- A. STANDBY
- B. Differential
- C. FULL
- D. CHECKSUM
- E. BULK_LOGGED
- F. CONTINUE_AFTER_ERROR
- G. SIMPLE
- H. DBO_ONLY
- I. COPY_ONLY
- J. SKIP
- K. RESTART
- L. Transaction log
- M. NO_CHECKSUM
- N. NORECOVERY

Answer: F

Explanation:

The CONTINUE_AFTER_ERROR option, of the Transact-SQL BACKUP command, instructs BACKUP to continue despite encountering errors such as invalid checksums or torn pages.

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/backup-transact-sql>

NEW QUESTION 57

- (Exam Topic 7)

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 question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are tuning the performance of a virtual machines that hosts a Microsoft SQL Server instance. The virtual machine originally had four CPU cores and now has 32 CPU cores.

The SQL Server instance uses the default settings and has an OLTP database named db1. The largest table in db1 is a key value store table named table1.

Several reports use the PIVOT statement and access more than 100 million rows in table1. You discover that when the reports run, there are PAGELATCH_IO waits on PFS pages 2:1:1, 2:2:1, 2:3:1, and 2:4:1 within the tempdb database. You need to prevent the PAGELATCH_IO waits from occurring. Solution: You add more files to db1. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

From SQL Server's perspective, you can measure the I/O latency from sys.dm_os_wait_stats. If you consistently see high waiting for PAGELATCH_IO, you can benefit from a faster I/O subsystem for SQL Server.

A cause can be poor design of your database - you may wish to split out data located on 'hot pages', which are accessed frequently and which you might identify as the causes of your latch contention. For example, if you have a currency table with a data page containing 100 rows, of which 1 is updated per transaction and you have a transaction rate of 200/sec, you could see page latch queues of 100 or more. If each page latch wait costs just 5ms before clearing, this represents a full half-second delay for each update. In this case, splitting out the currency rows into different tables might prove more performant (if less normalized and logically structured).

References: <https://www.mssqltips.com/sqlservertip/3088/Explanation:-of-sql-server-io-and-latches/>

NEW QUESTION 62

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database named Contoso on a server named Server01. You need to be notified immediately when fatal errors occur on Server01.

What should you create?

- A. A Database Audit Specification
- B. A Policy
- C. An Alert
- D. A SQL Profiler Trace
- E. A Resource Pool
- F. An Extended Event session
- G. A Server Audit Specification

Answer: C

Explanation:

SQL Server has alerts that get more important based on the severity of the alert. Anything of severity 16 or below tends to refer to the database and deals with issues that are tied to syntax errors, violations of foreign keys, etc. While those errors are typically important, they don't refer to anything with regards to overall health of the SQL Server. Alerts 17 through 25 do. Those are the ones your health checks are probably firing on.

Severity Level	Meaning
17	Insufficient Resources
18	Nonfatal Internal Error Detected
19	SQL Server Error in Resource
20	SQL Server Fatal Error in Current Process
21	SQL Server Fatal Error in Database (dbid) Process
22	SQL Server Fatal Error Table Integrity Suspect
23	SQL Server Fatal Error: Database Integrity Suspect
24	Hardware Error
25	(no description)

References: <https://www.mssqltips.com/sqlservertip/3384/configuring-critical-sql-server-alerts/>

NEW QUESTION 64

- (Exam Topic 7)

You have an on-premises Microsoft SQL server that has a database named DB1. DB1 contains several tables that are stretched to Microsoft Azure.

A network administrator upgrades the hardware firewalls on the network. You need to verify whether data migration still runs successfully.

Which stored procedure should you run?

- A. Sys_sp_testlinkedserver
- B. Sys_sp_rda_test_connection
- C. Sys_sp_rda_reauthorized_db
- D. Sp_set_firewall_rule

Answer: B

Explanation:

The Sys_sp_rda_test_connection cmdlet tests the connection from SQL Server to the remote Azure server and reports problems that may prevent data migration.

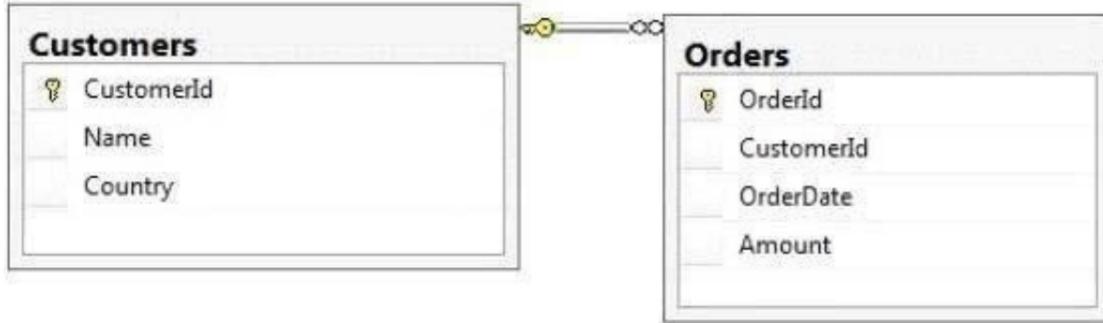
References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-rda-test-connection-tr>

NEW QUESTION 69

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database named ContosoDb. Tables are defined as shown in the exhibit. (Click the Exhibit button.)



You need to display rows from the Orders table for the Customers row having the CustomerId value set to 1 in the following XML format.

```

<Customers Name="Customer A" Country="Australia">
  <OrderId>1</OrderId>
  <OrderDate>2000-01-01T00:00:00</OrderDate>
  <Amount>3400.00</Amount>
</Customers>
<Customers Name="Customer A" Country="Australia">
  <OrderId>2</OrderId>
  <OrderDate>2001-01-01T00:00:00</OrderDate>
  <Amount>4300.00</Amount>
</Customers>
  
```

Which Transact-SQL query should you use?

- A. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML RAW
- B. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML RAW, ELEMENTS
- C. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML AUTO
- D. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId=Customers.CustomerId WHERE Customers.CustomerId= 1 FOR XML AUTO, ELEMENTS
- E. SELECT Name, Country, OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId= Customers.CustomerId WHERE Customers.CustomerId= 1 FOR XML AUTO, ELEMENTS
- F. SELECT Name, Country, CrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId= Customers.CustomerId WHERE Customers.CustomerId= FOR XML AUTO, ELEMENTS
- G. SELECT Name AS '@Name', Country AS '@Country', OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId= Customers.CustomerId WHERE Customers.CustomerId= 1 FOR XML PATH ('Customers')
- H. SELECT Name AS 'Customers/Name', Country AS 'Customers/Country', OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId= Customers.CustomerId WHERE Customers.CustomerId= 1 FOR XML PATH ('Customers')

Answer: E

NEW QUESTION 72

- (Exam Topic 7)

You manage an on-premises Microsoft SQL server that has a database named DB1. An application named App1 retrieves customer information for DB1. Users report that App1 takes an unacceptably long time to retrieve customer records. You need to find queries that take longer than 400 ms to run. Which statement should you execute?

A)

```

SELECT      qp.query_plan,
            qs.*
FROM        (
            SELECT TOP 50 *
            FROM sys.dm_exec_query_stats
            ORDER BY total_worker_time DESC
            ) AS qs
CROSS APPLY sys.dm_exec_query_plan(qs.plan_handle) AS qp
WHERE (qs.max_worker_time > 400
       OR qs.max_elapsed_time > 400)
  
```

B)

```
SELECT pa.DatabaseID, SUM(qs.total_worker_time/100) AS [CPU_Time_Ms]
FROM sys.dm_exec_query_stats AS qs
CROSS APPLY (SELECT CONVERT(int, value) AS [DatabaseID]
FROM sys.dm_exec_plan_attributes(qs.plan_handle)
WHERE attribute = N'dbid') AS pa
GROUP BY pa.DatabaseID
HAVING SUM(qs.total_worker_time/1000) > 400
ORDER BY 2 DESC
```

C)

```
SELECT qp.query_plan,
qs.*
FROM (
SELECT TOP 50 *
FROM sys.dm_exec_query_stats
ORDER BY total_worker_time DESC
) AS qs
CROSS APPLY sys.dm_exec_query_plan(qs.plan_handle) AS qp
WHERE (qs.max_logical_reads > 400
OR qs.max_logical_reads > 400)
```

D)

```
SELECT TOP 50 *
FROM sys.dm_exec_query_stats as qs
WHERE (qs.max_physical)_reads > 400
OR qs.max_physical_reads > 400)
ORDER BY total_worker_time DESC
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

Total_worker_time: Total amount of CPU time, reported in microseconds (but only accurate to milliseconds), that was consumed by executions of this plan since it was compiled.

NEW QUESTION 76

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance named SQL2012. You are in the process of migrating a database from a SQL Server 2008 instance named SQL2008 to the SQL2012 instance.

You have upgraded a database from the SQL2008 instance by using the side-by-side migration technique. You need to migrate the SQL Server logins from the SQL2008 instance to the SQL2012 instance.

What should you do?

- A. Back up the master database on the SQL2008 instanc
- B. Restore the master database on the SQL2012 instance
- C. Use the Transfer Logins task in a Microsoft SQL Server Integrated Services package
- D. Use sp_grantlogin
- E. Use xp_logininfo.

Answer: C

Explanation:

sp_grantlogin creates a SQL Server login.

NEW QUESTION 81

- (Exam Topic 7)

You plan to install Microsoft SQL Server 2014 for a web hosting company.

The company plans to host multiple web sites, each supported by a SQL Server database.

You need to select an edition of SQL Server that features backup compression of databases, basic data integration features, and low total cost of ownership.

Which edition should you choose?

- A. Express Edition with Tools
- B. Standard Edition
- C. Web Edition
- D. Express Edition with Advanced Services

Answer: B

NEW QUESTION 83

- (Exam Topic 7)

You have a server named Server1 that is hosted in an Azure virtual machine. Server1 contains the following:

One instance of SQL Server 2016 Enterprise

10 databases

500 stored procedures

You have a database named Database1 that is hosted on Server1.

Database1 contains 100 queries that are executed dynamically from web applications. You plan to remove data from the procedure cache on Database1.

You have the following requirements:

Changes to Database1 must not affect other databases that are hosted on Server1

Changes to Database1 must not affect the performance of queries that are stored in other databases.

The solution must minimize administrative effort.

You need to remove the data from the procedure cache as quickly as possible. What should you do?

- A. Run DBCC FREEPROCCACHE.
- B. Run ALTER DATABASE SCOPED CONFIGURATION CLEAR PROCEDURE CACHE in the context of Database 1.
- C. Run DBCC DROPCLEANBUFFERS.
- D. Write a script that iterates through each stored procedure definition and add WITH RECOMPILE to the definition.

Answer: B

Explanation:

You should run ALTER DATABASE SCOPED CONFIGURATION CLEAR PROCEDURE CACHE in the context of Database1. This statement lets you change the settings of a database without affecting other databases that are installed on the instance of SQL Server 2016.

NEW QUESTION 86

- (Exam Topic 7)

You administer a Windows Azure SQL Database database named Inventory that contains a stored procedure named p_AddInventory.

Users need to be able to SELECT from all tables in the database and execute the stored procedure. You need to grant only the necessary permissions.

What should you do?

- A. Grant EXECUTE permission on p_AddInventory to all user
- B. Grant VIEW DEFINITION to all users.
- C. Grant EXECUTE permission on p_AddInventory to all user
- D. Add all users to the db_datawriter role.
- E. Add all users to the db_owner role.
- F. Grant EXECUTE permission on p_AddInventory to all user
- G. Add all users to the db_datareader role.

Answer: D

Explanation:

Members of the db_datareader fixed database role can run a SELECT statement against any table or view in the database.

References: [https://technet.microsoft.com/en-us/library/ms188629\(v=sql.90\).aspx](https://technet.microsoft.com/en-us/library/ms188629(v=sql.90).aspx)

NEW QUESTION 87

- (Exam Topic 7)

Database DB1 must use two CPU cores.

Queries that were running on database DB2 prior to migration do not complete. You need to configure the databases.

In the table below, identify the parameter that must be configured for each databases. Select one option for DB1, and one option for DB2. Select one option for each column.

Parameter	DB1	DB2
MAXDOP	<input type="radio"/>	<input type="radio"/>
LEGACY_CARDINALITY_ESTIMATION	<input type="radio"/>	<input type="radio"/>
PARAMETER_SNIFFING	<input type="radio"/>	<input type="radio"/>
QUERY_OPTIMIZER_HOTFIXES	<input type="radio"/>	<input type="radio"/>
CLEAR PROCEDURE_CACHE	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Query_optimizer_hotfixes DB1: MAXDOP

You can use the max degree of parallelism (MAXDOP) option to limit the number of processors to use in parallel plan execution.

DB2: LEGACY_CARDINALITY_ESTIMATION

The CE (Cardinality Estimation) predicts how many rows your query will likely return. The cardinality prediction is used by the Query Optimizer to generate the optimal query plan. With more accurate estimations, the Query Optimizer can usually do a better job of producing a more optimal query plan.

Legacy CE: For a SQL Server database set at compatibility level 120 and above, the CE version 70 can be activated by using the at the database level by using the ALTER DATABASE SCOPED CONFIGURATION.

Example:

```
ALTER DATABASE SCOPED CONFIGURATION SET LEGACY_CARDINALITY_ESTIMATION = ON; GO
```

NEW QUESTION 91

- (Exam Topic 7)

Background

You manage the Microsoft SQL Server environment for a company that manufactures and sells automobile parts.

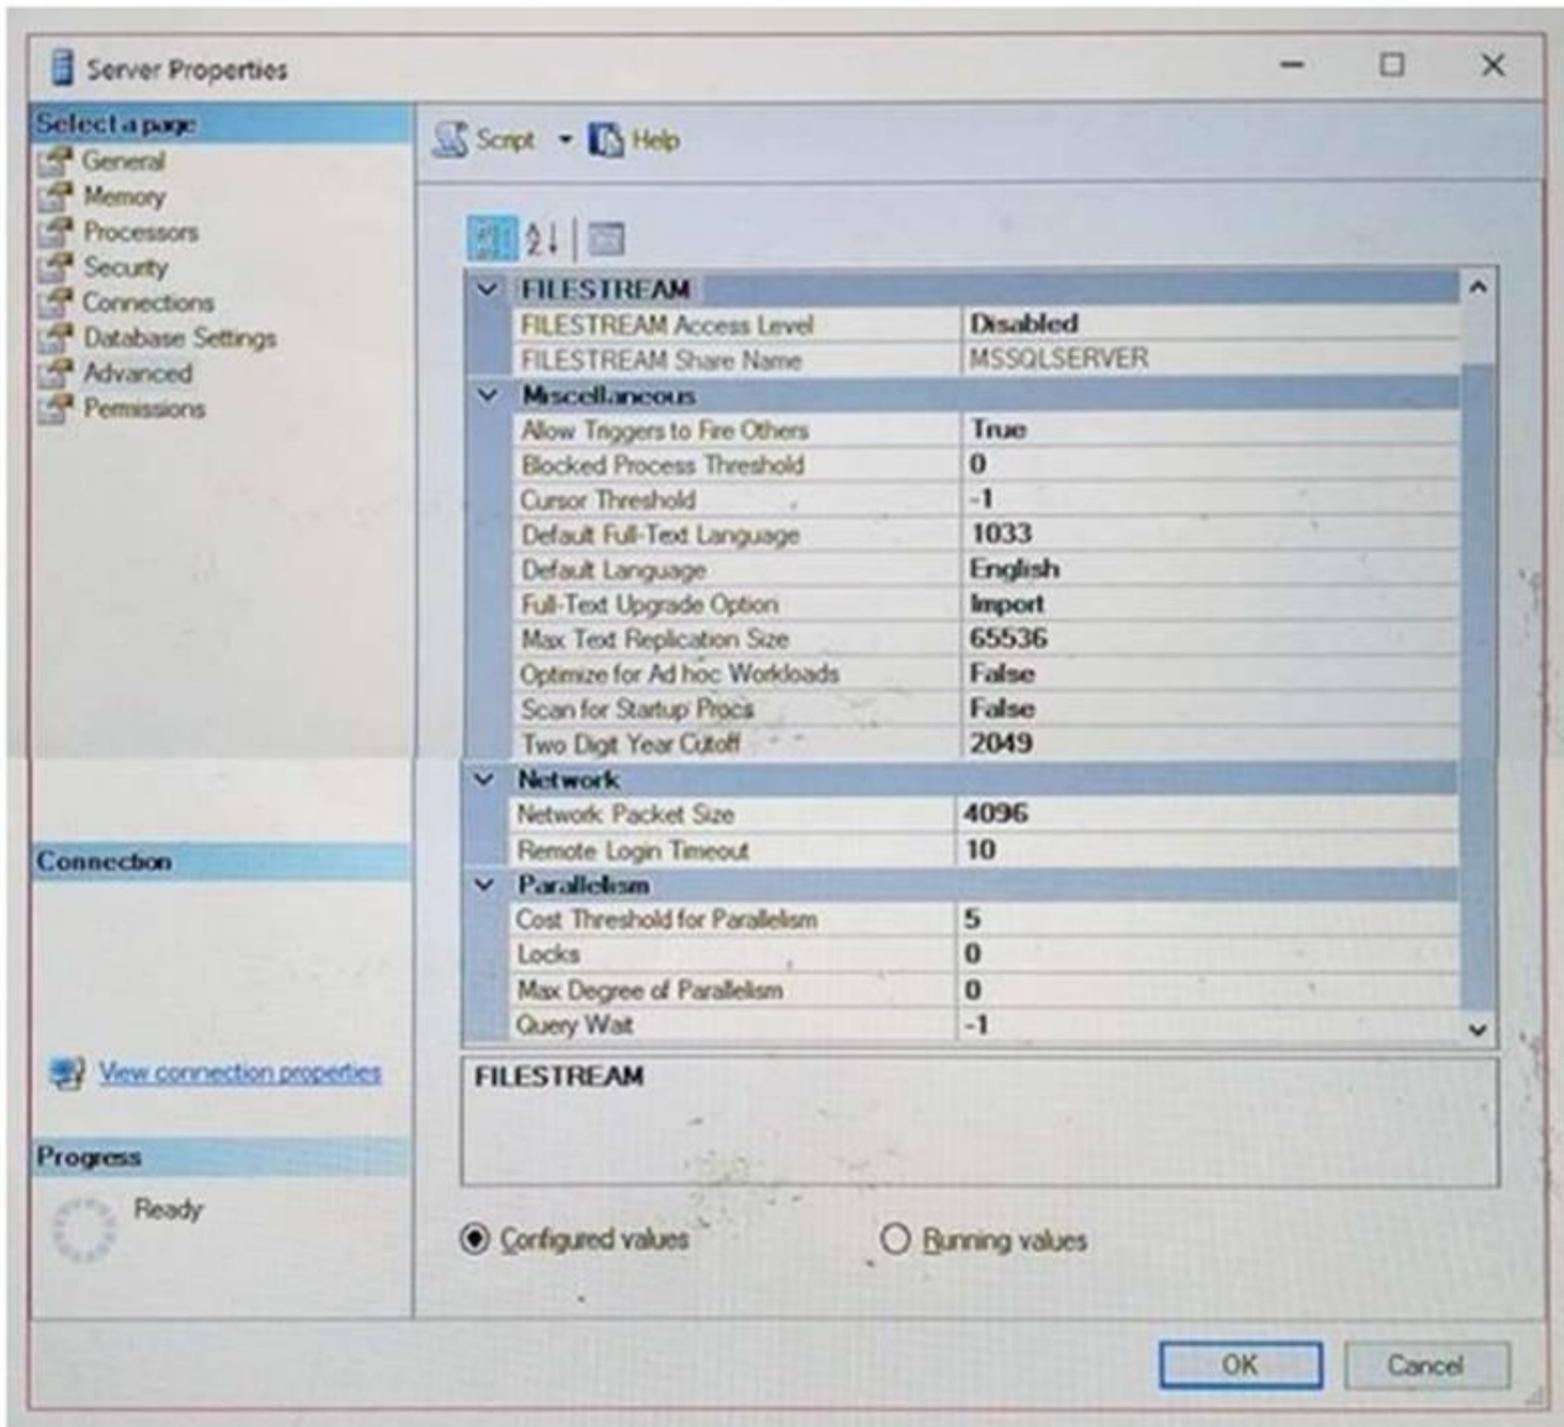
The environment includes the following servers: SRV1 and SRV2. SRV1 has 16 logical cores and hosts a SQL Server instance that supports a mission-critical application. The application has approximately 30,000 concurrent users and relies heavily on the use of temporary tables.

The environment also includes the following databases: DB1, DB2, and Reporting. The Reporting database is protected with Transparent Data Encryption (TDE). You plan to migrate this database to a new server. You detach the database and copy it to the new server.

You are performing tuning on a SQL Server database instance. The application which uses the database was written using an object relationship mapping (ORM) tool which maps tables as objects within the application code. There are 30 stored procedures that are regularly used by the application.

After reviewing the plan cache you have identified that a large number of simple queries are using parallelism, and that execution plans are not being kept in the plan cache for very long.

You review the properties of the instance (Click the Exhibit button). Exhibit:



You need to restore the Reporting database to SRV2. What should you do? To answer, drag the appropriate options to the correct locations. Each option may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Select and Place:

Values

master encryption key on the master database
service master key
server certificate
Reporting database .mdf file
master key password

Answer area

1. Copy the certificate and private key backups from the old server to the new server.
2. Create:
3. Restore:
4. Attach the Reporting database.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 2: Create: server certificate

Recreate the server certificate by using the original server certificate backup file.

Note: The password must be the same as the password that was used when the backup was created. Step 3: Restore: Reporting database .mdf file.

-- Attach the database that is being moved.

-- The path of the database files must be the location where you have stored the database files. Example:

```
CREATE DATABASE [CustRecords] ON
( FILENAME = N'C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA
\CustRecords.mdf' ),
( FILENAME = N'C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA
\CustRecords_log.LDF' ) FOR ATTACH ;
GO
```

From scenario: The Reporting database is protected with Transparent Data Encryption (TDE). You plan to migrate this database to a new server. You detach the database and copy it to the new server.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/move-a-tdeprotected-database-to-a>

NEW QUESTION 92

- (Exam Topic 7)

You are a database developer for an application hosted on a Microsoft SQL Server 2014 server. The database contains two tables that have the following definitions:

```
CREATE TABLE Customer
(CustomerID int NOT NULL PRIMARY KEY,
 CustomerName varchar(50) NOT NULL)

CREATE TABLE Orders
(OrderID int NOT NULL PRIMARY KEY,
 CustomerID int NOT NULL FOREIGN KEY REFERENCES Customer (CustomerID),
 OrderAmount money NOT NULL,
 ShippingCountry varchar(50) NOT NULL)
```

Global customers place orders from several countries. You need to view the country from which each customer has placed the most orders. Which Transact-SQL query do you use?

- A. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer c INNER JOIN (SELECT CustomerID, ShippingCountry, RANK() OVER (PARTITION BY CustomerID ORDER BY COUNT(OrderAmount) DESC) AS Rnk FROM Orders GROUP BY CustomerID, ShippingCountry) AS o ON c.CustomerID = o.CustomerID WHERE o.Rnk = 1
- B. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM (SELECT c.CustomerID, c.CustomerName, o.ShippingCountry, RANK() OVER (PARTITION BY CustomerID ORDER BY COUNT(o.OrderAmount) ASC) AS Rnk FROM Customer c INNER JOIN Orders o ON c.CustomerID = o.CustomerID GROUP BY c.CustomerID, c.CustomerName, o.ShippingCountry) cs WHERE Rnk = 1
- C. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer c INNER JOIN (SELECT CustomerID, ShippingCountry, RANK() OVER (PARTITION BY CustomerID ORDER BY OrderAmount DESC) AS Rnk FROM Orders GROUP BY CustomerID, ShippingCountry) AS o ON c.CustomerID = o.CustomerID WHERE o.Rnk = 1
- D. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer c INNER JOIN (SELECT CustomerID, ShippingCountry, COUNT(OrderAmount) DESC) AS OrderAmount FROM Orders GROUP BY CustomerID, ShippingCountry) AS o ON c.CustomerID = o.CustomerID ORDER BY OrderAmount DESC

Answer: A

Explanation:

Use descending (DESC) ordering.

To order by the number of orders we use ORDER BY COUNT(OrderAmount). Finally a WHERE clause is needed: WHERE o.Rnk = 1

NEW QUESTION 93

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft Azure virtual machine.

You suspect that the current SQL Server indexes cause queries to execute slowly.
 You need to identify which indexes must be created to reduce the query execution time.
 Which three dynamic management views should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. sys.dm_db_index_physical_stats
- B. sys.dm_db_missing_index_group_stats
- C. sys.indexes
- D. sys.dm_db_index_usage_stats
- E. sys.dm_db_missing_index_groups
- F. sys.dm_db_index_operational_stats
- G. sys.dm_db_missing_index_details
- H. sys.sysindexkeys

Answer: BEG

Explanation:

The missing indexes feature consists of the following components:
 A set of dynamic management objects that can be queried to return information about missing indexes.
 The Missing Indexes element in XML Showplans, which correlate indexes that the query optimizer considers missing with the queries for which they are missing.
 Dynamic Management Objects
 After running a typical workload on SQL Server, you can retrieve information about missing indexes by querying the dynamic management objects listed in the following table. These dynamic management objects are stored in the master database.
 sys.dm_db_missing_index_group_stats
 Returns summary information about missing index groups, for example, the performance improvements that could be gained by implementing a specific group of missing indexes.
 sys.dm_db_missing_index_groups
 Returns information about a specific group of missing indexes, such as the group identifier and the identifiers of all missing indexes that are contained in that group.
 sys.dm_db_missing_index_details
 Returns detailed information about a missing index; for example, it returns the name and identifier of the table where the index is missing, and the columns and column types that should make up the missing index.
 sys.dm_db_missing_index_columns
 Returns information about the database table columns that are missing an index. References: [https://technet.microsoft.com/en-us/library/ms345524\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms345524(v=sql.105).aspx)

NEW QUESTION 94

- (Exam Topic 7)

You use Microsoft Azure Resource Manager to deploy two new Microsoft SQL Server instances in an Azure virtual machine (VM). VM has 28 gigabytes (GB) of memory. The instances are named Instance1 and Instance2, respectively.
 The various databases on the instances have the following characteristics:

Instance name	Aggregate database size	Daily working set	Concurrent users
Instance1	200 GB	25 GB	2,000
Instance2	300 GB	10 GB	2,000

You run the following Transact-SQL statements:

```
sp_configure 'show advanced options', 1;
GO
RECONFIGURE;
GO
```

You need to configure each SQL Server instance to correctly allocate memory. What should you do?

- A. On Instance1, run the following Transact-SQL code: On Instance2, run the following Transact-SQL code:
- B. On Instance1, run the following Transact-SQL code: On Instance2, run the following Transact-SQL code:
- C. On Instance1, run the following Transact-SQL code: On Instance2, run the following Transact-SQL code:
- D. On Instance1, run the following Transact-SQL code: On Instance2, run the following Transact-SQL code:

Answer: D

NEW QUESTION 97

- (Exam Topic 7)

You administer a Microsoft SQL Server 2016 instance.
 You need to configure a new database to support FILETABLES. What should you do? Choose all that apply.

- A. Disable FILESTREAM on the Database.
- B. Enable FILESTREAM on the Server Instance.
- C. Configure the Database for Partial Containment.
- D. Create a non-empty FILESTREAM file group.
- E. Enable Contained Databases on the Server Instance.
- F. Set the FILESTREAM directory name on the Database.

Answer: BDF

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/blob/enable-the-prerequisites-for-filetable>

NEW QUESTION 101

- (Exam Topic 7)

You have Microsoft SQL server on a Microsoft Azure virtual machine. The virtual machine has 200 GB of data.

User report a slow response time when querying the database.

You need to identify whether the storage subsystem causes the performance issue. Which performance monitor counter should you view?

- A. Data sec/Write
- B. Avg.disk Read Queue Length
- C. % Disk Read Time
- D. Disk sec/Read

Answer: B

NEW QUESTION 105

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance named SQL2012 that hosts an OLTP database of 1 terabyte in size.

The database is modified by users only from Monday through Friday from 09:00 hours to 17:00 hours. Users modify more than 30 percent of the data in the database during the week.

Backups are performed as shown in the following schedule:

Type	Frequency
Full	Sunday at 20:00 hours
Differential	Monday through Friday at 20:00 hours
Log	Monday through Friday between 08:00 hours and 18:00 hours

The Finance department plans to execute a batch process every Saturday at 09:00 hours. This batch process will take a maximum of 8 hours to complete.

The batch process will update three tables that are 10 GB in size. The batch process will update these tables multiple times.

When the batch process completes, the Finance department runs a report to find out whether the batch process has completed correctly.

You need to ensure that if the Finance department disapproves the batch process, the batch operation can be rolled back in the minimum amount of time. What should you do on Saturday?

- A. Perform a differential backup at 08:59 hours.
- B. Record the LSN of the transaction log at 08:59 hour
- C. Perform a transaction log backup at 17:01 hours.
- D. Create a database snapshot at 08:59 hours.
- E. Record the LSN of the transaction log at 08:59 hour
- F. Perform a transaction log backup at 08:59 hours.
- G. Create a marked transaction in the transaction log at 08:59 hour
- H. Perform a transaction log backup at 17:01 hours.
- I. Create a marked transaction in the transaction log at 08:59 hour
- J. Perform a transaction log backup at 08:59 hours.

Answer: C

Explanation:

References: <https://docs.microsoft.com/en-us/sql/relational-databases/databases/database-snapshots-sql-server>

NEW QUESTION 108

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database instance.

You plan to migrate the database to Windows Azure SQL Database. You verify that all objects contained in the database are compatible with Windows Azure SQL Database.

You need to ensure that database users and required server logins are migrated to Windows Azure SQL Database.

What should you do?

- A. Use the copy database wizard
- B. Use the Database Transfer wizard
- C. Use SQL Server Management Studio to deploy the database to Windows Azure SQL Database
- D. Backup the database from the local server and restore it to Windows Azure SQL Database

Answer: C

Explanation:

You would need to use either the SQL Server Management Studio or Transact-SQL.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-cloud-migrate>

NEW QUESTION 113

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database named Contoso on a server named Server01.

You need to collect data for a long period of time to troubleshoot wait statistics when querying Contoso. You also need to ensure minimum impact to the server.

What should you create?

- A. An Alert
- B. A Resource Pool

- C. An Extended Event session
- D. A Server Audit Specification
- E. A SQL Profiler Trace
- F. A Database Audit Specification
- G. A Policy
- H. A Data Collector Set

Answer: C

Explanation:

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. A SQL Server Extended Events session is created in the SQL Server process hosting the Extended Events engine.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/extended-events/extended-events>

NEW QUESTION 117

- (Exam Topic 7)

You have Microsoft SQL Server on a DS-series Microsoft Azure virtual machine. The virtual machine has 28 GB of memory.

You discover the following performance statistics on the server:

The average Page life expectancy is 30.

The server has excessive PAGELATCH_IO waits.

You need to decrease the PAGELATCH_IO waits. What should you do?

- A. Enable large-page support.
- B. Enable lock pages in memory.
- C. Configure buffer pool extensions.
- D. Add more tempdb files.

Answer: D
Explanation:
References:

NEW QUESTION 122

- (Exam Topic 7)

User report that a query takes a long time to execute. The query has the following wait statistics.

```
<WaitStats>
  <Wait WaitType="MEMORY_ALLOCATION_EXT" WaitTimeMs="186" WaitCount="112046" />
  <Wait WaitType="PAGEIOLATCH_SH" WaitTimeMs="37001" WaitCount="183" />
  <Wait WaitType="SOS_SCHEDULER_YIELD" WaitTimeMs="399" WaitCount="12321" />
  <Wait WaitType="WRITELOG" WaitTimeMs="1632" WaitCount="627" />
  <Wait WaitType="IO_COMPLETION" WaitTimeMs="100287" WaitCount="5300" />
  <Wait WaitType="PAGEIOLATCH_UP" WaitTimeMs="59652" WaitCount="21027" />
  <Wait WaitType="PAGEIOLATCH_EX" WaitTimeMs="1116329" WaitCount="1840528" />
</WaitStats>
```

Which resource causes the issue?

- A. processor
- B. disk
- C. blocking
- D. network

Answer: B

Explanation:

PAGEIOLATCH Wait time and WaitCount are both high.

One of the most common wait type seen on SQL Server and definitely one that causes a lot of troubles to less experienced database administrators is the PAGEIOLATCH_SH wait type. This is one of those wait types that clearly indicates one thing, but which background and potential causes are much subtler and may lead to erroneous conclusions and worse, incorrect solutions

The Microsoft definition of this wait type is:

Occurs when a task is waiting on a latch for a buffer that is in an I/O request. The latch request is in Shared mode. Long waits may indicate problems with the disk subsystem.

References: https://www.sqlshack.com/handling-excessive-sql-server-pageiolatch_sh-wait-types/

NEW QUESTION 127

- (Exam Topic 7)

You administer all the deployments of Microsoft SQL Server 2014 in your company. You have two servers in the same data center that hosts your production database.

You need to ensure that the database remains available if a catastrophic server failure or a disk failure occurs. You also need to maintain transactional consistency of the data across both servers.

You need to achieve these goals without manual intervention. Which configuration should you use?

- A. Two servers configured in a Windows Failover Cluster in the same data center SQL Server configured as a clustered instance
- B. SQL Server that includes an application database configured to perform transactional replication
- C. Two servers configured in the same data center A primary server configured to perform log-shipping every 10 minutes A backup server configured as a warm standby
- D. Two servers configured in different data centers SQL Server Availability Group configured in Synchronous-Commit Availability Mode One server configured as an Active Secondary
- E. Two servers configured in the same data center SQL Server Availability Group configured in Asynchronous-Commit Availability Mode One server configured as an Active Secondary

- F. Two servers configured in different data centers SQL Server Availability Group configured in Asynchronous-Commit Availability Mode
- G. SQL Server that includes an application database configured to perform snapshot replication
- H. Two servers configured on the same subnet SQL Server Availability Group configured in Synchronous-Commit Availability Mode

Answer: H

Explanation:

Always On availability groups supports two availability modes— asynchronous-commit mode and synchronous-commit mode. Synchronous-commit mode emphasizes high availability over performance, at the cost of increased transaction latency. References: <https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-mode>

NEW QUESTION 129

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database named Contoso on a server named Server01. You need to write messages to the Application Log when users are added to or removed from a fixed server role in Server01. What should you create?

- A. A Database Audit Specification
- B. A Policy
- C. An Alert
- D. A SQL Profiler Trace
- E. A Resource Pool
- F. An Extended Event session
- G. A Server Audit Specification

Answer: G

Explanation:

The SQL Server Audit feature enables you to audit server-level and database-level groups of events and individual events. Audits can have the following categories of actions: Server-level. These actions include server operations, such as management changes, such as in this question, and logon and logoff operations. Database-level. These actions encompass data manipulation languages (DML) and data definition language (DDL) operations. Audit-level. These actions include actions in the auditing process. References: [http://technet.microsoft.com/en-us/library/cc280663\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/cc280663(v=sql.105).aspx)

NEW QUESTION 132

- (Exam Topic 7)

You are migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 databases that consume a total of 2 TB of disk space. The instance sustains more than 30,000 transactions per second. You need to provision storage for the virtual machine. The storage must be able to support the same load as the on-premises deployment. Solution: You use drive D on the virtual machine to store the database files. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

The D drive should only be used for temporary data.

NEW QUESTION 137

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database. You want to make a full backup of the database to a file on disk. In doing so, you need to output the progress of the backup. Which backup option should you use?

- A. STATS
- B. COMPRESSION
- C. CHECKSUM
- D. IN IT

Answer: A

Explanation:

STATS is a monitoring option of the BACKUP command. STATS [[=percentage]] Displays a message each time another percentage completes, and is used to gauge progress. If percentage is omitted, SQL Server displays a message after each 10 percent is completed. The STATS option reports the percentage complete as of the threshold for reporting the next interval. This is at approximately the specified percentage; for example, with STATS=10, if the amount completed is 40 percent, the option might display 43 percent. For large backup sets, this is not a problem, because the percentage complete moves very slowly between completed I/O calls. References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/backup-transact-sql>

NEW QUESTION 139

- (Exam Topic 7)

You administer a Microsoft SQL Server database named Sales. The database is 3 terabytes in size. The Sales database is configured as shown in the following table.

Filegroup	File
PRIMARY	<ul style="list-style-type: none"> • Sales.mdf
XACTIONS	<ul style="list-style-type: none"> • Sales_1.ndf • Sales_2.ndf • Sales_3.ndf
ARCHIVES	<ul style="list-style-type: none"> • SalesArch_1.ndf • SalesArch_2.ndf

You discover that all files except Sales_2.ndf are corrupt.

You need to recover the corrupted data in the minimum amount of time. What should you do?

- A. Perform a file restore.
- B. Perform a transaction log restore.
- C. Perform a restore from a full backup.
- D. Perform a filegroup restore.

Answer: A

Explanation:

In a file restore, the goal is to restore one or more damaged files without restoring the whole database. References: <https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/file-restores-simple-recovery-model>

NEW QUESTION 143

- (Exam Topic 7)

You have an on-premises server that runs Windows Server 2012 R2. The server has a Microsoft SQL Server 2016 instance that has one user database. The database is 2 TB.

Your company has a Win32 application installed on 1,000 computers. The application connects to the database by using a network name of server1.contoso.local. You need to migrate the database to SQL Server 2016 on a Microsoft Azure virtual machine that runs Windows Server 2016. The solution must minimize outages to the application.

What should you do?

- A. Copy the database files and update the records in DNS.
- B. Implement an availability group and update the records in DNS.
- C. Implement database mirroring and update the records in DNS.
- D. Implement database mirroring and change the connection string.

Answer: B

Explanation:

SQL Server high availability and disaster recovery (HADR) technologies that are supported in Azure include: References:

NEW QUESTION 147

- (Exam Topic 7)

You use Microsoft SQL Server 2014 to develop a database application. You need to implement a computed column that references a lookup table by using an INNER JOIN against another table.

What should you do?

- A. Reference a user-defined function within the computed column.
- B. Create a BEFORE trigger that maintains the state of the computed column.
- C. Add a default constraint to the computed column that implements hard-coded values.
- D. Add a default constraint to the computed column that implements hard-coded CASE statements.

Answer: A

Explanation:

A common way to define a computed column is by using a user-defined function (UDF) to encapsulate the calculation logic. References: <https://blogs.msdn.microsoft.com/sqlcat/2011/11/28/a-computed-column-defined-with-a-user-define>

NEW QUESTION 152

- (Exam Topic 7)

You use Microsoft SQL Server 2014 to develop a database application. You need to create an object that meets the following requirements: Which object should you use?

- A. Scalar-valued function
- B. Inline function
- C. User-defined data type
- D. Stored procedure

Answer: D

Explanation:

Stored procedures accept input parameters and return multiple values in the form of output parameters to the calling program. They cannot be used in views. References: <https://docs.microsoft.com/en-us/sql/relational-databases/stored-procedures/stored-procedures-datab>

NEW QUESTION 153

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

70-765 Practice Exam Features:

- * 70-765 Questions and Answers Updated Frequently
- * 70-765 Practice Questions Verified by Expert Senior Certified Staff
- * 70-765 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 70-765 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 70-765 Practice Test Here](#)