

70-765 Dumps

Provisioning SQL Databases (beta)

<https://www.certleader.com/70-765-dumps.html>



NEW QUESTION 1

- (Topic 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 is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in a development environment. Each VM has a dedicated disk for backups.

You need to backup a database to the local disk on a VM. The backup must be replicated to another region.

Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 diskstorage
- C. Premium P30 disk storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage
- F. Standard zone redundant blob storage
- G. Standard locally redundant blob storage
- H. Standard geo-redundant blob storage

Answer: E

Explanation:

Note: SQL Database automatically creates a database backups and uses Azure read- access geo-redundant storage (RA-GRS) to provide geo-redundancy.

These backups are created automatically and at no additional charge. You don't need to do anything to make them happen. Database backups are an essential part of any business continuity and disaster recovery strategy because they protect your data from accidental corruption or deletion.

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

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

DRAG DROP - (Topic 2)

You deploy a new Microsoft Azure SQL Database instance to support a variety of mobile applications and public websites. You plan to create a new security principal named User1.

The principal must have access to select all current and future objects in a database named Reporting. The activity and authentication of the database user must

be limited to the Reporting database.

You need to create the new security principal.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

| Actions | Answer Area |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| In SQL Server Management Studio, create a connection to the Reporting database on the Azure SQL Server instance. | |
| In SQL Server Management Studio, create a connection to the master database on the Azure SQL Server instance. | |
| Run the following Transact-SQL statement: EXEC sp_addrolemember 'db_datareader', 'User1' | |
| Run the following Transact_SQL statement: CREATE LOGIN User1 WITH password='Pa\$\$w0rd' | |
| Run the following Transact_SQL statement: CREATE USER User1 WITH password='Pa\$\$w0rd' | |
| Run the following Transact_SQL statements: EXEC sp_migrate_user_to_contained @username = N'User1', @rename = N'keep_name', @disablelogin = N'disable_login' | |
| Run the following Transact_SQL statement: CREATE LOGIN User1 FROM EXTERNAL PROVIDER | |
| Select the Reporting database and run the following Transact-SQL statements: CREATE USER User1 from LOGIN User1 GRANT SELECT TO User1 | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1, Step 2:

First you need to create a login for SQL Azure, it's syntax is as follows: CREATE LOGIN username WITH password='password'; This command needs to run in master db. Only afterwards can you run commands to create a user in the database.

Step 3:

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

NEW QUESTION 4

- (Topic 2)

You plan to deploy 20 Microsoft Azure SQL Database instances to an elastic pool in Azure to support a batch processing application.

Two of the databases in the pool reach their peak workload threshold at the same time every day. This leads to inconsistent performance for batch completion.

You need to ensure that all batches perform consistently. What should you do?

- A. Create an In-Memory table.
- B. Increase the storage limit in the pool.
- C. Implement a readable secondary database.
- D. Increase the total number of elastic Database Transaction Units (eDTUs) in the pool.

Answer: D

Explanation:

In SQL Database, the relative measure of a database's ability to handle resource demands is expressed in Database Transaction Units (DTUs) for single databases and elastic DTUs (eDTUs) for databases in an elastic pool.

A pool is given a set number of eDTUs, for a set price. Within the pool, individual databases are given the flexibility to auto-scale within set parameters. Under heavy load, a database can consume more eDTUs to meet demand.

Additional eDTUs can be added to an existing pool with no database downtime. References:<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool>

NEW QUESTION 5

- (Topic 2)

You manage a Microsoft SQL Server environment in a Microsoft Azure virtual machine. You must enable Always Encrypted for columns in a database.

You need to configure the key store provider.

What should you do?

- A. Manually specify the column master key.
- B. Modify the connection string for applications.
- C. Auto-generate a column master key.
- D. Use the Windows certificate store.

Answer: D

Explanation:

Always Encrypted supports multiple key stores for storing Always Encrypted column master keys. A column master key can be a certificate stored in Windows Certificate Store.

References: <https://msdn.microsoft.com/en-us/library/mt723359.aspx>

NEW QUESTION 6

- (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 LEGACY_CARDINALITY_ESTIMATION option for the databases. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

LEGACY_CARDINALITY_ESTIMATION = { ON | OFF | PRIMARY }

Enables you to set the query optimizer cardinality estimation model to the SQL Server 2012 and earlier version independent of the compatibility level of the database. This is equivalent to Trace Flag 9481.

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

NEW QUESTION 7

- (Topic 2)

Note: This questions 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.

Your company has several Microsoft Azure SQL Database instances.

Data encryption should be allowed to be implemented by the client applications that access the data. Encryption keys should not be made available to the database engine.

You need to configure the database. What should you implement?

- A. transport-level encryption
- B. cell-level encryption
- C. Transparent Data Encryption
- D. Always Encrypted
- E. Encrypting File System
- F. BitLocker
- G. dynamic data masking

Answer: A

Explanation:

Using encryption during transit with Azure File Shares

Azure File Storage supports HTTPS when using the REST API, but is more commonly used as an SMB file share attached to a VM.

HTTPS is a transport-level security protocol.

NEW QUESTION 8

- (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 configure the Resource Governor to set the MAXDOP parameter to 0 for all queries against the database.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

SQL Server will consider parallel execution plans for queries, index data definition language (DDL) operations, and static and keyset-driven cursor population.

You can override the max degree of parallelism value in queries by specifying the MAXDOP query hint in the query statement.

References: [https://technet.microsoft.com/en-us/library/ms181007\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms181007(v=sql.105).aspx)

NEW QUESTION 9

- (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 10

- (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 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 14

- (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 17

- (Exam Topic 7)

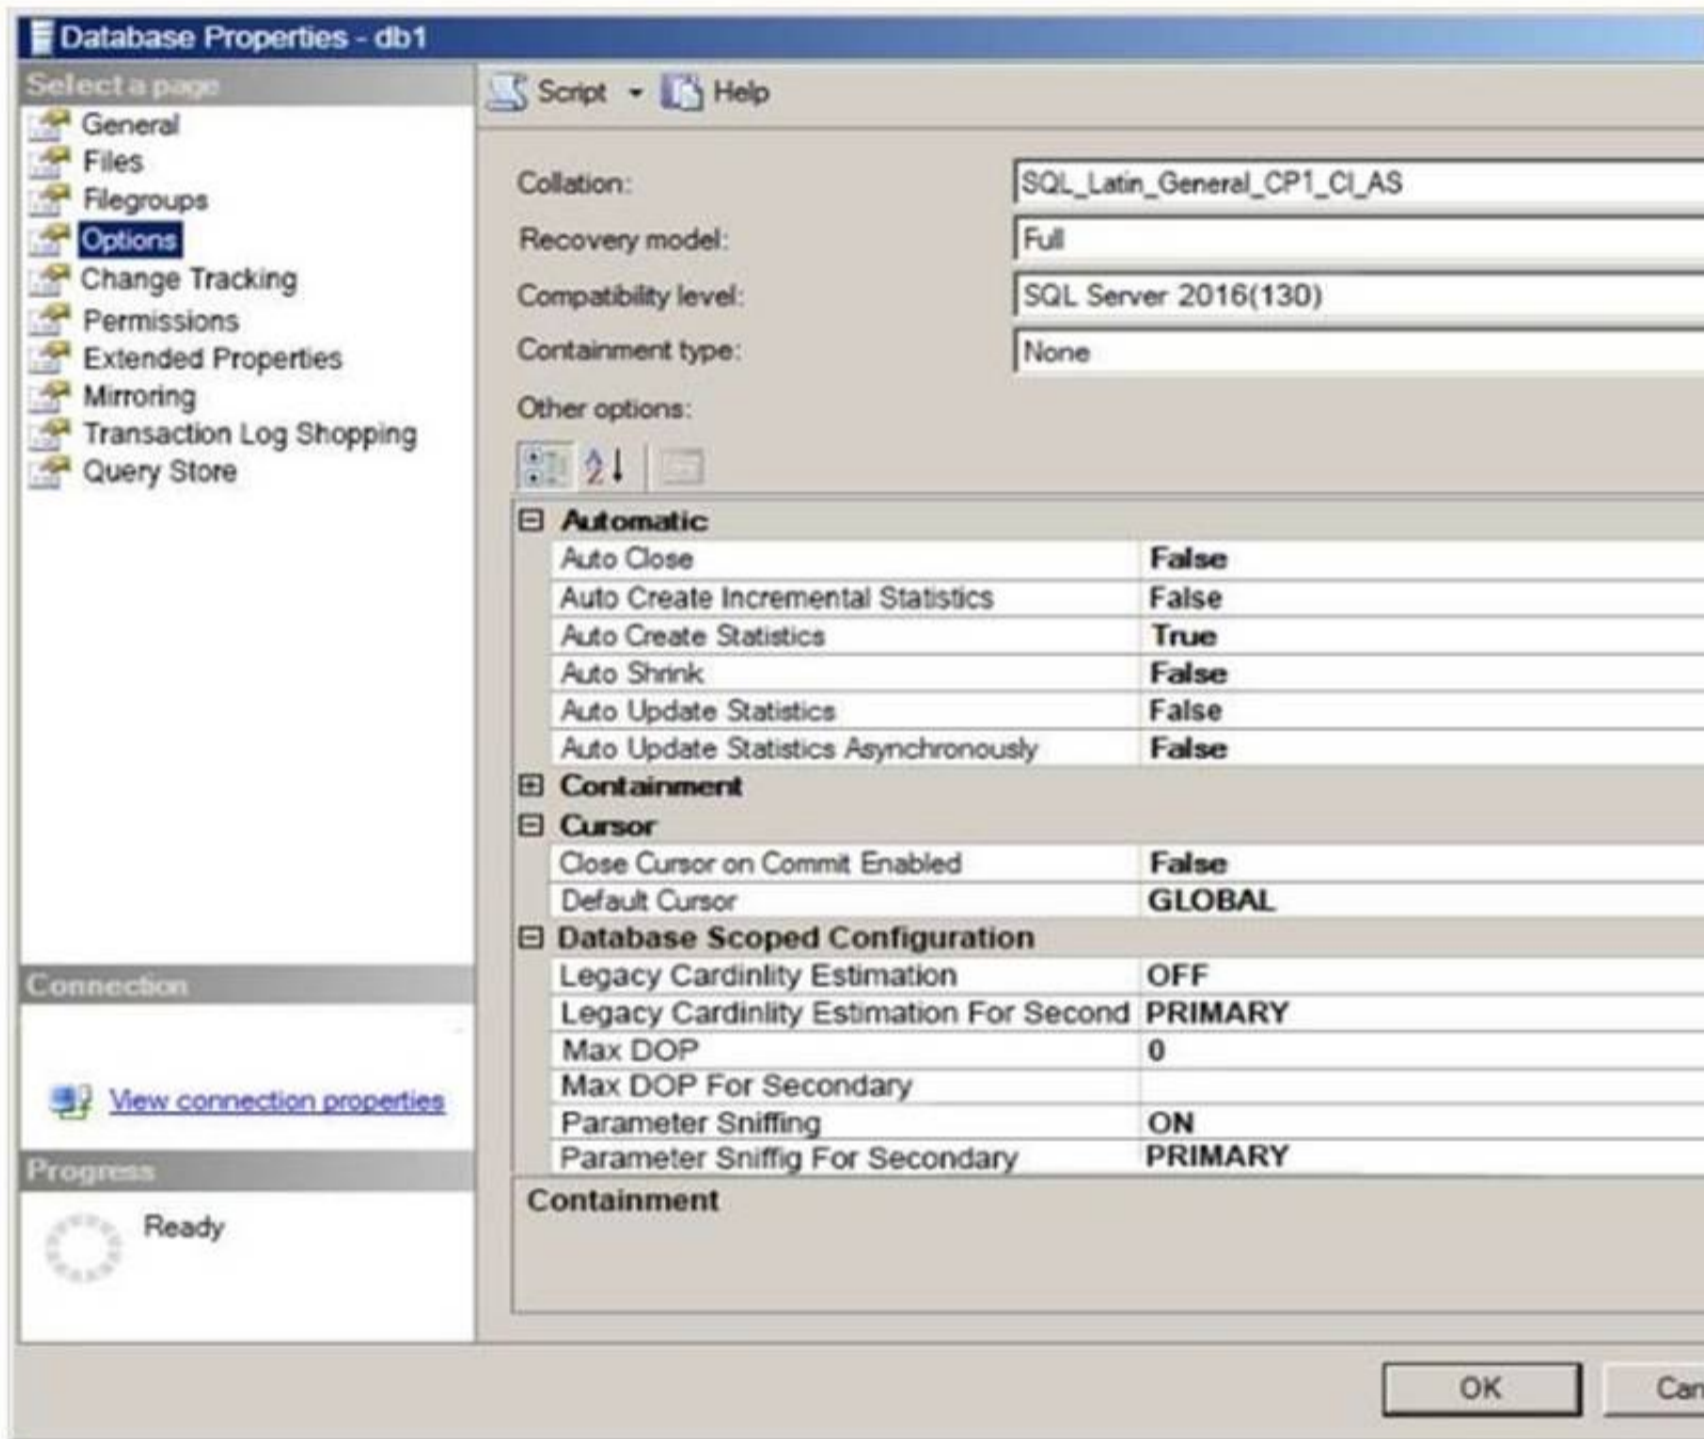
You have Microsoft SQL Server on a Microsoft Azure virtual machine. The virtual machine has a database named DB1. DB1 contains a table named Table1 that has 4 billion rows.

Users report that a query using Table1 takes longer than expected to execute.

You review the execution plan for the query and discover that the expected number of returned rows is one, while the actual number of returned rows is 1 million.

You need to reduce the amount of time it takes for the query to execute. The solution must prevent additional performance issues from being introduced.

Hot Area:



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

When you set the AUTO_CREATE_STATISTICS option on, the Query Optimizer creates statistics on individual columns used in a predicate, if these statistics are not already available. These statistics are necessary to generate the query plan.

References:

<https://www.mssqltips.com/sqlservertip/2766/sql-server-auto-update-and-auto-create-statisticsoptions/>

NEW QUESTION 19

- (Exam Topic 7)

You deploy a new Microsoft Azure SQL database instance to support a variety of mobile application and public websites. You configure geo-replication with regions in Brazil and Japan.

You need to implement real-time encryption of the database and all backups. Solution: you enable Dynamic Data Masking on the primary replica.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

SQL Database dynamic data masking does not encrypt the data. Transparent Data Encryption (TDE) would provide a solution.

Note: SQL Database dynamic data masking limits sensitive data exposure by masking it to non-privileged users.

Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer.

References:

<https://azure.microsoft.com/en-us/blog/how-to-configure-azure-sql-database-geo-dr-with-azure-key-vault/>

NEW QUESTION 24

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database that contains a table named AccountTransaction. You discover that query performance on the table is poor due to fragmentation on the

IDX_AccountTransaction_AccountCode non-clustered index. You need to defragment the index. You also need to ensure that user queries are able to use the index during the defragmenting process.

Which Transact-SQL batch should you use?

- A. ALTER INDEX IDX_AccountTransaction_AccountCode ON AccountTransaction.AccountCode REORGANIZE
- B. ALTER INDEX ALL ON AccountTransaction REBUILD
- C. ALTER INDEX IDX_AccountTransaction_AccountCode ON AccountTransaction.AccountCode REBUILD
- D. CREATE INDEX IDXAccountTransactionAccountCode ON AccountTransaction.AccountCode WITH DROP EXISTING

Answer: A

Explanation:

Reorganize: This option is more lightweight compared to rebuild. It runs through the leaf level of the index, and as it goes it fixes physical ordering of pages and also compacts pages to apply any previously set fillfactor settings. This operation is always online, and if you cancel it then it's able to just stop where it is (it doesn't have a giant operation to rollback).

References: <https://www.brentozar.com/archive/2013/09/index-maintenance-sql-server-rebuild-reorganize/>

NEW QUESTION 28

- (Exam Topic 7)

Settings Value VM size D3

Storage Location Drive E Storage type Standard Tempdb location Drive C

The workload on this instance has of the tempdb load.

You need to maximize the performance of the tempdb database.

Solution: You use a GS- Series VM and store the tempdb database on attached Premium storage. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

For VMs that support Premium Storage (DS-series, DSv2-series, and GS-series), we recommend storing TempDB on a disk that supports Premium Storage with read caching enabled. There is one exception to this recommendation; if your TempDB usage is write-intensive, you can achieve higher performance by storing TempDB on the local D drive, which is also SSD-based on these machine sizes.

References:

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

NEW QUESTION 30

- (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 33

- (Exam Topic 7)

You plan to deploy an AlwaysOn failover cluster in Microsoft Azure. The cluster has a Service Level Agreement (SLA) that requires an uptime of at least 99.95 percent.

You need to ensure that the cluster meets the SLA.

Which cmdlet should you run before you deploy the virtual machine?

- 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: B

Explanation:

On Azure virtual machines, a SQL Server Availability Group requires a load balancer. The load balancer holds the IP address for the Availability Group listener. The New-AzureRmLoadBalancer cmdlet creates an Azure load balancer.

References:

NEW QUESTION 35

- (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 39

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 environment. One of the SQL Server 2014 instances contains a database named Sales.

You plan to migrate Sales to Windows Azure SQL Database. To do so, you need to implement a contained database.

What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Set database containment to AZURE.
- B. Enable server property contained database authentication.
- C. Disable server property cross db ownership chaining.
- D. Set database containment to PARTIAL.
- E. Disable server property contained database authentication.
- F. database containment to FULL.

Answer: BD

Explanation:

A contained database is a database that is isolated from other databases and from the instance of SQL Server that hosts the database.

B: In the contained database user model, the login in the master database is not present. Instead, the authentication process occurs at the user database, and the database user in the user database does not have an associated login in the master database.

SQL Database and SQL Data Warehouse support Azure Active Directory identities as contained database users.

D: The contained database feature is currently available only in a partially contained state. A partially contained database is a contained database that allows the use of uncontained features.

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

NEW QUESTION 43

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database named Contoso on a server named Server01. You need to prevent users from disabling server audits 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: B

Explanation:

Writing to the Windows Security log requires the SQL Server service account to be added to the Generate security audits policy. By default, the Local System, Local Service, and NetworkService are part of this policy. This setting can be configured by using the security policy snap-in (secpol.msc). Additionally, the Audit object access security policy must be enabled for both Success and Failure.

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

NEW QUESTION 44

- (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 48

- (Exam Topic 7)

Settings Value VM size D3

Storage Location Drive E Storage type Standard Tempdb location Drive C

The workload on this instance has of the tempdb load.

You need to maximize the performance of the tempdb database.

Solution: You use a D- Series VM and store the tempdb database on drive D. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

For D-series, Dv2-series, and G-series VMs, the temporary drive on these VMs is SSD-based. If your workload makes heavy use of TempDB (such as temporary objects or complex joins), storing TempDB on the D drive could result in higher TempDB throughput and lower TempDB latency.

References:

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

NEW QUESTION 53

- (Exam Topic 7)

A company has an on-premises Microsoft SQL Server 2016 environment. All futures databases must meet the following requirements:

The recovery model must be set to simple.

The compatibility level must be set to SQL server 2014 (120).

Your need to configure the SQL server 2016 environment.

In the table below, identify the database you must modify for each requirement.

| Answer Area | | |
|-----------------|----------------|-----------------------|
| System database | Recovery model | Compatibility level |
| Master | simple | <input type="radio"/> |
| Msdb | simple | <input type="radio"/> |
| Model | full | <input type="radio"/> |
| Resource | | <input type="radio"/> |
| Tempdb | simple | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Model: Change from full to simple Recovery Model

Newly created user databases use the same recovery model as the model database.

The model database is used as the template for all databases created on an instance of SQL Server. Because tempdb is created every time SQL Server is started, the model database must always exist on a SQL Server system. The entire contents of the model database, including database options, are copied to the new database.

Model: Set compatibility level to 120

For all installations of SQL Server, the default compatibility level is set to the version of the Database Engine. Databases are set to this level unless the model database has a lower compatibility level.

References:

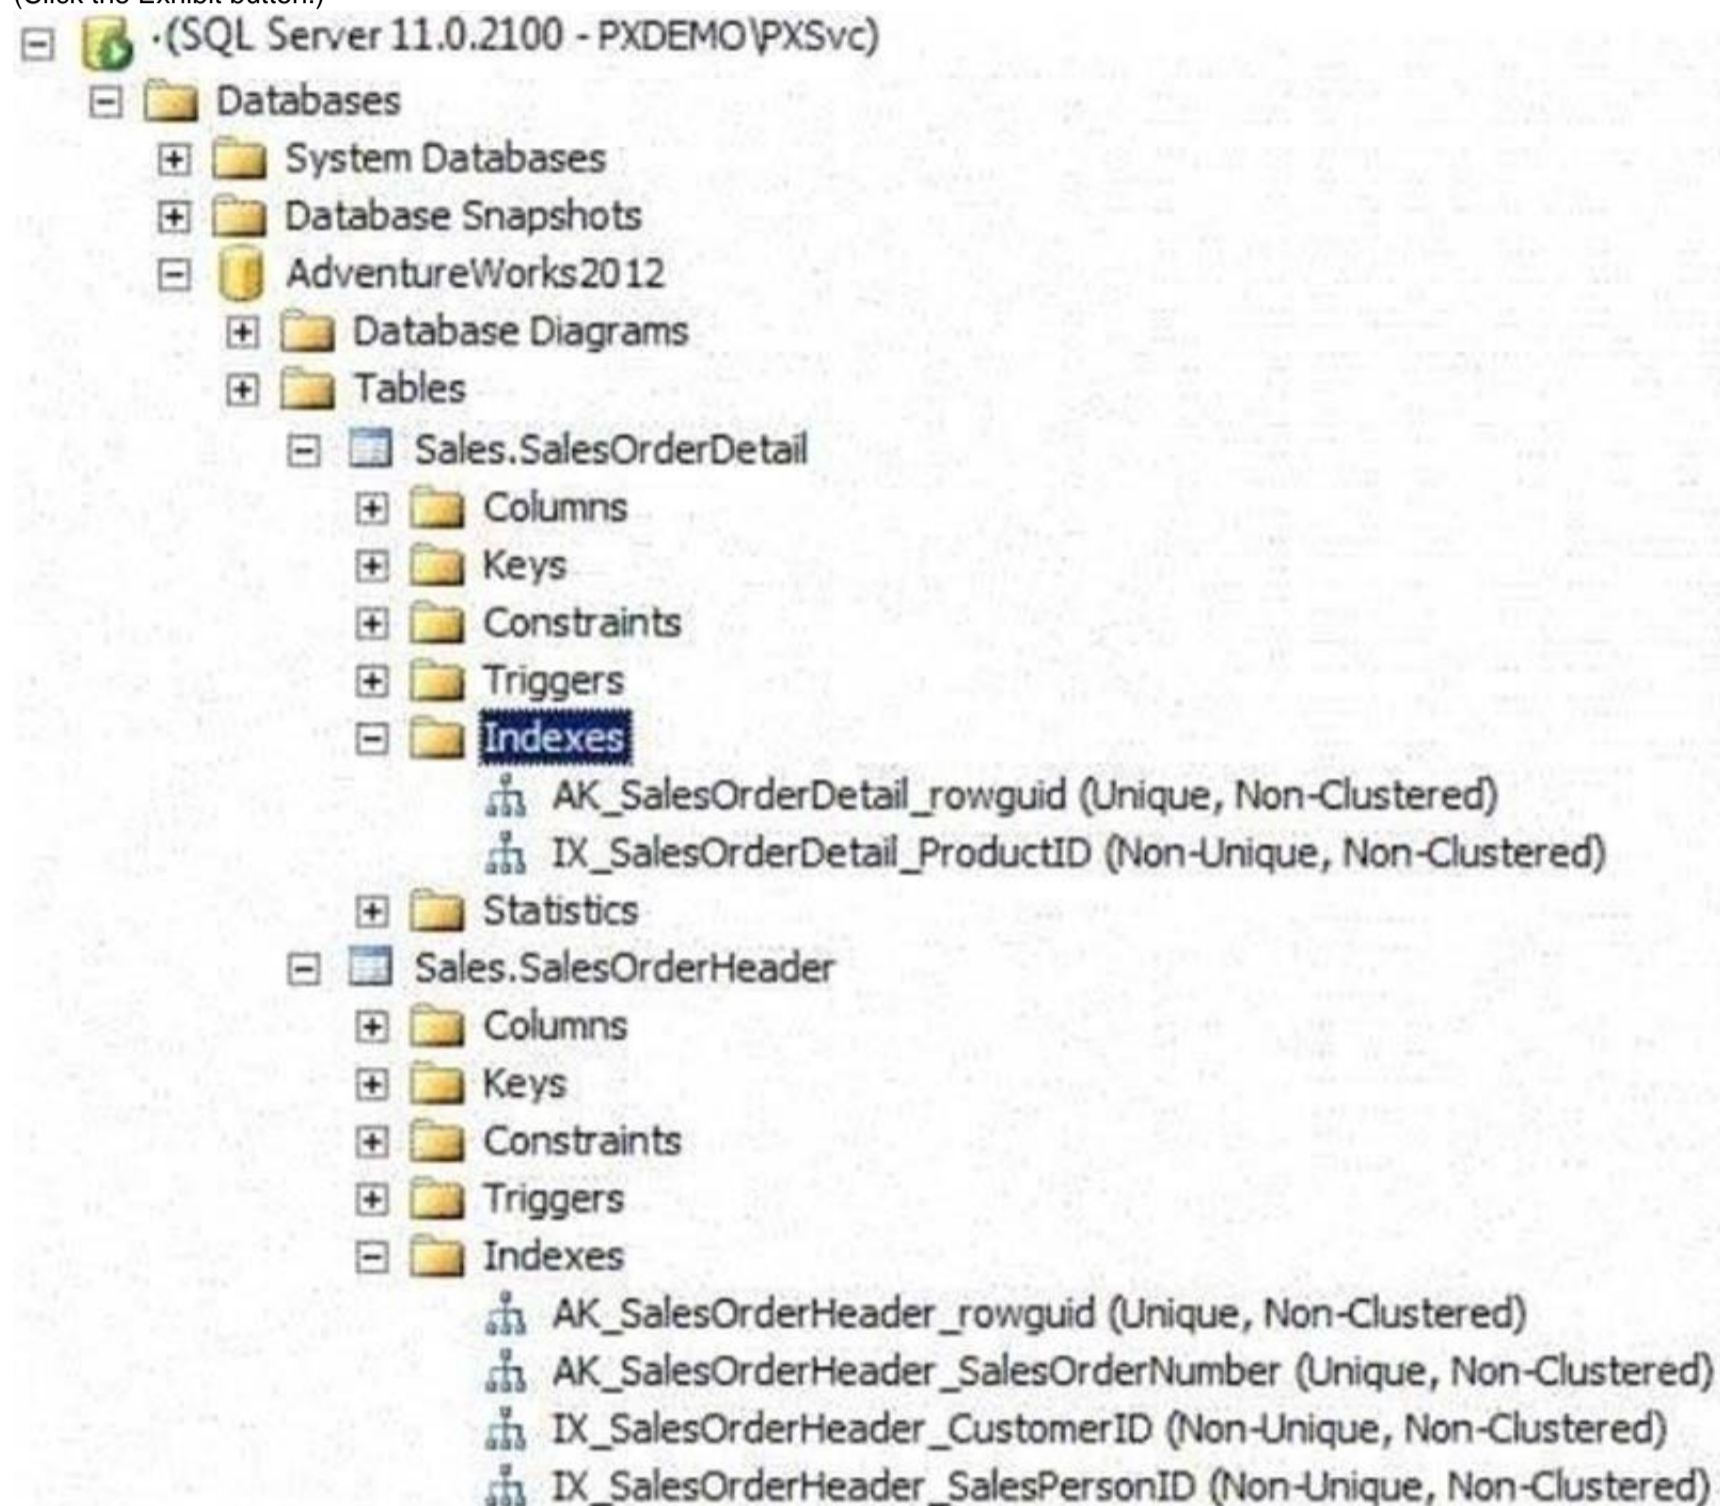
<https://docs.microsoft.com/en-us/sql/relational-databases/databases/model-database?view=sql-server-2017> <https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-transact-sql-compatibility-level?view=sql-se>

NEW QUESTION 57

- (Exam Topic 7)

You use a Microsoft SQL Server 2014 database that contains two tables named SalesOrderHeader and SalesOrderDetail. The indexes on the tables are as shown in the exhibit.

(Click the Exhibit button.)



You write the following Transact-SQL query:

```
SELECT h.SalesOrderID, h.TotalDue, d.OrderQty
FROM Sales.SalesOrderHeader AS h
    INNER JOIN Sales.SalesOrderDetail AS d
    ON h.SalesOrderID = d.SalesOrderID
WHERE h.TotalDue > 100
AND (d.OrderQty > 5 OR d.LineTotal < 1000.00);
```

You discover that the performance of the query is slow. Analysis of the query plan shows table scans where the estimated rows do not match the actual rows for SalesOrderHeader by using an unexpected index on SalesOrderDetail.

You need to improve the performance of the query. What should you do?

- A. Use a FORCESCAN hint in the query.
- B. Add a clustered index on SalesOrderID in SalesOrderHeader.
- C. Use a FORCESEEK hint in the query.
- D. Update statistics on SalesOrderID on both tables.

Answer: D

Explanation:

New statistics would be useful.

The UPDATE STATISTICS command updates query optimization statistics on a table or indexed view. By default, the query optimizer already updates statistics as necessary to improve the query plan; in some cases you can improve query performance by using UPDATE STATISTICS or the stored procedure sp_updatestats to update statistics more frequently than the default updates.

References:

<http://msdn.microsoft.com/en-us/library/ms187348.aspx>

NEW QUESTION 61

- (Exam Topic 7)

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 rewrite the queries to use aggregates instead of PIVOT statements. Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Instead you can add more files to the database.

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

NEW QUESTION 62

- (Exam Topic 7)

You plan to deploy an on-premises SQL Server 2014 database to Azure SQL Database. You have the following requirements:

Maximum database size of 500 GB

A point-in-time-restore of 35 days

Maximum database transaction units (DTUs) of 500

You need to choose the correct service tier and performance level. Which service tier should you choose?

A. Standard S3

B. Premium P4

C. Standard SO

D. Basic

Answer: B

Explanation:

You should choose Premium P4. The Premium tier is the highest Azure SQL Database tier offered. This tier is used for databases and application that require the highest level of performance and recovery. The P4 level supports a maximum of 500 DTUs, a maximum database size of 500 GB, and a point-in-time-restore to anypoint in the last 35 days.

NEW QUESTION 63

- (Exam Topic 7)

You administer a SQL Server 2014 server that contains a database named SalesDB. SalesDb contains a

schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales.

UserA is granted the Select permission on the Regions table. The Sales role is granted the Select permission on the Customers schema.

You need to ensure that the Sales role, including UserA, is disallowed to select from any of the tables in the Customers schema.

Which Transact-SQL statement should you use?

A. REVOKE SELECT ON Schema::Customers FROM UserA

B. DENY SELECT ON Object::Regions FROM UserA

C. EXEC sp_addrolemember 'Sales', 'UserA'

D. DENY SELECT ON Object::Regions FROM Sales

E. REVOKE SELECT ON Object::Regions FROM UserA

F. DENY SELECT ON Schema::Customers FROM Sales

G. DENY SELECT ON Schema::Customers FROM UserA

H. EXEC sp_droprolemember 'Sales', 'UserA'

I. REVOKE SELECT ON Object::Regions FROM Sales

J. REVOKE SELECT ON Schema::Customers FROM Sales

Answer: F

Explanation:

Use SQL Data Warehouse or Parallel Data Warehouse GRANT and DENY statements to grant or deny a permission (such as UPDATE) on a securable (such as a database, table, view, etc.) to a security principal (a login, a database user, or a database role).

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/permissions-grant-deny-revoke-azure-sql-data-warehouse->

NEW QUESTION 68

- (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 69

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database that contains a table named OrderDetail.

You discover that the NCI_OrderDetail_CustomerID non-clustered index is fragmented. You need to reduce fragmentation. You need to achieve this goal without taking the index offline.

Which Transact-SQL batch should you use?

- A. CREATE INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID WITH DROP EXISTING
- B. ALTER INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID REORGANIZE
- C. ALTER INDEX ALL ON OrderDetail REBUILD
- D. ALTER INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID REBUILD

Answer: B

Explanation:

REORGANIZE specifies to reorganize the index leaf level. The REORGANIZE operation is always performed online. This means long-term blocking table locks are not held and queries or updates to the underlying table can continue during the ALTER INDEX REORGANIZE transaction.

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

NEW QUESTION 70

- (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/azuresqldbssupport/2018/01/15/automating-azure-sql-db-index-and-statistic-m>

NEW QUESTION 74

- (Exam Topic 7)

A company runs Microsoft SQL Server 2017 in an on-premises environment. The databases are memory-optimized.

An integrity check of a database has failed.

You need to ensure that the data is healthy and passes an integrity check. What should you do?

- A. Run the checktable Transact-SQL statement.
- B. Clear the buffer of the database.
- C. Restore from a verified backup.
- D. Run the cleantable Transact-SQL statement.

Answer: C

Explanation:

To verify the integrity of the on-disk checkpoint files, perform a backup of the MEMORY_OPTIMIZED_DATA filegroup.

NEW QUESTION 77

- (Exam Topic 7)

You plan to migrate on-premises Microsoft SQL Server to SQL Server on a Microsoft Azure virtual machine. You need to ensure that the Azure virtual machine can handle the workload.

Which tool should you use for each environment? To answer, drag the appropriate tools to the correct options. Each tool may be used once. More than once, or

not at all.

| Tools, Select from these. | Answer Area |
|------------------------------|---------------------------------------|
| Distributed Replay | Tool to use on-premises: <Place here> |
| Performance Monitor | Tool to use in Azure: <Place here> |
| SQL Server Profiler | |
| SQL Server Extended Events | |
| SQL Server Data Tools (SSDT) | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

| Tools, Select from these. | Answer Area |
|------------------------------|----------------------------------------------------|
| Distributed Replay | Tool to use on-premises: SQL Server Profiler |
| Performance Monitor | Tool to use in Azure: SQL Server Data Tools (SSDT) |
| SQL Server Profiler | |
| SQL Server Extended Events | |
| SQL Server Data Tools (SSDT) | |

NEW QUESTION 78

- (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 83

- (Exam Topic 7) You have a database named DB1. You discover that DB1 is corrupt.

You run DBCC CHECKDB and receive an error message within a few seconds. No pages are listed in the error message.

You need to repair the database corruption as quickly as possible. The solution must minimize data loss.

What should you do?

- A. Run DBCC CHECKDB ('db1', REPAIR_ALLOW_DATA_LOSS).
- B. Run DBCC CHECKDB ('db1', REPAIR_FAST).
- C. Delete the transaction logs and restart the Microsoft SQL Server instance.
- D. Run DBCC CHECKDB ('db1', REPAIR_REBUILD).
- E. Restore the database from a backup.

Answer: C

Explanation:

REPAIR_REBUILD

Performs repairs that have no possibility of data loss. This can include quick repairs, such as repairing missing rows in non-clustered indexes, and more time-consuming repairs, such as rebuilding an index.

NEW QUESTION 88

- (Exam Topic 7)

You deploy a new Microsoft Azure SQL database instance to support a variety of mobile application and public websites. You configure geo-replication with regions in Brazil and Japan.

You need to implement real-time encryption of the database and all backups.

Solution: You password protect all azure SQL backups and enable azure active directory authentication for all azure SQL server instances.
Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Password protection does not encrypt the data.

Transparent Data Encryption (TDE) would provide a solution. References:

<https://azure.microsoft.com/en-us/blog/how-to-configure-azure-sql-database-geo-dr-with-azure-key-vault/>

NEW QUESTION 89

- (Exam Topic 7)

You administer two Microsoft SQL Server 2014 servers named ProdSrv1 and ProdSrv2. ProdSrv1 is configured as a Distributor. Both servers are configured to use the Windows NT Service virtual accounts for all SQL Services.

You are configuring snapshot replication from ProdSrv1 to ProdSrv2 by using ProdSrv2 as a pull subscriber. The distribution agent on ProdSrv2 regularly fails, displaying the following error message:

"Cannot access the file. Operating system error code 5 (Access is denied.)." You need to configure the distribution agent by granting only the minimum required access to all accounts.

What should you do?

- A. Configure the Subscriber to use the Local System account.
- B. Configure the SQL Server Agent service to run under the Local System account
- C. Configure the Subscriber to use the SQL Server Agent service account.
- D. Configure the SQL Server Agent service to run under a Windows domain account
- E. Configure the Subscriber to use the SQL Server Agent service account
- F. Grant FULL CONTROL access for the domain account to the Repldata share on ProdSrv1.
- G. Configure the Subscriber to use a Windows domain account
- H. Grant READ access for the domain account to the Repldata share on ProdSrv1.

Answer: D

Explanation:

Confirm that distribution agent has read privileges, full control access is not required, to the folder in question.

References:

<http://stackoverflow.com/questions/14555262/cannot-bulk-load-operating-system-error-code-5-access-is-denied>

NEW QUESTION 93

- (Exam Topic 7)

You plan to deploy a Microsoft SQL Server database that will use FILESTREAM. The database will store 4 TB of FILESTREAM data on a single Windows partition.

You need to configure the hard disk that will support the FILESTREAM data. The solution must provide the fastest read and write access to the data.

How should you configure the disk? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer area

| | |
|-----------------------|---------------------------------------------------------------------|
| File system: | <div><div></div><div>FAT32</div><div>FAT</div><div>NTFS</div></div> |
| 8.3 filename support: | <div><div></div><div>Enabled</div><div>Disabled</div></div> |
| Indexing: | <div><div></div><div>Enabled</div><div>Disabled</div></div> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

File System: NTFS

8.3 filename support: Disabled Indexing: Disabled

NTFS is required.

Disable generation of 8.3 names on all NTFS volumes used for FILESTREAM data storage.

Check that search indexing is not enabled on FILESTREAM volumes, under the Volume Properties window, unchecking the “Allow files on this drive to have contents indexed in addition to file properties” box.

References:

<https://blogs.msdn.microsoft.com/blogdoezequiel/2011/02/11/best-practices-on-filestreamimplementations/>

NEW QUESTION 95

- (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 100

- (Exam Topic 7)

You use a contained database named ContosoDb within a domain.

You need to create a user who can log on to the ContosoDb database. You also need to ensure that you can port the database to different database servers within the domain without additional user account configurations.

Which type of user should you create?

- A. SQL user without login
- B. User mapped to an asymmetric key
- C. Domain user
- D. login mapped to a virtual account

Answer: C

Explanation:

If the service must interact with network services, access domain resources like file shares or if it uses linked server connections to other computers running SQL Server, you might use a minimally-privileged domain account. Many server-to-server activities can be performed only with a domain user account.

References: <https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-windows-servic>

NEW QUESTION 103

- (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 105

- (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 migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 databased 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 create one storage account that has one container. You create multiple VHDs in the container. Does this meet the goal?

- A. Yes
- B. No

Answer: B**Explanation:**

Each Storage Account handles up to 20,000 IOPS, and 500TB of data.

References: <https://www.tech-coffee.net/understand-microsoft-azure-storage-for-virtual-machines/>

NEW QUESTION 108

- (Exam Topic 7)

You manage a Microsoft SQL Server environment in a Microsoft Azure virtual machine.

You must enable Always Encrypted for columns in a database. You need to configure the key store provider.

What should you do?

- A. Use the Randomized encryption type
- B. Modify the connection string for applications.
- C. Auto-generate a column master key.
- D. Use the Azure Key Vault.

Answer: D

Explanation:

There are two high-level categories of key stores to consider - Local Key Stores, and Centralized Key Stores.

Centralized Key Stores - serve applications on multiple computers. An example of a centralized key store is Azure Key Vault.

Local Key Stores References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/create-and-storecolumn-master-ke>

NEW QUESTION 113

- (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 116

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014.

A process that normally runs in less than 10 seconds has been running for more than an hour. You examine the application log and discover that the process is using session ID 60.

You need to find out whether the process is being blocked. Which Transact-SQL statement should you use?

- A. EXEC sp_who 60
- B. SELECT * FROM sys.dm_exec_sessions WHERE sessionid = 60
- C. EXEC sp_helpdb 60
- D. DBCC INPUTBUFFER (60)

Answer: A

Explanation:

sp_who provides information about current users, sessions, and processes in an instance of the Microsoft SQL Server Database Engine. The information can be filtered to return only those processes that are not idle, that belong to a specific user, or that belong to a specific session.

Example: Displaying a specific process identified by a session ID EXEC sp_who '10' --specifies the process_id;

References:<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-who-transact->

NEW QUESTION 117

- (Exam Topic 7)

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

You need to diagnose deadlocks that happen when executing a specific set of stored procedures by recording events and playing them back on a different test server.

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: D

Explanation:

Use SQL Server Profiler to identify the cause of a deadlock. A deadlock occurs when there is a cyclic dependency between two or more threads, or processes, for some set of resources within SQL Server. Using SQL Server Profiler, you can create a trace that records, replays, and displays deadlock events for analysis.

References:

<http://msdn.microsoft.com/en-us/library/ms188246.aspx>

NEW QUESTION 118

- (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 120

- (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.

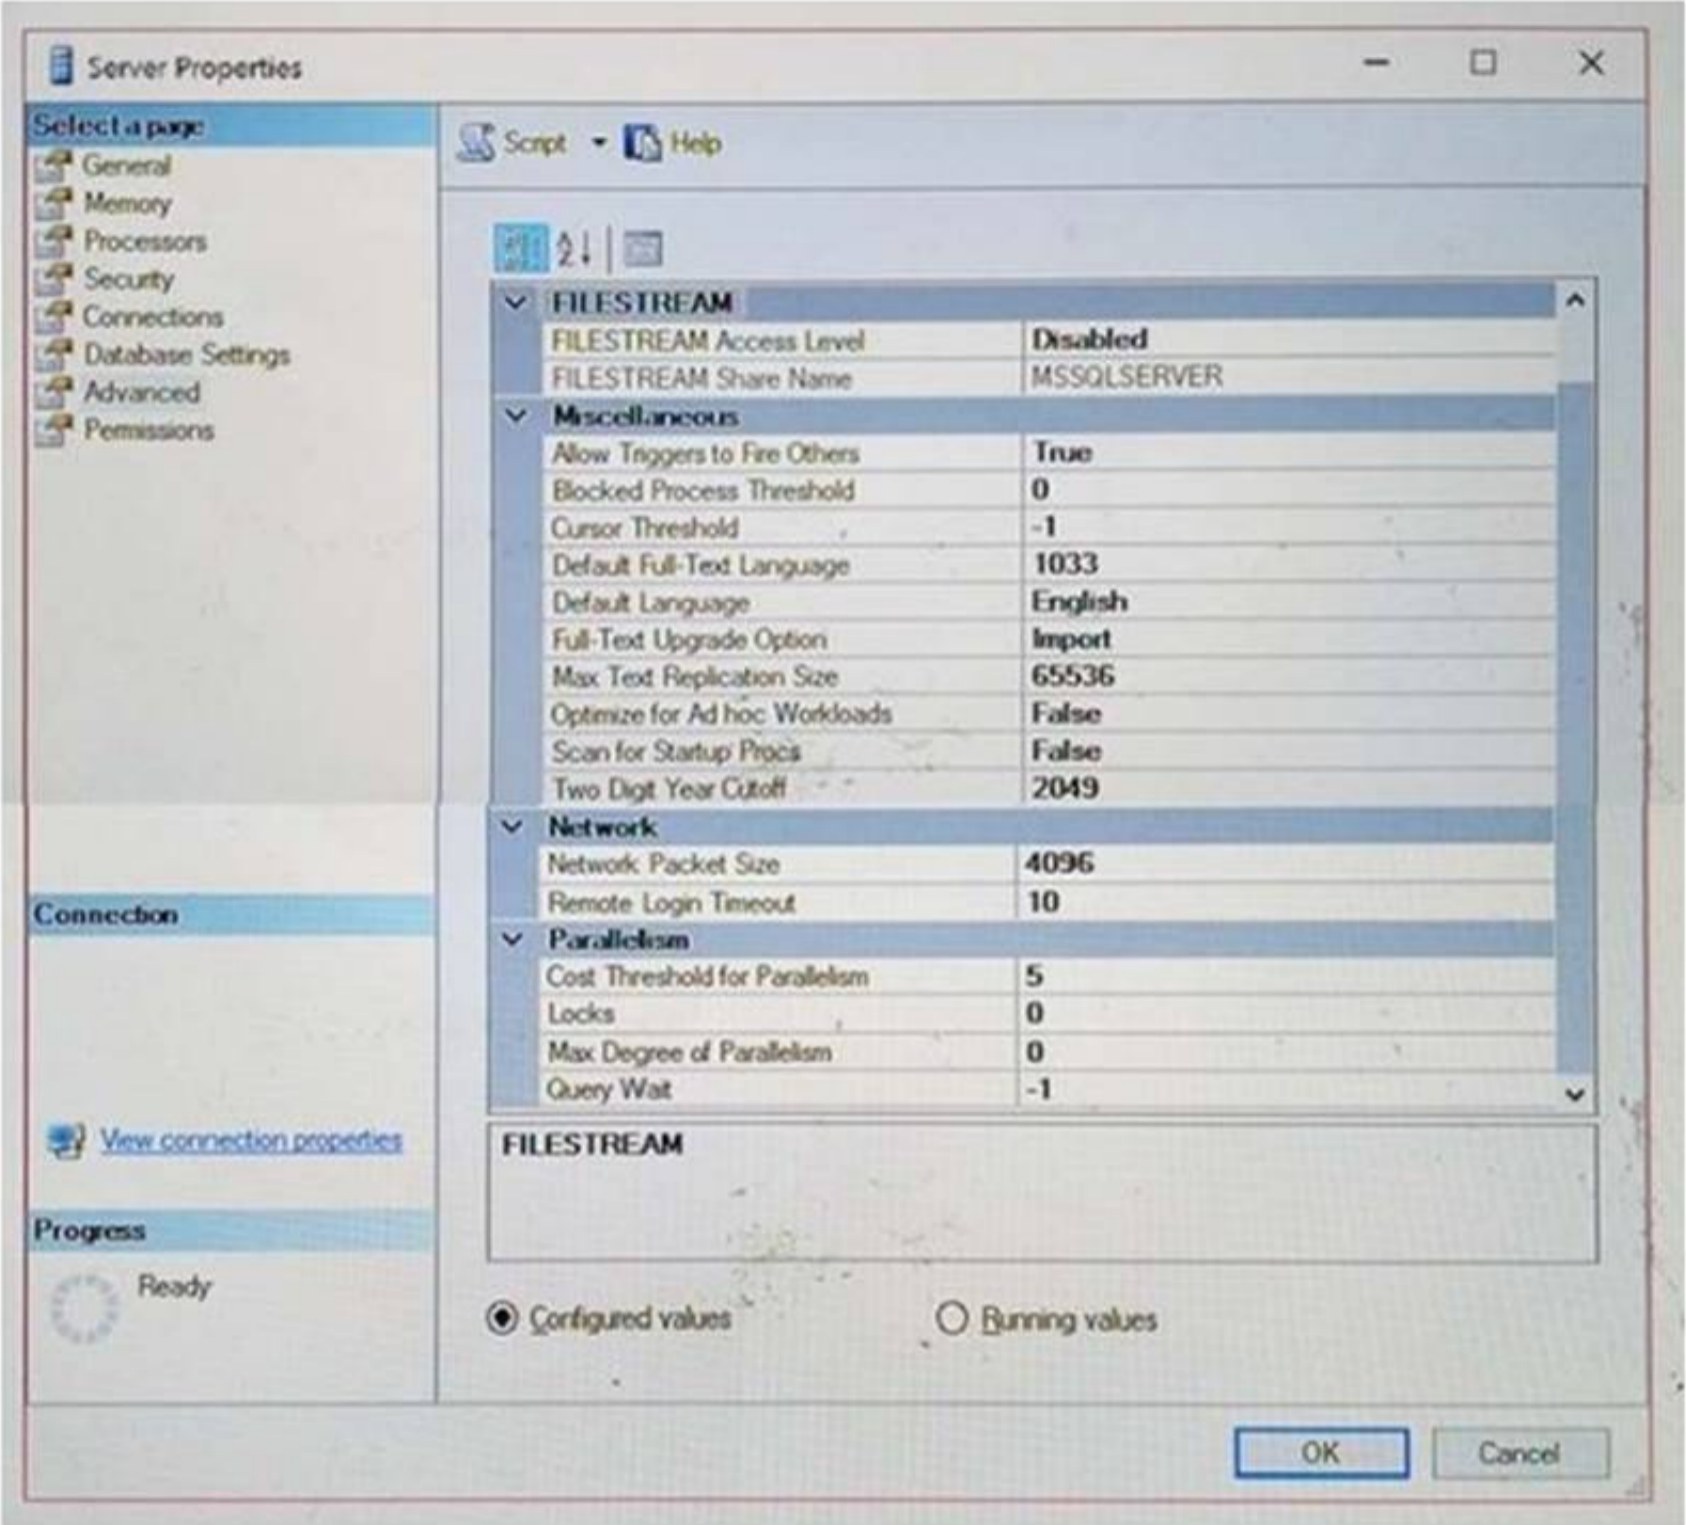
The environment also includes the following databases: DB1, DB2, and Reporting. The Reporting database is protected with Transparent Data Encryption (TDE).

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 | Answer area |
|----------------------------------------------|----------------------------------------------------------------------------------------|
| master encryption key on the master database | 1. Copy the certificate and private key backups from the old server to the new server. |
| service master key | 2. Create: <input type="text"/> |
| server certificate | 3. Restore: <input type="text"/> |
| Reporting database .mdf file | 4. Attach the Reporting database. |
| master key password | |

- 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 125

- (Exam Topic 7)

You plan to migrate a Microsoft sql server instance between physical servers. You must migrate the metadata associated with the database instance.

You need to ensure that the new instance retains the existing jobs and alerts. Solutions: You restore the model database.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

The model database does not handle alerts and jobs. It is used as the template for all databases created on an instance of SQL Server.

The msdb database is used by SQL Server Agent for scheduling alerts and jobs and by other features such as SQL Server Management Studio, Service Broker and Database Mail.

References:

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

NEW QUESTION 130

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft Azure Virtual machine that has a 4-TB database.

You plan to configure daily backups for the database. A single full backup will be approximately 1.5 TB of compressed data.

You need to ensure that the last backups are retained. Where should you store the daily backups?

A. Local storage

B. Page blob storage

C. Virtual disks

D. Block blob storage.

Answer: D

Explanation:

When backing up to Microsoft Azure blob storage, SQL Server 2016 supports backing up to multiple blobs to enable backing up large databases, up to a maximum of 12.8 TB. This is done through Block Blobs.

References:

NEW QUESTION 135

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance that has several SQL Server Agent jobs configured. When SQL Server Agent jobs fail, the error messages returned by the job steps do not provide the required detail.

The following error message is an example error message:

"The job failed. The Job was invoked by User CONTOSO\ServiceAccount. The last step to run was step 1 (Subplan_1)."

You need to ensure that all available details of the job step failures for SQL Server Agent jobs are retained. What should you do?

A. Configure output files.

B. Expand agent logging to include information from all events.

C. Disable the Limit size of job history log feature.

D. Configure event forwarding.

Answer: B

Explanation:

References:

<http://msdn.microsoft.com/en-us/library/ms175488.aspx>

NEW QUESTION 140

- (Exam Topic 7)

You use a Microsoft Azure SQL database as a data warehouse. The database is in the Standard service tier and has 400 elastic database throughput units (eDTUs).

You load data to the database by using Azure Data Factory. You need to reduce the amount of time it takes to load the data.

Solution: You move the database to a Standard database pool that has 800 eDTUs. Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:

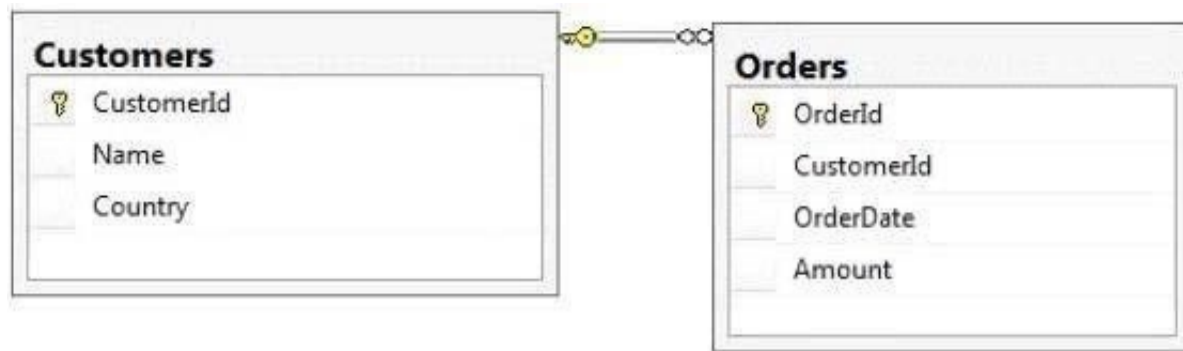
We need at least 400 eDTUs and the use of a Standard database pool.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dtu-resource-limits>

NEW QUESTION 143

- (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
- F. SELECT Name, Country, OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 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: G

NEW QUESTION 146

- (Exam Topic 7)

You administer a SQL 2012 server that contains a database named SalesDB. SalesDb contains a schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales.

UserA is granted the Select permission on the Regions table. The Sales role is granted the Select permission on the Customers schema.

You need to remove the Select permission for UserA on the Regions table. You also need to ensure that UserA can still access all the tables in the Customers schema, including the Regions table, through the Sales role permissions.

Which Transact-SQL statement should you use?

- A. REVOKE SELECT ON Schema::Customers FROM UserA
- B. DENY SELECT ON Object::Regions FROM UserA
- C. EXEC sp_addrolemember 'Sales', 'UserA'
- D. DENY SELECT ON Object::Regions FROM Sales
- E. REVOKE SELECT ON Object::Regions FROM UserA
- F. DENY SELECT ON Schema::Customers FROM Sales
- G. DENY SELECT ON Schema::Customers FROM UserA
- H. EXEC sp_droprolemember 'Sales', 'UserA'
- I. REVOKE SELECT ON Object::Regions FROM Sales
- J. REVOKE SELECT ON Schema::Customers FROM Sales

Answer: E

Explanation:

Use REVOKE to remove the grant or deny of a permission.

References: [https://docs.microsoft.com/en-us/sql/t-sql/statements/permissions-grant-deny-revoke-azure-sql-data-](https://docs.microsoft.com/en-us/sql/t-sql/statements/permissions-grant-deny-revoke-azure-sql-data-warehouse)

NEW QUESTION 149

- (Exam Topic 7)

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 log files to DB1.
Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

This problem is related to the disk response time, not to the number of log files.

References:

<https://www.mssqltips.com/sqlservertip/4131/troubleshooting-sql-server-transaction-log-related-wait-types/>

NEW QUESTION 152

- (Exam Topic 7)

You manage an on-premises, multi-tier application that has the following configuration:

Two SQL Server 2012 databases named SQL1 and SQL2

Two application servers named AppServer1 and AppServer2 that run IIS You plan to move your application to Azure.

You need to ensure that during an Azure update cycle or a hardware failure, the application remains available.

Which two deployment configurations should you implement? Each correct answer presents part of the solution.

- A. Deploy AppServer1 and AppServer2 in a single availability set.
- B. Deploy all servers in a single availability set.
- C. Deploy SQL1 and AppServer1 in a single availability set.
- D. Deploy SQL2 and AppServer2 in a single availability set.
- E. Deploy SQL1 and SQL2 in a single availability set.

Answer: AE

Explanation:

You should deploy AppServer1 and AppServer2 in a single availability set. You should deploy SQL1 and SQL2 in a single availability set.

Note: Using availability sets allows you to build in redundancy for your Azure services. By grouping related virtual machines and services (tiers) into an availability set (in this case, deploying both of your databases into an availability set), you ensure that if there is a planned or unplanned outage, your services will remain available. At the most basic level, virtual machines in an availability set are put into a different fault domain and update domain. An update domain allows virtual machines to have updates installed and then the virtual machines are rebooted together.

If you have two virtual machines in an availability set, each in its own update domain, a rebooting of one server does not bring down all of the servers in a given tier. A fault domain operates in the same manner, so if there is a physical problem with a server, rack, network, or other service, both machines are separated, and services will continue.

NEW QUESTION 153

- (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 156

- (Exam Topic 7)

You manage a Microsoft SQL Server instance named SQL1 that has 32 gigabytes (GB) of total memory. The instance supports an app named App1 that only uses a single thread. App1 frequently queries the database using the same index. The operating system and App1 combined require 8 GB of memory to function.

You need to ensure that the SQL Server does not limit the performance of App1. What configuration option should you set?

- A. min memory per query to 4 GB
- B. index create memory to 16 GB
- C. max worker threads to 1
- D. max server memory to 16 GB

Answer: B

Explanation:

The index creates memory option controls the maximum amount of memory initially allocated for sort operations when creating indexes. The default value for this option is 0 (self-configuring). If more memory is later needed for index creation and the memory is available, the server will use it; thereby, exceeding the setting of this option. If additional memory is not available, the index creation will continue using the memory already allocated.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-the-indexcreate-memory-ser>

NEW QUESTION 161

- (Exam Topic 7)

A company has an on-premises Microsoft SQL Server 2017 infrastructure. The storage area network (SAN) that supports the SQL infrastructure has reached maximum capacity.

You need to recommend a solution to reduce on-premises storage use without changing the application. What should you do?

- A. Configure an Express Route connection to Microsoft Azure.
- B. Configure a Microsoft Azure Key Vault.
- C. Configure geo-replication on the SAN.
- D. Configure SQL Server Stretch Database in Microsoft Azure.

Answer: D

Explanation:

Stretch warm and cold transactional data dynamically from SQL Server to Microsoft Azure with SQL Server Stretch Database. Unlike typical cold data storage, your data is always online and available to query. Benefit from the low cost of Azure rather than scaling expensive, on-premises storage.

References:

<https://docs.microsoft.com/en-us/sql/sql-server/stretch-database/stretch-database?view=sql-server-2017>

NEW QUESTION 164

- (Exam Topic 7)

You have a database named DB1 that contains a table named Table1. Table1 has a non-clustered index named index1.

You discover that index1 is corrupt. You need to repair index1.

Which statement should you execute?

- A. DBCC CHECKDB ('db1', REPAIR_FAST)
- B. ALTER INDEX idx1 ON table1 REBUILD WITH (ONLINE=ON)
- C. ALTER INDEX index1 ON table1 REORGANIZE
- D. DBCC CHECKDB ('db1', DATA_PURITY)

Answer: B

Explanation:

If REBUILD is performed online (ON) the data in this table is available for queries and data modification during the index operation.

NEW QUESTION 165

- (Exam Topic 7)

You are a database administrator for a Microsoft SQL Server 2014 environment.

You want to deploy a new application that will scale out the workload to at least five different SQL Server instances.

You need to ensure that for each copy of the database, users are able to read and write data that will then be synchronized between all of the database instances.

Which feature should you use?

- A. Database Mirroring
- B. Peer-to-Peer Replication
- C. Log Shipping
- D. Availability Groups

Answer: B

Explanation:

Peer-to-peer replication provides a scale-out and high-availability solution by maintaining copies of data across multiple server instances, also referred to as nodes. Built on the foundation of transactional replication, peer-to-peer replication propagates transactionally consistent changes in near real-time. This enables applications that require scale-out of read operations to distribute the reads from clients across multiple nodes. Because data is maintained across the nodes in near real-time, peer-to-peer replication provides data redundancy, which increases the availability of data.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/replication/transactional/peer-to-peer-trans>

NEW QUESTION 166

- (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 169

- (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 172

- (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 175

- (Exam Topic 7)

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

You need to track all SELECT statements issued in the Contoso database only by users in a role named Sales. 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: F

Explanation:

To audit users in a role use a Database Audit Specification.

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-audit-specification-transact-sql>

NEW QUESTION 176

- (Exam Topic 7)

You administer two Microsoft SQL Server 2014 servers. Each server resides in a different, untrusted domain. You plan to configure database mirroring.

You need to be able to create database mirroring endpoints on both servers. What should you do?

- A. Configure the SQL Server service account to use Network Service.
- B. Use a server certificate.
- C. Use a database certificate.
- D. Configure the SQL Server service account to use Local System.

Answer: B

Explanation:

To enable certificate authentication for database mirroring on a given server instance, the system administrator must configure each server instance to use certificates on both outbound and inbound connections.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/database-mirroring/use-certificates-for-a-database-mirrorin>

NEW QUESTION 177

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance.

The instance contains a database that supports a retail sales application. The application generates hundreds of transactions per second and is online 24 hours per day and 7 days per week.

You plan to define a backup strategy for the database. You need to ensure that the following requirements are met:

No more than 5 minutes worth of transactions are lost. Data can be recovered by using the minimum amount of administrative effort.

What should you do? Choose all that apply.

- A. Configure the database to use the SIMPLE recovery model.
- B. Create a DIFFERENTIAL database backup every 4 hours.
- C. Create a LOG backup every 5 minutes.
- D. Configure the database to use the FULL recovery model.
- E. Create a FULL database backup every 24 hours.
- F. Create a DIFFERENTIAL database backup every 24 hours.

Answer: BCDE

Explanation:

The full recovery model uses log backups to prevent data loss in the broadest range of failure scenarios, and backing and restoring the transaction log (log backups) is required. The advantage of using log backups is that they let you restore a database to any point of time that is contained within a log backup (point-in-time

recovery). You can use a series of log backups to roll a database forward to any point in time that is contained in one of the log backups. Be aware that to minimize your restore time, you can supplement each full backup with a series of differential backups of the same data.

References: [https://technet.microsoft.com/en-us/library/ms190217\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms190217(v=sql.105).aspx)

NEW QUESTION 178

- (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 183

- (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 184

- (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 185

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance that has multiple databases. You have a two-node SQL Server failover cluster. The cluster uses a storage area network (SAN). You discover I/O issues. The SAN is at capacity and additional disks cannot be added.

You need to reduce the I/O workload on the SAN at a minimal cost. What should you do?

A. Move user databases to a local disk.

B. Expand the tempdb data and log files

C. Modify application code to use table variables

D. Move the tempdb files to a local disk

Answer: D

Explanation:

The use of local disks for TempDB allows us to have more flexibility when configuring for optimal performance. It is a common performance recommendation to create the TempDB database on the fastest storage available. With the capability to utilize local disk for TempDB placement we can easily utilize disks that are larger, have a higher rotational speed or use SSD disks.

References: <https://www.mssqltips.com/sqlservertip/2817/sql-server-2012-cluster-with-tempdb-on-local-disk/>

NEW QUESTION 190

- (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 192

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database. The database is currently configured to log ship to a secondary server.

You are preparing to cut over to the secondary server by stopping log-shipping and bringing the secondary database online. You want to perform a tail-log backup.

You need to leave the primary database in a restoring state.

Which option of the BACKUP LOG command should you use?

A. NO_TRUNCATE

B. NORECOVERY

C. STANDBY

D. FORMAT

Answer: B

Explanation:

It is recommended that you take a tail-log backup in the following scenarios:

* If the database is online and you plan to perform a restore operation on the database, begin by backing up the tail of the log. To avoid an error for an online

database, you must use the ... WITH NORECOVERY option of the BACKUP Transact-SQL statement.

Note: A tail-log backup captures any log records that have not yet been backed up (the tail of the log) to prevent work loss and to keep the log chain intact. Before you can recover a SQL Server database to its latest point in time, you must back up the tail of its transaction log. The tail-log backup will be the last backup of interest in the recovery plan for the database.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/tail-log-backups-sql-server>

NEW QUESTION 193

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database instance. You create a new user named UserA. You need to ensure that UserA is able to create SQL Server Agent jobs and execute SQL Server agent jobs owned by UserA

To which role should you add UserA?

- A. DatabaseMailUserRole
- B. ServerGroupAdministratorGroup
- C. SQLAgentUserRole
- D. Securityadmin

Answer: C

Explanation:

SQLAgentUserRole is the least privileged of the SQL Server Agent fixed database roles. It has permissions on only operators, local jobs, and job schedules. Members of SQLAgentUserRole have permissions on only local jobs and job schedules that they own. Members can create local jobs.

References: <https://docs.microsoft.com/en-us/sql/ssms/agent/sql-server-agent-fixed-database-roles>

NEW QUESTION 195

- (Exam Topic 7)

You plan to create an AlwaysOn availability group that will have two replicas in Microsoft Azure and two on premises replicas.

You need to configure the network to support the availability group listener. 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: B

Explanation:

An availability group listener is a virtual network name that clients connect to for database access. On Azure virtual machines, a load balancer holds the IP address for the listener. The load balancer routes traffic to the instance of SQL Server that is listening on the probe port. Usually, an availability group uses an internal load balancer.

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windowsportal-sql-ps-al>

NEW QUESTION 199

- (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: Move the transaction logs to a faster disk. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Section: Deploy and migrate applications

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 204

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 70-765 Exam with Our Prep Materials Via below:

<https://www.certleader.com/70-765-dumps.html>