



Microsoft

Exam Questions 70-765

Provisioning SQL Databases (beta)

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 a virtual machine (VM) in Microsoft Azure, which has a 2 terabyte (TB) database. Microsoft SQL Server backups are performed by using Backup to URL. You need to provision the storage account for the backups while minimizing costs. Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 disk storage
- 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: G

Explanation:

A URL specifies a Uniform Resource Identifier (URI) to a unique backup file. The URL is used to provide the location and name of the SQL Server backup file. The URL must point to an actual blob, not just a container. If the blob does not exist, it is created. If an existing blob is specified, BACKUP fails, unless the "WITH FORMAT" option is specified to overwrite the existing backup file in the blob.

LOCALLY REDUNDANT STORAGE (LRS) makes multiple synchronous copies of your data within a single datacenter.

NEW QUESTION 2

- (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 a GS-series virtual machine (VM) in Microsoft Azure. You plan to deploy Microsoft SQL Server. You need to deploy a 30 megabyte (MB) database that requires 100 IOPS to be guaranteed while minimizing costs. Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 disk storage
- 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: A

Explanation:

Premium Storage Disks Limits

When you provision a disk against a Premium Storage account, how much input/output operations per second (IOPS) and throughput (bandwidth) it can get depends on the size of the disk. Currently, there are three types of Premium Storage disks: P10, P20, and P30. Each one has specific limits for IOPS and throughput as specified in the following table:

Premium Storage Disk Type	P10	P20	P30
Disk Size	128 GiB	512 GiB	1024 GiB (1 TB)
IOPS per disk	500	2300	5000
Throughput per disk	100 MB per second	150 MB per second	200 MB per second

References:<https://docs.microsoft.com/en-us/azure/storage/storage-premium-storage>

NEW QUESTION 3

- (Topic 1)

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.

Your company plans to use Microsoft Azure Resource Manager templates for all future deployments of SQL Server on Azure virtual machines. You need to create the templates.

Solution: You create the desired SQL Server configuration in an Azure Resource Group, then export the Resource Group template and save it to the Templates Library.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Azure Resource Manager template consists of JSON, and expressions that you can use to construct values for your deployment.

A good JSON editor, not a Resource Group template, can simplify the task of creating templates.

Note: In its simplest structure, a Azure Resource Manager template contains the following elements:

```
{
"$schema": "http://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
"contentVersion": "", "parameters": {},
"variables": {},
"resources": [],
"outputs": {}
}
```

References: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-authoring-templates>

NEW QUESTION 4

- (Topic 1)

You plan to migrate a database To Microsoft Azure SQL Database. The database requires 500 gigabytes (GB) of storage. The database must support 50 concurrent logins. You must minimize the cost associated with hosting the database. You need to create the database. Which pricing tier should you use?

- A. Standard S3 pricing tier
- B. Premium P2tier
- C. Standard S2 pricing tier
- D. Premium P1 tier

Answer: D

Explanation:

For a database size of 500 GB the Premium tier is required. Both P1 and P2 are adequate. P1 is preferred as it is cheaper.

Note:

Premium service tier

Service tier	P1	P2	P4	P6	P11	P15
Max DTUs	125	250	500	1000	1750	4000
Max database size*	500 GB	500 GB	500 GB	500 GB	1 TB	1 TB
Max in-memory OLTP storage	1 GB	2 GB	4 GB	8 GB	14 GB	32 GB
Max concurrent workers	200	400	800	1600	2400	6400
Max concurrent logins	200	400	800	1600	2400	6400
Max concurrent sessions	30000	30000	30000	30000	30000	30000

NEW QUESTION 5

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: <code>EXEC sp_addrolemember 'db_datareader', 'User1'</code>	
Run the following Transact_SQL statement: <code>CREATE LOGIN User1 WITH password='Pa\$\$w0rd'</code>	
Run the following Transact_SQL statement: <code>CREATE USER User1 WITH password='Pa\$\$w0rd'</code>	
Run the following Transact_SQL statements: <code>EXEC sp_migrate_user_to_contained @username = N'User1', @rename = N'keep_name', @disablelogin = N'disable_login'</code>	
Run the following Transact_SQL statement: <code>CREATE LOGIN User1 FROM EXTERNAL PROVIDER</code>	
Select the Reporting database and run the following Transact-SQL statements: <code>CREATE USER User1 from LOGIN User1 GRANT SELECT TO User1</code>	

- 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 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 QUERY_OPTIMIZER_HOTFIXES option for the databases. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

QUERY_OPTIMIZER_HOTFIXES = { ON | OFF | PRIMARY } enables or disables query optimization hotfixes regardless of the compatibility level of the database. This is equivalent to Trace Flag 4199.

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

NEW QUESTION 7

- (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 8

- (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 FileSystem
- 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 9

DRAG DROP - (Topic 2)

A new Azure Active Directory security principal named ReportUser@contoso.onmicrosoft.com should have access to select all current and future objects in the Reporting database. You should not grant the principal any other

permissions. You should use your Active Directory Domain Services (AD DS) account to authenticate to the Azure SQL 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

Create a connection to the **master** database on the Azure SQL Server instance by using your Active Directory authenticated account.

Create a connection to the **Reporting** database on the Azure SQL Server instance by using your Active Directory authenticated account.

Run the following Transact-SQL statement:

```
EXEC sp_addrolemember 'db_datareader', 'reportuser@contoso.onmicrosoft.com'
```

Run the following Transact-SQL statement:

```
CREATE USER [reportuser@contoso.onmicrosoft.com] FROM EXTERNAL PROVIDER
```

Run the following Transact-SQL statements:

```
USE Reporting
CREATE USER [reportuser@contoso.onmicrosoft.com] FOR LOGIN [reportuser@contoso.onmicrosoft.com]
GRANT SELECT TO [reportuser@contoso.onmicrosoft.com]
```

Create a connection to the **Reporting** database on the Azure SQL Server instance by using your SQL Server authenticated account.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1:
 To provision an Azure AD-based contained database user (other than the server administrator that owns the database), connect to the database (here the Reporting database) with an Azure AD identity (not with a SQL Server account) that has access to the database.

Step 2: CREATE USER ... FROM EXTERNAL PROVIDER
 To create an Azure AD-based contained database user (other than the server administrator that owns the database), connect to the database with an Azure AD identity, as a user with at least the ALTER ANY USER permission. Then use the following Transact-SQL syntax:
 CREATE USER <Azure_AD_principal_name> FROM EXTERNAL PROVIDER;

Step 3:
 Grant the proper reading permissions.
 References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-aad-authentication>

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 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 10

HOTSPOT - (Topic 6)

You need to set up the service accounts that the database engine and SQL Server Agent services will use. How should you design the solution? To answer, select the appropriate configuration options in the answer area.

Answer Area

Design option	Configuration option			
Account type	<table border="1" style="width: 100%;"> <tr><td>Domain account</td></tr> <tr><td>Local machine account</td></tr> <tr><td>Local system account</td></tr> </table>	Domain account	Local machine account	Local system account
Domain account				
Local machine account				
Local system account				
Group membership	<table border="1" style="width: 100%;"> <tr><td>Domain administrators</td></tr> <tr><td>Local administrators</td></tr> <tr><td>Domain users</td></tr> </table>	Domain administrators	Local administrators	Domain users
Domain administrators				
Local administrators				
Domain users				
Password management	<table border="1" style="width: 100%;"> <tr><td>Manually-managed passwords</td></tr> <tr><td>Managed Service accounts</td></tr> </table>	Manually-managed passwords	Managed Service accounts	
Manually-managed passwords				
Managed Service accounts				

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Domain Account
 The service startup account defines the Microsoft Windows account in which SQL Server Agent runs and its network permissions. SQL Server Agent runs as a specified user account. You select an account for the SQL Server Agent service by using SQL Server Configuration Manager, where you can choose from the following options:

* Built-in account. You can choose from a list of the following built-in Windows service accounts: Local System account.

* This account. Lets you specify the Windows domain account in which the SQL Server Agent service runs.

Box2: Domain users

Microsoft recommends choosing a Windows user account that is not a member of the Windows Administrators group.

Box 3: Managed Service Accounts

When resources external to the SQL Server computer are needed, Microsoft recommends using a Managed Service Account (MSA), configured with the minimum privileges necessary.

Note: A Managed Service Account (MSA) can run services on a computer in a secure and easy to maintain manner, while maintaining the capability to connect to network resources as a specific user principal.

References:<https://msdn.microsoft.com/en-us/library/ms191543.aspx>

NEW QUESTION 12

- (Exam Topic 7)

You have a server named server1-contoso.database.windows.net that has a Microsoft Azure SQL database.

You need to create a group of Azure SQL databases that share the resources on server1-contoso.database.windows.net.

Which cmdlet should you run before you create the database?

- 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: D

Explanation:

The New-AzureRmSqlElasticPool cmdlet creates an elastic database pool for an Azure SQL Database.

SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single Azure SQL Database server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

NEW QUESTION 13

- (Exam Topic 7)

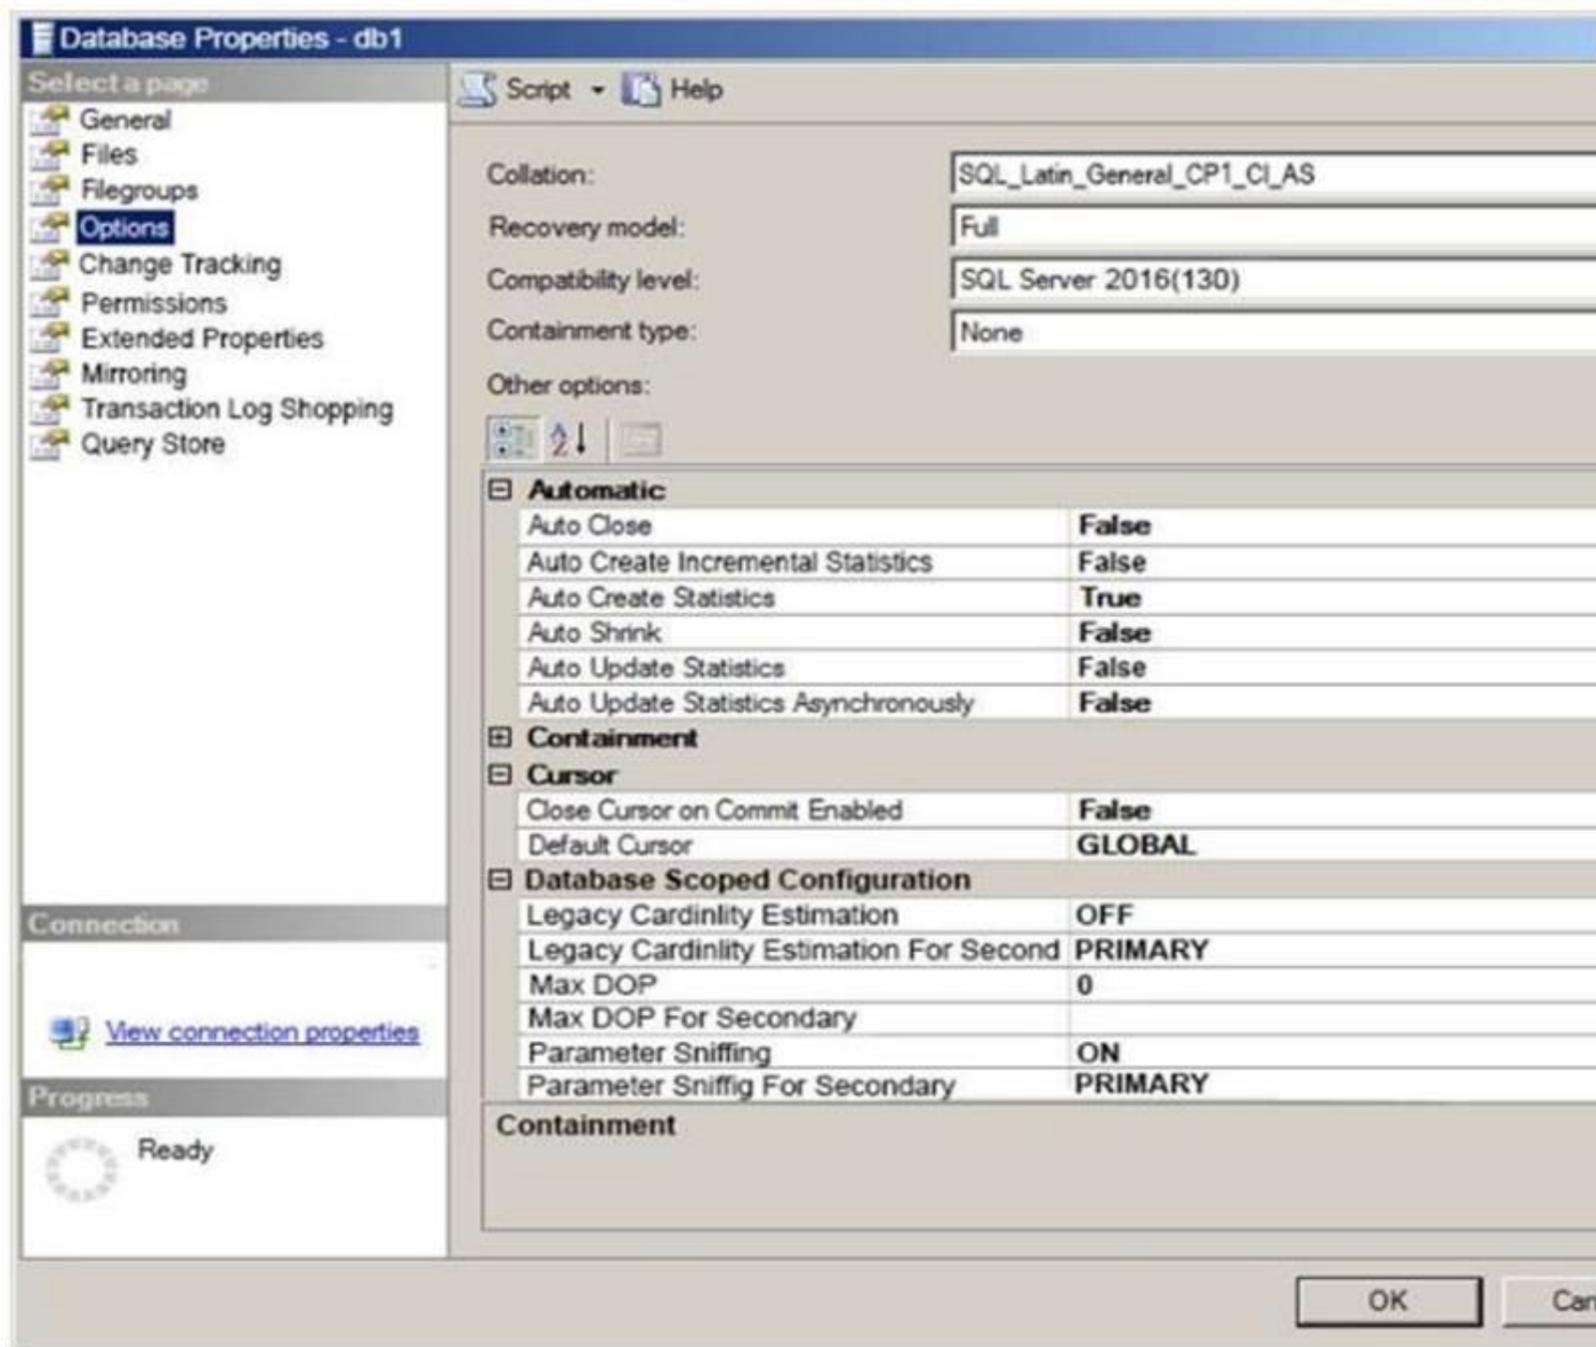
You have Microsoft SQL Server on a Microsoft Azure virtual machine. The virtual machine has a database named DB1. DB1 contains a table named Table1 that has 4 billion rows.

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 18

- (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 22

- (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 ONAccountTransaction.AccountCode REORGANIZE
- B. ALTER INDEX ALL ON AccountTransaction REBUILD
- C. ALTER INDEX IDX_AccountTransaction_AccountCode ONAccountTransaction.AccountCode REBUILD
- D. CREATE INDEX IDXAccountTransactionAccountCode ONAccountTransaction.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 27

- (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 28

- (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 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 create one storage account that has 30 containers. You create a VHD in each 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 30

- (Exam Topic 7)

You administer a SQL Server 2014 database instance.

You need to configure the SQL Server Database Engine service on a failover cluster. Which user account should you use?

- A. A domain user
- B. The BUILTIN\SYSTEM account
- C. A local user with Run as Service permissions
- D. The SQLBrowser account

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/install/create-a-new-sql-server-failover-cluster-s>

NEW QUESTION 32

- (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 35

- (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 40

- (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 41

- (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 44

- (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 49

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 failover cluster.

You need to ensure that a failover occurs when the server diagnostics returns query_processing error. Which server configuration property should you set?

- A. SqlOumperDumpFlags
- B. FailureConditionLevel
- C. HealthCheckTimeout
- D. SqlDumperDumpPath

Answer: B

Explanation:

Use the FailureConditionLevel property to set the conditions for the Always On Failover Cluster Instance (FCI) to fail over or restart.

The failure conditions are set on an increasing scale. For levels 1-5, each level includes all the conditions from the previous levels in addition to its own conditions.

Note: The system stored procedure sp_server_diagnostics periodically collects component diagnostics on the SQL instance. The diagnostic information that is collected is surfaced as a row for each of the following components and passed to the calling thread.

The system, resource, and query process components are used for failure detection. The io_subsystem and events components are used for diagnostic purposes only.

References: <https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/configure-failurecondition>

NEW QUESTION 52

- (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 56

- (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 61

- (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 63

- (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 66

- (Exam Topic 7)

You plan to deploy Microsoft SQL Server on a Microsoft Azure Virtual machine. The virtual machine will have a 30-TB database and will have 10 1-TB VHDs for the database.

You need to configure the storage to meet the following requirements:

Evenly distribute read and write operations across the VHDs.

Minimize the read and write time.

Which storage configuration should you use?

- A. a parity storage pool
- B. a simple storage pool
- C. a mirrored storage pool
- D. a striped volume
- E. a RAID-5 volume

Answer: D

Explanation:

Data that is written to a striped volume is interleaved to all disks at the same time instead of sequentially. Therefore, disk performance is the fastest on a RAID 0 volume as compared to any other type of disk configuration.

Reference:

<https://support.microsoft.com/en-us/help/323433/how-to-establish-a-striped-volume-raid-0-inwindows-server-20>

NEW QUESTION 71

- (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 the backup size is as small as possible. Which backup should you perform every two hours?

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

Answer: J

Explanation:

Minimally, you must have created at least one full backup before you can create any log backups. After that, the transaction log can be backed up at any time unless the log is already being backed up.

References:

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

NEW QUESTION 75

- (Exam Topic 7)

You administer two instances of Microsoft SQL Server 2014. You deploy an application that uses a database on the named instance.

The application is unable to connect to the database on the named instance. You need to ensure that the application can connect to the named instance. What should you do?

- A. Configure the application as data-tiered.
- B. Open port 1433 on the Windows firewall on the server.
- C. Configure the named SQL Server instance to use an account that is a member of the Domain Admins group.
- D. Start the SQL Server Browser Service.

Answer: D

Explanation:

The SQL Server Browser program runs as a Windows service. SQL Server Browser listens for incoming requests for Microsoft SQL Server resources and provides information about SQL Server instances installed on the computer. SQL Server Browser contributes to the following actions:

References: [https://technet.microsoft.com/en-us/library/ms181087\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms181087(v=sql.105).aspx)

NEW QUESTION 78

- (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 80

- (Exam Topic 7)

You are the administrator of a Microsoft SQL Server 2014 server.

Some applications consume significant resources. You need to manage the server workload by restricting resource-intensive applications

You need to dynamically limit resource consumption. What should you do?

- A. Configure Resource Pools, Workload Groups, and Classifier Function, and then enable the Resource Governor
- B. Set up Service Broker to ensure that application are not allowed to consume more than the specified amount of resource

- C. Create a new rule for each application that sets the resource limit allowed
- D. Create a new plan Guide with a Scope Type of sql and define the resource limits for each application

Answer: A

Explanation:

In the SQL Server Resource Governor, a resource pool represents a subset of the physical resources of an instance of the Database Engine. Resource Governor enables you to specify limits on the amount of CPU, physical IO, and memory that incoming application requests can use within the resource pool. Each resource pool can contain one or more workload groups. When a session is started, the Resource Governor classifier assigns the session to a specific workload group, and the session must run using the resources assigned to the workload group.

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

NEW QUESTION 84

- (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:

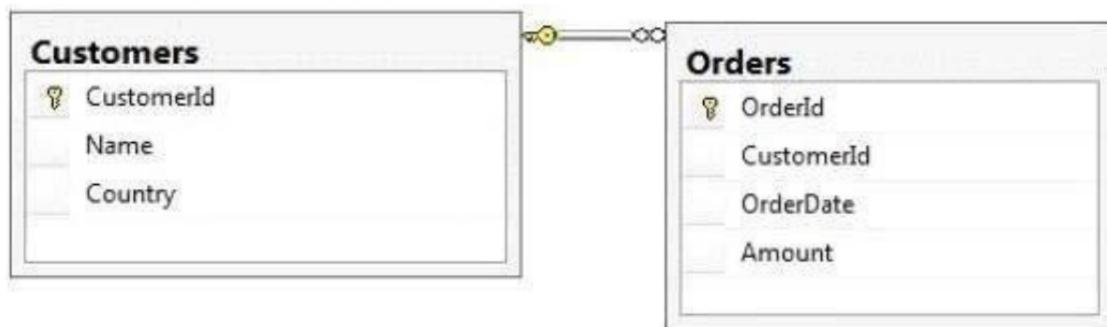
We need at least 400 eDTUs and the use of a Standard database pool.

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

NEW QUESTION 85

- (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, CrderId, 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 86

- (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 91

- (Exam Topic 7)

You are a database developer of a Microsoft SQL Server 2014 database. You are designing a table that will store Customer data from different sources. The table will include a column that contains the CustomerID from the source system and a column that contains the SourceID. A sample of this data is as shown in the following table.

SourceID	CustomerID	Customer Name
1	234	John Smith
3	7345	Jason Warren
3	4402	Susan Burk
2	866	Michael Allen

You need to ensure that the table has no duplicate CustomerID within a SourceID. You also need to ensure that the data in the table is in the order of SourceID and then CustomerID. Which Transact- SQL statement should you use?

- A. CREATE TABLE Customer(SourceID int NOT NULL IDENTITY, CustomerID int NOT NULL IDENTITY, CustomerName varchar(255) NOT NULL);
- B. CREATE TABLE Customer(SourceID int NOT NULL, CustomerID int NOT NULL PRIMARY KEY CLUSTERED, CustomerName varchar(255) NOT NULL);
- C. CREATE TABLE Customer(SourceID int NOT NULL PRIMARY KEY CLUSTERED, CustomerID int NOT NULL UNIQUE, CustomerName varchar(255) NOT NULL);
- D. CREATE TABLE Customer(SourceID int NOT NULL, CustomerID int NOT NULL, CustomerName varchar(255) NOT NULL, CONSTRAINT PK_Customer PRIMARY KEY CLUSTERED(SourceID, CustomerID));

Answer: D

NEW QUESTION 94

- (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 95

- (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 index1 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 98

- (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 102

- (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 30 storage accounts that each has one container. You create a VHD in each container. Does this meet the goal?

- A. Yes
- B. No

Answer: A

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 106

- (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 108

- (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 109

- (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 110

- (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 112

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance. After a routine shutdown, the drive that contains tempdb fails.

You need to be able to start the SQL Server. What should you do?

- A. Modify tempdb location in startup parameters.
- B. Start SQL Server in minimal configuration mode.
- C. Start SQL Server in single-user mode.
- D. Configure SQL Server to bypass Windows application logging.

Answer: B

Explanation:

If you have configuration problems that prevent the server from starting, you can start an instance of Microsoft SQL Server by using the minimal configuration startup option.

When you start an instance of SQL Server in minimal configuration mode, note the following: Only a single user can connect, and the CHECKPOINT process is not executed.

Remote access and read-ahead are disabled. Startup stored procedures do not run.

tempdb is configured at the smallest possible size.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/start-sql-server-with-minimal-configur>

NEW QUESTION 114

- (Exam Topic 7)

You are using dynamic management views to monitor an SQL Server server named SQL1. A database administrator named Dbal must monitor the health of SQL1.

You need to ensure that Dbal can access dynamic management views for SQL1. The solution must use the principle of least privilege.

Which permissions should you assign to Dbal?

- A. VIEW ANY DEFINITION
- B. VIEW SERVER STATE
- C. VIEW DEFINITION

D. CONTROL SERVER

Answer: B

Explanation:

To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission. There are two types of dynamic management views and functions:

Server-scoped dynamic management views and functions. These require VIEW SERVER STATE permission on the server.

Database-scoped dynamic management views and functions. These require VIEW DATABASE STATE permission on the database.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/system-dynamic->

NEW QUESTION 115

- (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 117

- (Exam Topic 7)

You create a new Microsoft Azure subscription.

You need to create a group of Azure SQL databases that share resources. 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: D

Explanation:

SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single Azure SQL Database server and share a set number of resources (elastic Database Transaction Units (eDTUs)) at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

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

NEW QUESTION 118

- (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 120

- (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 124

- (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 129

.....

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](#)