

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 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 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 3

- (Topic 1)

You have a Microsoft SQL Server 2014 named SRV2014 that has a single tempdb database file. The tempdb database file is eight gigabytes (GB) in size.

You install a SQL Server 2016 instance named SQL Server 2016 by using default settings. The new instance has eight logical processor cores.

You plan to migrate the databases from SRV2014 to SRV2016.

You need to configure the tempdb database on SRV2016. The solution must minimize the number of future tempdb autogrowth events.

What should you do?

- A. Increase the size of the tempdb datafile to 8 G
- B. In the tempdb database, set the value of the MAXDOP property to 8.
- C. Increase the size of the tempdb data files to 1 GB.
- D. Add seven additional tempdb data file
- E. In the tempdb database, set the value of the MAXDOP property to 8.
- F. Set the value for the autogrowth setting for the tempdb data file to 128 megabytes (MB). Add seven additional tempdb data files and set the autogrowth value to 128 MB.

Answer: B

Explanation:

In an effort to simplify the tempdb configuration experience, SQL Server 2016 setup has been extended to configure various properties for tempdb for multi-processor environments.

1. A new tab dedicated to tempdb has been added to the Database Engine Configuration step of setup workflow.

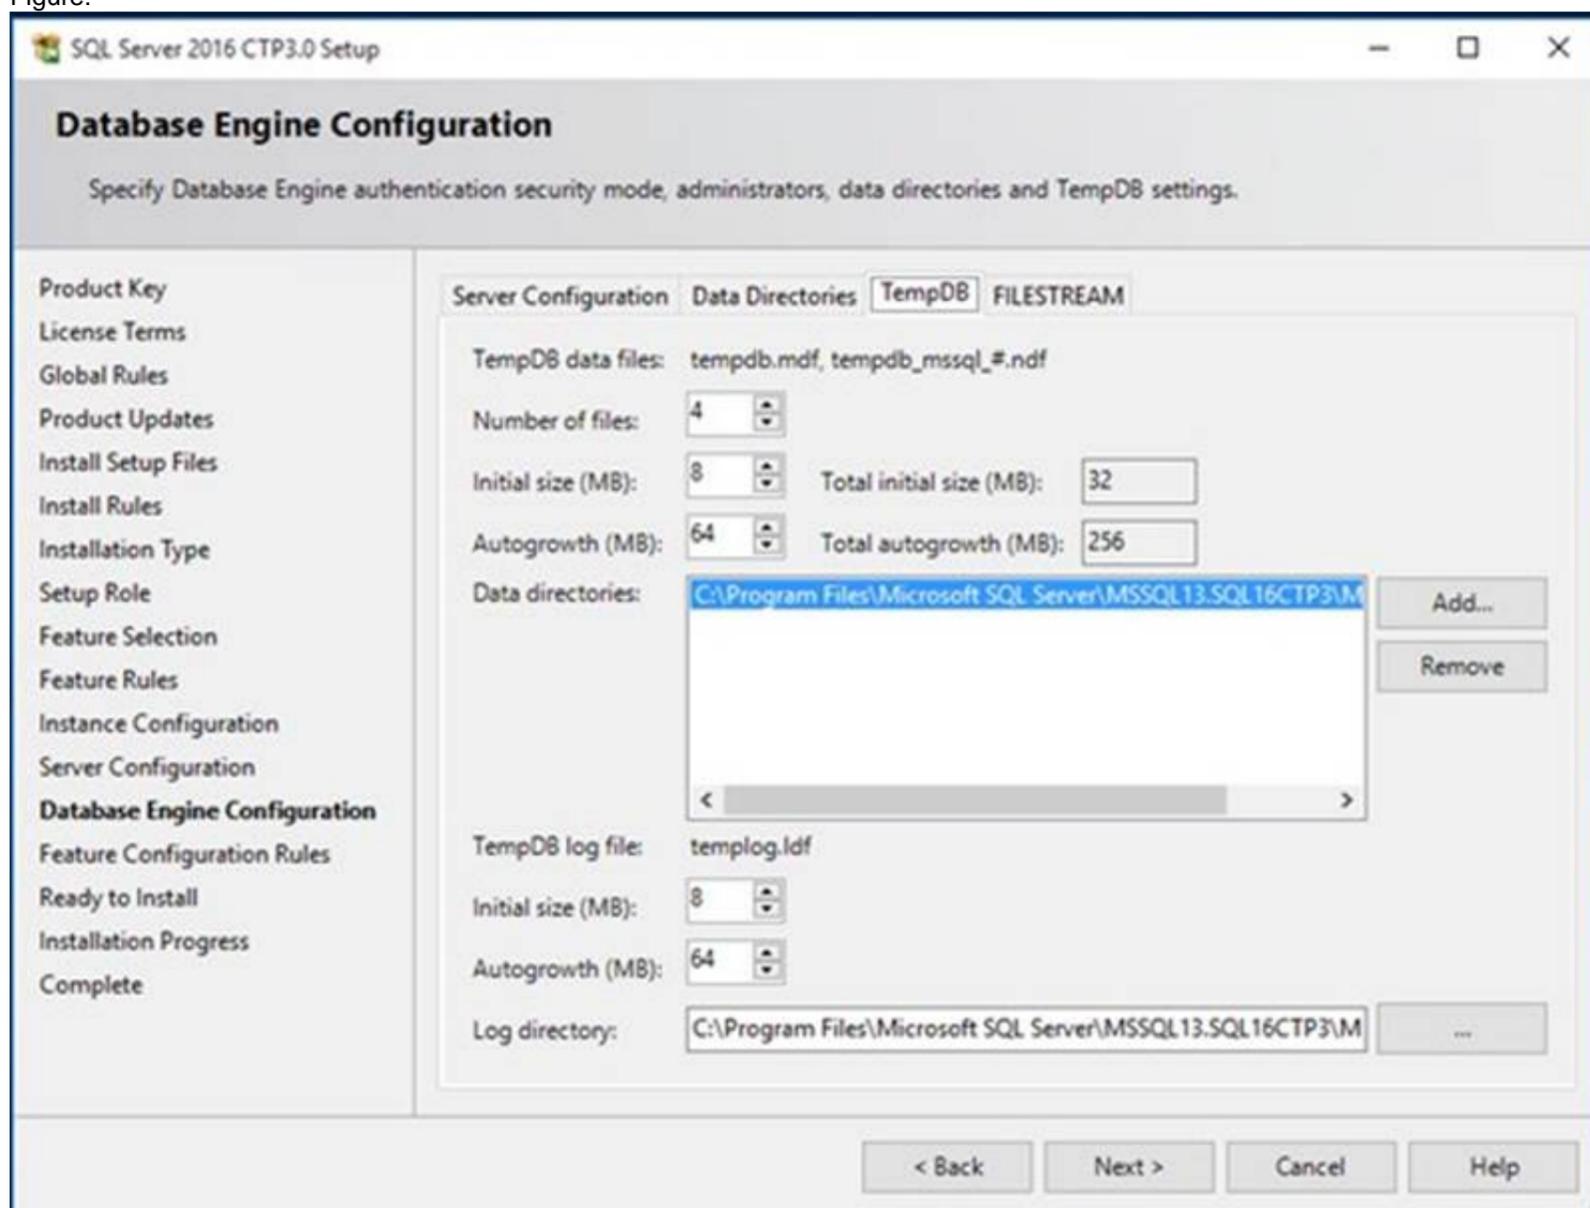
2. Configuration options: Data Files

* Number of files – this will default to the lower value of 8 or number of logical cores as detected by setup.

* Initial size – is specified in MB and applies to each tempdb data file. This makes it easier to configure all files of same size. Total initial size is the cumulative tempdb data file size (Number of files * Initial Size) that will be created.

* Autogrowth – is specified in MB (fixed growth is preferred as opposed to a non-linear percentage based growth) and applies to each file. The default value of

64MB was chosen to cover one PFS interval.
Figure:



References: <https://blogs.msdn.microsoft.com/psssql/2016/03/17/sql-2016-it-just-runs-faster-automatic-tempdb-configuration/>

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 P2 tier
- 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

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 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 PARAMETER_SNIFFING option for the databases. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

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

NEW QUESTION 7

- (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 8

- (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 9

- (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 10

- (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 10

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 13

- (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.

You deploy Microsoft SQL Server to a virtual machine in Azure. You distribute the database files and filegroups across multiple Azure storage disks.

You must be able to manage the databases as individual entities by using SQL Server Management Studio. All data in the databases must be stored encrypted.

Backups must be encrypted by using the same key as the live copy of the database.

You need to secure the data. 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: C

Explanation:

Transparent data encryption (TDE) encrypts your databases, associated backups, and transaction log files at rest without requiring changes to your applications. TDE encrypts the storage of an entire database by using a symmetric key called the database encryption key. In SQL Database the database encryption key is protected by a built-in server certificate. The built-in server certificate is unique for each SQL Database server.

References: <https://msdn.microsoft.com/en-us/library/dn948096.aspx>

NEW QUESTION 14

- (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 limit the amount of memory, CPU, and IOPS used for the pool of all queries that the Reporting_user login can run concurrently.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

SQL Server Resource Governor is a feature that you can use to manage SQL Server workload and system resource consumption. Resource Governor enables you to specify limits on the amount of CPU, physical IO, and memory that incoming application requests can use.

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

NEW QUESTION 15

- (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 18

HOTSPOT - (Topic 4)

You need to optimize SRV1.

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

Answer Area

How should you modify the tempdb configuration?

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

How should you reconfigure the tempdb database?

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

Box 2: Add additional tempdb files.

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

NEW QUESTION 23

HOTSPOT - (Topic 5)

You need to create the contosodb1 database.

How should you complete the Azure PowerShell command? To answer, select the appropriate Azure PowerShell segments in the answer area.

Answer Area

▼
New-AzureSqlDatabase
New-AzureRmSqlDatabase
Set-AzureRmSqlDatabase

- ResourceGroupName "contosodbrg"

- ServerName "contososrv"

- DatabaseName "contosodbl"

- Edition

▼
Basic
Standard
Premium

- RequestedServiceObjectName S2

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: New-AzureRmSqlDatabase

New-AzureRmSqlDatabase creates a database or an elastic database.

New-AzureRmSqlDatabase is a command with the Azure Resource Manager (AzureRM) module. Azure Resource Manager enables you to work with the resources in your solution as a group.

NEW QUESTION 27

HOTSPOT - (Topic 5)

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Data Entry: Geo-replicated database only

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

Business intelligence: Elastic database pools only

From Contoso scenario: For the business intelligence application, corporate executives must be able to view all data in near real-time with low network latency.

SQL DB elastic pools provide a simple cost effective solution to manage the performance goals for multiple databases that have widely varying and unpredictable usage patterns.

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

NEW QUESTION 31

HOTSPOT - (Topic 6)

You need to open the firewall ports for use with SQL Server environment. In table below, identify the firewall port that you must use for each service.

NOTE: Make only one selection in each column.

Answer Area

Port number	Report Server	SQL Server Browser service for SSAS
80	<input type="radio"/>	<input type="radio"/>
135	<input type="radio"/>	<input type="radio"/>
1433	<input type="radio"/>	<input type="radio"/>
2382	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Report Server: 80

By default, the report server listens for HTTP requests on port 80.

NEW QUESTION 33

- (Exam Topic 7)

You need to create an Elastic Database job to rebuild indexes across 10 Microsoft Azure SQL databases. Which powershell cmdlet should you run?

- A. New-AzureSqlJob
- B. New-AzureWebsiteJob
- C. New-AzureBatchJob
- D. New-ScheduledJobOption
- E. New-JobTrigger

Answer: A

Explanation:

The New-AzureSqlJob cmdlet, in the ElasticDatabaseJobs module, creates a job definition to be used for subsequent job runs.

References:

<https://docs.microsoft.com/en-us/powershell/module/elasticdatabasejobs/new-azuresqljob?view=azureelasticdbj>

NEW QUESTION 34

- (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 37

- (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 40

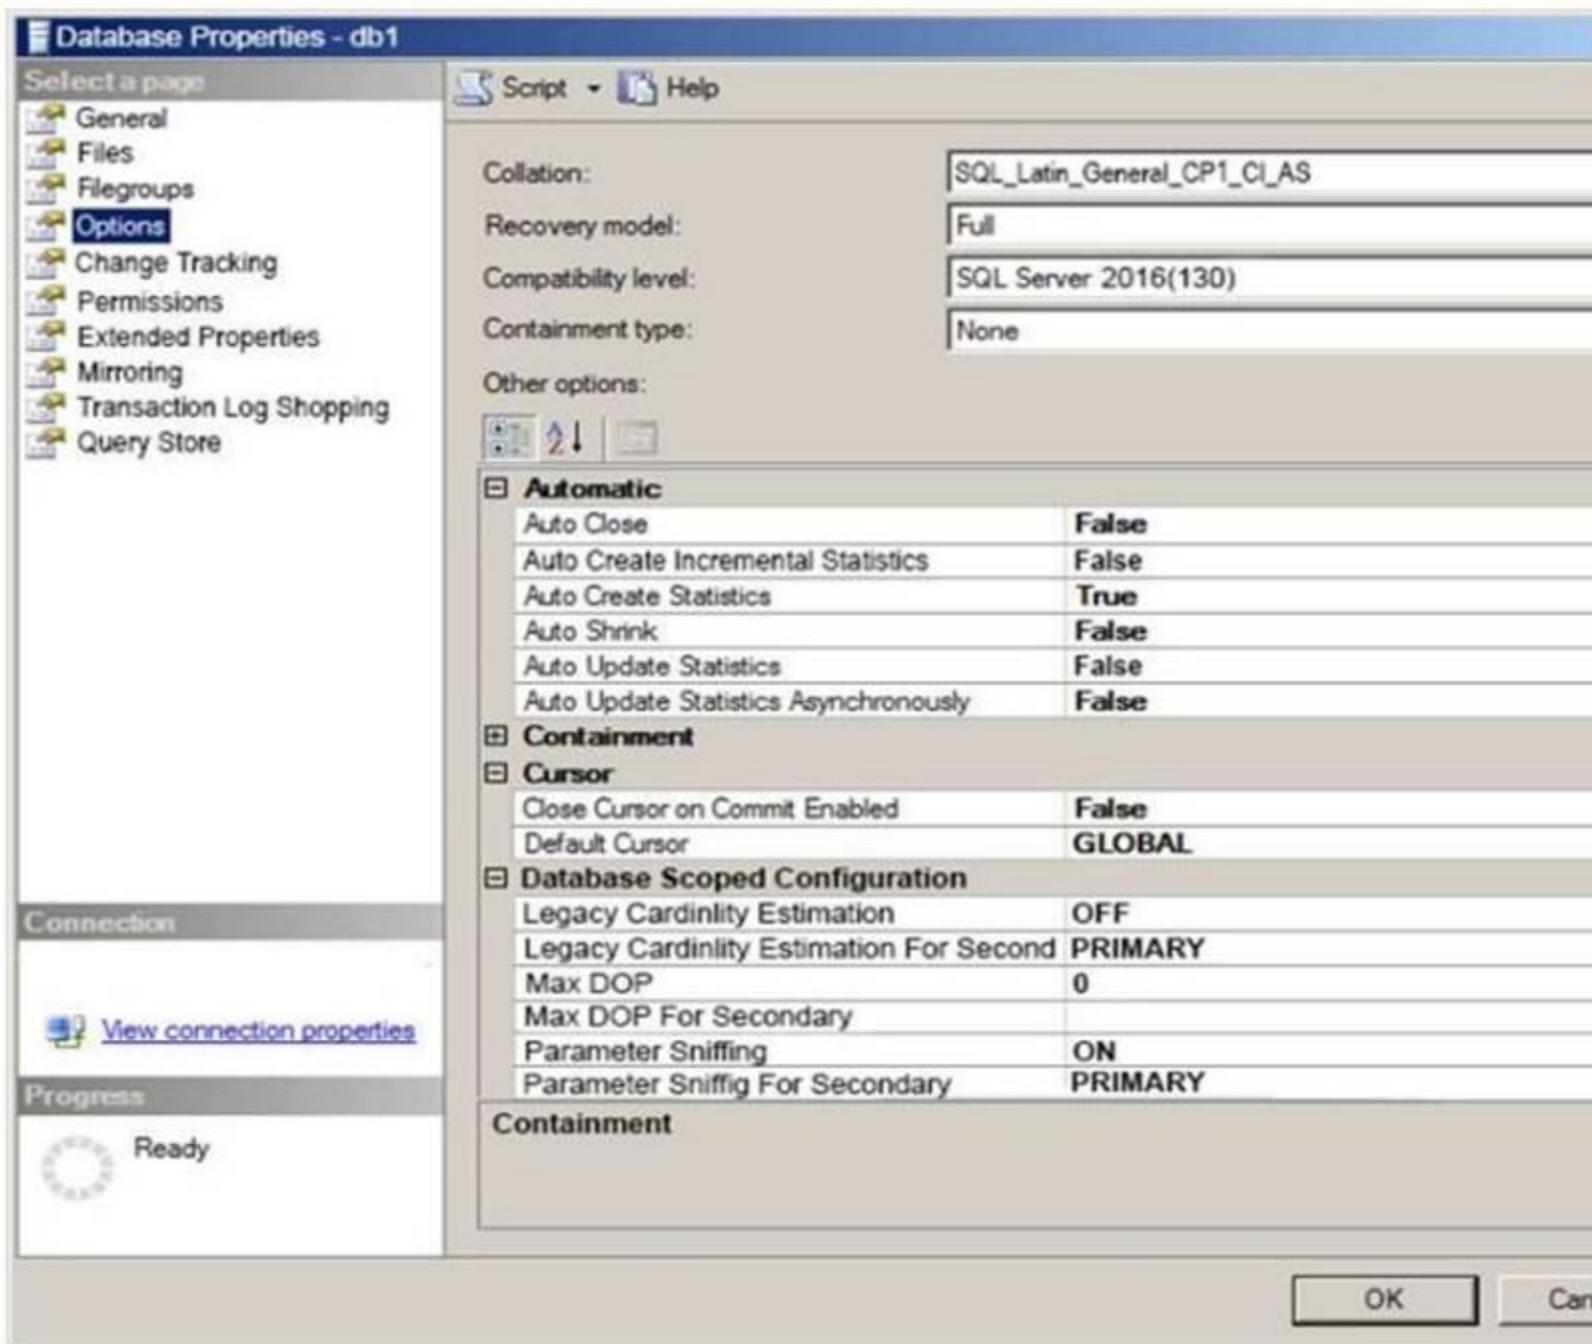
- (Exam Topic 7)

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 41

- (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 44

- (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 46

- (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 49

- (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 52

- (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 54

- (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 57

- (Exam Topic 7)

You have an on-premises Microsoft SQL Server named Server1.

You provision a Microsoft Azure SQL Database server named Server2. On Server1, you create a database named DB1.

You need to enable the Stretch Database feature for DB1.

Which five 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 master key in the master database | 1. |
| -Create a firewall rule in Azure | 2. |
| -Create a master key in DB1 | 3. |
| -Enable the remote data archive option in DB1 | 4. |
| -Create a database scoped credential for authentication to Azure. | 5. |
| -Create a server-level credential for authentication to Azure. | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Enable the remote data archive option in DB1 Prerequisite: Enable Stretch Database on the server
Before you can enable Stretch Database on a database or a table, you have to enable it on the local server. To enable Stretch Database on the server manually, run sp_configure and turn on the remote data archive option.

Step 2: Create a firewall rule in Azure

On the Azure server, create a firewall rule with the IP address range of the SQL Server that lets SQL Server communicate with the remote server.

Step 3: Create a master key in the master database

To configure a SQL Server database for Stretch Database, the database has to have a database master key. The database master key secures the credentials that Stretch Database uses to connect to the remote database.

Step 4: Create a database scoped credential for authentication to Azure

When you configure a database for Stretch Database, you have to provide a credential for Stretch Database to use for communication between the on premises SQL Server and the remote Azure server. You have two options.

Step 5: Create a server-level credential for authentication to Azure.

To configure a database for Stretch Database, run the ALTER DATABASE command. For the SERVER argument, provide the name of an existing Azure server, including the

.d atabase.windows.net portion of the name - for example, MyStretchDatabaseServer.database.windows.net.

Provide an existing administrator credential with the CREDENTIAL argument, or specify FEDERATED_SERVICE_ACCOUNT = ON. The following example provides an existing credential.

```
ALTER DATABASE <database name> SET REMOTE_DATA_ARCHIVE = ON (
SERVER = '<server_name>',
CREDENTIAL = <db_scoped_credential_name>
); GO
```

References:

<https://docs.microsoft.com/en-us/sql/sql-server/stretch-database/enable-stretch-database-for-a-database?view=sq>

NEW QUESTION 58

- (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 62

- (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 65

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database named Orders. Orders contains a table named OrderShip that is defined as follows:

```
CREATE TABLE OrderShip
(OrderID bigint NOT NULL PRIMARY KEY,
 CustomerID int NOT NULL,
 ShipAddress nvarchar(500) NOT NULL,
 CountryCode tinyint NULL)
```

A NULL value represents a domestic order. Ninety percent of the values in CountryCode are NULL. Customers require a procedure that will return orders for all customers from a specified country. You create a new procedure:

```
CREATE PROCEDURE p_GetIntlOrders
(@countrycode tinyint)
AS
SELECT DISTINCT CustomerID, ShipAddress
FROM OrderShip
WHERE CountryCode = @countrycode
GO
```

Performance on this procedure is slow.

You need to alter the schema to optimize this query. Objects created must use a minimum amount of resources.

Which Transact-SQL statement should you use?

- A. CREATE NONCLUSTERED INDEX IX_CountryCode ON Ordership (CountryCode) WHERE CountryCode IS NOT NULL
- B. CREATE STATISTICS ST_CountryCode ON OrderShip (CountryCode) WHERE CountryCode IS NOT NULL
- C. CREATE CLUSTERED INDEX IX_CountryCode ON OrderShip (CountryCode)
- D. CREATE INDEX IX_CountryCode ON OrderShip (CustomerID) WHERE CountryCode IS NOT NULL

Answer: B

Explanation:

Here creating statistics is relevant. The CREATE STATISTICS command creates query optimization statistics on one or more columns of a table, an indexed view, or an external table. For most queries, the query optimizer already generates the necessary statistics for a high-quality query plan; in a few cases, you need to create additional statistics with CREATE STATISTICS or modify the query design to improve query performance.

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

NEW QUESTION 66

- (Exam Topic 7)

You administer a Windows 2008 server hosting an instance of Microsoft SQL Server 2014 Standard Edition. The server hosts a database named Orders.

Users report that a query that filters on OrderDate is taking an exceptionally long time. You discover that an index named IX_OrderDate on the CustomerOrder table is heavily fragmented.

You need to improve the performance of the IX_OrderDate index. The index should remain online during the operation.

Which Transact-SQL command should you use?

- A. ALTER INDEX IX_OrderDate ON CustomerOrder DISABLE
- B. ALTER INDEX IX_OrderDate ON CustomerOrder ENABLE
- C. ALTER INDEX IX_OrderDate ON CustomerOrder REORGANIZE
- D. ALTER INDEX IX_OrderDate ON CustomerOrder REBUILD

Answer: C

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 68

- (Exam Topic 7)

You administer a single server that contains a Microsoft SQL Server 2014 default instance. You plan to install a new application that requires the deployment of a database on the server. The application login requires sysadmin permissions.

You need to ensure that the application login is unable to access other production databases. What should you do?

- A. Use the SQL Server default instance and configure an affinity mask.
- B. Install a new named SQL Server instance on the server.
- C. Use the SQL Server default instance and enable Contained Databases.
- D. Install a new default SQL Server instance on the server.

Answer: B

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/sql-server/install/work-with-multiple-versions-and-instances-of-sql-server>

NEW QUESTION 69

- (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 74

- (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 76

- (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 77

- (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 80

- (Exam Topic 7)

You have a Microsoft Azure SQL database in the US West region. You need to create a replica in the US East region. Which cmdlet should you run first?

- A. `New-AzureRmAvailabilitySet`
- B. `New-AzureRmLoadBalancer`
- C. `New-AzureRmSqlDatabaseSecondary`
- D. `New-AzureRmSqlElasticPool`
- E. `New-AzureRmVM`
- F. `New-AzureRmSqlServer`
- G. `New-AzureRmSqlDatabaseCopy`
- H. `New-AzureRmSqlServerCommunicationLink`

Answer: G

Explanation:

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

References:

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

NEW QUESTION 82

- (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 87

- (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 write/sec'` parameter. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 89

- (Exam Topic 7)

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 set the value of the `MAXDOP` parameter to 2.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

When an instance of SQL Server runs on a computer that has more than one microprocessor or CPU, it detects the best degree of parallelism, that is, the number of processors employed to run a single statement, for each parallel plan execution. You can use the max degree of parallelism (`MAXDOP`) option to limit the number of processors to use in parallel plan execution.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-the-max-degree-of-parallelism>

NEW QUESTION 92

- (Exam Topic 7)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

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

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

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

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

You need to prevent the PAGELATCH_IO waits from occurring. Solution: You add more files to db1.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

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

NEW QUESTION 93

- (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 Basic database pool that has 1,600 eDTUs. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

We need the use of a Standard database pool.

NEW QUESTION 98

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 server. The MSSQLSERVER service uses a domain account named CONTOSO\SQLService.

You plan to configure Instant File Initialization.

You need to ensure that Data File Autogrow operations use Instant File Initialization. What should you do? Choose all that apply.

- A. Restart the SQL Server Agent Service.
- B. Disable snapshot isolation.
- C. Restart the SQL Server Service.
- D. Add the CONTOSO\SQLService account to the Perform Volume Maintenance Tasks local security policy.
- E. Add the CONTOSO\SQLService account to the Server Operators fixed server role.
- F. Enable snapshot isolation.

Answer: CD

Explanation:

How To Enable Instant File Initialization References:

<http://msdn.microsoft.com/en-us/library/ms175935.aspx>

NEW QUESTION 102

- (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 103

- (Exam Topic 7)

You are building the database platform for a multi-tenant application. The application will have one database per tenant and will have at least 30 tenants. Each tenant will have a separate resource group for billing purposes.

The application will require at least 10 GB of clustered columnstore indexes for each database.

You need to implement the database platform for the application. The solution must minimize costs. What should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Service tier:

	▼
Basic	
Standard	
Premium	
Premium RS	

Database implementation:

	▼
One individual Azure SQL database	
Thirty individual Azure SQL databases	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The Standard tier service allows for 1TB of data. Here 30 x 10 GB, 0.3 TB, is required.

NEW QUESTION 108

- (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 109

- (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 114

- (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 119

- (Exam Topic 7)

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

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

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

Which edition should you choose?

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

Answer: B

NEW QUESTION 122

- (Exam Topic 7)

You have a Microsoft SQL Server instance that has a database named DB1. DB1 has data files on drive E and transaction logs on drive L.

You perform full backups of DB1 daily and transaction log backups hourly. Drive E fails and is replaced.

You need to recover DB1 and prevent any data loss.

Which four 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.

Select and Place:

Select and Place:

Actions	Answer Area
Restore the tail-log backup.	1
Restore a full backup.	2
Perform a tail-log backup.	3
Restore the log backups.	4
Truncate the log of DB1.	
Delete DB1.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Section: Deploy and migrate applications Step 1: Perform a tail-log backup.

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.

Step 2: Restore a full backup.

Backups must be restored in the order in which they were created. Before you can restore a particular transaction log backup, you must first restore the following previous backups without rolling back uncommitted transactions, that is WITH NORECOVERY:

The full database backup and the last differential backup, if any, taken before the particular transaction log backup.

Step 3: Restore the log backups.

Log backups must be applied in the sequence in which they were created, without any gaps in the log chain. Step 4: Restore the tail-log backups.

Reference:

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

NEW QUESTION 126

- (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 130

- (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 132

- (Exam Topic 7)

Background

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

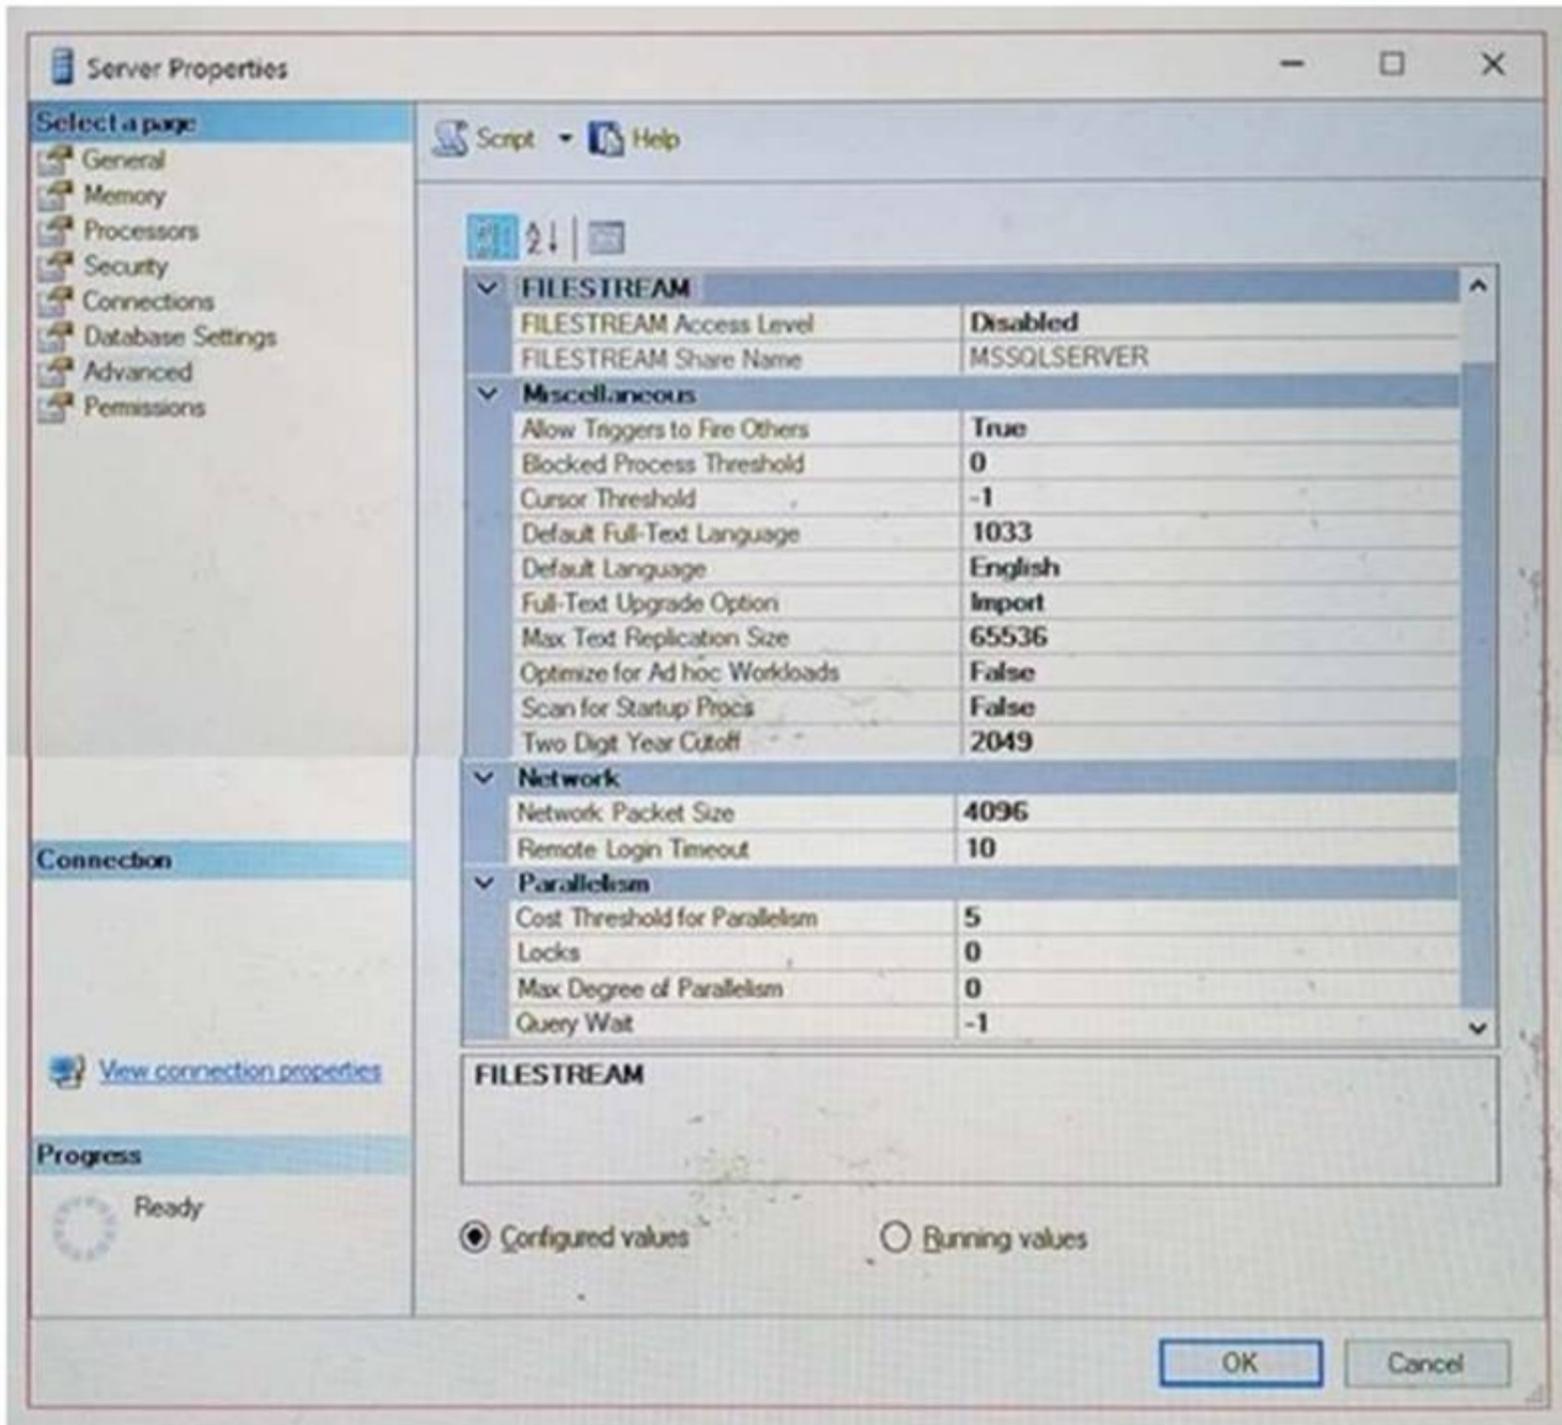
The environment includes the following servers: SRV1 and SRV2. SRV1 has 16 logical cores and hosts a SQL Server instance that supports a mission-critical application. The application has approximately 30,000 concurrent users and relies heavily on the use of temporary tables.

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

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

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

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



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

Values

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

Answer area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 2: Create: server certificate

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

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

-- Attach the database that is being moved.

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

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

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

References:

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

NEW QUESTION 133

- (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 136

- (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 139

- (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:

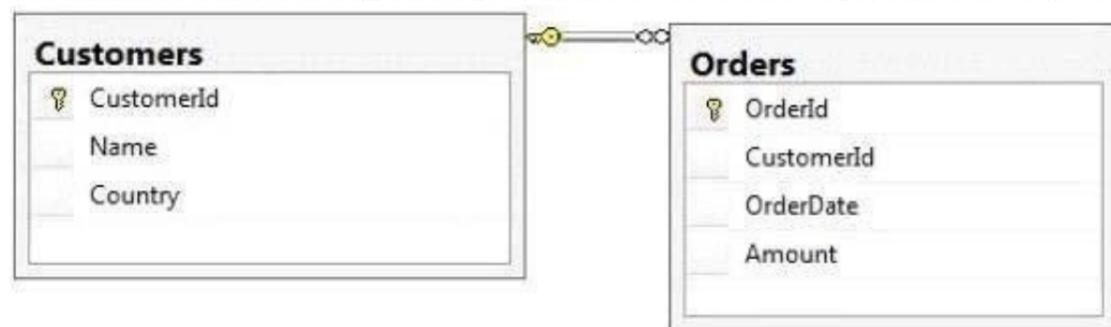
References:

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

NEW QUESTION 141

- (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 143

- (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 minimum amount of data is lost. Which recovery model should the database use?

- 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: B

Explanation:

The full recovery model requires log backups. No work is lost due to a lost or damaged data file. Can recover to a specific point in time, assuming that your backups are complete up to that point in time.

NEW QUESTION 148

- (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 indx1 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 151

- (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 156

- (Exam Topic 7)

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

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

Backups are performed as shown in the following schedule:

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 158

- (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 161

- (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 an AB compute-intensive instance and store the tempdb database in Standard storage. Does this meet the goal?

- A. Yes
- B. No

Answer: B

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 164

- (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 166

- (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 170

- (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: DExplanation:References:

NEW QUESTION 173

- (Exam Topic 7)

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

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

Which resource causes the issue?

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

Answer: B

Explanation:

PAGEIOLATCH Wait time and WaitCount are both high.

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

The Microsoft definition of this wait type is:

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

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

NEW QUESTION 178

- (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 179

- (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 180

- (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 181

- (Exam Topic 7)

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 an Always On Availability Group.

You expect to have less than 1 million IO transaction per month.

You need to recommend a storage solution for the SQL Servers. The solution must minimize 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:

P10 has 500 IOPS per disk, which provides for more than 1 million IOPS per month.

Note: $3600 \times 30 \times 500$ is 54 million IOPS/month.

References: <https://azure.microsoft.com/en-us/pricing/details/managed-disks/>

NEW QUESTION 185

- (Exam Topic 7)

You are the database administrator for your company. Your company has one main office and two branch offices. You plan to create three databases named DB1, DB2, and DB3 that will be hosted on one Azure SQL Database server. You have the following requirements:

The main office must be able to connect to all three databases.

The branch offices must be able to connect to DB2 and DB3.

The branch offices must not be able to access DB1.

You need to configure transparent data encryption (TDE) for DB1. Which two actions should you perform? Each correct answer presents part of the solution.

- A. Run `CREATE CERTIFICATE cert1 WITH Subject = TDE Cert1` on DB1.
- B. Connect to DB1.
- C. Run `ALTER DATABASE DB1 SET ENCRYPTION ON;`
- D. Connect to the master database.
- E. Run `CREATE MASTER KEY` on the master database.

Answer: BC

Explanation:

You should connect to DB1. To encrypt DB1, you connect directly to DB1. When you connect to DB1. You use your dbmanager or administrative credentials.

You should run `ALTER DATABASE DB1 SET ENCRYPTION ON`.

You use the `ALTER DATABASE DB1 SET ENCRYPTION ON` statement to encrypt the database. This is the statement that turns on TDE for Azure SQL Database.

NEW QUESTION 190

- (Exam Topic 7)

You create an availability group that has replicas named HA/Server01 and HA/Server02. Currently, HA/Server01 is the primary replica.

You have multiple queries that read data and produce reports from the database.

You need to offload the reporting workload to the secondary replica when HA/Server01 is the primary replica. What should you do?

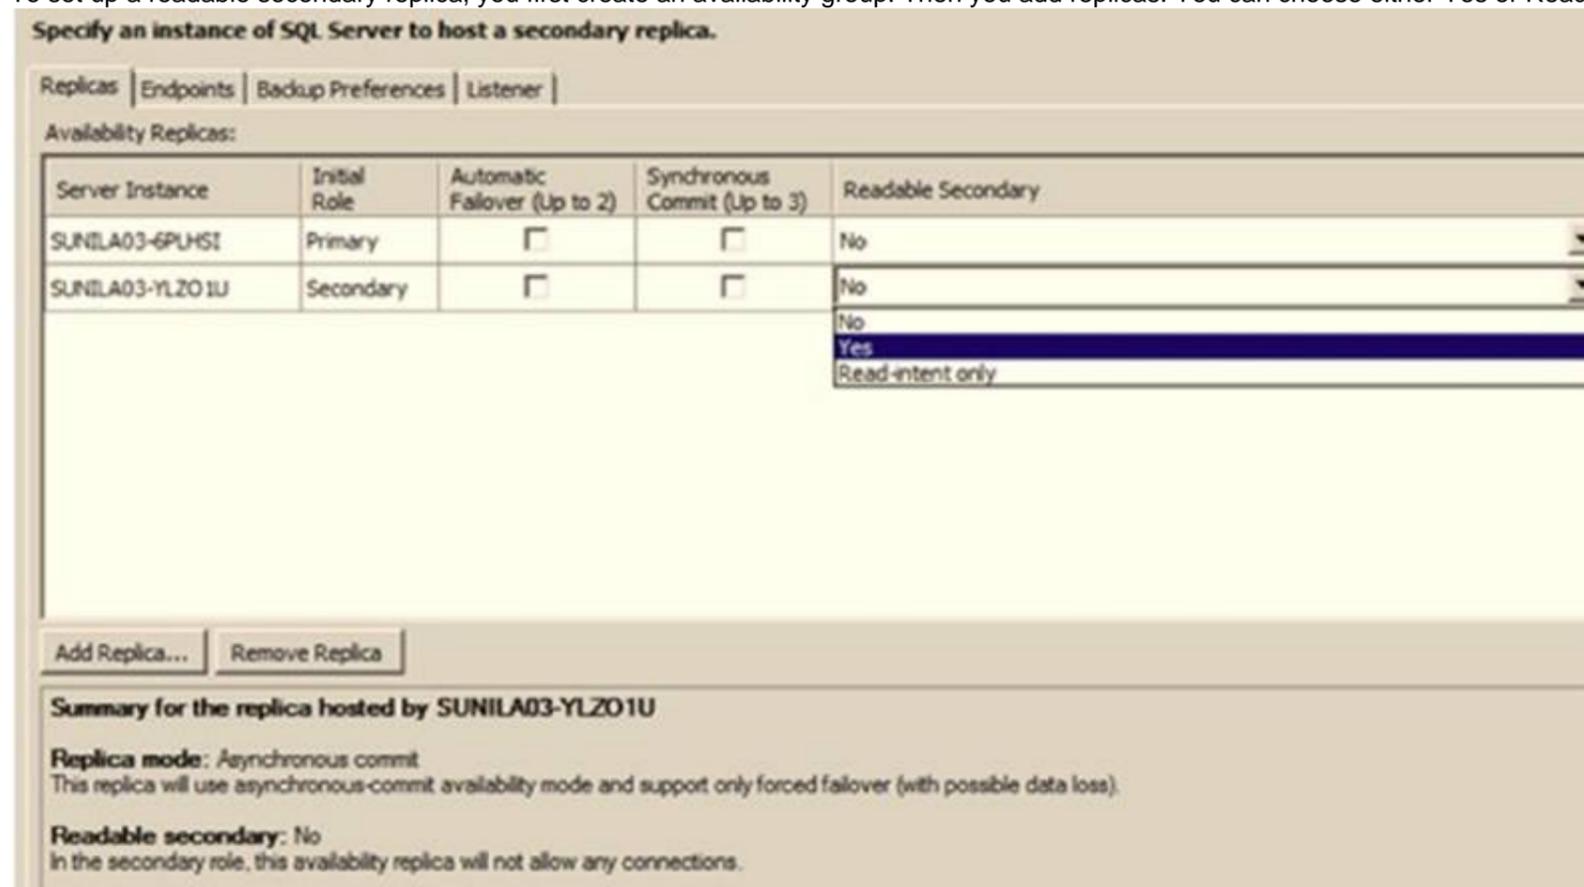
- A. Set the Availability Mode property of HA/Server02 to Asynchronous commit.
- B. Set the Readable Secondary property of HA/Server02 to Read-intent only.
- C. Set the Connections in Primary Role property of HA/Server01 to Allow read/write connections.

D. Set the Availability Mode property of HA/Server01 to Asynchronous commit.

Answer: B

Explanation:

To set up a readable secondary replica, you first create an availability group. Then you add replicas. You can choose either Yes or Read-intent only options.



References: <http://msdn.microsoft.com/en-us/library/jj542414.aspx>

NEW QUESTION 191

- (Exam Topic 7)

You plan to install a Microsoft SQL Server 2014 instance.

The instance will support a database that has the following requirements: Store Excel workbooks on the file system.

Access the workbooks through Transact-SQL.

Include the workbooks in database backups.

During installation, you need to ensure that the requirements will be met. Which feature should you use?

- A. Excel Services
- B. FILESTREAM
- C. SQL Server Integration Services (SSIS)
- D. OpenXML

Answer: B

Explanation:

FILESTREAM enables SQL Server-based applications to store unstructured data, such as documents and images, on the file system. Applications can leverage the rich streaming APIs and performance of the file system and at the same time maintain transactional consistency between the unstructured data and corresponding structured data.

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

NEW QUESTION 196

- (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 201

- (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 202

- (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 206

- (Exam Topic 7)

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

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

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

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

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

Answer: A

Explanation:

In a file restore, the goal is to restore one or more damaged files without restoring the whole database. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/file-restores-simple-recovery-model>

NEW QUESTION 211

- (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 216

- (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 219

- (Exam Topic 7)

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

What should you do?

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

Answer: A

Explanation:

A common way to define a computed column is by using a user-defined function (UDF) to encapsulate the calculation logic.

References: <https://blogs.msdn.microsoft.com/sqlcat/2011/11/28/a-computed-column-defined-with-a-user-define>

NEW QUESTION 222

- (Exam Topic 7)

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

Which object should you use?

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

Answer: D

Explanation:

Stored procedures accept input parameters and return multiple values in the form of output parameters to the calling program. They cannot be used in views.

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

NEW QUESTION 223

- (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 service master key.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

The Service Master Key is the root of the SQL Server encryption hierarchy. It does not handle alerts and jobs. 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 225

.....

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>