

70-765 Dumps

Provisioning SQL Databases (beta)

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



NEW QUESTION 1

- (Topic 1)

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

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

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

Which storage option should you use?

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

Answer: E

Explanation:

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

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

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

NEW QUESTION 2

- (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 3

- (Topic 1)

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in a development environment.

You need to provide storage to the environment that minimizes 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: D

NEW QUESTION 4

- (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 to8.

- 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.

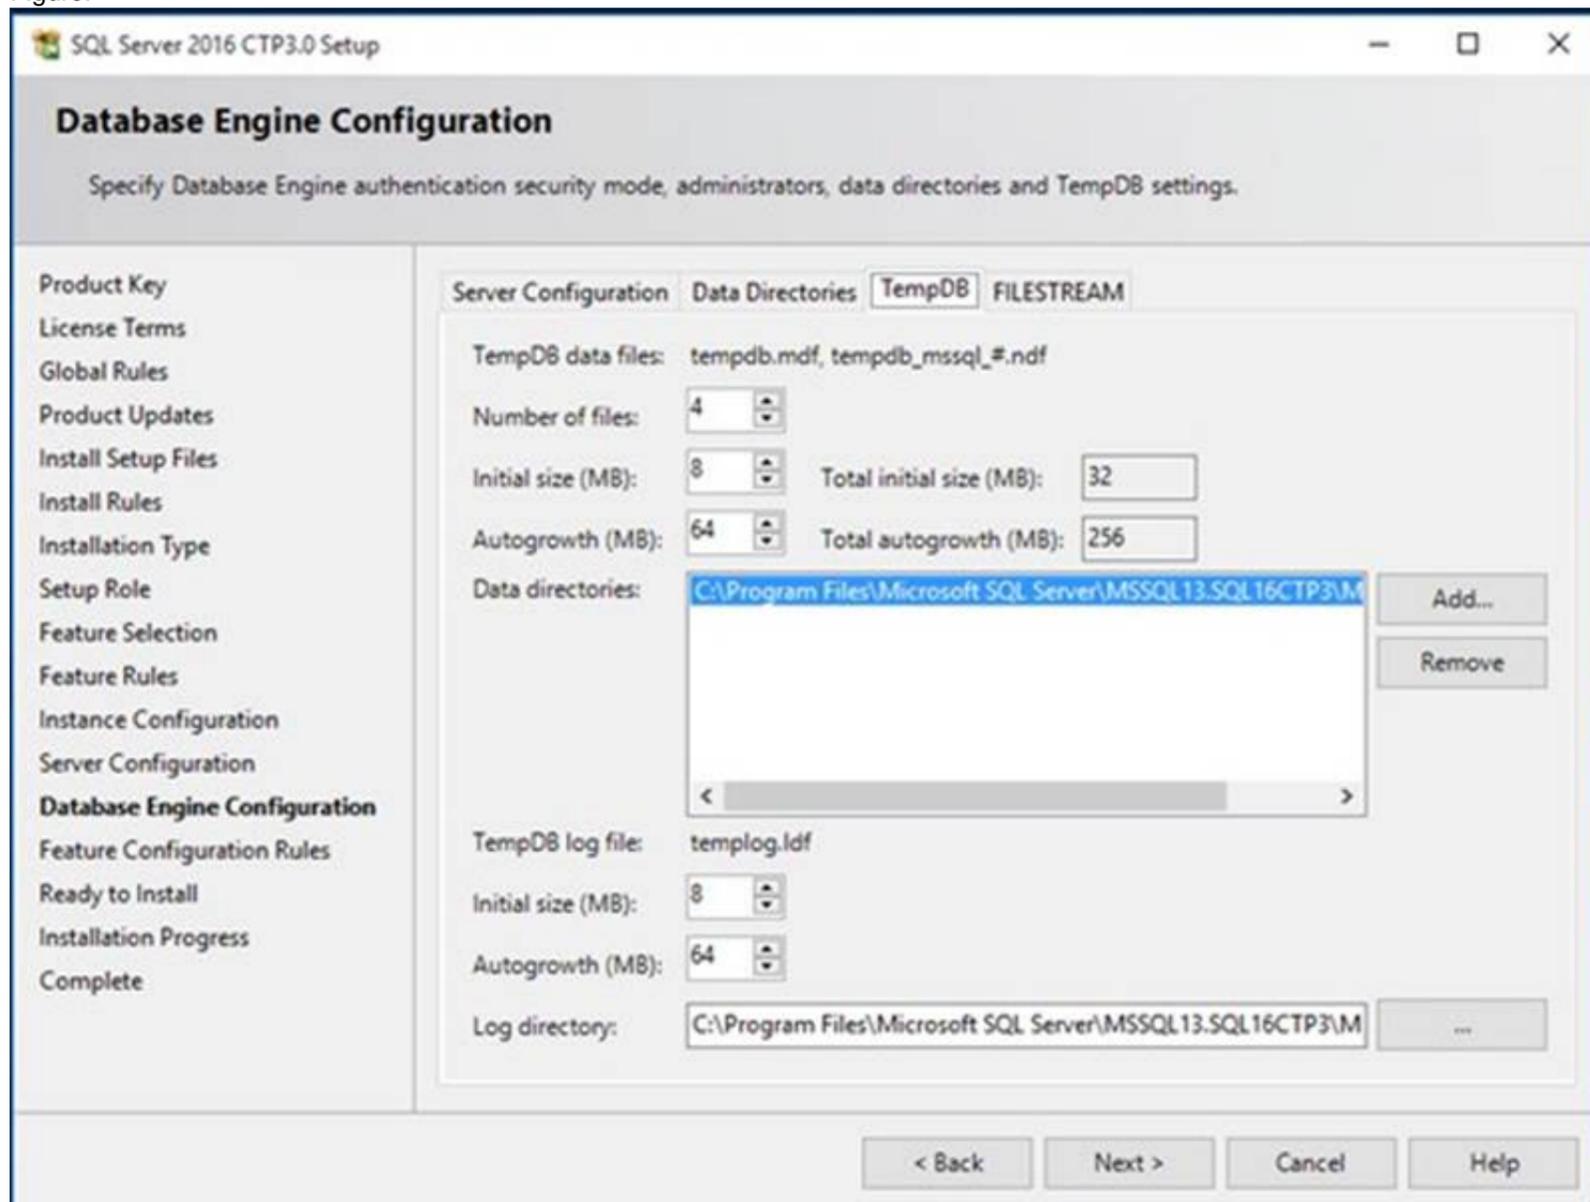
2. Configuration options: Data Files

* 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 5

- (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 use Visual Studio to create a XAML template that defines the deployment and configuration settings for the SQL Server environment.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Azure Resource Manager template consists of JSON, not XAML, and expressions that you can use to construct values for your deployment.

A good JSON editor can simplify the task of creating templates.

Note: In its simplest structure, an 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 6

DRAG DROP - (Topic 1)

You are building a new Always On Availability Group in Microsoft Azure. The corporate domain controllers (DCs) are attached to a virtual network named ProductionNetwork. The DCs are part of an availability set named ProductionServers1.

You create the first node of the availability group and add it to an availability set named ProductionServers2. The availability group node is a virtual machine (VM) that runs Microsoft SQL Server. You attach the node to ProductionNetwork.

The servers in the availability group must be directly accessible only by other company VMs in Azure.

You need to configure the second SQL Server VM for the availability group.

How should you configure the VM? To answer, drag the appropriate configuration settings to the correct target locations. Each configuration setting 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.

NOTE: Each correct selection is worth one point.

Configuration settings

None/Not Assigned

ProductionServers1

ProductionNetwork

ProductionServers2

Create a new Object

VM settings page

Settings
— □ X

Storage
Disk type

Standard Premium (SSD)

* Storage account >

(new) sqlstorage3

Network

* Virtual network >

setting

* Subnet >

ProductionServers (10.1.0.0/24)

* Public IP address >

setting

* Network security group >

(new) SQLServers

Extensions

Extensions >

No extensions

Monitoring

Diagnostics

Disabled Enabled

Availability

* Availability set >

setting

OK

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

; Box 1: ProductionNetwork
The virtual network is named ProductionNetwork.

Box 2: None /Not Assigned

As the servers in the availability group must be directly accessible only by other company VMs in Azure, there should be no Public IP address.

Box 3: ProductionServer2

You create the first node of the availability group and add it to an availability set named ProductionServers2. The availability group node is a virtual machine (VM) that runs Microsoft SQL Server.

NEW QUESTION 7

- (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 8

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: <pre>EXEC sp_addrolemember 'db_datareader', 'reportuser@contoso.onmicrosoft.com'</pre>	
Run the following Transact-SQL statement: <pre>CREATE USER [reportuser@contoso.onmicrosoft.com] FROM EXTERNAL PROVIDER</pre>	
Run the following Transact-SQL statements: <pre>USE Reporting CREATE USER [reportuser@contoso.onmicrosoft.com] FOR LOGIN [reportuser@contoso.onmicrosoft.com] GRANT SELECT TO [reportuser@contoso.onmicrosoft.com]</pre>	
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 9

- (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 10

- (Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals.

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

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

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

Solution: You configure the Resource Governor to set the MAXDOP parameter to 0 for all queries against the database.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

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

NEW QUESTION 10

- (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 15

- (Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals.

You have a mission-critical application that stores data in a Microsoft SQL Server instance. The application runs several financial reports. The reports use a SQL

Server-authenticated login named Reporting_User. All queries that write data to the database use Windows authentication. Users report that the queries used to provide data for the financial reports take a long time to complete. The queries consume the majority of CPU and memory resources on the database server. As a result, read-write queries for the application also take a long time to complete. You need to improve performance of the application while still allowing the report queries to finish. Solution: You create a snapshot of the database. You configure all report queries to use the database snapshot. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Use a Resource Governor instead.
References: <https://msdn.microsoft.com/en-us/library/bb933866.aspx>

NEW QUESTION 18

HOTSPOT - (Topic 4)

You need to resolve the identified issues.

Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic.

Answer Area

<p>What setting would you change to reduce the number of execution plans in the plan cache?</p>	<p>Optimize for Ad Hoc workload ▼ Max Degree of Parallelism Query Wait</p>
<p>What setting would you change to which value to reduce the number of queries which are using parallelism?</p>	<p>Max Degree of Parallelism to 4 ▼ Cost Threshold for Parallelism to 50 Locks to 100</p>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

From exhibit we see:
Cost Threshold of Parallelism: 5 Optimize for Ad Hoc Workloads: false
Max Degree of Parallelism: 0 (This is the default setting, which enables the server to determine the maximum degree of parallelism. It is fine.)
Locks: 0
Query Wait: -1
Box 1: Optimize for Ad Hoc Workload
Change the Optimize for Ad Hoc Workload setting from false to 1/True.
The optimize for ad hoc workloads option is used to improve the efficiency of the plan cache for workloads that contain many single use ad hoc batches. When this option is set to 1, the Database Engine stores a small compiled plan stub in the plan cache when a batch is compiled for the first time, instead of the full compiled plan. This helps to relieve memory pressure by not allowing the plan cache to become filled with compiled plans that are not reused.

NEW QUESTION 21

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 23

DRAG DROP - (Topic 6)

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Reporting

The user is to be created in the Reporting database.

Box 2: CREATE USER

Box 3: FOR LOGIN [BIAppUser]

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

CREATE USER readonlyuser FROM LOGIN readonlylogin;

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

NEW QUESTION 25

HOTSPOT - (Topic 6)

You need to maximize performance of writes to each database without requiring changes to existing database tables.

In the table below, identify the database setting that you must configure for each database. NOTE: Make only one selection in each column. Each correct selection is worth one point.

Answer Area

Database setting	DB1	DB2
DELAYED_DURABILITY = FORCED	<input type="radio"/>	<input type="radio"/>
DELAYED_DURABILITY = ALLOWED	<input type="radio"/>	<input type="radio"/>
ALLOW_SNAPSHOT_ISOLATION ON	<input type="radio"/>	<input type="radio"/>
ALLOW_SNAPSHOT_ISOLATION ON and READ_COMMITTED_SNAPSHOT ON	<input type="radio"/>	<input type="radio"/>
AUTO_UPDATE_STATISTICS_ASYNC ON	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

DB1: DELAYED_DURABILITY=FORCED

From scenario: Thousands of records are inserted into DB1 or updated each second. Inserts are made by many different external applications that your company's developers do not control. You observe that transaction log write latency is a bottleneck in performance. Because of the transient nature of all the data in this database, the business can tolerate some data loss in the event of a server shutdown.

With the DELAYED_DURABILITY=FORCED setting, every transaction that commits on the database is delayed durable.

With the DELAYED_DURABILITY= ALLOWED setting, each transaction's durability is determined at the transaction level.

Note: Delayed transaction durability reduces both latency and contention within the system because:

* The transaction commit processing does not wait for log IO to finish and return control to the client.

* Concurrent transactions are less likely to contend for log IO; instead, the log buffer can be flushed to disk in larger chunks, reducing contention, and increasing throughput.

DB2: ALLOW_SNAPSHOT_ISOLATION ON and READ_COMMITTED_SNAPSHOT ON

Snapshot isolation enhances concurrency for OLTP applications.

Snapshot isolation must be enabled by setting the ALLOW_SNAPSHOT_ISOLATION ON database option before it is used in transactions.

The following statements activate snapshot isolation and replace the default READ COMMITTED behavior with SNAPSHOT:

```
ALTER DATABASE MyDatabase
SET ALLOW_SNAPSHOT_ISOLATION ON
ALTER DATABASE MyDatabase
SET READ_COMMITTED_SNAPSHOT ON
```

Setting the READ_COMMITTED_SNAPSHOT ON option allows access to versioned rows under the default READ COMMITTED isolation level.

From scenario: The DB2 database was migrated from SQLServer 2012 to SQL Server 2016. Thousands of records are updated or inserted per second. You observe that the WRITELOG wait type is the highest aggregated wait type. Most writes must have no tolerance for data loss in the event of a server shutdown. The business has identified certain write queries where data loss is tolerable in the event of a server shutdown.

References:

<https://msdn.microsoft.com/en-us/library/dn449490.aspx> [https://msdn.microsoft.com/en-us/library/tcbchxcb\(v=vs.110\).aspx](https://msdn.microsoft.com/en-us/library/tcbchxcb(v=vs.110).aspx)

NEW QUESTION 30

- (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 32

- (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 msdb database.

Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:

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 36

- (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 37

- (Exam Topic 7)

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

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

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

A. use SQL Profiler to trace all queries that are processing on the serve

B. Filter queries that have a Duration value of more than 1,000.

C. Use sp_configure to set a value for blocked process threshol

D. Create an extended event session.

E. Use the Job Activity monitor to review all processes that are actively runnin

F. Review the Job History to find out the duration of each step.

G. Run the sp_who command from a query window.

H. Run the DBCC TRACEON 1222 command from a query window and review the SQL Server event log.

Answer: A

NEW QUESTION 40

- (Exam Topic 7)

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

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

You are migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 databased that consume a total of 2 TB of disk space.

The instance sustains more than 30,000 transactions per second.

You need to provision storage for the virtual machine. The storage must be able to support the same load as the on-premises deployment.

Solution: You create one storage account that has 30 containers. You create a VHD in each container. Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

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

NEW QUESTION 43

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 environment. One of the SQL Server 2014 instances contains a database named Sales. You plan to migrate Sales to Windows Azure SQL Database. To do so, you need to implement a contained database. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

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

Answer: BD

Explanation:

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

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

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

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

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

NEW QUESTION 48

- (Exam Topic 7)

You administer a SQL Server 2014 database instance.

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

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

Answer: A

Explanation:

References:

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

NEW QUESTION 52

- (Exam Topic 7)

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

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

You need to resume Stretch Database operations.

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

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

Answer: CD

Explanation:

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

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

NEW QUESTION 56

- (Exam Topic 7)

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

The recovery model must be set to simple.

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

Your need to configure the SQL server 2016 environment.

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Model: Change from full to simple Recovery Model

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

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

Model: Set compatibility level to 120

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

References:

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

NEW QUESTION 58

- (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

.database.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:

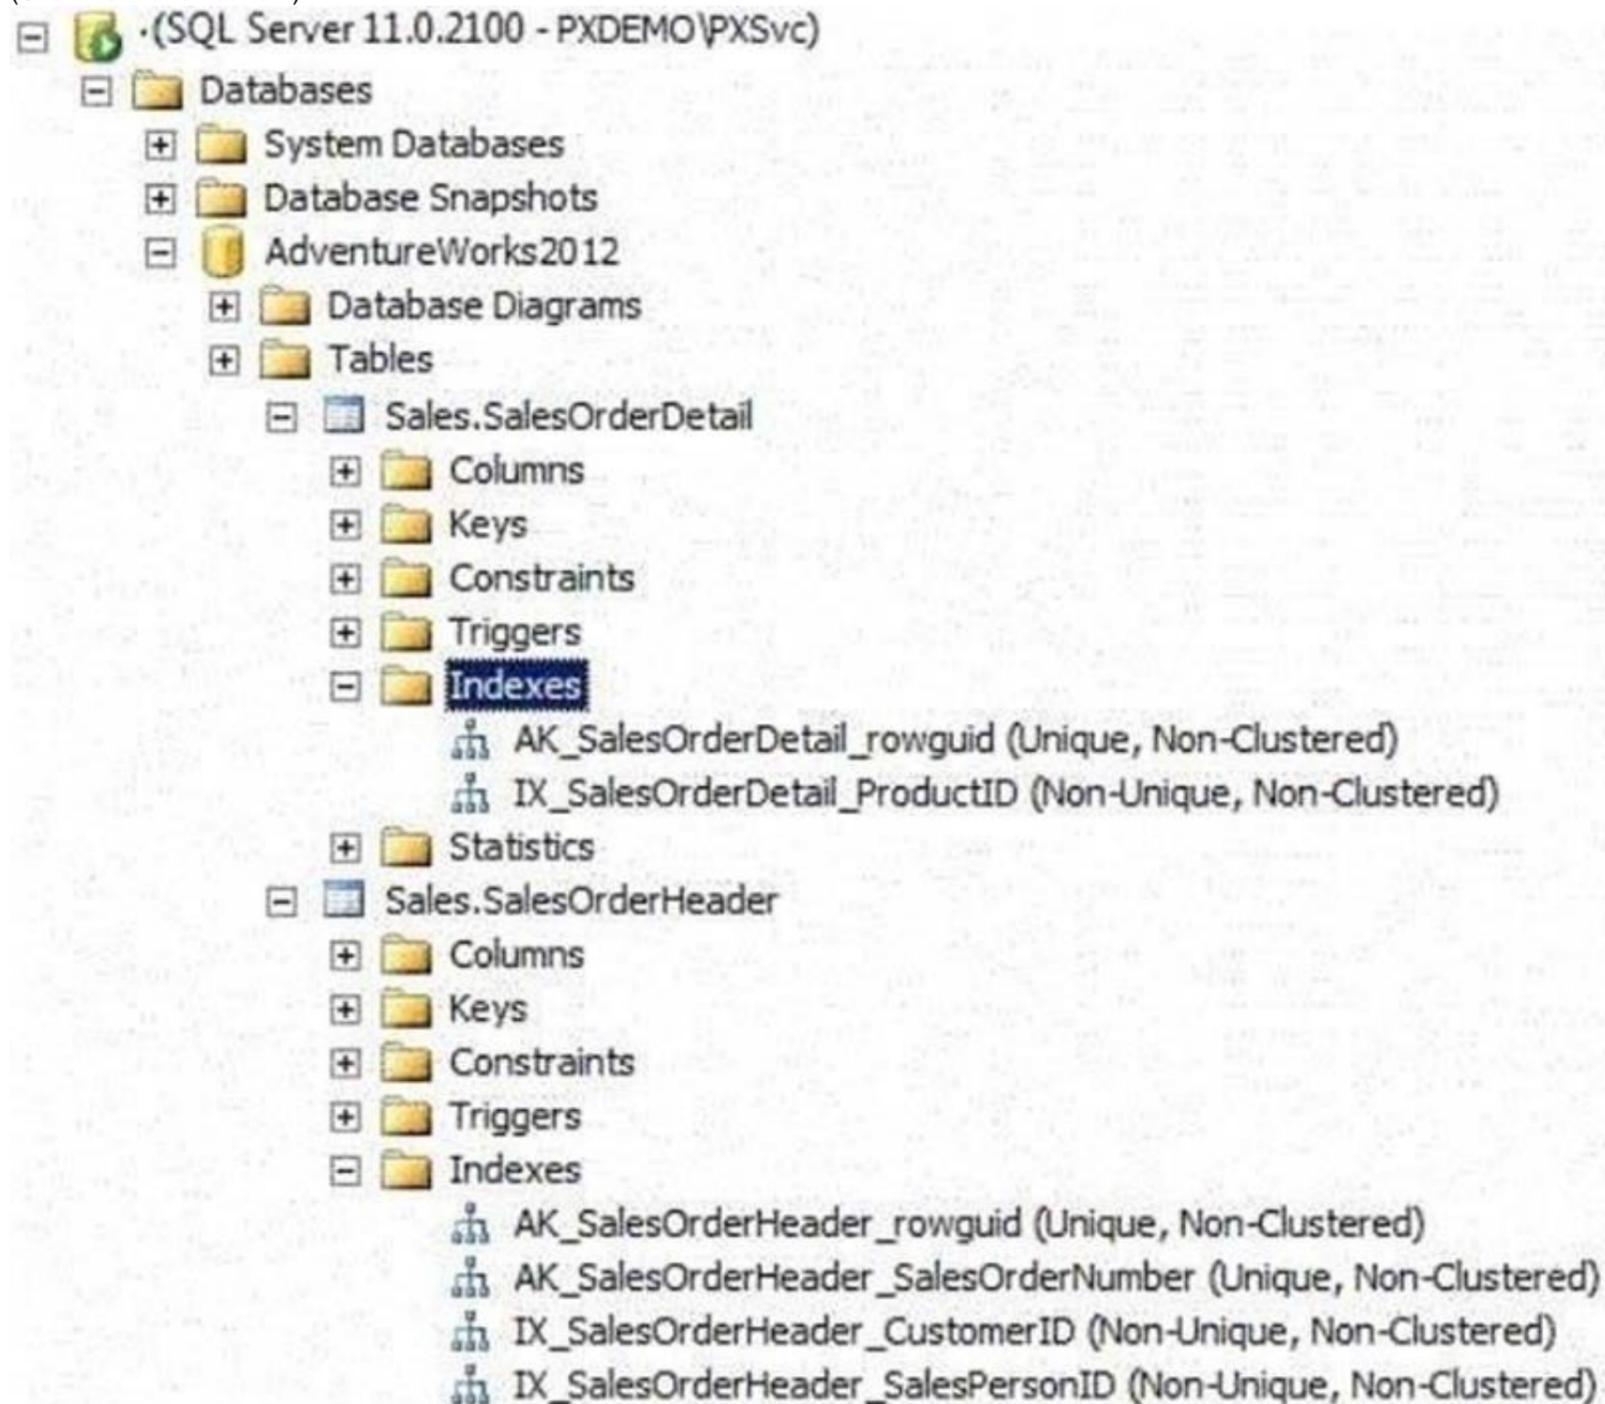
<https://docs.microsoft.com/en-us/sql/sql-server/stretch-database/enable-stretch-database-for-a-database?view=sq>

NEW QUESTION 60

- (Exam Topic 7)

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

(Click the Exhibit button.)



You write the following Transact-SQL query:

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

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

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

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

Answer: D

Explanation:

New statistics would be useful.

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

References:

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

NEW QUESTION 64

- (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 67

- (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 71

- (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 73

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database.

You configure Transparent Data Encryption (TDE) on the Orders database by using the following statements: CREATE MASTER KEY ENCRYPTION BY PASSWORD = 'MyPassword1!'

CREATE CERTIFICATE TDE_Certificate WITH SUBJECT = 'TDE Certificate'; BACKUP CERTIFICATE TDE_Certificate TO FILE = "d:\TDE_Certificate.cer" WITH PRIVATE KEY (FILE = 'D:\TDE_Certificate.key', ENCRYPTION BY PASSWORD = 'MyPassword1!'); CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES_256

ENCRYPTION BY SERVER CERTIFICATE TDE_Certificate;

ALTER DATABASE Orders SET ENCRYPTION ON;

You attempt to restore the Orders database and the restore fails. You copy the encryption file to the original location.

A hardware failure occurs and so a new server must be installed and configured.

After installing SQL Server to the new server, you restore the Orders database and copy the encryption files to their original location. However, you are unable to access the database.

You need to be able to restore the database.

Which Transact-SQL statement should you use before attempting the restore?

- A. ALTER DATABASE Master SET ENCRYPTION OFF;
- B. CREATE CERTIFICATE TDE_Certificate FROM FILE = 'd:\TDE_Certificate.cer' WITH PRIVATE KEY (FILE = 'D:\TDE_Certificate.key', DECRYPTION BY PASSWORD = 'MyPassword1!');
- C. CREATE CERTIFICATE TDE_Certificate WITH SUBJECT = 'TDE Certificate'; USE Orders; CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES_256 ENCRYPTION BY SERVER CERTIFICATE TDE_Certificate;
- D. CREATE CERTIFICATE TDE_Certificate FROM FILE = 'd:\TDE_Certificate.cer';

Answer: B

Explanation:

The CREATE CERTIFICATE command adds a certificate to a database in SQL Server. Creating a certificate from a file

The following example creates a certificate in the database, loading the key pair from files. Code

```
Copy
USE AdventureWorks2012; CREATE CERTIFICATE Shipping11
FROM FILE = 'c:\Shipping\Certs\Shipping11.cer'
WITH PRIVATE KEY (FILE = 'c:\Shipping\Certs\Shipping11.pvk', DECRYPTION BY PASSWORD = 'sldkflk34et6gs%53#v00');
GO
```

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

NEW QUESTION 76

- (Exam Topic 7)

You have a Microsoft Azure SQL Database server named server1-contoso.database.windows.net in a resource group named RG1.

You need to create an elastic pool.

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

```
$server =  -ServerName "server1-contoso" -ResourceGroupName "RG1"
```

```
$server =  -ElasticPoolName "Pool1" -Edition Premium
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: New-AzureRmSqlServer

Create an Azure SQL Database logical server using the New-AzureRmSqlServer command. A logical server contains a group of databases managed as a group.

Example:

```
New-AzureRmSqlServer -ResourceGroupName $resourcegroupname `
-ServerName $servername `
-Location $location ` etc.
```

Box 2: New-AzureRmSqlElasticPool

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

```
New-AzureRmSqlElasticPool -ResourceGroupName "ResourceGroup01" -ServerName "Server01"
-ElasticPoolName "ElasticPool01" -Edition "Standard"
```

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-get-started-powershell> <https://docs.microsoft.com/en-us/powershell/module/azurermsql/new-azurermsqlelasticpool?view=azurermps-6>

NEW QUESTION 77

- (Exam Topic 7)

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

An integrity check of a database has failed.

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

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

Answer: C

Explanation:

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

NEW QUESTION 78

- (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.

On Wednesday at 10:00 hours, the development team requests you to refresh the database on a development server by using the most recent version.

You need to perform a full database backup that will be restored on the development server. Which backup option should you 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: L

Explanation:

COPY_ONLY specifies that the backup is a copy-only backup, which does not affect the normal sequence of backups. A copy-only backup is created independently of your regularly scheduled, conventional backups. A copy-only backup does not affect your overall backup and restore procedures for the database.

References:

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

NEW QUESTION 80

- (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

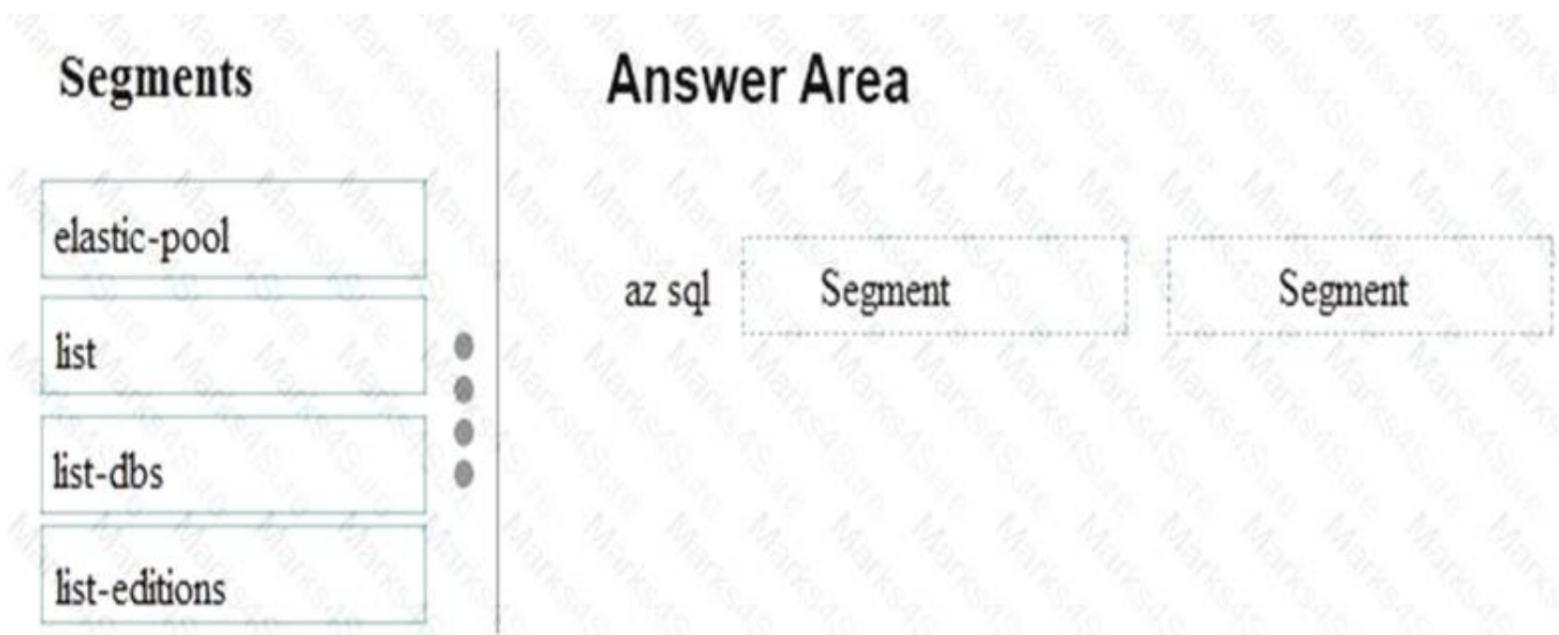
NEW QUESTION 83

- (Exam Topic 7)

Your company has several Microsoft Azure SQL Database instances used within an elastic pool. You need to obtain a list of databases in the pool.

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

NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/cli/azure/sql/elastic-pool?view=azure-cli-latest#az-sql-elastic-pool-list-dbs>

NEW QUESTION 85

- (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 88

- (Exam Topic 7)

You have a database named DB1 that uses simple recovery mode.

Full backups of DB1 are taken daily and DB1 is checked for corruption before each backup. There was no corruption when the last backup was complete.

You run the sys.columns catalog view and discover corrupt pages.

You need to recover the database. The solution must minimize data loss. What should you do?

- A. Run RESTORE DATABASE WITH RECOVERY.
- B. Run RESTORE DATABASE WITH PAGE.
- C. Run DBCC CHECKDB and specify the REPAIR_ALLOW_DATA_LOSS parameter.
- D. Run DBCC CHECKDB and specify the REPAIR_REBUILD parameter.

Answer: B

Explanation:

A page restore is intended for repairing isolated damaged pages. Restoring and recovering a few individual pages might be faster than a file restore, reducing the amount of data that is offline during a restore operation.

RESTORE DATABASE WITH PAGE

Restores individual pages. Page restore is available only under the full and bulk-logged recovery models. References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

NEW QUESTION 89

- (Exam Topic 7)

You have a database named DB1 that contains a table named Table1. Table1 has 1 billion rows.

You import 10 million rows of data into Table1. After the import, users report that queries take longer than usual to execute.

You need to identify whether an out-of-date execution plan is causing the performance issue. Which dynamic management view should you use?

- A. sys.dm_xtp_transaction_stats
- B. sys.dm_exec_input_buffer
- C. sys.dm_db_index_operational_stats

D. sys.dm_db_stats_properties

Answer: C

Explanation:

sys.dm_db_index_operational_stats dynamic management function provides us the current low-level I/O, locking, latching, and access method for each partition of the table. This information is really useful to troubleshoot SQL Server performance issues.

Reference:

https://basitaalishan.com/2013/03/19/using-sys-dm_db_index_operational_stats-to-analyse-howindexes-are-utili

NEW QUESTION 92

- (Exam Topic 7)

You are the database administrator in your company. You plan to create 10 identical environments that use SQL Server 2016 as a database engine. Each environment has the following custom requirements:

Three user databases must be preinstalled.

The tempdb database must contain eight data files that are 1024 MB each.

Trace flag 2371 must be turned at the instance level.

The solution must meet the following requirements:

The instance must be preconfigured.

No other database features are required in the future.

The solution must use the minimum administrative effort.

You need to prepare the environments. What should you do?

A. Provision 10 Azure virtual machines that each contain SQL Server 2016, installed by using the default settings.

B. Create an installation configuration file and perform unattended installations of SQL Server 2016.

C. Create a virtual machine template by using a prepared instance of SQL Server 2016.

D. Create a virtual machine template by using a complete instance of SQL Server 2016.

Answer: D

Explanation:

You should create a virtual machine template by using a complete instance of SQL Server 2016. You use the sysprep tool to prepare a complete instance of SQL Server 2016. By using a complete instance, SQL Server, the network, and the users are all created, and the system cannot be reconfigured during the installation process.

NEW QUESTION 97

- (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 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: G

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

NEW QUESTION 99

- (Exam Topic 7)

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

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

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

Answer area

File system:

	▼
FAT32	
FAT	
NTFS	

8.3 filename support:

	▼
Enabled	
Disabled	

Indexing:

	▼
Enabled	
Disabled	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

File System: NTFS

8.3 filename support: Disabled Indexing: Disabled

NTFS is required.

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

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

References:

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

NEW QUESTION 102

- (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 following requirements are met: 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: J

Explanation:

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

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

NEW QUESTION 106

- (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 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 113

- (Exam Topic 7)

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

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

Evenly distribute read and write operations across the VHDs.

Minimize the read and write time.

Which storage configuration should you use?

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

Answer: D

Explanation:

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

Reference:

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

NEW QUESTION 117

- (Exam Topic 7)

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

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

What should you do?

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

Answer: D

Explanation:

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

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

Local Key Stores References:

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

NEW QUESTION 122

- (Exam Topic 7)

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

One instance of SQL Server 2016 Enterprise

10 databases

500 stored procedures

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

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

You have the following requirements:

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

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

The solution must minimize administrative effort.

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

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

Answer: B

Explanation:

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

NEW QUESTION 127

- (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 132

- (Exam Topic 7)

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

References:

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

NEW QUESTION 134

- (Exam Topic 7)

You have database that contains a 400-GB table that is read-only. You need to enable the Stretch Database feature.

How should you complete the statement? To answer, drag the appropriate values to the correct targets. Each value 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.

NOTE: Each correct selection is worth one point. Select and Place:

Values

- OUTBOUND
- PARTITION
- PAUSED
- REMOTE_DATA_ARCHIVE

Answer area

```
ALTER TABLE table1 SET ( [ ] = ON
    FILTER_PREDICATE dbo.fn_stretchpredicate(order_date)
    MIGRATION_STATE = [ ]
))
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Section: Deploy and migrate applications

To configure an existing table for Stretch Database, run the ALTER TABLE command. Here's an example that migrates the entire table and begins data migration immediately. USE <Stretch-enabled database name>;

GO

ALTER TABLE <table name>

SET (REMOTE_DATA_ARCHIVE = ON (MIGRATION_STATE = OUTBOUND)) ; GO

References: <https://docs.microsoft.com/en-us/sql/sql-server/stretch-database/enable-stretch-database-for-atable>

NEW QUESTION 135

- (Exam Topic 7)

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

The following error message is an example error message:

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

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

- A. Configure output files.
- B. Expand agent logging to include information from all events.
- C. Disable the Limit size of job history log feature.
- D. Configure event forwarding.

Answer: B

Explanation:

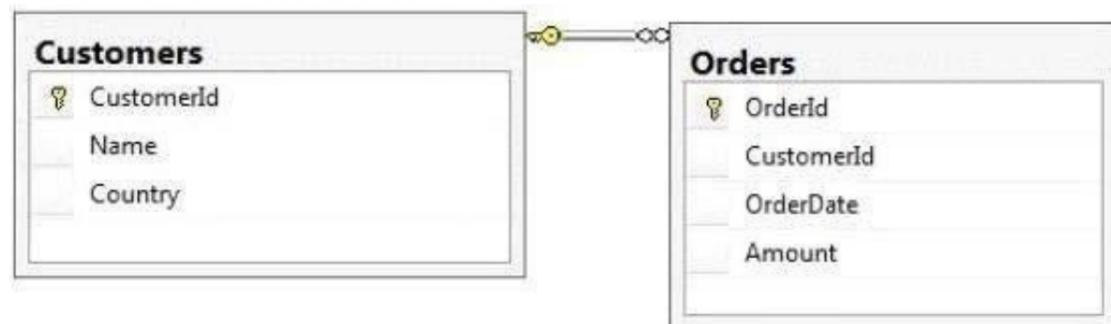
References:

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

NEW QUESTION 140

- (Exam Topic 7)

You 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, CountryFROM OrdersINNER JOIN CustomersON Orders.CustomerId = Customers.CustomerIdWHERE Customers.CustomerId = 1FOR XML RAW
- B. SELECT OrderId, OrderDate, Amount, Name, CountryFROM Orders INNER JOIN CustomersON Orders.CustomerId = Customers.CustomerIdWHERE Customers.CustomerId = 1FOR XML RAW, ELEMENTS
- C. SELECT OrderId, OrderDate, Amount, Name, CountryFROM OrdersINNER JOIN CustomersON Orders.CustomerId = Customers.CustomerIdWHERE Customers.CustomerId = 1FOR XML AUTO
- D. SELECT OrderId, OrderDate, Amount, Name, CountryFROM OrdersINNER JOIN CustomersON Orders.CustomerId = Customers.CustomerIdWHERE Customers.CustomerId= 1FOR XML AUTO, ELEMENTS
- E. SELECT Name, Country, OrderId, OrderDate, AmountFROM OrdersINNER JOIN CustomersON Orders.CustomerId= Customers.CustomerIdWHERE Customers.CustomerId= FOR XML AUTO
- F. SELECT Name, Country, CrderId, OrderDate, AmountFROM OrdersINNER JOIN CustomersON Orders.CustomerId= Customers.CustomerIdWHERE Customers.CustomerId= FOR XML AUTO, ELEMENTS
- G. SELECT Name AS `@Name`, Country AS `@Country`, OrderId, OrderDate, AmountFROM OrdersINNER JOIN CustomersON Orders.CustomerId= Customers.CustomerIdWHERE Customers.CustomerId = 1FOR XML PATH (`Customers`)

H. SELECT Name AS `Customers/Name`, CountryAS `Customers/Country`, OrderId, OrderDate, AmountFROM OrdersINNER JOIN CustomersON Orders.CustomerId= Customers.CustomerIdWHERE Customers.CustomerId= 1FOR XML PATH (`Customers`)

Answer: G

NEW QUESTION 145

- (Exam Topic 7)

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

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

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

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

- A. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer cINNER JOIN(SELECT CustomerID, ShippingCountry,RANK() OVER (PARTITION BY CustomerIDORDER BY COUNT(OrderAmount) DESC) AS RnkFROM OrdersGROUP BY CustomerID, ShippingCountry) AS oON c.CustomerID = o.CustomerIDWHERE o.Rnk = 1
- B. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM(SELECT c.CustomerID, c.CustomerName, o.ShippingCountry, RANK()OVER (PARTITION BY CustomerIDORDER BY COUNT(o.OrderAmount) ASC) AS RnkFROM Customer cINNER JOIN Orders oON c.CustomerID = o.CustomerIDGROUP BY c.CustomerID, c.CustomerName, o.ShippingCountry) cs WHERE Rnk = 1
- C. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer cINNER JOIN(SELECT CustomerID, ShippingCountry,RANK() OVER (PARTITION BY CustomerIDORDER BY OrderAmount DESC) AS RnkFROM OrdersGROUP BY CustomerID, ShippingCountry) AS oON c.CustomerID = o.CustomerIDWHERE o.Rnk = 1
- D. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer cINNER JOIN(SELECT CustomerID, ShippingCountry,COUNT(OrderAmount) DESC) AS OrderAmountFROM OrdersGROUP BY CustomerID, ShippingCountry) AS oON c.CustomerID = o.CustomerIDORDER BY OrderAmount DESC

Answer: A

Explanation:

Use descending (DESC) ordering.

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

NEW QUESTION 149

- (Exam Topic 7)

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

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

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

Which Transact-SQL statement should you use?

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

Answer: E

Explanation:

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

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

NEW QUESTION 152

- (Exam Topic 7)

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

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

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

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

Answer: D

NEW QUESTION 155

- (Exam Topic 7)

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

The various databases on the instances have the following characteristics:

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

You run the following Transact-SQL statements:

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

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

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

Answer: D

NEW QUESTION 160

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance.

You need to configure a new database to support FILETABLES. What should you do? Choose all that apply.

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

Answer: BDF

Explanation:

B: FileTables extend the capabilities of the FILESTREAM feature of SQL Server. Therefore you have to enable FILESTREAM for file I/O access at the Windows level and on the instance of SQL Server before you can create and use FileTables.

D: Before you can create FileTables in a database, the database must have a FILESTREAM filegroup. F: Specifying a Directory for FileTables at the Database Level

When you enable non-transactional access to files at the database level, you can optionally provide a directory name at the same time by using the DIRECTORY_NAME option. If you do not provide a directory name when you enable non-transactional access, then you have to provide it later before you can create FileTables in the database.

References:

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

NEW QUESTION 164

- (Exam Topic 7)

You manage a Microsoft SQL Server instance named SQL1 that has 32 gigabytes (GB) of total memory. The instance supports an app named App1 that only uses a single thread. App1 frequently queries the database using the same index. The operating system and App1 combined require 8 GB of memory to function. You need to ensure that the SQL Server does not limit the performance of App1. What configuration option should you set?

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 165

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 server. You plan to deploy new features to an application. You need to evaluate existing and potential clustered and non-clustered indexes that will improve performance. What should you do?

- A. Query the sys.dm_db_index_usage_stats DMV.
- B. Query the sys.dm_db_missing_index_details DMV.
- C. Use the Database Engine Tuning Advisor.
- D. Query the sys.dm_db_missing_index_columns DMV.

Answer: C

Explanation:

The Microsoft Database Engine Tuning Advisor (DTA) analyzes databases and makes recommendations that you can use to optimize query performance. You can use the Database Engine Tuning Advisor to select and create an optimal set of indexes, indexed views, or table partitions without having an expert understanding of the database structure or the internals of SQL Server.

NEW QUESTION 170

- (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 171

- (Exam Topic 7)

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

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

You are migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 databases that consume a total of 2 TB of disk space.

The instance sustains more than 30,000 transactions per second.

You need to provision storage for the virtual machine. The storage must be able to support the same load as the on-premises deployment.

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

- A. Yes
- B. No

Answer: A

Explanation:

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

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

NEW QUESTION 173

- (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 174

- (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 177

- (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 182

- (Exam Topic 7)

You administer two Microsoft SQL Server 2014 servers. Each server resides in a different, untrusted domain. You plan to configure database mirroring. You need to be able to create database mirroring endpoints on both servers. What should you do?

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 185

- (Exam Topic 7)

You are using dynamic management views to monitor an SQL Server server named SQL1. A database administrator named Dba1 must monitor the health of SQL1.

You need to ensure that Dba1 can access dynamic management views for SQL1. The solution must use the principle of least privilege.

Which permissions should you assign to Dba1?

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

Answer: B

Explanation:

To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission.

There are two types of dynamic management views and functions:

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

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

References:

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

NEW QUESTION 189

- (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.

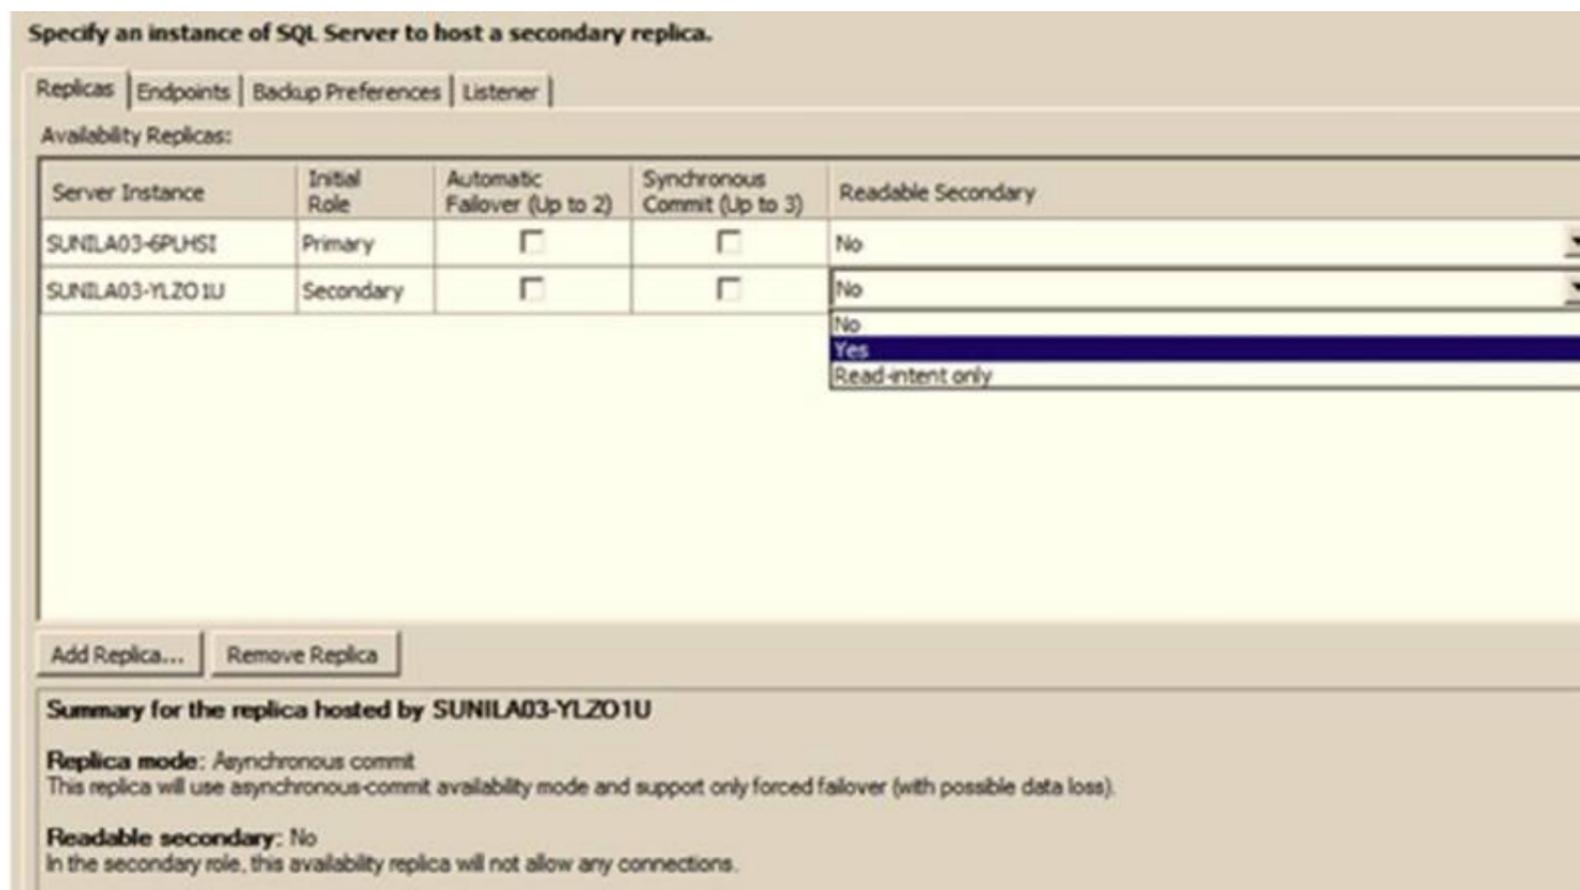
You need to offload the reporting workload to the secondary replica when HA/Server01 is the primary replica. What should you do?

- 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 193

- (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 196

- (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 200

- (Exam Topic 7)

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

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

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

Which option of the BACKUP LOG command should you use?

- A. NO_TRUNCATE
- B. NORECOVERY
- C. STANDBY
- D. FORMAT

Answer: B

Explanation:

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

* If the database is online and you plan to perform a restore operation on the database, begin by backing up the tail of the log. To avoid an error for an online database, you must use the ... WITH NORECOVERY option of the BACKUP Transact-SQL statement.

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

References:

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

NEW QUESTION 205

- (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 207

- (Exam Topic 7)

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

You need to configure the network to support the availability group listener. Which cmdlet should you run first?

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 210

- (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 212

- (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 214

.....

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>