

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 deployed a GS-series virtual machine (VM) in Microsoft Azure. You plan to deploy Microsoft SQL Server.

You need to deploy a 30 megabyte (MB) database that requires 100 IOPS to be guaranteed while minimizing costs.

Which storage option should you use?

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

Answer: A

Explanation:

Premium Storage Disks Limits

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

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

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

NEW QUESTION 3

- (Topic 1)

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

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

You need to create the templates.

Solution: You use Visual Studio to create a JSON template that defines the deployment and configuration settings for the SQL Server environment.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

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 4

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 5

HOTSPOT - (Topic 1)

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

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

You have the following requirements:

You need to deploy the SQL instance.

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

Answer area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Enable read caching on the disk(s) hosting the data files and TempDB.

Do not enable caching on disk(s) hosting the log file. Host caching is not used for log files.

NEW QUESTION 6

DRAG DROP - (Topic 2)

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

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

You need to create the new security principal.

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1, Step 2:

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

Step 3:

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

`CREATE USER readonlyuser FROM LOGIN readonlylogin`; References: <https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/>

NEW QUESTION 7

- (Topic 2)

You plan to deploy 20 Microsoft Azure SQL Database instances to an elastic pool in Azure to support a batch processing application. Two of the databases in the pool reach their peak workload threshold at the same time every day. This leads to inconsistent performance for batch completion. You need to ensure that all batches perform consistently. What should you do?

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 8

- (Topic 2)

Note: This 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 manage on-premises and Microsoft Azure SQL Database instances for a company. Your environment must support the Microsoft SQL Server 2012 ODBS

driver.

You need to encrypt only specific columns in the database. What should you implement?

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

Answer: D

Explanation:

To encrypt columns you can configure Always Encrypted.

SQL Server Management Studio (SSMS) provides a wizard that helps you easily configure Always Encrypted by setting up the column master key, column encryption key, and encrypted columns for you.

Always Encrypted allows client applications to encrypt sensitive data and never reveal the data or the encryption keys to SQL Server or Azure SQL Database. An Always Encrypted enabled driver, such as the ODBC Driver 13.1 for SQL Server, achieves this by transparently encrypting and decrypting sensitive data in the client application.

Note: The ODBC driver automatically determines which query parameters correspond to sensitive database columns (protected using Always Encrypted), and encrypts the values of those parameters before passing the data to SQL Server or Azure SQL Database. Similarly, the driver transparently decrypts data retrieved from encrypted database columns in query results.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault#encrypt-columns-configure-always-encrypted>
[https://msdn.microsoft.com/en-us/library/mt637351\(v=sql.110\).aspx](https://msdn.microsoft.com/en-us/library/mt637351(v=sql.110).aspx)

NEW QUESTION 9

- (Topic 2)

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

You need to configure the key store provider.

What should you do?

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

Answer: D

Explanation:

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

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

NEW QUESTION 10

- (Topic 2)

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

You manage a Microsoft SQL Server environment with several databases.

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

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

- A. Yes
- B. No

Answer: B

Explanation:

LEGACY_CARDINALITY_ESTIMATION = { ON | OFF | PRIMARY }

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

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

NEW QUESTION 10

DRAG DROP - (Topic 2)

A new Azure Active Directory security principal named ReportUser@contoso.onmicrosoft.com should have access to select all current and future objects in the Reporting database. You should not grant the principal any other permissions. You should use your Active Directory Domain Services (AD DS) account to authenticate to the Azure SQL database.

You need to create the new security principal.

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

Actions

Answer Area

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

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

Run the following Transact-SQL statement:

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

Run the following Transact-SQL statement:

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

Run the following Transact-SQL statements:

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1:

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

Step 2: CREATE USER ... FROM EXTERNAL PROVIDER

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

```
CREATE USER <Azure_AD_principal_name> FROM EXTERNAL PROVIDER;
```

Step 3:

Grant the proper reading permissions.

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

NEW QUESTION 13

- (Topic 2)

You are deploying a Microsoft SQL Server database that will support a mixed OLTP and OLAP workload. The target virtual machine has four CPUs. You need to ensure that reports do not use all available system resources. What should you do?

- A. Enable Auto Close.
- B. Increase the value for the Minimum System Memory setting.
- C. Set MAXDOP to half the number of CPUs available.
- D. Increase the value for the Minimum Memory per query setting.

Answer: C

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 option to limit the number of processors to use in parallel plan execution.

NEW QUESTION 14

- (Topic 3)

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

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

Users report that the queries used to provide data for the financial reports take a long time to complete. The queries consume the majority of CPU and memory

resources on the database server. As a result, read-write queries for the application also take a long time to complete. You need to improve performance of the application while still allowing the report queries to finish.
Solution: You configure the Resource Governor to 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 16

- (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 6)

You need to open the firewall ports for use with SQL Server environment. In table below, identify the firewall port that you must use for each service. NOTE: Make only one selection in each column.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 22

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

- Reporting
- master
- CREATE USER
- ALTER LOGIN
- ALTER USER
- FOR LOGIN [BIAppUser]
- FOR USER [BIAppUser]
- WITH LOGIN = [BIAppUser]

Answer area

```
USE [ Code segment ]
GO
Code segment [BIAppUser] Code segment
GO
```

- 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 26

- (Exam Topic 7)

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

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

Answer: A

Explanation:

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

References:

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

NEW QUESTION 29

- (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)

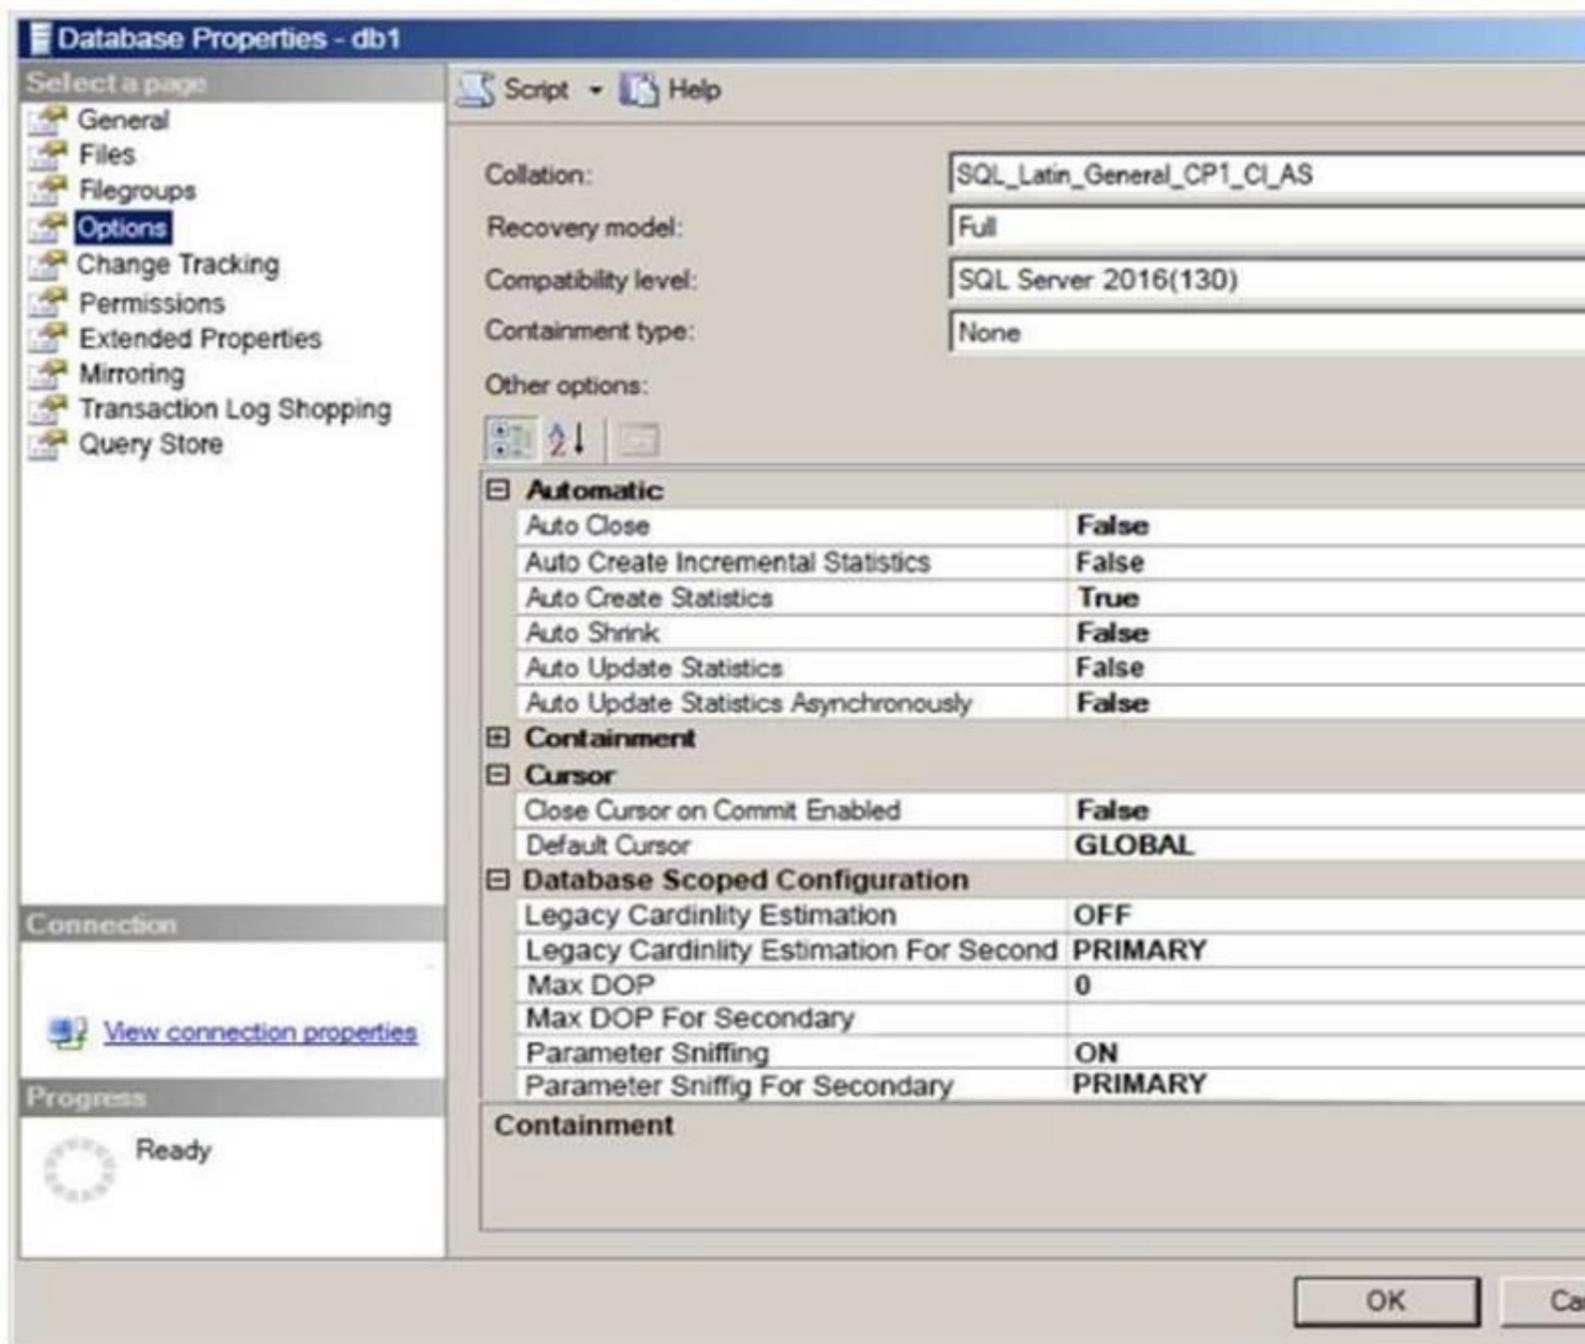
NEW QUESTION 32

- (Exam Topic 7)

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

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

You review the execution plan for the query and discover that the expected number of returned rows is one, while the actual number of returned rows is 1 million. You need to reduce the amount of time it takes for the query to execute. The solution must prevent additional performance issues from being introduced. Hot Area:



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

References:

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

NEW QUESTION 34

- (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 36

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database.

You have a SQL Server Agent job instance that runs using the service account. You have a job step within the job that requires elevated privileges. You need to ensure that the job step can run using a different user account. What should you use?

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

Answer: D

Explanation:

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

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

NEW QUESTION 39

- (Exam Topic 7)

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

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

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

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

Answer: BD

Explanation:

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

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

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

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

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

NEW QUESTION 44

- (Exam Topic 7)

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

What should you create?

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

Answer: B

Explanation:

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

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

NEW QUESTION 49

- (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 52

- (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.

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:

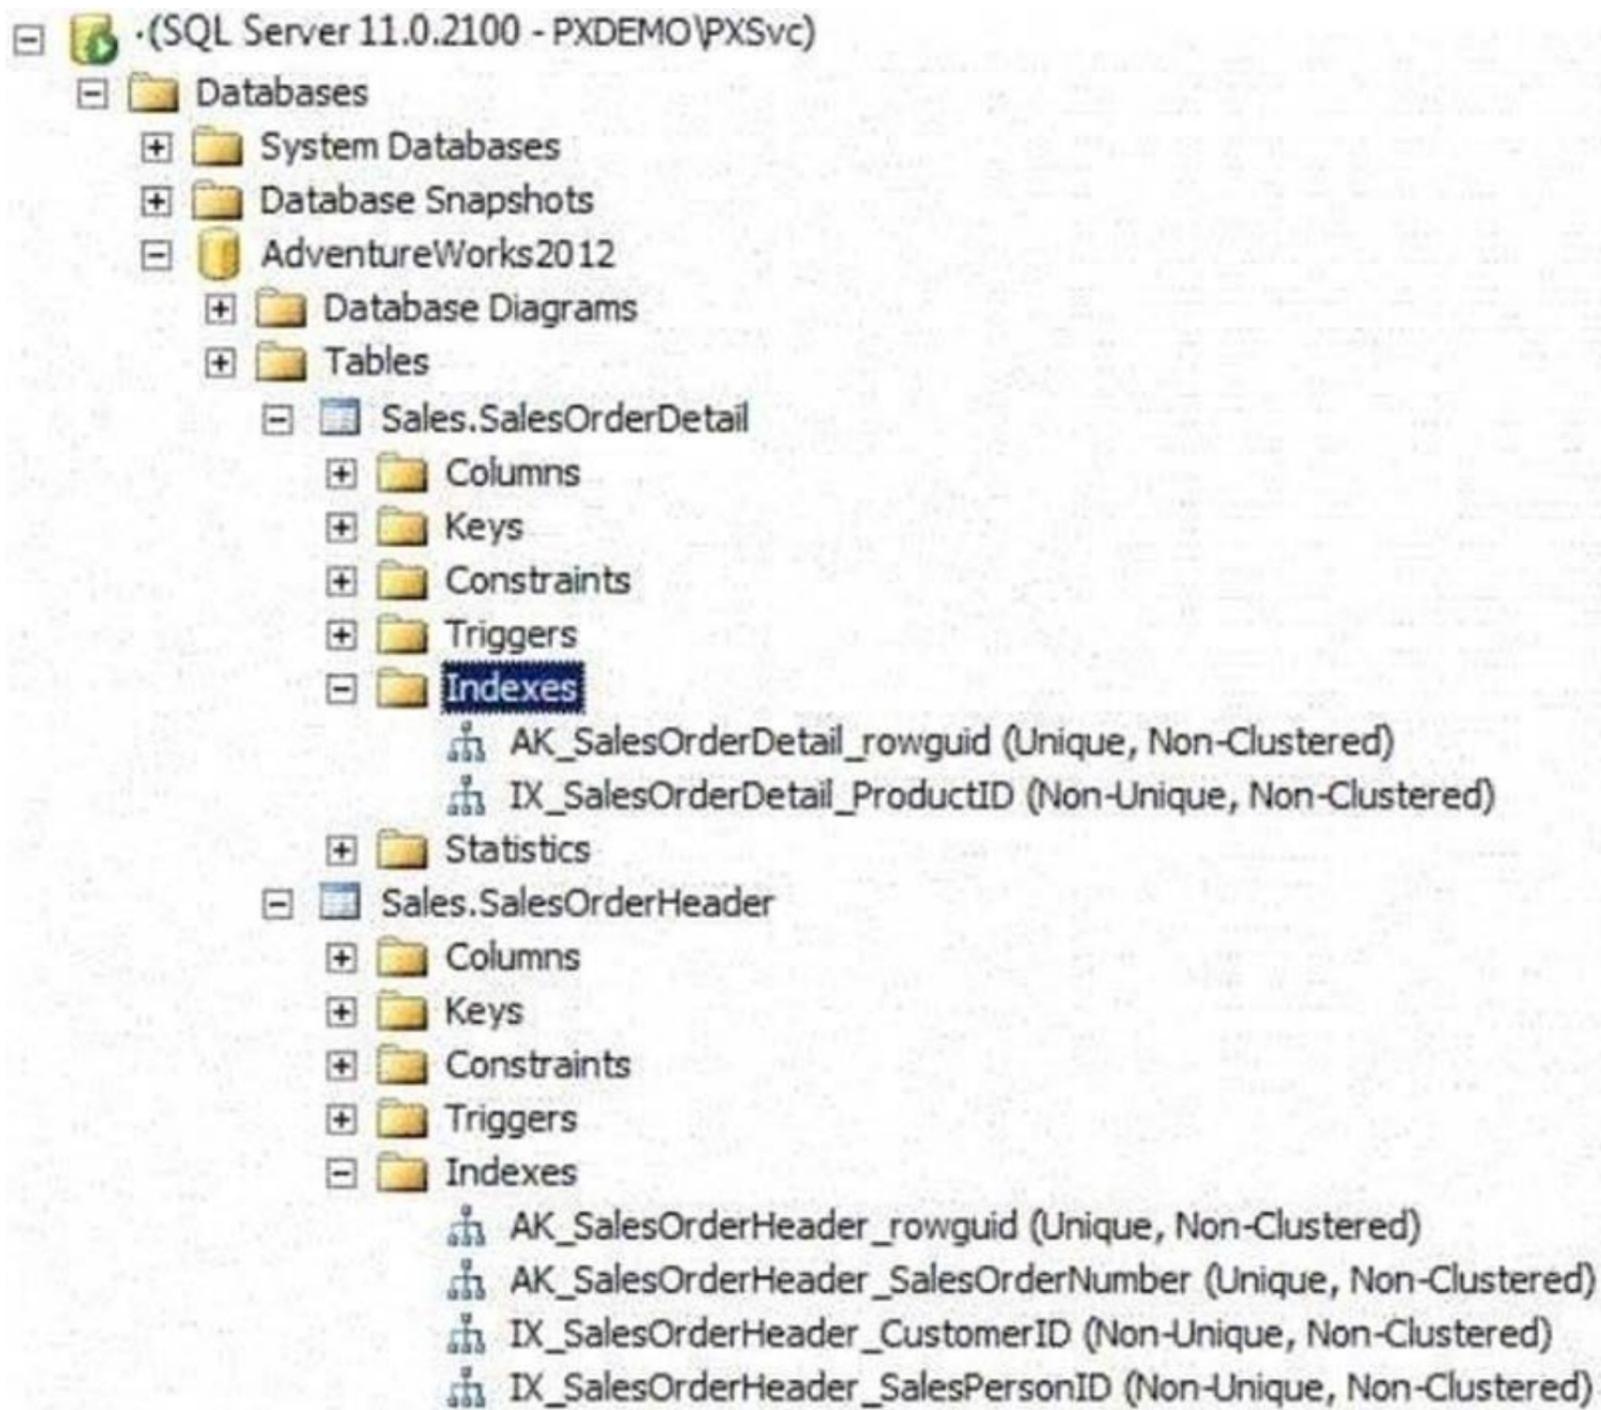
<https://docs.microsoft.com/en-us/sql/relational-databases/databases/model-database?view=sql-server-2017> <https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-transact-sql-compatibility-level?view=sql-se>

NEW QUESTION 53

- (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 54

- (Exam Topic 7)

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

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

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

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

Performance on this procedure is slow.

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

Which Transact-SQL statement should you use?

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

Answer: B

Explanation:

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

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

NEW QUESTION 59

- (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 64

- (Exam Topic 7)

You have been hired as a Database Consultant by ABC.com to design a SQL Server 2014 database solution. You are tasked with designing a scale-out and high-availability SQL Server 2014 Online Transaction

Processing (OLTP) database solution that will maintain copies of data across two server instances.

Your solution must provide scale-out of read operations by distributing the reads from clients across two SQL Server 2014 nodes. The data in both SQL Server nodes needs to be indexed.

What should you include in your solution?

- A. You should include a primary database with scheduled log shipping to the secondary database configured.
- B. You should include two servers configured in an Active-Passive SQL Server 2014 Cluster.
- C. You should include a primary SQL Server 2014 database that uses transactional replication to replicate data to the secondary database.
- D. You should include two servers in an Asynchronous-Commit Availability Mode Availability Group.
- E. You should include two servers in a Synchronous-Commit Availability Mode Availability Group.

Answer: C

Explanation:

Peer-to-peer replication provides a scale-out and high-availability solution by maintaining copies of data across multiple server instances, also referred to as nodes.

Built on the foundation of transactional replication, peer-to-peer replication propagates transactionally consistent changes in near real-time. This enables applications that require scale-out of read operations to distribute the reads from clients across multiple nodes. Because data is maintained across the nodes in near real-time, peer-to-peer replication provides data redundancy, which increases the availability of data.

References: [https://msdn.microsoft.com/en-us/library/ms151196\(v=sql.110\)](https://msdn.microsoft.com/en-us/library/ms151196(v=sql.110))

NEW QUESTION 69

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft Azure virtual machine that has 12 databases. All database files are in the same Azure Blob storage account. You need to receive an email notification if I/O operations to the database files exceed 800 MB/s for more than five minutes.

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

Does this meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 73

- (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 75

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 failover cluster.

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 78

- (Exam Topic 7)

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

Which cmdlet should you run first?

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

Answer: G

Explanation:

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

References:

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

NEW QUESTION 82

- (Exam Topic 7)

You 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 87

- (Exam Topic 7)

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

What should you create?

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 92

- (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 94

- (Exam Topic 7)

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

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

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

What should you do?

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

Answer: D

Explanation:

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

References:

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

NEW QUESTION 95

- (Exam Topic 7)

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

You plan to configure Instant File Initialization.

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

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

Answer: CD

Explanation:

How To Enable Instant File Initialization References:

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

NEW QUESTION 99

- (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 Add-AzureRmMetricAlertRule cmdlet and specify the -MetricName 'Network Out' parameter.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 103

- (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-in-windows-server-20>

NEW QUESTION 107

- (Exam Topic 7)

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

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

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

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

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

Solution: You create one storage account that has one container. You create multiple VHDs in the container. Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

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

NEW QUESTION 108

- (Exam Topic 7)

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

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

What should you do?

A. Back up the master database on the SQL2008 instance

B. Restore the master database on the SQL2012 instance

C. Use the Transfer Logins task in a Microsoft SQL Server Integrated Services package

D. Use sp_grantlogin

E. Use xp_logininfo.

Answer: C

Explanation:

sp_grantlogin creates a SQL Server login.

NEW QUESTION 110

- (Exam Topic 7)

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

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

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

Which edition should you choose?

A. Express Edition with Tools

B. Standard Edition

C. Web Edition

D. Express Edition with Advanced Services

Answer: B

NEW QUESTION 112

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014.

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

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

A. EXEC sp_who 60

B. SELECT * FROM sys.dm_exec_sessions WHERE sessionid = 60

C. EXEC sp_helpdb 60

D. DBCC INPUTBUFFER (60)

Answer: A

Explanation:

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

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

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

NEW QUESTION 114

- (Exam Topic 7)

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 118

- (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 119

- (Exam Topic 7)

Background

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

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

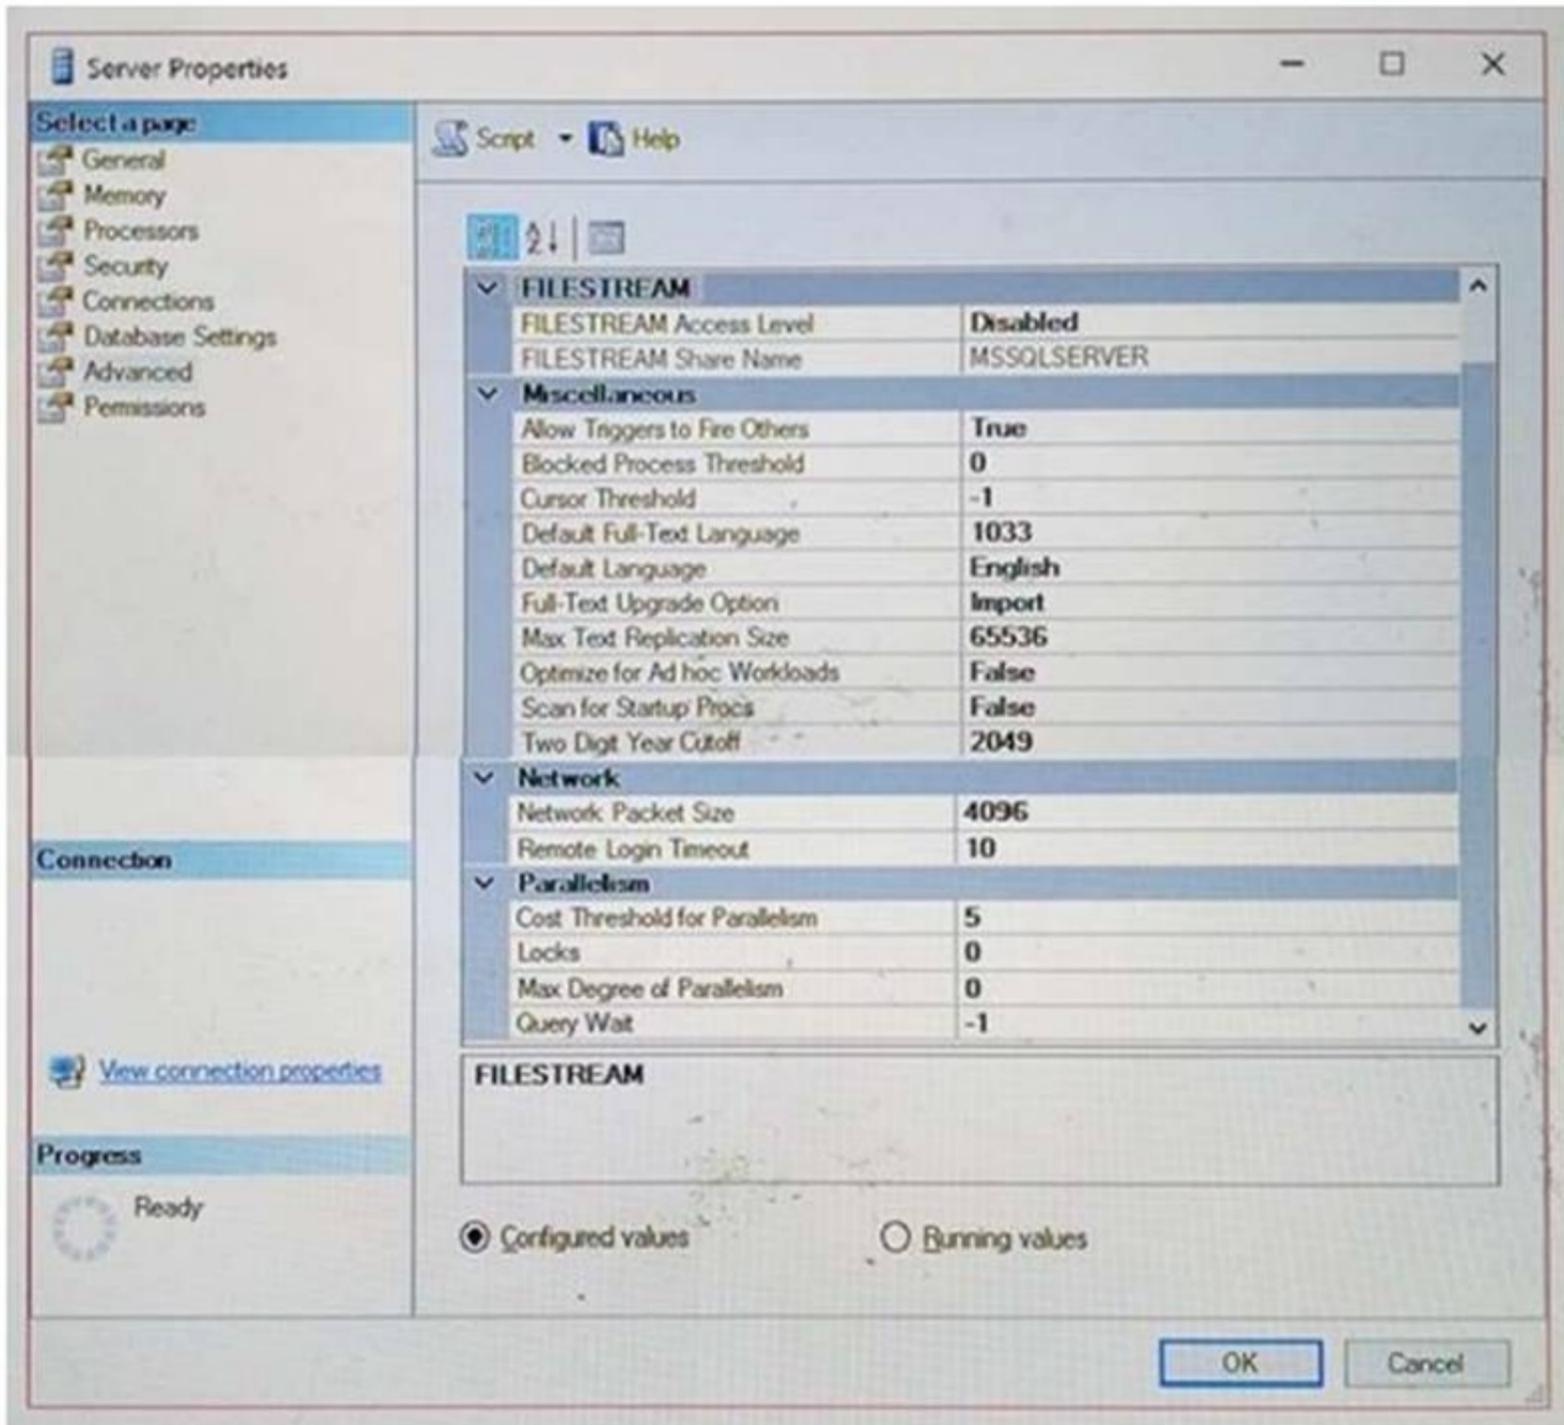
The environment also includes the following databases: DB1, DB2, and Reporting. The Reporting database is protected with Transparent Data Encryption (TDE).

You plan to migrate this database to a new server. You detach the database and copy it to the new server.

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

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

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



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

Values

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

Answer area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 2: Create: server certificate

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

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

-- Attach the database that is being moved.

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

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

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

References:

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

NEW QUESTION 124

- (Exam Topic 7)

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

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

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

- A. Configure Resource Pools, Workload Groups, and Classifier Function, and then enable the Resource Governor
- B. Set up Service Broker to ensure that application are not allowed to consume more than the specified amount of resource
- C. Create a new rule for each application that sets the resource limit allowed
- D. Create a new plan Guide with a Scope Type of sql and define the resource limits for each application

Answer: A

Explanation:

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

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

NEW QUESTION 128

- (Exam Topic 7)

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

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

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

- A. Local storage
- B. Page blob storage
- C. Virtual disks
- D. Block blob storage.

Answer: D

Explanation:

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

References:

NEW QUESTION 129

- (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 131

- (Exam Topic 7)

You are designing a Windows Azure SQL Database for an order fulfillment system. You create a table named Sales.Orders with the following script.

```
CREATE TABLE Sales.Orders
(
    OrderID int IDENTITY(1,1) NOT NULL PRIMARY KEY,
    OrderDate datetimeoffset NOT NULL,
    CustomerID int NOT NULL
);
```

Each order is tracked by using one of the following statuses:

Fulfilled

Shipped
Ordered
Received

You need to design the database to ensure that that you can retrieve the following information:

- The current status of an order
- The previous status of an order.
- The date when the status changed.
- The solution must minimize storage.

More than one answer choice may achieve the goal. Select the BEST answer.

- A. To the Sales.Orders table, add three columns named Status, PreviousStatus and ChangeDat
- B. Update rows as the order status changes.
- C. Create a new table named Sales.OrderStatus that contains three columns named OrderID, StatusDate, and Statu
- D. Insert new rows into the table as the order status changes.
- E. Implement change data capture on the Sales.Orders table.
- F. To the Sales.Orders table, add three columns named FulfilledDate, ShippedDate, and ReceivedDate.Update the value of each column from null to the appropriate date as the order status changes.

Answer: A

Explanation:

This stores only the minimal information required.

NEW QUESTION 135

- (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 137

- (Exam Topic 7)

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 140

- (Exam Topic 7)

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

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

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

are longer than 50 ms.

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

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

NEW QUESTION 142

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database instance.

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

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

What should you do?

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

Answer: C

Explanation:

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

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

NEW QUESTION 145

- (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 148

- (Exam Topic 7)

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

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

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

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

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

Answer: H

Explanation:

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

NEW QUESTION 153

- (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 156

- (Exam Topic 7)

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

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

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

The branch offices must not be able to access DB1.

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

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

Answer: BC

Explanation:

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

You should run ALTER DATABASE DB1 SET ENCRYPTION ON.

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

NEW QUESTION 161

- (Exam Topic 7)

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

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

- A. Move user databases to a local disk.
- B. Expand the tempdb data and log files
- C. Modify application code to use table variables
- D. Move the tempdb files to a local disk

Answer: D

Explanation:

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

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

NEW QUESTION 164

- (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 167

- (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 168

- (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 171

- (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 175

- (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 178

- (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 181

- (Exam Topic 7)

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

You need to ensure that the new instance retains the existing jobs and alerts. Solutions: You restore the service master key.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

The Service Master Key is the root of the SQL Server encryption hierarchy. It does not handle alerts and jobs. The msdb database is used by SQL Server Agent for scheduling alerts and jobs and by other features such as

SQL Server Management Studio, Service Broker and Database Mail.

References:

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

NEW QUESTION 182

.....

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>