

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 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)

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 6

- (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 ResourceManager 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 7

- (Topic 1)

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

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

Answer: D

Explanation:

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

Premium service tier

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

NEW QUESTION 8

- (Topic 2)

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

You manage a Microsoft SQL Server environment with several databases.

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 9

DRAG DROP - (Topic 2)

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

You need to create the new security principal.

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

Actions

Answer Area

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

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

Run the following Transact-SQL statement:

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

Run the following Transact-SQL statement:

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

Run the following Transact-SQL statements:

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1:

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

Step 2: CREATE USER ... FROM EXTERNAL PROVIDER

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

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

Step 3:

Grant the proper reading permissions.

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

NEW QUESTION 10

- (Topic 3)

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

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

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

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

Solution: You configure the Resource Governor to limit the amount of memory, CPU, and IOPS used for the pool of all queries that the Reporting_user login can run concurrently.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

NEW QUESTION 10

- (Topic 3)

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

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

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

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

Solution: You create a snapshot of the database. You configure all report queries to use the database snapshot.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Use a Resource Governor instead.

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

NEW QUESTION 13

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

What setting would you change to reduce the number of execution plans in the plan cache?

- Optimize for Ad Hoc workload ▼
- Max Degree of Parallelism
- Query Wait

What setting would you change to which value to reduce the number of queries which are using parallelism?

- Max Degree of Parallelism to 4 ▼
- Cost Threshold for Parallelism to 50
- Locks to 100

- 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 18

HOTSPOT - (Topic 4)

You need to optimize SRV1.

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

Answer Area

How should you modify the tempdb configuration? ▼

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

How should you reconfigure the tempdb database? ▼

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

Box 2: Add additional tempdb files.

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

NEW QUESTION 22

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

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 27

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 29

- (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 30

- (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 31

- (Exam Topic 7)

Settings Value VM size D3

Storage Location Drive E Storage type Standard Tempdb location Drive C

The workload on this instance has of the tempdb load.

You need to maximize the performance of the tempdb database.

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

A. Yes

B. No

Answer: B

Explanation:

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

References:

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

NEW QUESTION 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 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 38

- (Exam Topic 7)

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

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

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

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

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 43

- (Exam Topic 7)

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

What should you create?

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

Answer: B

Explanation:

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

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

NEW QUESTION 47

- (Exam Topic 7)

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

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

Solution: You move the database to a Premium database pool that has 125 eDTUs. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

We need at least 400 eDTUs.

NEW QUESTION 49

- (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	100
Msdb	simple	100
Model	full	100
Resource	simple	100
Tempdb	simple	100

- 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 52

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 Enterprise Edition server that uses 64 cores.

You discover performance issues when large amounts of data are written to tables under heavy system load. You need to limit the number of cores that handle I/O.

What should you configure?

- A. Processor affinity
- B. Lightweight pooling
- C. Max worker threads
- D. I/O affinity

Answer: D

Explanation:

The affinity Input-Output (I/O) mask Server Configuration Option.

To carry out multitasking, Microsoft Windows 2000 and Windows Server 2003 sometimes move process threads among different processors. Although efficient from an operating system point of view, this activity can reduce Microsoft SQL Server performance under heavy system loads, as each processor cache is repeatedly reloaded with data. Assigning processors to specific threads can improve performance under these conditions by eliminating processor reloads; such an association between a thread and a processor is called processor affinity.

References:

<http://msdn.microsoft.com/en-us/library/ms189629.aspx>

NEW QUESTION 56

- (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 60

- (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 63

- (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 67

- (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.

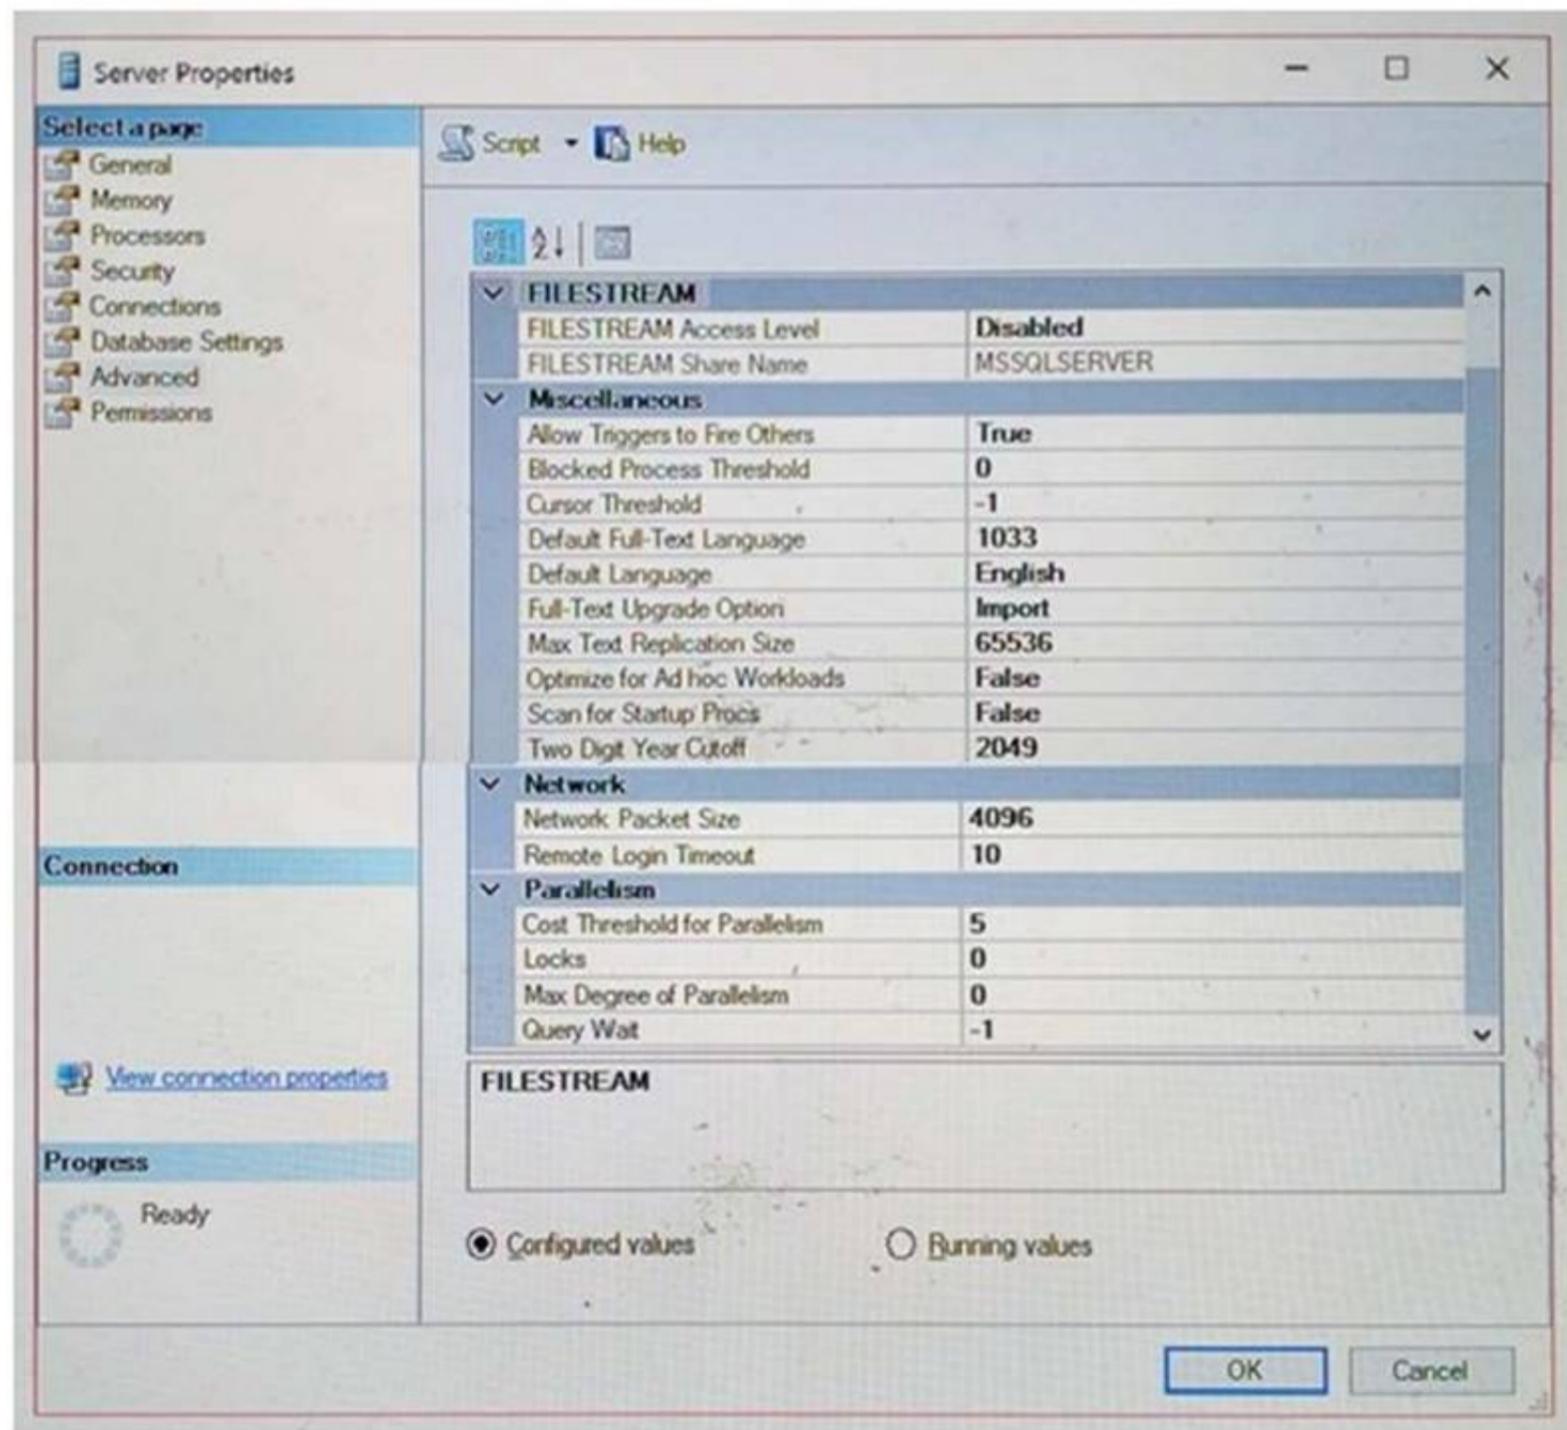
The environment also includes the following databases: DB1, DB2, and Reporting. The Reporting database is protected with Transparent Data Encryption (TDE).

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 set the size of the log files for the tempdb database on SRV1.
How should you complete the Transact-SQL statement? To answer, select the appropriate Transact-SQL segments in the answer area.
Hot Area:

Answer Area

[tempdb] (NAME =N'templog', SIZE = 6553

UPDATE	MODIFY FILE
ALTER	UPDATE FILE

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The ALTER DATABASE with MODIFY FILE command can make a file size bigger (but not smaller). Example:

ALTER DATABASE AdventureWorks2012 MODIFY FILE

(NAME = test1dat3, SIZE = 200MB); Note: MODIFY FILE

Specifies the file that should be modified. Only one <filespec> property can be changed at a time. NAME must always be specified in the <filespec> to identify the file to be modified. If SIZE is specified, the new size must be larger than the current file size.

References:

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

NEW QUESTION 69

- (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 73

- (Exam Topic 7)

You have an on-premises database.

You plan to migrate the database to Microsoft SQL Server on a Microsoft Azure virtual machine.

You move the database files to Azure.

You need to attach the database files to the SQL Server instance on the virtual machine. The solution must ensure that you can run file snapshot backups.

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

Answer area

```
USE (master)
```

```
GO
```

```
CREATE DATABASE [Production_DB]
```

```
( [ ] = N'https://proddbstorage=contoso.blob.core.windows.net/datafiles/proddb.mdf'
```

- DISK
- NAME
- FILEGROUP
- FILENAME

- ON PRIMARY;
- ON COLLATE;

```
GO
```

```
CREATE
```

A. Mastered

B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 76

- (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 81

- (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 85

- (Exam Topic 7)

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

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

You have Microsoft SQL Server on a Microsoft Azure virtual machine that has a database named DB1. You discover that DB1 experiences WRITE_LOG waits that are longer than 50 ms.

You need to reduce the WRITE_LOG wait time. Solution: Add additional data files to DB1. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

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

NEW QUESTION 86

- (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 89

- (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 94

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database that includes a table named Application.Events. Application.Events contains millions of records about user activity in an application.

Records in Application.Events that are more than 90 days old are purged nightly. When records are purged, table locks are causing contention with inserts.

You need to be able to modify Application.Events without requiring any changes to the applications that utilize Application.Events.

Which type of solution should you use?

- A. Partitioned tables
- B. Online index rebuild
- C. Change data capture
- D. Change tracking

Answer: A

Explanation:

Partitioning large tables or indexes can have manageability and performance benefits including:

You can perform maintenance operations on one or more partitions more quickly. The operations are more efficient because they target only these data subsets, instead of the whole table.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/partitions/partitioned-tables-and-indexes>

NEW QUESTION 95

- (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 98

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance that contains a financial database hosted on a storage area network (SAN).

The financial database has the following characteristics:

The database is continually modified by users during business hours from Monday through Friday between 09:00 hours and 17:00 hours. Five percent of the existing data is modified each day.

The Finance department loads large CSV files into a number of tables each business day at 11:15 hours and 15:15 hours by using the BCP or BULK INSERT commands. Each data load adds 3 GB of data to the database.

These data load operations must occur in the minimum amount of time.

A full database backup is performed every Sunday at 10:00 hours. Backup operations will be performed every two hours (11:00, 13:00, 15:00, and 17:00) during business hours.

You need to ensure that your backup will continue if any invalid checksum is encountered. Which backup option should you use?

- A. STANDBY
- B. Differential
- C. FULL
- D. CHECKSUM
- E. BULK_LOGGED
- F. CONTINUE_AFTER_ERROR
- G. SIMPLE
- H. DBO_ONLY
- I. COPY_ONLY
- J. SKIP
- K. RESTART
- L. Transaction log
- M. NO_CHECKSUM

N. NORECOVERY

Answer: F

Explanation:

The CONTINUE_AFTER_ERROR option, of the Transact-SQL BACKUP command, instructs BACKUP to continue despite encountering errors such as invalid checksums or torn pages.

References:

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

NEW QUESTION 99

- (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 103

- (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 106

- (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 107

- (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 109

- (Exam Topic 7)

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

Service tier:

	▼
Basic	
Standard	
Premium	
Premium RS	

Database implementation:

	▼
One individual Azure SQL database	
Thirty individual Azure SQL databases	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 114

- (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 115

- (Exam Topic 7)

Your database contains a table named Purchases. The table includes a DATETIME column named PurchaseTime that stores the date and time each purchase is made. There is a non- clustered index on the PurchaseTime column. The business team wants a report that displays the total number of purchases made on the current day. You need to write a query that will return the correct results in the most efficient manner.

Which Transact-SQL query should you use?

- A. `SELECT COUNT(*)FROM PurchasesWHERE PurchaseTime = CONVERT(DATE, GETDATE())`
- B. `SELECT COUNT(*)FROM PurchasesWHERE PurchaseTime = GETDATE()`
- C. `SELECT COUNT(*)FROM PurchasesWHERE CONVERT(VARCHAR, PurchaseTime, 112)=CONVERT(VARCHAR, GETDATE(), 112)`
- D. `SELECT COUNT(*)FROM PurchasesWHERE PurchaseTime >= CONVERT(DATE, GETDATE())AND PurchaseTime <DATEADD(DAY, 1, CONVERT(DATE, GETDATE()))`

Answer: D

Explanation:

To compare a time with date we must use >= and > operators, and not the = operator.

NEW QUESTION 118

- (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 119

- (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 Database1. 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 123

- (Exam Topic 7)

You have a SQL Server 2016 database named DB1.

You plan to import a large number of records from a SQL Azure database to DB1.

You need to recommend a solution to minimize the amount of space used in the transaction log during the import operation.

What should you include in the recommendation?

- A. The bulk-logged recovery model
- B. The full recovery model
- C. A new partitioned table
- D. A new log file
- E. A new file group

Answer: A

Explanation:

Compared to the full recovery model, which fully logs all transactions, the bulk-logged recovery model minimally logs bulk operations, although fully logging other transactions. The bulk-logged recovery model protects against media failure and, for bulk operations, provides the best performance and least log space usage.

Note: The bulk-logged recovery model is a special-purpose recovery model that should be used only intermittently to improve the performance of certain large-scale bulk operations, such as bulk imports of large amounts of data.

References: [https://technet.microsoft.com/en-us/library/ms190692\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms190692(v=sql.105).aspx)

NEW QUESTION 125

- (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 129

- (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 131

- (Exam Topic 7)

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

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

You need to recover DB1 and prevent any data loss.

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

Select and Place:

Select and Place:

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

Step 2: Restore a full backup.

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

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

Step 3: Restore the log backups.

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

Reference:

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

NEW QUESTION 135

- (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 140

- (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 143

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft Azure virtual machine that has a database named DB1. You discover that DB1 experiences WRITE_LOG waits that are longer than 50 ms.

You need to reduce the WRITE_LOG wait time. Solution: Add additional log files to DB1.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

References:

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

NEW QUESTION 144

- (Exam Topic 7)

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

Two SQL Server 2012 databases named SQL1 and SQL2

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

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

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

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

Answer: AE

Explanation:

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

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

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

NEW QUESTION 147

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft Azure virtual machine.

You suspect that the current SQL Server indexes cause queries to execute slowly.

You need to identify which indexes must be created to reduce the query execution time.

Which three dynamic management views should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. sys.dm_db_index_physical_stats
- B. sys.dm_db_missing_index_group_stats
- C. sys.indexes
- D. sys.dm_db_index_usage_stats
- E. sys.dm_db_missing_index_groups
- F. sys.dm_db_index_operational_stats
- G. sys.dm_db_missing_index_details
- H. sys.sysindexkeys

Answer: BEG

Explanation:

The missing indexes feature consists of the following components:

A set of dynamic management objects that can be queried to return information about missing indexes.

The Missing Indexes element in XML Showplans, which correlate indexes that the query optimizer considers missing with the queries for which they are missing. Dynamic Management Objects

After running a typical workload on SQL Server, you can retrieve information about missing indexes by querying the dynamic management objects listed in the following table. These dynamic management objects are stored in the master database.

sys.dm_db_missing_index_group_stats

Returns summary information about missing index groups, for example, the performance improvements that could be gained by implementing a specific group of missing indexes.

sys.dm_db_missing_index_groups

Returns information about a specific group of missing indexes, such as the group identifier and the identifiers of all missing indexes that are contained in that group.

sys.dm_db_missing_index_details

Returns detailed information about a missing index; for example, it returns the name and identifier of the table where the index is missing, and the columns and column types that should make up the missing index.

sys.dm_db_missing_index_columns

Returns information about the database table columns that are missing an index. References: [https://technet.microsoft.com/en-us/library/ms345524\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms345524(v=sql.105).aspx)

NEW QUESTION 152

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance that contains a financial database hosted on a storage area network (SAN).

The financial database has the following characteristics:

The database is continually modified by users during business hours from Monday through Friday between 09:00 hours and 17:00 hours. Five percent of the existing data is modified each day.

The Finance department loads large CSV files into a number of tables each business day at 11:15 hours and 15:15 hours by using the BCP or BULK INSERT commands. Each data load adds 3 GB of data to the database.

These data load operations must occur in the minimum amount of time.

A full database backup is performed every Sunday at 10:00 hours. Backup operations will be performed every two hours (11:00, 13:00, 15:00, and 17:00) during business hours.

You need to ensure that the minimum amount of data is lost. Which recovery model should the database use?

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

Answer: B

Explanation:

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

NEW QUESTION 154

- (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 159

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft Azure virtual machine.

You have two Windows accounts named serviceAccount1 and ServiceAccount2. The SQL Server Agent runs as ServiceAccount1.

You need to run SQL Server Agent job steps by using ServiceAccount2. Which cmdlet should you run first?

- A. Set-ADServiceAccount
- B. Set-SqlCredential
- C. New-ADServiceAccount
- D. New-SqlCredential

Answer: C

Explanation:

The New-ADServiceAccount command creates a new Active Directory managed service account or group managed service account object.

NEW QUESTION 164

- (Exam Topic 7)

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

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

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

Which feature should you use?

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

Answer: B

Explanation:

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

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

NEW QUESTION 166

- (Exam Topic 7)

You administer a Windows Azure SQL Database database named Orders. You need to create a copy of Orders named Orders_Reporting.

Which Transact-SQL command should you use?

- A. BACKUP DATABASE Orders TO DISK = 'D:\Orders.bak'RESTORE DATABASEOrders_ReportingFROM DISK = 'D:\Orders.bak
- B. BACKUP DATABASE Orders TO DISK = 'D:\Orders.bak'CREATE DATABASEOrders_ReportingFROM DISK = 'D:\Orders.bak
- C. CREATE DATABASE Orders_Reporting AS COPY OF Orders
- D. BACKUP DATABASE Orders TO DISK = 'D:\Orders.bak'MIRROR TO DISK = 'Orders_Reporting

Answer: C

Explanation:

BACKUP DATABASE ...AS COPY OF [source_server_name.]source_database_name Is used for copying a database to the same or a different SQL Database server.

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

NEW QUESTION 170

- (Exam Topic 7)

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

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

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

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

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

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

- A. Yes
- B. No

Answer: A

Explanation:

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

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

NEW QUESTION 173

- (Exam Topic 7)

Settings Value VM size D3

Storage Location Drive E Storage type Standard Tempdb location Drive C

The workload on this instance has of the tempdb load.

You need to maximize the performance of the tempdb database.

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

- A. Yes
- B. No

Answer: B

Explanation:

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

References:

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

NEW QUESTION 174

- (Exam Topic 7)

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

You need to collect data for a long period of time to troubleshoot wait statistics when querying Contoso. You also need to ensure minimum impact to the server. What should you create?

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

Answer: C

Explanation:

SQL Server Extended Events has a highly scalable and highly configurable architecture that allows users to collect as much or as little information as is necessary to troubleshoot or identify a performance problem.

Extended Events is a light weight performance monitoring system that uses very few performance resources. A SQL Server Extended Events session is created in the SQL Server process hosting the Extended Events engine.

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

NEW QUESTION 177

- (Exam Topic 7)

You have Microsoft SQL Server on a DS-series Microsoft Azure virtual machine. The virtual machine has 28 GB of memory.

You discover the following performance statistics on the server:

The average Page life expectancy is 30.

The server has excessive PAGELATCH_IO waits.

You need to decrease the PAGELATCH_IO waits. What should you do?

- A. Enable large-page support.
- B. Enable lock pages in memory.
- C. Configure buffer pool extensions.
- D. Add more tempdb files.

Answer: D
Explanation: References:

NEW QUESTION 178

- (Exam Topic 7)

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

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

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

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

Answer: H

Explanation:

Always On availability groups supports two availability modes— asynchronous-commit mode and synchronous-commit mode

Synchronous-commit mode emphasizes high availability over performance, at the cost of increased transaction latency.

References: <https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-mode>

NEW QUESTION 182

- (Exam Topic 7)

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

You have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in an Always On Availability Group.

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

You need to recommend a storage solution for the SQL Servers. The solution must minimize costs. Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 disk storage
- C. Premium P30 disk storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage

- F. Standard zone redundant blob storage
- G. Standard locally redundant blob storage
- H. Standard geo-redundant blob storage

Answer: A

Explanation:

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

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

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

NEW QUESTION 187

- (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 188

- (Exam Topic 7)

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 190

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database. 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 195

- (Exam Topic 7)

You plan to create an AlwaysOn availability group that will have two replicas in Microsoft Azure and two on premises replicas. You need to configure the network to support the availability group listener. Which cmdlet should you run first?

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 197

- (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 202

- (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 206

.....

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>