

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. 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 2

- (Topic 1)

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

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

You need to create the templates.

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 3

- (Topic 1)

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

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

You plan to migrate the databases from SRV2014 to SRV2016.

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

What should you do?

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

Answer: B

Explanation:

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

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


2. Configuration options: Data Files

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

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

* Autogrowth – is specified in MB (fixed growth is preferred as opposed to a non-linear percentage based growth) and applies to each file. The default value of 64MB was chosen to cover one PFS interval.

Figure:



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

NEW QUESTION 4

- (Topic 1)

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

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

Answer: D

Explanation:

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

Premium service tier

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

NEW QUESTION 5

DRAG DROP - (Topic 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 6

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 7

- (Topic 2)

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

You manage a Microsoft SQL Server environment with several databases.

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

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

- A. Yes
- B. No

Answer: A

Explanation:

PARAMETER_SNIFFING = { ON | OFF | PRIMARY} enables or disables parameter sniffing. This is equivalent to Trace Flag 4136. SQL server uses a process called parameter sniffing when executing queries or stored procedures that use parameters. During compilation, the value passed into the parameter is evaluated and used to create an execution plan. That value is also stored with the execution plan in the plan cache. Future executions of the plan will re-use the plan that was compiled with that reference value. References: <https://msdn.microsoft.com/en-us/library/mt629158.aspx>

NEW QUESTION 8

- (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 ODBC 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)

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

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

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

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

You need to secure the data. What should you implement?

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

Answer: C

Explanation:

Transparent data encryption (TDE) encrypts your databases, associated backups, and transaction log files at rest without requiring changes to your applications.

TDE encrypts the storage of an entire database by using a symmetric key called the database encryption key. In SQL Database the database encryption key is protected by a built-in server certificate. The built-in server certificate is unique for each SQL Database server.

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

NEW QUESTION 10

- (Topic 3)

A company has an on-premises Microsoft SQL Server 2014 environment. The company has a main office in Seattle, and remote offices in Amsterdam and Tokyo.

You plan to deploy a Microsoft Azure SQL Database instance to support a new application. You expect to have 100 users from each office.

In the past, users at remote sites reported issues when they used applications hosted at the Seattle office.

You need to optimize performance for users running reports while minimizing costs. What should you do?

- A. Implement an elastic pool.
- B. Implement a standard database with readable secondaries in Asia and Europe, and then migrate the application.
- C. Implement replication from an on-premises SQL Server database to the Azure SQL Database instance.
- D. Deploy a database from the Premium service tier.

Answer: B

Explanation:

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-geo-replication-transact-sql#add-secondary-database>

NEW QUESTION 10

HOTSPOT - (Topic 5)

You need to create the contosodb1 database.

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

Answer Area

▼
New-AzureSqlDatabase
New-AzureRmSqlDatabase
Set-AzureRmSqlDatabase

- ResourceGroupName “contosodbrg”

- ServerName “contososrv”

- DatabaseName “contosodbl”

- Edition

▼
Basic
Standard
Premium

- RequestedServiceObjectName S2

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: New-AzureRmSqlDatabase

New-AzureRmSqlDatabase creates a database or an elastic database.

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

NEW QUESTION 15

HOTSPOT - (Topic 6)

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

DB1: DELAYED_DURABILITY=FORCED

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

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

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

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

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

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

DB2: ALLOW_SNAPSHOT_ISOLATION ON and READ_COMMITTED_SNAPSHOT ON

Snapshot isolation enhances concurrency for OLTP applications.

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

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

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

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

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

References:

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

NEW QUESTION 16

- (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 18

- (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 23

- (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 27

- (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 28

- (Exam Topic 7)

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

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

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

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

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 32

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database.

The database contains a Product table created by using the following definition:

```
CREATE TABLE dbo.Product
(ProductID INT PRIMARY KEY,
Name VARCHAR(50) NOT NULL,
Color VARCHAR(15) NOT NULL,
Size VARCHAR(5) NOT NULL,
Style CHAR(2) NULL,
Weight DECIMAL(8,2) NULL);
```

You need to ensure that the minimum amount of disk space is used to store the data in the Product table. What should you do?

- A. Convert all indexes to Column Store indexes.
- B. Implement Unicode Compression.
- C. Implement row-level compression.
- D. Implement page-level compression.

Answer: D

NEW QUESTION 34

- (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 37

- (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 42

- (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 43

- (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 47

- (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 50

- (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 55

- (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_OrderDateON CustomerOrder DISABLE
- B. ALTER INDEX IX_OrderDateON CustomerOrder ENABLE
- C. ALTER INDEX IX_OrderDateON CustomerOrder REORGANIZE
- D. ALTER INDEX IX_OrderDateON 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 59

- (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 would 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 60

- (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 63

- (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 65

- (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 67

- (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 72

- (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 73

- (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 76

- (Exam Topic 7)

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

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

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

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

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

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

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

Does this meet the goal?

A. Yes

B. No

Answer: A

Explanation:

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

A cause can be poor design of your database - you may wish to split out data located on 'hot pages', which are accessed frequently and which you might identify as the causes of your latch contention. For example, if you have a currency table with a data page containing 100 rows, of which 1 is updated per transaction and

you have a transaction rate of 200/sec, you could see page latch queues of 100 or more. If each page latch wait costs just 5ms before clearing, this represents a full half-second delay for each update. In this case, splitting out the currency rows into different tables might prove more performant (if less normalized and logically structured).

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

NEW QUESTION 78

- (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 79

- (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 80

- (Exam Topic 7)

You deploy a new Microsoft Azure SQL database instance to support a variety of mobile application and public websites. You configure geo-replication with regions in Brazil and Japan.

You need to implement real-time encryption of the database and all backups.

Solution: You password protect all azure SQL backups and enable azure active directory authentication for all azure SQL server instances.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Password protection does not encrypt the data.

Transparent Data Encryption (TDE) would provide a solution. References:

<https://azure.microsoft.com/en-us/blog/how-to-configure-azure-sql-database-geo-dr-with-azure-key-vault/>

NEW QUESTION 85

- (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:	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="background-color: #e0e0e0; height: 20px;"> </td><td style="text-align: right;">▼</td></tr> <tr><td colspan="2">FAT32</td></tr> <tr><td colspan="2">FAT</td></tr> <tr><td colspan="2">NTFS</td></tr> </table>		▼	FAT32		FAT		NTFS	
	▼								
FAT32									
FAT									
NTFS									
8.3 filename support:	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="background-color: #e0e0e0; height: 20px;"> </td><td style="text-align: right;">▼</td></tr> <tr><td colspan="2">Enabled</td></tr> <tr><td colspan="2">Disabled</td></tr> </table>		▼	Enabled		Disabled			
	▼								
Enabled									
Disabled									
Indexing:	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="background-color: #e0e0e0; height: 20px;"> </td><td style="text-align: right;">▼</td></tr> <tr><td colspan="2">Enabled</td></tr> <tr><td colspan="2">Disabled</td></tr> </table>		▼	Enabled		Disabled			
	▼								
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 86

- (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 91

- (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 96

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database named ContosoDb. Tables are defined as shown in the exhibit. (Click the Exhibit button.)



You need to display rows from the Orders table for the Customers row having the CustomerId value set to 1 in the following XML format.

```
<Customers Name="Customer A" Country="Australia">
  <OrderId>1</OrderId>
  <OrderDate>2000-01-01T00:00:00</OrderDate>
  <Amount>3400.00</Amount>
</Customers>
<Customers Name="Customer A" Country="Australia">
  <OrderId>2</OrderId>
  <OrderDate>2001-01-01T00:00:00</OrderDate>
  <Amount>4300.00</Amount>
</Customers>
```

Which Transact-SQL query should you use?

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

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

Answer: E

NEW QUESTION 98

- (Exam Topic 7)

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

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

Evenly distribute read and write operations across the VHDs.

Minimize the read and write time.

Which storage configuration should you use?

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

Answer: D

Explanation:

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

Reference:

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

NEW QUESTION 101

- (Exam Topic 7)

You manage an on-premises Microsoft SQL server that has a database named DB1. An application named App1 retrieves customer information for DB1.

Users report that App1 takes an unacceptably long time to retrieve customer records. You need to find queries that take longer than 400 ms to run.

Which statement should you execute?

A)

```
SELECT      qp.query_plan,
            qs.*
FROM        (
            SELECT TOP 50 *
            FROM sys.dm_exec_query_stats
            ORDER BY total_worker_time DESC
            ) AS qs
CROSS APPLY sys.dm_exec_query_plan(qs.plan_handle) AS qp
WHERE (qs.max_worker_time > 400
       OR qs.max_elapsed_time > 400)
```

B)

```
SELECT pa.DatabaseID, SUM(qs.total_worker_time/100) AS [CPU_Time_Ms]
FROM sys.dm_exec_query_stats AS qs
CROSS APPLY (SELECT CONVERT(int, value) AS [DatabaseID]
            FROM sys.dm_exec_plan_attributes(qs.plan_handle)
            WHERE attribute = N'dbid') AS pa
GROUP BY pa.DatabaseID
HAVING SUM(qs.total_worker_time/1000) > 400
ORDER BY 2 DESC
```

C)

```
SELECT      qp.query_plan,
            qs.*
FROM        (
            SELECT TOP 50 *
            FROM sys.dm_exec_query_stats
            ORDER BY total_worker_time DESC
            ) AS qs
CROSS APPLY sys.dm_exec_query_plan(qs.plan_handle) AS qp
WHERE      (qs.max_logical_reads > 400
            OR qs.max_logical_reads > 400)
```

D)

```
SELECT TOP 50 *
FROM sys.dm_exec_query_stats as qs
WHERE (qs.max_physical)_reads > 400
      OR qs.max_physical_reads > 400)
ORDER BY total_worker_time DESC
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

Total_worker_time: Total amount of CPU time, reported in microseconds (but only accurate to milliseconds), that was consumed by executions of this plan since it was compiled.

NEW QUESTION 104

- (Exam Topic 7)

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

The financial database has the following characteristics:

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

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

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

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

You need to ensure that the backup size is as small as possible. Which backup should you perform every two hours?

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

Answer: J

Explanation:

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

References:

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

NEW QUESTION 107

- (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)

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 115

- (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 117

- (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 121

- (Exam Topic 7)

You administer a Windows Azure SQL Database database named Human_Resources. The database contains 2 tables named Employees and SalaryDetails. You add two Windows groups as logins for the server:

You need to grant users access according to the following requirements: What should you do?

- A. Create a database role called Employees.Add CORP\Employees to the db_datareader rol
- B. Add all company employees except HR administrators to the Employees rol
- C. Deny SELECT access to the SalaryDetails table to the Employees role.
- D. Create a database role called HRAdmins.Add all company employees except HR administrators to the db_datareader rol

- E. Add all HR administrators to the HRAdmins rol
- F. Grant SELECT access to the SalaryDetails table to the HRAdmins role.Deny SELECT access to the SalaryDetails table to the db_datareader role.
- G. Create two database roles: Employees and HRAdmin
- H. Add all company employees to the Employees role.Add HR administrators to the HRAdmins rol
- I. Grant SELECT access to all tables except SalaryDetails to the Employees rol
- J. Grant SELECT access to the SalaryDetails table to the HRAdmins rol
- K. Deny SELECT access to the SalaryDetails table to the Employees role.
- L. Create a database role called Employees.Add all HR administrators to the db_datareader rol
- M. Add all company employees to the Employees rol
- N. Grant SELECT access to all tables except the SalaryDetails table to the Employees rol
- O. Deny SELECT access to the SalaryDetails table to the Employees role.

Answer: D

Explanation:

Members of the db_datareader fixed database role can run a SELECT statement against any table or view in the database.

References: [https://technet.microsoft.com/en-us/library/ms188629\(v=sql.90\).aspx](https://technet.microsoft.com/en-us/library/ms188629(v=sql.90).aspx)

NEW QUESTION 122

- (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 126

- (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 128

- (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 133

- (Exam Topic 7)

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

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

Solution: You move the database to a Standard database pool that has 800 eDTUs. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

NEW QUESTION 136

- (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 137

- (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 142

- (Exam Topic 7)

A company has an on-premises Microsoft SQL Server 2017 infrastructure. The storage area network (SAN) that supports the SQL infrastructure has reached maximum capacity.

You need to recommend a solution to reduce on-premises storage use without changing the application. What should you do?

- A. Configure an Express Route connection to Microsoft Azure.
- B. Configure a Microsoft Azure Key Vault.
- C. Configure geo-replication on the SAN.
- D. Configure SQL Server Stretch Database in Microsoft Azure.

Answer: D

Explanation:

Stretch warm and cold transactional data dynamically from SQL Server to Microsoft Azure with SQL Server Stretch Database. Unlike typical cold data storage, your data is always online and available to query. Benefit from the low cost of Azure rather than scaling expensive, on-premises storage.

References:

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

NEW QUESTION 146

- (Exam Topic 7)

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

What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 148

- (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 151

- (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 154

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance.

The instance contains a database that supports a retail sales application. The application generates hundreds of transactions per second and is online 24 hours per day and 7 days per week.

You plan to define a backup strategy for the database. You need to ensure that the following requirements are met:

No more than 5 minutes worth of transactions are lost. Data can be recovered by using the minimum amount of administrative effort.

What should you do? Choose all that apply.

- A. Configure the database to use the SIMPLE recovery model.
- B. Create a DIFFERENTIAL database backup every 4 hours.
- C. Create a LOG backup every 5 minutes.
- D. Configure the database to use the FULL recovery model.
- E. Create a FULL database backup every 24 hours.
- F. Create a DIFFERENTIAL database backup every 24 hours.

Answer: BCDE

Explanation:

The full recovery model uses log backups to prevent data loss in the broadest range of failure scenarios, and backing and restoring the transaction log (log backups) is required. The advantage of using log backups is that they let you restore a database to any point of time that is contained within a log backup (point-in-time

recovery). You can use a series of log backups to roll a database forward to any point in time that is contained in one of the log backups. Be aware that to minimize your restore time, you can supplement each full backup with a series of differential backups of the same data.

References: [https://technet.microsoft.com/en-us/library/ms190217\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms190217(v=sql.105).aspx)

NEW QUESTION 155

- (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 160

- (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 163

- (Exam Topic 7)

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

You need to write messages to the Application Log when users are added to or removed from a fixed server role 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: G

Explanation:

The SQL Server Audit feature enables you to audit server-level and database-level groups of events and individual events.

Audits can have the following categories of actions:

Server-level. These actions include server operations, such as management changes, such as in this question, and logon and logoff operations.

Database-level. These actions encompass data manipulation languages (DML) and data definition language (DDL) operations.

Audit-level. These actions include actions in the auditing process.

References:

[http://technet.microsoft.com/en-us/library/cc280663\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/cc280663(v=sql.105).aspx)

NEW QUESTION 164

- (Exam Topic 7)

You create a new Microsoft Azure subscription.

You need to create a group of Azure SQL databases that share resources. Which cmdlet should you run first?

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

Answer: D

Explanation:

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

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

NEW QUESTION 166

- (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 168

- (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 169

- (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 171

- (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 174

- (Exam Topic 7)

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

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

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

You need to reduce the WRITE_LOG wait time. Solution: Move the transaction logs to a faster disk. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Section: Deploy and migrate applications

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

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

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

NEW QUESTION 177

- (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 181

.....

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>