

Exam Questions 70-765

Provisioning SQL Databases (beta)

<https://www.2passeasy.com/dumps/70-765/>



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

HOTSPOT - (Topic 1)

You plan to migrate a Microsoft SQL Server workload from an on-premises server to a Microsoft Azure virtual machine (VM). The current server contains 4 cores with an average CPU workload of 6 percent and a peak workload of 10 percent when using 2.4Ghz processors.

You gather the following metrics:

	Minimum IOPS	Average IOPS	Maximum IOPS
Data Drive	100	938	7253
Transaction Log Drive	12	145	350
TempDB Drive	300	900	1900

You need to design a SQL Server VM to support the migration while minimizing costs.

For each setting, which value should you use? To answer, select the appropriate storage option from each list in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

VM setting	Value ▼
Data drive	Local storage Premium storage Standard storage
Transaction log drive	Local storage Premium storage Standard storage
TempDB drive	Local storage Premium storage Standard storage
VM size	A3 D3 DS3

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Data drive: Premium Storage Transaction log drive: Standard Storage TempDB drive: Premium Storage
Note: A standard disk is expected to handle 500 IOPS or 60MB/s. A P10 Premium disk is expected to handle 500 IOPS.
A P20 Premium disk is expected to handle 2300 IOPS. A P30 Premium disk is expected to handle 5000 IOPS.
VM size: A3
Max data disk throughput is 8x500 IOPS
References:<https://docs.microsoft.com/en-us/azure/virtual-machines/virtual-machines-windows-sizes>

NEW QUESTION 6

DRAG DROP - (Topic 1)

You are building a new Always On Availability Group in Microsoft Azure. The corporate domain controllers (DCs) are attached to a virtual network named ProductionNetwork. The DCs are part of an availability set named ProductionServers1.
You create the first node of the availability group and add it to an availability set named ProductionServers2. The availability group node is a virtual machine (VM) that runs Microsoft SQL Server. You attach the node to ProductionNetwork.
The servers in the availability group must be directly accessible only by other company VMs in Azure.
You need to configure the second SQL Server VM for the availability group.
How should you configure the VM? To answer, drag the appropriate configuration settings to the correct target locations. Each configuration setting may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Configuration settings

None/Not Assigned

ProductionServers1

ProductionNetwork

ProductionServers2

Create a new Object

VM settings page

Settings
— □ X

Storage

Disk type

Standard Premium (SSD)

* Storage account >

(new) sqlstorage3

Network

* Virtual network >

setting

* Subnet >

ProductionServers (10.1.0.0/24)

* Public IP address >

setting

* Network security group >

(new) SQLServers

Extensions

Extensions >

No extensions

Monitoring

Diagnostics

Disabled Enabled

Availability

* Availability set >

setting

OK

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Box 2: None /Not Assigned

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

Box 3: ProductionServer2

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

NEW QUESTION 7

- (Topic 2)

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

You manage a Microsoft SQL Server environment with several databases.

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

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

A. Yes

B. No

Answer: A

Explanation:

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

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

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

NEW QUESTION 8

- (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

HOTSPOT - (Topic 2)

You need to ensure that a user named Admin2 can manage logins.

How should you complete the Transact-SQL statements? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

▼	Admin2 WITH password = 'Pa\$\$wOrd';	
CREATE USER		
ALTER SERVER ROLE		
CREATE LOGIN		

▼	Admin2User FROM	▼	Admin2
CREATE USER		WINDOWS	
ALTER SERVER ROLE		EXTERNAL PROVIDER	
CREATE LOGIN		LOGIN	

ALTER ROLE '	▼
	loginmanager
	dbmanager
	bd_ddladmin

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: CREATE LOGIN

First you need to create a login for SQL Azure, it's syntax is as follows: CREATE LOGIN username WITH password='password';

Step 2, CREATE USER Step 3: LOGIN

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

CREATE USER readonlyuser FROM LOGIN readonlylogin; Step 4: loginmanager

Members of the loginmanager role can create new logins in the master database.

References:

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

NEW QUESTION 10

DRAG DROP - (Topic 2)

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

You need to create the new security principal.

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

Actions

Answer Area

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

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

Run the following Transact-SQL statement:

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

Run the following Transact-SQL statement:

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

Run the following Transact-SQL statements:

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1:

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

Step 2: CREATE USER ... FROM EXTERNAL PROVIDER

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

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

Step 3:

Grant the proper reading permissions.

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

NEW QUESTION 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 13

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 14

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 19

- (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 24

- (Exam Topic 7)


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

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

You review the execution plan for the query and discover that the expected number of returned rows is one, while the actual number of returned rows is 1 million.

You need to reduce the amount of time it takes for the query to execute. The solution must prevent additional performance issues from being introduced.

Hot Area:



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

References:

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

NEW QUESTION 29

- (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 33

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database that contains a table named AccountTransaction. You discover that query performance on the table is poor due to fragmentation on the

IDX_AccountTransaction_AccountCode non-clustered index. You need to defragment the index. You also need to ensure that user queries are able to use the index during the defragmenting process.

Which Transact-SQL batch should you use?

- A. ALTER INDEX IDX_AccountTransaction_AccountCode ON AccountTransaction.AccountCode REORGANIZE
- B. ALTER INDEX ALL ON AccountTransaction REBUILD
- C. ALTER INDEX IDX_AccountTransaction_AccountCode ON AccountTransaction.AccountCode REBUILD
- D. CREATE INDEX IDXAccountTransactionAccountCode ON AccountTransaction.AccountCode WITH DROP EXISTING

Answer: A

Explanation:

Reorganize: This option is more lightweight compared to rebuild. It runs through the leaf level of the index, and as it goes it fixes physical ordering of pages and also compacts pages to apply any previously set fillfactor settings. This operation is always online, and if you cancel it then it's able to just stop where it is (it doesn't have a giant operation to rollback).

References: <https://www.brentozar.com/archive/2013/09/index-maintenance-sql-server-rebuild-reorganize/>

NEW QUESTION 38

- (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 43

- (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 44

- (Exam Topic 7)

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

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

Solution: You use the Always Encrypted wizard to encrypt all possible for the tables in the primary instance. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Always Encrypted does not support geo replication. Transparent Data Encryption (TDE) would provide a solution.

Note: Use the Always Encrypted Wizard to help protect sensitive data stored in a SQL Server database. Always Encrypted allows clients to encrypt sensitive data inside client applications and never reveal the encryption keys to SQL Server.

References:

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

<http://blog.pragmaticworks.com/sql-server-2016-data-masking-and-always-encrypted>

NEW QUESTION 45

- (Exam Topic 7)

You are the administrator for a SQL Server 2016 instance that stores the data for an online transaction processing sales system. The company takes full backups every week; differential backups on the days with no full backups; and hourly transaction backups.

These backups are stored on a backup server in the company's data center.

Every week, the company places the full backup on a tape and sends it to a third-party backup storage system. The company is worried that a disaster might occur that could destroy their computer center and cause them to lose orders.

You need to determine the best method for providing the smallest amount of data loss and downtime without leasing or purchasing additional physical locations.

What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Set up SQL Server Always On with a SQL Azure database as a replica.
- B. Set up SQL Server Always On by using a SQL Server on a Windows Azure Virtual Machine.
- C. Put the differential backup on tape and send it to the third-party backup storage system.
- D. Use the Microsoft SQL Server Backup to Microsoft Windows Azure Tool to direct all backups to a different geographical location.

Answer: D

Explanation:

Microsoft SQL Server Backup to Microsoft Azure Tool enables backup to Azure Blob Storage and encrypts and compresses SQL Server backups stored locally or in the cloud.

References: <https://www.microsoft.com/en-us/download/details.aspx?id=40740>

NEW QUESTION 50

- (Exam Topic 7)

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

Drive L fails and is replaced.

You need to recover DB1. The solution must minimize data loss.

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

Statements, Select from these	Statements, place here
ALTER DATABASE DB1 SET EMERGENCY, SINGLE_USER	
ALTER DATABASE DB1 SET ONLINE, ROLLBACK IMMEDIATE	
DBCC CHECKED('DB1', REPAIR_REBUILD)	
ALTER DATABASE DB1 SET ONLINE, MULTI_USER	
ALTER DATABASE db1 SET EMERGENCY, ROLLBACK IMMEDIATE	
ALTER DATABASE db1 SET SINGLE_USER WITH ROLLBACK IMMEDIATE	
DBCC CHECKDB('DB1', REPAIR_ALLOW_DATA_LOSS)	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

ALTER DATABASE '<your_database>' SET SINGLE_USER WITH ROLLBACK IMMEDIATE GO

DBCC CHECKDB ('<your_database>', REPAIR_REBUILD) GO

ALTER DATABASE '<your_database>' SET MULTI_USER GO

NEW QUESTION 55

- (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 60

- (Exam Topic 7)

A company has an on-premises Microsoft SQL Server environment with a SQL-Server named SQL01. You need to create a local sysadmin account on SQL01 named Admin1.

How should you complete the Transact-SQL statements? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer area

	[Admin] WITH PASSWORD=N'Pa\$\$w0rd'
CREATE USER	
CREATE LOGIN	
	[sysadmin] ADD MEMBER [Admin1]
ALTER DATABASE	
ALTER ROLE	
ALTER SERVER ROLE	
	[Admin1] FOR LOGIN [Admin1]
CREATE LOGIN	
GRANT LOGIN	
CREATE USER	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

B: First we create a login with the CREATE LOGIN command. E: Then we add it to the sysadmin role.

1. To add a member to a fixed server role
2. In Object Explorer, connect to an instance of Database Engine.
3. On the Standard bar, click New Query.

Copy and paste the following example into the query window and click Execute. ALTER SERVER ROLE diskadmin ADD [Domain\Juan] ; GO

G: Finally we add a database user for the login we created.

References: [https://technet.microsoft.com/en-us/library/aa337562\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/aa337562(v=sql.110).aspx)

NEW QUESTION 65

- (Exam Topic 7)

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

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

Does this meet the goal?

- A. Yes

B. No

Answer: A

NEW QUESTION 66

- (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 71

- (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 72

- (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 77

- (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 82

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft azure virtual machine that has 12 databases. All database files are in the same Azure Blob storage account.

You need to receive an email notification if I/O operations to the database files exceed 800 MB/s for more than five minutes.

Solution: You run the Get-Counter cmdlet and specify the -counter '\physicaldisk:disk write/sec' parameter. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 85

- (Exam Topic 7) You have a database named DB1. You discover that DB1 is corrupt.

You run DBCC CHECKDB and receive an error message within a few seconds. No pages are listed in the error message. You need to repair the database corruption as quickly as possible. The solution must minimize data loss. What should you do?

- A. Run DBCC CHECKDB ('db1', REPAIR_ALLOW_DATA_LOSS).
- B. Run DBCC CHECKDB ('db1', REPAIR_FAST).
- C. Delete the transaction logs and restart the Microsoft SQL Server instance.
- D. Run DBCC CHECKDB ('db1', REPAIR_REBUILD).
- E. Restore the database from a backup.

Answer: C

Explanation:

REPAIR_REBUILD

Performs repairs that have no possibility of data loss. This can include quick repairs, such as repairing missing rows in non-clustered indexes, and more time-consuming repairs, such as rebuilding an index.

NEW QUESTION 90

- (Exam Topic 7)

You manage a Microsoft SQL Server environment with several databases.

You need to ensure that queries use statistical data and do not initialize values for local variables. Solution: you set the value of the MAXDOP parameter to 2. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

References:

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

NEW QUESTION 93

- (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 96

- (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 97

- (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 101

- (Exam Topic 7)

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

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

Solution: You move the database to a Basic database pool that has 1,600 eDTUs. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

We need the use of a Standard database pool.

NEW QUESTION 106

- (Exam Topic 7)

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

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

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

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

Answer: C

Explanation:

sys.dm_db_index_operational_stats dynamic management function provides us the current low-level I/O, locking, latching, and access method for each partition of

the table. This information is really useful to troubleshoot SQL Server performance issues.

Reference:

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

NEW QUESTION 111

- (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 112

- (Exam Topic 7)

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

Three user databases must be preinstalled.

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

Trace flag 2371 must be turned at the instance level.

The solution must meet the following requirements:

The instance must be preconfigured.

No other database features are required in the future.

The solution must use the minimum administrative effort.

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

- A. Provision 10 Azure virtual machines that each contain SQL Server 2016, installed by using the default settings.
- B. Create an installation configuration file and perform unattended installations of SQL Server 2016.
- C. Create a virtual machine template by using a prepared instance of SQL Server 2016.
- D. Create a virtual machine template by using a complete instance of SQL Server 2016.

Answer: D

Explanation:

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

NEW QUESTION 116

- (Exam Topic 7)

You administer a SQL Server 2014 server that contains a database named SalesDB. SalesDB contains a schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales.

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

You need to ensure that UserA is disallowed to select from any of the tables in the Customers schema. Which Transact-SQL statement should you use?

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

Answer: G

Explanation:

Use SQL Data Warehouse or Parallel Data Warehouse GRANT and DENY statements to grant or deny a permission (such as UPDATE) on a securable (such as a database, table, view, etc.) to a security principal (a login, a database user, or a database role).

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

NEW QUESTION 118

- (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 119

- (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; text-align: right;">▼</td></tr> <tr><td>FAT32</td></tr> <tr><td>FAT</td></tr> <tr><td>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; text-align: right;">▼</td></tr> <tr><td>Enabled</td></tr> <tr><td>Disabled</td></tr> </table>	▼	Enabled	Disabled	
▼					
Enabled					
Disabled					
Indexing:	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="background-color: #e0e0e0; text-align: right;">▼</td></tr> <tr><td>Enabled</td></tr> <tr><td>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 122

- (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 Transparent Data Encryption (TDE) on the primary instance. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Azure SQL Database and Data Warehouse offer encryption-at-rest by providing Transparent Data Encryption (TDE) for all data written to disk, including

databases, log files and backups. This protects data in case of unauthorized access to hardware. TDE provides a TDE Protector that is used to encrypt the Database Encryption Key (DEK), which in turn is used to encrypt the data. With the TDE and Bring Your Own Key (BYOK) offering currently in preview, customers can take control of the TDE Protector in Azure Key Vault.

Taking advantage of TDE with BYOK for databases that are geo-replicated to maintain high availability requires to configure and test the scenario carefully.


References:

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

NEW QUESTION 124

- (Exam Topic 7)

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



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

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

Which Transact-SQL query should you use?

- A. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML RAW
- B. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML RAW, ELEMENTS
- C. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML AUTO
- D. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId * Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML AUTO, ELEMENTS
- E. SELECT Name, Country, OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML AUTO, ELEMENTS
- F. SELECT Name, Country, OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = FOR XML AUTO, ELEMENTS
- G. SELECT Name AS '@Name', Country AS '@Country', OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML PATH ('Customers')
- H. SELECT Name AS 'Customers/Name', Country AS 'Customers/Country', OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML PATH ('Customers')

Answer: E

NEW QUESTION 127

- (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 131

- (Exam Topic 7)

You use a contained database named ContosoDb within a domain.

You need to create a user who can log on to the ContosoDb database. You also need to ensure that you can port the database to different database servers within the domain without additional user account configurations.

Which type of user should you create?

- A. SQL user without login
- B. User mapped to an asymmetric key

- C. Domain user
- D. login mapped to a virtual account

Answer: C

Explanation:

If the service must interact with network services, access domain resources like file shares or if it uses linked server connections to other computers running SQL Server, you might use a minimally-privileged domain account. Many server-to-server activities can be performed only with a domain user account.
References: <https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-windows-servic>

NEW QUESTION 132

- (Exam Topic 7)

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

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

- Evenly distribute read and write operations across the VHDs.
- Minimize the read and write time.

Which storage configuration should you use?

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

Answer: D

Explanation:

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

Reference:

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

NEW QUESTION 135

- (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 140

- (Exam Topic 7)

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

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

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

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

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

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

- A. Yes
- B. No

Answer: B**Explanation:**

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

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

NEW QUESTION 145

- (Exam Topic 7)

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

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

What should you do?

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

Answer: D**Explanation:**

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

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

Local Key Stores References:

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

NEW QUESTION 147

- (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 150

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

NEW QUESTION 154

- (Exam Topic 7)

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

You need to diagnose deadlocks that happen when executing a specific set of stored procedures by recording events and playing them back on a different test server.

What should you create?

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

Answer: D

Explanation:

Use SQL Server Profiler to identify the cause of a deadlock. A deadlock occurs when there is a cyclic dependency between two or more threads, or processes, for some set of resources within SQL Server. Using SQL Server Profiler, you can create a trace that records, replays, and displays deadlock events for analysis.

References:

<http://msdn.microsoft.com/en-us/library/ms188246.aspx>

NEW QUESTION 158

- (Exam Topic 7)

You administer a Windows Azure SQL Database database named Inventory that contains a stored procedure named p_AddInventory.

Users need to be able to SELECT from all tables in the database and execute the stored procedure. You need to grant only the necessary permissions.

What should you do?

- A. Grant EXECUTE permission on p_AddInventory to all user
- B. Grant VIEW DEFINITION to all users.
- C. Grant EXECUTE permission on p_AddInventory to all user
- D. Add all users to the db_datawriter role.
- E. Add all users to the db_owner role.
- F. Grant EXECUTE permission on p_AddInventory to all user
- G. Add all users to the db_datareader 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 160

- (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 162

- (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 165

- (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 170

- (Exam Topic 7)

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

The following error message is an example error message:

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

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

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

Answer: B

Explanation:

References:

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

NEW QUESTION 175

- (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 179

- (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 180

- (Exam Topic 7)

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



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

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

Which Transact-SQL query should you use?

- A. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML RAW
- B. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML RAW, ELEMENTS
- C. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML AUTO
- D. SELECT OrderId, OrderDate, Amount, Name, Country FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML AUTO, ELEMENTS
- E. SELECT Name, Country, OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML AUTO
- F. SELECT Name, Country, OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML AUTO, ELEMENTS
- G. SELECT Name AS '@Name', Country AS '@Country', OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML PATH ('Customers')
- H. SELECT Name AS 'Customers/Name', Country AS 'Customers/Country', OrderId, OrderDate, Amount FROM Orders INNER JOIN Customers ON Orders.CustomerId = Customers.CustomerId WHERE Customers.CustomerId = 1 FOR XML PATH ('Customers')

Answer: G

NEW QUESTION 184

- (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 186

- (Exam Topic 7)

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

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

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

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

Answer: D

NEW QUESTION 191

- (Exam Topic 7)

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

The various databases on the instances have the following characteristics:

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

You run the following Transact-SQL statements:

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

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

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

Answer: D

NEW QUESTION 195

- (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 199

- (Exam Topic 7)

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

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 204

- (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 205

- (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 209

- (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 213

- (Exam Topic 7)

You administer a Microsoft SQL Server 2016 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:

References:

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

NEW QUESTION 214

- (Exam Topic 7)

You have Microsoft SQL server on a Microsoft Azure virtual machine. The virtual machine has 200 GB of data.

User report a slow response time when querying the database.

You need to identify whether the storage subsystem causes the performance issue. Which performance monitor counter should you view?

- A. Data sec/Write
- B. Avg.disk Read Queue Length
- C. % Disk Read Time
- D. Disk sec/Read

Answer: B

NEW QUESTION 218

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database instance.

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

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

What should you do?

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

Answer: C

Explanation:

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

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

NEW QUESTION 222

- (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 227

- (Exam Topic 7)

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

You need to track all SELECT statements issued in the Contoso database only by users in a role named Sales. What should you create?

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

Answer: F

Explanation:

To audit users in a role use a Database Audit Specification.

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

NEW QUESTION 229

- (Exam Topic 7)

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

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 233

- (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 237

- (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 240

- (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 241

- (Exam Topic 7)

You are migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 databases that consume a total of 2 TB of disk space. The instance sustains more than 30,000 transactions per second.

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

Solution: You use drive D on the virtual machine to store the database files. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

The D drive should only be used for temporary data.

NEW QUESTION 243

- (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 245

- (Exam Topic 7)

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

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

You discover that all files except Sales_2.ndf are corrupt.
 You need to recover the corrupted data in the minimum amount of time. What should you do?

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

Answer: A

Explanation:

In a file restore, the goal is to restore one or more damaged files without restoring the whole database. References:
<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/file-restores-simple-recovery-model>

NEW QUESTION 248

- (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 253

- (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 258

- (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 262

- (Exam Topic 7)

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

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

Answer: D

Explanation:

Stored procedures accept input parameters and return multiple values in the form of output parameters to the calling program. They cannot be used in views.
References:<https://docs.microsoft.com/en-us/sql/relational-databases/stored-procedures/stored-procedures-datab>

NEW QUESTION 264

- (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 268

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-765 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-765 Product From:

<https://www.2passeasy.com/dumps/70-765/>

Money Back Guarantee

70-765 Practice Exam Features:

- * 70-765 Questions and Answers Updated Frequently
- * 70-765 Practice Questions Verified by Expert Senior Certified Staff
- * 70-765 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * 70-765 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year