

Microsoft

Exam Questions 70-765

Provisioning SQL Databases (beta)



NEW QUESTION 1

- (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 2

- (Topic 1)

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

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

You need to create the templates.

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 3

- (Topic 1)

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

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

You need to create the templates.

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

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

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

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

```
}
```

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

NEW QUESTION 4

- (Topic 2)

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

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

Answer: D

Explanation:

In SQL Database, the relative measure of a database's ability to handle resource demands is expressed in Database Transaction Units (DTUs) for single databases and elastic DTUs (eDTUs) for databases in an elastic pool. A pool is given a set number of eDTUs, for a set price. Within the pool, individual databases are given the flexibility to auto-scale within set parameters. Under heavy load, a database can consume more eDTUs to meet demand. Additional eDTUs can be added to an existing pool with no database downtime. References:<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool>

NEW QUESTION 5

- (Topic 2)

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.

Your company has several Microsoft Azure SQL Database instances.

Data encryption should be allowed to be implemented by the client applications that access the data. Encryption keys should not be made available to the database engine.

You need to configure the database. What should you implement?

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

Answer: A

Explanation:

Using encryption during transit with Azure File Shares
Azure File Storage supports HTTPS when using the REST API, but is more commonly used as an SMB file share attached to a VM. HTTPS is a transport-level security protocol.

NEW QUESTION 6

- (Topic 3)

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

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 7

HOTSPOT - (Topic 5)

You need to create the contosodb1 database.

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

Answer Area

▼
New-AzureSqlDatabase
New-AzureRmSqlDatabase
Set-AzureRmSqlDatabase

- ResourceGroupName “contosodbrg”

- ServerName “contososrv”

- DatabaseName “contosodbl”

- Edition

▼
Basic
Standard
Premium

- RequestedServiceObjectName S2

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: New-AzureRmSqlDatabase

New-AzureRmSqlDatabase creates a database or an elastic database.

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

NEW QUESTION 8

- (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 tempdb databases.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

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 9

- (Exam Topic 7)

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

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

Answer: A

Explanation:

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

References:

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

NEW QUESTION 10

- (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 10

- (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 13

- (Exam Topic 7)

You plan to deploy an AlwaysOn failover cluster in Microsoft Azure. The cluster has a Service Level Agreement (SLA) that requires an uptime of at least 99.95 percent.

You need to ensure that the cluster meets the SLA.

Which cmdlet should you run before you deploy the virtual machine?

- 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:

On Azure virtual machines, a SQL Server Availability Group requires a load balancer. The load balancer holds the IP address for the Availability Group listener. The New-AzureRmLoadBalancer cmdlet creates an Azure load balancer.

References:

NEW QUESTION 18

- (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 21

- (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 24

- (Exam Topic 7)

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

The recovery model must be set to simple.

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

Your need to configure the SQL server 2016 environment.

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Model: Change from full to simple Recovery Model

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

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

Model: Set compatibility level to 120

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

References:

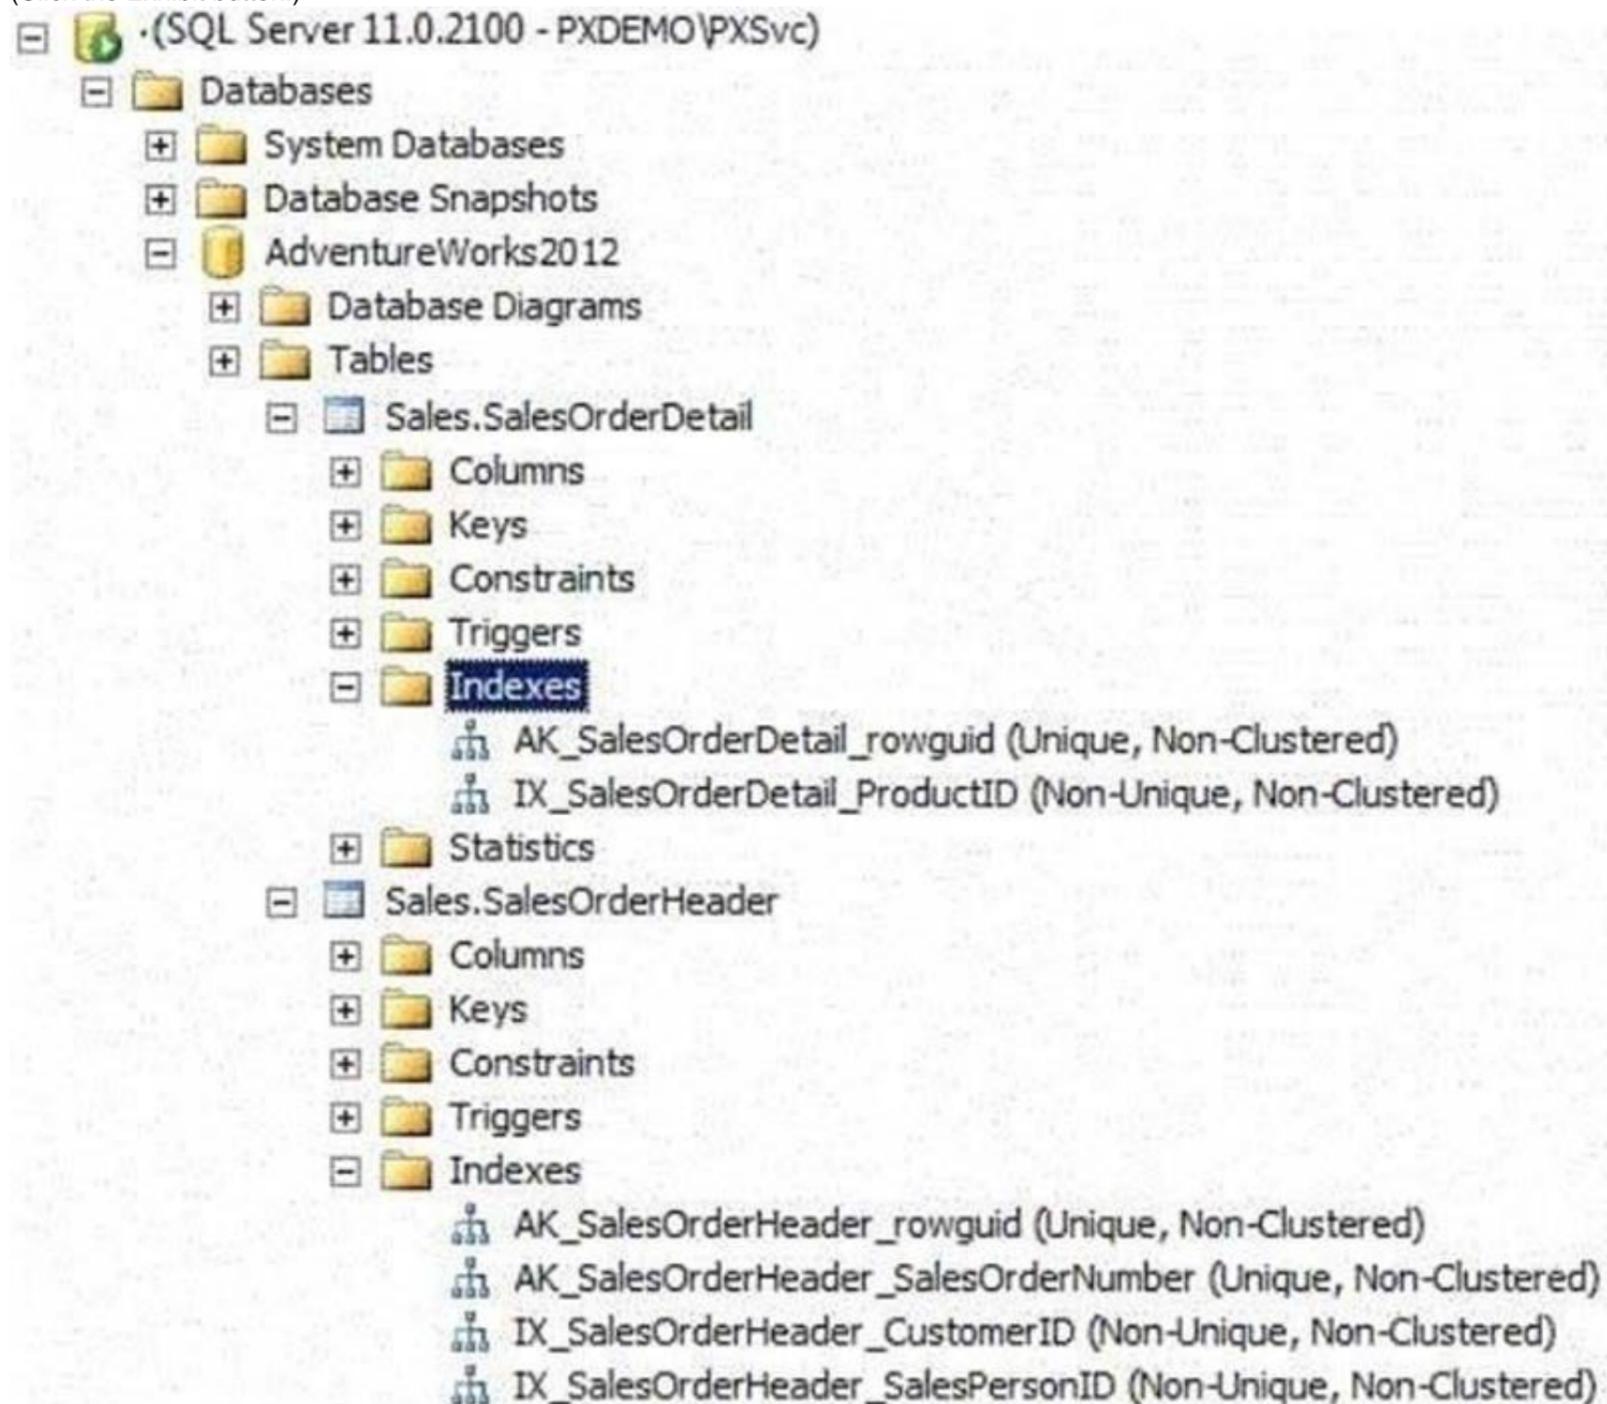
<https://docs.microsoft.com/en-us/sql/relational-databases/databases/model-database?view=sql-server-2017> <https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-transact-sql-compatibility-level?view=sql-se>

NEW QUESTION 27

- (Exam Topic 7)

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

(Click the Exhibit button.)



You write the following Transact-SQL query:

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

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

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

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

Answer: D

Explanation:

New statistics would be useful.

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

to update statistics more frequently than the default updates.

References:

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

NEW QUESTION 29

- (Exam Topic 7)

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

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

Several reports use the PIVOT statement and access more than 100 million rows in table1. You discover that when the reports run, there are PAGELATCH_IO waits on PFS pages 2:1:1, 2:2:1, 2:3:1, and 2:4:1 within the tempdb database.

You need to prevent the PAGELATCH_IO waits from occurring.

Solution: You rewrite the queries to use aggregates instead of PIVOT statements. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead you can add more files to the database.

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

NEW QUESTION 31

- (Exam Topic 7)

You plan to deploy an on-premises SQL Server 2014 database to Azure SQL Database. You have the following requirements:

Maximum database size of 500 GB

A point-in-time-restore of 35 days

Maximum database transaction units (DTUs) of 500

You need to choose the correct service tier and performance level. Which service tier should you choose?

- A. Standard S3
- B. Premium P4
- C. Standard SO
- D. Basic

Answer: B

Explanation:

You should choose Premium P4. The Premium tier is the highest Azure SQL Database tier offered. This tier is used for databases and application that require the highest level of performance and recovery. The P4 level supports a maximum of 500 DTUs, a maximum database size of 500 GB, and a point-in-time-restore to anypoint in the last 35 days.

NEW QUESTION 36

- (Exam Topic 7)

You administer a SQL Server 2014 server that contains a database named SalesDB. SalesDb contains a

schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales.

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

You need to ensure that the Sales role, including 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: F

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-warehouse->

NEW QUESTION 39

- (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 41

- (Exam Topic 7)

Background

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

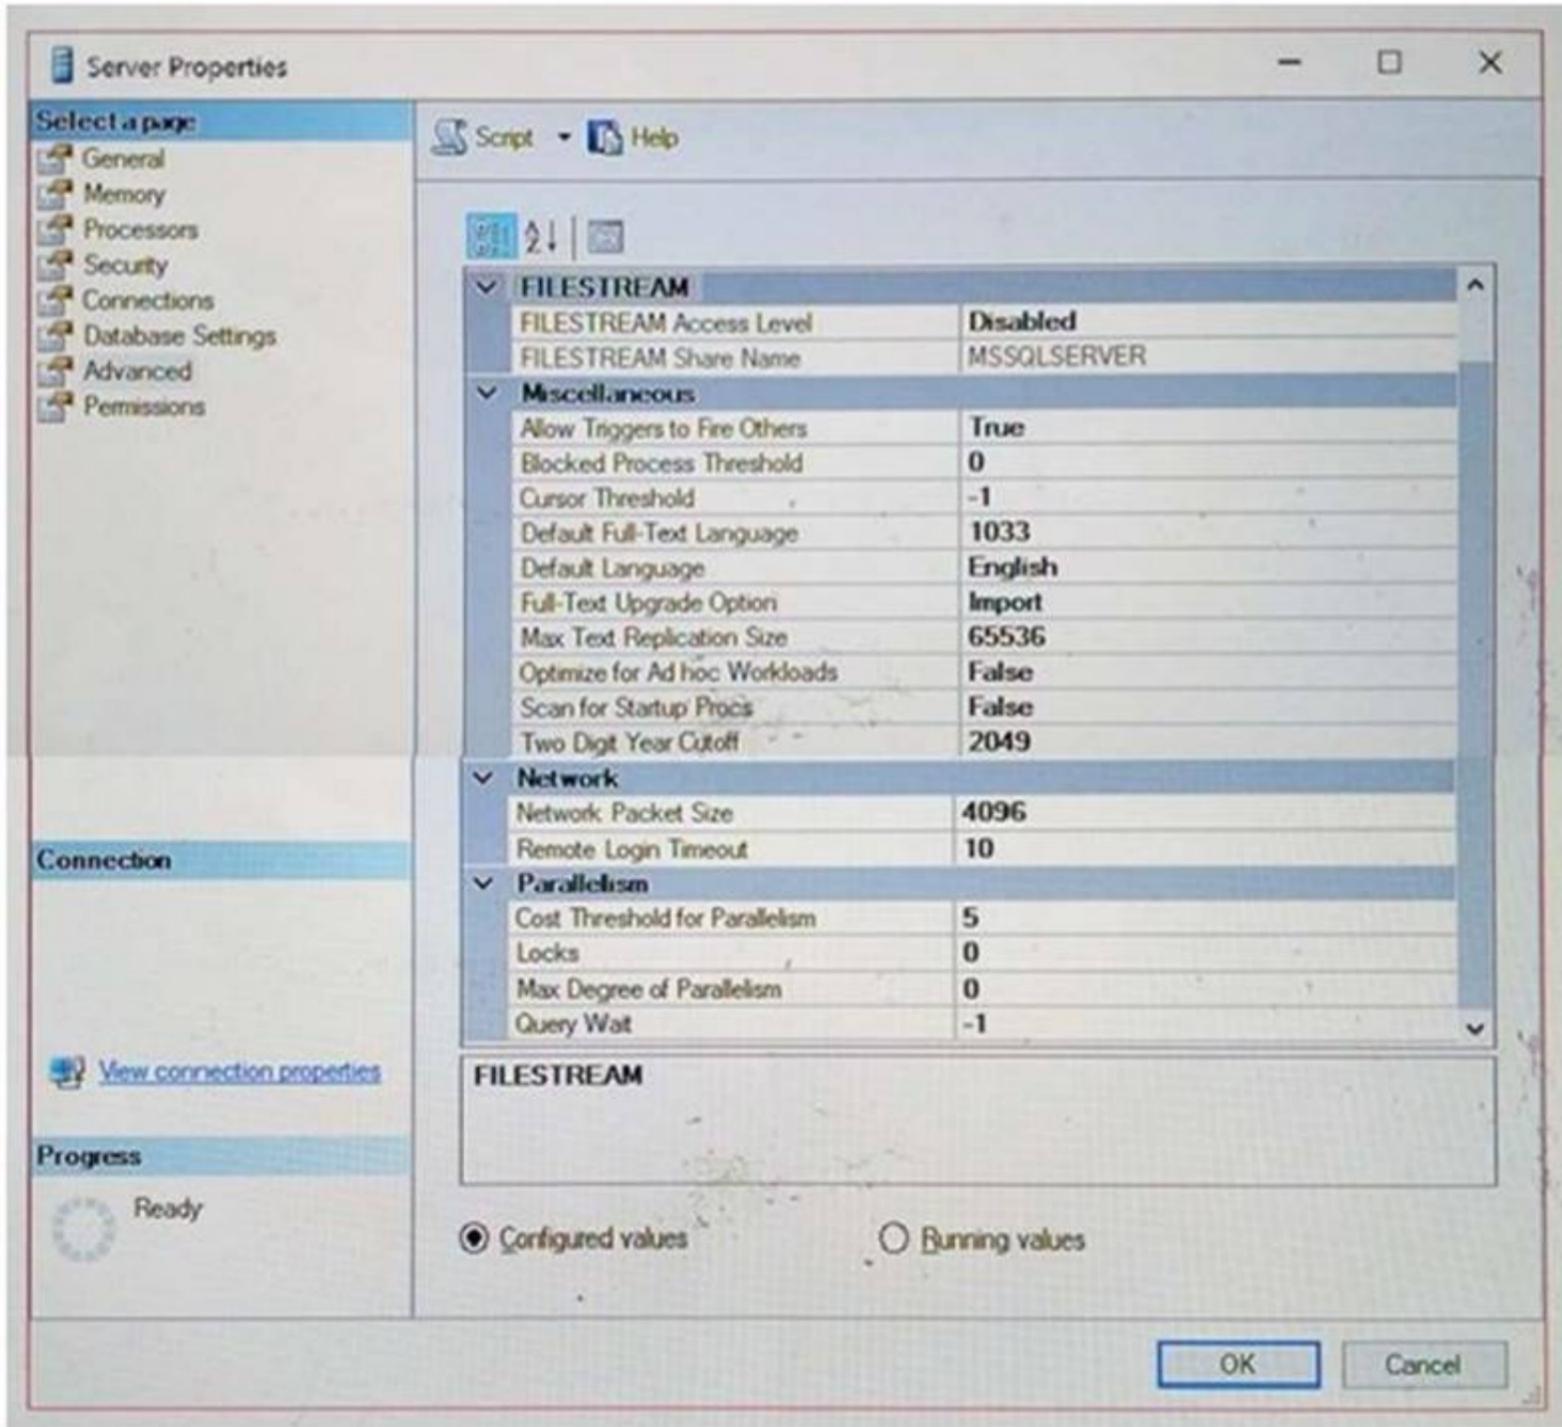
The environment includes the following servers: SRV1 and SRV2. SRV1 has 16 logical cores and hosts a SQL Server instance that supports a mission-critical application. The application has approximately 30,000 concurrent users and relies heavily on the use of temporary tables.

The environment also includes the following databases: DB1, DB2, and Reporting. The Reporting database is protected with Transparent Data Encryption (TDE). You plan to migrate this database to a new server. You detach the database and copy it to the new server.

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

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

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



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

Answer Area

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

UPDATE	MODIFY FILE
ALTER	UPDATE FILE

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

ALTER DATABASE AdventureWorks2012 MODIFY FILE
 (NAME = test1dat3, SIZE = 200MB); Note: MODIFY FILE

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

References:

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

NEW QUESTION 45

- (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]
(
    (
        [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 50

- (Exam Topic 7)

You plan to migrate on-premises Microsoft SQL Server to SQL Server on a Microsoft Azure virtual machine. You need to ensure that the Azure virtual machine can handle the workload.

Which tool should you use for each environment? To answer, drag the appropriate tools to the correct options. Each tool may be used once. More than once, or not at all.

Tools, Select from these.

Answer Area

- Distributed Replay
- Performance Monitor
- SQL Server Profiler
- SQL Server Extended Events
- SQL Server Data Tools (SSDT)

Tool to use on-premises: <Place here>
 Tool to use in Azure: <Place here>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Tools, Select from these.

Answer Area

- Distributed Replay
- Performance Monitor
- SQL Server Profiler
- SQL Server Extended Events
- SQL Server Data Tools (SSDT)

Tool to use on-premises: SQL Server Profiler
 Tool to use in Azure: SQL Server Data Tools (SSDT)

NEW QUESTION 54

- (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 59

- (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 61

- (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 65

- (Exam Topic 7)

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

A network administrator upgrades the hardware firewalls on the network. You need to verify whether data migration still runs successfully.

Which stored procedure should you run?

- A. Sys_sp_testlinkedserver
- B. Sys_sp_rda_test_connection
- C. Sys_sp_rda_reauthorized_db
- D. Sp_set_firewall_rule

Answer: B

Explanation:

The Sys_sp_rda_test_connection cmdlet tests the connection from SQL Server to the remote Azure server and reports problems that may prevent data migration.

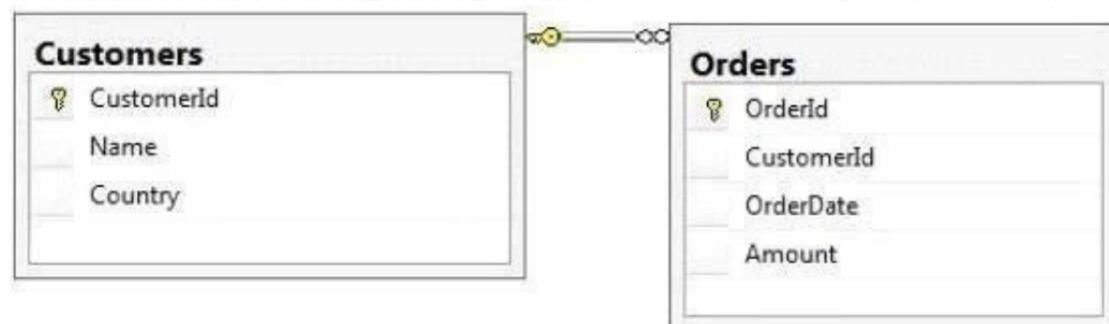
References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-rda-test-connection-tr>

NEW QUESTION 70

- (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 = 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: E

NEW QUESTION 75

- (Exam Topic 7)

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

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

What should you do?

- A. Back up the master database on the SQL2008 instance
- B. Restore the master database on the SQL2012 instance
- C. Use the Transfer Logins task in a Microsoft SQL Server Integrated Services package
- D. Use sp_grantlogin
- E. Use xp_logininfo.

Answer: C

Explanation:

sp_grantlogin creates a SQL Server login.

NEW QUESTION 78

- (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 82

- (Exam Topic 7)

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

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

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

D. Start the SQL Server Browser Service.

Answer: D

Explanation:

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

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

NEW QUESTION 85

- (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 88

- (Exam Topic 7)

Database DB1 must use two CPU cores.

Queries that were running on database DB2 prior to migration do not complete. You need to configure the databases.

In the table below, identify the parameter that must be configured for each databases. Select one option for DB1, and one option for DB2. Select one option for each column.

Parameter	DB1	DB2
MAXDOP	<input type="radio"/>	<input type="radio"/>
LEGACY_CARDINALITY_ESTIMATION	<input type="radio"/>	<input type="radio"/>
PARAMETER_SNIFFING	<input type="radio"/>	<input type="radio"/>
QUERY_OPTIMIZER_HOTFIXES	<input type="radio"/>	<input type="radio"/>
CLEAR PROCEDURE_CACHE	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Query_optimizer_hotfixes DB1: MAXDOP

You can use the max degree of parallelism (MAXDOP) option to limit the number of processors to use in parallel plan execution.

DB2: LEGACY_CARDINALITY_ESTIMATION

The CE (Cardinality Estimation) predicts how many rows your query will likely return. The cardinality prediction is used by the Query Optimizer to generate the optimal query plan. With more accurate estimations, the Query Optimizer can usually do a better job of producing a more optimal query plan.

Legacy CE: For a SQL Server database set at compatibility level 120 and above, the CE version 70 can be activated by using the at the database level by using the ALTER DATABASE SCOPED CONFIGURATION.

Example:

```
ALTER DATABASE SCOPED CONFIGURATION SET LEGACY_CARDINALITY_ESTIMATION = ON; GO
```

NEW QUESTION 91

- (Exam Topic 7)

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 95

- (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 99

- (Exam Topic 7)

You have a database named DB1 that contains a table named Table1. Table1 has a non-clustered index named index1.

You discover that index1 is corrupt. You need to repair index1.

Which statement should you execute?

- A. DBCC CHECKDB ('db1', REPAIR_FAST)
- B. ALTER INDEX indx1 ON table1 REBUILD WITH (ONLINE=ON)
- C. ALTER INDEX index1 ON table1 REORGANIZE
- D. DBCC CHECKDB ('db1', DATA_PURITY)

Answer: B

Explanation:

If REBUILD is performed online (ON) the data in this table is available for queries and data modification during the index operation.

NEW QUESTION 102

- (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 103

- (Exam Topic 7)

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

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

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

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

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

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

- A. Yes
- B. No

Answer: A

Explanation:

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

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

NEW QUESTION 105

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance named SQL2012 that hosts an OLTP database of 1 terabyte in size.

The database is modified by users only from Monday through Friday from 09:00 hours to 17:00 hours. Users modify more than 30 percent of the data in the database during the week.

Backups are performed as shown in the following schedule:

Type	Frequency
Full	Sunday at 20:00 hours
Differential	Monday through Friday at 20:00 hours
Log	Monday through Friday between 08:00 hours and 18:00 hours

The Finance department plans to execute a batch process every Saturday at 09:00 hours. This batch process will take a maximum of 8 hours to complete.

The batch process will update three tables that are 10 GB in size. The batch process will update these tables multiple times.

When the batch process completes, the Finance department runs a report to find out whether the batch process has completed correctly.

You need to ensure that if the Finance department disapproves the batch process, the batch operation can be rolled back in the minimum amount of time. What should you do on Saturday?

- A. Perform a differential backup at 08:59 hours.
- B. Record the LSN of the transaction log at 08:59 hour
- C. Perform a transaction log backup at 17:01 hours.
- D. Create a database snapshot at 08:59 hours.
- E. Record the LSN of the transaction log at 08:59 hour
- F. Perform a transaction log backup at 08:59 hours.
- G. Create a marked transaction in the transaction log at 08:59 hour
- H. Perform a transaction log backup at 17:01 hours.
- I. Create a marked transaction in the transaction log at 08:59 hour
- J. Perform a transaction log backup at 08:59 hours.

Answer: C

Explanation:

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

NEW QUESTION 106

- (Exam Topic 7)

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

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

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

Answer: C

Explanation:

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

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

the SQL Server process hosting the Extended Events engine.

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

NEW QUESTION 107

- (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 110

- (Exam Topic 7)

You develop a Microsoft SQL Server 2014 database that contains a heap named OrdersHistorical. You write the following Transact-SQL query:

```
INSERT INTO OrdersHistorical SELECT * FROM CompletedOrders
```

You need to optimize transaction logging and locking for the statement. Which table hint should you use?

- A. HOLDLOCK
- B. ROWLOCK
- C. XLOCK
- D. UPDLOCK
- E. TABLOCK

Answer: E

Explanation:

When importing data into a heap by using the INSERT INTO SELECT <columns> FROM statement, you can enable optimized logging and locking for the statement by specifying the TABLOCK hint for the target table.

References:<https://docs.microsoft.com/en-us/sql/t-sql/queries/hints-transact-sql-table>

NEW QUESTION 113

- (Exam Topic 7)

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

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

Which permissions should you assign to Dbal?

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

Answer: B

Explanation:

To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission. There are two types of dynamic management views and functions:

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

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

References:

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

NEW QUESTION 118

- (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 121

- (Exam Topic 7)

You create a new Microsoft Azure subscription.

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 122

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database. You want to make a full backup of the database to a file on disk.

In doing so, you need to output the progress of the backup. Which backup option should you use?

- A. STATS
- B. COMPRESSION
- C. CHECKSUM
- D. IN IT

Answer: A

Explanation:

STATS is a monitoring option of the BACKUP command. STATS [=percentage]

Displays a message each time another percentage completes, and is used to gauge progress. If percentage is omitted, SQL Server displays a message after each 10 percent is completed.

The STATS option reports the percentage complete as of the threshold for reporting the next interval. This is at approximately the specified percentage; for example, with STATS=10, if the amount completed is 40 percent, the option might display 43 percent. For large backup sets, this is not a problem, because the percentage complete moves very slowly between completed I/O calls.

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

NEW QUESTION 125

- (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 128

- (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 131

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

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 First Try
- * 70-765 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 70-765 Practice Test Here](#)