



MCSA / MCSE for Microsoft SQL Server 2016

Exam 70-765 Provisioning SQL Databases

Version 15.35 (233 Questions)

Microsoft
CERTIFIED

Solutions Associate

SQL 2016
Database Development

Microsoft
CERTIFIED

Solutions Expert

Data Management and
Analytics

NO.1 You are deploying a Microsoft SQL Server database that will support a mixed OLTP and OLAP workload.

The target virtual machine has four CPUs.

You need to ensure that reports do not use all available system resources.

What should you do?

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

Answer: C

Explanation

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

NO.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.

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

NO.3 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>

NO.4 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals.

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

You need to create the templates.

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation

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

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

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

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

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

NO.5 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>

NO.6 You use Resource Manager to deploy a new Microsoft SQL Server instance in a Microsoft Azure virtual machine (VM) that uses Premium storage. The combined initial size of the SQL Server user database files is expected to be over 200 gigabytes (GB). You must maximize performance for the database files and the log file.

You add the following additional drive volumes to the VM:

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

You have the following requirements:

You need to deploy the SQL instance.

In the table below, identify the drive where you must store each SQL Server file type.

NOTE: Make only one selection in each column. Each correct selection is worth one point.

Answer area

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

Answer:**Answer area**

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

Explanation

Answer area

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

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

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

NO.7 Note: This questions is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

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

You must be able to manage the databases as individual entities by using SQL Server Management Studio. All data in the databases must be stored encrypted. Backups must be encrypted by using the same key as the live copy of the database.

You need to secure the data.

What should you implement?

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

Answer: C

Explanation

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

TDE encrypts the storage of an entire database by using a symmetric key called the database encryption key.

In SQL Database the database encryption key is protected by a built-in server certificate. The built-in server certificate is unique for each SQL Database server.

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

NO.8 You manage a Microsoft SQL Server environment in a Microsoft Azure virtual machine.

You must enable Always Encrypted for columns in a database.

You need to configure the key store provider.

What should you do?

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

Answer: D

Explanation

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

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

NO.9 Note: This question is part of a series of questions that present the same scenario. Each

question in the series contains a unique solution. Determine whether the solution meets stated goals.

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

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation

Use a Resource Governor instead.

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

NO.10 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>

NO.11 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals.

You have a mission-critical application that stores data in a Microsoft SQL Server instance. The application runs several financial reports. The reports use a SQL Server-authenticated login named

Reporting_User. All queries that write data to the database use Windows authentication. Users report that the queries used to provide data for the financial reports take a long time to complete. The queries consume the majority of CPU and memory resources on the database server. As a result, read-write queries for the application also take a long time to complete. You need to improve performance of the application while still allowing the report queries to finish. Solution: You configure the Resource Governor to set the MAXDOP parameter to 0 for all queries against the database.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation

SQL Server will consider parallel execution plans for queries, index data definition language (DDL) operations, and static and keyset-driven cursor population.

You can override the max degree of parallelism value in queries by specifying the MAXDOP query hint in the query statement.

References: [https://technet.microsoft.com/en-us/library/ms181007\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms181007(v=sql.105).aspx)

NO.12 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.

Answer:

Actions

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.

Answer Area

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:

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

Run the following Transact-SQL statement:

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

Explanation

Answer Area

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:

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

Run the following Transact-SQL statement:

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

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>

NO.13 You have a Microsoft SQL Server 2014 named SRV2014 that has a single tempdb database file. The tempdb database file is eight gigabytes (GB) in size.

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

You plan to migrate the databases from SRV2014 to SRV2016.

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

What should you do?

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

Answer: B

Explanation

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

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

2. Configuration options:

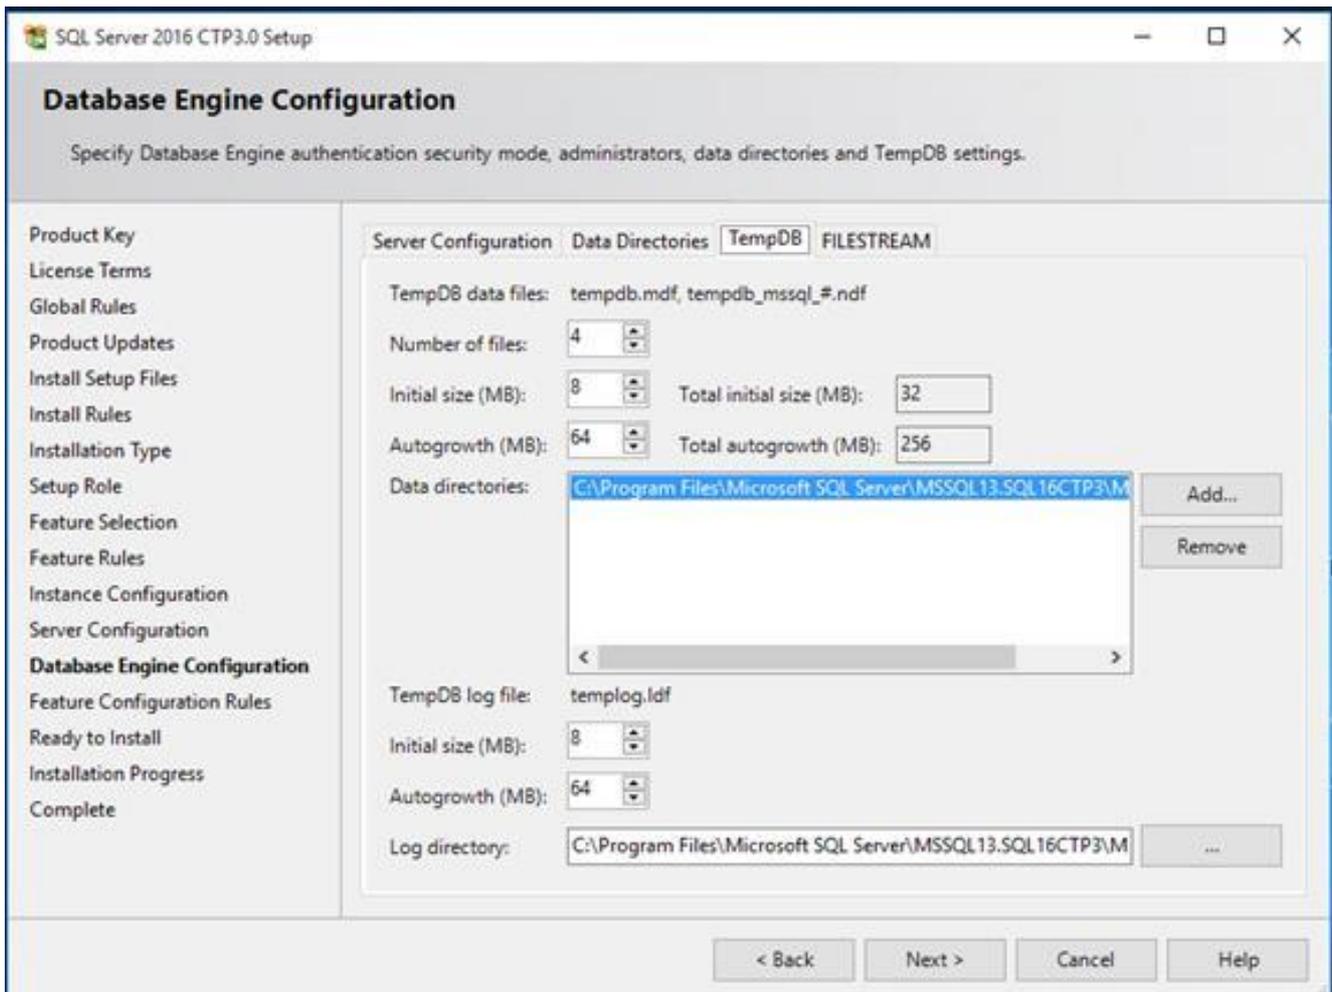
Data Files

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

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

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

Figure:



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

NO.14 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 a virtual machine (VM) in Microsoft Azure, which has a 2 terabyte (TB) database. Microsoft SQL Server backups are performed by using Backup to URL.

You need to provision the storage account for the backups 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: G

Explanation

A URL specifies a Uniform Resource Identifier (URI) to a unique backup file. The URL is used to provide the location and name of the SQL Server backup file. The URL must point to an actual blob, not just a container.

If the blob does not exist, it is created. If an existing blob is specified, BACKUP fails, unless the "WITH FORMAT" option is specified to overwrite the existing backup file in the blob.

LOCALLY REDUNDANT STORAGE (LRS) makes multiple synchronous copies of your data within a single datacenter.

NO.15 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals.

You manage a Microsoft SQL Server environment with several databases.

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

Solution: You enable the LEGACY_CARDINALITY_ESTIMATION option for the databases.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation

LEGACY_CARDINALITY_ESTIMATION = { ON | OFF | PRIMARY }

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

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

NO.16 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

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 >

ProductionNetwork

* Subnet >

ProductionServers (10.1.0.0/24)

* Public IP address >

None/Not Assigned

* Network security group >

(new) SQLServers

Extensions

Extensions >

No extensions

Monitoring

Diagnostics

Disabled Enabled

Availability

* Availability set >

ProductionServers2

OK