

## Microsoft MCSA Certification 70-765 Exam



- **Vendor: Microsoft**
- **Exam Code: 70-765**
- **Exam Name: Provisioning SQL Databases**

**Get Complete Version Exam 70-765 Dumps with VCE and PDF Here**



<https://www.passleader.com/70-765.html>

**QUESTION 1**

Hotspot Question

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:

- Maximize performance of the SQL Server instance.
- Use Premium storage when possible.

You need to deploy the SQL instance. In the table below, identify the drive where you must store each SQL Server file type. (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:**

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.

**QUESTION 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.**

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.

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

**QUESTION 3**

Drag and Drop Question

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. 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

**Answer:**

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

**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.

**QUESTION 4**

You have a Microsoft SQL Server 2014 named SRV2014 that has a single tempdb database file.

The tempdb database file is 8 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 8GB.  
In the tempdb database, set the value of the MAXDOP property to 8.
- B. Increase the size of the tempdb data files to 1GB.
- 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 128MB.

**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**

- a. **Number of files** - this will default to the lower value of 8 or number of logical cores as detected by setup. The value can be increased as required for a specific workload. The file names for secondary data files will follow the tempdb\_mssql\_#.ndf naming convention where # represents the file ordinal for each additional file.
- b. **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.
- c. **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. Since **TF 1117 is enabled by default for tempdb**, all files will grow at the same time by the specified size. Total autogrowth reflects the cumulative size the database will grow each time an autogrow is triggered.

**Note:** If you specify a very large initial size or autogrowth value, ensure that **Instant File Initialization** is enabled for SQL Server service startup account.

- d. **Data directories** - specify multiple folders/drives to spread the data files across several volumes. Each file will be placed in a round-robin manner. For example: if you have specified 8 data files and 3 volumes - files 1,4,7 will go to vol 1; files 2,5,8 will go to vol 2; files 3,6 will go to vol 3.

**Log File**

- a. **Initial size** - is specified in MB and log file with that size is created.
  - b. **Autogrowth** - is specified in MB (fixed growth is preferred as opposed to a non-linear percentage based growth)
- A default value of 64MB is provided to so that the number of Virtual Log Files (VLFs) during initial creation is a small and manageable number and with appropriate size so that the unused log space can be reclaimed easily.

The screenshot shows the 'Database Engine Configuration' window for SQL Server 2016 CTP3.0 Setup. The window has a left sidebar with navigation options: Product Key, License Terms, Global Rules, Product Updates, Install Setup Files, Install Rules, Installation Type, Setup Role, Feature Selection, Feature Rules, Instance Configuration, Server Configuration, Database Engine Configuration (selected), Feature Configuration Rules, Ready to Install, Installation Progress, and Complete. The main area is titled 'Database Engine Configuration' and includes the instruction: 'Specify Database Engine authentication security mode, administrators, data directories and TempDB settings.' There are four tabs: Server Configuration, Data Directories, TempDB, and FILESTREAM. The 'TempDB' tab is active, showing the following settings: TempDB data files: tempdb.mdf, tempdb\_mssql\_#.ndf; Number of files: 4; Initial size (MB): 8; Total initial size (MB): 32; Autogrowth (MB): 64; Total autogrowth (MB): 256; Data directories: C:\Program Files\Microsoft SQL Server\MSSQL13.SQL16CTP3\M (with an 'Add...' button); TempDB log file: templog.ldf; Initial size (MB): 8; Autogrowth (MB): 64; Log directory: C:\Program Files\Microsoft SQL Server\MSSQL13.SQL16CTP3\M (with an '...' button). At the bottom are buttons for '< Back', 'Next >', 'Cancel', and 'Help'.

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

#### QUESTION 5

**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.

**QUESTION 6**

**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**

**QUESTION 7**

Hotspot Question

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. 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

**Answer:**

**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

**Explanation:**

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

**QUESTION 8**

**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:**

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

### QUESTION 9

Hotspot Question

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.)

<div style="border: 1px solid black; padding: 2px;">▼</div> <div style="border: 1px solid black; padding: 2px;">CREATE USER</div> <div style="border: 1px solid black; padding: 2px;">ALTER SERVER ROLE</div> <div style="border: 1px solid black; padding: 2px;">CREATE LOGIN</div>	Admin2 WITH password = 'Pa\$\$w0rd';		
<div style="border: 1px solid black; padding: 2px;">▼</div> <div style="border: 1px solid black; padding: 2px;">CREATE USER</div> <div style="border: 1px solid black; padding: 2px;">ALTER SERVER ROLE</div> <div style="border: 1px solid black; padding: 2px;">CREATE LOGIN</div>	Admin2User FROM	<div style="border: 1px solid black; padding: 2px;">▼</div> <div style="border: 1px solid black; padding: 2px;">WINDOWS</div> <div style="border: 1px solid black; padding: 2px;">EXTERNAL PROVIDER</div> <div style="border: 1px solid black; padding: 2px;">LOGIN</div>	Admin2
ALTER ROLE '	<div style="border: 1px solid black; padding: 2px;">▼</div> <div style="border: 1px solid black; padding: 2px;">loginmanager</div> <div style="border: 1px solid black; padding: 2px;">dbmanager</div> <div style="border: 1px solid black; padding: 2px;">bd_ddladmin</div>		

**Answer:**

▼	Admin2 WITH password = 'Pa\$\$w0rd';	
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	

**Explanation:**

<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>

**QUESTION 10**

**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.

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

**QUESTION 11**

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.

**QUESTION 12**

**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:**

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

**QUESTION 13**

Drag and Drop Question

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:**

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:

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.

Step 3:

Grant the proper reading permissions.

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

**QUESTION 14**

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.

**[70-765 Exam Dumps](#) [70-765 Exam Questions](#) [70-765 PDF Dumps](#) [70-765 VCE Dumps](#)**

**[Back to the Source of this PDF & Get More Free Braindumps -- www.microsoftbraindumps.com](#)**



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

#### **QUESTION 15**

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

#### **QUESTION 16**

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

#### **QUESTION 17**

**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 manage on-premises and Microsoft Azure SQL Database instances for a company. Your environment must support the Microsoft SQL Server 2012 ODBC driver. You need to encrypt only specific columns in the database. What should you implement?

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

**Answer: D**

**Explanation:**

To encrypt columns you can configure Always Encrypted. SQL Server Management Studio (SSMS) provides a wizard that helps you easily configure Always Encrypted by setting up the column master key, column encryption key, and encrypted columns for you. Always Encrypted allows client applications to encrypt sensitive data and never reveal the data or the encryption keys to SQL Server or Azure SQL Database. An Always Encrypted enabled driver, such as the ODBC Driver 13.1 for SQL Server, achieves this by transparently encrypting and decrypting sensitive data in the client application.

**Note:**

The ODBC driver automatically determines which query parameters correspond to sensitive database columns (protected using Always Encrypted), and encrypts the values of those parameters before passing the data to SQL Server or Azure SQL Database. Similarly, the driver transparently decrypts data retrieved from encrypted database columns in query results.

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault#encrypt-columns-configure-always-encrypted>

[https://msdn.microsoft.com/en-us/library/mt637351\(v=sql.110\).aspx](https://msdn.microsoft.com/en-us/library/mt637351(v=sql.110).aspx)

#### **QUESTION 18**

**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.**

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 FileSystem
- 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.

**QUESTION 19**

Drag and Drop Question

You deploy a new Microsoft Azure SQL Database instance to support a variety of mobile applications and public websites. You plan to create a new security principal named User1. The principal must have access to select all current and future objects in a database named Reporting. The activity and authentication of the database user must be limited to the Reporting 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
In SQL Server Management Studio, create a connection to the <b>Reporting</b> database on the Azure SQL Server instance.	
In SQL Server Management Studio, create a connection to the master database on the Azure SQL Server instance.	
Run the following Transact-SQL statement:  EXEC sp_addrolemember 'db_datareader', 'User1'	
Run the following Transact-SQL statement:  CREATE LOGIN User1 WITH password='Pa\$\$w0rd'	
Run the following Transact-SQL statement:  CREATE USER User1 WITH password='Pa\$\$w0rd'	
Run the following Transact-SQL statements:  EXEC sp_migrate_user_to_contained @username = N'User1', @rename = N'keep_name', @disablelogin = N'disable_login'	
Run the following Transact-SQL statement:  CREATE LOGIN User1 FROM EXTERNAL PROVIDER	
Select the Reporting database and run the following Transact-SQL statements:  CREATE USER User1 from LOGIN User1 GRANT SELECT TO User1	

**Answer:**

Actions	Answer Area
In SQL Server Management Studio, create a connection to the <b>Reporting</b> database on the Azure SQL Server instance.	In SQL Server Management Studio, create a connection to the master database on the Azure SQL Server instance.
In SQL Server Management Studio, create a connection to the master database on the Azure SQL Server instance.	Run the following Transact_SQL statement:
Run the following Transact-SQL statement:  EXEC sp_addrolemember 'db_datareader', 'User1'	CREATE LOGIN User1 WITH password='Pa\$\$w0rd'
Run the following Transact_SQL statement:  CREATE LOGIN User1 WITH password='Pa\$\$w0rd'	Select the Reporting database and run the following Transact-SQL statements:
Run the following Transact_SQL statement:  CREATE USER User1 WITH password='Pa\$\$w0rd'	CREATE USER User1 from LOGIN User1 GRANT SELECT TO User1
Run the following Transact_SQL statements:  EXEC sp_migrate_user_to_contained @username = N'User1', @rename = N'keep_name', @disablelogin = N'disable_login'	
Run the following Transact_SQL statement:  CREATE LOGIN User1 FROM EXTERNAL PROVIDER	
Select the Reporting database and run the following Transact-SQL statements:  CREATE USER User1 from LOGIN User1 GRANT SELECT TO User1	

**Explanation:**

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

**QUESTION 20**

**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.

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

**QUESTION 21**

**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.

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

**QUESTION 22**

**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.

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

**QUESTION 23**

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**

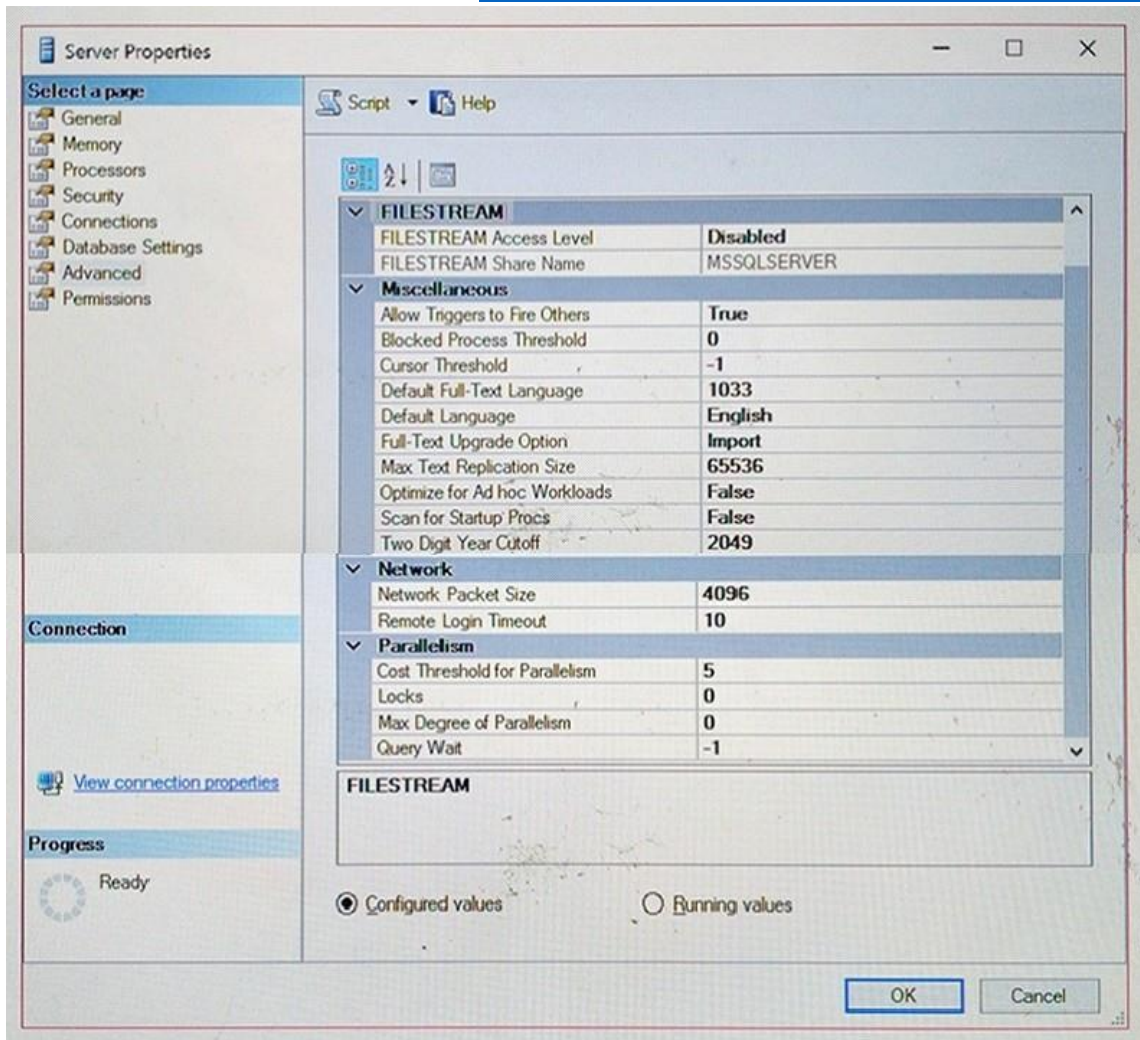
**Explanation:**

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

## **Case Study 1 - Automobile Parts**

### **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.



#### QUESTION 24

##### Hotspot Question

You need to resolve the identified issues. Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic.

##### Answer Area

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

Optimize for Ad Hoc workload  
Max Degree of Parallelism  
Query Wait

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

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

**Answer:**



**Answer Area**

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

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

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

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

**Explanation:**

From exhibit we see:

- Cost Threshold of Parallelism: 5
- Optimize for Ad Hoc Workloads: false
- Max Degree of Parallelism: 0 (This is the default setting, which enables the server to determine the maximum degree of parallelism. It is fine.)
- Locks: 0
- Query Wait: -1

Box 1: Optimize for Ad Hoc Workload. Change the Optimize for Ad Hoc Workload setting from false to 1/True. The optimize for ad hoc workloads option is used to improve the efficiency of the plan cache for workloads that contain many single use ad hoc batches. When this option is set to 1, the Database Engine stores a small compiled plan stub in the plan cache when a batch is compiled for the first time, instead of the full compiled plan. This helps to relieve memory pressure by not allowing the plan cache to become filled with compiled plans that are not reused.

**QUESTION 25**

Hotspot Question

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.

**Answer:**

**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.

**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.

## Case Study 2 - Contoso, Ltd.

### Background

You are the database administrator for Contoso, Ltd. The company has 200 offices around the world. The company has corporate executives that are located in offices in London, New York, Toronto, Sydney, and Tokyo. Contoso, Ltd. has a Microsoft Azure SQL Database environment. You plan to deploy a new Azure SQL Database to support a variety of mobile applications and public websites. The company is deploying a multi-tenant environment. The environment will host Azure SQL Database instances. The company plans to make the instances available to internal departments and partner companies. Contoso is in the final stages of setting up networking and communications for the environment. Existing Contoso and Customer instances need to be migrated to Azure virtual machines (VM) according to the following requirements:

- Contoso instances - should use the method requiring the least administrative effort to migrate instances to Azure Vms.
- Customer instances - should use a method that allows customers to bring their own licenses to Azure VMs. Customers have approved down time for the migration.

The company plans to deploy a new order entry application and a new business intelligence and analysis application. Each application will be supported by a new database. Contoso creates a new Azure SQL database named Reporting. The database will be used to support the company's financial reporting requirements. You associate the database with the Contoso Azure Active Directory domain. Each location database for the data entry application may have an unpredictable amount of activity. Data must be replicated to secondary databases in Azure datacenters in

different regions. To support the application, you need to create a database named contosodb1 in the existing environment.

#### Objects

Parameter	Name
Logical Server	contososrv
Resource Group	contosodbrg

#### Database

The contosodb1 database must support the following requirements:

- a size of at least 200 gigabytes (GB).
- 1,000 concurrent sessions.
- point-in-time restore to any point in the two weeks prior to a failure.
- minimize costs.

#### Application

For the business intelligence application, corporate executives must be able to view all data in near real-time with low network latency. Contoso has the following security, networking, and communications requirements:

- Multi-Location Load Balancing - to ensure customers have access to their tenants at multiple Azure locations across the world.
- Secure Message/Data Flow - to securely support communication between Azure and on-premises applications and services.
- Accounts should support accessing external domain resources and be configured in the most secure and lowest-maintenance way possible, including meeting the company policy of regular service account password changes.

#### QUESTION 26

##### Hotspot Question

You need to configure the data entry and business intelligence databases. In the table below, identify the option that you must use for each database. (Make only one selection in each column.)

##### Answer Area

Option	Data entry	Business intelligence
Elastic database pools only	<input type="radio"/>	<input type="radio"/>
Geo-replicated database only	<input type="radio"/>	<input type="radio"/>
Elastic database pools and geo-replicated databases	<input type="radio"/>	<input type="radio"/>

**Answer:**



**Answer Area**

Option	Data entry	Business intelligence
Elastic database pools only	<input type="radio"/>	<input checked="" type="radio"/>
Geo-replicated database only	<input checked="" type="radio"/>	<input type="radio"/>
Elastic database pools and geo-replicated databases	<input type="radio"/>	<input type="radio"/>

**Explanation:**

Data Entry: Geo-replicated database only. From Contoso scenario: Each location database for the data entry application may have an unpredictable amount of activity. Data must be replicated to secondary databases in Azure datacenters in different regions.

Business intelligence: Elastic database pools only. From Contoso scenario: For the business intelligence application, corporate executives must be able to view all data in near real-time with low network latency. SQL DB elastic pools provide a simple cost effective solution to manage the performance goals for multiple databases that have widely varying and unpredictable usage patterns.

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

**QUESTION 27**

Hotspot Question

You need to create the contososdb1 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

**Answer:**

**Answer Area**

	▼
New-AzureSqlDatabase	
New-AzureRmSqlDatabase	
Set-AzureRmSqlDatabase	

- ResourceGroupName "contosodbrg"

- ServerName "contososrv"

- DatabaseName "contosodbl"

- Edition

	▼
Basic	
Standard	
Premium	

- RequestedServiceObjectName S2

**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.

**Case Study 3 - SQL Server Reporting**

**Background**

You manage a Microsoft SQL Server environment that includes the following databases: DB1, DB2.

**Reporting.**

The environment also includes SQL Reporting Services (SSRS) and SQL Server Analysis Services (SSAS). All SSRS and SSAS servers use named instances. You configure a firewall rule for SSAS.

**Databases**

**DB1.**

This database was migrated from SQL Server 2012 to SQL Server 2016. 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.

**DB2.**

This database was migrated from SQL Server 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.

### Reporting.

You create a SQL Server-authenticated login named BIAppUser on the SQL Server instance to support users of the Reporting database. The BIAppUser login is not a member of the sysadmin role. You plan to configure performance-monitoring alerts for this instance by using SQL Agent Alerts.

### QUESTION 28

Drag and Drop Question

You create a login named BIAppUser. The login must be able to access the Reporting database. You need to grant access to the BIAppUser login in the database. How should you complete the Transact-SQL statements? (To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.)

Code segments	Answer area
Reporting	USE [ Code segment ]
master	GO
CREATE USER	Code segment [BIAppUser] Code segment
ALTER LOGIN	GO
ALTER USER	
FOR LOGIN [BIAppUser]	
FOR USER [BIAppUser]	
WITH LOGIN = [BIAppUser]	

**Answer:**

Code segments	Answer area
Reporting	USE [ Reporting ]
master	GO
CREATE USER	CREATE USER [BIAppUser] FOR LOGIN [BIAppUser]
ALTER LOGIN	GO
ALTER USER	
FOR LOGIN [BIAppUser]	
FOR USER [BIAppUser]	
WITH LOGIN = [BIAppUser]	

### Explanation:

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

### QUESTION 29

Hotspot Question

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. (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"/>

**Answer:**

**Answer Area**

Database setting	DB1	DB2
DELAYED_DURABILITY = FORCED	<input checked="" 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 checked="" type="radio"/>
AUTO_UPDATE_STATISTICS_ASYNC ON	<input type="radio"/>	<input type="radio"/>

**Explanation:**

<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)

**QUESTION 30**

Hotspot Question

You need to set up the service accounts that the database engine and SQL Server Agent services will use. How should you design the solution? (To answer, select the appropriate configuration options in the answer area.)

**Answer Area**

Design option	Configuration option
Account type	Domain account
	Local machine account
	Local system account
Group membership	Domain administrators
	Local administrators
	Domain users
Password management	Manually-managed passwords
	Managed Service accounts

**Answer:**

**Answer Area**

Design option	Configuration option
Account type	Domain account
	Local machine account
	Local system account
Group membership	Domain administrators
	Local administrators
	Domain users
Password management	Manually-managed passwords
	Managed Service accounts

**Explanation:**

<https://msdn.microsoft.com/en-us/library/ms191543.aspx>

**QUESTION 31**

Hotspot Question

You need to open the firewall ports for use with SQL Server environment. In table below, identify the firewall port that you must use for each service. (Make only one selection in each column.)

**Answer Area**

Port number	Report Server	SQL Server Browser service for SSAS
80	<input type="radio"/>	<input type="radio"/>
135	<input type="radio"/>	<input type="radio"/>
1433	<input type="radio"/>	<input type="radio"/>
2382	<input type="radio"/>	<input type="radio"/>

**Answer:**

**Answer Area**

Port number	Report Server	SQL Server Browser service for SSAS
80	<input checked="" type="radio"/>	<input type="radio"/>
135	<input type="radio"/>	<input type="radio"/>
1433	<input type="radio"/>	<input checked="" type="radio"/>
2382	<input type="radio"/>	<input type="radio"/>

**Explanation:**

Report Server: 80. By default, the report server listens for HTTP requests on port 80.

**QUESTION 32**

**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

<https://docs.microsoft.com/en-us/azure/storage/storage-premium-storage>

### QUESTION 33

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

- A. Standard S3 pricing tier
- B. Premium P2 tier
- C. Standard S2 pricing tier
- D. Premium P1 tier

**Answer: D**

**Explanation:**

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

Premium service tier

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

### QUESTION 34

**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

[70-765 Exam Dumps](#) [70-765 Exam Questions](#) [70-765 PDF Dumps](#) [70-765 VCE Dumps](#)

[Back to the Source of this PDF & Get More Free Braindumps -- www.microsoftbraindumps.com](#)

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.

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

#### QUESTION 35

**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:**

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.

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-automated-backups>

#### QUESTION 36

**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:**

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

**QUESTION 37**

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 any point in the last 35 days.

Incorrect:

Not D. The Basic service tier only supports a maximum database size of 2 GB, a maximum of 5 DTUs, and a point-in-time-restore of 7 days.

Not C. The Standard tier with a performance level of SO only supports a maximum database size of 250 GB. A maximum of 10 DTUs, and a point-in-time-restore of 35 days.

Not A. The Standard tier with a performance level of S3 only supports a maximum database size of 250 GB. A maximum of 100 DTUs, and a point-in-time-restore of 35 days.

**QUESTION 38**

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. Choose two.)

- 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. You should not run CREATE MASTER KEY on the master database. You would execute CREATE MASTER KEY when you have to encrypt an on-premises database. You should not run CREATE CERTIFICATE cert1 WITH Subject = TDE Cert' on DB1. You do not need to create a certificate to encrypt an Azure SQL database. This would be a part of the solution when you encrypt an on-premises database. You should not connect to the master database. To encrypt DB1, you need to connect directly to DB1. not to the master database.

**QUESTION 39**

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. Choose two.)

- 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 SQL1 and SQL2 in a single availability set. 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.

- You should deploy AppServer1 and AppServer2 in a single availability set. Using availability sets allows to build in redundancy for your Azure services. By grouping related virtual machines and services (tiers) an availability set (in this case, deploying both of your application servers into an availability set), ensure that if there is a planned or unplanned outage, your services will remain available.

- You should not deploy SQL1 and AppServer1 in a single availability set. Microsoft recommends grouping similar sets of services in an availability set. This will help ensure that each tier of your application is protected by an availability set. SQL1 and AppServer1 provide different services, so they should not be grouped together. You should not deploy SQL2 and AppServer2 in a single availability set. SQL2 and AppServer2 provide different services, so they should not be grouped together.

- You should not deploy all servers in a single availability set. This will not provide the fault tolerance needed, as all machines would be rebooted (or suffer a hardware failure) together.

**QUESTION 40**

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 1024MB 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.

- You should not create a virtual machine template by using a prepared instance of SQL Server 2016. A virtual machine template that contains a prepared instance of SQL Server could be modified during the installation process. For example, some features could be added or removed during the installation process. The prepared instance contains a preconfigured version of SQL Server without network and user configurations.

- You should not create an installation configuration file and perform unattended installations of SQL Server 2016. This could be a valid option, but you would need to prepare 10 installations, and this would require a lot of effort. There are also some other administrative actions that must be performed that could also increase the amount of effort required, such as user database configuration.

- You should not provision 10 Azure virtual machines that each contain SQL Server 2016, installed by using the default settings. A virtual machine created by using the Azure library provides you with a standard installation, and the requirements call for a custom installation. The instance will need some specific parameters for your custom installation. Therefore, you would have to customize each Azure virtual machine.

#### **QUESTION 41**

You have a server named Server1 that is hosted in an Azure virtual machine. Server1 contains the following:

- One instance of SQL Server 2016 Enterprise.
- 10 databases.
- 500 stored procedures.

You have a database named Database1 that is hosted on Server1. Database1 contains 100 queries that are executed dynamically from web applications. You plan to remove data from the procedure cache on Database1. You have the following requirements:

- Changes to Database1 must not affect other databases that are hosted on Server1.
- Changes to Database1 must not affect the performance of queries that are stored in other databases.
- The solution must minimize administrative effort.

You need to remove the data from the procedure cache as quickly as possible. What should you

do?

- A. Run DBCC FREEPROCCACHE.
- B. Run ALTER DATABASE SCOPED CONFIGURATION CLEAR PROCEDURE CACHE in the context of Database1.
- C. Run DBCC DROPCLEANBUFFERS.
- D. Write a script that iterates through each stored procedure definition and add WITH RECOMPILE to the definition.

**Answer: B**

**Explanation:**

- You should run ALTER DATABASE SCOPED CONFIGURATION CLEAR PROCEDURE CACHE in the context of Database1. This statement lets you change the settings of a database without affecting other databases that are installed on the instance of SQL Server 2016.

- You should not run DBCC FREEPROCCACHE. DBCC FREEPROCCACHE would clean the entire plan cache and would affect all databases. It is possible to remove a single plan from the cache by using the plan\_handle argument to DBCC FREEPROCCACHE, but you would have to identify all plans that are related to Database1, which requires a lot more administrative effort.

- You should not run DBCC DROPCLEANBUFFERS. DBCC DROPCLEANBUFFERS will remove the clean pages from the buffer cache. Columnstore pages are removed from the columnstore cache.

- You should not write a script that will iterate through each stored procedure definition and add WITH RECOMPILE to the definition. Each time the procedure is called, it will be recompiled, and this might degrade the server's performance. This approach would require additional administrative effort to produce the script. In addition, some calls are made from the web application and the script would not have any control over these calls. In earlier versions of SQL Server, prior to the availability of the ALTER DATABASE statement, this option would have been the way to avoid affecting other databases.

#### **QUESTION 42**

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.

What should you do? 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 ChangeDate. Update rows as the order status changes.
- B. Create a new table named Sales.OrderStatus that contains three columns named OrderID, StatusDate, and Status.



Insert new rows into the table as the order status changes.

- C. Implement change data capture on the Sales.Orders table.
- D. 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: B**

**QUESTION 43**

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.

**QUESTION 44**

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**

**QUESTION 45**

You have a SQL Server 2016 database named DB1. You plan to import a large number of records from a SQL Azure database to DB1. You need to recommend a solution to minimize the amount of space used in the transaction log during the import operation. What should you include in the recommendation?

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

E. A new file group

**Answer: A**

**Explanation:**

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

**Note:**

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

**QUESTION 46**

.....

**Get Complete Version Exam 70-765 Dumps with VCE and PDF Here**



<https://www.passleader.com/70-765.html>