

## 70-765 Dumps

### Provisioning SQL Databases (beta)

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



**NEW QUESTION 1**

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

You have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in a development environment. Each VM has a dedicated disk for backups.

You need to backup a database to the local disk on a VM. The backup must be replicated to another region.

Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 diskstorage
- C. Premium P30 disk storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage
- F. Standard zone redundant blob storage
- G. Standard locally redundant blob storage
- H. Standard geo-redundant blob storage

**Answer:** E

**Explanation:** Note: SQL Database automatically creates a database backups and uses Azure read- access geo-redundant storage (RA-GRS) to provide geo-redundancy. These backups are created automatically and at no additional charge. You don't need to do anything to make them happen. Database backups are an essential part of any business continuity and disaster recovery strategy because they protect your data from accidental corruption or deletion.

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

**NEW QUESTION 2**

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

You have deployed a GS-series virtual machine (VM) in Microsoft Azure. You plan to deploy Microsoft SQL Server.

You need to deploy a 30 megabyte (MB) database that requires 100 IOPS to be guaranteed while minimizing costs.

Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 disk storage
- C. Premium P30 disk storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage
- F. Standard zone redundant blob storage
- G. Standard locally redundant blob storage
- H. Standard geo-redundant blob storage

**Answer:** A

**Explanation:** Premium Storage Disks Limits

When you provision a disk against a Premium Storage account, how much input/output operations per second (IOPS) and throughput (bandwidth) it can get depends on the size of the disk. Currently, there are three types of Premium Storage disks: P10, P20, and P30. Each one has specific limits for IOPS and throughput as specified in the following table:

Premium Storage Disk Type	P10	P20	P30
Disk Size	128 GiB	512 GiB	1024 GiB (1 TB)
IOPS per disk	500	2300	5000
Throughput per disk	100 MB per second	150 MB per second	200 MB per second

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

**NEW QUESTION 3**

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

You have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in a development environment.

You need to provide storage to the environment that minimizes costs. Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 disk storage
- C. Premium P30 disk storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage
- F. Standard zone redundant blob storage
- G. Standard locally redundant blob storage
- H. Standard geo-redundant blob storage

**Answer:** D

#### NEW QUESTION 4

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

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

You need to create the templates.

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

Does the solution meet the goal?

A. Yes

B. No

**Answer: A**

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

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

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

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

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

#### NEW QUESTION 5

HOTSPOT

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

You gather the following metrics:

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

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

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

NOTE: Each correct selection is worth one point.

#### Answer Area

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

**Answer:**

**Explanation:** Data drive: Premium Storage Transaction log drive: Standard Storage TempDB drive: Premium Storage

Note: A standard disk is expected to handle 500 IOPS or 60MB/s. A P10 Premium disk is expected to handle 500 IOPS.

A P20 Premium disk is expected to handle 2300 IOPS. A P30 Premium disk is expected to handle 5000 IOPS.

VM size: A3

Max data disk throughput is 8x500 IOPS

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

### NEW QUESTION 6

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>

### NEW QUESTION 7

You plan to migrate a database To Microsoft Azure SQL Database. The database requires 500 gigabytes (GB) of storage.

The database must support 50 concurrent logins. You must minimize the cost associated with hosting the database.

You need to create the database. Which pricing tier should you use?

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

**Answer: D**

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

Note:

#### Premium service tier

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

### NEW QUESTION 8

#### DRAG DROP

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

**Answer:**

**Explanation: ;**

Box 1: ProductionNetwork

The virtual network is named ProductionNetwork.

Box 2: None /Not Assigned

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

Box 3: ProductionServer2

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

#### NEW QUESTION 9

##### HOTSPOT

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:**

**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.

Topic 2, Manage databases and instances

#### NEW QUESTION 10

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 theWindows 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>

#### NEW QUESTION 10

##### HOTSPOT

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

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

Answer Area

<div><div>▼</div><div>CREATE USER</div><div>ALTER SERVER ROLE</div><div>CREATE LOGIN</div></div>	Admin2 WITH password = 'Pa\$\$w0rd';		
<div><div>▼</div><div>CREATE USER</div><div>ALTER SERVER ROLE</div><div>CREATE LOGIN</div></div>	Admin2User FROM	<div><div>▼</div><div>WINDOWS</div><div>EXTERNAL PROVIDER</div><div>LOGIN</div></div>	Admin2
ALTER ROLE '	<div><div>▼</div><div>loginmanager</div><div>dbmanager</div><div>bd_ddladmin</div></div>		

**Answer:**

**Explanation:** Step 1: CREATE LOGIN

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

Step 2, CREATE USER Step 3: LOGIN

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

CREATE USER readonlyuser FROM LOGIN readonlylogin; Step 4: loginmanager

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

References:

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

## NEW QUESTION 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:** 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>

## NEW QUESTION 17

DRAG DROP

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:

#### Explanation: Step 1:

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

#### Step 2: CREATE USER ... FROM EXTERNAL PROVIDER

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

CREATE USER <Azure\_AD\_principal\_name> FROM EXTERNAL PROVIDER;

#### Step 3:

Grant the proper reading permissions.

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

### NEW QUESTION 19

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.

### NEW 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.

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)

#### NEW QUESTION 24

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

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

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

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

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

**Answer:** B

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

#### NEW QUESTION 29

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>

Topic 4, automobile parts Case Study 1Background

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

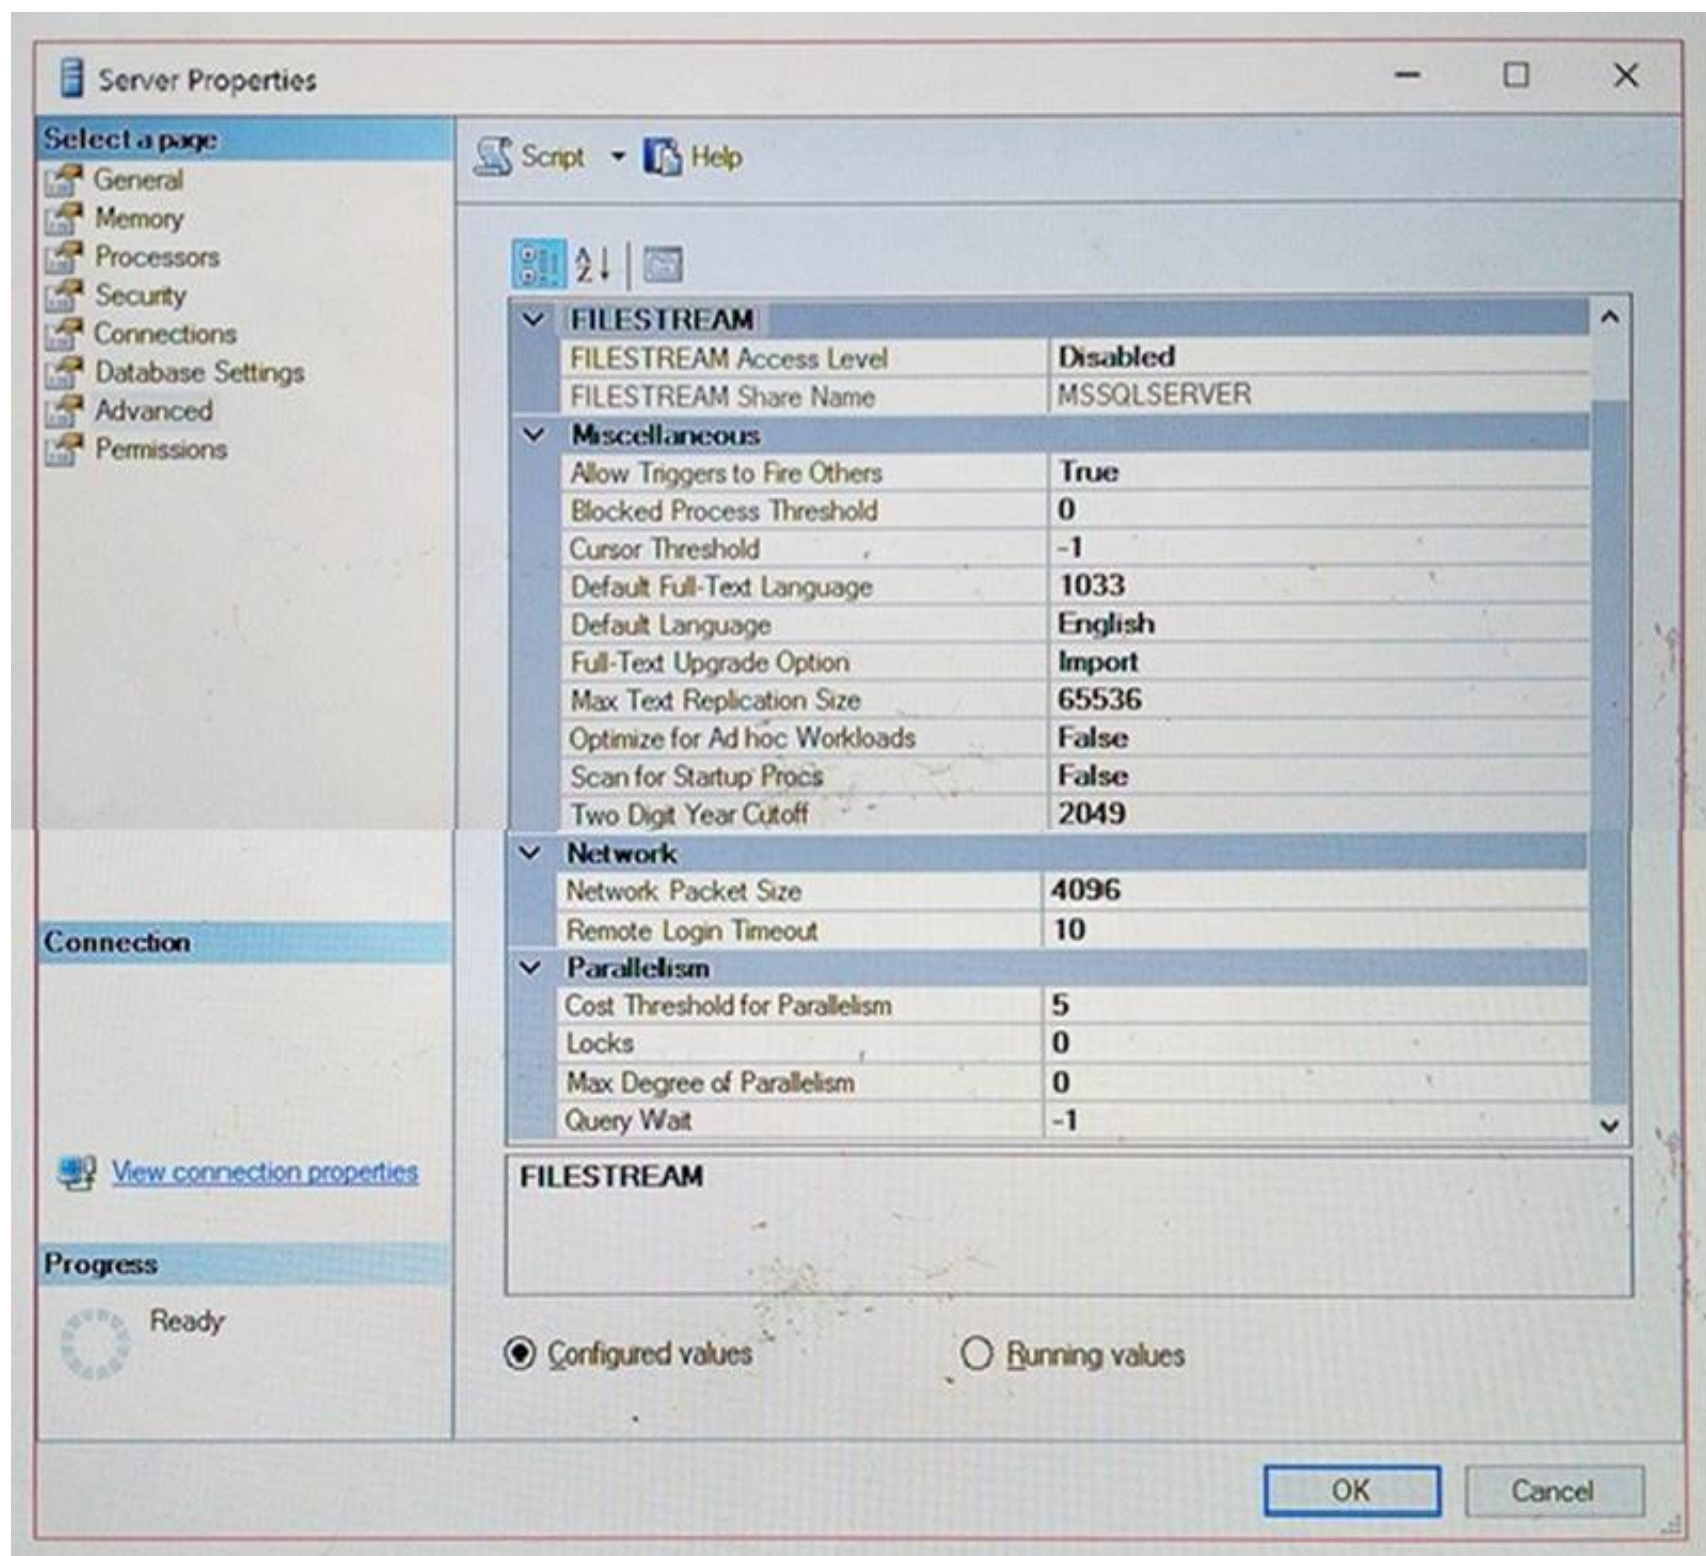
The environment includes the following servers: SRV1 and SRV2. SRV1 has 16 logical cores and hosts a SQL Server instance that supports a mission-critical application. The application has approximately 30,000 concurrent users and relies heavily on the use of temporary tables.

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.

Exhibit



#### NEW QUESTION 34

##### HOTSPOT

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:**

**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.



**NEW QUESTION 36**

**HOTSPOT**

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. NOTE: 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:**

**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.

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

Topic 6, SQL Server ReportingBackground

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 Database Name:

DB1

Notes:

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.

Database Name: DB2

Notes:

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.

Database Name: Reporting

Notes:

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.

**NEW QUESTION 39**

**HOTSPOT**

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

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

## Answer Area

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

### Answer:

#### Explanation: DB1: DELAYED\_DURABILITY=FORCED

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

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

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

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

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

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

DB2: ALLOW\_SNAPSHOT\_ISOLATION ON and READ\_COMMITTED\_SNAPSHOT ON

Snapshot isolation enhances concurrency for OLTP applications.

Snapshot isolation must be enabled by setting the ALLOW\_SNAPSHOT\_ISOLATION ON database option before it is used in transactions.

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

```
ALTER DATABASE MyDatabase
```

```
SET ALLOW_SNAPSHOT_ISOLATION ON
```

```
ALTER DATABASE MyDatabase
```

```
SET READ_COMMITTED_SNAPSHOT ON
```

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

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

References:

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

### NEW QUESTION 40

You administer a Microsoft SQL Server 2014 database.

You need to ensure that the size of the transaction log file does not exceed 2 GB. What should you do?

- A. Execute sp\_configure 'max log size', 2G.
- B. use the ALTER DATABASE...SET LOGFILE command along with the maxsize parameter.
- C. In SQL Server Management Studio, right-click the instance and select Database Setting
- D. Set the maximum size of the file for the transaction log.
- E. in SQL Server Management Studio, right-click the database, select Properties, and then click Files. Open the Transaction log Autogrowth window and set the maximum size of the file.

### Answer: B

**Explanation:** You can use the ALTER DATABASE (Transact-SQL) statement to manage the growth of a transaction log file

To control the maximum the size of a log file in KB, MB, GB, and TB units or to set growth to UNLIMITED, use the MAXSIZE option. However, there is no SET LOGFILE subcommand.

References: [https://technet.microsoft.com/en-us/library/ms365418\(v=sql.110\).aspx#ControlGrowth](https://technet.microsoft.com/en-us/library/ms365418(v=sql.110).aspx#ControlGrowth)



#### NEW QUESTION 41

You plan to migrate a Microsoft SQL server instance between physical servers.

You must migrate the metadata associated with the database instance.

You need to ensure that the new instance retains the existing jobs and alerts. Solutions: You restore the msdb database.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:** The msdb database is used by SQL Server Agent for scheduling alerts and jobs and by other features such as SQL Server Management Studio, Service Broker and Database Mail.

References:

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

#### NEW QUESTION 42

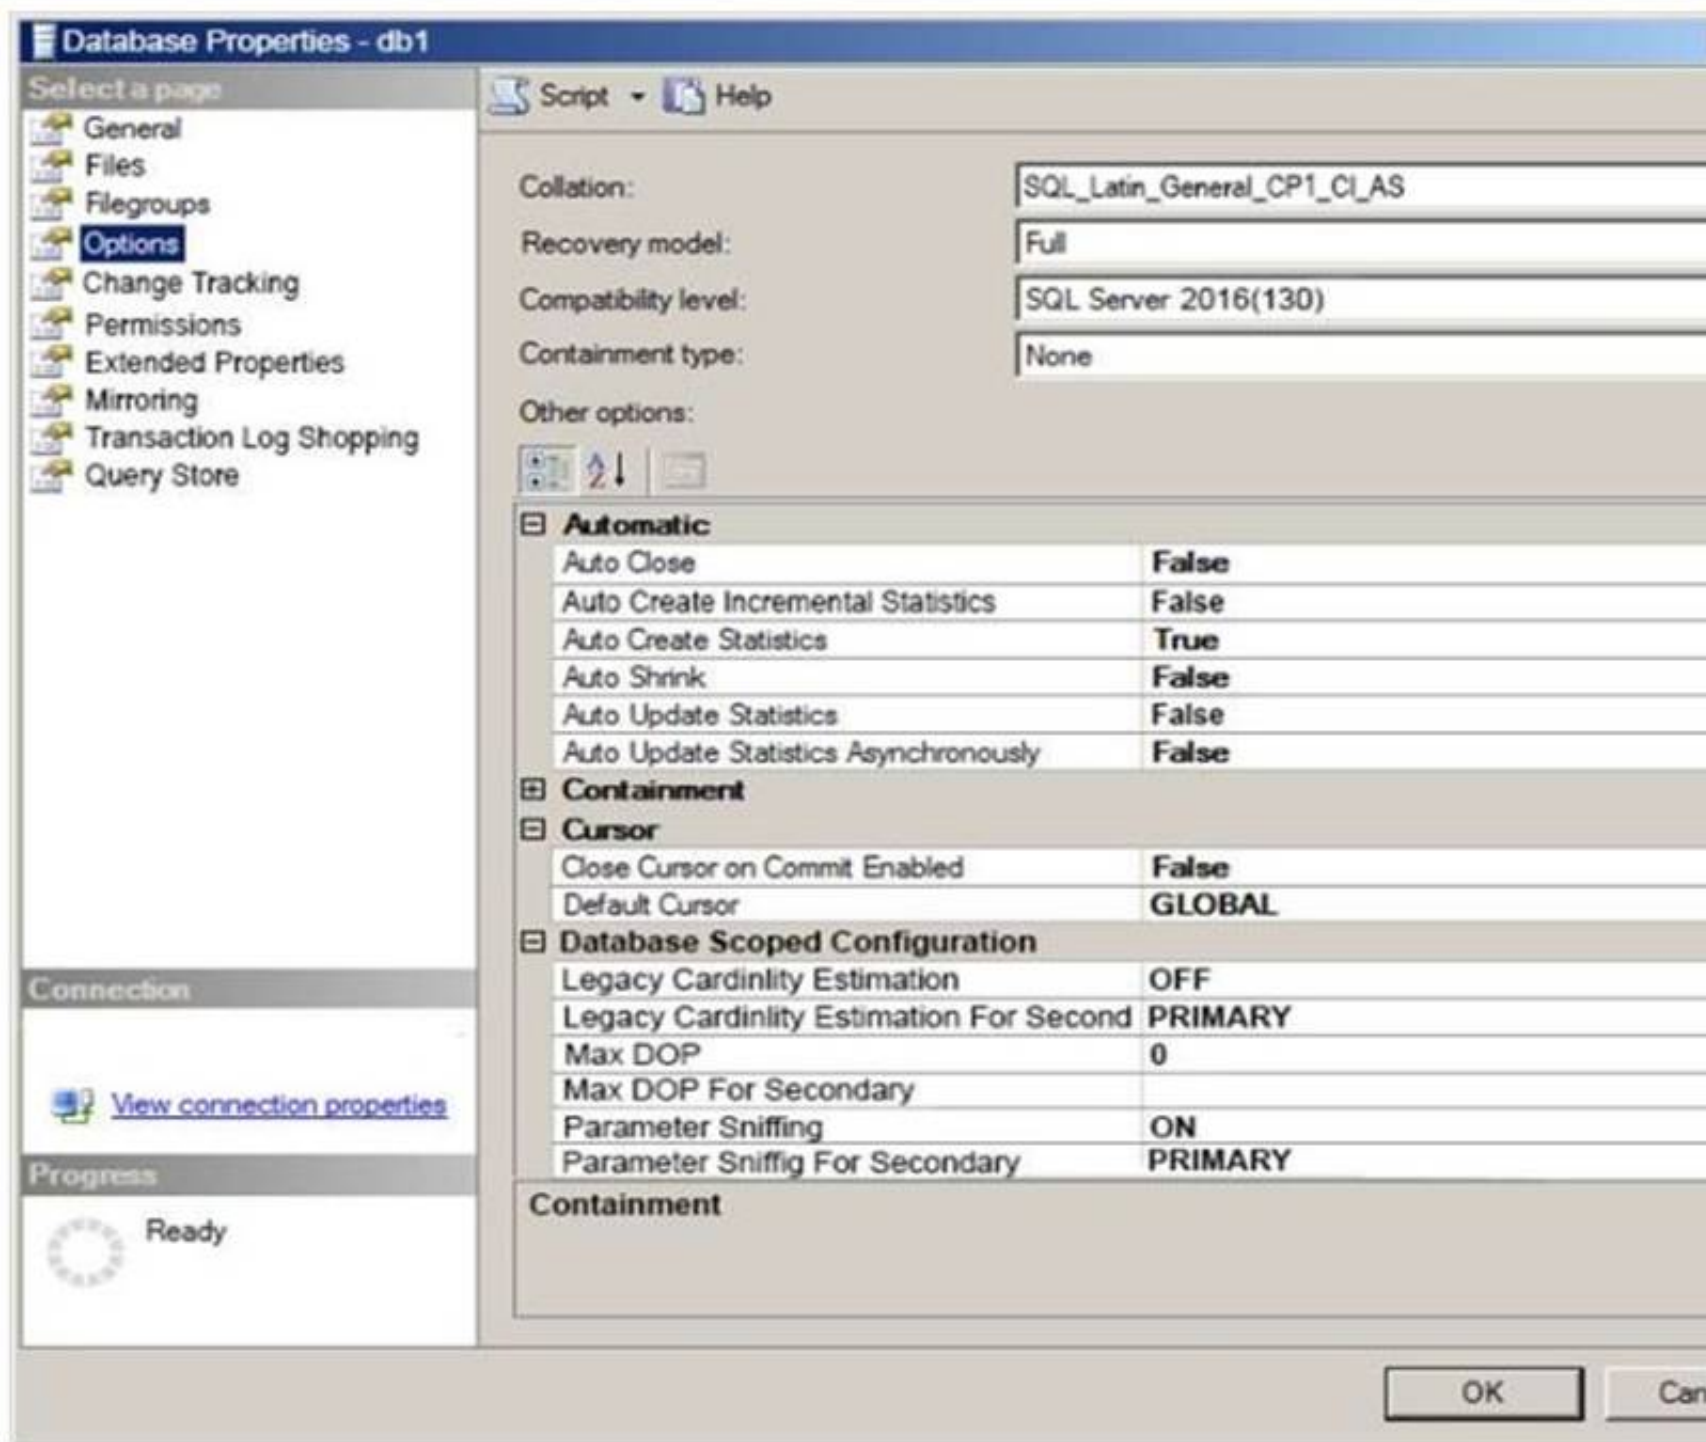
You have Microsoft SQL Server on a Microsoft Azure virtual machine. The virtual machine has a database named DB1. DB1 contains a table named Table1 that has 4 billion rows.

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

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

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

Hot Area:



**Answer:**

**Explanation:** When you set the AUTO\_CREATE\_STATISTICS option on, the Query Optimizer creates statistics on individual columns used in a predicate, if these statistics are not already available. These statistics are necessary to generate the query plan.

References:

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

#### NEW QUESTION 44

You administer a Microsoft SQL Server 2014 failover cluster that contains two nodes named Node A and Node B. A single instance of SQL Server is installed on the cluster.

An additional node named Node C has been added to the existing cluster.  
You need to ensure that the SQL Server instance can use all nodes of the cluster. What should you do?

- A. Run the New SQL Server stand-alone installation Wizard on Node C.
- B. Run the Add Node to SQL Server Failover Cluster Wizard on Node C.
- C. Use Node B to install SQL Server on Node C.
- D. Use Node A to install SQL Server on Node C.

**Answer: B**

**Explanation:** To add a node to an existing SQL Server failover cluster, you must run SQL Server Setup on the node that is to be added to the SQL Server failover cluster instance. Do not run Setup on the active node.

The Installation Wizard will launch the SQL Server Installation Center. To add a node to an existing failover cluster instance, click Installation in the left-hand pane. Then, select Add node to a SQL Server failover cluster.

References:

<http://technet.microsoft.com/en-us/library/ms191545.aspx>

#### NEW QUESTION 49

Settings Value VM size D3

Storage Location Drive E Storage type Standard Tempdb location Drive C

The workload on this instance has of the tempdb load.

You need to maximize the performance of the tempdb database.

Solution: You use a GS- Series VM and store the tempdb database on attached Premium storage. Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:** For VMs that support Premium Storage (DS-series, DSv2-series, and GS-series), we recommend storing TempDB on a disk that supports Premium Storage with read caching enabled. There is one exception to this recommendation; if your TempDB usage is write-intensive, you can achieve higher performance by storing TempDB on the local D drive, which is also SSD-based on these machine sizes.

References:

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

#### NEW QUESTION 53

You administer a Microsoft SQL Server 2014 database.

You have a SQL Server Agent job instance that runs using the service account. You have a job step within the job that requires elevated privileges.

You need to ensure that the job step can run using a different user account. What should you use?

- A. a schedule
- B. an alert
- C. an operator
- D. a proxy

**Answer: D**

**Explanation:** A SQL Server Agent proxy defines the security context for a job step. A proxy provides SQL Server Agent with access to the security credentials for a Microsoft Windows user. Each proxy can be associated with one or more subsystems. A job step that uses the proxy can access the specified subsystems by using the security context of the Windows user. Before SQL Server Agent runs a job step that uses a proxy, SQL Server Agent impersonates the credentials defined in the proxy, and then runs the job step by using that security context.

References:[https://technet.microsoft.com/en-us/library/ms189064\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms189064(v=sql.105).aspx)

#### NEW QUESTION 56

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

You need to ensure that the cluster meets the SLA.

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

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

**Answer: B**

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

References:

#### NEW QUESTION 57

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

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

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

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

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

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

- A. Yes
- B. No

**Answer: B**

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

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

#### NEW QUESTION 59

You administer a Microsoft SQL Server 2014 database.

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

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

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

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

**Answer: D**

#### NEW QUESTION 61

You administer a Microsoft SQL Server 2014 environment. One of the SQL Server 2014 instances contains a database named Sales.

You plan to migrate Sales to Windows Azure SQL Database. To do so, you need to implement a contained database.

What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Set database containment to AZURE.
- B. Enable server property contained database authentication.
- C. Disable server property cross db ownership chaining.
- D. Set database containment to PARTIAL.
- E. Disable server property contained database authentication.
- F. database containment to FULL.

**Answer: BD**

**Explanation:** A contained database is a database that is isolated from other databases and from the instance of SQL Server that hosts the database.

B: In the contained database user model, the login in the master database is not present. Instead, the authentication process occurs at the user database, and the database user in the user database does not have an associated login in the master database.

SQL Database and SQL Data Warehouse support Azure Active Directory identities as contained database users.

D: The contained database feature is currently available only in a partially contained state. A partially contained database is a contained database that allows the use of uncontained features.

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

#### NEW QUESTION 66

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

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

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

- A. Yes
- B. No

**Answer: B**

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

Note: Use the Always Encrypted Wizard to help protect sensitive data stored in a SQL Server database. Always Encrypted allows clients to encrypt sensitive data

inside client applications and never reveal the encryption keys to SQL Server.

References:

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

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

#### NEW QUESTION 71

You administer a SQL Server 2014 database instance.

You need to configure the SQL Server Database Engine service on a failover cluster. Which user account should you use?

- A. A domain user
- B. The BUILTIN\SYSTEM account
- C. A local user with Run as Service permissions
- D. The SQLBrowser account

**Answer:** A

**Explanation:** References:

<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/install/create-a-new-sql-server-failover-cluster-s>

#### NEW QUESTION 73

You administer a Microsoft SQL Server 2014 database named Contoso on a server named Server01. You need to prevent users from disabling server audits in Server01.

What should you create?

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

**Answer:** B

**Explanation:** Writing to the Windows Security log requires the SQL Server service account to be added to the Generate security audits policy. By default, the Local System, Local Service, and NetworkService are part of this policy. This setting can be configured by using the security policy snap-in (secpol.msc). Additionally, the Audit object access security policy must be enabled for both Success and Failure.

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

#### NEW QUESTION 75

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

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

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

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

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

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

**Answer:** D

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

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

#### NEW QUESTION 79

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

From SQL Server Management Studio (SSMS), a junior database administrator accidentally deletes several rows from the Azure SQL database and breaks the connection to Azure.

You need to resume Stretch Database operations.

Which two stored procedures should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. sys.sp\_rda\_reconcile\_batch
- B. sys.sp\_rda\_reconcile\_indexes
- C. sys.sp\_rda\_reauthorize\_db
- D. sys.sp\_rda\_reconcile\_columns
- E. sys.sp\_rda\_set\_rpo\_duration

**Answer:** CD

**Explanation:** sys.sp\_rda\_reauthorize\_db restores the authenticated connection between a local database enabled for Stretch and the remote database.



If you have accidentally deleted columns from the remote table, run sp\_rda\_reconcile\_columns to add columns to the remote table that exist in the Stretch-enabled SQL Server table but not in the remote table.

#### NEW QUESTION 81

You administer a Microsoft SQL Server 2014 Enterprise Edition server that uses 64 cores.

You discover performance issues when large amounts of data are written to tables under heavy system load. You need to limit the number of cores that handle I/O.

What should you configure?

- A. Processor affinity
- B. Lightweight pooling
- C. Max worker threads
- D. I/O affinity

**Answer:** D

**Explanation:** The affinity Input-Output (I/O) mask Server Configuration Option.

To carry out multitasking, Microsoft Windows 2000 and Windows Server 2003 sometimes move process threads among different processors. Although efficient from an operating system point of view, this activity can reduce Microsoft SQL Server performance under heavy system loads, as each processor cache is repeatedly reloaded with data. Assigning processors to specific threads can improve performance under these conditions by eliminating processor reloads; such an association between a thread and a processor is called processor affinity.


References:

<http://msdn.microsoft.com/en-us/library/ms189629.aspx>

#### NEW QUESTION 86

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

(Click the Exhibit button.)



You write the following Transact-SQL query:

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

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

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

**Answer: D**

**Explanation:** New statistics would be useful.

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

References:

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

#### NEW QUESTION 88

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

Drive L fails and is replaced.

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

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

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

**Answer:**

**Explanation:** ALTER DATABASE '<your\_database>' SET SINGLE\_USER WITH ROLLBACK IMMEDIATE GO

DBCC CHECKDB ('<your\_database>', REPAIR\_REBUILD) GO

ALTER DATABASE '<your\_database>' SET MULTI\_USER GO

#### NEW QUESTION 89

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

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

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

You need to prevent the PAGELATCH\_IO waits from occurring.

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

- A. Yes
- B. No

**Answer: B**

**Explanation:** Instead you can add more files to the database.

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

### NEW QUESTION 92

You have been hired as a Database Consultant by ABC.com to design a SQL Server 2014 database solution. You are tasked with designing a scale-out and high-availability SQL Server 2014 Online Transaction

Processing (OLTP) database solution that will maintain copies of data across two server instances.

Your solution must provide scale-out of read operations by distributing the reads from clients across two SQL Server 2014 nodes. The data in both SQL Server nodes needs to be indexed.

What should you include in your solution?

- A. You should include a primary database with scheduled log shipping to the secondary database configured.
- B. You should include two servers configured in an Active-Passive SQL Server 2014 Cluster.
- C. You should include a primary SQL Server 2014 database that uses transactional replication to replicate data to the secondary database.
- D. You should include two servers in an Asynchronous-Commit Availability Mode Availability Group.
- E. You should include two servers in a Synchronous-Commit Availability Mode Availability Group.

**Answer:** C

**Explanation:** Peer-to-peer replication provides a scale-out and high-availability solution by maintaining copies of data across multiple server instances, also referred to as nodes.

Built on the foundation of transactional replication, peer-to-peer replication propagates transactionally consistent changes in near real-time. This enables applications that require scale-out of read operations to distribute the reads from clients across multiple nodes. Because data is maintained across the nodes in near real-time, peer-to-peer replication provides data redundancy, which increases the availability of data.

References:[https://msdn.microsoft.com/en-us/library/ms151196\(v=sql.110\)](https://msdn.microsoft.com/en-us/library/ms151196(v=sql.110))

### NEW QUESTION 95

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

Maximum database size of 500 GB

A point-in-time-restore of 35 days

Maximum database transaction units (DTUs) of 500

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

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

**Answer:** B

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

### NEW QUESTION 100

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

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

#### Answer area

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

**Answer:**

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

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



Copy and paste the following example into the query window and click Execute. ALTER SERVER ROLE diskadmin ADD [Domain\Juan] ;  
GO  
G: Finally we add a database user for the login we created.  
References: [https://technet.microsoft.com/en-us/library/aa337562\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/aa337562(v=sql.110).aspx)

#### NEW QUESTION 102

Your company has many Microsoft SQL Server instances hosted in a data center. You also manage five Microsoft Azure SQL Database instances that are hosted on a single server in Azure.  
You need to minimize costs associated with Azure resources while maintaining the current performance levels of each Azure SQL Database instance.  
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

Determine the maximum DTU usage of all Azure SQL databases.
Add databases to the pool.
Add a new pool to the Azure SQL server.
Create an Azure Availability group.
Create an Azure Blob store.
Create an Azure SQL server.

**Answer:**

**Explanation:** SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single Azure SQL Database server and share a set number of resources at a set price.  
References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool>

#### NEW QUESTION 107

You have Microsoft SQL Server on a Microsoft azure virtual machine that has 12 databases. All database files are in the same Azure Blob storage account.  
You need to receive an email notification if I/O operations to the database files exceed 800 MB/s for more than five minutes.  
Solution: You run the Get-Counter cmdlet and specify the -counter '\physicaldisk:disk Transfers/sec' parameter.  
Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

#### NEW QUESTION 111

You administer a Microsoft SQL Server 2014 database that contains a table named OrderDetail.  
You discover that the NCI\_OrderDetail\_CustomerID non-clustered index is fragmented. You need to reduce fragmentation. You need to achieve this goal without taking the index offline.  
Which Transact-SQL batch should you use?

- A. CREATE INDEX NCI\_OrderDetail\_CustomerID ON OrderDetail.CustomerID WITH DROP EXISTING
- B. ALTER INDEX NCI\_OrderDetail\_CustomerID ON OrderDetail.CustomerID REORGANIZE
- C. ALTER INDEX ALL ON OrderDetail REBUILD
- D. ALTER INDEX NCI\_OrderDetail\_CustomerID ON OrderDetail.CustomerID REBUILD

**Answer:** B

**Explanation:** REORGANIZE specifies to reorganize the index leaf level. The REORGANIZE operation is always performed online. This means long-term blocking table locks are not held and queries or updates to the underlying table can continue during the ALTER INDEX REORGANIZE transaction.  
References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-index-transact-sql>



#### NEW QUESTION 116

You have an on-premises database that runs several maintenance jobs. You move the database to a Microsoft Azure SQL database.

You need to ensure that the maintenance jobs on indexes continue to run after the move.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

#### Actions, Select from these

#### Answer Area, Place here

Create a runbook

1.

Create an Automation Account

2.

Configure a schedule

3.

Create a credential

4.

Publish a runbook

5.

**Answer:**

**Explanation:** General steps for a solution to automate the maintenance you Azure SQL DB statistics: References:

<https://blogs.msdn.microsoft.com/azuresqldb-support/2018/01/15/automating-azure-sql-db-index-and-statistic-m>

#### NEW QUESTION 121

You administer a Microsoft SQL Server 2014 database.

You configure Transparent Data Encryption (TDE) on the Orders database by using the following statements: CREATE MASTER KEY ENCRYPTION BY PASSWORD = 'MyPassword1!'

CREATE CERTIFICATE TDE\_Certificate WITH SUBJECT = 'TDE Certificate'; BACKUP CERTIFICATE TDE\_Certificate TO FILE = "d:\TDE\_Certificate.cer" WITH PRIVATE KEY (FILE = 'D:\TDE\_Certificate.key',

ENCRYPTION BY PASSWORD = 'MyPassword1!'); CREATE DATABASE ENCRYPTION KEY

WITH ALGORITHM = AES\_256

ENCRYPTION BY SERVER CERTIFICATE TDE\_Certificate;

ALTER DATABASE Orders SET ENCRYPTION ON;

You attempt to restore the Orders database and the restore fails. You copy the encryption file to the original location.

A hardware failure occurs and so a new server must be installed and configured.

After installing SQL Server to the new server, you restore the Orders database and copy the encryption files to their original location. However, you are unable to access the database.

You need to be able to restore the database.

Which Transact-SQL statement should you use before attempting the restore?

A. ALTER DATABASE Master SET ENCRYPTION OFF;

B. CREATE CERTIFICATE TDE\_Certificate FROM FILE = 'd:\TDE\_Certificate.cer' WITH PRIVATE KEY (FILE = 'D:\TDE\_Certificate.key', DECRYPTION BY PASSWORD = 'MyPassword1!');

C. CREATE CERTIFICATE TDE\_Certificate WITH SUBJECT = 'TDE Certificate'; USE Orders; CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES\_256 ENCRYPTION BY SERVER CERTIFICATE TDE\_Certificate;

D. CREATE CERTIFICATE TDE\_Certificate FROM FILE = 'd:\TDE\_Certificate.cer';

**Answer:** B

**Explanation:** The CREATE CERTIFICATE command adds a certificate to a database in SQL Server. Creating a certificate from a file

The following example creates a certificate in the database, loading the key pair from files. Code

Copy

USE AdventureWorks2012; CREATE CERTIFICATE Shipping11

FROM FILE = 'c:\Shipping\Certs\Shipping11.cer'

WITH PRIVATE KEY (FILE = 'c:\Shipping\Certs\Shipping11.pvk', DECRYPTION BY PASSWORD = 'sldkflk34et6gs%53#v00');

GO

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

#### NEW QUESTION 123

You have a Microsoft Azure SQL Database server named server1-contoso.database.windows.net in a resource group named RG1.

You need to create an elastic pool.

How should you complete the script? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

\$server = 

▼
Get-AzureRmSqlElasticPool
Get-AzureRmSqlServer
New-AzureRmSqlServer
Set-AzureRmSqlElasticPool
Set-AzureRmSqlServer

 -ServerName "server1-contoso" -ResourceGroupName "RG1"

\$server = 

▼
Get-AzureRmSqlElasticPool
Get-AzureRmSqlServer
New-AzureRmSqlServer
Set-AzureRmSqlElasticPool
Set-AzureRmSqlServer
New-AzureRmSqlElasticPool

 -ElasticPoolName "Pool1" -Edition Premium

**Answer:**

**Explanation:** Box 1: New-AzureRmSqlServer

Create an Azure SQL Database logical server using the New-AzureRmSqlServer command. A logical server contains a group of databases managed as a group. Example:

New-AzureRmSqlServer -ResourceGroupName \$resourcegroupname ` -ServerName \$servername ` -Location \$location `etc.

Box 2: New-AzureRmSqlElasticPool

The New-AzureRmSqlElasticPool cmdlet creates an elastic database pool for an Azure SQL Database. Example:

New-AzureRmSqlElasticPool -ResourceGroupName "ResourceGroup01" -ServerName "Server01" -ElasticPoolName "ElasticPool01" -Edition "Standard"

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-get-started-powershell> <https://docs.microsoft.com/en-us/powershell/module/azurerm.sql/new-azurermsqlelasticpool?view=azurermps-6>

#### NEW QUESTION 127

You have an on-premises database.

You plan to migrate the database to Microsoft SQL Server on a Microsoft Azure virtual machine.

You move the database files to Azure.

You need to attach the database files to the SQL Server instance on the virtual machine. The solution must ensure that you can run file snapshot backups.

How should you complete the statement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

#### Answer area

```
USE (master)
GO
CREATE DATABASE [Production_DB]
(
  (
    ▼ = N'https://proddbstorage=contoso.blob.core.windows.net/datafiles/proddb.mdf'
    DISK
    NAME
    FILEGROUP
    FILENAME
  )
  (
    ▼
    ON PRIMARY;
    ON COLLATE;
  )
)
GO
CREATE
```

**Answer:**

**Explanation:** References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-sql-server-transact-sql>

#### NEW QUESTION 130

A company runs Microsoft SQL Server 2017 in an on-premises environment. The databases are memory-optimized.

An integrity check of a database has failed.

You need to ensure that the data is healthy and passes an integrity check. What should you do?

- A. Run the checktable Transact-SQL statement.
- B. Clear the buffer of the database.
- C. Restore from a verified backup.

D. Run the cleantable Transact-SQL statement.

**Answer:** C

**Explanation:** To verify the integrity of the on-disk checkpoint files, perform a backup of the MEMORY\_OPTIMIZED\_DATA filegroup.

#### NEW QUESTION 131

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

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

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

You need to reduce the WRITE\_LOG wait time. Solution: Add additional data files to DB1. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

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

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

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

#### NEW QUESTION 135

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

The financial database has the following characteristics:

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

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

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

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

On Wednesday at 10:00 hours, the development team requests you to refresh the database on a development server by using the most recent version.

You need to perform a full database backup that will be restored on the development server. Which backup option should you use?

- A. NORECOVERY
- B. FULL
- C. NO\_CHECKSUM
- D. CHECKSUM
- E. Differential
- F. BULK\_LOGGED
- G. STANDBY
- H. RESTART
- I. SKIP
- J. Transaction log
- K. DBO ONLY
- L. COPY\_ONLY
- M. SIMPLE
- N. CONTINUE AFTER ERROR

**Answer:** L

**Explanation:** COPY\_ONLY specifies that the backup is a copy-only backup, which does not affect the normal sequence of backups. A copy-only backup is created independently of your regularly scheduled, conventional backups. A copy-only backup does not affect your overall backup and restore procedures for the database.

References:

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

#### NEW QUESTION 138

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

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

Tools, Select from these.	Answer Area
Distributed Replay	Tool to use on-premises: <Place here>
Performance Monitor	Tool to use in Azure: <Place here>
SQL Server Profiler	
SQL Server Extended Events	
SQL Server Data Tools (SSDT)	

**Answer:**

**Explanation:** Tools, Select from these.

**Answer Area**

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

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

#### NEW QUESTION 141

You administer a Microsoft SQL Server 2014 database that includes a table named Application.Events. Application.Events contains millions of records about user activity in an application.

Records in Application.Events that are more than 90 days old are purged nightly. When records are purged, table locks are causing contention with inserts.

You need to be able to modify Application.Events without requiring any changes to the applications that utilize Application.Events.

Which type of solution should you use?

- A. Partitioned tables
- B. Online index rebuild
- C. Change data capture
- D. Change tracking

**Answer:** A

**Explanation:** Partitioning large tables or indexes can have manageability and performance benefits including:

You can perform maintenance operations on one or more partitions more quickly. The operations are more efficient because they target only these data subsets, instead of the whole table.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/partitions/partitioned-tables-and-indexes>

#### NEW QUESTION 142

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

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

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

- A. Yes
- B. No

**Answer:** B

#### NEW QUESTION 143

You have a database named DB1. You discover that DB1 is corrupt.

You run DBCC CHECKDB and receive an error message within a few seconds. No pages are listed in the error message.

You need to repair the database corruption as quickly as possible. The solution must minimize data loss.

What should you do?

- A. Run DBCC CHECKDB ('db1', REPAIR\_ALLOW\_DATA\_LOSS).
- B. Run DBCC CHECKDB ('db1', REPAIR\_FAST).
- C. Delete the transaction logs and restart the Microsoft SQL Server instance.
- D. Run DBCC CHECKDB ('db1', REPAIR\_REBUILD).
- E. Restore the database from a backup.

**Answer:** C

**Explanation:**

REPAIR\_REBUILD

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

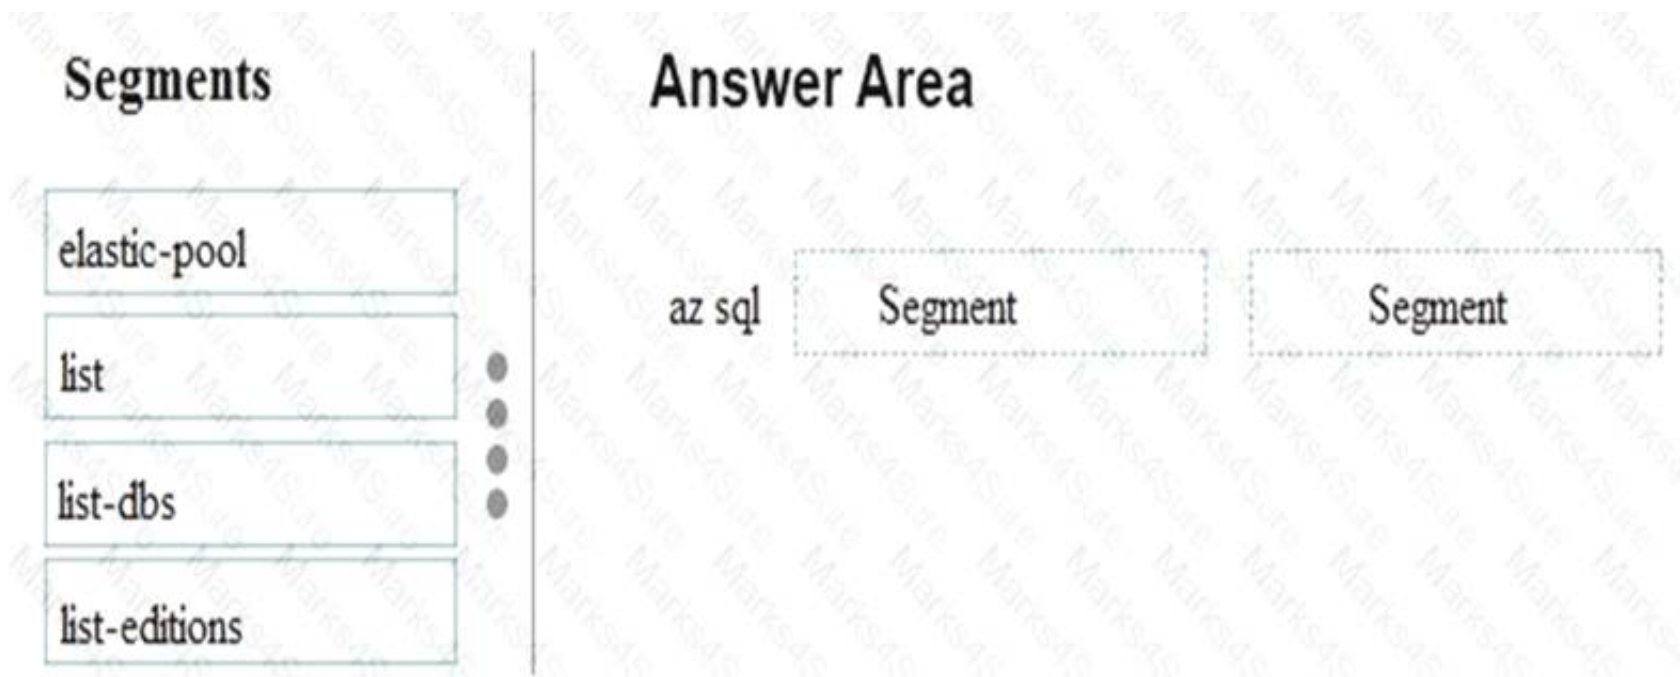
#### NEW QUESTION 144

Your company has several Microsoft Azure SQL Database instances used within an elastic pool. You need to obtain a list of databases in the pool.

How should you complete the commands? To answer, drag the appropriate segments to the correct targets. Each 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.

NOTE: Each correct selection is worth one point.





**Answer:**

**Explanation:** References:

<https://docs.microsoft.com/en-us/cli/azure/sql/elastic-pool?view=azure-cli-latest#az-sql-elastic-pool-list-dbs>

#### NEW QUESTION 148

You manage a Microsoft SQL Server environment with several databases.

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

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

References:

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

#### NEW QUESTION 149

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

The financial database has the following characteristics:

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

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

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

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

You need to ensure that your backup will continue if any invalid checksum is encountered. Which backup option should you use?

- A. STANDBY
- B. Differential
- C. FULL
- D. CHECKSUM
- E. BULK\_LOGGED
- F. CONTINUE\_AFTER\_ERROR
- G. SIMPLE
- H. DBO\_ONLY
- I. COPY\_ONLY
- J. SKIP
- K. RESTART
- L. Transaction log
- M. NO\_CHECKSUM
- N. NORECOVERY

**Answer:** F

**Explanation:** The CONTINUE\_AFTER\_ERROR option, of the Transact-SQL BACKUP command, instructs BACKUP to continue despite encountering errors such as invalid checksums or torn pages.

References:

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

**NEW QUESTION 150**

You have a database named DB1 that uses simple recovery mode.

Full backups of DB1 are taken daily and DB1 is checked for corruption before each backup. There was no corruption when the last backup was complete.

You run the sys.columns catalog view and discover corrupt pages.

You need to recover the database. The solution must minimize data loss. What should you do?

- A. Run RESTORE DATABASE WITH RECOVERY.
- B. Run RESTORE DATABASE WITH PAGE.
- C. Run DBCC CHECKDB and specify the REPAIR\_ALLOW\_DATA\_LOSS parameter.
- D. Run DBCC CHECKDB and specify the REPAIR\_REBUILD parameter.

**Answer:** B

**Explanation:** A page restore is intended for repairing isolated damaged pages. Restoring and recovering a few individual pages might be faster than a file restore, reducing the amount of data that is offline during a restore operation.

RESTORE DATABASE WITH PAGE

Restores individual pages. Page restore is available only under the full and bulk-logged recovery models. References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

**NEW QUESTION 153**

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

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

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

- A. Yes
- B. No

**Answer:** B

**Explanation:** We need the use of a Standard database pool.

**NEW QUESTION 154**

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

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

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

- A. sys.dm\_xtp\_transaction\_stats
- B. sys.dm\_exec\_input\_buffer
- C. sys.dm\_db\_index\_operational\_stats
- D. sys.dm\_db\_stats\_properties

**Answer:** C

**Explanation:** sys.dm\_db\_index\_operational\_stats dynamic management function provides us the current low-level I/O, locking, latching, and access method for each partition of the table. This information is really useful to troubleshoot SQL Server performance issues.

Reference:

<https://basitaalishan.com/2013/03/19/using-sys-dm-db-index-operational-stats-to-analyse-howindexes-are-utili>

**NEW QUESTION 156**

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:** Password protection does not encrypt the data.

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

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

**NEW QUESTION 161**

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

Three user databases must be preinstalled.

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

Trace flag 2371 must be turned at the instance level.

The solution must meet the following requirements:

The instance must be preconfigured.

No other database features are required in the future.

The solution must use the minimum administrative effort.  
You need to prepare the environments. What should you do?

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

**Answer:** D

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

#### NEW QUESTION 166

You administer a SQL Server 2014 server that contains a database named SalesDB. SalesDb contains a schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales. UserA is granted the Select permission on the Regions table. The Sales role is granted the Select permission on the Customers schema. You need to ensure that UserA is disallowed to select from any of the tables in the Customers schema. Which Transact-SQL statement should you use?

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

**Answer:** G

**Explanation:** Use SQL Data Warehouse or Parallel Data Warehouse GRANT and DENY statements to grant or deny a permission (such as UPDATE) on a securable (such as a database, table, view, etc.) to a security principal (a login, a database user, or a database role).  
References: [https://docs.microsoft.com/en-us/sql/t-sql/statements/permissions-grant-deny-revoke-azure-sql-data-](https://docs.microsoft.com/en-us/sql/t-sql/statements/permissions-grant-deny-revoke-azure-sql-data-warehouse)

#### NEW QUESTION 168

You administer two Microsoft SQL Server 2014 servers named ProdSrv1 and ProdSrv2. ProdSrv1 is configured as a Distributor. Both servers are configured to use the Windows NT Service virtual accounts for all SQL Services.

You are configuring snapshot replication from ProdSrv1 to ProdSrv2 by using ProdSrv2 as a pull subscriber. The distribution agent on ProdSrv2 regularly fails, displaying the following error message:

"Cannot access the file. Operating system error code 5 (Access is denied.)." You need to configure the distribution agent by granting only the minimum required access to all accounts.

What should you do?

- A. Configure the Subscriber to use the Local System account.
- B. Configure the SQL Server Agent service to run under the Local System account
- C. Configure the Subscriber to use the SQL Server Agent service account.
- D. Configure the SQL Server Agent service to run under a Windows domain account
- E. Configure the Subscriber to use the SQL Server Agent service account
- F. Grant FULL CONTROL access for the domain account to the ReplData share on ProdSrv1.
- G. Configure the Subscriber to use a Windows domain account
- H. Grant READ access for the domain account to the ReplData share on ProdSrv1.

**Answer:** D

**Explanation:** Confirm that distribution agent has read privileges, full control access is not required, to the folder in question.

References:

<http://stackoverflow.com/questions/14555262/cannot-bulk-load-operating-system-error-code-5-access-is-denied>

#### NEW QUESTION 169

You plan to deploy a Microsoft SQL Server database that will use FILESTREAM. The database will store 4 TB of FILESTREAM data on a single Windows partition.

You need to configure the hard disk that will support the FILESTREAM data. The solution must provide the fastest read and write access to the data.

How should you configure the disk? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.



## Answer area

File system:

	▼
FAT32	
FAT	
NTFS	

8.3 filename support:

	▼
Enabled	
Disabled	

Indexing:

	▼
Enabled	
Disabled	

**Answer:**

**Explanation:** File System: NTFS

8.3 filename support: Disabled Indexing: Disabled

NTFS is required.

Disable generation of 8.3 names on all NTFS volumes used for FILESTREAM data storage.

Check that search indexing is not enabled on FILESTREAM volumes, under the Volume Properties window, unchecking the “Allow files on this drive to have contents indexed in addition to file properties” box.

References:

<https://blogs.msdn.microsoft.com/blogdoezequiel/2011/02/11/best-practices-on-filestreamimplementations/>

### NEW QUESTION 170

You administer a SQL Server 2014 server that contains a database named SalesDb. SalesDb contains a schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales. UserA is granted the Select permission on the Regions table. The Sales role is granted the Select permission on the Customers schema.

You need to ensure that the following requirements are met: Which Transact-SQL statement should you use?

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

**Answer:** J

**Explanation:** Use REVOKE to remove the grant or deny of a permission.

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

### NEW QUESTION 171

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

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

Which stored procedure should you run?

- A. Sys\_sp\_testlinkedserver
- B. Sys\_sp\_rda\_test\_connection
- C. Sys\_sp\_rda\_reauthorized\_db
- D. Sp\_set\_firewall\_rule

**Answer:** B

**Explanation:** The Sys\_sp\_rda\_test\_connection cmdlet tests the connection from SQL Server to the remote Azure server and reports problems that may prevent data migration.

References:

[https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-rda-test-connection-tr](https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-rda-test-connection-troubleshooting)

#### NEW QUESTION 176

You are building the database platform for a multi-tenant application. The application will have one database per tenant and will have at least 30 tenants. Each tenant will have a separate resource group for billing purposes. The application will require at least 10 GB of clustered columnstore indexes for each database. You need to implement the database platform for the application. The solution must minimize costs. What should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

## Answer Area

Service tier:

	▼
Basic	
Standard	
Premium	
Premium RS	

Database implementation:

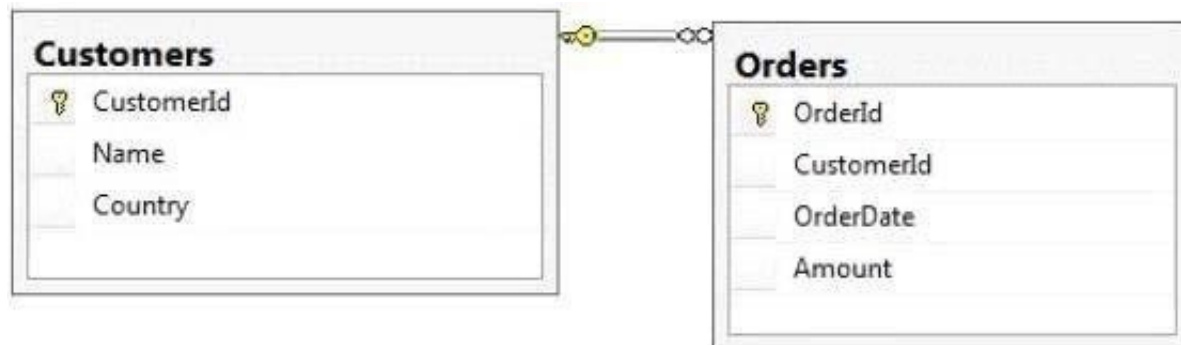
	▼
One individual Azure SQL database	
Thirty individual Azure SQL databases	

Answer:

**Explanation:** The Standard tier service allows for 1TB of data. Here 30 x 10 GB, 0.3 TB, is required.

#### NEW QUESTION 178

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



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

```

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

Which Transact-SQL query should you use?

- A. SELECT OrderId, OrderDate, Amount, Name, CountryFROM OrdersINNER JOIN CustomersON Orders.CustomerId = Customers.CustomerIdWHERE Customers.CustomerId = 1FOR XML RAW
- B. SELECT OrderId, OrderDate, Amount, Name, CountryFROM OrdersINNER JOIN CustomersON Orders.CustomerId = Customers.CustomerIdWHERE Customers.CustomerId = 1FOR XML RAW, ELEMENTS
- C. SELECT OrderId, OrderDate, Amount, Name, CountryFROM OrdersINNER JOIN CustomersON Orders.CustomerId = Customers.CustomerIdWHERE Customers.CustomerId = 1FOR XML AUTO
- D. SELECT OrderId, OrderDate, Amount, Name, CountryFROM OrdersINNER JOIN CustomersON Orders.CustomerId=Customers.CustomerIdWHERE Customers.CustomerId= 1FOR XML AUTO, ELEMENTS
- E. SELECT Name, Country, OrderId, OrderDate, AmountFROM OrdersINNER JOIN CustomersON Orders.CustomerId= Customers.CustomerIdWHERE Customers.CustomerId= 1FOR XML AUTO, ELEMENTS

F. SELECT Name, Country, OrderId, OrderDate, AmountFROM OrdersINNER JOIN CustomersON Orders.CustomerId= Customers.CustomerIdWHERE Customers.CustomerId= FOR XML AUTO, ELEMENTS

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

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

**Answer:** E

#### NEW QUESTION 182

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

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

Solution: You run the Add-AzureRmMetricAlertRule cmdlet and specify the –MetricName ‘Network Out’ parameter.

Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

#### NEW QUESTION 185

You use a contained database named ContosoDb within a domain.

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

Which type of user should you create?

- A. SQL user without login
- B. User mapped to an asymmetric key
- C. Domain user
- D. login mapped to a virtual account

**Answer:** C

**Explanation:** If the service must interact with network services, access domain resources like file shares or if it uses linked server connections to other computers running SQL Server, you might use a minimally-privileged domain account. Many server-to-server activities can be performed only with a domain user account.

References:<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-windows-servic>

#### NEW QUESTION 186

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

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

Evenly distribute read and write operations across the VHDs.

Minimize the read and write time.

Which storage configuration should you use?

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

**Answer:** D

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

Reference:

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

#### NEW QUESTION 187

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

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

Which statement should you execute?

- A)



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

B)

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

C)

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

D)

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

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

**Answer: B**

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

#### NEW QUESTION 188

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

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

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

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

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

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

- A. Yes
- B. No

**Answer:** B

**Explanation:** Each Storage Account handles up to 20,000 IOPS, and 500TB of data.  
References: <https://www.tech-coffee.net/understand-microsoft-azure-storage-for-virtual-machines/>

#### NEW QUESTION 193

You administer a Microsoft SQL Server 2014 instance that contains a financial database hosted on a storage area network (SAN). The financial database has the following characteristics:  
The database is continually modified by users during business hours from Monday through Friday between 09:00 hours and 17:00 hours. Five percent of the existing data is modified each day.  
The Finance department loads large CSV files into a number of tables each business day at 11:15 hours and 15:15 hours by using the BCP or BULK INSERT commands. Each data load adds 3 GB of data to the database.  
These data load operations must occur in the minimum amount of time.  
A full database backup is performed every Sunday at 10:00 hours. Backup operations will be performed every two hours (11:00, 13:00, 15:00, and 17:00) during business hours.  
You need to ensure that the backup size is as small as possible. Which backup should you perform every two hours?

- A. NORECOVERY
- B. FULL
- C. NO\_CHECKSUM
- D. CHECKSUM
- E. Differential
- F. BULK\_LOGGED
- G. STANDBY
- H. RESTART
- I. SKIP
- J. Transaction log
- K. DBO ONLY
- L. COPY\_ONLY
- M. SIMPLE
- N. CONTINUE AFTER ERROR

**Answer:** J

**Explanation:** Minimally, you must have created at least one full backup before you can create any log backups. After that, the transaction log can be backed up at any time unless the log is already being backed up.  
References: <https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/transaction-log-backups-sql-server>

#### NEW QUESTION 194

You plan to install Microsoft SQL Server 2014 for a web hosting company. The company plans to host multiple web sites, each supported by a SQL Server database. You need to select an edition of SQL Server that features backup compression of databases, basic data integration features, and low total cost of ownership. Which edition should you choose?

- A. Express Edition with Tools
- B. Standard Edition
- C. Web Edition
- D. Express Edition with Advanced Services

**Answer:** B

#### NEW QUESTION 198

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

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

**Answer:** A

**Explanation:** Compared to the full recovery model, which fully logs all transactions, the bulk-logged recovery model minimally logs bulk operations, although fully logging other transactions. The bulk-logged recovery model protects against media failure and, for bulk operations, provides the best performance and least log space usage.  
Note: The bulk-logged recovery model is a special-purpose recovery model that should be used only intermittently to improve the performance of certain large-scale bulk operations, such as bulk imports of large amounts of data.  
References: [https://technet.microsoft.com/en-us/library/ms190692\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms190692(v=sql.105).aspx)

#### NEW QUESTION 199

You administer a Microsoft SQL Server 2014.

A process that normally runs in less than 10 seconds has been running for more than an hour. You examine the application log and discover that the process is using session ID 60.

You need to find out whether the process is being blocked. Which Transact-SQL statement should you use?

- A. EXEC sp\_who 60
- B. SELECT \* FROM sys.dm\_exec\_sessions WHERE sessionid = 60
- C. EXEC sp\_helpdb 60
- D. DBCC INPUTBUFFER (60)

**Answer:** A

**Explanation:** sp\_who provides information about current users, sessions, and processes in an instance of the Microsoft SQL Server Database Engine. The information can be filtered to return only those processes that are not idle, that belong to a specific user, or that belong to a specific session.

Example: Displaying a specific process identified by a session ID EXEC sp\_who '10' --specifies the process\_id;

References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-who-transact->

#### NEW QUESTION 203

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

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

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

**Answer:** D

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

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

#### NEW QUESTION 207

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

You perform full backups of DB1 daily and transaction log backups hourly. Drive E fails and is replaced.

You need to recover DB1 and prevent any data loss.

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

Select and Place:

**Select and Place:**

#### Actions

Restore the tail-log backup.

Restore a full backup.

Perform a tail-log backup.

Restore the log backups.

Truncate the log of DB1.

Delete DB1.

#### Answer Area

1

2

3

4

**Answer:**

**Explanation:** Section: Deploy and migrate applications Step 1: Perform a tail-log backup.

A tail-log backup captures any log records that have not yet been backed up (the tail of the log) to prevent work loss and to keep the log chain intact. Before you can recover a SQL Server database to its latest point in time, you must back up the tail of its transaction log. The tail-log backup will be the last backup of interest in the recovery plan for the database.

Step 2: Restore a full backup.

Backups must be restored in the order in which they were created. Before you can restore a particular transaction log backup, you must first restore the following previous backups without rolling back uncommitted transactions, that is WITH NORECOVERY:

The full database backup and the last differential backup, if any, taken before the particular transaction log backup.

Step 3: Restore the log backups.

Log backups must be applied in the sequence in which they were created, without any gaps in the log chain. Step 4: Restore the tail-log backups.

Reference:

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



**NEW QUESTION 212**

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

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

What should you create?

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

**Answer:** D

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

References:

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

**NEW QUESTION 216**

You administer a Windows Azure SQL Database database named Inventory that contains a stored procedure named p\_AddInventory.

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

What should you do?

- A. Grant EXECUTE permission on p\_AddInventory to all user
- B. Grant VIEW DEFINITION to all users.
- C. Grant EXECUTE permission on p\_AddInventory to all user
- D. Add all users to the db\_datawriter role.
- E. Add all users to the db\_owner role.
- F. Grant EXECUTE permission on p\_AddInventory to all user
- G. Add all users to the db\_datareader role.

**Answer:** D

**Explanation:** Members of the db\_datareader fixed database role can run a SELECT statement against any table or view in the database.

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

**NEW QUESTION 221**

Database DB1 must use two CPU cores.

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

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

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

**Answer:**

**Explanation:** Query\_optimizer\_hotfixes DB1: MAXDOP

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

DB2: LEGACY\_CARDINALITY\_ESTIMATION

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

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

Example:

ALTER DATABASE SCOPED CONFIGURATION SET LEGACY\_CARDINALITY\_ESTIMATION = ON; GO

**NEW QUESTION 226**

Background

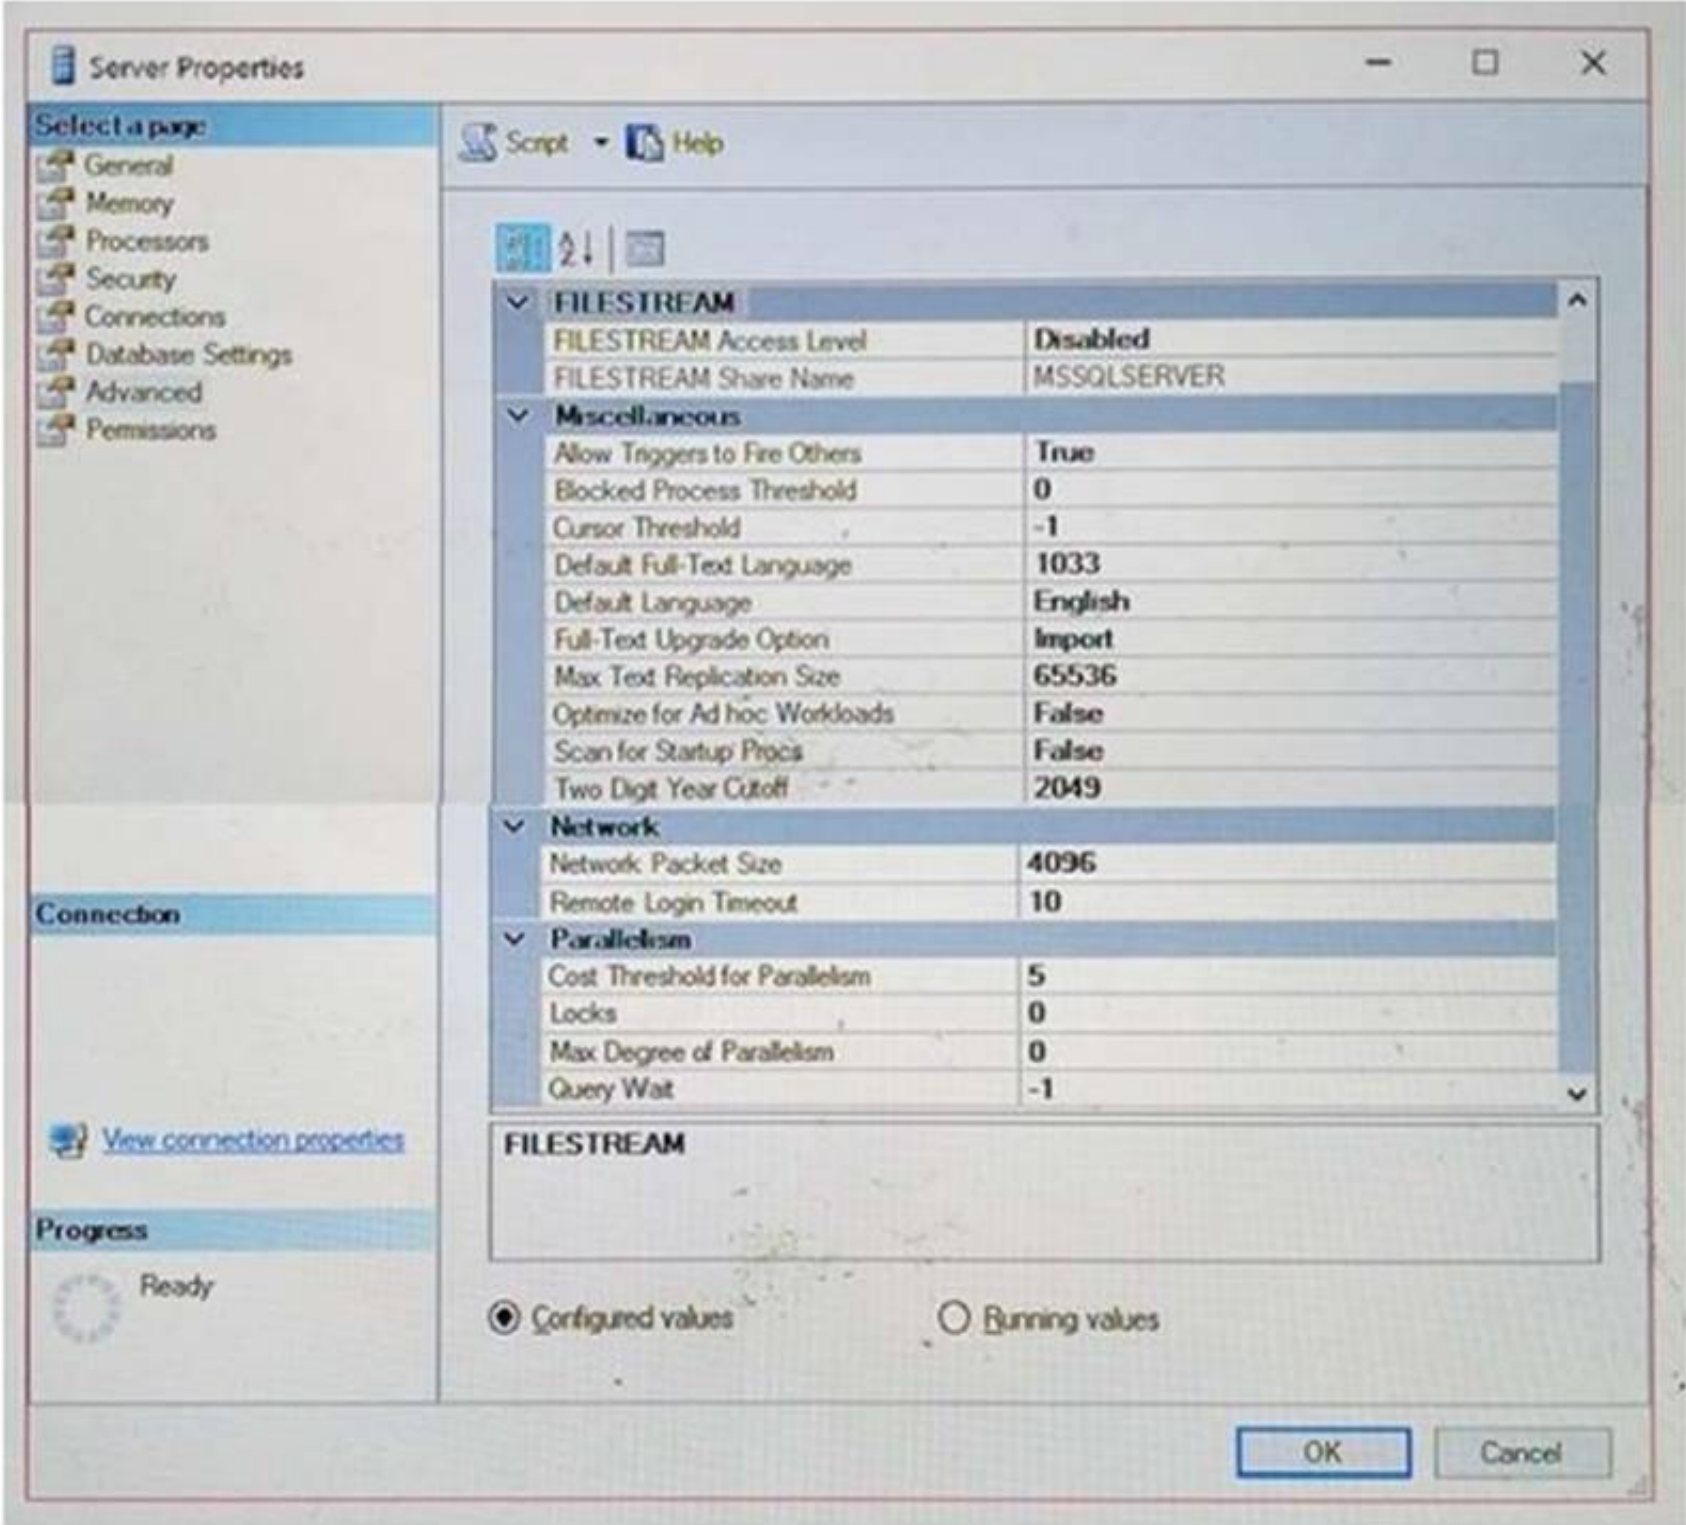
You manage the Microsoft SQL Server environment for a company that manufactures and sells automobile parts.

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

The environment also includes the following databases: DB1, DB2, and Reporting. The Reporting database is protected with Transparent Data Encryption (TDE).

You plan to migrate this database to a new server. You detach the database and copy it to the new server.

You are performing tuning on a SQL Server database instance. The application which uses the database was written using an object relationship mapping (ORM) tool which maps tables as objects within the application code. There are 30 stored procedures that are regularly used by the application. After reviewing the plan cache you have identified that a large number of simple queries are using parallelism, and that execution plans are not being kept in the plan cache for very long. You review the properties of the instance (Click the Exhibit button). Exhibit:



You need to restore the Reporting database to SRV2. What should you do? To answer, drag the appropriate options to the correct locations. Each option 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. Select and Place:

### Values

master encryption key on the master database
service master key
server certificate
Reporting database .mdf file
master key password

### Answer area

1. Copy the certificate and private key backups from the old server to the new server.
2. Create:
3. Restore:
4. Attach the Reporting database.

Answer:

**Explanation:** Step 2: Create: server certificate

Recreate the server certificate by using the original server certificate backup file.

Note: The password must be the same as the password that was used when the backup was created. Step 3: Restore: Reporting database .mdf file.

-- Attach the database that is being moved.

-- The path of the database files must be the location where you have stored the database files. Example:

CREATE DATABASE [CustRecords] ON

( FILENAME = N'C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA

\CustRecords.mdf' ),

( FILENAME = N'C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA

\CustRecords\_log.LDF' ) FOR ATTACH ;  
GO

From scenario: 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.

References:

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

#### NEW QUESTION 229

You are the administrator of a Microsoft SQL Server 2014 server.

Some applications consume significant resources. You need to manage the server workload by restricting resource-intensive applications

You need to dynamically limit resource consumption. What should you do?

- A. Configure Resource Pools, Workload Groups, and Classifier Function, and then enable the Resource Governor
- B. Set up Service Broker to ensure that application are not allowed to consume more than the specified amount of resource
- C. Create a new rule for each application that sets the resource limit allowed
- D. Create a new plan Guide with a Scope Type of sql and define the resource limits for each application

**Answer:** A

**Explanation:** In the SQL Server Resource Governor, a resource pool represents a subset of the physical resources of an instance of the Database Engine.

Resource Governor enables you to specify limits on the amount of CPU, physical IO, and memory that incoming application requests can use within the resource pool. Each resource pool can contain one or more workload groups. When a session is started, the Resource Governor classifier assigns the session to a specific workload group, and the session must run using the resources assigned to the workload group.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor-resou>

#### NEW QUESTION 234

You plan to migrate a Microsoft sql server instance between physical servers. You must migrate the metadata associated with the database instance.

You need to ensure that the new instance retains the existing jobs and alerts. Solutions: You restore the model database.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:** The model database does not handle alerts and jobs. It is used as the template for all databases created on an instance of SQL Server.

The msdb database is used by SQL Server Agent for scheduling alerts and jobs and by other features such as SQL Server Management Studio, Service Broker and Database Mail.

References:

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

#### NEW QUESTION 238

You plan to migrate a Microsoft SQL server instance between physical servers. You must migrate the metadata associated with the database instance.

You need to ensure that the new instance retains the existing jobs and alerts. Solutions: You restore the master database.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:** The master database does not handle alerts and jobs. It records all the system-level information for a SQL Server system. This includes instance-wide metadata such as logon accounts, endpoints, linked servers, and system configuration settings.

The msdb database is used by SQL Server Agent for scheduling alerts and jobs and by other features such as SQL Server Management Studio, Service Broker and Database Mail.

References:

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

#### NEW QUESTION 242

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

The following error message is an example error message:

"The job failed. The Job was invoked by User CONTOSO\ServiceAccount. The last step to run was step 1 (Subplan\_1)."

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

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

**Answer:** B

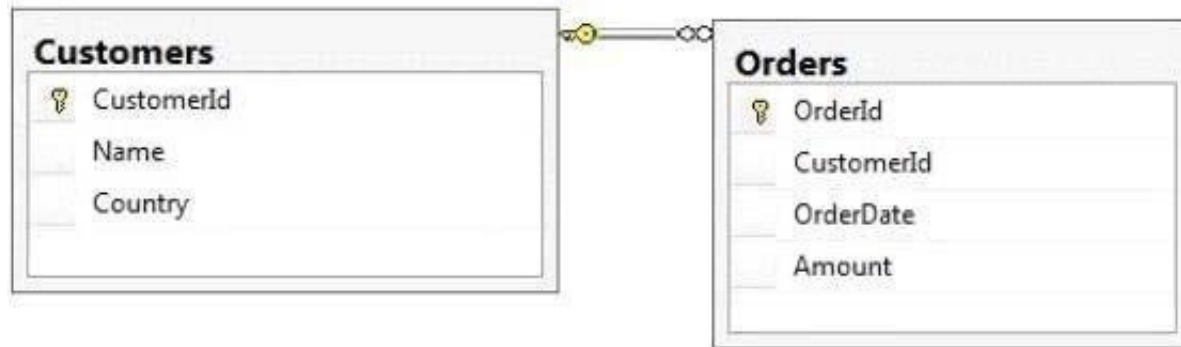
**Explanation:** References:

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



#### NEW QUESTION 246

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



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

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

Which Transact-SQL query should you use?

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

**Answer: G**

#### NEW QUESTION 251

You administer a SQL 2012 server that contains a database named SalesDB. SalesDb contains a schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales.

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

You need to remove the Select permission for UserA on the Regions table. You also need to ensure that UserA can still access all the tables in the Customers schema, including the Regions table, through the Sales role permissions.

Which Transact-SQL statement should you use?

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

**Answer: E**

**Explanation:** Use REVOKE to remove the grant or deny of a permission.

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

#### NEW QUESTION 254

You have Microsoft SQL Server on a Microsoft Azure virtual machine that has a database named DB1. You discover that DB1 experiences WRITE\_LOG waits that

are longer than 50 ms.  
You need to reduce the WRITE\_LOG wait time. Solution: Add additional log files to DB1.  
Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:** This problem is related to the disk response time, not to the number of log files.

References:

<https://www.mssqltips.com/sqlservertip/4131/troubleshooting-sql-server-transaction-log-related-wait-types/>

#### NEW QUESTION 256

You manage an on-premises, multi-tier application that has the following configuration:

Two SQL Server 2012 databases named SQL1 and SQL2

Two application servers named AppServer1 and AppServer2 that run IIS You plan to move your application to Azure.

You need to ensure that during an Azure update cycle or a hardware failure, the application remains available.

Which two deployment configurations should you implement? Each correct answer presents part of the solution.

- A. Deploy AppServer1 and AppServer2 in a single availability set.
- B. Deploy all servers in a single availability set.
- C. Deploy SQL1 and AppServer1 in a single availability set.
- D. Deploy SQL2 and AppServer2 in a single availability set.
- E. Deploy SQL1 and SQL2 in a single availability set.

**Answer:** AE

**Explanation:** You should deploy AppServer1 and AppServer2 in a single availability set. You should deploy SQL1 and SQL2 in a single availability set.

Note: Using availability sets allows you to build in redundancy for your Azure services. By grouping related virtual machines and services (tiers) into an availability set (in this case, deploying both of your databases into an availability set), you ensure that if there is a planned or unplanned outage, your services will remain available. At the most basic level, virtual machines in an availability set are put into a different fault domain and update domain. An update domain allows virtual machines to have updates installed and then the virtual machines are rebooted together.

If you have two virtual machines in an availability set, each in its own update domain, a rebooting of one server does not bring down all of the servers in a given tier. A fault domain operates in the same manner, so if there is a physical problem with a server, rack, network, or other service, both machines are separated, and services will continue.

#### NEW QUESTION 258

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

The financial database has the following characteristics:

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

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

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

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

You need to ensure that the minimum amount of data is lost. Which recovery model should the database use?

- A. NORECOVERY
- B. FULL
- C. NO\_CHECKSUM
- D. CHECKSUM
- E. Differential
- F. BULK\_LOGGED
- G. STANDBY
- H. RESTART
- I. SKIP
- J. Transaction log
- K. DBO ONLY
- L. COPY\_ONLY
- M. SIMPLE
- N. CONTINUE AFTER ERROR

**Answer:** B

**Explanation:** The full recovery model requires log backups. No work is lost due to a lost or damaged data file. Can recover to a specific point in time, assuming that your backups are complete up to that point in time.

#### NEW QUESTION 259

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

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

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

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

**Answer: D**

#### NEW QUESTION 260

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

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

You run the following Transact-SQL statements:

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

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

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

**Answer: D**

#### NEW QUESTION 261

You administer a Microsoft SQL Server 2014 instance.

You need to configure a new database to support FILETABLES. What should you do? Choose all that apply.

- A. Disable FILESTREAM on the Database.
- B. Enable FILESTREAM on the Server Instance.
- C. Configure the Database for Partial Containment.
- D. Create a non-empty FILESTREAM file group.
- E. Enable Contained Databases on the Server Instance.
- F. Set the FILESTREAM directory name on the Database.

**Answer: BDF**

**Explanation:** B: FileTables extend the capabilities of the FILESTREAM feature of SQL Server. Therefore you have to enable FILESTREAM for file I/O access at the Windows level and on the instance of SQL Server before you can create and use FileTables.

D: Before you can create FileTables in a database, the database must have a FILESTREAM filegroup. F: Specifying a Directory for FileTables at the Database Level

When you enable non-transactional access to files at the database level, you can optionally provide a directory name at the same time by using the DIRECTORY\_NAME option. If you do not provide a directory name when you enable non-transactional access, then you have to provide it later before you can create FileTables in the database.

References:

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

#### NEW QUESTION 262

You manage a Microsoft SQL Server instance named SQL1 that has 32 gigabytes (GB) of total memory. The instance supports an app named App1 that only uses



a single thread. App1 frequently queries the database using the same index. The operating system and App1 combined require 8 GB of memory to function. You need to ensure that the SQL Server does not limit the performance of App1. What configuration option should you set?

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

**Answer:** B

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

References:

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

#### NEW QUESTION 264

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

What should you do?

- A. Query the sys.dm\_db\_index\_usage\_stats DMV.
- B. Query the sys.dm\_db\_missing\_index\_details DMV.
- C. Use the Database Engine Tuning Advisor.
- D. Query the sys.dm\_db\_missing\_index\_columns DMV.

**Answer:** C

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

#### NEW QUESTION 267

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

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

Which statement should you execute?

- A. DBCC CHECKDB ('db1', REPAIR\_FAST)
- B. ALTER INDEX index1 ON table1 REBUILD WITH (ONLINE=ON)
- C. ALTER INDEX index1 ON table1 REORGANIZE
- D. DBCC CHECKDB ('db1', DATA\_PURITY)

**Answer:** B

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

#### NEW QUESTION 271

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your 70-765 Exam with Our Prep Materials Via below:**

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