

## 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 a virtual machine (VM) in Microsoft Azure, which has a 2 terabyte (TB) database. Microsoft SQL Server backups are performed by using Backup to URL. You need to provision the storage account for the backups while minimizing costs. Which storage option should you use?

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

**Answer:** G

**Explanation:** A URL specifies a Uniform Resource Identifier (URI) to a unique backup file. The URL is used to provide the location and name of the SQL Server backup file. The URL must point to an actual blob, not just a container. If the blob does not exist, it is created. If an existing blob is specified, BACKUP fails, unless the "WITH FORMAT" option is specified to overwrite the existing backup file in the blob.

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

**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 present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals.

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

You need to create the templates.

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

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

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

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

```
{  
"$schema": "http://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",  

```

```
"contentVersion": "", "parameters": { },
"variables": { },
"resources": [ ],
"outputs": { }
}
```

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

#### NEW QUESTION 4

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

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

You need to create the templates.

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

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

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

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

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

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

#### NEW QUESTION 5

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 6

DRAG DROP

You deploy a new Microsoft Azure SQL Database instance to support a variety of mobile applications and public websites. You plan to create a new security principal named User1.

The principal must have access to select all current and future objects in a database named Reporting. The activity and authentication of the database user must be limited to the Reporting database.

You need to create the new security principal.

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

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

**Answer:**

**Explanation:** Step 1, Step 2:

First you need to create a login for SQL Azure, it's syntax is as follows: CREATE LOGIN username WITH password='password'; This command needs to run in master db. Only afterwards can you run commands to create a user in the database.

Step 3:

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

#### NEW QUESTION 7

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

You manage a Microsoft SQL Server environment with several databases.

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

Solution: You enable the PARAMETER\_SNIFFING option for the databases. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:** PARAMETER\_SNIFFING = { ON | OFF | PRIMARY} enables or disables parameter sniffing. This is equivalent to Trace Flag 4136.

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

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

#### NEW QUESTION 8

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

You manage on-premises and Microsoft Azure SQL Database instances for a company. Your environment must support the Microsoft SQL Server 2012 ODBS driver.

You need to encrypt only specific columns in the database. What should you implement?

- A. transport-level encryption
- B. cell-level encryption



- C. Transparent Data Encryption
- D. Always Encrypted
- E. Encrypting File System
- F. BitLocker
- G. dynamic data masking

**Answer:** D

**Explanation:** To encrypt columns you can configure Always Encrypted.

SQL Server Management Studio (SSMS) provides a wizard that helps you easily configure Always Encrypted by setting up the column master key, column encryption key, and encrypted columns for you.

Always Encrypted allows client applications to encrypt sensitive data and never reveal the data or the encryption keys to SQL Server or Azure SQL Database. An Always Encrypted enabled driver, such as the ODBC Driver 13.1 for SQL Server, achieves this by transparently encrypting and decrypting sensitive data in the client application.

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

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault#encrypt-columns-configure-always-encrypted>  
[https://msdn.microsoft.com/en-us/library/mt637351\(v=sql.110\).aspx](https://msdn.microsoft.com/en-us/library/mt637351(v=sql.110).aspx)

## NEW QUESTION 9

### 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 style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">▼</div> <div style="border: 1px solid black; padding: 2px;">CREATE USER</div> <div style="border: 1px solid black; padding: 2px;">ALTER SERVER ROLE</div> <div style="border: 1px solid black; padding: 2px;">CREATE LOGIN</div>	Admin2 WITH password = 'Pa\$\$w0rd';
<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">▼</div> <div style="border: 1px solid black; padding: 2px;">CREATE USER</div> <div style="border: 1px solid black; padding: 2px;">ALTER SERVER ROLE</div> <div style="border: 1px solid black; padding: 2px;">CREATE LOGIN</div>	Admin2User FROM
	<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">▼</div> <div style="border: 1px solid black; padding: 2px;">WINDOWS</div> <div style="border: 1px solid black; padding: 2px;">EXTERNAL PROVIDER</div> <div style="border: 1px solid black; padding: 2px;">LOGIN</div>
ALTER ROLE '	Admin2
	<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">▼</div> <div style="border: 1px solid black; padding: 2px;">loginmanager</div> <div style="border: 1px solid black; padding: 2px;">dbmanager</div> <div style="border: 1px solid black; padding: 2px;">bd_ddladmin</div>

**Answer:**

**Explanation:** Step 1: CREATE LOGIN

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

Step 2, CREATE USER Step 3: LOGIN

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

CREATE USER readonlyuser FROM LOGIN readonlylogin; Step 4: loginmanager

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

References:

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

## NEW QUESTION 10

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

Your company has several Microsoft Azure SQL Database instances.

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

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

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

**Answer:** A

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

**NEW QUESTION 10**

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 15**

HOTSPOT

You need to create the contosodb1 database.

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

### Answer Area

	▼
New-AzureSqlDatabase	
New-AzureRmSqlDatabase	
Set-AzureRmSqlDatabase	

- ResourceGroupName “contosodbrg”

- ServerName “contososrv”

- DatabaseName “contosodbl”

- Edition

	▼
Basic	
Standard	
Premium	

- RequestedServiceObjectName S2

**Answer:**

**Explanation:** Box 1: New-AzureRmSqlDatabase

New-AzureRmSqlDatabase creates a database or an elastic database.

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

### NEW QUESTION 17

DRAG DROP

You create a login named BIAppUser. The login must be able to access the Reporting database.

You need to grant access to the BIAppUser login in the database.

How should you complete the Transact-SQL statements? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

### Code segments

Reporting
master
CREATE USER
ALTER LOGIN
ALTER USER
FOR LOGIN [BIAppUser]
FOR USER [BIAppUser]
WITH LOGIN = [BIAppUser]

### Answer area

```
USE [ Code segment ]
GO
Code segment [BIAppUser] Code segment
GO
```

**Answer:**

**Explanation:** Box 1: Reporting

The user is to be created in the Reporting database.

Box 2: CREATE USER

Box 3: FOR LOGIN [BIAppUser]

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. Here is some sample Transact-SQL that creates a user:

CREATE USER readonlyuser FROM LOGIN readonlylogin;

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

## NEW QUESTION 21

### 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 24

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

A. New-AzureSqlJob



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

**Answer:** A

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

References:

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

#### NEW QUESTION 28

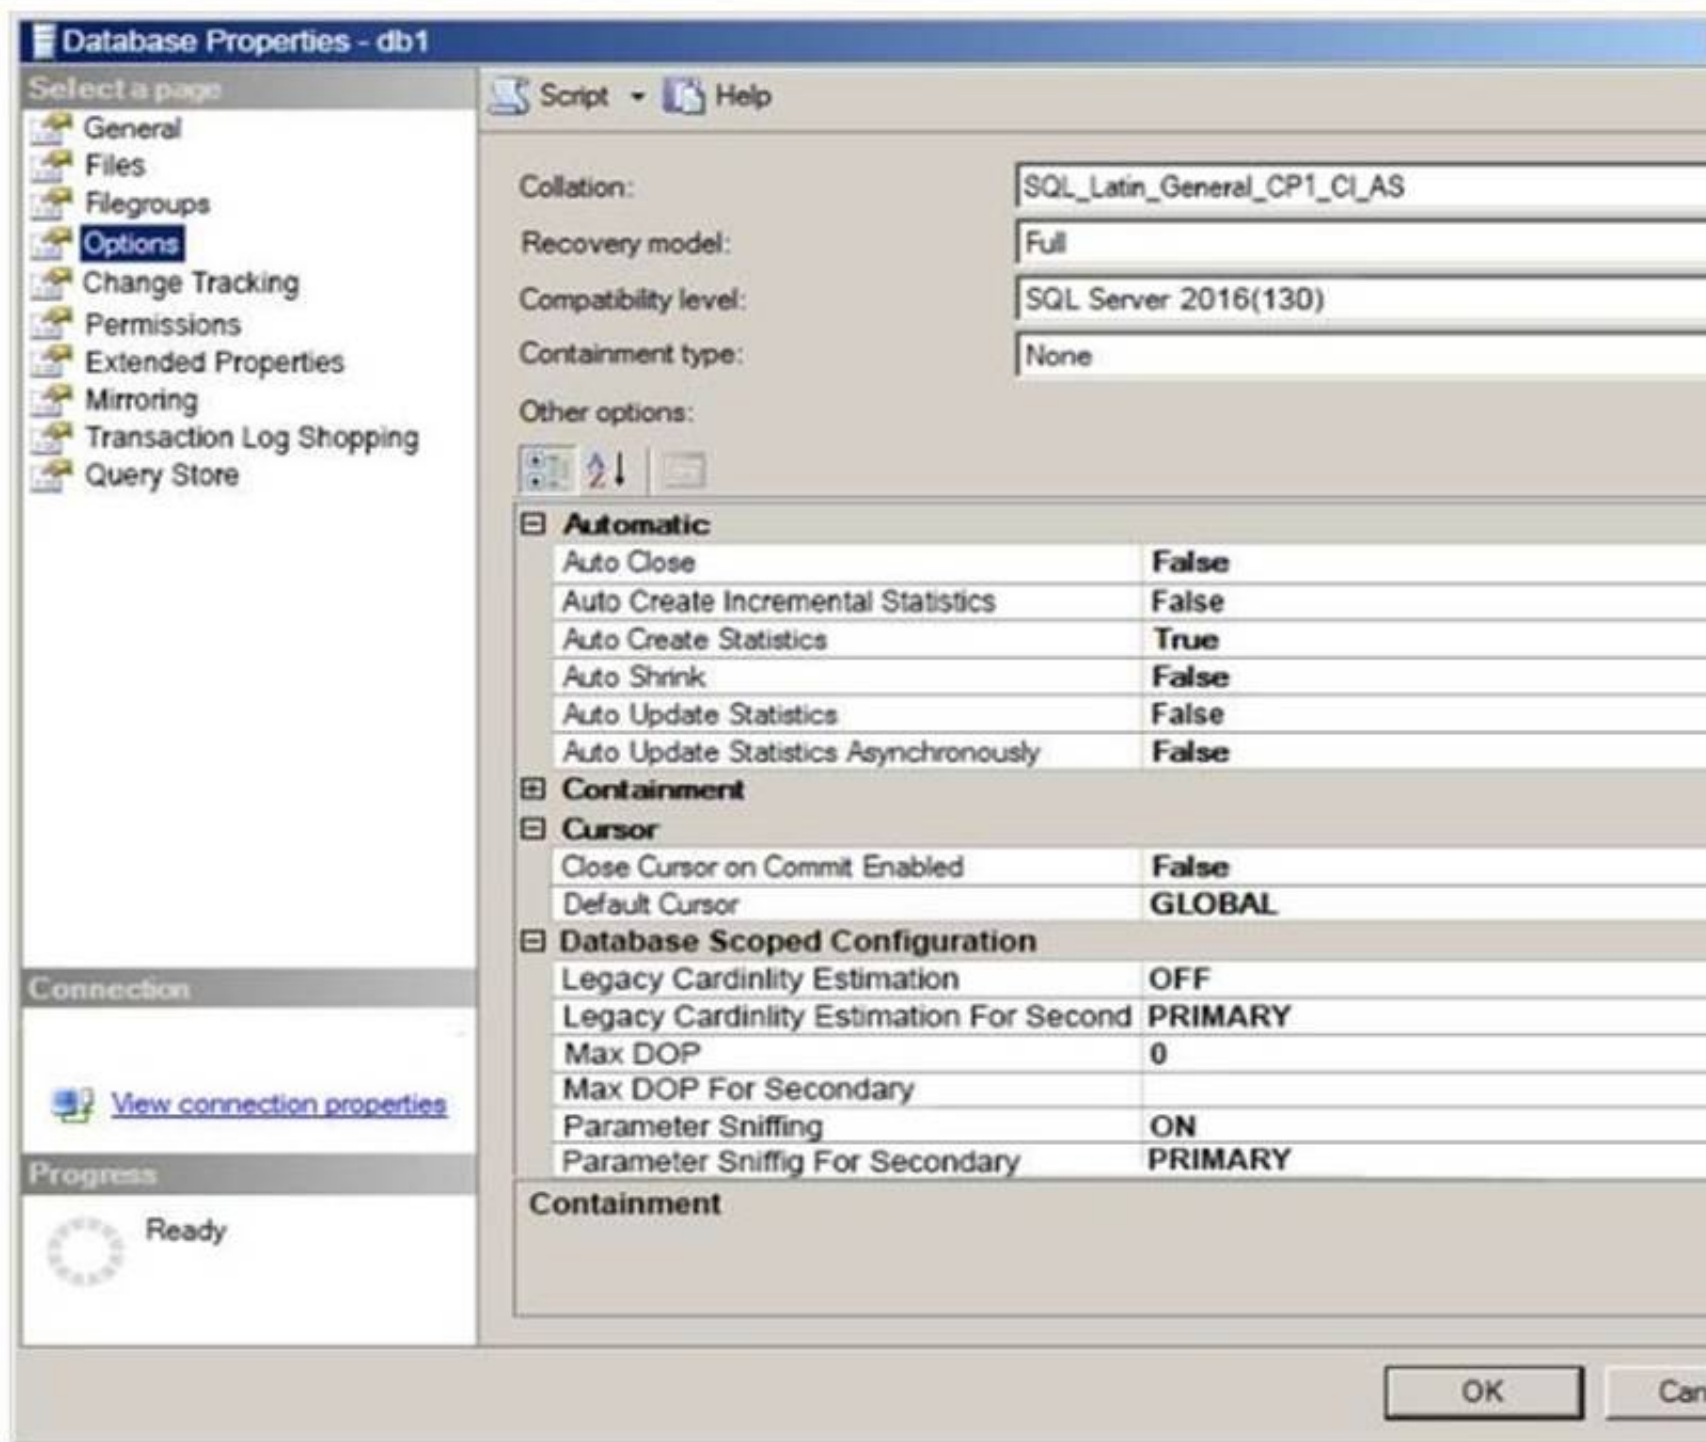
You have Microsoft SQL Server on a Microsoft Azure virtual machine. The virtual machine has a database named DB1. DB1 contains a table named Table1 that has 4 billion rows.

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 30

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

You need to implement real-time encryption of the database and all backups. Solution: you enable Dynamic Data Masking on the primary replica.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:** SQL Database dynamic data masking does not encrypt the data. Transparent Data Encryption (TDE) would provide a solution.

Note: SQL Database dynamic data masking limits sensitive data exposure by masking it to non-privileged users.

Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer.

References:

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

#### NEW QUESTION 31

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 36

You administer a Microsoft SQL Server 2014 server. One of the databases on the server supports a highly active OLTP application.

Users report abnormally long wait times when they submit data into the application.

You need to identify which queries are taking longer than 1 second to run over an extended period of time. What should you do?

- A. use SQL Profiler to trace all queries that are processing on the server
- B. Filter queries that have a Duration value of more than 1,000.
- C. Use sp\_configure to set a value for blocked process threshold
- D. Create an extended event session.
- E. Use the Job Activity monitor to review all processes that are actively running
- F. Review the Job History to find out the duration of each step.
- G. Run the sp\_who command from a query window.
- H. Run the DBCC TRACEON 1222 command from a query window and review the SQL Server event log.

**Answer:** A

#### NEW QUESTION 38

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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

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

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

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

Solution: You create one storage account that has 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 43

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 47

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 Premium database pool that has 125 eDTUs. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:** We need at least 400 eDTUs.

#### NEW QUESTION 51

You administer a Microsoft SQL Server 2014 database named Orders. Orders contains a table named OrderShip that is defined as follows:

```
CREATE TABLE OrderShip
(OrderID bigint NOT NULL PRIMARY KEY,
 CustomerID int NOT NULL,
 ShipAddress nvarchar(500) NOT NULL,
 CountryCode tinyint NULL)
```

A NULL value represents a domestic order. Ninety percent of the values in CountryCode are NULL. Customers require a procedure that will return orders for all customers from a specified country. You create a new procedure:

```
CREATE PROCEDURE p_GetIntlOrders
(@countrycode tinyint)
AS
SELECT DISTINCT CustomerID, ShipAddress
FROM OrderShip
WHERE CountryCode = @countrycode
GO
```

Performance on this procedure is slow.

You need to alter the schema to optimize this query. Objects created must use a minimum amount of resources.

Which Transact-SQL statement should you use?

- A. CREATE NONCLUSTERED INDEX IX\_CountryCode ON Ordership (CountryCode) WHERE CountryCode IS NOT NULL
- B. CREATE STATISTICS ST\_CountryCode ON OrderShip (CountryCode) WHERE CountryCode IS NOT NULL
- C. CREATE CLUSTERED INDEX IX\_CountryCode ON OrderShip (CountryCode)
- D. CREATE INDEX IX\_CountryCode ON OrderShip (CustomerID) WHERE CountryCode IS NOT NULL

**Answer:** B

**Explanation:** Here creating statistics is relevant. The CREATE STATISTICS command creates query optimization statistics on one or more columns of a table, an indexed view, or an external table. For most queries, the query optimizer already generates the necessary statistics for a high-quality query plan; in a few cases, you need to create additional statistics with CREATE STATISTICS or modify the query design to improve query performance.

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

#### NEW QUESTION 55

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 60

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/azuresqldbssupport/2018/01/15/automating-azure-sql-db-index-and-statistic-m>

**NEW QUESTION 64**

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 66**

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/azurermsql/new-azurermsqlelasticpool?view=azurermps-6>

#### NEW QUESTION 67

You administer a Microsoft SQL Server 2014 failover cluster.

You need to ensure that a failover occurs when the server diagnostics returns query\_processing error. Which server configuration property should you set?

- A. SqlOumperDumpFlags
- B. FailureConditionLevel
- C. HealthCheckTimeout
- D. SqlDumperDumpPath

**Answer: B**

**Explanation:** Use the FailureConditionLevel property to set the conditions for the Always On Failover Cluster Instance (FCI) to fail over or restart.

The failure conditions are set on an increasing scale. For levels 1-5, each level includes all the conditions from the previous levels in addition to its own conditions.

Note: The system stored procedure sp\_server\_diagnostics periodically collects component diagnostics on the SQL instance. The diagnostic information that is collected is surfaced as a row for each of the following components and passed to the calling thread.

The system, resource, and query process components are used for failure detection. The io\_subsystem and events components are used for diagnostic purposes only.

References:<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/configure-failurecondition>

#### NEW QUESTION 69

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

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

Tools, Select from these.

Answer Area

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

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

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 72

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 77

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.

Segments

elastic-pool  
list  
list-dbs  
list-editions

Answer Area

az sql Segment Segment

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 81

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.  
After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.  
You are tuning the performance of a virtual machines that hosts a Microsoft SQL Server instance. The virtual machine originally had four CPU cores and now has 32 CPU cores.  
The SQL Server instance uses the default settings and has an OLTP database named db1. The largest table in db1 is a key value store table named table1.  
Several reports use the PIVOT statement and access more than 100 million rows in table1.  
You discover that when the reports run, there are PAGELATCH\_IO waits on PFS pages 2:1:1, 2:2:1, 2:3:1, and 2:4:1 within the tempdb database.  
You need to prevent the PAGELATCH\_IO waits from occurring. Solution: You add more files to db1.  
Does this meet the goal?

- A. Yes  
B. No

**Answer:** A

**Explanation:** From SQL Server's perspective, you can measure the I/O latency from sys.dm\_os\_wait\_stats. If you consistently see high waiting for PAGELATCH\_IO, you can benefit from a faster I/O subsystem for SQL Server.

A cause can be poor design of your database - you may wish to split out data located on 'hot pages', which are accessed frequently and which you might identify as the causes of your latch contention. For example, if you have a currency table with a data page containing 100 rows, of which 1 is updated per transaction and you have a transaction rate of 200/sec, you could see page latch queues of 100 or more. If each page latch wait costs just 5ms before clearing, this represents a full half-second delay for each update. In this case, splitting out the currency rows into different tables might prove more performant (if less normalized and logically structured).

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

#### NEW QUESTION 83

You administer a Microsoft SQL Server 2014 database named Contoso on a server named Server01. You need to be notified immediately when fatal errors occur on Server01.

What should you create?

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

**Answer:** C

**Explanation:** SQL Server has alerts that get more important based on the severity of the alert. Anything of severity 16 or below tends to refer to the database and deals with issues that are tied to syntax errors, violations of foreign keys, etc. While those errors are typically important, they don't refer to anything with regards to overall health of the SQL Server. Alerts 17 through 25 do. Those are the ones your health checks are probably firing on.

Severity Level	Meaning
17	Insufficient Resources
18	Nonfatal Internal Error Detected
19	SQL Server Error in Resource
20	SQL Server Fatal Error in Current Process
21	SQL Server Fatal Error in Database (dbid) Process
22	SQL Server Fatal Error Table Integrity Suspect
23	SQL Server Fatal Error: Database Integrity Suspect
24	Hardware Error
25	(no description)

References: <https://www.mssqltips.com/sqlservertip/3384/configuring-critical-sql-server-alerts/>

#### NEW QUESTION 87

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 89

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 91

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 WarehouseGRANT 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 92

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:	<div>▼</div> <div>FAT32</div> <div>FAT</div> <div>NTFS</div>
8.3 filename support:	<div>▼</div> <div>Enabled</div> <div>Disabled</div>
Indexing:	<div>▼</div> <div>Enabled</div> <div>Disabled</div>

**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 96

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->

#### NEW QUESTION 98

You administer a Microsoft SQL Server 2014 server. The MSSQLSERVER service uses a domain account named CONTOSO\SQLService.

You plan to configure Instant File Initialization.

You need to ensure that Data File Autogrow operations use Instant File Initialization. What should you do? Choose all that apply.

- A. Restart the SQL Server Agent Service.
- B. Disable snapshot isolation.
- C. Restart the SQL Server Service.
- D. Add the CONTOSO\SQLService account to the Perform Volume Maintenance Tasks local security policy.
- E. Add the CONTOSO\SQLService account to the Server Operators fixed server role.
- F. Enable snapshot isolation.

**Answer:** CD

**Explanation:** How To Enable Instant File Initialization References:

<http://msdn.microsoft.com/en-us/library/ms175935.aspx>

#### NEW QUESTION 102

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 105

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 108

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

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

What should you do?

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

**Answer:** D

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

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

Local Key Stores References:

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

#### NEW QUESTION 113

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

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

What should you do?

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

**Answer:** C

**Explanation:** sp\_grantlogin creates a SQL Server login.

#### NEW QUESTION 117

Your database contains a table named Purchases. The table includes a DATETIME column named PurchaseTime that stores the date and time each purchase is made. There is a non-clustered index on the PurchaseTime column. The business team wants a report that displays the total number of purchases made on the current day. You need to write a query that will return the correct results in the most efficient manner.

Which Transact-SQL query should you use?

- A. SELECT COUNT(\*)FROM PurchasesWHERE PurchaseTime = CONVERT(DATE, GETDATE())
- B. SELECT COUNT(\*)FROM PurchasesWHERE PurchaseTime = GETDATE()
- C. SELECT COUNT(\*)FROM PurchasesWHERE CONVERT(VARCHAR, PurchaseTime, 112)=CONVERT(VARCHAR, GETDATE(), 112)
- D. SELECT COUNT(\*)FROM PurchasesWHERE PurchaseTime >= CONVERT(DATE, GETDATE())AND PurchaseTime <DATEADD(DAY, 1, CONVERT(DATE, GETDATE()))

**Answer:** D

**Explanation:** To compare a time with date we must use >= and > operators, and not the = operator.

#### NEW QUESTION 121

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 122**

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

One instance of SQL Server 2016 Enterprise

10 databases

500 stored procedures

You have a database named Database1 that is hosted on Server1.

Database1 contains 100 queries that are executed dynamically from web applications. You plan to remove data from the procedure cache on Database1.

You have the following requirements:

Changes to Database1 must not affect other databases that are hosted on Server1

Changes to Database1 must not affect the performance of queries that are stored in other databases.

The solution must minimize administrative effort.

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

- A. Run DBCC FREEPROCCACHE.
- B. Run ALTER DATABASE SCOPED CONFIGURATION CLEAR PROCEDURE CACHE in the context of Database 1.
- C. Run DBCC DROPCLEANBUFFERS.
- D. Write a script that iterates through each stored procedure definition and add WITH RECOMPILE to the definition.

**Answer:** B

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

**NEW QUESTION 124**

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 129**

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 133**

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 137**

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 138**

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 141**

You have Microsoft SQL Server on a Microsoft Azure Virtual machine that has a 4-TB database.

You plan to configure daily backups for the database. A single full backup will be approximately 1.5 TB of compressed data.

You need to ensure that the last backups are retained. Where should you store the daily backups?

- A. Local storage
- B. Page blob storage
- C. Virtual disks
- D. Block blob storage.

**Answer:** D

**Explanation:** When backing up to Microsoft Azure blob storage, SQL Server 2016 supports backing up to multiple blobs to enable backing up large databases, up to a maximum of 12.8 TB. This is done through Block Blobs.

References:

**NEW QUESTION 146**

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 150**

You are designing a Windows Azure SQL Database for an order fulfillment system. You create a table named Sales.Orders with the following script.



```
CREATE TABLE Sales.Orders
(
    OrderID int IDENTITY(1,1) NOT NULL PRIMARY KEY,
    OrderDate datetimeoffset NOT NULL,
    CustomerID int NOT NULL
);
```

Each order is tracked by using one of the following statuses:

Fulfilled  
Shipped  
Ordered  
Received

You need to design the database to ensure that that you can retrieve the following information:

The current status of an order  
The previous status of an order.  
The date when the status changed.  
The solution must minimize storage.

More than one answer choice may achieve the goal. Select the BEST answer.

- A. To the Sales.Orders table, add three columns named Status, PreviousStatus and ChangeDat
- B. Update rows as the order status changes.
- C. Create a new table named Sales.OrderStatus that contains three columns named OrderID, StatusDate, and Statu
- D. Insert new rows into the table as the order status changes.
- E. Implement change data capture on the Sales.Orders table.
- F. To the Sales.Orders table, add three columns named FulfilledDate, ShippedDate, and ReceivedDate.Update the value of each column from null to the appropriate date as the order status changes.

**Answer:** A

**Explanation:** This stores only the minimal information required.

#### NEW QUESTION 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 Standard database pool that has 800 eDTUs. Does the solution meet the goal?

- A. Yes
- B. No

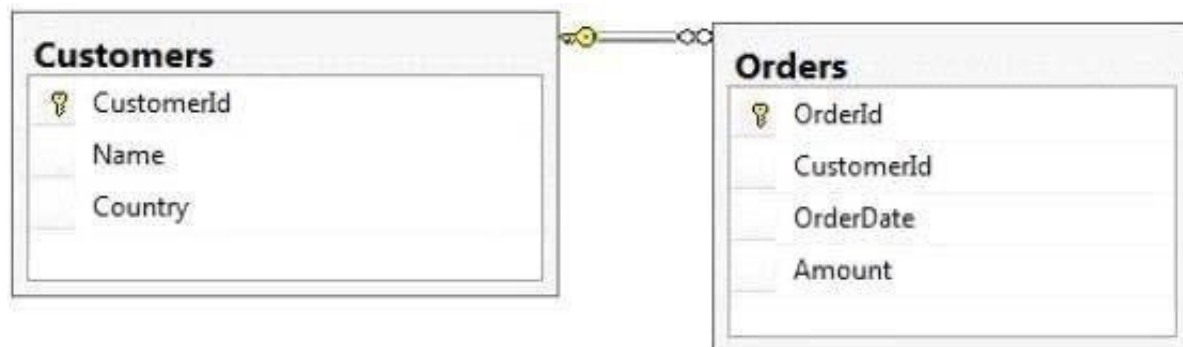
**Answer:** A

**Explanation:** We need at least 400 eDTUs and the use of a Standard database pool.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dtu-resource-limits>

#### NEW QUESTION 154

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 Orders INNER 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= FOR XML AUTO
- 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`, Country AS `@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:** G

#### NEW QUESTION 156

You are a database developer for an application hosted on a Microsoft SQL Server 2014 server. The database contains two tables that have the following definitions:

```
CREATE TABLE Customer
(CustomerID int NOT NULL PRIMARY KEY,
 CustomerName varchar(50) NOT NULL)
```

```
CREATE TABLE Orders
(OrderID int NOT NULL PRIMARY KEY,
 CustomerID int NOT NULL FOREIGN KEY REFERENCES Customer (CustomerID),
 OrderAmount money NOT NULL,
 ShippingCountry varchar(50) NOT NULL)
```

Global customers place orders from several countries. You need to view the country from which each customer has placed the most orders. Which Transact-SQL query do you use?

- A. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer cINNER JOIN(SELECT CustomerID, ShippingCountry,RANK() OVER (PARTITION BY CustomerIDORDER BY COUNT(OrderAmount) DESC) AS RnkFROM OrdersGROUP BY CustomerID, ShippingCountry) AS oON c.CustomerID = o.CustomerIDWHERE o.Rnk = 1
- B. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM(SELECT c.CustomerID, c.CustomerName, o.ShippingCountry, RANK()OVER (PARTITION BY CustomerIDORDER BY COUNT(o.OrderAmount) ASC) AS RnkFROM Customer cINNER JOIN Orders oON c.CustomerID = o.CustomerIDGROUP BY c.CustomerID, c.CustomerName, o.ShippingCountry) cs WHERE Rnk = 1
- C. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer cINNER JOIN(SELECT CustomerID, ShippingCountry,RANK() OVER (PARTITION BY CustomerIDORDER BY OrderAmount DESC) AS RnkFROM OrdersGROUP BY CustomerID, ShippingCountry) AS oON c.CustomerID = o.CustomerIDWHERE o.Rnk = 1
- D. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer cINNER JOIN(SELECT CustomerID, ShippingCountry,COUNT(OrderAmount) DESC) AS OrderAmountFROM OrdersGROUP BY CustomerID, ShippingCountry) AS oON c.CustomerID = o.CustomerIDORDER BY OrderAmount DESC

**Answer:** A

**Explanation:** Use descending (DESC) ordering.

To order by the number of orders we use ORDER BY COUNT(OrderAmount). Finally a WHERE clause is needed: WHERE o.Rnk = 1

#### NEW QUESTION 160

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 165**

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 166**

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

You suspect that the current SQL Server indexes cause queries to execute slowly.

You need to identify which indexes must be created to reduce the query execution time.

Which three dynamic management views should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. sys.dm\_db\_index\_physical\_stats
- B. sys.dm\_db\_missing\_index\_group\_stats
- C. sys.indexes
- D. sys.dm\_db\_index\_usage\_stats
- E. sys.dm\_db\_missing\_index\_groups
- F. sys.dm\_db\_index\_operational\_stats
- G. sys.dm\_db\_missing\_index\_details
- H. sys.sysindexkeys

**Answer:** BEG

**Explanation:** The missing indexes feature consists of the following components:

A set of dynamic management objects that can be queried to return information about missing indexes.

The Missing Indexes element in XML Showplans, which correlate indexes that the query optimizer considers missing with the queries for which they are missing.

Dynamic Management Objects

After running a typical workload on SQL Server, you can retrieve information about missing indexes by querying the dynamic management objects listed in the following table. These dynamic management objects are stored in the master database.

sys.dm\_db\_missing\_index\_group\_stats

Returns summary information about missing index groups, for example, the performance improvements that could be gained by implementing a specific group of missing indexes.

sys.dm\_db\_missing\_index\_groups

Returns information about a specific group of missing indexes, such as the group identifier and the identifiers of all missing indexes that are contained in that group.

sys.dm\_db\_missing\_index\_details

Returns detailed information about a missing index; for example, it returns the name and identifier of the table where the index is missing, and the columns and column types that should make up the missing index.

sys.dm\_db\_missing\_index\_columns

Returns information about the database table columns that are missing an index. References: [https://technet.microsoft.com/en-us/library/ms345524\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms345524(v=sql.105).aspx)

**NEW QUESTION 169**

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 171**



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

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

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

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

**Answer:** C

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

#### NEW QUESTION 173

A company has an on-premises Microsoft SQL Server 2017 infrastructure. The storage area network (SAN) that supports the SQL infrastructure has reached maximum capacity.

You need to recommend a solution to reduce on-premises storage use without changing the application. What should you do?

- A. Configure an Express Route connection to Microsoft Azure.
- B. Configure a Microsoft Azure Key Vault.
- C. Configure geo-replication on the SAN.
- D. Configure SQL Server Stretch Database in Microsoft Azure.

**Answer:** D

**Explanation:** Stretch warm and cold transactional data dynamically from SQL Server to Microsoft Azure with SQL Server Stretch Database. Unlike typical cold data storage, your data is always online and available to query. Benefit from the low cost of Azure rather than scaling expensive, on-premises storage.

References:

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

#### NEW QUESTION 174

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 178

You are a database administrator for a Microsoft SQL Server 2014 environment.

You want to deploy a new application that will scale out the workload to at least five different SQL Server instances.

You need to ensure that for each copy of the database, users are able to read and write data that will then be synchronized between all of the database instances.

Which feature should you use?

- A. Database Mirroring
- B. Peer-to-Peer Replication
- C. Log Shipping
- D. Availability Groups

**Answer:** B

**Explanation:** Peer-to-peer replication provides a scale-out and high-availability solution by maintaining copies of data across multiple server instances, also referred to as nodes. Built on the foundation of transactional replication, peer-to-peer replication propagates transactionally consistent changes in near real-time. This enables applications that require scale-out of read operations to distribute the reads from clients across multiple nodes. Because data is maintained across the nodes in near real-time, peer-to-peer replication provides data redundancy, which increases the availability of data.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/replication/transactional/peer-to-peer-trans>

#### NEW QUESTION 183

You are a database administrator for a Microsoft SQL Server 2014 database named AdventureWorks2012. You create an Availability Group defined by the following schema. (Line numbers are included for reference only.)



```
01 CREATE AVAILABILITY GROUP Group1
02 FOR DATABASE AdventureWorks2012
03 REPLICA ON 'SecondaryServer'
04 WITH (
05     ENDPOINT_URL = 'TCP://SecondaryServer:5022',
06     ...
07 );
```

You need to implement an AlwaysOnAvailability Group that will meet the following conditions:

Production transactions should be minimally affected. The secondary server should allow reporting queries to be performed. If the primary server goes offline, the secondary server should not automatically take over.

Which Transact-SQL statement should you insert at line 06?

- A. AVAILABILITY\_MODE = SYNCHRONOUS\_COMMIT,FAILOVER\_MODE = MANUALESECONDARY\_ROLE (ALLOW\_CONNECTIONS = READ\_ONLY,READ\_ONLY\_ROUTING\_URL = 'TCP://SecondaryServer:1433') PRIMARY\_ROLE (ALLOW\_CONNECTIONS = READ\_WRITE,READ\_ONLY\_ROUTING\_LIST = NONE)
- B. AVAILABILITY\_MODE = SYNCHRONOUS\_COMMIT,FAILOVER\_MODE = MANUALESECONDARY\_ROLE (ALLOW\_CONNECTIONS =READ\_ONLY,READ\_ONLY\_ROUTING\_URL = 'TCP://SecondaryServer:1433')
- C. AVAILABILITY\_MODE = ASYNCHRONOUS\_COMMIT,FAILOVER\_MODE = MANUALESECONDARY\_ROLE (ALLOW\_CONNECTIONS =READ\_ONLY,READ\_ONLY\_ROUTING\_URL = 'TCP://SecondaryServer:1433')
- D. AVAILABILITY\_MODE = ASYNCHRONOUS\_COMMIT,FAILOVER\_MODE = MANUALESECONDARY\_ROLE (ALLOW\_CONNECTIONS =YES,READ\_ONLY\_ROUTING\_URL = 'TCP://SecondaryServer:1433')

**Answer:** C

**Explanation:** As production transaction should be MINIMALLY affected we should use asynchronous-commit mode.

#### NEW QUESTION 186

You administer a Windows Azure SQL Database database named Orders. You need to create a copy of Orders named Orders\_Reporting.

Which Transact-SQL command should you use?

- A. BACKUP DATABASE Orders TO DISK = 'D:\Orders.bak'RESTORE DATABASEOrders\_ReportingFROM DISK = 'D:\Orders.bak'
- B. BACKUP DATABASE Orders TO DISK = 'D:\Orders.bak'CREATE DATABASEOrders\_ReportingFROM DISK = 'D:\Orders.bak'
- C. CREATE DATABASE Orders\_Reporting AS COPY OF Orders
- D. BACKUP DATABASE Orders TO DISK = 'D:\Orders.bak'MIRROR TO DISK = 'Orders\_Reporting'

**Answer:** C

**Explanation:** BACKUP DATABASE ...AS COPY OF [source\_server\_name.]source\_database\_name Is used for copying a database to the same or a different SQL Database server.

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

#### NEW QUESTION 189

You have an on-premises SQL Server database named DB1 that contains a table named TB1. TB1 is stretched to Microsoft Azure.

A catastrophic hardware failure occurs on the on-premises SQL server.

You deploy a new on-premises server and restore all databases to the new server. You need to resume Stretch Database operations to Azure.

Which statements should you execute?

```
A. ALTER TABLE tb1
    SET (REMOTE_DATA_ARCHIVE ( MIGRATION_STATE = PAUSE ) ) ;
GO
sp_rda_get_rpo_duration

B. EXEC sp_rda_reauthorize_db @credential = <credential>;
GO
ALTER TABLE tb1
    SET ( REMOTE_DATA_ARCHIVE = ON (
        FILTER_PREDICATE = dbo.fn_stretchpredicae(),
        MIGRATION_STATE = OUTBOUND) );

C. sp_rda_deauthorize_db
GO
ALTER TABLE tb1
    SET ( REMOTE_DATA_ARCHIVE ( MIGRATION_STATE = PAUSE ) ) ;
    FILTER_PREDICATE = db

GO
EXEC sp_rda_reauthorize_db @credential = <credential>;
GO

D. GO
EXEC sp_rda_reauthorize_db @credential = <credential>;
GO
CREATE TABLE tb1
    ...
    WITH ( REMOTE_DATA_ARCHIVE = ON ( MIGRATION_STATE = OUTBOUND) );
```

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

**Answer:** B

**Explanation:** Use ALTER TABLE, not CREATE TABLE.

When you enable Stretch for a table by specifying ON, you also have to specify MIGRATION\_STATE = OUTBOUND to begin migrating data immediately, or MIGRATION\_STATE = PAUSED to postpone data migration.

Syntax:

```
<stretch_configuration> ::=
{
SET ( REMOTE_DATA_ARCHIVE
{
= ON ( <table_stretch_options> )
| = OFF_WITHOUT_DATA_RECOVERY ( MIGRATION_STATE = PAUSED )
| ( <table_stretch_options> [, ...n] )
}
)
}
```

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-table-transact-sql?view=sql-server-2017>

#### NEW QUESTION 192

You administer a Microsoft SQL Server 2014 database instance.

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

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

What should you do?

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

**Answer:** C

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

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

#### NEW QUESTION 196

You administer a Microsoft SQL Server 2014 database named Contoso on a server named Server01. You need to collect data for a long period of time to troubleshoot wait statistics when querying Contoso. You also need to ensure minimum impact to the server. What should you create?

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

**Answer: C**

**Explanation:** SQL Server Extended Events has a highly scalable and highly configurable architecture that allows users to collect as much or as little information as is necessary to troubleshoot or identify a performance problem. Extended Events is a light weight performance monitoring system that uses very few performance resources. A SQL Server Extended Events session is created in the SQL Server process hosting the Extended Events engine.  
References: <https://docs.microsoft.com/en-us/sql/relational-databases/extended-events/extended-events>

#### NEW QUESTION 198

User report that a query takes a long time to execute. The query has the following wait statistics.

```
<WaitStats>
  <Wait WaitType="MEMORY_ALLOCATION_EXT" WaitTimeMs="186" WaitCount="112046" />

  <Wait WaitType="PAGEIOLATCH_SH" WaitTimeMs="37001" WaitCount="183" />
  <Wait WaitType="SOS_SCHEDULER_YIELD" WaitTimeMs="399" WaitCount="12321" />
  <Wait WaitType="WRITELOG" WaitTimeMs="1632" WaitCount="627" />
  <Wait WaitType="IO_COMPLETION" WaitTimeMs="100287" WaitCount="5300" />
  <Wait WaitType="PAGEIOLATCH_UP" WaitTimeMs="59652" WaitCount="21027" />
  <Wait WaitType="PAGEIOLATCH_EX" WaitTimeMs="1116329" WaitCount="1840528" />
</WaitStats>
```

Which resource causes the issue?

- A. processor
- B. disk
- C. blocking
- D. network

**Answer: B**

**Explanation:** PAGEIOLATCH Wait time and WaitCount are both high. One of the most common wait type seen on SQL Server and definitely one that causes a lot of troubles to less experienced database administrators is the PAGEIOLATCH\_SH wait type. This is one of those wait types that clearly indicates one thing, but which background and potential causes are much subtler and may lead to erroneous conclusions and worse, incorrect solutions. The Microsoft definition of this wait type is: Occurs when a task is waiting on a latch for a buffer that is in an I/O request. The latch request is in Shared mode. Long waits may indicate problems with the disk subsystem.  
References: [https://www.sqlshack.com/handling-excessive-sql-server-pageiolatch\\_sh-wait-types/](https://www.sqlshack.com/handling-excessive-sql-server-pageiolatch_sh-wait-types/)

#### NEW QUESTION 202

You plan to deploy two new Microsoft Azure SQL Database instances. One instance will support a data entry application. The other instance will support the company's business intelligence efforts. The databases will be accessed by mobile applications from public IP addresses.

You need to ensure that the database instances meet the following requirements:

The database administration team must receive alerts for any suspicious activity in the data entry database, including potential SQL injection attacks.

Executives around the world must have access to the business intelligence application.

Sensitive data must never be transmitted. Sensitive data must not be stored in plain text in the database. In the table below, identify the feature that you must implement for each database.

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



# Answer Area

Option

Data entry

Business  
intelligence

Transparent Data Encryption

☐
☐

Dynamic Data Masking

☐
☐

Always Encrypted

☐
☐

Database-level firewall rules

☐
☐

Threat Detection

☐
☐

**Answer:**

**Explanation:** Data entry: Threat Detection

SQL Threat Detection provides a new layer of security, which enables customers to detect and respond to potential threats as they occur by providing security alerts on anomalous activities. Users receive an alert upon suspicious database activities, potential vulnerabilities, and SQL injection attacks, as well as anomalous database access patterns.

Business intelligence: Dynamic Data Masking

Dynamic data masking limits (DDM) sensitive data exposure by masking it to non-privileged users. It can be used to greatly simplify the design and coding of security in your application.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-threat-detection> <https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking>

## NEW QUESTION 206

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

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

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

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

**Answer:** E

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

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

## NEW QUESTION 208

You administer all the deployments of Microsoft SQL Server 2014 in your company. You have two servers in the same data center that hosts your production database.

You need to ensure that the database remains available if a catastrophic server failure or a disk failure occurs. You also need to maintain transactional consistency of the data across both servers.

You need to achieve these goals without manual intervention. Which configuration should you use?

- A. Two servers configured in a Windows Failover Cluster in the same data center SQL Server configured as a clustered instance
- B. SQL Server that includes an application database configured to perform transactional replication
- C. Two servers configured in the same data center A primary server configured to perform log-shipping every 10 minutes A backup server configured as a warm standby
- D. Two servers configured in different data centers SQL Server Availability Group configured in Synchronous-Commit Availability Mode One server configured as an Active Secondary



- E. Two servers configured in the same data centerSQL Server Availability Group configured in Asynchronous-Commit Availability Mode One server configured as an Active Secondary
- F. Two servers configured in different data centersSQL Server Availability Group configured in Asynchronous-Commit Availability Mode
- G. SQL Server that includes an application database configured to perform snapshot replication
- H. Two servers configured on the same subnetSQL Server Availability Group configured in Synchronous-Commit Availability Mode

**Answer:** H

**Explanation:** Always On availability groups supports two availability modes—asynchronous-commit mode and synchronous-commit mode. Synchronous-commit mode emphasizes high availability over performance, at the cost of increased transaction latency.

References:<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-mode>

#### NEW QUESTION 212

You are the database administrator for your company. Your company has one main office and two branch offices. You plan to create three databases named DB1, DB2, and DB3 that will be hosted on one Azure SQL Database server. You have the following requirements:

- The main office must be able to connect to all three databases.
- The branch offices must be able to connect to DB2 and DB3.
- The branch offices must not be able to access DB1.

You need to configure transparent data encryption (TDE) for DB1. Which two actions should you perform? Each correct answer presents part of the solution.

- A. Run `CREATE CERTIFICATE cert1 WITH Subject = TDE Cert1` on DB1.
- B. Connect to DB1.
- C. Run `ALTER DATABASE DB1 SET ENCRYPTION ON;`
- D. Connect to the master database.
- E. Run `CREATE MASTER KEY` on the master database.

**Answer:** BC

**Explanation:** You should connect to DB1. To encrypt DB1, you connect directly to DB1. When you connect to DB1, you use your dbmanager or administrative credentials.

You should run `ALTER DATABASE DB1 SET ENCRYPTION ON`.

You use the `ALTER DATABASE DB1 SET ENCRYPTION ON` statement to encrypt the database. This is the statement that turns on TDE for Azure SQL Database.

#### NEW QUESTION 214

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

You need to write messages to the Application Log when users are added to or removed from a fixed server role 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:** G

**Explanation:** The SQL Server Audit feature enables you to audit server-level and database-level groups of events and individual events.

Audits can have the following categories of actions:

Server-level. These actions include server operations, such as management changes, such as in this question, and logon and logoff operations.

Database-level. These actions encompass data manipulation languages (DML) and data definition language (DDL) operations.

Audit-level. These actions include actions in the auditing process.

References:

[http://technet.microsoft.com/en-us/library/cc280663\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/cc280663(v=sql.105).aspx)

#### NEW QUESTION 217

You create an availability group that has replicas named HA/Server01 and HA/Server02. Currently, HA/Server01 is the primary replica.

You have multiple queries that read data and produce reports from the database.

You need to offload the reporting workload to the secondary replica when HA/Server01 is the primary replica. What should you do?

- A. Set the Availability Mode property of HA/Server02 to Asynchronous commit.
- B. Set the Readable Secondary property of HA/Server02 to Read-intent only.
- C. Set the Connections in Primary Role property of HA/Server01 to Allow read/write connections.
- D. Set the Availability Mode property of HA/Server01 to Asynchronous commit.

**Answer:** B

**Explanation:** To set up a readable secondary replica, you first create an availability group. Then you add replicas. You can choose either Yes or Read-intent only options.

Specify an instance of SQL Server to host a secondary replica.

Replicas

Endpoints

Backup Preferences

Listener

Availability Replicas:

Server Instance	Initial Role	Automatic Failover (Up to 2)	Synchronous Commit (Up to 3)	Readable Secondary
SUNILA03-6PUH5I	Primary	<input type="checkbox"/>	<input type="checkbox"/>	No
SUNILA03-YLZO1U	Secondary	<input type="checkbox"/>	<input type="checkbox"/>	No
				No
				Yes
				Read-intent only

Add Replica...

Remove Replica

Summary for the replica hosted by SUNILA03-YLZO1U

Replica mode:

Asynchronous commit  
This replica will use asynchronous-commit availability mode and support only forced failover (with possible data loss).

Readable secondary:

No  
In the secondary role, this availability replica will not allow any connections.

References: <http://msdn.microsoft.com/en-us/library/jj542414.aspx>

**NEW QUESTION 221**

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

- A. Yes
- B. No

**Answer:** B

**Explanation:** The D drive should only be used for temporary data.

**NEW QUESTION 223**

You plan to create an AlwaysOn availability group that will have two replicas in Microsoft Azure and two on premises replicas. You need to configure the network to support the availability group listener. Which cmdlet should you run first?

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

**Answer:** B

**Explanation:** An availability group listener is a virtual network name that clients connect to for database access. On Azure virtual machines, a load balancer holds the IP address for the listener. The load balancer routes traffic to the instance of SQL Server that is listening on the probe port. Usually, an availability group uses an internal load balancer.

References:  
<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windowsportal-sql-ps-al>

**NEW QUESTION 227**

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: Move the transaction logs to a faster disk. Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:** Section: Deploy and migrate applications

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 228

.....

## 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>