

Exam Questions 70-765

Provisioning SQL Databases (beta)

<https://www.2passeasy.com/dumps/70-765/>



NEW QUESTION 1

- (Topic 1)

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

You have deployed 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 2

- (Topic 1)

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

You have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in a development environment. You need to provide storage to the environment that minimizes costs. Which storage option should you use?

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

Answer: D

NEW QUESTION 3

- (Topic 2)

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

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 4

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 5

- (Topic 2)

You manage a Microsoft SQL Server environment in a Microsoft Azure virtual machine. You must enable Always Encrypted for columns in a database. You need to configure the key store provider. What should you do?

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

Answer: D

Explanation:

Always Encrypted supports multiple key stores for storing Always Encrypted column master keys. A column master key can be a certificate stored in Windows Certificate Store. References: <https://msdn.microsoft.com/en-us/library/mt723359.aspx>

NEW QUESTION 6

- (Topic 2)

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

Your company has several Microsoft Azure SQL Database instances.

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 7

- (Topic 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.

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

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

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

Solution: You configure the Resource Governor to set the MAXDOP parameter to 0 for all queries against the database.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

You can override the max degree of parallelism value in queries by specifying the MAXDOP query hint in the query statement.
 References: [https://technet.microsoft.com/en-us/library/ms181007\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms181007(v=sql.105).aspx)

NEW QUESTION 8

HOTSPOT - (Topic 4)

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

- A. Mastered
- B. Not Mastered

Answer: A

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 9

HOTSPOT - (Topic 5)

You need to create the contosodb1 database.

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

Answer Area

▼
New-AzureSqlDatabase
New-AzureRmSqlDatabase
Set-AzureRmSqlDatabase

- ResourceGroupName “contosodbrg”

- ServerName “contososrv”

- DatabaseName “contosodbl”

- Edition

▼
Basic
Standard
Premium

- RequestedServiceObjectName S2

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: New-AzureRmSqlDatabase

New-AzureRmSqlDatabase creates a database or an elastic database.

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

NEW QUESTION 10

HOTSPOT - (Topic 6)

You need to open the firewall ports for use with SQL Server environment. In table below, identify the firewall port that you must use for each service.

NOTE: Make only one selection in each column.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Report Server: 80

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

NEW QUESTION 10

- (Exam Topic 7)

Settings Value VM size D3

Storage Location Drive E Storage type Standard Tempdb location Drive C

The workload on this instance has of the tempdb load.

You need to maximize the performance of the tempdb database.

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

A. Yes

B. No

Answer: B

Explanation:

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

References:

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

NEW QUESTION 11

- (Exam Topic 7)

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

You need to ensure that the cluster meets the SLA.

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

A. New-AzureRmAvailabilitySet

B. New-AzureRmLoadBalancer

C. New-AzureRmSqlDatabaseSecondary

D. New-AzureRmSqlElasticPool

E. New-AzureRmVM

F. New-AzureRmSqlServer

G. New-AzureRmSqlDatabaseCopy

H. New-AzureRmSqlServerCommunicationLink

Answer: B

Explanation:

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

References:

NEW QUESTION 14

- (Exam Topic 7)

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

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

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

A. Yes

B. No

Answer: B

Explanation:

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

Note: Use the Always Encrypted Wizard to help protect sensitive data stored in a SQL Server database. Always Encrypted allows clients to encrypt sensitive data inside client applications and never reveal the encryption keys to SQL Server.

References:

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

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

NEW QUESTION 16

- (Exam Topic 7)

You administer a SQL Server 2014 database instance.

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

A. A domain user

B. The BUILTIN\SYSTEM account

C. A local user with Run as Service permissions

D. The SQLBrowser account

Answer: A

Explanation:

References:

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

NEW QUESTION 18

- (Exam Topic 7)

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

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

Solution: You move the database to a 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 19

- (Exam Topic 7)

You have an on-premises Microsoft SQL Server named Server1.

You provision a Microsoft Azure SQL Database server named Server2. On Server1, you create a database named DB1.

You need to enable the Stretch Database feature for DB1.

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

Actions	Answer Area
-Create a master key in the master database	1.
-Create a firewall rule in Azure	2.
-Create a master key in DB1	3.
-Enable the remote data archive option in DB1	4.
-Create a database scoped credential for authentication to Azure.	5.
-Create a server-level credential for authentication to Azure.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Enable the remote data archive option in DB1 Prerequisite: Enable Stretch Database on the server

Before you can enable Stretch Database on a database or a table, you have to enable it on the local server. To enable Stretch Database on the server manually, run sp_configure and turn on the remote data archive option.

Step 2: Create a firewall rule in Azure

On the Azure server, create a firewall rule with the IP address range of the SQL Server that lets SQL Server communicate with the remote server.

Step 3: Create a master key in the master database

To configure a SQL Server database for Stretch Database, the database has to have a database master key. The database master key secures the credentials that Stretch Database uses to connect to the remote database.

Step 4: Create a database scoped credential for authentication to Azure

When you configure a database for Stretch Database, you have to provide a credential for Stretch Database to use for communication between the on premises SQL Server and the remote Azure server. You have two options.

Step 5: Create a server-level credential for authentication to Azure.

To configure a database for Stretch Database, run the ALTER DATABASE command. For the SERVER argument, provide the name of an existing Azure server, including the

.database.windows.net portion of the name - for example, MyStretchDatabaseServer.database.windows.net.

Provide an existing administrator credential with the CREDENTIAL argument, or specify FEDERATED_SERVICE_ACCOUNT = ON. The following example provides an existing credential.

```
ALTER DATABASE <database name> SET REMOTE_DATA_ARCHIVE = ON (
SERVER = '<server_name>',
CREDENTIAL = <db_scoped_credential_name>
); GO
```

References:

<https://docs.microsoft.com/en-us/sql/sql-server/stretch-database/enable-stretch-database-for-a-database?view=sq>

NEW QUESTION 21

- (Exam Topic 7)

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

Drive L fails and is replaced.

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
ALTER DATABASE '<your_database>' SET SINGLE_USER WITH ROLLBACK IMMEDIATE GO
DBCC CHECKDB ('<your_database>', REPAIR_REBUILD) GO
ALTER DATABASE '<your_database>' SET MULTI_USER GO
```

NEW QUESTION 23

- (Exam Topic 7)

You administer a Windows 2008 server hosting an instance of Microsoft SQL Server 2014 Standard Edition. The server hosts a database named Orders. Users report that a query that filters on OrderDate is taking an exceptionally long time. You discover that an index named IX_OrderDate on the CustomerOrder table is heavily fragmented.

You need to improve the performance of the IX_OrderDate index. The index should remain online during the operation.

Which Transact-SQL command should you use?

- A. ALTER INDEX IX_OrderDate ON CustomerOrder DISABLE
- B. ALTER INDEX IX_OrderDate ON CustomerOrder ENABLE
- C. ALTER INDEX IX_OrderDate ON CustomerOrder REORGANIZE
- D. ALTER INDEX IX_OrderDate ON CustomerOrder REBUILD

Answer: C

Explanation:

Reorganize: This option is more lightweight compared to rebuild. It runs through the leaf level of the index, and as it goes it fixes physical ordering of pages and also compacts pages to apply any previously set fillfactor settings. This operation is always online, and if you cancel it then it's able to just stop where it is (it doesn't have a giant operation to rollback).

References: <https://www.brentozar.com/archive/2013/09/index-maintenance-sql-server-rebuild-reorganize/>

NEW QUESTION 24

- (Exam Topic 7)

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

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

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

What should you include in your solution?

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 28

- (Exam Topic 7)

You administer a single server that contains a Microsoft SQL Server 2014 default instance. You plan to install a new application that requires the deployment of a database on the server. The application login requires sysadmin permissions.

You need to ensure that the application login is unable to access other production databases. What should you do?

- A. Use the SQL Server default instance and configure an affinity mask.
- B. Install a new named SQL Server instance on the server.
- C. Use the SQL Server default instance and enable Contained Databases.
- D. Install a new default SQL Server instance on the server.

Answer: B

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/sql-server/install/work-with-multiple-versions-and-instances-of-sql-server>

NEW QUESTION 31

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database that contains a table named OrderDetail.

You discover that the NCI_OrderDetail_CustomerID non-clustered index is fragmented. You need to reduce fragmentation. You need to achieve this goal without taking the index offline.

Which Transact-SQL batch should you use?

- A. CREATE INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID WITH DROP EXISTING
- B. ALTER INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID REORGANIZE
- C. ALTER INDEX ALL ON OrderDetail REBUILD
- D. ALTER INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID REBUILD

Answer: B

Explanation:

REORGANIZE specifies to reorganize the index leaf level. The REORGANIZE operation is always performed online. This means long-term blocking table locks are not held and queries or updates to the underlying table can continue during the ALTER INDEX REORGANIZE transaction.

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

NEW QUESTION 34

- (Exam Topic 7)

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

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

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

Which type of solution should you use?

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 36

- (Exam Topic 7)

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

The financial database has the following characteristics:

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

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

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

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

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

- A. STANDBY
- B. Differential
- C. FULL
- D. CHECKSUM
- E. BULK_LOGGED
- F. CONTINUE_AFTER_ERROR
- G. SIMPLE
- H. DBO_ONLY
- I. COPY_ONLY
- J. SKIP
- K. RESTART
- L. Transaction log
- M. NO_CHECKSUM
- N. NORECOVERY

Answer: F

Explanation:

The CONTINUE_AFTER_ERROR option, of the Transact-SQL BACKUP command, instructs BACKUP to continue despite encountering errors such as invalid checksums or torn pages.

References:

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

NEW QUESTION 39

- (Exam Topic 7)

You administer a SQL Server 2014 server that contains a database named SalesDB. SalesDb contains a schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales.

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

You need to ensure that UserA is disallowed to select from any of the tables in the Customers schema. Which Transact-SQL statement should you use?

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

Answer: G

Explanation:

Use SQL Data Warehouse or Parallel Data Warehouse GRANT and DENY statements to grant or deny a permission (such as UPDATE) on a securable (such as a database, table, view, etc.) to a security principal (a login, a database user, or a database role).

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

NEW QUESTION 42

- (Exam Topic 7)

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:	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="background-color: #e0e0e0; text-align: right;">▼</td></tr> <tr><td>FAT32</td></tr> <tr><td>FAT</td></tr> <tr><td>NTFS</td></tr> </table>	▼	FAT32	FAT	NTFS
▼					
FAT32					
FAT					
NTFS					
8.3 filename support:	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="background-color: #e0e0e0; text-align: right;">▼</td></tr> <tr><td>Enabled</td></tr> <tr><td>Disabled</td></tr> </table>	▼	Enabled	Disabled	
▼					
Enabled					
Disabled					
Indexing:	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="background-color: #e0e0e0; text-align: right;">▼</td></tr> <tr><td>Enabled</td></tr> <tr><td>Disabled</td></tr> </table>	▼	Enabled	Disabled	
▼					
Enabled					
Disabled					

- A. Mastered
- B. Not Mastered

Answer: A

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 45

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft Azure virtual machine that has 12 databases. All database files are in the same Azure Blob storage account. You need to receive an email notification if I/O operations to the database files exceed 800 MB/s for more than five minutes.

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

Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 49

- (Exam Topic 7)

You manage an on-premises Microsoft SQL server that has a database named DB1. An application named App1 retrieves customer information for DB1. Users report that App1 takes an unacceptably long time to retrieve customer records. You need to find queries that take longer than 400 ms to run.

Which statement should you execute?

A)

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

B)

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

C)

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

D)

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

- A. Option A
- B. Option B

- C. Option C
- D. Option D

Answer: B

Explanation:

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

NEW QUESTION 51

- (Exam Topic 7)

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

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

You need to recover DB1 and prevent any data loss.

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

Select and Place:

Select and Place:

Actions	Answer Area
Restore the tail-log backup.	1
Restore a full backup.	2
Perform a tail-log backup.	3
Restore the log backups.	4
Truncate the log of DB1.	
Delete DB1.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

Step 2: Restore a full backup.

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

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

Step 3: Restore the log backups.

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

Reference:

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

NEW QUESTION 53

- (Exam Topic 7)

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.

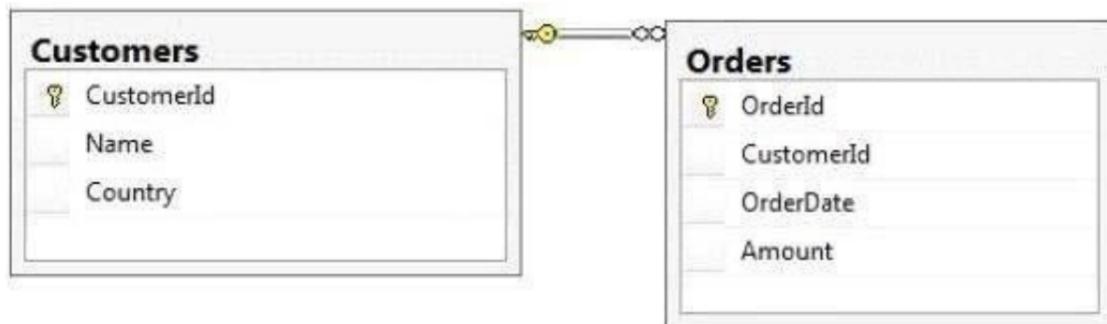
References:

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

NEW QUESTION 58

- (Exam Topic 7)

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



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

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

Which Transact-SQL query should you use?

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

Answer: G

NEW QUESTION 61

- (Exam Topic 7)

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->

NEW QUESTION 65

- (Exam Topic 7)

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

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

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

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

Answer: D

NEW QUESTION 70

- (Exam Topic 7)

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

The various databases on the instances have the following characteristics:

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

You run the following Transact-SQL statements:

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

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

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

Answer: D

NEW QUESTION 73

- (Exam Topic 7)

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

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

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

Which feature should you use?

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

Answer: B

Explanation:

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

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

NEW QUESTION 76

- (Exam Topic 7)

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

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

NEW QUESTION 77

- (Exam Topic 7)

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 79

- (Exam Topic 7)

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

You need to track all SELECT statements issued in the Contoso database only by users in a role named Sales. What should you create?

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

Answer: F

Explanation:

To audit users in a role use a Database Audit Specification.

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

NEW QUESTION 81

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance. After a routine shutdown, the drive that contains tempdb fails.

You need to be able to start the SQL Server. What should you do?

- A. Modify tempdb location in startup parameters.
- B. Start SQL Server in minimal configuration mode.
- C. Start SQL Server in single-user mode.
- D. Configure SQL Server to bypass Windows application logging.

Answer: B

Explanation:

If you have configuration problems that prevent the server from starting, you can start an instance of Microsoft SQL Server by using the minimal configuration startup option.

When you start an instance of SQL Server in minimal configuration mode, note the following: Only a single user can connect, and the CHECKPOINT process is not executed.

Remote access and read-ahead are disabled. Startup stored procedures do not run.

tempdb is configured at the smallest possible size.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/start-sql-server-with-minimal-configur>

NEW QUESTION 84

- (Exam Topic 7)

You plan to install a Microsoft SQL Server 2014 instance.

The instance will support a database that has the following requirements: Store Excel workbooks on the file system.

Access the workbooks through Transact-SQL.
Include the workbooks in database backups.
During installation, you need to ensure that the requirements will be met. Which feature should you use?

- A. Excel Services
- B. FILESTREAM
- C. SQL Server Integration Services (SSIS)
- D. OpenXML

Answer: B

Explanation:

FILESTREAM enables SQL Server-based applications to store unstructured data, such as documents and images, on the file system. Applications can leverage the rich streaming APIs and performance of the file system and at the same time maintain transactional consistency between the unstructured data and corresponding structured data.

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

NEW QUESTION 88

- (Exam Topic 7)

You use Microsoft SQL Server 2014 to develop a database application. You need to implement a computed column that references a lookup table by using an INNER JOIN against another table.

What should you do?

- A. Reference a user-defined function within the computed column.
- B. Create a BEFORE trigger that maintains the state of the computed column.
- C. Add a default constraint to the computed column that implements hard-coded values.
- D. Add a default constraint to the computed column that implements hard-coded CASE statements.

Answer: A

Explanation:

A common way to define a computed column is by using a user-defined function (UDF) to encapsulate the calculation logic.

References:<https://blogs.msdn.microsoft.com/sqlcat/2011/11/28/a-computed-column-defined-with-a-user-define>

NEW QUESTION 89

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-765 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-765 Product From:

<https://www.2passeasy.com/dumps/70-765/>

Money Back Guarantee

70-765 Practice Exam Features:

- * 70-765 Questions and Answers Updated Frequently
- * 70-765 Practice Questions Verified by Expert Senior Certified Staff
- * 70-765 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 70-765 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year