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 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
- (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 3
- (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 4
- (Topic 1)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals.
Your company plans to use Microsoft Azure Resource Manager templates for all future deployments of SQL Server on Azure virtual machines.
You need to create the templates.
Solution: You use Visual Studio to create a JSON template that defines the deployment and configuration settings for the SQL Server environment.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
Azure Resource Manager template consists of JSON, not XAML, and expressions that you can use to construct values for your deployment.
A good JSON editor can simplify the task of creating templates.
Note: In its simplest structure, an Azure Resource Manager template contains the following elements:
{
"$schema": "http://schema.management.azure.com/schemas/2015-01- 01/deploymentTemplate.json#",
"contentVersion": "", "parameters": { },
"variables": { },
"resources": [ ],
"outputs": { }
}
References:https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-authoring-templates

NEW QUESTION 5
- (Topic 1)
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 6
HOTSPOT - (Topic 2)
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

A. Mastered
B. Not Mastered

Answer: A

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 7
- (Topic 2)
You are deploying a Microsoft SQL Server database that will support a mixed OLTP and OLAP workload. The target virtual machine has four CPUs.
You need to ensure that reports do not use all available system resources. What should you do?

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

Answer: C

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

NEW QUESTION 8
- (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

NEW QUESTION 9
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	○	○
135	○	○
1433	○	○
2382	○	○

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
DRAG DROP - (Topic 6)
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

A. Mastered
B. Not Mastered

Answer: A

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 10
HOTSPOT - (Topic 6)
You need to set up the service accounts that the database engine and SQL Server Agent services will use.
How should you design the solution? To answer, select the appropriate configuration options in the answer area.

Answer Area

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Domain Account
The service startup account defines the Microsoft Windows account in which SQL Server Agent runs and its network permissions. SQL Server Agent runs as a specified user account. You select an account for the SQL Server Agent service by using SQL Server Configuration Manager, where you can choose fromthe following options:
* Built-in account. You can choose from a list of the following built-in Windows service accounts: Local System account.
* This account. Lets you specify the Windows domain account in which the SQL Server Agent service runs.
Box2: Domain users
Microsoft recommends choosing a Windows user account that is not a member of the Windows Administrators group.
Box 3: Managed Service Accounts
When resources external to the SQL Server computer are needed, Microsoft recommends using a Managed Service Account (MSA), configured with the minimum privileges necessary.
Note: A Managed Service Account (MSA) can run services on a computer in a secure and easy to maintain manner, while maintaining the capability to connect to network resources as a specific user principal.
References:https://msdn.microsoft.com/en-us/library/ms191543.aspx

NEW QUESTION 14
- (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 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 tempdb databases.
Does this meet the goal?

A. Yes
B. No

Answer: B

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 18
- (Exam Topic 7)
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 serve
B. Filter queries that have a Duration value of more than 1,000.
C. Use sp_configure to set a value for blocked process threshol
D. Create an extended event session.
E. Use the Job Activity monitor to review all processes that are actively runnin
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 19
- (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 22
- (Exam Topic 7)
You administer a Microsoft SQL Server 2014 database.
The database contains a Product table created by using the following definition:

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

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

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

Answer: D

NEW QUESTION 26
- (Exam Topic 7)
You administer a Microsoft SQL Server 2014 environment. One of the SQL Server 2014 instances contains a database named Sales.
You plan to migrate Sales to Windows Azure SQL Database. To do so, you need to implement a contained database.
What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

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

E. Disable server property contained database authentication.
F. database containment to FULL.

Answer: BD

Explanation:
A contained database is a database that is isolated from other databases and from the instance of SQL Server that hosts the database.
B: In the contained database user model, the login in the master database is not present. Instead, the authentication process occurs at the user database, and the database user in the user database does not have an associated login in the master database.
SQL Database and SQL Data Warehouse support Azure Active Directory identities as contained database users.
D: The contained database feature is currently available only in a partially contained state. A partially contained database is a contained database that allows the use of uncontained features.
References: https://docs.microsoft.com/en-us/sql/relational-databases/databases/contained-databases

NEW QUESTION 28
- (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 29
- (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
.d atabase.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 32
- (Exam Topic 7)
You use a Microsoft SQL Server 2014 database that contains two tables named SalesOrderHeader and SalesOrderDetail. The indexes on the tables are as shown in the exhibit.
(Click the Exhibit button.)

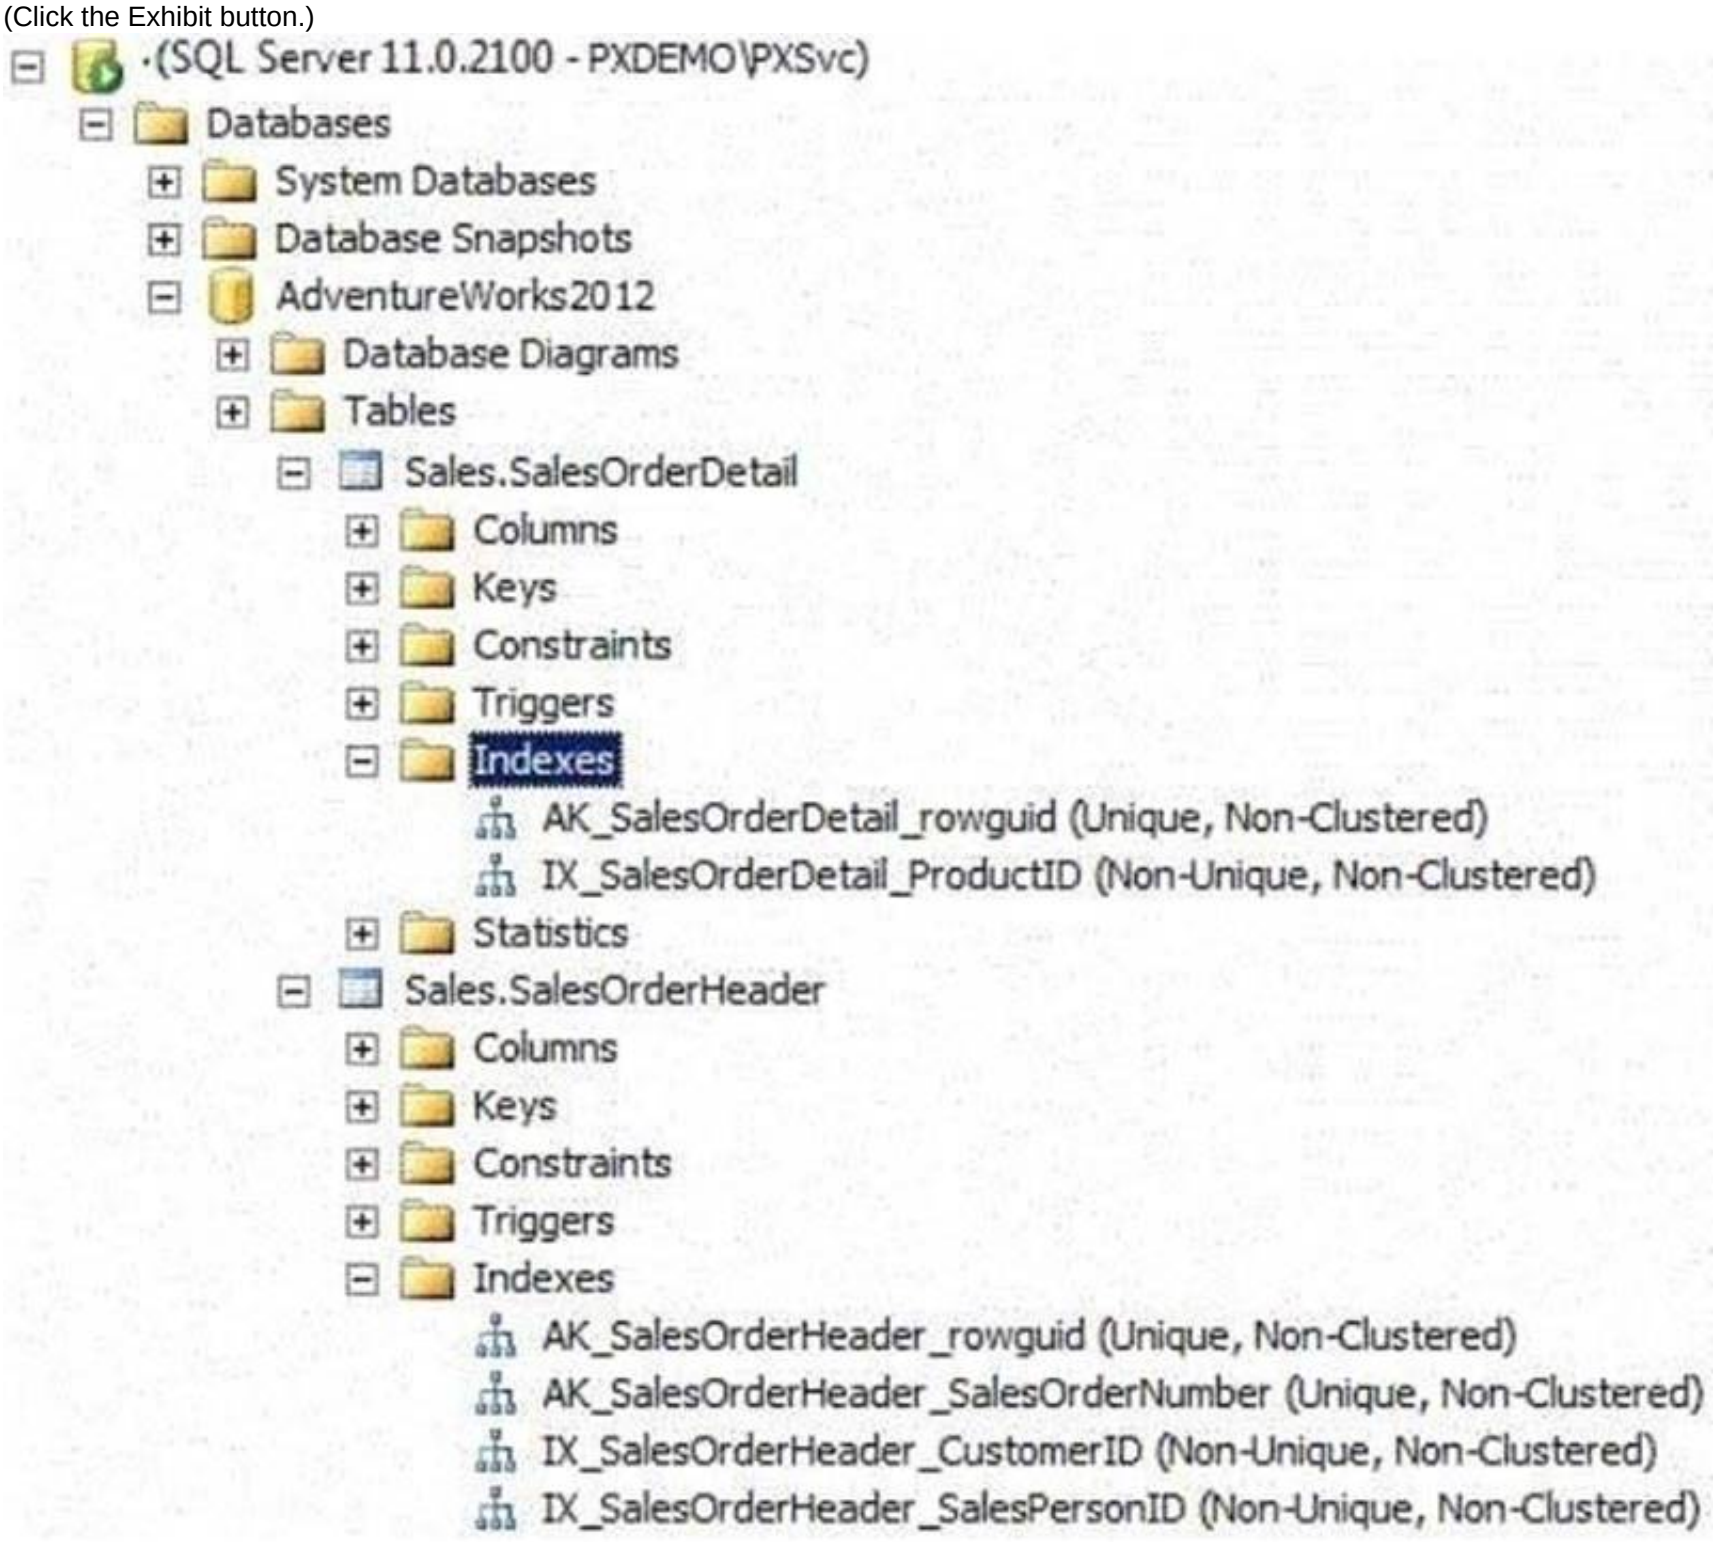

You write the following Transact-SQL query:

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

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

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

Answer: D

Explanation:
New statistics would be useful.
The UPDATE STATISTICS command updates query optimization statistics on a table or indexed view. By default, the query optimizer already updates statistics as necessary to improve the query plan; in some cases you can improve query performance by using UPDATE STATISTICS or the stored procedure sp_updatestats to update statistics more frequently than the default updates.
References:
http://msdn.microsoft.com/en-us/library/ms187348.aspx

NEW QUESTION 37
- (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 the Sales role, including 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: F

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-warehouse-

NEW QUESTION 39
- (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 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 41
- (Exam Topic 7)
Background
You manage the Microsoft SQL Server environment for a company that manufactures and sells automobile parts.
The environment includes the following servers: SRV1 and SRV2. SRV1 has 16 logical cores and hosts a SQL Server instance that supports a mission-critical application. The application has approximately 30,000 concurrent users and relies heavily on the use of temporary tables.
The environment also includes the following databases: DB1, DB2, and Reporting. The Reporting database is protected with Transparent Data Encryption (TDE).
You plan to migrate this database to a new server. You detach the database and copy it to the new server.
You are performing tuning on a SQL Server database instance. The application which uses the database was written using an object relationship mapping (ORM) tool which maps tables as objects within the application code. There are 30 stored procedures that are regularly used by the application.
After reviewing the plan cache you have identified that a large number of simple queries are using parallelism, and that execution plans are not being kept in the plan cache for very long.
You review the properties of the instance (Click the Exhibit button). Exhibit:

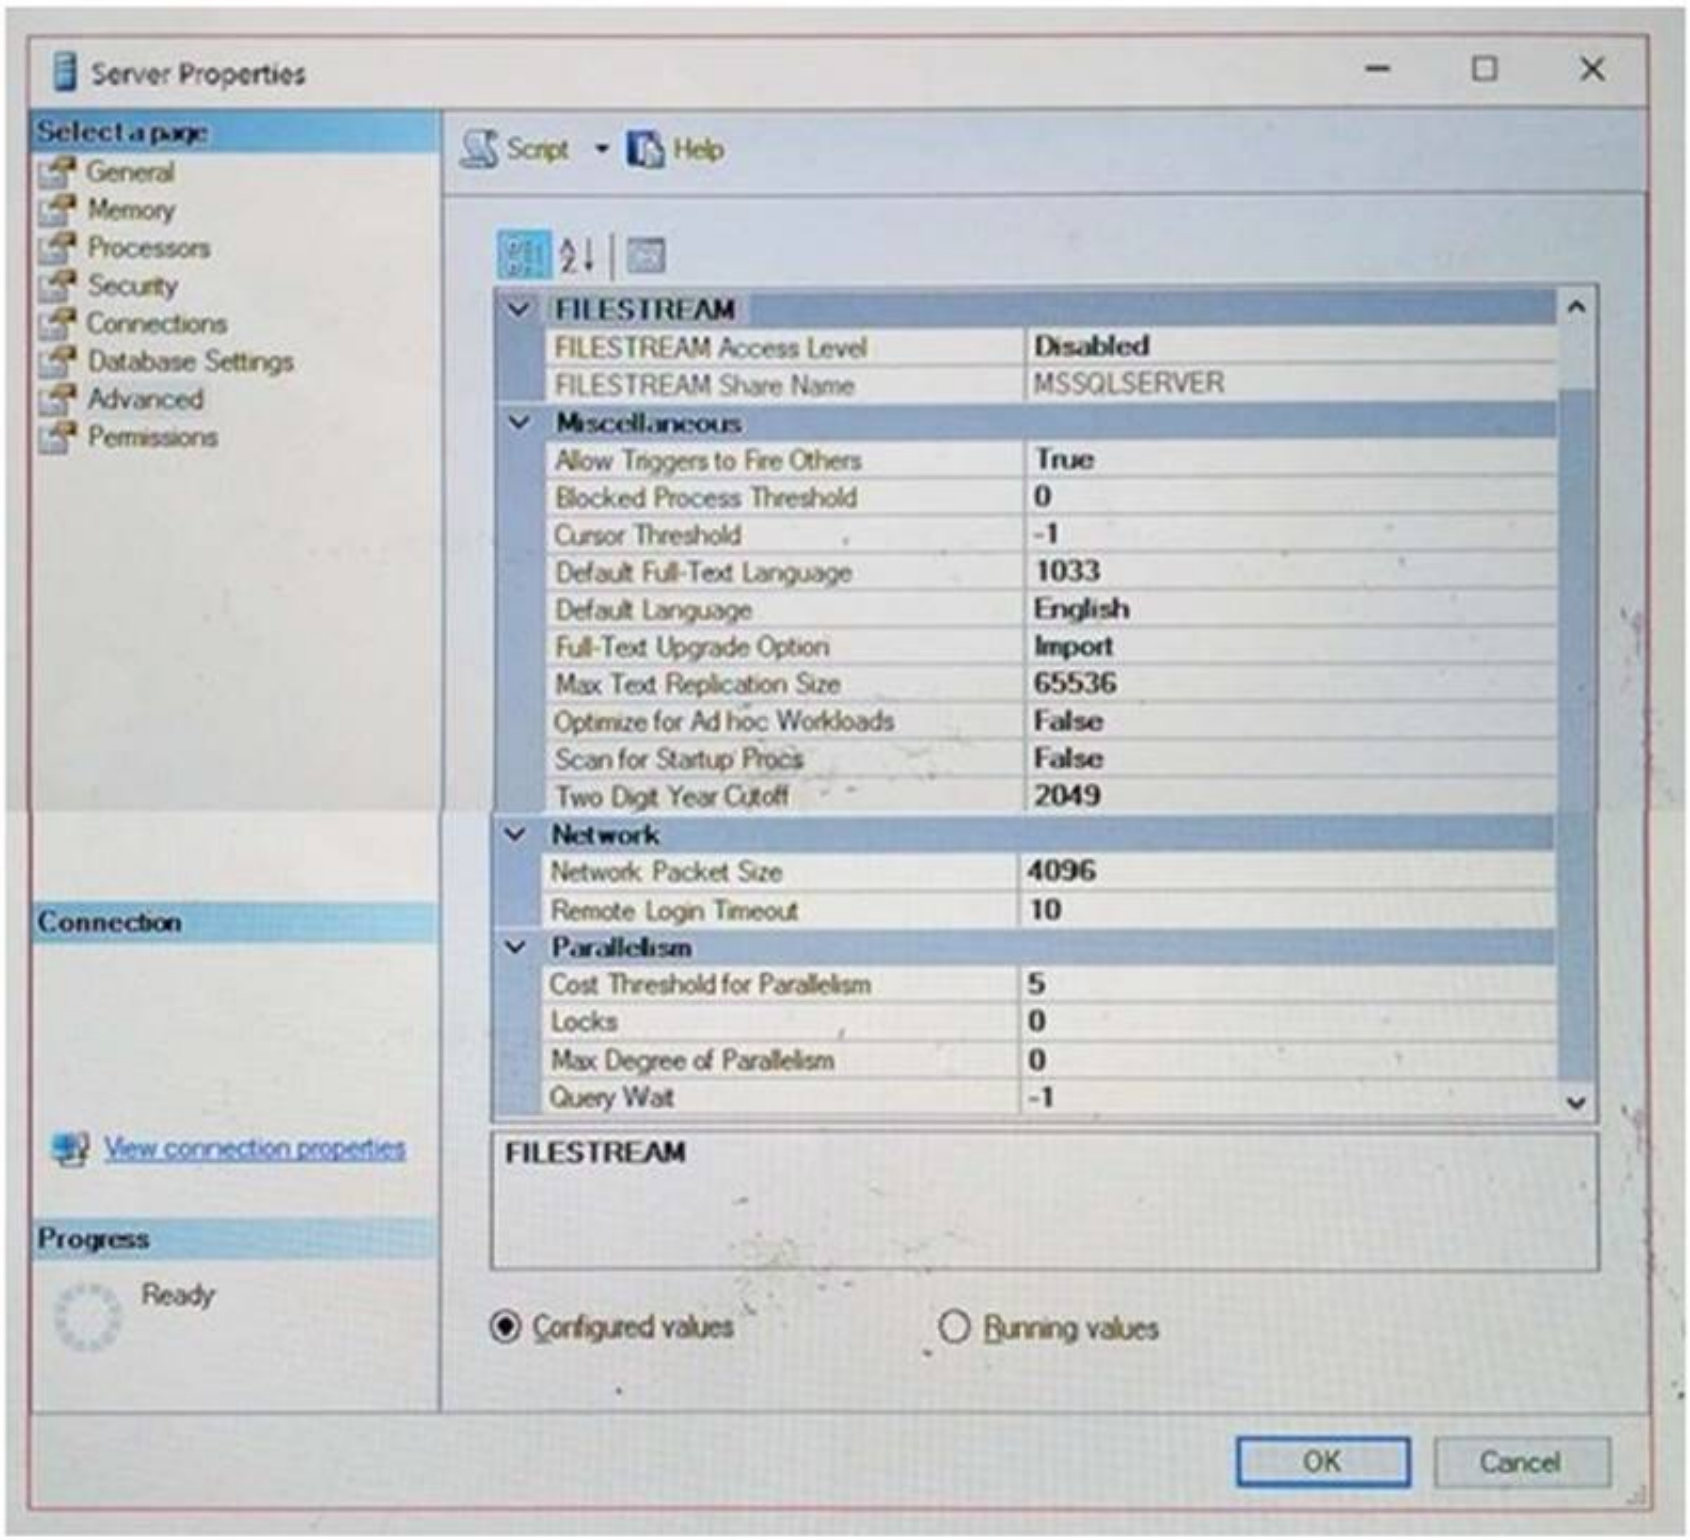

You need to set the size of the log files for the tempdb database on SRV1.
How should you complete the Transact-SQL statement? To answer, select the appropriate Transact-SQL segments in the answer area.
Hot Area:

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The ALTER DATABASE with MODIFY FILE command can make a file size bigger (but not smaller). Example:
ALTER DATABASE AdventureWorks2012 MODIFY FILE
(NAME = test1dat3, SIZE = 200MB); Note: MODIFY FILE
Specifies the file that should be modified. Only one <filespec> property can be changed at a time. NAME must always be specified in the <filespec> to identify the file to be modified. If SIZE is specified, the new size must be larger than the current file size.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/move-a-tdeprotected-database-to-a

NEW QUESTION 43
- (Exam Topic 7)
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_subsytem 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 45
- (Exam Topic 7)
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.

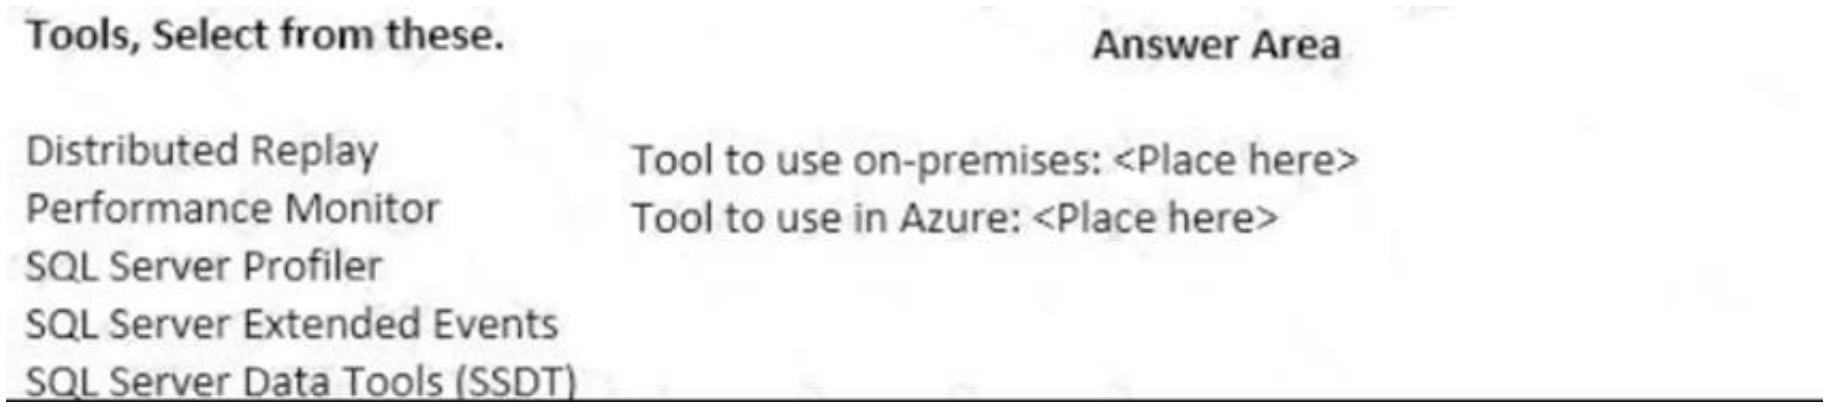

A. Mastered
B. Not Mastered

Answer: A

Explanation:

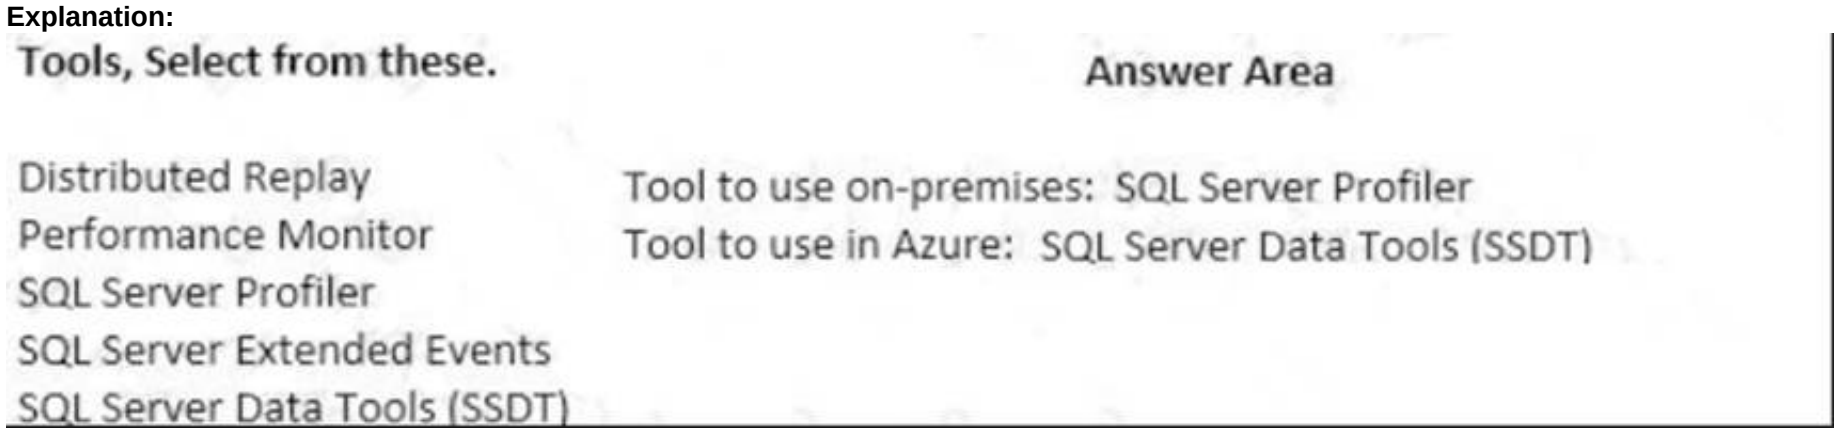

NEW QUESTION 46
- (Exam Topic 7)
You administer a Microsoft SQL Server 2014 database that includes a table named Application.Events. Application.Events contains millions of records about user activity in an application.
Records in Application.Events that are more than 90 days old are purged nightly. 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 49
- (Exam Topic 7)
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 50
- (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

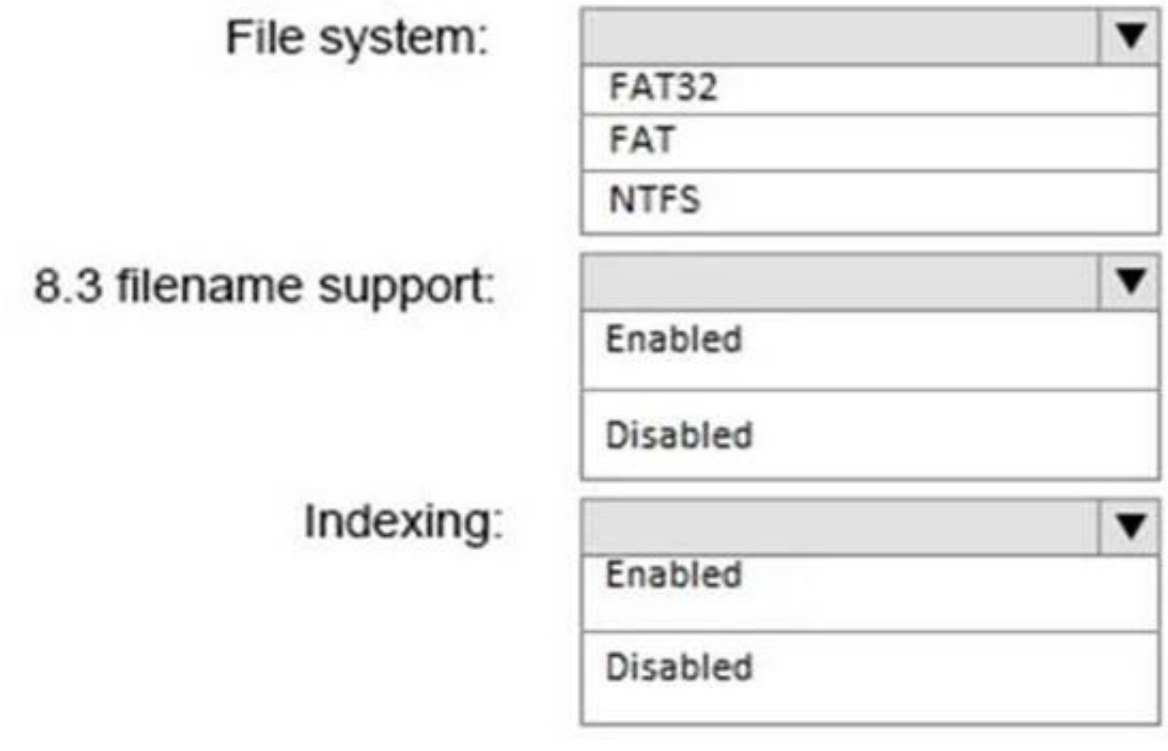

A. Mastered
B. Not Mastered

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 54
- (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 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 59
- (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 enable Transparent Data Encryption (TDE) on the primary instance.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
Azure SQL Database and Data Warehouse offer encryption-at-rest by providing Transparent Data Encryption (TDE) for all data written to disk, including databases, log files and backups. This protects data in case of unauthorized access to hardware. TDE provides a TDE Protector that is used to encrypt the Database Encryption Key (DEK), which in turn is used to encrypt the data. With the TDE and Bring Your Own Key (BYOK) offering currently in preview, customers can take control of the TDE Protector in Azure Key Vault.
Taking advantage of TDE with BYOK for databases that are geo-replicated to maintain high availability requires to configure and test the scenario carefully.
References:
https://azure.microsoft.com/en-us/blog/how-to-configure-azure-sql-database-geo-dr-with-azure-key-vault/

NEW QUESTION 63
- (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 67
- (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 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 68
- (Exam Topic 7)
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

NEW QUESTION 69
- (Exam Topic 7)
You administer a Microsoft SQL Server 2014.
A process that normally runs in less than 10 seconds has been running for more than an hour. You examine the application log and discover that the process is using session ID 60.
You need to find out whether the process is being blocked. Which Transact-SQL statement should you use?

A. EXEC sp_who 60
B. SELECT * FROM sys.dm_exec_sessions WHERE sessionid = 60
C. EXEC sp_helpdb 60
D. DBCC INPUTBUFFER (60)

Answer: A

Explanation:
sp_who provides information about current users, sessions, and processes in an instance of the Microsoft SQL Server Database Engine. The information can be filtered to return only those processes that are not idle, that belong to a specific user, or that belong to a specific session.
Example: Displaying a specific process identified by a session ID EXEC sp_who '10' --specifies the process_id;
References:https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-who-transact-

NEW QUESTION 74
- (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:

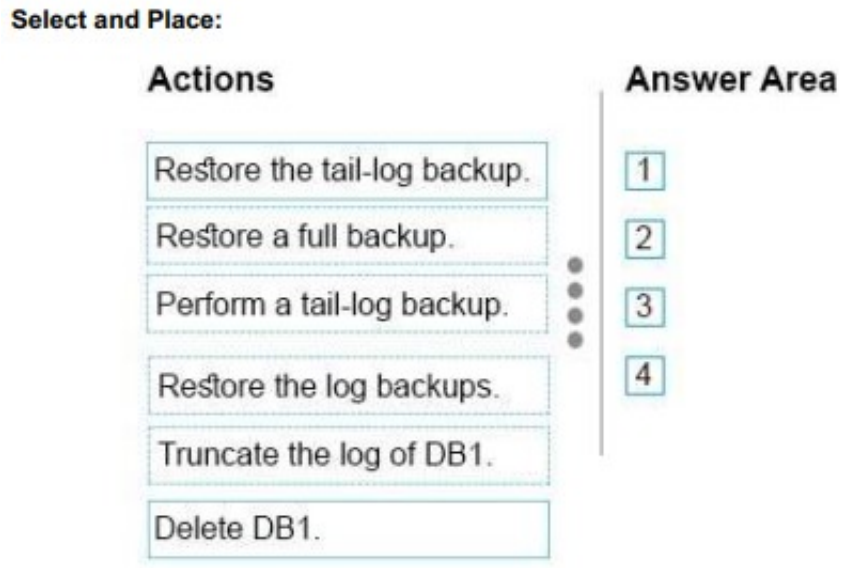

A. Mastered
B. Not Mastered

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 79
- (Exam Topic 7)
You administer a Windows Azure SQL Database database named Human_Resources. The database contains 2 tables named Employees and SalaryDetails. You add two Windows groups as logins for the server:
You need to grant users access according to the following requirements: What should you do?

A. Create a database role called Employees.Add CORP\Employees to the db_datareader rol
B. Add all company employees except HR administrators to the Employees rol
C. Deny SELECT access to the SalaryDetails table to the Employees role.
D. Create a database role called HRAdmins.Add all company employees except HR administrators to the db_datareader rol
E. Add all HR administrators to the HRAdmins rol

F. Grant SELECT access to the SalaryDetails table to the HRAdmins role.Deny SELECT access to the SalaryDetails table to the db_datareader role.
G. Create two database roles: Employees and HRAdmin
H. Add all company employees to the Employees role.Add HR administrators to the HRAdmins rol
I. Grant SELECT access to all tables except SalaryDetails to the Employees rol
J. Grant SELECT access to the SalaryDetails table to the HRAdmins rol
K. Deny SELECT access to the SalaryDetails table to the Employees role.
L. Create a database role called Employees.Add all HR administrators to the db_datareader rol
M. Add all company employees to the Employees rol
N. Grant SELECT access to all tables except the SalaryDetails table to the Employees rol
O. Deny SELECT access to the SalaryDetails table to the Employees 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

NEW QUESTION 81
- (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 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 83
- (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 87
- (Exam Topic 7)
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 close is needed: WHERE o.Rnk = 1

NEW QUESTION 90
- (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 the minimum amount of data is lost. Which recovery model should the database use?

A. NORECOVERY
B. FULL
C. NO_CHECKSUM
D. CHECKSUM
E. Differential
F. BULK_LOGGED
G. STANDBY
H. RESTART
I. SKIP
J. Transaction log
K. DBO ONLY
L. COPY_ONLY
M. SIMPLE
N. CONTINUE AFTER ERROR

Answer: B

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

NEW QUESTION 92
- (Exam Topic 7)
You administer a Microsoft SQL Server 2014 server. You plan to deploy new features to an application. You need to evaluate existing and potential clustered and non-clustered indexes that will improve
performance.
What should you do?

A. Query the sys.dm_db_index_usage_stats DMV.
B. Query the sys.dm_db_missing_index_details DMV.
C. Use the Database Engine Tuning Advisor.
D. Query the sys.dm_db_missing_index_columns DMV.

Answer: C

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

NEW QUESTION 93
- (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 97
- (Exam Topic 7)
You administer a Microsoft SQL Server 2014 instance named SQL2012 that hosts an OLTP database of 1 terabyte in size.
The database is modified by users only from Monday through Friday from 09:00 hours to 17:00 hours. Users modify more than 30 percent of the data in the database during the week.
Backups are performed as shown in the following schedule:

Type	Frequency
Full	Sunday at 20:00 hours
Differential	Monday through Friday at 20:00 hours
Log	Monday through Friday between 08:00 hours and 18:00 hours

The Finance department plans to execute a batch process every Saturday at 09:00 hours. This batch process will take a maximum of 8 hours to complete.
The batch process will update three tables that are 10 GB in size. The batch process will update these tables multiple times.
When the batch process completes, the Finance department runs a report to find out whether the batch process has completed correctly.
You need to ensure that if the Finance department disapproves the batch process, the batch operation can be rolled back in the minimum amount of time. What should you do on Saturday?

A. Perform a differential backup at 08:59 hours.
B. Record the LSN of the transaction log at 08:59 hour
C. Perform a transaction log backup at17:01 hours.
D. Create a database snapshot at 08:59 hours.
E. Record the LSN of the transaction log at 08:59 hour
F. Perform a transaction log backup at08:59 hours.
G. Create a marked transaction in the transaction log at 08:59 hour
H. Perform a transaction log backup at 17:01 hours.
I. Create a marked transaction in the transaction log at 08:59 hour
J. Perform a transaction log backup at 08:59 hours.

Answer: C

Explanation:
References:https://docs.microsoft.com/en-us/sql/relational-databases/databases/database-snapshots-sql-server

NEW QUESTION 99
- (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 tembdb load.
You need to maximize the performance of the tempdb database.
Solution: You use an AB compute-intensive instance and store the tempdb database in Standard storage. Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
For D-series, Dv2-series, and G-series VMs, the temporary drive on these VMs is SSD-based. If your workload makes heavy use of TempDB (such as temporary objects or complex joins), storing TempDB on the D drive could result in higher TempDB throughput and lower TempDB latency.
References:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-sql-performan

NEW QUESTION 102
- (Exam Topic 7)
You have Microsoft SQL Server on a DS-series Microsoft Azure virtual machine. The virtual machine has 28 GB of memory.
You discover the following performance statistics on the server:
The average Page life expectancy is 30.
The server has excessive PAGELATCH_IO waits.
You need to decrease the PAGELATCH_IO waits. What should you do?

A. Enable large-page support.
B. Enable lock pages in memory.
C. Configure buffer pool extensions.
D. Add more tempdb files.

Answer: DExplanation:References:

NEW QUESTION 107
- (Exam Topic 7)

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/

NEW QUESTION 108
- (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 110
- (Exam Topic 7)
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 114
- (Exam Topic 7)
You administer a Microsoft SQL Server 2014 database. The database is currently configured to log ship to a secondary server.
You are preparing to cut over to the secondary server by stopping log-shipping and bringing the secondary database online. You want to perform a tail-log backup.
You need to leave the primary database in a restoring state.
Which option of the BACKUP LOG command should you use?

A. NO_TRUNCATE
B. NORECOVERY

C. STANDBY
D. FORMAT

Answer: B

Explanation:
It is recommended that you take a tail-log backup in the following scenarios:
* If the database is online and you plan to perform a restore operation on the database, begin by backing up the tail of the log. To avoid an error for an online database, you must use the … WITH NORECOVERY option of the BACKUP Transact-SQL statement.
Note: 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.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/tail-log-backups-sql-server

NEW QUESTION 116
- (Exam Topic 7)
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 119
- (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 124
.....

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 FirstTry

* 70-765 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year