DUMPSARENA

Provisioning SQL Databases

Microsoft 70-765

Version Demo

Total Demo Questions: 15

Total Premium Questions: 223

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Implementing SQL in Azure	72
Topic 2, Manage databases and instances	72
Topic 3, Manage Storage	79
Total	223

QUESTION NO: 1

You create a new Microsoft Azure subscription.

You need to create a group of Azure SQL databases that share resources.

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: D

Explanation:

SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single Azure SQL Database server and share a set number of resources (elastic Database Transaction Units (eDTUs)) at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

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

QUESTION NO: 2

You have the databases configured as shown in the following table.

Database name	Location	Platform
DB1	Microsoft Azure	Microsoft Azure SQL Database
DB2	Microsoft Azure	Microsoft SQL Server 2016 on a Microsoft Azure virtual machine
DB3	On-premises	Microsoft SQL Server 2016
DB4	On-premises	Microsoft SQL Server 2014

Which two databases can use the Stretch Database feature? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. DB1

B. DB2

C. DB3

D. DB4

ANSWER: A C

Explanation:

References: https://docs.microsoft.com/en-us/sql/sql-server/stretch-database/enable-stretch-database-for-a-database?view=sql-server-2017

QUESTION NO: 3

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

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

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

Which storage option should you use?

A. Premium P10 disk storage

B. Premium P20 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: E

Explanation:

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

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

QUESTION NO: 4 - (HOTSPOT)

HOTSPOT

Background

You are the database administrator for Contoso, Ltd. The company has 200 offices around the world. The company has corporate executives that are located in offices in London, New York, Toronto, Sydney, and Tokyo.

Contoso, Ltd. has a Microsoft Azure SQL Database environment. You plan to deploy a new Azure SQL Database to support a variety of mobile applications and public websites.

The company is deploying a multi-tenant environment. The environment will host Azure SQL Database instances. The company plans to make the instances available to internal departments and partner companies. Contoso is in the final stages of setting up networking and communications for the environment.

Existing Contoso and Customer instances need to be migrated to Azure virtual machines (VM) according to the following requirements:

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

The company plans to deploy a new order entry application and a new business intelligence and analysis application. Each application will be supported by a new database. Contoso creates a new Azure SQL database named Reporting. The database will be used to support the company's financial reporting requirements. You associate the database with the Contoso Azure Active Directory domain.

Each location database for the data entry application may have an unpredictable amount of activity. Data must be replicated to secondary databases in Azure datacenters in different regions.

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

Objects

Parameter	Name
Logical Server	contososrv
Resource Group	contosodbrg

Database

The contosodb1 database must support the following requirements:

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

Application

For the business intelligence application, corporate executives must be able to view all data in near real-time with low network latency.

Contoso has the following security, networking, and communications requirements:

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

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.

Hot Area:

Answer Area

▼
New-AzureSqlDatabase
New-AzureRmSqlDatabase
Set-AzureRmSqlDatabase

- ResourceGroupName "contosodbrg"

- ServerName "contososrv"

-DatabaseName "contosodb1"

- Edition

▼
Basic
Standard
Premium

-RequestedServiceObjectName S2

ANSWER:

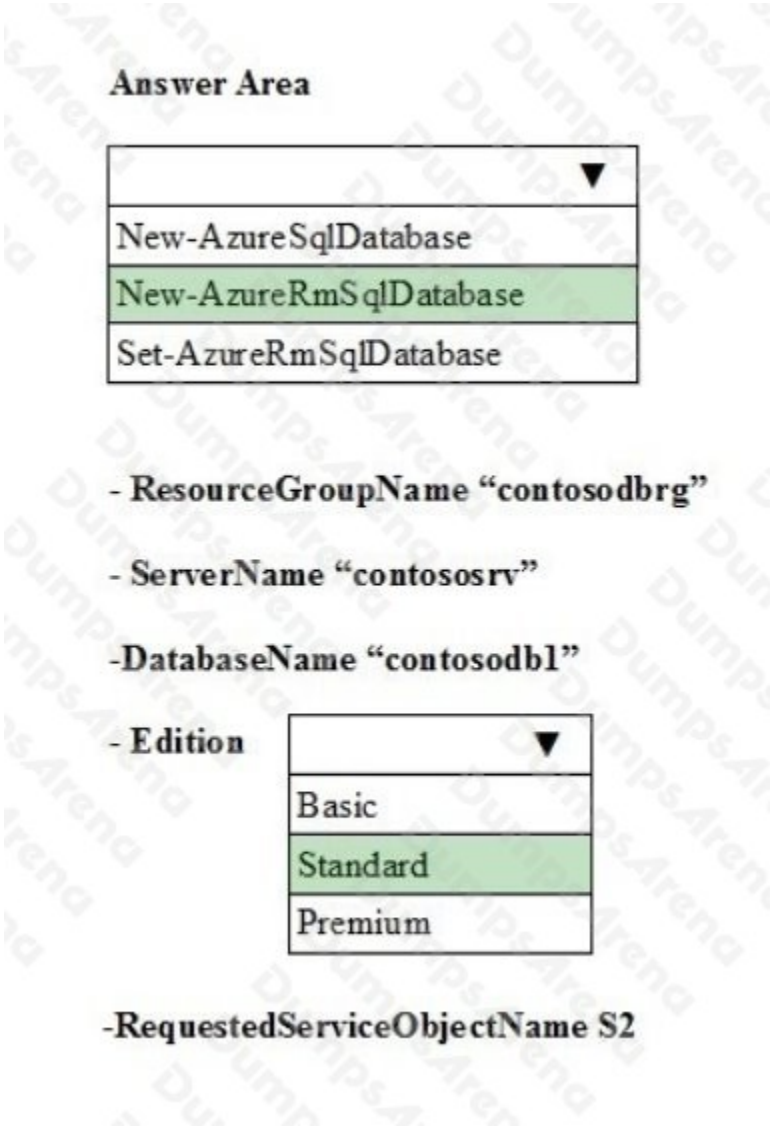

Explanation:

Box 1: New-AzureRmSqlDatabase

New-AzureRmSqlDatabase creates a database or an elastic database.

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

Incorrect Answers:

Not New-AzureSqlDatabase: New-AzureSqlDatabase cannot used be used for Resource Groups.

Box 2: Standard

The maximum database size for the Standard edition is 250 GB, while 200 GB is required

The maximum concurrents sessions for the Standard edition is 1200 for S2 and 2400 for S3, while 1000 concurrent sessions is required.

From the scenario: The contosodb1 database must support the following requirements:

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

Incorrect Answers:

Not Basic: The maximum database size for the Basic edition is only 2 GB, while 200 GB is required.

Not Premium: Standard Edition meets the requirements, and Premium would be a more expensive solution.

References: https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tiers

QUESTION NO: 5

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

You discover that queries take longer than expected to execute. You also find that the top wait type for the SQL Server instance is ASYNC_NETWORK_IO.

You need to resolve the performance issue.

What should you do?

A. Increase the number of processor cores for the virtual machine

B. Create any missing indexes

C. Modify the application code to remove row-by-row processing

D. Allocate more memory to the SQL Server instance

ANSWER: B

Explanation:

References:

https://www.sqlservercentral.com/articles/speeding-up-data-access-part-1-missing-indexes-1

QUESTION NO: 6

You plan to deploy Microsoft SQL Server on a Microsoft Azure virtual machine. The virtual machine will have two databases. Each database will reside on a separate VHD and will be between 600 and 800 GB.

Each database will have the I/O requirements shown in the following table.

Database name	Maximum IOPS
DB1	4,000
DB2	1,200

You are evaluating whether to use the P30 storage disk type.

What is the minimum number of disks required for each database when using P30 storage disk type? (Select two.)

A. DB1: 0

B. DB1: 1

C. DB1: 2

D. DB1: 3

E. DB1: 4

F. DB2: 0

G. DB2: 1

H. DB2: 2

I. DB2: 3

J. DB2: 4

ANSWER: C G

Explanation:

P30 stats: Disk size is 1024 GB (1 TB), IOPS per disk is 5000.

Recommendation: Use a minimum of 2 P30 disks (1 for log files and 1 for data files and TempDB; or stripe two or more disks and store all files in a single volume).

References: https://docs.microsoft.com/en-us/azure/virtual-machines/windows/premium-storage#scalability-and-performance-targets

QUESTION NO: 7

You have a database named DB1 that is 3 TB. DB1 contains a fact table that is 1.2 TB.

You load 200 GB of new data to the fact table from a line-of-business application.

Users of DB1 notice that reports render more slowly after you loaded the data.

What are two possible causes of the performance issue? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. transaction log growth

B. out-of-date statistics

C. page corruption

D. index fragmentation

ANSWER: B D

Explanation:

Reference: https://www.sqlshack.com/gathering-sql-server-indexes-statistics-and-usage-information/

QUESTION NO: 8

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

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

You are migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 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 30 storage accounts that each has one container. You create a VHD in each container.

Does this meet the goal?

A. Yes

B. No

ANSWER: A

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/

QUESTION NO: 9 - (DRAG DROP)

DRAG DROP

You upgrade a database named DB1 to Microsoft SQL Server 2016.

Users report that performance for several queries is degraded. You determine that the query optimizer is incorrectly estimating the number of rows that the queries will return.

You need to resolve the performance issues.

How should you complete the Transact-SQL statement? To answer, drag the 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.

Select and Place:

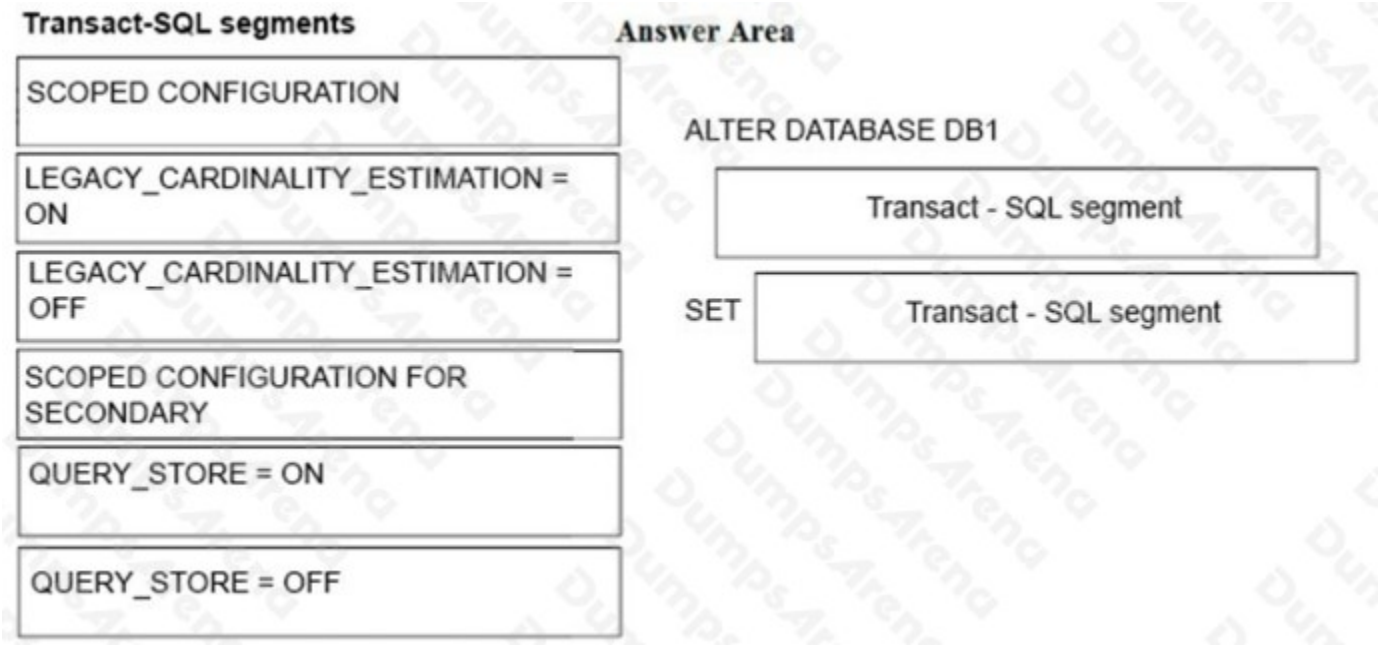

ANSWER:

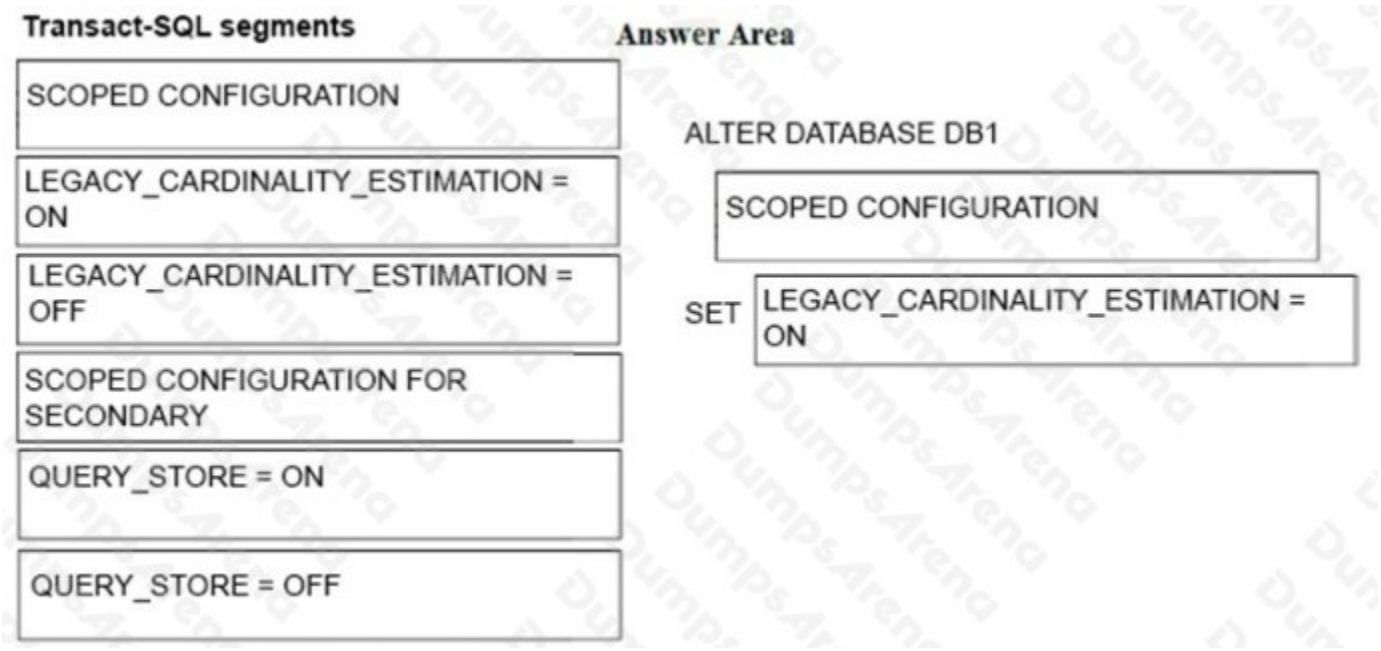

Explanation:

References:

https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql?view=sql-server-2017

QUESTION NO: 10

You have just completed a new Microsoft SQL Server installation.

You need to configure a new SQL Server Agent alert to send an email to the DBA team for severity 20 errors. Which three actions should you perform? Each correct answer presents part of the solution. (Choose three.)

A. Set up SQL Mail.

B. Define an operator.

C. Configure a credential object.

D. Define the alert settings.

E. Configure a proxy.

F. Set up an External Events collector.

G. Set up Database Mail.

ANSWER: B D G

Explanation:

Reference: https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-database-mail?view=sql-server-2017

QUESTION NO: 11

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

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

You deploy a new on-premises server and restore all databases to the new server.

You need to resume Stretch Database operations to Azure.

Which statements should you execute?

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

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

C.
```
sp_rda_deauthorize_db
GO
ALTER TABLE tb1
     SET ( REMOTE_DATA_ARCHIVE( MIGRATION_STATE = PAUSE ) ) ;
          FILTER PREDICATE = db
GO
EXEC sp_rda_reauthorize_db @credential = <credential>;
GO
```

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

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

Use ALTER TABLE, not CREATE TABLE.

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

Syntax:

::=

{

SET (

REMOTE_DATA_ARCHIVE

{

= ON ( )

| = OFF_WITHOUT_DATA_RECOVERY (MIGRATION_STATE = PAUSED)

| ([, ...n])

}

)

}

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

QUESTION NO: 12 - (DRAG DROP)

DRAG DROP

You are responsible for deploying and maintaining the Microsoft SQL Server virtual machines (VMs) in Microsoft Azure. You use Azure Resource Group templates to deploy and modify the VMs.

With the exception of their names, all deployed VMs must be identical. You may need to add additional data disks to the VMs in the future.

You need to streamline the template creation and deployment process as much as possible.

For each requirement, what should you do? To answer, drag the appropriate actions or features to the correct locations. Each action or feature 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.

Select and Place:

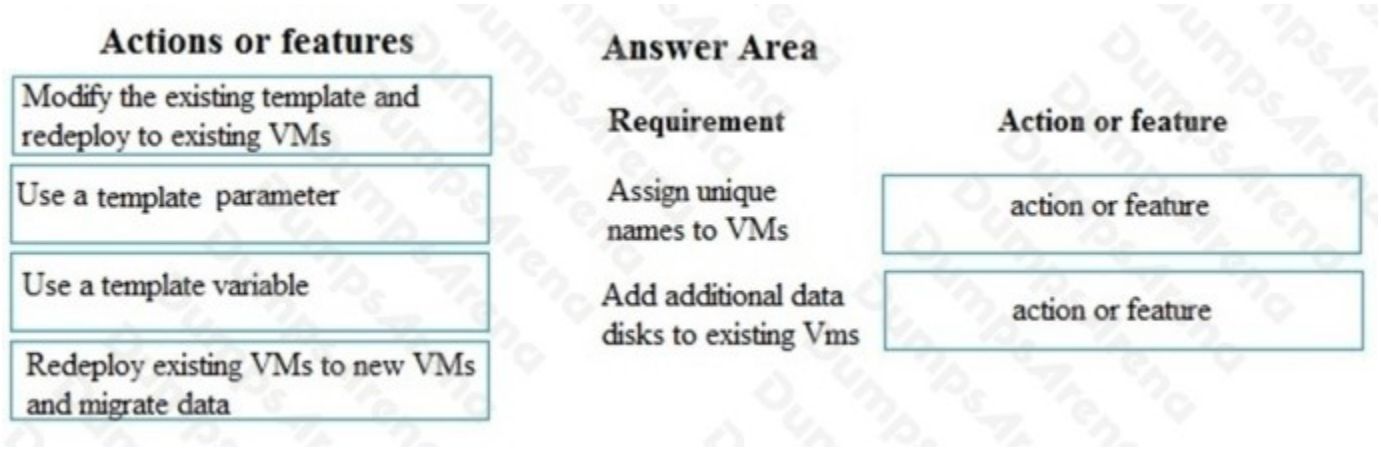

ANSWER:

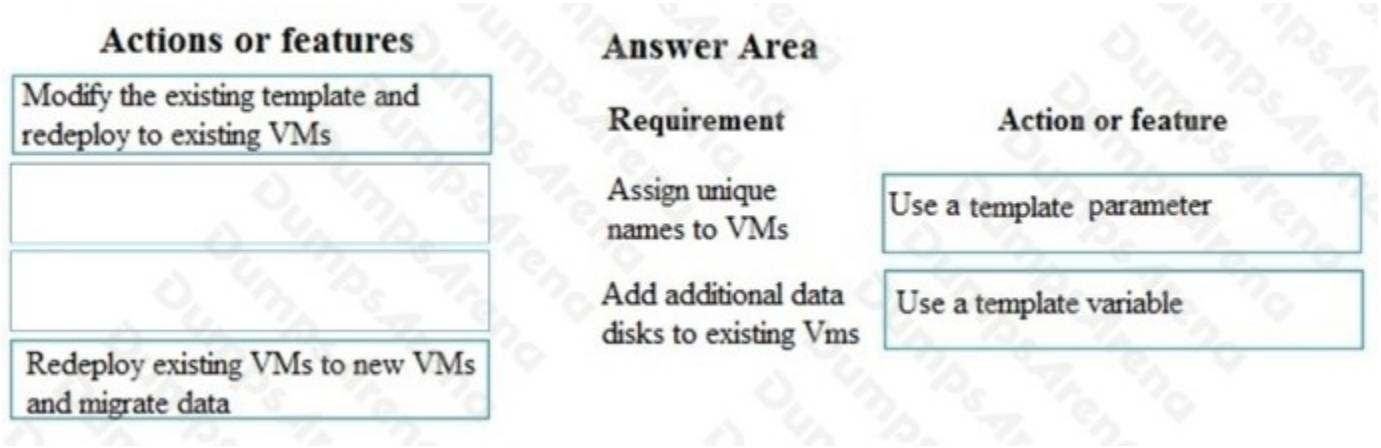

Explanation:

References: https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-manager-templates-parameters
https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-manager-templates-variables

QUESTION NO: 13

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

You discover that index1 is corrupt.

You need to repair index1.

Which statement should you execute?

A. DBCC CHECKDB ('db1', REPAIR_FAST)

B. ALTER INDEX indx1 ON table1 REBUILD WITH (ONLINE-ON)

C. ALTER INDEX index1 ON table1 REORGANIZE

D. DBCC CHECKDB ('db1', DATA_PURITY)

ANSWER: B

Explanation:

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

Incorrect Answers:

A: REPAIR_FAST maintains syntax for backward compatibility only. No repair actions are performed.

D: DATA_PURITY causes DBCC CHECKDB to check the database for column values that are not valid or out-of-range.

References: https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-index-transact-sql?view=sql-server-2017
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql?view=sql-server-2017

QUESTION NO: 14

You deploy a new Microsoft Azure SQL Database instance to support a variety of mobile applications and public websites.

You have the following requirements:

- You must be able to replicate the database to Azure datacenters in other regions by using geo-replication.
- Clients must only connect to the Azure SQL Database by using contained database users.
- You must be able to move the database to other servers in the future.
- Only applications in the 23.96.52.0-23.96.52.255 IP range are permitted to access the database.
- The firewall settings for the instance must not provide access to any other Azure services.

You need to configure the firewall settings for the environment.

Which three actions should you perform? Each correct answer presents part of the solution.

A. In the Azure portal, enable Allow access to Azure services for the Azure SQL Database server

B. Run the following Transact-SQL statement:
EXECUTE sp_set_database_firewall_rule N'Applications', '23.96.52.0', '23.96.52.255'

C. Run the following Transact-SQL statement:
EXECUTE sp_set_firewall_rule N'Allow Windows Azure', '0.0.0.0', '0.0.0.0'

D. Run the following Transact-SQL statement:
EXECUTE sp_set_firewall_rule N'Applications', '23.96.52.0', '23.96.52.255'

E. Run the following Transact-SQL statement:
EXECUTE sp_set_database_firewall_rule N'Allow Windows Azure', '0.0.0.0', '0.0.0.0'

F. In the Azure portal, disable Allow access to Azure services for the Azure SQL Database server

ANSWER: B C F

Explanation:

References:

https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-set-database-firewall-rule-azure-sql-database?view=azuresqldb-current https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-set-firewall-rule-azure-sql-database?view=azuresqldb-current https://docs.microsoft.com/en-us/azure/sql-database/sql-database-vnet-service-endpoint-rule-overview

QUESTION NO: 15 - (HOTSPOT)

HOTSPOT

You have an on-premises database.

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

You move the database files to Azure.

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

How should you complete the statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

ANSWER:

Explanation:

Reference: https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-sql-server-transact-sql