DUMPSARENA

Administering Relational Databases on Microsoft Azure

Microsoft DP-300

Version Demo

Total Demo Questions: 15

Total Premium Questions: 343

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 2, New Update	146
Topic 3, Case Study 1	2
Topic 4, Case Study 2	3
Topic 5, Case Study 3	2
Topic 6, Case Study 4	3
Topic 7, Case Study 5	2
Topic 8, Case Study 6	2
Topic 9, Case Study 7	3
Topic 10, Mixed Questions	180
Total	343

QUESTION NO: 1 - (DRAG DROP)

You have a database named db1.

The log for db1 contains the following entry.

```
Date 10/5/2021 10:57:08 AM
Log SQL Server (Current - 10/5/2021 11:26:00 AM)

Source spid1595

Message
The transaction log for database 'db1' is full due to 'AVAILABILITY_REPLICA'
```

You need to ensure That db1 can process transactions.

Actions

- Add db1 back to the availability group.
- Shrink db1.
- Shrink the transaction log file.
- Remove db1 from the availability group.
- Back up the transaction log file.

ANSWER:

Explanation:

QUESTION NO: 2 - (DRAG DROP)

DRAG DROP

You have an Azure Active Directory (Azure AD) tenant named contoso.com that contains a user named user1@contoso.com and an Azure SQL managed instance named SQLMI1.

You need to ensure that user1@contoso.com can create logins in SQLMI1 that map to Azure AD service principals.

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

Select and Place:

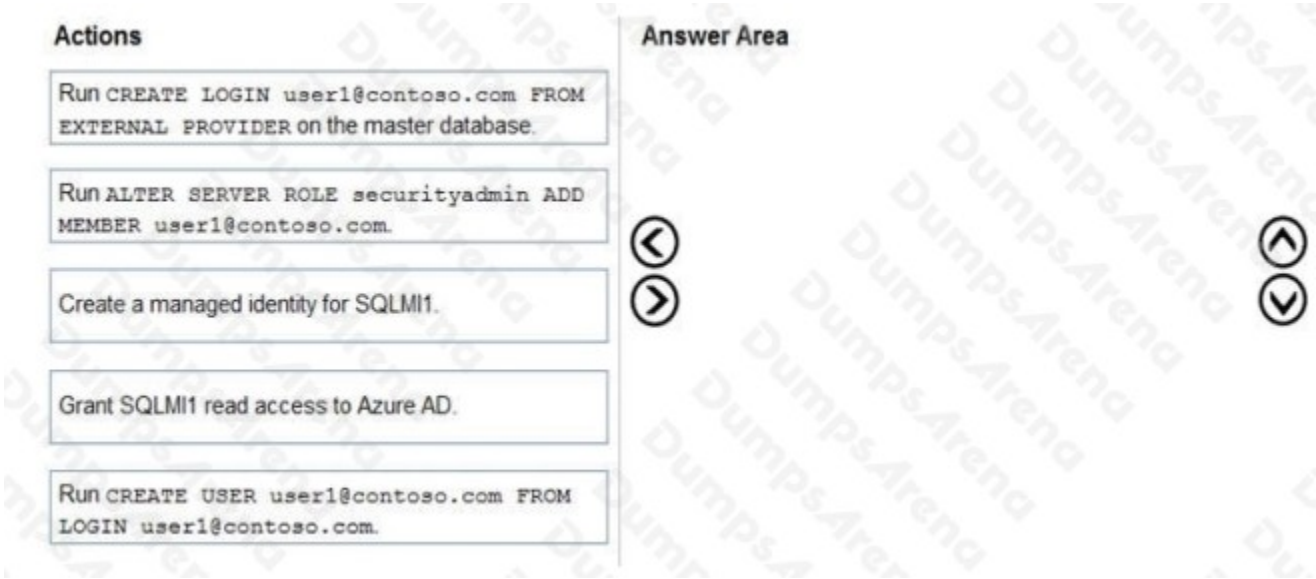

ANSWER:

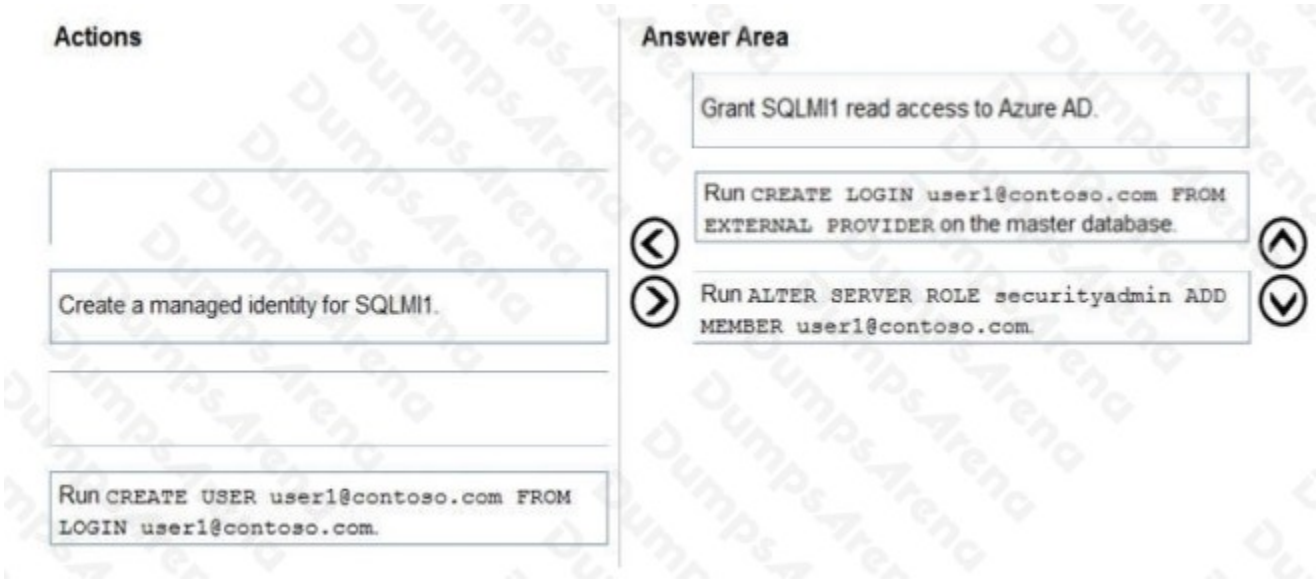

Explanation:

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/aad-security-configure-tutorial

QUESTION NO: 3

You have an Azure SQL database named db1 on a server named server1.

The Intelligent Insights diagnostics log identifies that several tables are missing indexes.

You need to ensure that indexes are created for the tables.

What should you do?

A. Run the DBCC SQLPERF command.

B. Run the DBCC DBREINDEX command.

C. Modify the automatic tuning settings for db1.

D. Modify the Query Store settings for db1.

ANSWER: C

Explanation:

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-overview

QUESTION NO: 4

You have a new Azure subscription.

You create an Azure SQL Database instance named DB1 on an Azure SQL Database server named Server1.

You need to ensure that users can connect to DB1 in the event of an Azure regional outage. In the event of an outage, applications that connect to DB1 must be able to connect without having to update the connection strings.

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

NOTE: Each correct selection is worth one point.

A. From the properties of DB1. configure geo-replication.

B. From the properties of Server1 add a failover group.

C. Create a new Azure SQL Database server named Server2.

D. From the properties of Server1 configure retention for DB1

E. Create a new Azure SQL Database instance named DB2.

ANSWER: B C

Explanation:

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview?tabs=azure-powershell#best-practices-for-sql-database

https://docs.microsoft.com/en-us/azure/azure-sql/database/failover-group-add-single-database-tutorial?tabs=azure-portal

QUESTION NO: 5

You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.

You need to migrate the databases hosted on SQL 1 to Azure. The solution must meet the following requirements:

The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.

Administrative effort must be minimized.

What should you use to host the databases?

A. a single Azure SQL database

B. an Azure SQL Database elastic pool

C. SQL Server on Azure Virtual Machines

D. Azure SQL Managed Instance

ANSWER: D

QUESTION NO: 6

You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.

You need to migrate the databases hosted on SQL 1 to Azure. The solution must meet the following requirements:

The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.

Administrative effort must be minimized.

What should you use to host the databases?

A. a single Azure SQL database

B. an Azure SQL Database elastic pool

C. SQL Server on Azure Virtual Machines

D. Azure SQL Managed Instance

ANSWER: D

QUESTION NO: 7

You are planning disaster recovery for the failover group of an Azure SQL Database managed instance.

Your company's SLA requires that the database in the failover group become available as quickly as possible if a major outage occurs.

You set the Read/Write failover policy to Automatic.

What are two results of the configuration? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. In the event of a datacenter or Azure regional outage, the databases will fail over automatically.

B. In the event of an outage, the databases in the primary instance will fail over immediately.

C. In the event of an outage, you can selectively fail over individual databases.

D. In the event of an outage, you can set a different grace period to fail over each database.

E. In the event of an outage, the minimum delay for the databases to fail over in the primary instance will be one hour.

ANSWER: A E

Explanation:

A: Auto-failover groups allow you to manage replication and failover of a group of databases on a server or all databases in a managed instance to another region.

E: Because verification of the scale of the outage and how quickly it can be mitigated involves human actions by the operations team, the grace period cannot be set below one hour. This limitation applies to all databases in the failover group regardless of their data synchronization state.

Incorrect Answers:

C: individual SQL Managed Instance databases cannot be added to or removed from a failover group.

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview

QUESTION NO: 8

You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1.

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1. The on-premises network connects to VNET1 by using an ExpressRoute connection.

You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service.

You need to configure VNET1 to support the migration.

What should you do?

A. Configure service endpoints.

B. Configure virtual network peering.

C. Deploy an Azure firewall.

D. Configure network security groups (NSGs).

ANSWER: A

Explanation:

Reference:

https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-to-managed-instance

QUESTION NO: 9

You have an Azure SQL database named DB1. DB1 has a table named Table1 that contains the following columns.

Name	Type
Column1	Ntext
Column2	Geometry
Column3	Image
Column4	Varchar
Column5	Datetime2

You plan to enable Always Encrypted for Table1.

Which two columns support encryption? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point

A. Column1

B. Column2

C. Column3

D. Column4

E. Column5

ANSWER: A D

QUESTION NO: 10

You are designing a security model for an Azure Synapse Analytics dedicated SQL pool that will support multiple companies.

You need to ensure that users from each company can view only the data of their respective company.

Which two objects should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a column encryption key

B. asymmetric keys

C. a function

D. a custom role-based access control (RBAC) role

E. a security policy

ANSWER: D E

Explanation:

Azure RBAC is used to manage who can create, update, or delete the Synapse workspace and its SQL pools, Apache Spark pools, and Integration runtimes.

Define and implement network security configurations for resources related to your dedicated SQL pool with Azure Policy.

Reference:

https://docs.microsoft.com/en-us/azure/synapse-analytics/security/synapse-workspace-synapse-rbac

https://docs.microsoft.com/en-us/security/benchmark/azure/baselines/synapse-analytics-security-baseline

QUESTION NO: 11 - (HOTSPOT)

HOTSPOT

You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.

You execute the Transact-SQL query shown in the following exhibit.

```
SELECT
   payment_type,
   SUM(fare_amount) AS fare_total
FROM OPENROWSET(
         BULK 'csv/busfare/tripdata_2020*.csv',
         DATA_SOURCE = 'BusData',
         FORMAT = 'CSV', PARSER_VERSION = '2.0',
         FIRSTROW = 2
   )
   WITH (
         payment_type INT 10,
         fare_amount FLOAT 11
   ) AS nyc
GROUP BY payment_type
ORDER BY payment_type;
```

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

Hot Area:

Answer Area

The query results include only **[answer choice]** in the csv/busfare folder.

- CSV files in the tripdata_2020 subfolder
- files that have files names beginning with "tripdata_2020"
- CSV files that have file names containing "tripdata_202"
- CSV files that have file named beginning with "tripdata_2020"

The query assumes that the first row in a CSV file is **[answer choice]** row.

- a header
- a data
- an empty

ANSWER:

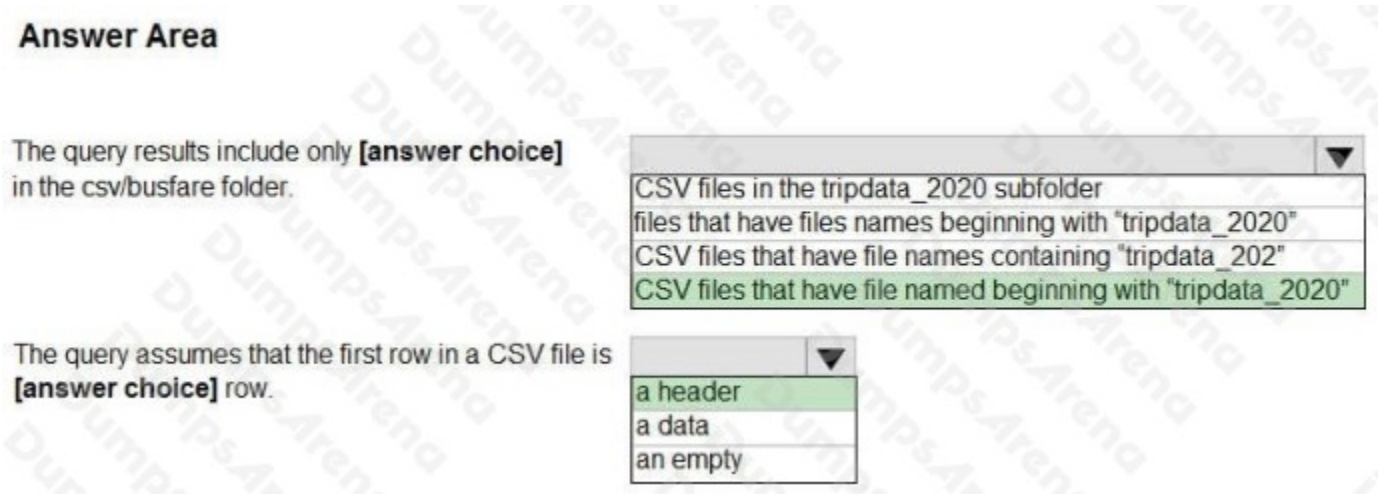

Explanation:

Box 1: CSV files that have file named beginning with "tripdata_2020"

Box 2: a header

FIRSTROW = 'first_row'

Specifies the number of the first row to load. The default is 1 and indicates the first row in the specified data file. The row numbers are determined by counting the row terminators. FIRSTROW is 1-based.

Reference: https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-openrowset

QUESTION NO: 12 - (SIMULATION)

You have two on-premises servers that run Windows Server 2019 and host a Microsoft SQL Server 2017 Always On availability group named AG1. AG1 contains a single database named DB1.

You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Linux.

You need to migrate DB1 to a SQL Server 2019 instance on VM1. The solution must minimize the downtime of DB1 during the migration.

What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

ANSWER: seetheanswerbelowinexplanation.

Explanation:

Answer as in image below.

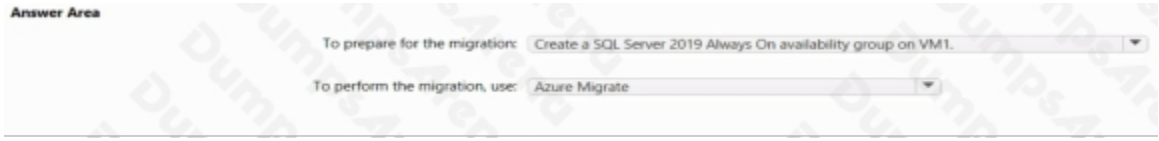

QUESTION NO: 13

A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

```
CREATE TABLE dbo.Employee
(
    EmployeeID      INT IDENTITY(1,1) PRIMARY KEY CLUSTERED NOT NULL,
    FirstName       VARCHAR(100) NOT NULL,
    LastName        VARCHAR(100) NOT NULL,
    Title           VARCHAR(100) NULL,
    LastHireDate    DATETIME NULL,
    StreetAddress1  VARCHAR(500) NOT NULL,
    StreetAddress2  VARCHAR(500) NOT NULL,
    StreetAddress3  VARCHAR(500) NOT NULL,
    City            VARCHAR(200) NOT NULL,
    StateName       VARCHAR(20) NOT NULL,
    Salary          VARCHAR(20) NULL,
    PhoneNumber     VARCHAR(20) NOT NULL
)
```

You need to recommend changes to the data types to reduce storage and improve performance.

Which two actions should you recommend? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Change Salary to the money data type.

B. Change PhoneNumber to the float data type.

C. Change LastHireDate to the datetime2(7) data type.

D. Change PhoneNumber to the bigint data type.

E. Change LastHireDate to the date data type.

ANSWER: A E

QUESTION NO: 14 - (HOTSPOT)

You have an Azure SQL managed instance.

You need to restore a database named DB1 by using Transact-SQL.

Which command should you run? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

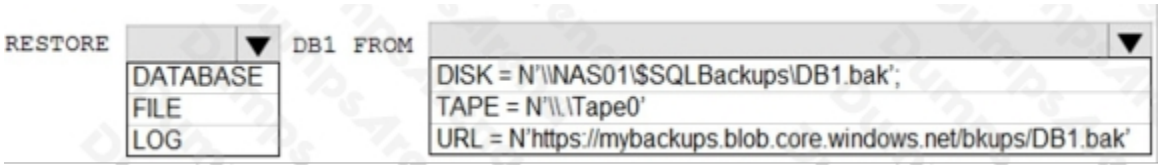

ANSWER:

Explanation:

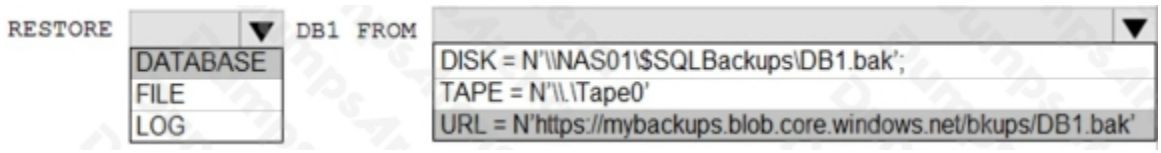

QUESTION NO: 15

You have an Azure Synapse Analytics workspace named WS1 that contains an Apache Spark pool named Pool1.

You plan to create a database named DB1 in Pool1.

You need to ensure that when tables are created in DB1, the tables are available automatically as external tables to the built-in serverless SQL pool.

Which format should you use for the tables in DB1?

A. JSON

B. CSV

C. Parquet

D. ORC

ANSWER: C

Explanation:

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.

For each Spark external table based on Parquet and located in Azure Storage, an external table is created in a serverless SQL pool database. As such, you can shut down your Spark pools and still query Spark external tables from serverless SQL pool.

Reference:

https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-storage-files-spark-tables