

# DUMPS ARENA

## Administering Relational Databases on Microsoft Azure (beta)

Microsoft DP-300

Version Demo

Total Demo Questions: 15

Total Premium Questions: 343

Buy Premium PDF

<https://dumpsarena.com>

[sales@dumpsarena.com](mailto:sales@dumpsarena.com)

dumpsarena.com

## 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**

Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10-Mbps internet connection.

You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the database.

What should you use?

- A. Azure Migrate
- B. Data Migration Assistant (DMA)
- C. Azure Data BOX
- D. Azure Database Migration Service

**ANSWER: C****Explanation:**

<https://www.techtarget.com/searchitoperations/tip/Easily-transfer-VMs-to-the-cloud-with-Microsoft-Azure-Migrate>

**QUESTION NO: 2**

You have SQL Server on an Azure virtual machine that contains a database named DB1.

You have an application that queries DB1 to generate a sales report.

You need to see the parameter values from the last time the query was executed.

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

NOTE: Each correct selection is worth one point.

- A. Enable Last\_Query\_Plan\_Stats in the master database
- B. Enable Lightweight\_Query\_Profiling in DB1
- C. Enable Last\_Query\_Plan\_Stats in DB1
- D. Enable Lightweight\_Query\_Profiling in the master database
- E. Enable PARAMETER\_SNIFFING in DB1

**ANSWER: A C****Explanation:**

Last\_Query\_Plan\_Stats allows you to enable or disable collection of the last query plan statistics (equivalent to an actual execution plan) in sys.dm\_exec\_query\_plan\_stats.

Lightweight profiling can be disabled at the database level using the LIGHTWEIGHT\_QUERY\_PROFILING database scoped configuration: ALTER DATABASE SCOPED CONFIGURATION SET LIGHTWEIGHT\_QUERY\_PROFILING = OFF;.

Incorrect Answers:

E: Parameter sensitivity, also known as "parameter sniffing", refers to a process whereby SQL Server "sniffs" the current parameter values during compilation or recompilation, and passes it along to the Query Optimizer so that they can be used to generate potentially more efficient query execution plans.

Parameter values are sniffed during compilation or recompilation for the following types of batches:

- Stored procedures
- Queries submitted via sp\_executesql
- Prepared queries

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure>

### QUESTION NO: 3

Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10-Mbps internet connection.

You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the database.

What should you use?

- A. Azure Migrate
- B. Data Migration Assistant (DMA)
- C. Azure Data BOX
- D. Azure Database Migration Service

**ANSWER: D**

**Explanation:**

<https://www.techtarget.com/searchitoperations/tip/Easily-transfer-VMs-to-the-cloud-with-Microsoft-Azure-Migrate>

### QUESTION NO: 4

You have an Azure SQL database named DB1 that contains a nonclustered index named index1.

End users report slow queries when they use index1.

You need to identify the operations that are being performed on the index.

Which dynamic management view should you use?

- A. `sys.dm_exec_query_plan_stats`
- B. `sys.dm_db_index_physical_stats`
- C. `sys.dm_db_index_operational_stats`
- D. `sys.dm_db_index_usage_stats`

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

**ANSWER: D**

#### QUESTION NO: 5 - (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 Area

To prepare for the migration:

To perform the migration, use:

**ANSWER: seetheanswerbelowinexplanation.**

#### Explanation:

Answer as in image below.

Answer Area

To prepare for the migration: Create a SQL Server 2019 Always On availability group on VM1.

To perform the migration, use: Azure Migrate

**QUESTION NO: 6**

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 have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1.

You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1.

You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.

You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.

Solution: In an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.

Does this meet the goal?

A. Yes

B. No

**ANSWER: B****Explanation:**

Instead use a serverless SQL pool to create an external table with the extra column.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables>

**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 a SQL pool in Azure Synapse that contains a table named `dbo.Customers`. The table contains a column name `Email`.

You need to prevent nonadministrative users from seeing the full email addresses in the `Email` column. The users must see values in a format of `aXXX@XXXX.com` instead.

What should you do?

- A.** From the Azure portal, set a mask on the `Email` column.
- B.** From the Azure portal, set a sensitivity classification of Confidential for the `Email` column.
- C.** From Microsoft SQL Server Management Studio, set an email mask on the `Email` column.
- D.** From Microsoft SQL Server Management Studio, grant the `SELECT` permission to the users for all the columns in the `dbo.Customers` table except `Email`.

**ANSWER: B**

**Explanation:**

The Email masking method, which exposes the first letter and replaces the domain with XXX.com using a constant string prefix in the form of an email address.

Example: aXX@XXXX.com

**QUESTION NO: 9**

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: 10**

You have 40 Azure SQL databases, each for a different customer. All the databases reside on the same Azure SQL Database server.

You need to ensure that each customer can only connect to and access their respective database.

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

NOTE: Each correct selection is worth one point.

- A. Implement row-level security (RLS).
- B. Create users in each database.
- C. Configure the database firewall.



- D. Configure the server firewall.
- E. Create logins in the master database.
- F. Implement Always Encrypted.

**ANSWER: B C**

**Explanation:**

Manage database access by adding users to the database, or allowing user access with secure connection strings.

Database-level firewall rules only apply to individual databases.

Incorrect Answers:

B: Server-level IP firewall rules apply to all databases within the same server.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/secure-database-tutorial>

**QUESTION NO: 11 - (HOTSPOT)**

**HOTSPOT**

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and an Azure Data Lake Storage Gen2 account named Account1.

You plan to access the files in Account1 by using an external table.

You need to create a data source in Pool1 that you can reference when you create the external table.

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

NOTE: Each correct selection is worth one point.

**Hot Area:**

## Answer Area

```
CREATE EXTERNAL DATA SOURCE source1
```

```
WITH
```

```
( LOCATION = 'https://account1.
```

▼
blob
dfs
table

```
PUSHDOWN = ON  
TYPE = BLOB_STORAGE  
TYPE = HADOOP
```

```
)
```

ANSWER:

## Answer Area

```
CREATE EXTERNAL DATA SOURCE source1
```

```
WITH
```

```
( LOCATION = 'https://account1.
```

▼
blob
dfs
table

```
PUSHDOWN = ON  
TYPE = BLOB_STORAGE  
TYPE = HADOOP
```

```
)
```

Explanation:

Box 1: blob

The following example creates an external data source for Azure Data Lake Gen2

```
CREATE EXTERNAL DATA SOURCE YellowTaxi
```

```
WITH ( LOCATION = 'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/', TYPE = HADOOP)
```

Box 2: HADOOP

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables>

### QUESTION NO: 12 - (DRAG DROP)

DRAG DROP

You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1. DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL Server Data Tools (SSDT) database project.

You need to populate a new resource group named App1Test with the DB1 database and an Azure SQL Server named TestServer1. The resources in App1Test must have the same configurations as the resources in App1Dev.

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:

#### Actions

Change the Active Directory Admin on TestServer1

Change the server name and related variables in the templates

From the database project, deploy the database schema and permissions

Add IP addresses to the firewall

From the Azure portal, export the Azure Resource Manager templates

From the Azure portal, deploy the templates.

#### Answer Area



**ANSWER:**

## Actions

Change the Active Directory Admin on TestServer1

Change the server name and related variables in the templates

From the database project, deploy the database schema and permissions

Add IP addresses to the firewall

From the Azure portal, export the Azure Resource Manager templates

From the Azure portal, deploy the templates.

## Answer Area

From the Azure portal, export the Azure Resource Manager templates

Change the server name and related variables in the templates

From the Azure portal, deploy the templates.

From the database project, deploy the database schema and permissions

Explanation:

**QUESTION NO: 13**

You have an Azure Stream Analytics job.

You need to ensure that the job has enough streaming units provisioned.

You configure monitoring of the SU % Utilization metric.

Which two additional metrics should you monitor? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Late Input Events
- B. Out of order Events
- C. Backlogged Input Events
- D. Watermark Delay
- E. Function Events

**ANSWER: C D**

Explanation:

To react to increased workloads and increase streaming units, consider setting an alert of 80% on the SU Utilization metric. Also, you can use watermark delay and backlogged events metrics to see if there is an impact.

Note: Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently non-zero, you should scale out your job, by increasing the SUs.

Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring>

### QUESTION NO: 14 - (HOTSPOT)

You need to use an Azure Resource Manager ARM) template to deploy an Azure virtual machine that will host a Microsoft SQL Server instance. The solution must maximize disk I/O

performance for the SQL Server database and log files

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

NOTE: Each correct selection is worth one point.



**ANSWER:**

**Explanation:**

Read only ReadWrite

### QUESTION NO: 15 - (DRAG DROP)

You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored

You need to ensure that SQL Server Agent jobs running on SQL 1 can notify when a failure occurs

Which three actions should you perform in sequence 7 TO answer. move the appropriate actions from the list Of actions to answer area and arrange them in correct order.

Actions

Create a Database Mail account.

Enable pager notifications upon failure.

Create a profile named AzureManagedInstance\_dbmail\_profile.

Enable email notifications upon failure.

Create a profile named application\_dbmail\_profile.

>

<

Answer Area

ANSWER:

Explanation:

Answer Area

Create a Database Mail account.

Create a profile named AzureManagedInstance\_dbmail\_profile.

Enable email notifications upon failure.