

# DUMPS ARENA

## Implementing an Azure Data Solution

Microsoft DP-200

Version Demo

Total Demo Questions: 15

Total Premium Questions: 372

Buy Premium PDF

<https://dumpsarena.co>

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

[sales@dumpsarena.co](mailto:sales@dumpsarena.co)  
[dumpsarena.co](https://dumpsarena.co)

## Topic Break Down

<b>Topic</b>	<b>No. of Questions</b>
Topic 1, Case Study 1	2
Topic 2, Case Study 2	5
Topic 3, Case Study 3	2
Topic 4, Case Study 4	2
Topic 5, Case Study 5	3
Topic 6, Case Study 6	2
Topic 7, Case Study 7	5
Topic 8, Case Study 8	2
Topic 9, Case Study 9	3
Topic 10, Case Study 10	3
Topic 11, Mixed Questions	343
<b>Total</b>	<b>372</b>

**QUESTION NO: 1**

You have an Azure Blob storage account.

The storage account has an alert that is configured to indicate when the Availability metric falls below 100 percent.

You receive an alert for the Availability metric. The logs for the storage account show that requests are failing because of a ServerTimeoutError error.

What does ServerTimeoutError indicate?

- A. Read and write storage requests exceeded capacity.
- B. A transient server timeout occurred while the service was moved to a different partition to load balance requests.
- C. A client application attempted to perform an operation and did not have valid credentials.
- D. There was excessive network latency between a client application and the storage account.

**ANSWER: D****QUESTION NO: 2**

You are the data engineer for your company. An application uses a NoSQL database to store data. The database uses the key-value and wide-column NoSQL database type.

Developers need to access data in the database using an API.

You need to determine which API to use for the database model and type.

Which two APIs should you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Table API
- B. MongoDB API
- C. Gremlin API
- D. SQL API
- E. Cassandra API

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

B: Azure Cosmos DB is the globally distributed, multimodel database service from Microsoft for mission-critical applications. It is a multimodel database and supports document, key-value, graph, and columnar data models.

E: Wide-column stores store data together as columns instead of rows and are optimized for queries over large datasets. The most popular are Cassandra and HBase.

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/graph-introduction> <https://www.mongodb.com/scale/types-of-nosql-databases>

### QUESTION NO: 3

You have JSON files stored in an Azure Data Lake Storage Gen2 account. The JSON file contains the FirstName and LastName of customers. You need to use Azure Data bricks to copy the data in the JSON files to an Azure data warehouse. A new column must be created which concatenates the FirstName and LastName values.

You have the following components in place in Azure:

- A destination table in the SQL Data Warehouse
- An Azure Blob storage container
- A service principal

Which of the following are actions you would perform to transfer the data onto the Azure SQL Data warehouse table? (Choose five.)

- A. Write the results onto Azure Data Lake Storage
- B. Drop the data frame
- C. Perform transformations on the data frame
- D. Mount the Data Lake Storage onto DBFS
- E. Perform transformations on the file
- F. Read the file into a data frame
- G. Specify a temporary folder to stage the data
- H. Write the results to a table in SQL Data Warehouse

**ANSWER: C D F G H**

#### Explanation:

To Connect to Azure Data Lake Gen2 storage, we first need to mount the file system. This is also given in the Azure data bricks documentation:

Azure Data Lake Storage Gen2

Azure Data Lake Storage Gen2 (also known as ADLS Gen2) is a next-generation data lake solution for big data analytics. Azure Data Lake Storage Gen2 builds Azure Data Lake Storage Gen1 capabilities – file system semantics, file-level security, and scale – into Azure Blob storage, with its low-cost tiered storage, high availability, and disaster recovery features.

Note:

The Azure Data Lake Storage Gen2 connector is supported in Databricks Runtime 5.2 and above with full support for Delta Lake in Databricks Runtime 5.5 and above.

There are three ways of accessing Azure Data Lake Storage Gen2:

1. Mount an Azure Data Lake Storage Gen2 filesystem to DBFS using a service principal and OAuth 2.0.
2. Use a service principal directly.
3. Use the Azure Data Lake Storage Gen2 storage account access key directly.

Next, we need to load the data from the Azure Data Lake Storage Gen2 account. An example is also given in the Microsoft documentation:

Extract data from the Azure Data Lake Storage Gen2 account

1. You can now load the sample json file as a data frame in Azure Databricks. Paste the following code in a new cell. Replace the placeholders shown in brackets with your values.

```
Scala Copy  
val df = spark.read.json("adfss://<file-system-name>@<storage-account-name>
```

Next, we need to transform the data in the data frame to combine the FirstName and LastName values.

You need to mention an Azure storage account as a temporary staging area.

Load data into Azure SQL Data Warehouse

In this section, you upload the transformed data into Azure SQL Data Warehouse. You use the Azure SQL Data Warehouse connector for Azure Databricks to directly upload a dataframe as a table in a SQL data warehouse.

As mentioned earlier, the SQL Data Warehouse connector uses Azure Blob storage as temporary storage to upload data between Azure Databricks and Azure SQL Data Warehouse. So, you start by providing the configuration to connect to the storage account. You must already have already created the account as part of the prerequisites for this article.

And then finally you copy the data onto the Azure SQL data warehouse.

Option “Write the results onto Azure Data Lake Storage” is incorrect since you don’t need to write the results onto Azure Data Lake storage. The results are written to the Azure SQL Data warehouse table. Option “Drop the data frame” is incorrect since we don’t need to drop the data frames.

Option “Perform transformations on the file” is incorrect since transformations need to be carried out on the data frames. Reference:

<https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-extract-load-sql-data-warehouse#load-data-into-azure-sql-data-warehouse>

## QUESTION NO: 4 - (SIMULATION)

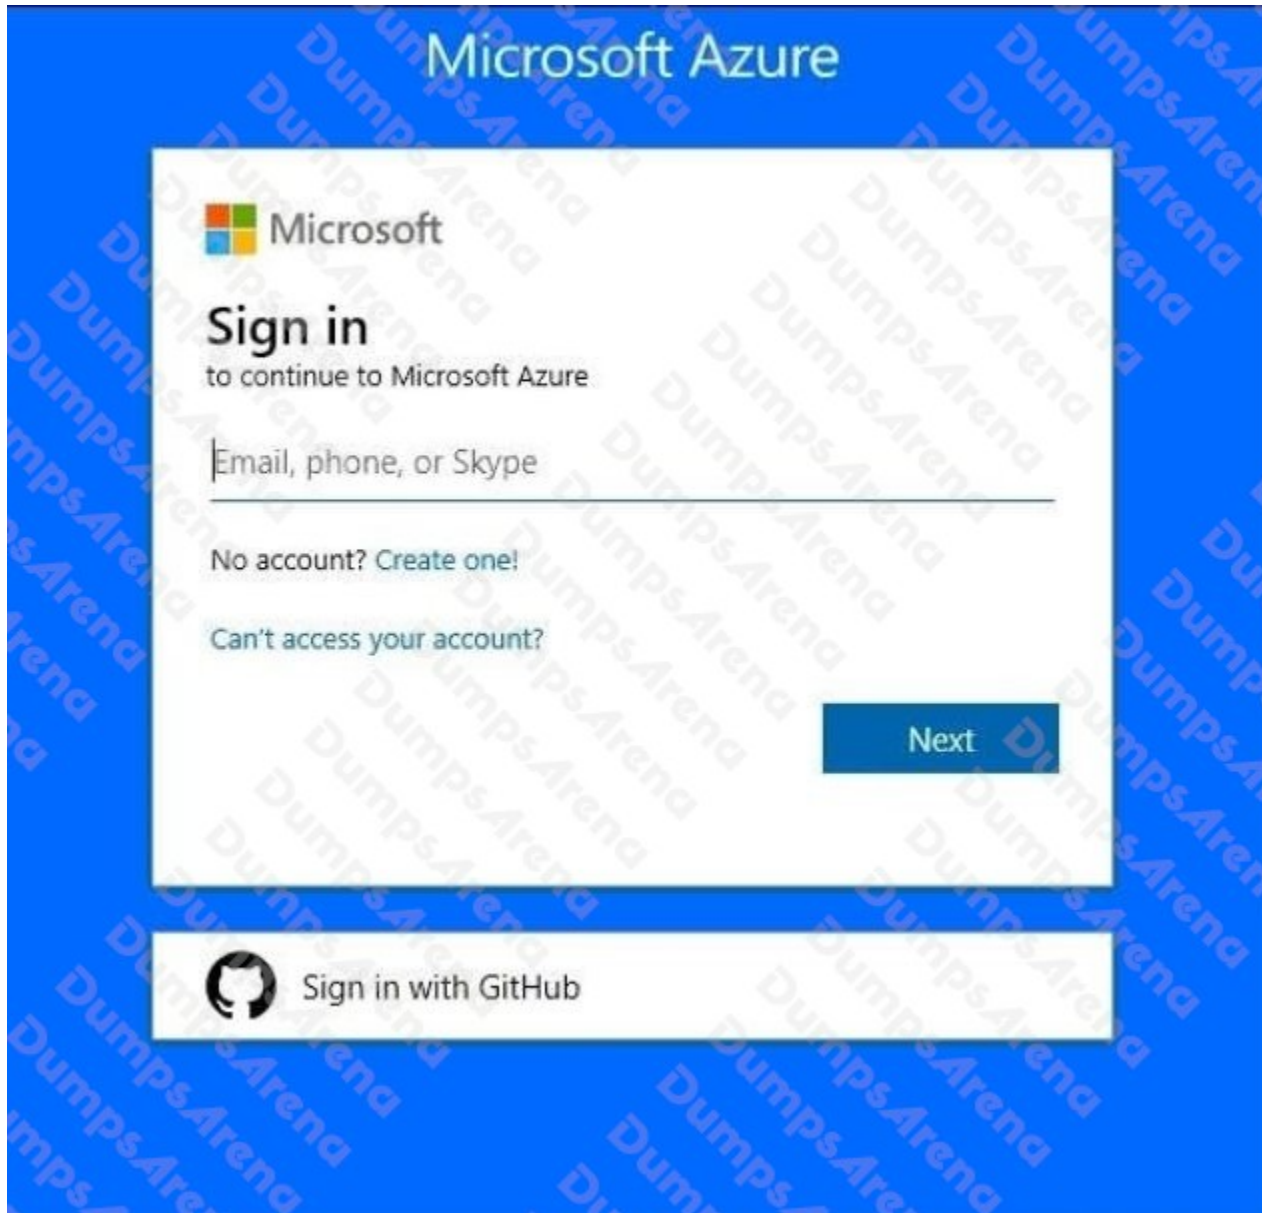
## SIMULATION

Use the following login credentials as needed:

Azure Username: xxxxx Azure Password: xxxxx

The following information is for technical support purposes only:

Lab Instance: 10543936



You need to create an elastic pool that contains an Azure SQL database named db2 and a new SQL database named db3.

To complete this task, sign in to the Azure portal.

**ANSWER: See the explanation below.**

**Explanation:**

Step 1: Create a new SQL database named db3

1. Select SQL in the left-hand menu of the Azure portal. If SQL is not in the list, select All services, then type SQL in the search box.
2. Select + Add to open the Select SQL deployment option page. Select Single Database. You can view additional information about the different databases by selecting Show details on the Databases tile.
3. Select Create:

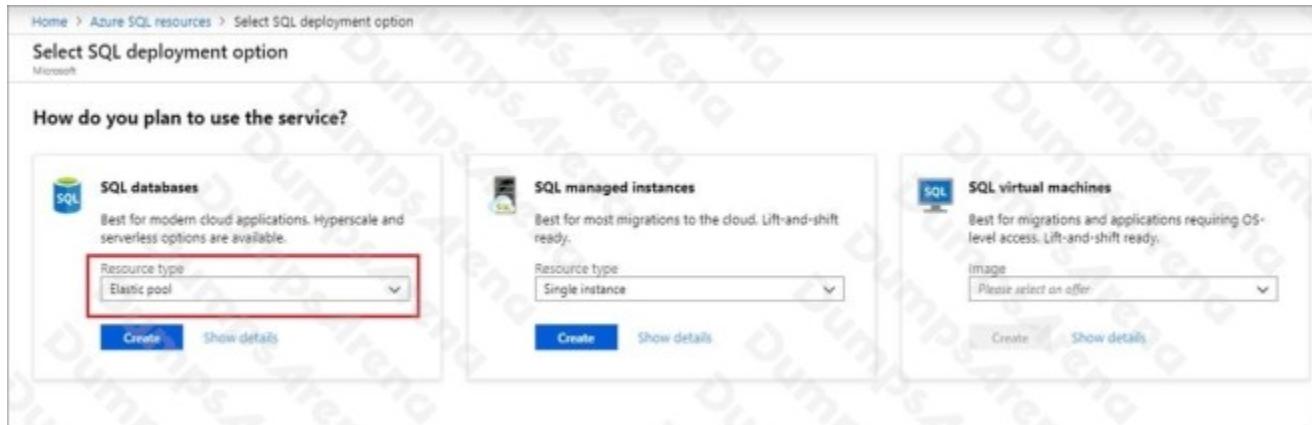


4. Enter the required fields if necessary.
5. Leave the rest of the values as default and select Review + Create at the bottom of the form.
6. Review the final settings and select Create. Use Db3 as database name.

On the SQL Database form, select Create to deploy and provision the resource group, server, and database.

Step 2: Create your elastic pool using the Azure portal.

1. Select Azure SQL in the left-hand menu of the Azure portal. If Azure SQL is not in the list, select All services, then type Azure SQL in the search box.
2. Select + Add to open the Select SQL deployment option page.
3. Select Elastic pool from the Resource type drop-down in the SQL Databases tile. Select Create to create your elastic pool.



4. Configure your elastic pool with the following values:

Name: Provide a unique name for your elastic pool, such as myElasticPool.

Subscription: Select your subscription from the drop-down.

ResourceGroup: Select the resource group.

Server: Select the server

Home > Azure SQL resources > Select SQL deployment option > Create SQL Elastic pool

## Create SQL Elastic pool

Microsoft

Basics • Tags Review + create

Create a SQL Elastic pool with your preferred configurations. Elastic pools provide a simple and cost effective solution for managing the performance of multiple databases within a fixed budget. Complete the Basic tab, then go to Review + Create to provision with smart defaults, or visit each tab to customize. [Learn more](#)

### Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

\* Subscription

\* Resource group  [Create new](#)

### Elastic pool details

Enter required settings for this pool, including picking a logical server and configuring the compute and storage resources.

\* Elastic Pool Name  ✓

\* Server  [Create new](#)

\* Compute + storage **GeneralPurpose**  
Gen5, 2 vCores, 32 GB, 0 databases  
[Configure elastic pool](#)

5. Select Configure elastic pool

6. On the Configure page, select the Databases tab, and then choose to Add database.



7. Add the Azure SQL database named db2, and the new SQL database named db3 that you created in Step 1.

8. Select Review + create to review your elastic pool settings and then select Create to create your elastic pool.

Reference:

<https://docs.microsoft.com/bs-latn-ba/azure/sql-database/sql-database-elastic-pool-failover-group-tutorial>

## QUESTION NO: 5 - (DRAG DROP)

DRAG DROP

You manage security for a database that supports a line of business application.

Private and personal data stored in the database must be protected and encrypted.

You need to configure the database to use Transparent Data Encryption (TDE).

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

**Select and Place:**

**Actions****Answer Area**

Create a database encryption key using a certificate generated with the master key.

Create a certificate and then create the master key using a password.

Set the context to the master database.

Create a master key using a password.

Set the context to the company database.

Enable encryption.

**ANSWER:**

**Actions****Answer Area**

Create a database encryption key using a certificate generated with the master key.

Create a master key using a password.

Create a certificate and then create the master key using a password.

Create a certificate and then create the master key using a password.

Set the context to the master database.

Set the context to the company database.

Create a master key using a password.

Create a database encryption key using a certificate generated with the master key.

Set the context to the company database.

Enable encryption.

Enable encryption.

**Explanation:**

Step 1: Create a master key

Step 2: Create or obtain a certificate protected by the master key

Step 3: Set the context to the company database

Step 4: Create a database encryption key and protect it by the certificate

Step 5: Set the database to use encryption

Example code:

```
USE master;
```

```
GO
```

```
CREATE MASTER KEY ENCRYPTION BY PASSWORD = ""; go
```

```
CREATE CERTIFICATE MyServerCert WITH SUBJECT = 'My DEK Certificate'; go
```

```
USE AdventureWorks2012;
```

```
GO
```

```
CREATE DATABASE ENCRYPTION KEY
```

```
WITH ALGORITHM = AES_128
```

```
ENCRYPTION BY SERVER CERTIFICATE MyServerCert;
```

```
GO
```

```
ALTER DATABASE AdventureWorks2012
```

```
SET ENCRYPTION ON; GO
```

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption>

**QUESTION NO: 6**

An application is currently making use of a database on the Azure platform. Below is a snippet of the code base:

```
private static readonly string XYZendpointUrl = ConfigurationManager.AppSettings["EndpointUrl"];
```

```
private static readonly SecureString XYZkey=ToSecureString(ConfigurationManager.AppSettings["AuthorizationKey"]); var  
XYZ_client= new DocumentClient(new Url(XYZendpointUrl), XYZkey);
```

```
Database database= await XYZ_client.CreateDatabaseAsync(new Database { Id="XYZdb" });
```

Which of the following is the key type used in the code?

- A. Resource token
- B. Master key
- C. Certificate
- D. Password

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

Here the master keys of the account are being used for authentication. The documentation mentions the following:

Master keys provide access to all the administrative resources for the database account.

Master keys:

- Provide access to accounts, databases, users, and permissions.
- Cannot be used to provide granular access to containers and documents.
- Are created during the creation of an account.
- Can be regenerated at any time.

Each account consists of two Master keys: a primary key and secondary key. The purpose of dual keys is so that you can regenerate, or roll keys, providing continuous access to your account and data.

Since this is clearly mentioned in the Microsoft documentation, all other options are incorrect. Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/secure-access-to-data>

**QUESTION NO: 7**

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

You develop a data ingestion process that will import data to an enterprise data warehouse in Azure Synapse Analytics. The data to be ingested resides in parquet files stored in an Azure Data Lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Data Warehouse.

Solution:

1. Create a remote service binding pointing to the Azure Data Lake Gen 2 storage account
2. Create an external file format and external table using the external data source
3. Load the data using the CREATE TABLE AS SELECT statement

Does the solution meet the goal?

- A. Yes
- B. No

**ANSWER: B**

**Explanation:**

You need to create an external file format and external table from an external data source, instead from a remote service binding pointing.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

**QUESTION NO: 8**

You need to enable Transparent Data Encryption for an Azure SQL database.

Which of the following steps would you perform for this requirement? (Choose five.)

- A. Create a database encryption key using a certificate
- B. Create a certificate protected by the master key
- C. Set the context to the master database
- D. Create a master key using a password
- E. Set the context to the company database

**F. Enable Encryption****ANSWER: A B D E F****Explanation:**

The list of steps is given in the Microsoft documentation:

Using Transparent Data Encryption To use TDE, follow these steps. Applies to: SQL Server.

- Create a master key
- Create or obtain a certificate protected by the master key
- Create a database encryption key and protect it by the certificate
- Set the database to use encryption

The following example illustrates encrypting and decrypting the AdventureWorks2012 database using a certificate installed on the server named MyServerCert.

```
SQL Copy  
USE master;  
GO  
CREATE MASTER KEY ENCRYPTION BY PASSWORD = '<UseStrongPasswordHere>';  
go  
CREATE CERTIFICATE MyServerCert WITH SUBJECT = 'My DEK Certificate';  
go  
USE AdventureWorks2012;  
GO  
CREATE DATABASE ENCRYPTION KEY  
WITH ALGORITHM = AES_128  
ENCRYPTION BY SERVER CERTIFICATE MyServerCert;  
GO  
ALTER DATABASE AdventureWorks2012  
SET ENCRYPTION ON;  
GO
```

Since this is clearly mentioned in the Microsoft documentation, all other options are incorrect. Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption?view=sql-server-ver15>

**QUESTION NO: 9 - (DRAG DROP)****DRAG DROP**

You need to ensure that phone-based polling data can be analyzed in the PollingData database.

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:**

**Actions**

**Answer Area**

Parameterize deployment by using Azure Integration Runtime

Configure an Azure Logic App to deploy the deployment artifact

Configure Azure DevOps to deploy the deployment artifact

Create a deployment artifact containing an extracted Azure Resource Manager template

Parameterize deployment by using the Azure Resource Manager template parameter file

Create a deployment artifact containing a SQL Server Integration Services (SSIS) package

**ANSWER:**

**Actions**

Parameterize deployment by using Azure Integration Runtime

Configure an Azure Logic App to deploy the deployment artifact

Configure Azure DevOps to deploy the deployment artifact

Create a deployment artifact containing an extracted Azure Resource Manager template

Parameterize deployment by using the Azure Resource Manager template parameter file

Create a deployment artifact containing a SQL Server Integration Services (SSIS) package

**Answer Area**

Create a deployment artifact containing an extracted Azure Resource Manager template

Parameterize deployment by using the Azure Resource Manager template parameter file

Configure Azure DevOps to deploy the deployment artifact

**Explanation:**

Scenario:

All deployments must be performed by using Azure DevOps. Deployments must use templates used in multiple environments No credentials or secrets should be used during deployments

**QUESTION NO: 10**

You need to migrate data from an Azure Blob storage account to an Azure SQL Data warehouse.

Which of the following actions do you need to implement for this requirement? (Choose four.)

- A. Provision an Azure SQL Data Warehouse instance
- B. Create an Azure Blob storage container
- C. Run the T-SQL statements to load the data
- D. Connect to the Azure SQL Data warehouse via SQL Server Management Studio
- E. Build external tables by using Azure portal

F. Build external tables by using SQL Server Management Studio

**ANSWER: A C D F**

**Explanation:**

You first need to create an Azure SQL Data Warehouse instance.

Then you need to connect to the data warehouse via SQL Server Management Studio.

Then create external tables to the Azure Blob storage account.

And then finally use T-SQL statements to load the data.

This is also given as an example in github as part of the Microsoft documentation on loading data from Azure Blob to an Azure SQL data warehouse.

This tutorial uses PolyBase to load New York Taxicab data from a global Azure blob to Azure SQL Data Warehouse. The tutorial uses the Azure portal and SQL Server Management Studio (SSMS) to: [!div class="checklist"]

- Create a data warehouse in the Azure portal
- Set up a server-level firewall rule in the Azure portal
- Connect to the data warehouse with SSMS
- Create a user designated for loading data
- Create external tables for data in Azure blob storage
- Use the CTAS T-SQL statement to load data into your data warehouse
- View the progress of data as it is loading
- Create statistics on the newly loaded data

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

**HOTSPOT**

You are a data engineer. You are designing a Hadoop Distributed File System (HDFS) architecture. You plan to use Microsoft Azure Data Lake as a data storage repository.

You must provision the repository with a resilient data schema. You need to ensure the resiliency of the Azure Data Lake Storage. What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Hot Area:**

**Answer Area**

**Requirement**

**Node**

Provide data access to clients.

DataNode	✓
NameNode	

Run operations on files and directories of the file system.

DataNode	✓
NameNode	

Perform block creation, deletion, and replication.

DataNode	✓
NameNode	

**ANSWER:**

**Answer Area**

**Requirement**

**Node**

Provide data access to clients.

DataNode	✓
NameNode	

Run operations on files and directories of the file system.

DataNode	✓
NameNode	

Perform block creation, deletion, and replication.

DataNode	✓
NameNode	

**Explanation:**

Box 1: NameNode

An HDFS cluster consists of a single NameNode, a master server that manages the file system namespace and regulates access to files by clients.

Box 2: DataNode

The DataNodes are responsible for serving read and write requests from the file system's clients.

Box 3: DataNode

The DataNodes perform block creation, deletion, and replication upon instruction from the NameNode.

Note: HDFS has a master/slave architecture. An HDFS cluster consists of a single NameNode, a master server that manages the file system namespace and regulates access to files by clients. In addition, there are a number of DataNodes, usually one per node in the cluster, which manage storage attached to the nodes that they run on. HDFS exposes a file system namespace and allows user data to be stored in files. Internally, a file is split into one or more blocks and these blocks are stored in a set of DataNodes. The NameNode executes file system namespace operations like opening, closing, and renaming files and directories. It also determines the mapping of blocks to DataNodes. The DataNodes are responsible for serving read and write requests from the file system's clients. The DataNodes also perform block creation, deletion, and replication upon instruction from the NameNode.

References: [https://hadoop.apache.org/docs/r1.2.1/hdfs\\_design.html#NameNode+and+DataNodes](https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html#NameNode+and+DataNodes)

**QUESTION NO: 12**

You have to design a Hadoop Distributed File System architecture. You are going to be using Microsoft Azure Data Lake as the data storage repository. You have to ensure the data repository has a resilient data schema.

Which of the following would be used to run operations on files and directories on the file system?

- A. DataNode
- B. NameNode
- C. PrimaryNode
- D. SecondaryNode

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

The file system namespace resides on the NameNode. The Hadoop documentation mentions the following:

The File System Namespace

HDFS supports a traditional hierarchical file organization. A user or an application can create directories and store files inside these directories. The file system namespace hierarchy is similar to most other existing file systems; one can create and remove files, move a file from one directory to another, or rename a file. HDFS does not yet implement user quotas. HDFS does not support hard links or soft links.

However, the HDFS architecture does not preclude implementing these features.

The NameNode maintains the file system namespace. Any change to the file system namespace or its properties is recorded by the NameNode. An application can specify the number of replicas of a file that should be maintained by HDFS. The number of copies of a file is called the replication factor of that file. This information is stored by the NameNode.

Since this is clear from the documentation, all other options are incorrect. Reference:

[https://hadoop.apache.org/docs/r1.2.1/hdfs\\_design.html#NameNode+and+DataNodes](https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html#NameNode+and+DataNodes)

**QUESTION NO: 13**

A company wants to implement a solution that would perform batch processing of geospatial data.

Which of the following can be used to develop batch processing applications that use Azure HDInsight?

- A. HDInsight Tools for Visual studio
- B. Hive View
- C. HDInsight REST API
- D. Azure Data Factory

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

You can use the HDInsight REST API to issue calls to the HDInsight cluster. The Microsoft documentation mentions the following:

Azure HDInsight REST API

Use these APIs to create and manage HDInsight resources through Azure Resource Manager. All task operations conform to the HTTP/1.1 protocol specification and each operation returns an x-msrequest-id header that can be used to obtain information about the request. Ensure that requests made to these resources are secure. For more information, see [Authenticating Azure Resource Manager requests](#).

Option "HDInsight Tools for Visual studio" is incorrect since this is a native tool that can be used to run interactive queries against HDInsight. Option "Hive View" is incorrect since this is Web UI tool that can be used to execute queries against Apache Hadoop in HDInsight. Option "Azure Data Factory" is incorrect since this can be used as an ETL or ELT tool. Reference:

<https://docs.microsoft.com/en-us/rest/api/hdinsight/>

**QUESTION NO: 14 - (SIMULATION)****SIMULATION**

Use the following login credentials as needed:

Azure Username: xxxxx Azure Password: xxxxx

The following information is for technical support purposes only:

Lab Instance: 10543936

You need to ensure that all REST API calls to an Azure Storage account named storage10543936 use HTTPS only.

To complete this task, sign in to the Azure portal.

**ANSWER: See the explanation below.**

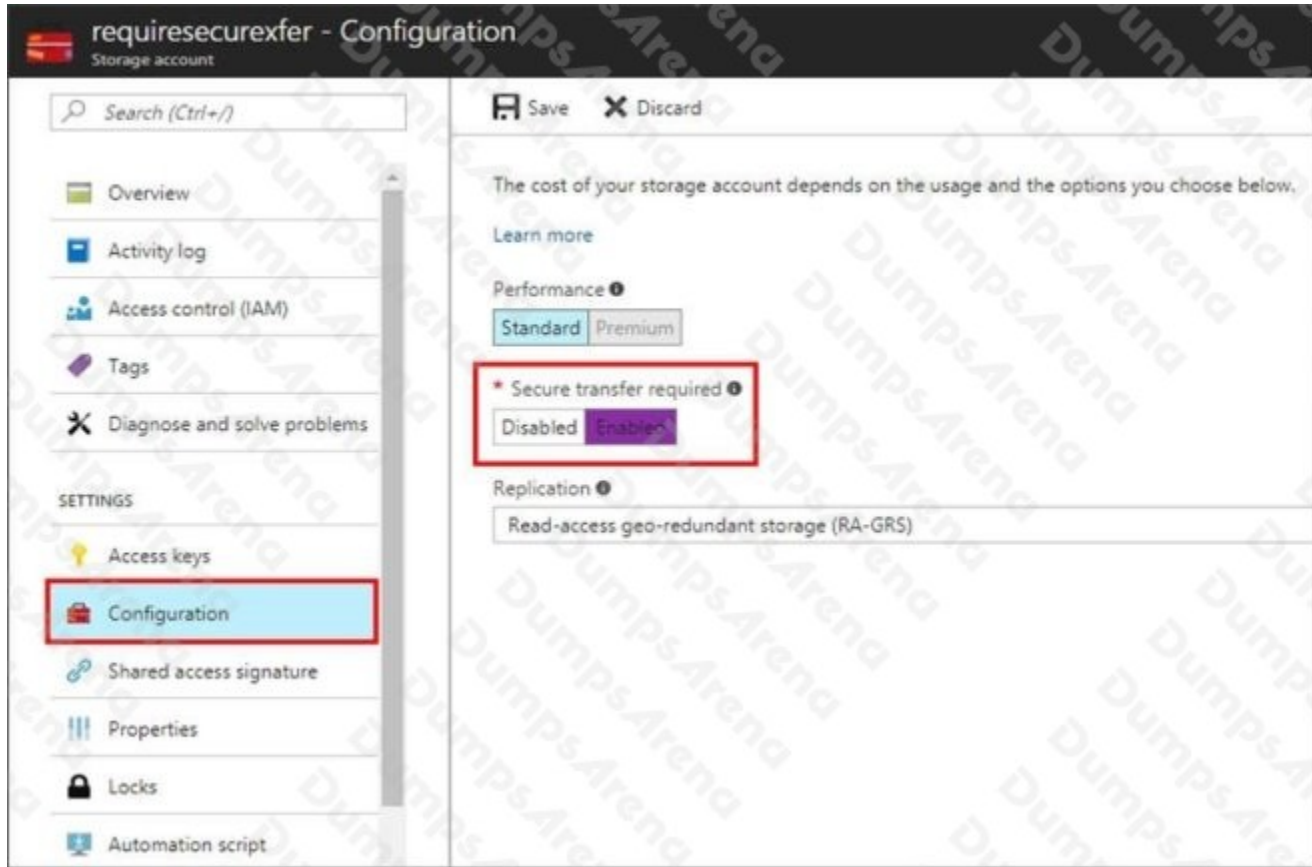
**Explanation:**

You can configure your storage account to accept requests from secure connections only by setting the Secure transfer required property for the storage account.

Require secure transfer for an existing storage account

1. Select the existing storage account storage10543936 in the Azure portal.

2. In the storage account menu pane, under SETTINGS, select Configuration.
3. Under Secure transfer required, select Enabled.



Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-require-secure-transfer>

## QUESTION NO: 15 - (HOTSPOT)

### HOTSPOT

You have an enterprise data warehouse in Azure Synapse Analytics that contains a table named FactOnlineSales. The table contains data from the start of 2009 to the end of 2012.

You need to improve the performance of queries against FactOnlineSales by using table partitions. The solution must meet the following requirements:

- Create four partitions based on the order date.
- Ensure that each partition contains all the orders placed during a given calendar year.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
CREATE TABLE [dbo].FactOnlineSales
([OnlineSalesKey] [int] NOT NULL,
[OrderDateKey] [datetime] NOT NULL,
[StoreKey] [int] NOT NULL,
[ProductKey] [int] NOT NULL,
[CustomerKey] [int] NOT NULL,
[SalesOrderNumber] [nvarchar](20) NOT NULL,
[SalesQuantity] [int] NOT NULL,
[SalesAmount] [money] NOT NULL,
[UnitPrice] [money] NULL)
WITH (CLUSTERED COLUMNSTORE INDEX)
PARTITION ([OrderDateKey] RANGE
```

	▼	FOR VALUES
RIGHT		
LEFT		

	▼	)
20090101,20121231		
201001101,20110101,20120101		
20090101,20100101,20110101,20120101		

ANSWER:

## Answer Area

```

CREATE TABLE [dbo].FactOnlineSales
([OnlineSalesKey] [int] NOT NULL,
[OrderDateKey] [datetime] NOT NULL,
[StoreKey] [int] NOT NULL,
[ProductKey] [int] NOT NULL,
[CustomerKey] [int] NOT NULL,
[SalesOrderNumber] [nvarchar](20) NOT NULL,
[SalesQuantity] [int] NOT NULL,
[SalesAmount] [money] NOT NULL,
[UnitPrice] [money] NULL)
WITH (CLUSTERED COLUMNSTORE INDEX)
PARTITION ([OrderDateKey] RANGE

```

	▼
RIGHT	
LEFT	

```

FOR VALUES
(

```

	▼
20090101,20121231	
201001101,20110101,20120101	
20090101,20100101,20110101,20120101	

```

)

```

**Explanation:**

Box 1: LEFT

RANGE LEFT: Specifies the boundary value belongs to the partition on the left (lower values). The default is LEFT.

Box 2: 20090101, 20100101, 20110101, 20120101

FOR VALUES ( boundary\_value [...n] ) specifies the boundary values for the partition. boundary\_value is a constant expression.

Reference: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse>