



DP-200^{Q&As}

Implementing an Azure Data Solution

Pass Microsoft DP-200 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/dp-200.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

- ⚙ **Instant Download** After Purchase
- ⚙ **100% Money Back** Guarantee
- ⚙ **365 Days** Free Update
- ⚙ **800,000+** Satisfied Customers



**QUESTION 1**

You develop data engineering solutions for a company.

A project requires an in-memory batch data processing solution.

You need to provision an HDInsight cluster for batch processing of data on Microsoft Azure.

How should you complete the PowerShell segment? To answer, select the appropriate option in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

1 ▼

New-AzureStorageContainer
New-AzureRmHDInsightClusterConfig
New-AzureRmHDInsightCluster

2 ▼

"Spark"
"Hadoop"

3 ▼

New-AzureRmHDInsightCluster
New-AzureRmHDInsightClusterConfig
New-AzureStorageContainer

4 ▼

"spark"
"Hadoop"
"HBase"
"Storm"

1

```
-Name $clusterName -Context $defaultStorageContext
$ObjectConfig = New-Object "system.collections.Generic.D
$ObjectConfig.Add 2 ,"2.3")

(
3
-ResourceGroupName $resourceGroupName
-ClusterName $clusterName
-Location $location
-ClusterSizeInNodes $clusterSizeInNodes
-ClusterType 4
-OSType $clusterOS
-Version $clusterVersion
```

Correct Answer:



1 ▾

New-AzureStorageContainer

New-AzureRmHDInsightClusterConfig

New-AzureRmHDInsightCluster

2 ▾

"Spark"

"Hadoop"

3 ▾

New-AzureRmHDInsightCluster

New-AzureRmHDInsightClusterConfig

New-AzureStorageContainer

4 ▾

"spark"

"Hadoop"

"HBase"

"Storm"

1

```
-Name $clusterName -Context $defaultStorageContext
$ObjectConfig = New-Object "system.collections.Generic.D
$ObjectConfig.Add 

2

 ,"2.3")

(


3


-ResourceGroupName $resourceGroupName
-ClusterName $clusterName
-Location $location
-ClusterSizeInNodes $clusterSizeInNodes
-ClusterType 

4


-OSType $clusterOS
-Version $clusterVersion
```

Box 1: New-AzStorageContainer

Example: Create a blob container. This holds the default data store for the cluster.

```
New-AzStorageContainer `
```

```
-Name $clusterName `
```

```
-Context $defaultStorageContext
```

```
$sparkConfig = New-Object "System.Collections.Generic.Dictionary`2[System.String,System.String]"
$sparkConfig.Add("spark", "2.3")
```

Box 2: Spark Spark provides primitives for in-memory cluster computing. A Spark job can load and cache data into memory and query it repeatedly. In-memory computing is much faster than disk-based applications than disk-based applications, such as Hadoop, which shares data through Hadoop distributed file system (HDFS).

Box 3: New-AzureRMHDInsightCluster # Create the HDInsight cluster. Example: New-AzHDInsightCluster `

```
-ResourceGroupName $resourceGroupName `
```

```
-ClusterName $clusterName `
```

```
-Location $location `
```

```
-ClusterSizeInNodes $clusterSizeInNodes `
```

```
-ClusterType $"Spark" `
```

```
-OSType "Linux" `
```

**Box 4: Spark**

HDInsight is a managed Hadoop service. Use it to deploy and manage Hadoop clusters in Azure. For batch processing, you can use Spark, Hive, Hive LLAP, MapReduce.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/hdinsight/spark/apache-spark-jupyter-spark-sql-use-powershell>

<https://docs.microsoft.com/bs-latn-ba/azure/hdinsight/spark/apache-spark-overview>

QUESTION 2**HOTSPOT**

You are implementing Azure Stream Analytics functions.

Which windowing function should you use for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Segment the data stream into distinct time segments that repeat but do not overlap:

	▼
Hopping	
Session	
Sliding	
Tumbling	

Segment the data stream into distinct time segments that repeat and can overlap:

	▼
Hopping	
Session	
Sliding	
Tumbling	

Segment the data stream to produce an output only when an event occurs:

	▼
Hopping	
Session	
Sliding	
Tumbling	

Correct Answer:



Answer Area

Segment the data stream into distinct time segments that repeat but do not overlap:

	▼
Hopping	
Session	
Sliding	
Tumbling	

Segment the data stream into distinct time segments that repeat and can overlap:

	▼
Hopping	
Session	
Sliding	
Tumbling	

Segment the data stream to produce an output only when an event occurs:

	▼
Hopping	
Session	
Sliding	
Tumbling	

Box 1: Tumbling

Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap,

and an event cannot belong to more than one tumbling window.




Tell me the count of tweets per time zone every 10 seconds



```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second,10)
```

Box 2: Hopping Hopping window functions hop forward in time by a fixed period. It may be easy to think of them as Tumbling windows that can overlap, so events can belong to more than one Hopping window result set. To make a Hopping window the same as a Tumbling window, specify the hop size to be the same as the window size.

Every 5 seconds give me the count of tweets over the last 10 seconds




```
SELECT Topic, COUNT(*) AS TotalTweets
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY Topic, HoppingWindow(second, 10 , 5)
```

Box 3: Sliding

Sliding window functions, unlike Tumbling or Hopping windows, produce an output only when an event occurs. Every window will have at least one event and the window continuously moves forward by an ϵ (epsilon). Like hopping windows,



Give me the count of tweets for a single topic in the last 10 seconds.



```
SELECT Topic, COUNT(*)  
FROM TwitterStream TIMESTAMP BY CreatedAt  
GROUP BY Topic, SlidingWindow(second, 10)
```

events can belong to more than one sliding window.

References: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-window-functions>

QUESTION 3

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: 10277521

You plan to create large data sets on db2.

You need to ensure that missing indexes are created automatically by Azure in db2. The solution must apply ONLY to db2.

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

Correct Answer: See the below.

Explanation:

Automatic tuning

Revert to defaults

Azure SQL Database built-in intelligence automatically tunes your databases to optimize performance. Click here to learn more about automatic tuning.

Inherit from:
 ☒ Azure defaults ☐ Don't inherit

The database is inheriting automatic tuning configuration from Azure defaults.

Configure the automatic tuning options:

OPTION	DESIRED STATE			CURRENT STATE
FORCE PLAN	<input type="checkbox"/> ON	<input type="checkbox"/> OFF	<input checked="" type="button" value="INHERIT"/>	OFF Forced by user
CREATE INDEX	<input type="checkbox"/> ON	<input type="checkbox"/> OFF	<input checked="" type="button" value="INHERIT"/>	OFF Forced by user
DROP INDEX	<input type="checkbox"/> ON	<input type="checkbox"/> OFF	<input checked="" type="button" value="INHERIT"/>	OFF Forced by user

The selected configuration will be applied to all the databases that inherit automatic tuning configuration from this server. Click to see the list of databases.

Apply

1. To enable automatic tuning on Azure SQL Database logical server, navigate to the server in Azure portal and then select Automatic tuning in the menu.

2.

Select database db2

3.

Click the Apply button

Reference: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-automatic-tuning-enable>

QUESTION 4

HOTSPOT



You have a SQL pool in Azure Synapse.

You plan to load data from Azure Blob storage to a staging table. Approximately 1 million rows of data will be loaded daily. The table will be truncated before each daily load.

You need to create the staging table. The solution must minimize how long it takes to load the data to the staging table.

How should you configure the table? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

Distribution:

	▼
Hash	
Replicated	
Round-robin	

Indexing:

	▼
Clustered	
Clustered columnstore	
Heap	

Partitioning:

	▼
Date	
None	

Correct Answer:



Answer Area

Distribution:

	▼
Hash	
Replicated	
Round-robin	

Indexing:

	▼
Clustered	
Clustered columnstore	
Heap	

Partitioning:

	▼
Date	
None	

Box 1: Hash

Hash-distributed tables improve query performance on large fact tables. they can have very large numbers of rows and still achieve high performance.

Incorrect:

Round-robin tables are useful for improving loading speed.

Box 2: Clustered columnstore



When creating partitions on clustered columnstore tables, it is important to consider how many rows belong to each partition. For optimal compression and performance of clustered columnstore tables, a minimum of 1 million rows per distribution and partition is needed.

Box 3: Date

Table partitions enable you to divide your data into smaller groups of data. In most cases, table partitions are created on a date column.

Partition switching can be used to quickly remove or replace a section of a table.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-partition>

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribute>

QUESTION 5

HOTSPOT

You need to receive an alert when Azure SQL Data Warehouse consumes the maximum allotted resources.

Which resource type and signal should you use to create the alert in Azure Monitor? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Resource type:

	▼
Resource group	
SQL server	
SQL data warehouse	
Subscription	

Signal:

	▼
CPU used	
Data IO percentage	
DWU limit	
DWU used	



Correct Answer:

Answer Area

Resource type:

	▼
Resource group	
SQL server	
SQL data warehouse	
Subscription	

Signal:

	▼
CPU used	
Data IO percentage	
DWU limit	
DWU used	

Resource type: SQL data warehouse

DWU limit belongs to the SQL data warehouse resource type.

Signal: DWU limit

SQL Data Warehouse capacity limits are maximum values allowed for various components of Azure SQL Data Warehouse.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-insights-alerts-portal>

[DP-200 PDF Dumps](#)

[DP-200 Study Guide](#)

[DP-200 Braindumps](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - Windows, Mac, Android, iPhone, iPod, iPad, Kindle

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.pass4itsure.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 One Year Free Update <p>Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 Money Back Guarantee <p>To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 Security & Privacy <p>We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © pass4itsure, All Rights Reserved.