



# DP-200<sup>Q&As</sup>

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.passapply.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

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 questions 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 need to implement diagnostic logging for Data Warehouse monitoring.

Which log should you use?

- A. RequestSteps
- B. DmsWorkers
- C. SqlRequests
- D. ExecRequests

Correct Answer: C

Scenario:

The Azure SQL Data Warehouse cache must be monitored when the database is being used.

Metric	Description
A	Low cache hit %, high cache usage %
B	Low cache hit %, low cache usage %
C	High cache hit %, high cache usage %

References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-sql-requests-transact-sq>

### QUESTION 2

You have an Azure SQL database named DB1 that contains a table named Table1. Table1 has a field named Customer\_ID that is varchar(22). You need to implement masking for the Customer\_ID field to meet the following requirements:

The first two prefix characters must be exposed. The last four prefix characters must be exposed. All other characters must be masked.

Solution: You implement data masking and use a credit card function mask.

Does this meet the goal?

- A. Yes



B. No

Correct Answer: B

We must use Custom Text data masking, which exposes the first and last characters and adds a custom padding string in the middle.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started>

---

### QUESTION 3

You have the Diagnostics settings of an Azure Storage account as shown in the following exhibit.



**i** It may take up to 30 seconds to persist diagnostic logs settings.

### Status **i**

Off  On

### Blob properties

### File Properties

### Table properties

### Queue properties

### Hour metrics

- Enable
  - Include API metrics
  - Delete data

after  days

### Minute metrics

- Enable

### Logging

Logging version **i**

- Read
- Write
- Delete
- Deleta data

How long will the logging data be retained?

- A. 7 days
- B. 365 days
- C. indefinitely



D. 90 days

Correct Answer: A

Reference: <https://docs.microsoft.com/en-us/azure/storage/common/storage-analytics-metrics>

**QUESTION 4**

You need to ensure polling data security requirements are met.

Which security technologies 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**

Context	Security technology
SQL Server	Azure Active Directory user
	Domain Active Directory user
	Managed Identity
PolyBase	Database scoped credential
	Database encryption key
	Application role

Correct Answer:



## Answer Area

Context	Security technology
SQL Server	Azure Active Directory user
	Domain Active Directory user
	Managed Identity
PolyBase	Database scoped credential
	Database encryption key
	Application role

Box 1: Azure Active Directory user

Scenario:

Access to polling data must set on a per-active directory user basis

Box 2: DataBase Scoped Credential

SQL Server uses a database scoped credential to access non-public Azure blob storage or Kerberos-secured Hadoop clusters with PolyBase.

PolyBase cannot authenticate by using Azure AD authentication.

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-scoped-credential-transact-sql>

## QUESTION 5

You plan to monitor the performance of Azure Blob storage by using Azure Monitor.

You need to be notified when there is a change in the average time it takes for a storage service or API operation type to process requests.

For which two metrics should you set up alerts? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. SuccessE2ELatency
- B. SuccessServerLatency
- C. UsedCapacity
- D. Egress



E. Ingress

Correct Answer: AB

Success E2E Latency: The average end-to-end latency of successful requests made to a storage service or the specified API operation. This value includes the required processing time within Azure Storage to read the request, send the response, and receive acknowledgment of the response.

Success Server Latency: The average time used to process a successful request by Azure Storage. This value does not include the network latency specified in SuccessE2ELatency.

Reference: <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-scalable-app-verify-metrics>

---

### QUESTION 6

You need to implement complex stateful business logic within an Azure Stream Analytics service. Which type of function should you create in the Stream Analytics topology?

- A. JavaScript user-define functions (UDFs)
- B. Azure Machine Learning
- C. JavaScript user-defined aggregates (UDA)

Correct Answer: C

Azure Stream Analytics supports user-defined aggregates (UDA) written in JavaScript, it enables you to implement complex stateful business logic. Within UDA you have full control of the state data structure, state accumulation, state decumulation, and aggregate result computation.

References: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-javascript-user-defined-aggregates>

---

### QUESTION 7

You plan to create a new single database instance of Microsoft Azure SQL Database.

The database must only allow communication from the data engineer's workstation. You must connect directly to the instance by using Microsoft SQL Server Management Studio.

You need to create and configure the Database. Which three Azure PowerShell cmdlets should you use to develop the solution? To answer, move the appropriate cmdlets from the list of cmdlets to the answer area and arrange them in the correct order.

Select and Place:





**Azure PowerShell cmdlets**

**Answer Area**

New-AzureRmSqlElasticPool
New-AzureRmSqlServerFirewallRule
New-AzureRmSqlServer
New-AzureRmSqlServerVirtualNetworkRule
New-AzureRmSqlDatabase


Correct Answer:

**Azure PowerShell cmdlets**

**Answer Area**

New-AzureRmSqlElasticPool
New-AzureRmSqlServerVirtualNetworkRule

New-AzureRmSqlServer
New-AzureRmSqlServerFirewallRule
New-AzureRmSqlDatabase

Step 1: New-AzureSqlServer

Create a server.

Step 2: New-AzureRmSqlServerFirewallRule

New-AzureRmSqlServerFirewallRule creates a firewall rule for a SQL Database server.

Can be used to create a server firewall rule that allows access from the specified IP range.

Step 3: New-AzureRmSqlDatabase

Example: Create a database on a specified server

```
PS C:\>New-AzureRmSqlDatabase -ResourceGroupName "ResourceGroup01" -ServerName "Server01"
-DatabaseName "Database01"
```

References: <https://docs.microsoft.com/en-us/azure/sql-database/scripts/sql-database-create-and-configure-database->





powershell?toc=%2fpowershell%2fmodule%2ftoc.json

---

### QUESTION 8

You have a SQL pool in Azure Synapse.

A user reports that queries against the pool take longer than expected to complete.

You need to add monitoring to the underlying storage to help diagnose the issue.

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

NOTE: Each correct selection is worth one point.

- A. Cache used percentage
- B. DWU Limit
- C. Snapshot Storage Size
- D. Active queries
- E. Cache hit percentage

Correct Answer: AE

A: Cache used is the sum of all bytes in the local SSD cache across all nodes and cache capacity is the sum of the storage capacity of the local SSD cache across all nodes.

E: Cache hits is the sum of all columnstore segments hits in the local SSD cache and cache miss is the columnstore segments misses in the local SSD cache summed across all nodes

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-concept-resource-utilization-query-activity>

---

### QUESTION 9

You develop data engineering solutions for a company.

A project requires analysis of real-time Twitter feeds. Posts that contain specific keywords must be stored and processed on Microsoft Azure and then displayed by using Microsoft Power BI. You need to implement the solution.

Which five 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**

Create an HDInsight cluster with the Hadoop cluster type.

Create a Jupyter Notebook.

Run a job that uses the Spark Streaming API to ingest data from Twitter.

Create a Runbook.

Create an HDInsight cluster with the Spark cluster type.

Create a table.

Load the hvac table into Power BI Desktop.

**Answer Area**






Correct Answer:

**Actions**

Create an HDInsight cluster with the Hadoop cluster type.


Create a Runbook.



**Answer Area**

Create an HDInsight cluster with the Spark cluster type.

Create a Jupyter Notebook.

Create a table.

Run a job that uses the Spark Streaming API to ingest data from Twitter.

Load the hvac table into Power BI Desktop.

Step 1: Create an HDInsight cluster with the Spark cluster type Step 2: Create a Jupyter Notebook Step 3: Create a table

The Jupyter Notebook that you created in the previous step includes code to create an hvac table. Step 4: Run a job that uses the Spark Streaming API to ingest data from Twitter Step 5: Load the hvac table into Power BI Desktop

You use Power BI to create visualizations, reports, and dashboards from the Spark cluster data.

References:

<https://acadgild.com/blog/streaming-twitter-data-using-spark>

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-use-with-data-lake-store>



## QUESTION 10

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 scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are developing a solution that will use Azure Stream Analytics. The solution will accept an Azure Blob storage file named Customers. The file will contain both in-store and online customer details. The online customers will provide a mailing address.

You have a file in Blob storage named LocationIncomes that contains median incomes based on location. The file rarely changes.

You need to use an address to look up a median income based on location. You must output the data to Azure SQL Database for immediate use and to Azure Data Lake Storage Gen2 for long-term retention.

Solution: You implement a Stream Analytics job that has two streaming inputs, one query, and two outputs.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

We need one reference data input for LocationIncomes, which rarely changes

Note: Stream Analytics also supports input known as reference data. Reference data is either completely static or changes slowly.

Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-add-inputs#stream-and-reference-inputs>

---

## QUESTION 11

DRAG DROP

You deploy an Azure SQL database named DB1 to an Azure SQL server named SQL1.

Currently, only the server admin has access to DB1.

An Azure Active Directory (Azure AD) group named Analysts contains all the users who must have access to DB1.

You have the following data security requirements:

The Analysts group must have read-only access to all the views and tables in the Sales schema of DB1.



A manager will decide who can access DB1. The manager will not interact directly with DB1.

Users must not have to manage a separate password solely to access DB1.

Which four actions should you perform in sequence to meet the data security requirements? 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

- From the Azure portal, set the Active Directory admin for SQL1.
- From Microsoft SQL Server Management Studio (SSMS), set the Active Directory admin by using Transact-SQL.
- On DB1, grant the SELECT and VIEW DEFINITION permissions on the Sales schema to the Analysts group user.
- Create a login for each Azure AD user in the Analysts group by using Transact-SQL.
- From Microsoft SQL Server Management Studio (SSMS), sign in to SQL1 by using the account set as the Active Directory admin.
- On DB1, create a contained user for the Analysts group by using Transact-SQL.



Correct Answer:



### Actions

From Microsoft SQL Server Management Studio (SSMS), set the Active Directory admin by using Transact-SQL.

Create a login for each Azure AD user in the Analysts group by using Transact-SQL.

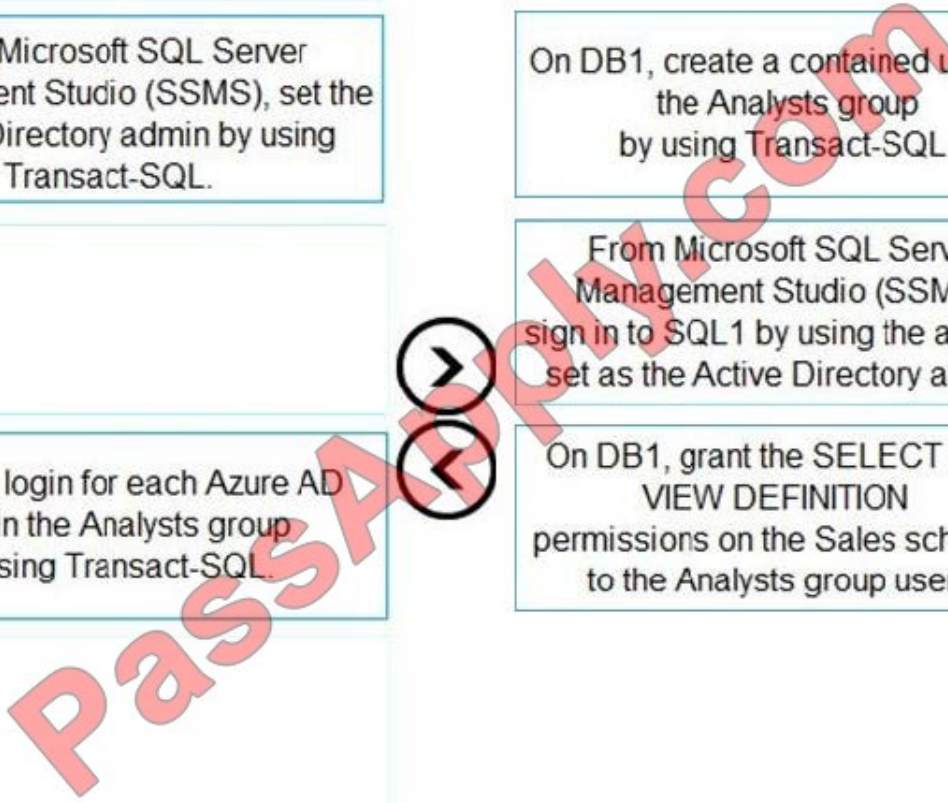
### Answer Area

From the Azure portal, set the Active Directory admin for SQL1.

On DB1, create a contained user for the Analysts group by using Transact-SQL.

From Microsoft SQL Server Management Studio (SSMS), sign in to SQL1 by using the account set as the Active Directory admin.

On DB1, grant the SELECT and VIEW DEFINITION permissions on the Sales schema to the Analysts group user.



### QUESTION 12

You have an Azure Databricks resource.

You need to log actions that relate to compute changes triggered by the Databricks resources.

Which Databricks services should you log?

- A. workspace
- B. SSH



C. DBFS

D. clusters

E. jobs

Correct Answer: D

An Azure Databricks cluster is a set of computation resources and configurations on which you run data engineering, data science, and data analytics workloads. Incorrect Answers:

A: An Azure Databricks workspace is an environment for accessing all of your Azure Databricks assets. The workspace organizes objects (notebooks, libraries, and experiments) into folders, and provides access to data and computational resources such as clusters and jobs.

B: SSH allows you to log into Apache Spark clusters remotely.

C: Databricks File System (DBFS) is a distributed file system mounted into an Azure Databricks workspace and available on Azure Databricks clusters.

E: A job is a way of running a notebook or JAR either immediately or on a scheduled basis.

Reference: <https://docs.microsoft.com/en-us/azure/databricks/clusters/>

[DP-200 VCE 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.passapply.com/allproducts>

## Need Help

Please provide as much detail as possible so we can best assist you.  
To update a previously submitted ticket:



 <p><b>One Year Free Update</b> 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>	 <p><b>Money Back Guarantee</b> To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p><b>Security &amp; Privacy</b> We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information &amp; 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 © passapply, All Rights Reserved.