



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.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

DRAG DROP

You have an Azure data factory.

You need to ensure that pipeline-run data is retained for 120 days. The solution must ensure that you can query the data by using the Kusto query language.

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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:

Actions

Answer Area

- Select the PipelineRuns category.
- Create a Log Analytics workspace that has Data Retention set to 120 days.
- Stream to an Azure event hub.
- Create an Azure Storage account that has a lifecycle policy.
- From the Azure portal, add a diagnostic setting.
- Send the data to a Log Analytics workspace.
- Select the TriggerRuns category.

Correct Answer:



Actions

Answer Area

Select the PipelineRuns category.	Create an Azure Storage account that has a lifecycle policy.
	Create a Log Analytics workspace that has Data Retention set to 120 days.
Stream to an Azure event hub.	From the Azure portal, add a diagnostic setting.
	Send the data to a Log Analytics workspace.
Select the TriggerRuns category.	

Step 1: Create an Azure Storage account that has a lifecycle policy

To automate common data management tasks, Microsoft created a solution based on Azure Data Factory. The service, Data Lifecycle Management, makes frequently accessed data available and archives or purges other data according to retention policies. Teams across the company use the service to reduce storage costs, improve app performance, and comply with data retention policies.

Step 2: Create a Log Analytics workspace that has Data Retention set to 120 days.

Data Factory stores pipeline-run data for only 45 days. Use Azure Monitor if you want to keep that data for a longer time. With Monitor, you can route diagnostic logs for analysis to multiple different targets, such as a Storage Account: Save

your diagnostic logs to a storage account for auditing or manual inspection. You can use the diagnostic settings to specify the retention time in days.

Step 3: From Azure Portal, add a diagnostic setting.

Step 4: Send the data to a log Analytics workspace,

Event Hub: A pipeline that transfers events from services to Azure Data Explorer.

Keeping Azure Data Factory metrics and pipeline-run data.



Configure diagnostic settings and workspace.

Create or add diagnostic settings for your data factory.

In the portal, go to Monitor. Select Settings > Diagnostic settings.

Select the data factory for which you want to set a diagnostic setting.

If no settings exist on the selected data factory, you'll be prompted to create a setting. Select Turn on diagnostics.

Give your setting a name, select Send to Log Analytics, and then select a workspace from Log Analytics Workspace.

Select Save.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/monitor-using-azure-monitor>

QUESTION 2

You manage an enterprise data warehouse in Azure Synapse Analytics.

Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries.

You need to monitor resource utilization to determine the source of the performance issues.

Which metric should you monitor?

- A. Data Warehouse Units (DWU) used
- B. DWU limit
- C. Cache hit percentage
- D. Data IO percentage

Correct Answer: C

The Azure Synapse Analytics storage architecture automatically tiers your most frequently queried columnstore segments in a cache residing on NVMe based SSDs designed for Gen2 data warehouses. Greater performance is realized when your queries retrieve segments that are residing in the cache. You can monitor and troubleshoot slow query performance by determining whether your workload is optimally leveraging the Gen2 cache.

Note: As of November 2019, Azure SQL Data Warehouse is now Azure Synapse Analytics. Reference:
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-how-to-monitor-cache>

<https://docs.microsoft.com/bs-latn-ba/azure/sql-data-warehouse/sql-data-warehouse-concept-resource-utilization-query-activity>

QUESTION 3

DRAG DROP



You need to replace the SSIS process by using Data Factory.

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	Answer Area
Create an Azure-SSIS integration runtime.	
Create a link service to each database.	
Add a stored procedure activity.	⬅
Create a pipeline.	➡
Add a copy activity.	
Create two datasets.	⬆
	⬇

Correct Answer:

Actions	Answer Area
Create an Azure-SSIS integration runtime.	Create a link service to each database.
	Create two datasets.
Add a stored procedure activity.	Create a pipeline.
	Add a copy activity.

Scenario: A daily process creates reporting data in REPORTINGDB from the data in SALESDB. The process is implemented as a SQL Server Integration Services (SSIS) package that runs a stored procedure from SALESDB.

Step 1: Create a linked service to each database

Step 2: Create two datasets

You can create two datasets: InputDataset and OutputDataset. These datasets are of type AzureBlob. They refer to the Azure Storage linked service that you created in the previous section.

Step 3: Create a pipeline



You create and validate a pipeline with a copy activity that uses the input and output datasets.

Step 4: Add a copy activity

References:

<https://docs.microsoft.com/en-us/azure/data-factory/quickstart-create-data-factory-portal>

QUESTION 4

You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL. Which switch should you use to switch between languages?

- A. %
- B. \[]
- C. \()
- D. @

Correct Answer: A

You can override the primary language by specifying the language magic command % at the beginning of a cell. The supported magic commands are: %python, %r, %scala, and %sql.

References: <https://docs.databricks.com/user-guide/notebooks/notebook-use.html#mix-languages>

QUESTION 5

A company plans to use Azure SQL Database to support a mission-critical application.

The application must be highly available without performance degradation during maintenance windows.

You need to implement the solution.

Which three technologies should you implement? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

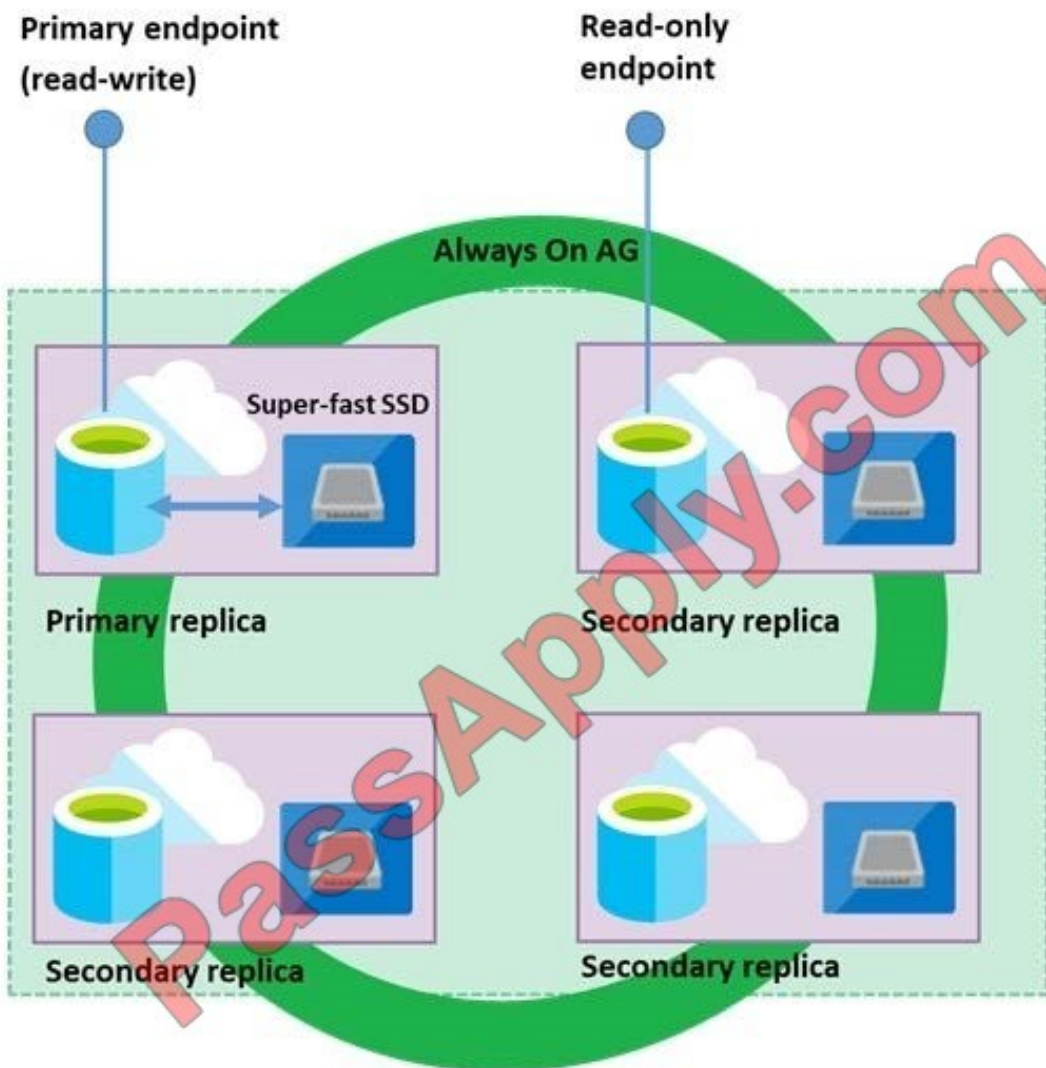
- A. Premium service tier
- B. Virtual machine Scale Sets
- C. Basic service tier
- D. SQL Data Sync
- E. Always On availability groups
- F. Zone-redundant configuration

Correct Answer: AEF



A: Premium/business critical service tier model that is based on a cluster of database engine processes. This architectural model relies on a fact that there is always a quorum of available database engine nodes and has minimal performance impact on your workload even during maintenance activities.

E: In the premium model, Azure SQL database integrates compute and storage on the single node. High availability in this architectural model is achieved by replication of compute (SQL Server Database Engine process) and storage (locally attached SSD) deployed in 4-node cluster, using technology similar to SQL Server Always On Availability Groups.



Business Critical service tier: collocated compute and storage

F: Zone redundant configuration By default, the quorum-set replicas for the local storage configurations are created in the same datacenter. With the introduction of Azure Availability Zones, you have the ability to place the different replicas in the quorum-sets to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-high-availability>



VCE & PDF

PassApply.com

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

2021 Latest passapply DP-200 PDF and VCE dumps Download

[DP-200 VCE Dumps](#)

[DP-200 Practice Test](#)

[DP-200 Study Guide](#)



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>One Year Free Update 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>Money Back Guarantee 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>Security & Privacy 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 © passapply, All Rights Reserved.