



DP-201^{Q&As}

Designing an Azure Data Solution

Pass Microsoft DP-201 Exam with 100% Guarantee

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

<https://www.pass4itsure.com/dp-201.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**

A company has an application that uses Azure SQL Database as the data store.

The application experiences a large increase in activity during the last month of each year.

You need to manually scale the Azure SQL Database instance to account for the increase in data write operations.

Which scaling method should you recommend?

- A. Scale up by using elastic pools to distribute resources.
- B. Scale out by sharding the data across databases.
- C. Scale up by increasing the database throughput units.

Correct Answer: C

As of now, the cost of running an Azure SQL database instance is based on the number of Database Throughput Units (DTUs) allocated for the database. When determining the number of units to allocate for the solution, a major contributing

factor is to identify what processing power is needed to handle the volume of expected requests.

Running the statement to upgrade/downgrade your database takes a matter of seconds.

Incorrect Answers:

A: Elastic pools is used if there are two or more databases.

References: https://www.skylinetechnologies.com/Blog/Skyline-Blog/August_2017/dynamically-scale-azure-sql-database

QUESTION 2

You plan to use Azure SQL Database to support a line of business app.

You need to identify sensitive data that is stored in the database and monitor access to the data.

Which three actions should you recommend? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable Data Discovery and Classification.
- B. Implement Transparent Data Encryption (TDE).
- C. Enable Auditing.
- D. Run Vulnerability Assessment.
- E. Use Advanced Threat Protection.



Correct Answer: CDE

QUESTION 3

You need to recommend a security solution for containers in Azure Blob storage. The solution must ensure that only read permissions are granted to a specific user for a specific container. What should you include in the recommendation?

- A. shared access signatures (SAS)
- B. an RBAC role in Azure Active Directory (Azure AD)
- C. public read access for blobs only
- D. access keys

Correct Answer: A

You can delegate access to read, write, and delete operations on blob containers, tables, queues, and file shares that are not permitted with a service SAS.

Note: A shared access signature (SAS) provides secure delegated access to resources in your storage account without compromising the security of your data. With a SAS, you have granular control over how a client can access your data.

You can control what resources the client may access, what permissions they have on those resources, and how long the SAS is valid, among other parameters.

Incorrect Answers:

C: You can enable anonymous, public read access to a container and its blobs in Azure Blob storage. By doing so, you can grant read-only access to these resources without sharing your account key, and without requiring a shared access signature (SAS).

Public read access is best for scenarios where you want certain blobs to always be available for anonymous read access.

References: <https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

QUESTION 4

You have a Windows-based solution that analyzes scientific data. You are designing a cloud-based solution that performs real-time analysis of the data.

You need to design the logical flow for the solution.

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

NOTE: Each correct selection is worth one point.

- A. Send data from the application to an Azure Stream Analytics job.
- B. Use an Azure Stream Analytics job on an edge device. Ingress data from an Azure Data Factory instance and build queries that output to Power BI.



C. Use an Azure Stream Analytics job in the cloud. Ingress data from the Azure Event Hub instance and build queries that output to Power BI.

D. Use an Azure Stream Analytics job in the cloud. Ingress data from an Azure Event Hub instance and build queries that output to Azure Data Lake Storage.

E. Send data from the application to Azure Data Lake Storage.

F. Send data from the application to an Azure Event Hub instance.

Correct Answer: CF

Stream Analytics has first-class integration with Azure data streams as inputs from three kinds of resources:

1.

Azure Event Hubs

2.

Azure IoT Hub

3.

Azure Blob storage

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

QUESTION 5

You need to design the data loading pipeline for Planning Assistance.

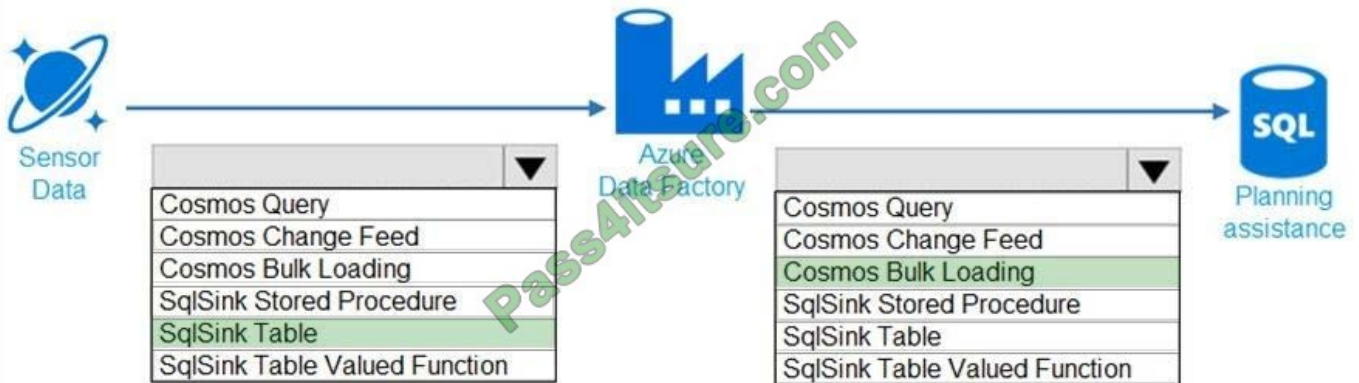
What should you recommend? To answer, drag the appropriate technologies to the correct locations. Each technology may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Correct Answer:

Answer Area

Box 1: SqlSink Table

Sensor data must be stored in a Cosmos DB named treydata in a collection named SensorData

Box 2: Cosmos Bulk Loading

Use Copy Activity in Azure Data Factory to copy data from and to Azure Cosmos DB (SQL API).

Scenario: Data from the Sensor Data collection will automatically be loaded into the Planning Assistance database once a week by using Azure Data Factory. You must be able to manually trigger the data load process.

Data used for Planning Assistance must be stored in a sharded Azure SQL Database.

References:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-cosmos-db>

QUESTION 6**HOTSPOT**

You use Azure Data Lake Storage Gen2 to store data that data scientists and data engineers will query by using Azure



Databricks interactive notebooks. The folders in Data Lake Storage will be secured, and users will have access only to the

folders that relate to the projects on which they work.

You need to recommend which authentication methods to use for Databricks and Data Lake Storage to provide the users with the appropriate access. The solution must minimize administrative effort and development effort

Which authentication method should you recommend for each Azure service? To answer, select the appropriate options in the answer area

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Databricks:

	▼
Azure Active Directory	
Azure Key Vault secrets	
Personal access tokens	

Data Lake Storage:

	▼
Azure Active Directory	
Shared access keys	
Shared access signatures	

Correct Answer:



Answer Area

Databricks:

	▼
Azure Active Directory	
Azure Key Vault secrets	
Personal access tokens	

Data Lake Storage:

	▼
Azure Active Directory	
Shared access keys	
Shared access signatures	

Databricks: Personal access tokens

To authenticate and access Databricks REST APIs, you use personal access tokens. Tokens are similar to passwords; you should treat them with care. Tokens expire and can be revoked.

Data Lake Storage: Azure Active Directory

Azure Data Lake Storage Gen1 uses Azure Active Directory for authentication.

References:

<https://docs.azuredatabricks.net/dev-tools/api/latest/authentication.html>

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lakes-store-authentication-using-azure-active-directory>

QUESTION 7

DRAG DROP

You are designing a Spark job that performs batch processing of daily web log traffic.

When you deploy the job in the production environment, it must meet the following requirements:

1.

Run once a day.

2.

Display status information on the company intranet as the job runs.



You need to recommend technologies for triggering and monitoring jobs.

Which technologies should you recommend? To answer, drag the appropriate technologies to the correct locations. Each technology may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll

to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Technologies

Livy

Beeline

Azure Logic App

Azure API App

Requirement

Triggering of jobs

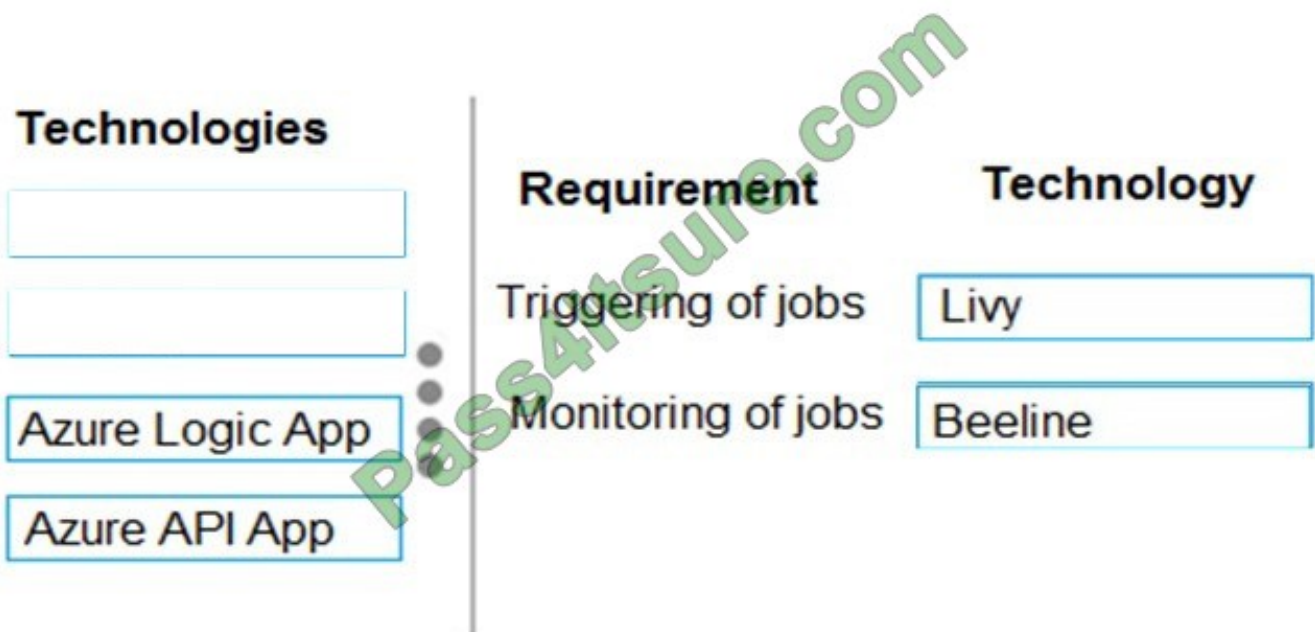
Monitoring of jobs

Technology

Correct Answer:



Answer Area



Box 1: Livy

You can use Livy to run interactive Spark shells or submit batch jobs to be run on Spark.

Box 2: Beeline

Apache Beeline can be used to run Apache Hive queries on HDInsight. You can use Beeline with Apache Spark.

Note: Beeline is a Hive client that is included on the head nodes of your HDInsight cluster. Beeline uses JDBC to connect to HiveServer2, a service hosted on your HDInsight cluster. You can also use Beeline to access Hive on HDInsight

remotely over the internet.

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-livy-rest-interface>

<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/apache-hadoop-use-hive-beeline>

QUESTION 8

You need to design the image processing solution to meet the optimization requirements for image tag data.

What should you configure? To answer, drag the appropriate setting to the correct drop targets.

Each source may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.



Select and Place:

Answer Area

	Configurations	Location	Configuration
		New York	<input type="text"/>
<input type="text" value="Write region"/>		Manchester	<input type="text"/>
<input type="text" value="Read region"/>		Singapore	<input type="text"/>
		Melbourne	<input type="text"/>

Correct Answer:

Answer Area

	Configurations	Location	Configuration
		New York	<input type="text" value="Write region"/>
<input type="text" value="Write region"/>		Manchester	<input type="text" value="Read region"/>
<input type="text" value="Read region"/>		Singapore	<input type="text" value="Read region"/>
		Melbourne	<input type="text" value="Read region"/>

Tagging data must be uploaded to the cloud from the New York office location.

Tagging data must be replicated to regions that are geographically close to company office locations.

QUESTION 9



You need to design a real-time stream solution that uses Azure Functions to process data uploaded to Azure Blob Storage. The solution must meet the following requirements:

1.

Support up to 1 million blobs.

2.

Scaling must occur automatically.

3.

Costs must be minimized. What should you recommend?

A. Deploy the Azure Function in an App Service plan and use a Blob trigger.

B. Deploy the Azure Function in a Consumption plan and use an Event Grid trigger.

C. Deploy the Azure Function in a Consumption plan and use a Blob trigger.

D. Deploy the Azure Function in an App Service plan and use an Event Grid trigger.

Correct Answer: C

Create a function, with the help of a blob trigger template, which is triggered when files are uploaded to or updated in Azure Blob storage. You use a consumption plan, which is a hosting plan that defines how resources are allocated to your function app. In the default Consumption Plan, resources are added dynamically as required by your functions. In this serverless hosting, you only pay for the time your functions run. When you run in an App Service plan, you must manage the scaling of your function app.

References: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function>

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

You plan to store delimited text files in an Azure Data Lake Storage account that will be organized into department folders.

You need to configure data access so that users see only the files in their respective department folder.

Solution: From the storage account, you disable a hierarchical namespace, and you use access control lists (ACLs).

Does this meet the goal?

A. Yes



B. No

Correct Answer: A

Azure Data Lake Storage implements an access control model that derives from HDFS, which in turn derives from the POSIX access control model.

Blob container ACLs does not support the hierarchical namespace, so it must be disabled.

References:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-known-issues>

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-access-control>

QUESTION 11

HOTSPOT

You manage an on-premises server named Server1 that has a database named Database1. The company purchases a new application that can access data from Azure SQL Database.

You recommend a solution to migrate Database1 to an Azure SQL Database instance.

What should you recommend? To answer, select the appropriate configuration in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

File type for exporting the on-premises database

BACPAC
DACPAC
VHDX

Azure storage type for exported data

Blob
Disk
Table
File

Correct Answer:



File type for exporting the on-premises database

BACPAC
DACPAC
VHDX

Azure storage type for exported data

Blob
Disk
Table
File

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

QUESTION 12

You are designing a recovery strategy for your Azure SQL Databases.

The recovery strategy must use default automated backup settings. The solution must include a Point-in time restore recovery strategy.

You need to recommend which backups to use and the order in which to restore backups.

What should you recommend? To answer, select the appropriate configuration in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

Restore order

Backup type

first

▼
full weekly backup
full daily backup
differential weekly backup
differential daily backup

second

▼
full daily backup
differential backup from the last 12 hours
all differential backups since the last full backup
all log backups since the last full backup

third

▼
all log backups since the last differential backup
differential backup from the last 12 hours
all differential backups since the last full backup
all log backups since the last full backup

Correct Answer:



Answer Area

Restore order

Backup type

first

▼
full weekly backup
full daily backup
differential weekly backup
differential daily backup

second

▼
full daily backup
differential backup from the last 12 hours
all differential backups since the last full backup
all log backups since the last full backup

third

▼
all log backups since the last differential backup
differential backup from the last 12 hours
all differential backups since the last full backup
all log backups since the last full backup

All Basic, Standard, and Premium databases are protected by automatic backups. Full backups are taken every week, differential backups every day, and log backups every 5 minutes.

References: <https://azure.microsoft.com/sv-se/blog/azure-sql-database-point-in-time-restore/>

[DP-201 PDF Dumps](#)

[DP-201 VCE Dumps](#)

[DP-201 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.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 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.	 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.	 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.
---	---	--

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.