

# Microsoft

## Exam Questions DP-201

Designing an Azure Data Solution



**NEW QUESTION 1**

- (Exam Topic 1)

You need to design the Planning Assistance database.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statement	Yes	No
Including a clustered columnstore index in the design will benefit performance.	<input type="radio"/>	<input type="radio"/>
Including a nonclustered columnstore index in the design will benefit performance.	<input type="radio"/>	<input type="radio"/>
Including an index on the License Plate column will benefit performance.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: No

Data used for Planning Assistance must be stored in a sharded Azure SQL Database. Box 2: Yes

Box 3: Yes

Planning Assistance database will include reports tracking the travel of a single vehicle

**NEW QUESTION 2**

- (Exam Topic 1)

You need to design the vehicle images storage solution. What should you recommend?

- A. Azure Media Services
- B. Azure Premium Storage account
- C. Azure Redis Cache
- D. Azure Cosmos DB

**Answer: B**

**Explanation:**

Premium Storage stores data on the latest technology Solid State Drives (SSDs) whereas Standard Storage stores data on Hard Disk Drives (HDDs). Premium Storage is designed for Azure Virtual Machine workloads which require consistent high IO performance and low latency in order to host IO intensive workloads like OLTP, Big Data, and Data Warehousing on platforms like SQL Server, MongoDB, Cassandra, and others. With Premium Storage, more customers will be able to lift-and-shift demanding enterprise applications to the cloud.

Scenario: Traffic sensors will occasionally capture an image of a vehicle for debugging purposes. You must optimize performance of saving/storing vehicle images. The impact of vehicle images on sensor data throughout must be minimized. References:

<https://azure.microsoft.com/es-es/blog/introducing-premium-storage-high-performance-storage-for-azure-virtual>

**NEW QUESTION 3**

- (Exam Topic 1)

You need to ensure that performance requirements for Backtrack reports are met.

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.

Technologies	Answer Area	
	Requirement	Technology
Cosmos DB TTL		
Cosmos DB indexes		
Cosmos DB transactions	Backtrack reporting	<input type="text"/>
Cosmos DB change feed	Privacy and security policy	<input type="text"/>
Cosmos DB stored procedures		

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Cosmos DB indexes

The report for Backtrack must execute as quickly as possible.

You can override the default indexing policy on an Azure Cosmos container, this could be useful if you want to tune the indexing precision to improve the query performance or to reduce the consumed storage.

Box 2: Cosmos DB TTL

This solution reports on all data related to a specific vehicle license plate. The report must use data from the SensorData collection. Users must be able to filter vehicle data in the following ways:

- vehicles on a specific road
- vehicles driving above the speed limit

Note: With Time to Live or TTL, Azure Cosmos DB provides the ability to delete items automatically from a container after a certain time period. By default, you can set time to live at the container level and override the value on a per-item basis. After you set the TTL at a container or at an item level, Azure Cosmos DB will automatically remove these items after the time period, since the time they were last modified.

#### NEW QUESTION 4

- (Exam Topic 1)

You need to design the runtime environment for the Real Time Response system. What should you recommend?

- A. General Purpose nodes without the Enterprise Security package
- B. Memory Optimized Nodes without the Enterprise Security package
- C. Memory Optimized nodes with the Enterprise Security package
- D. General Purpose nodes with the Enterprise Security package

**Answer: B**

#### NEW QUESTION 5

- (Exam Topic 1)

You need to recommend an Azure SQL Database pricing tier for Planning Assistance. Which pricing tier should you recommend?

- A. Business critical Azure SQL Database single database
- B. General purpose Azure SQL Database Managed Instance
- C. Business critical Azure SQL Database Managed Instance
- D. General purpose Azure SQL Database single database

**Answer: B**

#### Explanation:

Azure resource costs must be minimized where possible.

Data used for Planning Assistance must be stored in a sharded Azure SQL Database. The SLA for Planning Assistance is 70 percent, and multiday outages are permitted.

#### NEW QUESTION 6

- (Exam Topic 1)

You need to design a sharding strategy for the Planning Assistance database. What should you recommend?

- A. a list mapping shard map on the binary representation of the License Plate column
- B. a range mapping shard map on the binary representation of the speed column
- C. a list mapping shard map on the location column
- D. a range mapping shard map on the time column

**Answer: A**

#### Explanation:

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

A shard typically contains items that fall within a specified range determined by one or more attributes of the data. These attributes form the shard key (sometimes referred to as the partition key). The shard key should be static. It shouldn't be based on data that might change.

References:

<https://docs.microsoft.com/en-us/azure/architecture/patterns/sharding>

#### NEW QUESTION 7

- (Exam Topic 2)

You plan to use an Azure SQL data warehouse to store the customer data. You need to recommend a disaster recovery solution for the data warehouse. What should you include in the recommendation?

- A. AzCopy
- B. Read-only replicas
- C. AdlCopy
- D. Geo-Redundant backups

**Answer: D**

#### Explanation:

References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore>

#### NEW QUESTION 8

- (Exam Topic 2)

You need to design the solution for analyzing customer data. What should you recommend?

- A. Azure Databricks
- B. Azure Data Lake Storage
- C. Azure SQL Data Warehouse
- D. Azure Cognitive Services
- E. Azure Batch

**Answer:** A

**Explanation:**

Customer data must be analyzed using managed Spark clusters. You create spark clusters through Azure Databricks. References: <https://docs.microsoft.com/en-us/azure/azure-databricks/quickstart-create-databricks-workspace-portal>

**NEW QUESTION 9**

- (Exam Topic 2)

You need to design a backup solution for the processed customer data. What should you include in the design?

- A. AzCopy
- B. AdlCopy
- C. Geo-Redundancy
- D. Geo-Replication

**Answer:** C

**Explanation:**

Scenario: All data must be backed up in case disaster recovery is required.

Geo-redundant storage (GRS) is designed to provide at least 99.99999999999999% (16 9's) durability of objects over a given year by replicating your data to a secondary region that is hundreds of miles away from the primary region. If your storage account has GRS enabled, then your data is durable even in the case of a complete regional outage or a disaster in which the primary region isn't recoverable. References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs>

**NEW QUESTION 10**

- (Exam Topic 2)

You need to design the image processing and storage solutions.

What should you recommend? To answer, select the appropriate configuration in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

Solution component	Tool
Image processing	<ul style="list-style-type: none"> <li>Azure HDInsight</li> <li>Azure Databricks</li> <li>Azure Batch</li> <li>Azure Cognitive Services</li> </ul>
data storage for tagging data	<ul style="list-style-type: none"> <li>Azure Blob Storage</li> <li>Azure Table Storage</li> <li>Azure Cosmos DB</li> <li>Azure SQL Database</li> </ul>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/batch-processing> <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tier-hyperscale>

**NEW QUESTION 10**

- (Exam Topic 3)

You need to recommend the appropriate storage and processing solution? What should you recommend?

- A. Enable auto-shrink on the database.
- B. Flush the blob cache using Windows PowerShell.
- C. Enable Apache Spark RDD (RDD) caching.
- D. Enable Databricks IO (DBIO) caching.
- E. Configure the reading speed using Azure Data Studio.

**Answer:** C

**Explanation:**

Scenario: You must be able to use a file system view of data stored in a blob. You must build an architecture that will allow Contoso to use the DB FS filesystem layer over a blob store.

Databricks File System (DBFS) is a distributed file system installed on Azure Databricks clusters. Files in DBFS persist to Azure Blob storage, so you won't lose data even after you terminate a cluster.

The Databricks Delta cache, previously named Databricks IO (DBIO) caching, accelerates data reads by creating copies of remote files in nodes' local storage using a fast intermediate data format. The data is cached automatically whenever a file has to be fetched from a remote location. Successive reads of the same data are then performed locally, which results in significantly improved reading speed.

**NEW QUESTION 15**

- (Exam Topic 3)

You are designing an Azure SQL Data Warehouse for a financial services company. Azure Active Directory will be used to authenticate the users. You need to ensure that the following security requirements are met:

- ▶ Department managers must be able to create new database.
- ▶ The IT department must assign users to databases.
- ▶ Permissions granted must be minimized.

Which role memberships should you recommend? To answer, drag the appropriate roles to the correct groups. Each role 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.

Roles	Group	Role
dbmanager	Department managers	<input type="text"/>
loginmanager	IT	<input type="text"/>
dc_admin		
db_securityadmin		
db_owner		
db_accessadmin		

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: dbmanager

Members of the dbmanager role can create new databases. Box 2: db\_accessadmin

Members of the db\_accessadmin fixed database role can add or remove access to the database for Windows logins, Windows groups, and SQL Server logins.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-manage-logins>

**NEW QUESTION 19**

- (Exam Topic 4)

You design data engineering solutions for a company.

You must integrate on-premises SQL Server data into an Azure solution that performs Extract-Transform-Load (ETL) operations have the following requirements:

- ▶ Develop a pipeline that can integrate data and run notebooks.
- ▶ Develop notebooks to transform the data.
- ▶ Load the data into a massively parallel processing database for later analysis. You need to recommend a solution.

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

NOTE: Each correct selection is worth one point.

Requirement	Service
Integrate the on-premises data into the cloud.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>
Develop notebooks to transform the data.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>
Run notebooks.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>
Load the data.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>
Store the transformed data.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Requirement	Service
Integrate the on-premises data into the cloud.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>
Develop notebooks to transform the data.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>
Run notebooks.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>
Load the data.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>
Store the transformed data.	<ul style="list-style-type: none"> <li>Azure Databricks</li> <li>Azure Data Factory</li> <li>Azure SQL Data Warehouse</li> <li>Azure Batch</li> </ul>

**NEW QUESTION 20**

- (Exam Topic 4)

You need to design the unauthorized data usage detection system. What Azure service should you include in the design?

- A. Azure Databricks
- B. Azure SQL Data Warehouse
- C. Azure Analysis Services
- D. Azure Data Factory

**Answer: B**

**NEW QUESTION 23**

- (Exam Topic 4)

You are designing a solution for a company. You plan to use Azure Databricks. You need to recommend workloads and tiers to meet the following requirements:

- ▶ Provide managed clusters for running production jobs.
- ▶ Provide persistent clusters that support auto-scaling for analytics processes.
- ▶ Provide role-based access control (RBAC) support for Notebooks.

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

NOTE: Each correct selection is worth one point.

Requirement	Workload	Tier
Provide managed clusters for running production jobs.	<div style="border: 1px solid black; padding: 2px;">                     ▼ Standard                      Data Engineering only                      Data Analytics only                      Data Engineering and Data Analytics                 </div>	
Provide persistent clusters that support auto-scaling for analytics processes.	<div style="border: 1px solid black; padding: 2px;">                     ▼                      Data Engineering only                      Data Analytics only                      Data Engineering and Data Analytics                 </div>	<div style="border: 1px solid black; padding: 2px;">                     ▼                      Standard                      Premium                 </div>
Provide role-based access control (RBAC) support for Notebooks.	<div style="border: 1px solid black; padding: 2px;">                     ▼                      Data Engineering only                      Data Analytics only                      Data Engineering and Data Analytics                 </div>	<div style="border: 1px solid black; padding: 2px;">                     ▼                      Standard                      Premium                 </div>

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Data Engineering Only

Box 2: Data Engineering and Data Analytics Box 3: Standard

Box 4: Data Analytics only Box 5: Premium

Premium required for RBAC. Data Analytics Premium Tier provide interactive workloads to analyze data collaboratively with notebooks

References:

<https://azure.microsoft.com/en-us/pricing/details/databricks/>

**NEW QUESTION 27**

- (Exam Topic 4)

A company is evaluating data storage solutions.

You need to recommend a data storage solution that meets the following requirements: Minimize costs for storing blob objects.

Optimize access for data that is infrequently accessed. Data must be stored for at least 30 days.

Data availability must be at least 99 percent. What should you recommend?

- A. Premium
- B. Cold
- C. Hot
- D. Archive

**Answer: B**

**Explanation:**

Azure's cool storage tier, also known as Azure cool Blob storage, is for infrequently-accessed data that needs to be stored for a minimum of 30 days. Typical use cases include backing up data before tiering to archival systems, legal data, media files, system audit information, datasets used for big data analysis and more.

The storage cost for this Azure cold storage tier is lower than that of hot storage tier. Since it is expected that the data stored in this tier will be accessed less frequently, the data access charges are high when compared to hot tier. There are no additional changes required in your applications as these tiers can be accessed using

APIs in the same manner that you access Azure storage. References:

<https://cloud.netapp.com/blog/low-cost-storage-options-on-azure>

**NEW QUESTION 29**

- (Exam Topic 4)

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.

Restore order	Backup type
first	<div style="border: 1px solid black; padding: 2px;">                     full weekly backup                      full daily backup                      differential weekly backup                      differential daily backup                 </div>
second	<div style="border: 1px solid black; padding: 2px;">                     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                 </div>
third	<div style="border: 1px solid black; padding: 2px;">                     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                 </div>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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/>

**NEW QUESTION 30**

- (Exam Topic 4)

You plan to deploy an Azure SQL Database instance to support an application. You plan to use the DTUbased purchasing model.

Backups of the database must be available for 30 days and point-in-time restoration must be possible. You need to recommend a backup and recovery policy.

What are two possible ways to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Use the Premium tier and the default backup retention policy.
- B. Use the Basic tier and the default backup retention policy.
- C. Use the Standard tier and the default backup retention policy.
- D. Use the Standard tier and configure a long-term backup retention policy.
- E. Use the Premium tier and configure a long-term backup retention policy.

**Answer:** DE

**Explanation:**

The default retention period for a database created using the DTU-based purchasing model depends on the service tier:

- Basic service tier is 1 week.
- Standard service tier is 5 weeks.
- Premium service tier is 5 weeks.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-long-term-retention>

**NEW QUESTION 32**

- (Exam Topic 4)

You are designing a solution for a company. The solution will use model training for objective classification. You need to design the solution.

What should you recommend?

- A. an Azure Cognitive Services application
- B. a Spark Streaming job
- C. interactive Spark queries
- D. Power BI models
- E. a Spark application that uses Spark MLlib.

**Answer:** E

**Explanation:**

Spark in SQL Server big data cluster enables AI and machine learning.

You can use Apache Spark MLlib to create a machine learning application to do simple predictive analysis on an open dataset.

MLlib is a core Spark library that provides many utilities useful for machine learning tasks, including utilities that are suitable for:

- Classification
- Regression
- Clustering
- Topic modeling
- Singular value decomposition (SVD) and principal component analysis (PCA)
- Hypothesis testing and calculating sample statistics

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-machine-learning-mllib-ipython>

**NEW QUESTION 34**

- (Exam Topic 4)

You are designing a data processing solution that will run as a Spark job on an HDInsight cluster. The solution will be used to provide near real-time information about online ordering for a retailer.

The solution must include a page on the company intranet that displays summary information. The summary information page must meet the following requirements:

- ▶ Display a summary of sales to date grouped by product categories, price range, and review scope.
- ▶ Display sales summary information including total sales, sales as compared to one day ago and sales as compared to one year ago.
- ▶ Reflect information for new orders as quickly as possible. You need to recommend a design for the solution.

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

Use case	Technology
Data abstraction	<div style="border: 1px solid black; padding: 2px;"> <span style="float: right;">▼</span>                     Resilient Distributed Dataset (RDD)                      Dataset                      DataFrame                 </div>
Data format	<div style="border: 1px solid black; padding: 2px;"> <span style="float: right;">▼</span>                     Avro                      parquet                 </div>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: DataFrame

DataFrames

Best choice in most situations.

Provides query optimization through Catalyst. Whole-stage code generation.

Direct memory access.

Low garbage collection (GC) overhead.

Not as developer-friendly as DataSets, as there are no compile-time checks or domain object programming. Box 2: parquet

The best format for performance is parquet with snappy compression, which is the default in Spark 2.x. Parquet stores data in columnar format, and is highly optimized in Spark.

**NEW QUESTION 38**

- (Exam Topic 4)

You are designing an Azure Databricks cluster that runs user-defined local processes. You need to recommend a cluster configuration that meets the following requirements:

- Minimize query latency.
- Reduce overall costs.
- Maximize the number of users that can run queries on the cluster at the same time. Which cluster type should you recommend?

- A. Standard with Autoscaling
- B. High Concurrency with Auto Termination
- C. High Concurrency with Autoscaling
- D. Standard with Auto Termination

**Answer:** C

**Explanation:**

High Concurrency clusters allow multiple users to run queries on the cluster at the same time, while minimizing query latency. Autoscaling clusters can reduce overall costs compared to a statically-sized cluster.

References:

<https://docs.azuredatabricks.net/user-guide/clusters/create.html> <https://docs.azuredatabricks.net/user-guide/clusters/high-concurrency.html#high-concurrency>

<https://docs.azuredatabricks.net/user-guide/clusters/terminate.html> <https://docs.azuredatabricks.net/user-guide/clusters/sizing.html#enable-and-configure-autoscaling>

**NEW QUESTION 39**

- (Exam Topic 4)

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.

A company is developing a solution to manage inventory data for a group of automotive repair shops. The solution will use Azure SQL Data Warehouse as the data store. Shops will upload data every 10 days.

Data corruption checks must run each time data is uploaded. If corruption is detected, the corrupted data must be removed.

You need to ensure that upload processes and data corruption checks do not impact reporting and analytics processes that use the data warehouse.

Proposed solution: Insert data from shops and perform the data corruption check in a transaction. Rollback transfer if corruption is detected.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Instead, create a user-defined restore point before data is uploaded. Delete the restore point after data corruption checks complete.

References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore>

**NEW QUESTION 44**

- (Exam Topic 4)

You are evaluating data storage solutions to support a new application.

You need to recommend a data storage solution that represents data by using nodes and relationships in graph structures.

Which data storage solution should you recommend?

- A. Blob Storage
- B. Cosmos DB
- C. Data Lake Store
- D. HDInsight

**Answer:** B

**Explanation:**

For large graphs with lots of entities and relationships, you can perform very complex analyses very quickly. Many graph databases provide a query language that you can use to traverse a network of relationships efficiently.

Relevant Azure service: Cosmos DB

References:

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview>

**NEW QUESTION 45**

- (Exam Topic 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
- C. Ingress data from an Azure Data Factory instance and build queries that output to Power BI.
- D. Use an Azure Stream Analytics job in the cloud
- E. Ingress data from the Azure Event Hub instance and build queries that output to Power BI.
- F. Use an Azure Stream Analytics job in the cloud
- G. Ingress data from an Azure Event Hub instance and build queries that output to Azure Data Lake Storage.
- H. Send data from the application to Azure Data Lake Storage.
- I. Send data from the application to an Azure Event Hub instance.

**Answer:** CF

**Explanation:**

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

Azure IoT Hub Azure Blob storage References:

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

**NEW QUESTION 46**

- (Exam Topic 4)

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.

Option	Value
File type for exporting the on-premises database	<input type="text" value="BACPAC"/> <input type="text" value="DACPAC"/> <input type="text" value="VHDX"/>
Azure storage type for exported data	<input type="text" value="Blob"/> <input type="text" value="Disk"/> <input type="text" value="Table"/> <input type="text" value="File"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

References:

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

**NEW QUESTION 51**

- (Exam Topic 4)

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 are designing an Azure SQL Database that will use elastic pools. You plan to store data about customers in a table. Each record uses a value for CustomerID. You need to recommend a strategy to partition data based on values in CustomerID. Proposed Solution: Separate data into shards by using horizontal partitioning. Does the solution meet the goal?

- A. Yes
- B. No

**Answer: A**

**Explanation:**

Horizontal Partitioning - Sharding: Data is partitioned horizontally to distribute rows across a scaled out data tier. With this approach, the schema is identical on all participating databases. This approach is also called "sharding". Sharding can be performed and managed using (1) the elastic database tools libraries or (2) selfsharding.

An elastic query is used to query or compile reports across many shards. References:  
<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-query-overview>

**NEW QUESTION 54**

- (Exam Topic 4)

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 are designing an Azure SQL Database that will use elastic pools. You plan to store data about customers in a table. Each record uses a value for CustomerID. You need to recommend a strategy to partition data based on values in CustomerID. Proposed Solution: Separate data into customer regions by using vertical partitioning. Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Vertical partitioning is used for cross-database queries. Instead we should use Horizontal Partitioning, which also is called charding.

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

**NEW QUESTION 58**

- (Exam Topic 4)

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:

- ▶ Run once a day.
- ▶ 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.

Technologies	Requirement	Technology
Livy	Triggering of jobs	
Beeline	Monitoring of jobs	
Azure Logic App		
Azure API App		

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

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>

**NEW QUESTION 61**

- (Exam Topic 4)

You need to design the system for notifying law enforcement officers about speeding vehicles.

How should you design the pipeline? To answer, drag the appropriate services to the correct locations. Each service 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.

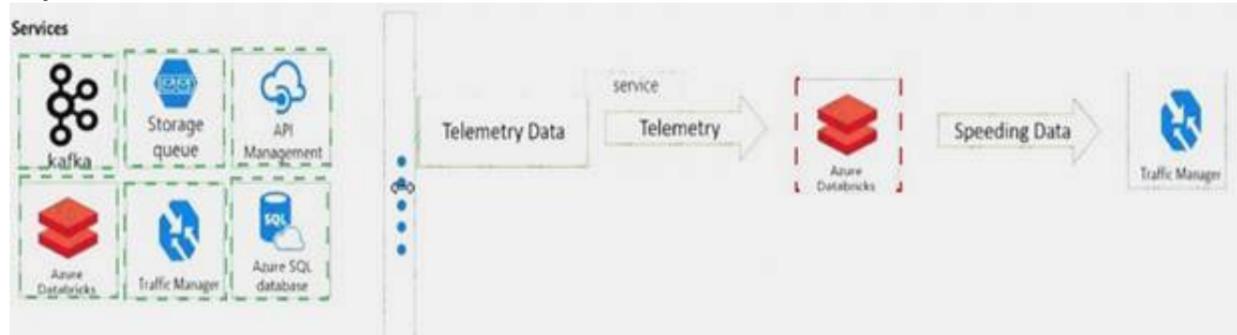
NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 63

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **DP-201 Practice Exam Features:**

- \* DP-201 Questions and Answers Updated Frequently
- \* DP-201 Practice Questions Verified by Expert Senior Certified Staff
- \* DP-201 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* DP-201 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The DP-201 Practice Test Here](#)**