



Actual exam question from Microsoft's DP-201

Question #: 1

Topic #: 1

[\[All DP-201 Questions\]](#)

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 HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage.

The solution requires POSIX permissions and enables diagnostics logging for auditing.

You need to recommend solutions that optimize storage.

Proposed Solution: Ensure that files stored are larger than 250MB.

Does the solution meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 2

Topic #: 1

[\[All DP-201 Questions\]](#)

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 HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage.

The solution requires POSIX permissions and enables diagnostics logging for auditing.

You need to recommend solutions that optimize storage.

Proposed Solution: Implement compaction jobs to combine small files into larger files.

Does the solution meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 3

Topic #: 1

[\[All DP-201 Questions\]](#)

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 HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage.

The solution requires POSIX permissions and enables diagnostics logging for auditing.

You need to recommend solutions that optimize storage.

Proposed Solution: Ensure that files stored are smaller than 250MB.

Does the solution meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 4

Topic #: 1

[\[All DP-201 Questions\]](#)

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

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 5

Topic #: 1

[\[All DP-201 Questions\]](#)

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 horizontal partitioning.

Does the solution meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 6

Topic #: 1

[\[All DP-201 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 7

Topic #: 1

[\[All DP-201 Questions\]](#)

HOTSPOT -

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.

Hot Area:

Answer Area

Use case	Technology
Data abstraction	<input type="text"/> Resilient Distributed Dataset (RDD) Dataset DataFrame
Data format	<input type="text"/> Avro parquet

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 8

Topic #: 1

[\[All DP-201 Questions\]](#)

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. Azure Cosmos DB
- C. Azure Data Lake Store
- D. HDInsight

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 9

Topic #: 1

[\[All DP-201 Questions\]](#)

HOTSPOT -

You have an on-premises data warehouse that includes the following fact tables. Both tables have the following columns: DataKey, ProductKey, RegionKey.

There are 120 unique product keys and 65 unique region keys.

Table	Comments
Sales	The table is 600 GB in size. DateKey is used extensively in the WHERE clause in queries. ProductKey is used extensively in join operations. RegionKey is used for grouping. Seventy-five percent of records relate to one of 40 regions.
Invoice	The table is 6 GB in size. DateKey and ProductKey are used extensively in the WHERE clause in queries. RegionKey is used for grouping.

Queries that use the data warehouse take a long time to complete.

You plan to migrate the solution to use Azure Synapse Analytics. You need to ensure that the Azure-based solution optimizes query performance and minimizes processing skew.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Table	Distribution type	Distribution column
Sales	<input type="text" value="Hash-distributed"/> <input type="text" value="Round-robin"/>	<input type="text" value="DateKey"/> <input type="text" value="ProductKey"/> <input type="text" value="RegionKey"/>
Invoices	<input type="text" value="Hash-distributed"/> <input type="text" value="Round-robin"/>	<input type="text" value="DateKey"/> <input type="text" value="ProductKey"/> <input type="text" value="RegionKey"/>

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 10

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a data processing solution that will implement the lambda architecture pattern. The solution will use Spark running on HDInsight for data processing.

You need to recommend a data storage technology for the solution.

Which two technologies should you recommend? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Azure Cosmos DB
- B. Azure Service Bus
- C. Azure Storage Queue
- D. Apache Cassandra
- E. Kafka HDInsight

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 11

Topic #: 1

[\[All DP-201 Questions\]](#)

A company manufactures automobile parts. The company installs IoT sensors on manufacturing machinery.

You must design a solution that analyzes data from the sensors.

You need to recommend a solution that meets the following requirements:

- ☞ Data must be analyzed in real-time.
- ☞ Data queries must be deployed using continuous integration.
- ☞ Data must be visualized by using charts and graphs.
- ☞ Data must be available for ETL operations in the future.
- ☞ The solution must support high-volume data ingestion.

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

NOTE: Each correct selection is worth one point.

- A. Use Azure Analysis Services to query the data. Output query results to Power BI.
- B. Configure an Azure Event Hub to capture data to Azure Data Lake Storage.
- C. Develop an Azure Stream Analytics application that queries the data and outputs to Power BI. Use Azure Data Factory to deploy the Azure Stream Analytics application.
- D. Develop an application that sends the IoT data to an Azure Event Hub.
- E. Develop an Azure Stream Analytics application that queries the data and outputs to Power BI. Use Azure Pipelines to deploy the Azure Stream Analytics application.
- F. Develop an application that sends the IoT data to an Azure Data Lake Storage container.

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 12

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing an Azure Databricks interactive cluster.

You need to ensure that the cluster meets the following requirements:

- ☞ Enable auto-termination
- ☞ Retain cluster configuration indefinitely after cluster termination.

What should you recommend?

- A. Start the cluster after it is terminated.
- B. Pin the cluster
- C. Clone the cluster after it is terminated.
- D. Terminate the cluster manually at process completion.

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 13

Topic #: 1

[\[All DP-201 Questions\]](#)

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.

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 14

Topic #: 1

[\[All DP-201 Questions\]](#)

A company stores data in multiple types of cloud-based databases.

You need to design a solution to consolidate data into a single relational database. Ingestion of data will occur at set times each day.

What should you recommend?

- A. SQL Server Migration Assistant
- B. SQL Data Sync
- C. Azure Data Factory
- D. Azure Database Migration Service
- E. Data Migration Assistant

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 15

Topic #: 1

[\[All DP-201 Questions\]](#)

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:

Answer Area

Option	Value				
File type for exporting the on-premises database	<table border="1"><tr><td>BACPAC</td></tr><tr><td>DACPAC</td></tr><tr><td>VHDX</td></tr></table>	BACPAC	DACPAC	VHDX	
BACPAC					
DACPAC					
VHDX					
Azure storage type for exported data	<table border="1"><tr><td>Blob</td></tr><tr><td>Disk</td></tr><tr><td>Table</td></tr><tr><td>File</td></tr></table>	Blob	Disk	Table	File
Blob					
Disk					
Table					
File					

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 16

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing an application. You plan to use Azure SQL Database to support the application.

The application will extract data from the Azure SQL Database and create text documents. The text documents will be placed into a cloud-based storage solution.

The text storage solution must be accessible from an SMB network share.

You need to recommend a data storage solution for the text documents.

Which Azure data storage type should you recommend?

- A. Queue
- B. Files
- C. Blob
- D. Table

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 17

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing an application that will have an Azure virtual machine. The virtual machine will access an Azure SQL database. The database will not be accessible from the Internet.

You need to recommend a solution to provide the required level of access to the database.

What should you include in the recommendation?

- A. Deploy an On-premises data gateway.
- B. Add a virtual network to the Azure SQL server that hosts the database.
- C. Add an application gateway to the virtual network that contains the Azure virtual machine.
- D. Add a virtual network gateway to the virtual network that contains the Azure virtual machine.

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 18

Topic #: 1

[\[All DP-201 Questions\]](#)

HOTSPOT -

You are designing an application that will store petabytes of medical imaging data

When the data is first created, the data will be accessed frequently during the first week. After one month, the data must be accessible within 30 seconds, but files will be accessed infrequently. After one year, the data will be accessed infrequently but must be accessible within five minutes.

You need to select a storage strategy for the data. The solution must minimize costs.

Which storage tier should you use for each time frame? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

First week:

	▼
Archive	
Cool	
Hot	

After one month:

	▼
Archive	
Cool	
Hot	

After one year:

	▼
Archive	
Cool	
Hot	

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 19

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a data store that will store organizational information for a company. The data will be used to identify the relationships between users. The data will be stored in an Azure Cosmos DB database and will contain several million objects.

You need to recommend which API to use for the database. The API must minimize the complexity to query the user relationships. The solution must support fast traversals.

Which API should you recommend?

- A. MongoDB
- B. Table
- C. Gremlin
- D. Cassandra

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 20

Topic #: 1

[\[All DP-201 Questions\]](#)

HOTSPOT -

You are designing a new application that uses Azure Cosmos DB. The application will support a variety of data patterns including log records and social media relationships.

You need to recommend which Cosmos DB API to use for each data pattern. The solution must minimize resource utilization.

Which API should you recommend for each data pattern? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Log records:

Cassandra
Gremlin
SQL

Social media mentions:

Cassandra
Gremlin
SQL

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 21

Topic #: 1

[\[All DP-201 Questions\]](#)

You need to recommend a storage solution to store flat files and columnar optimized files. The solution must meet the following requirements:

- ⇒ Store standardized data that data scientists will explore in a curated folder.
- ⇒ Ensure that applications cannot access the curated folder.
- ⇒ Store staged data for import to applications in a raw folder.
- ⇒ Provide data scientists with access to specific folders in the raw folder and all the content the curated folder.

Which storage solution should you recommend?

- A. Azure Synapse Analytics
- B. Azure Blob storage
- C. Azure Data Lake Storage Gen2
- D. Azure SQL Database

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 22

Topic #: 1

[\[All DP-201 Questions\]](#)

Your company is an online retailer that can have more than 100 million orders during a 24-hour period, 95 percent of which are placed between 16:30 and 17:00.

All the orders are in US dollars. The current product line contains the following three item categories:

- ⇒ Games with 15,123 items
- ⇒ Books with 35,312 items
- ⇒ Pens with 6,234 items

You are designing an Azure Cosmos DB data solution for a collection named Orders Collection. The following documents is a typical order in Orders Collection.

```
"OrderTime": "16:35",
"id": " d0379ca2-f912-5h7f-k159-340ffa1z18e4"
"Item": {
  "id": "08g17u57-1j58-6511-4x65-
2qb5bf723u5s",
  "Title": "Living the Data Dream",
  "Category": "Books",
  "PurchasePrice": 12.56,
  "Currency": "USD"
}
```

Orders Collection is expected to have a balanced read/write-intensive workload.

Which partition key provides the most efficient throughput?

- A. Item/Category
- B. OrderTime
- C. Item/Currency
- D. Item/id

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 23

Topic #: 1

[\[All DP-201 Questions\]](#)

You have a MongoDB database that you plan to migrate to an Azure Cosmos DB account that uses the MongoDB API.

During testing, you discover that the migration takes longer than expected.

You need to recommend a solution that will reduce the amount of time it takes to migrate the data.

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

NOTE: Each correct selection is worth one point.

- A. Increase the Request Units (RUs).
- B. Turn off indexing.
- C. Add a write region.
- D. Create unique indexes.
- E. Create compound indexes.

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 24

Topic #: 1

[\[All DP-201 Questions\]](#)

You need to recommend a storage solution for a sales system that will receive thousands of small files per minute. The files will be in JSON, text, and CSV formats. The files will be processed and transformed before they are loaded into a data warehouse in Azure Synapse Analytics. The files must be stored and secured in folders. Which storage solution should you recommend?

- A. Azure Data Lake Storage Gen2
- B. Azure Cosmos DB
- C. Azure SQL Database
- D. Azure Blob storage

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 25

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing an Azure Cosmos DB database that will support vertices and edges.

Which Cosmos DB API should you include in the design?

- A. SQL
- B. Cassandra
- C. Gremlin
- D. Table

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 26

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a big data storage solution. The solution must meet the following requirements:

- ⇒ Provide unlimited account sizes.
- ⇒ Support a hierarchical file system.
- ⇒ Be optimized for parallel analytics workloads.

Which storage solution should you use?

- A. Azure Data Lake Storage Gen2
- B. Azure Blob storage
- C. Apache HBase in Azure HDInsight
- D. Azure Cosmos DB

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 27

Topic #: 1

[\[All DP-201 Questions\]](#)

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 enable a hierarchical namespace, and you use RBAC.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 28

Topic #: 1

[\[All DP-201 Questions\]](#)

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 RBAC (role-based access control).

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 29

Topic #: 1

[\[All DP-201 Questions\]](#)

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

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 30

Topic #: 1

[\[All DP-201 Questions\]](#)

You plan to store 100 GB of data used by a line-of-business (LOB) app.

You need to recommend a data storage solution for the data. The solution must meet the following requirements:

- ⇒ Minimize storage costs.
- ⇒ Natively support relational queries.
- ⇒ Provide a recovery time objective (RTO) of less than one minute.

What should you include in the recommendation?

- A. Azure Cosmos DB
- B. Azure SQL Database
- C. Azure Synapse Analytics
- D. Azure Blob storage

Show Suggested Answer



Actual exam question from Microsoft's DP-201

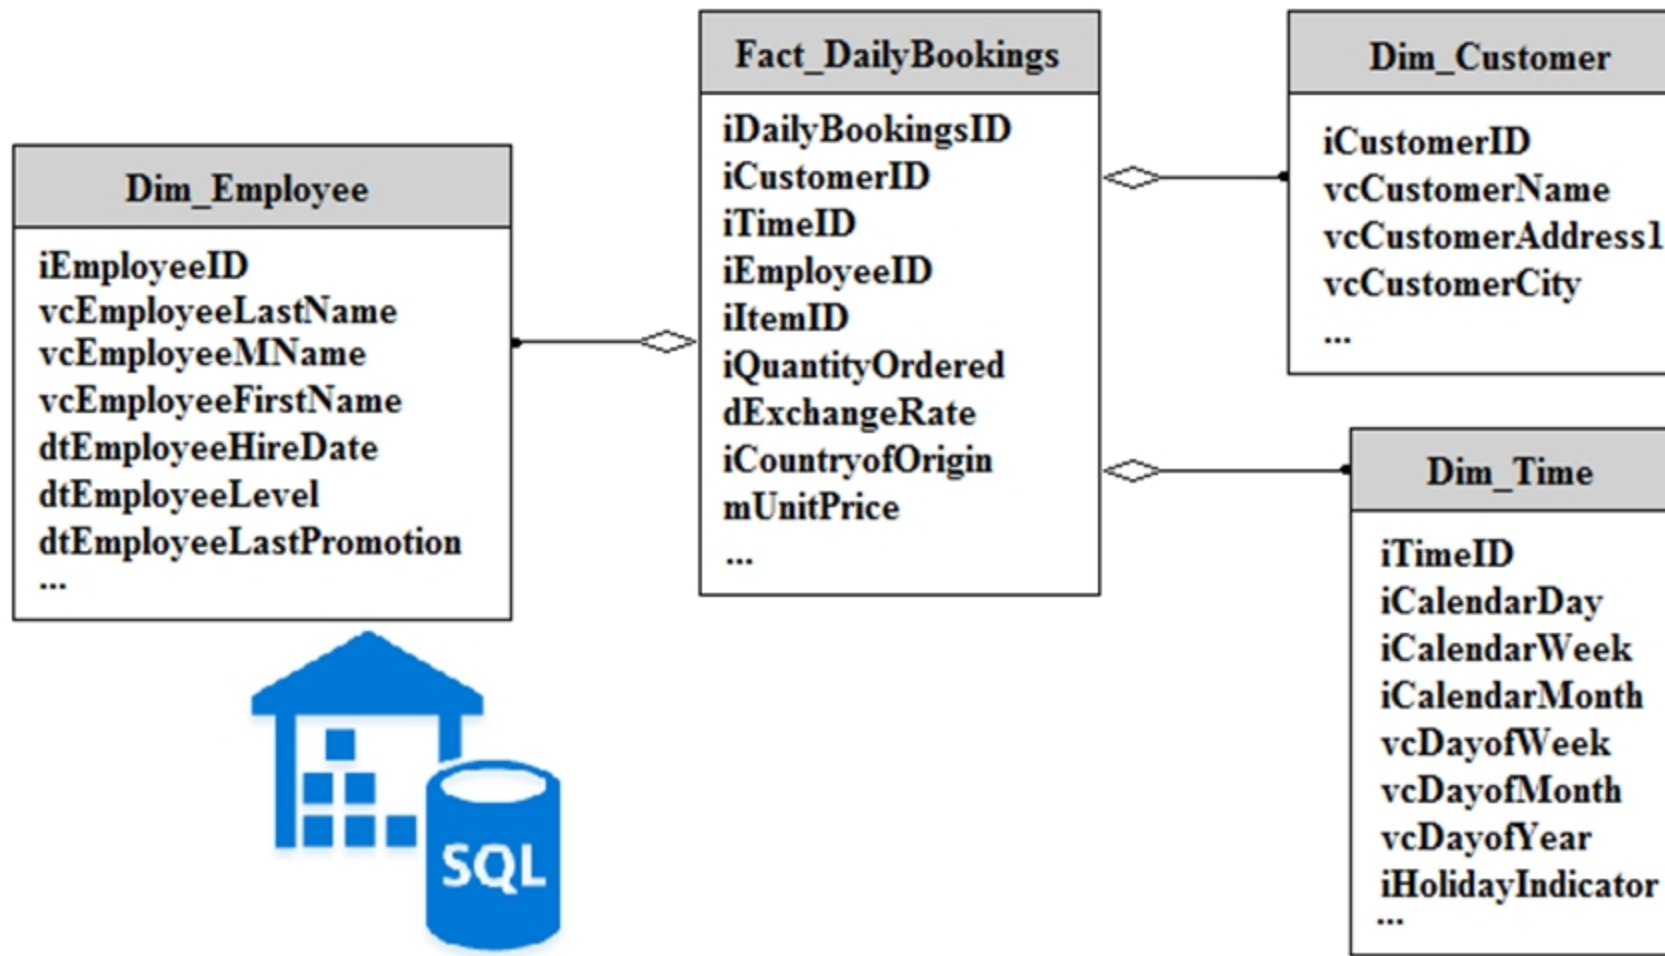
Question #: 31

Topic #: 1

[\[All DP-201 Questions\]](#)

HOTSPOT -

You have a data model that you plan to implement in a data warehouse in Azure Synapse Analytics as shown in the following exhibit.



Azure Synapse Analytics

All the dimension tables will be less than 2 GB after compression, and the fact table will be approximately 6 TB.

Which type of table should you use for each table? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Dim_Customer:

	▼
Hash distributed	
Round-robin	
Replicated	

Dim_Employee:

	▼
Hash distributed	
Round-robin	
Replicated	

Dim_Time:

	▼
Hash distributed	
Round-robin	
Replicated	

Fact_DailyBookings:

	▼
Hash distributed	
Round-robin	
Replicated	

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 32

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a data storage solution for a database that is expected to grow to 50 TB. The usage pattern is singleton inserts, singleton updates, and reporting. Which storage solution should you use?

- A. Azure SQL Database elastic pools
- B. Azure Synapse Analytics
- C. Azure Cosmos DB that uses the Gremlin API
- D. Azure SQL Database Hyperscale

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 33

Topic #: 1

[\[All DP-201 Questions\]](#)

HOTSPOT -

You are designing a solution that will use Azure Table storage. The solution will log records in the following entity.

```
DepartmentName+EmployeeID (string)
Year+Month+Day+Hour+EventID (string)
FirstName (string)
LastName (string)
EventType (string)
EventTimestamp (datetime)
EvenText (string)
```

You are evaluating which partition key to use based on the following two scenarios:

- ⇒ Scenario1: Minimize hotspots under heavy write workloads.
- ⇒ Scenario2: Ensure that date lookups are as efficient as possible for read workloads.

Which partition key should you use for each scenario? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Scenario1:

DepartmentName+EmployeeID
EvenText (string)
EventTimestamp (datetime)
EventType (string)
Year+Month+Day+Hour+EventID (string)

Scenario2:

DepartmentName+EmployeeID
EvenText (string)
EventTimestamp (datetime)
EventType (string)
Year+Month+Day+Hour+EventID (string)

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 34

Topic #: 1

[\[All DP-201 Questions\]](#)

DRAG DROP -

You have data on the 75,000 employees of your company. The data contains the properties shown in the following table.

Name	Data populated	Unique values	Distinct values
Employee ID	100%	Each value is unique	75,000
Employee Surname	99%	50% of values are unique	40,000
Employee Given Name	98%	40% of values are unique	20,000
Employee Birth Date	99%	20% of values are unique	200
Employment Start Date	100%	10% of values are unique	100
Current Department	100%	0% of values are unique	25

You need to store the employee data in an Azure Cosmos DB container. Most queries on the data will filter by the Current Department and the Employee Surname properties.

Which partition key and item ID should you use for the container? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Select and Place:

Employee ID

Employee Surname

Current Department

Answer Area

Partition key:

Item ID:

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 35

Topic #: 1

[\[All DP-201 Questions\]](#)

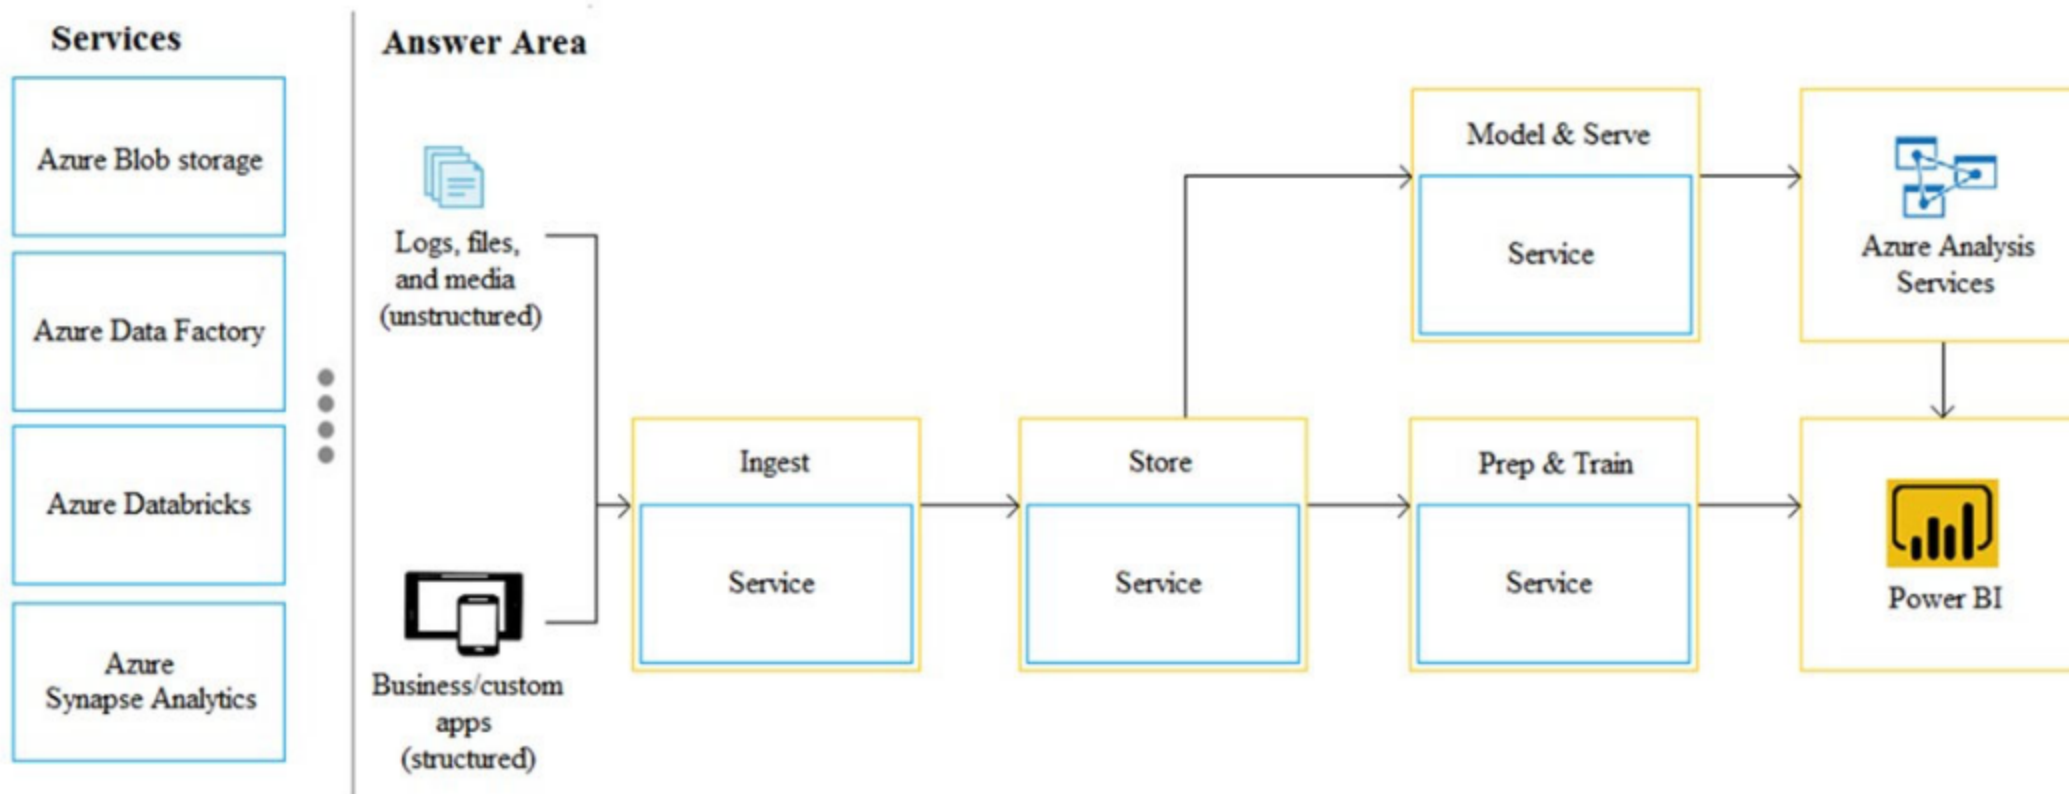
DRAG DROP -

You need to design a data architecture to bring together all your data at any scale and provide insights into all your users through the use of analytical dashboards, operational reports, and advanced analytics.

How should you complete the architecture? To answer, drag the appropriate Azure services to the correct locations in the architecture. 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.

Select and Place:



Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 36

Topic #: 1

[\[All DP-201 Questions\]](#)

HOTSPOT -

You are designing an enterprise data warehouse in Azure Synapse Analytics that will store website traffic analytic in a star schema.

You plan to have a fact table for website visits. The table will be approximately 5 GB.

You need to recommend which distribution type and index type to use for the table. The solution must provide the fastest query performance.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Distribution:

	▼
Hash	
Round robin	
Replicated	

Index:

	▼
Clustered columnstore	
Clustered	
Nonclustered	

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 37

Topic #: 1

[\[All DP-201 Questions\]](#)

You plan to deploy a reporting database to Azure. The database will contain 30 GB of data. The amount of data will increase by 300 MB each year. Rarely will the database be accessed during the second and third weeks of each month. During the first and fourth week of each month, new data will be loaded each night.

You need to recommend a solution for the planned database. The solution must meet the following requirements:

- ⇒ Minimize costs.
- ⇒ Minimize administrative effort.

What should you recommend?

- A. an Azure HDInsight cluster
- B. Azure SQL Database Hyperscale
- C. Azure SQL Database Business Critical
- D. Azure SQL Database serverless

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 38

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a solution for the ad hoc analysis of data in Azure Databricks notebooks. The data will be stored in Azure Blob storage.

You need to ensure that Blob storage will support the recovery of the data if the data is overwritten accidentally.

What should you recommend?

- A. Enable soft delete.
- B. Add a resource lock.
- C. Enable diagnostics logging.
- D. Use read-access geo-redundant storage (RA-GRS).

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 39

Topic #: 1

[\[All DP-201 Questions\]](#)

You are planning a solution that combines log data from multiple systems. The log data will be downloaded from an API and stored in a data store.

You plan to keep a copy of the raw data as well as some transformed versions of the data. You expect that there will be at least 2 TB of log files. The data will be used by data scientists and applications.

You need to recommend a solution to store the data in Azure. The solution must minimize costs.

What storage solution should you recommend?

- A. Azure Data Lake Storage Gen2
- B. Azure Synapse Analytics
- C. Azure SQL Database
- D. Azure Cosmos DB

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 40

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a serving layer for data. The design must meet the following requirements:

- ⇒ Authenticate users by using Azure Active Directory (Azure AD).
- ⇒ Serve as a hot path for data.
- ⇒ Support query scale out.
- ⇒ Support SQL queries.

What should you include in the design?

- A. Azure Data Lake Storage
- B. Azure Cosmos DB
- C. Azure Blob storage
- D. Azure Synapse Analytics

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 41

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a storage solution for streaming data that is processed by Azure Databricks. The solution must meet the following requirements:

- ⇒ The data schema must be fluid.
- ⇒ The source data must have a high throughput.
- ⇒ The data must be available in multiple Azure regions as quickly as possible.

What should you include in the solution to meet the requirements?

- A. Azure Cosmos DB
- B. Azure Synapse Analytics
- C. Azure SQL Database
- D. Azure Data Lake Storage

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 42

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a log storage solution that will use Azure Blob storage containers.

CSV log files will be generated by a multi-tenant application. The log files will be generated for each customer at five-minute intervals. There will be more than 5,000 customers. Typically, the customers will query data generated on the day the data was created.

You need to recommend a naming convention for the virtual directories and files. The solution must minimize the time it takes for the customers to query the log files.

What naming convention should you recommend?

- A. {year}/{month}/{day}/{hour}/{minute}/{CustomerID}.csv
- B. {year}/{month}/{day}/{CustomerID}/{hour}/{minute}.csv
- C. {minute}/{hour}/{day}/{month}/{year}/{CustomerID}.csv
- D. {CustomerID}/{year}/{month}/{day}/{hour}/{minute}.csv

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 43

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing an Azure Cosmos DB database that will contain news articles.

The articles will have the following properties: Category, Created Datetime, Publish Datetime, Author, Headline, Body Text, and Publish Status. Multiple articles will be published in each category daily, but no two stories in a category will be published simultaneously.

Headlines may be updated over time. Publish Status will have the following values: draft, published, updated, and removed. Most articles will remain in the published or updated status. Publish Datetime will be populated only when Publish Status is set to published.

You will serve the latest articles to websites for users to consume.

You need to recommend a partition key for the database container. The solution must ensure that the articles are served to the websites as quickly as possible.

Which partition key should you recommend?

- A. Publish Status
- B. Category + Created Datetime
- C. Headline
- D. Publish Date + random suffix

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 44

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a product catalog for a customer. The product data will be stored in Azure Cosmos DB. The product properties will be different for each product and additional properties will be added to products as needed.

Which Cosmos DB API should you use to provision the database?

- A. Cassandra API
- B. Core (SQL) API
- C. Gremlin API

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 45

Topic #: 1

[\[All DP-201 Questions\]](#)

You work for a finance company.

You need to design a business network analysis solution that meets the following requirements:

- ⇒ Analyzes the flow of transactions between the Azure environments of the company's various partner organizations
- ⇒ Supports Gremlin (graph) queries

What should you include in the solution?

- A. Azure Cosmos DB
- B. Azure Synapse
- C. Azure Analysis Services
- D. Azure Data Lake Storage Gen2

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 46

Topic #: 1

[\[All DP-201 Questions\]](#)


HOTSPOT -

You are evaluating the use of an Azure Cosmos DB account for a new database.

The proposed account will be configured as shown in the following exhibit.

[Home](#) > [New](#) > [Azure Cosmos DB](#) > Create Azure Cosmos DB Account

Create Azure Cosmos DB Account

 Create a new Azure Cosmos DB account with multi-region writes in any region by February 29, 2020 and receive up to 33% off for the life

PROJECT DETAILS

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

* Subscription

* Resource Group [Create new](#)

INSTANCE DETAILS

* Account Name

* API

Apache Spark [Sign up for Apache Spark Preview](#)

* Location

Geo-Redundancy

Multi-region Writes

'Up to 33% off multi-region writes is available to qualifying new accounts only Accounts must be created between December 1, 2019 and February 29, 2020. Offer limited to accounts with both account locations and geo-redundancy, and applies only to multi-region writes in those same regions. Both Geo-Redundancy and Multi-region Writes must be enabled on account settings. Actual discount will vary based on number of qualifying regions selected.

[Review + create](#) [Previous](#) [Next: Networking](#)

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

The data in the account will be organized and queried as:

- documents
- rows and columns
- vertices and edges

If no additional Azure regions are added by default, read access to the data in the account will be available in:

- US East
- US West 2
- US East 2

[Show Suggested Answer](#)

Actual exam question from Microsoft's DP-201

Question #: 47

Topic #: 1

[\[All DP-201 Questions\]](#)

You are designing a streaming solution that must meet the following requirements:

- ⇒ Accept input data from an Azure IoT hub.
- ⇒ Write aggregated data to Azure Cosmos DB.
- ⇒ Calculate minimum, maximum, and average sensor readings every five minutes.
- ⇒ Define calculations by using a SQL query.
- ⇒ Deploy to multiple environments by using Azure Resource Manager templates.

What should you include in the solution?

- A. Azure Functions
- B. Azure HDInsight with Spark Streaming
- C. Azure Databricks
- D. Azure Stream Analytics

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 48

Topic #: 1

[\[All DP-201 Questions\]](#)

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 enable a hierarchical namespace, and you use access control lists (ACLs).

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 49

Topic #: 1

[\[All DP-201 Questions\]](#)

You need to design a solution to support the storage of datasets. The solution must meet the following requirements:

- ⇒ Send email alerts when new datasets are added.
 - ⇒ Control access to collections of datasets by using Azure Active Directory groups.
- Support the storage of Microsoft Excel, Comma Separated Values (CSV), and zip files.

▪

What should you include in the solution?

- A. Azure SQL Database
- B. Azure Storage
- C. Azure Cosmos DB
- D. Azure HDInsight

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 1

Topic #: 2

[\[All DP-201 Questions\]](#)

You are designing an Azure Databricks interactive cluster. The cluster will be used infrequently and will be configured for auto-termination. You need to ensure that the cluster configuration is retained indefinitely after the cluster is terminated. The solution must minimize costs. What should you do?

- A. Clone the cluster after it is terminated.
- B. Terminate the cluster manually when processing completes.
- C. Create an Azure runbook that starts the cluster every 90 days.
- D. Pin the cluster.

Show Suggested Answer



Actual exam question from Microsoft's DP-201

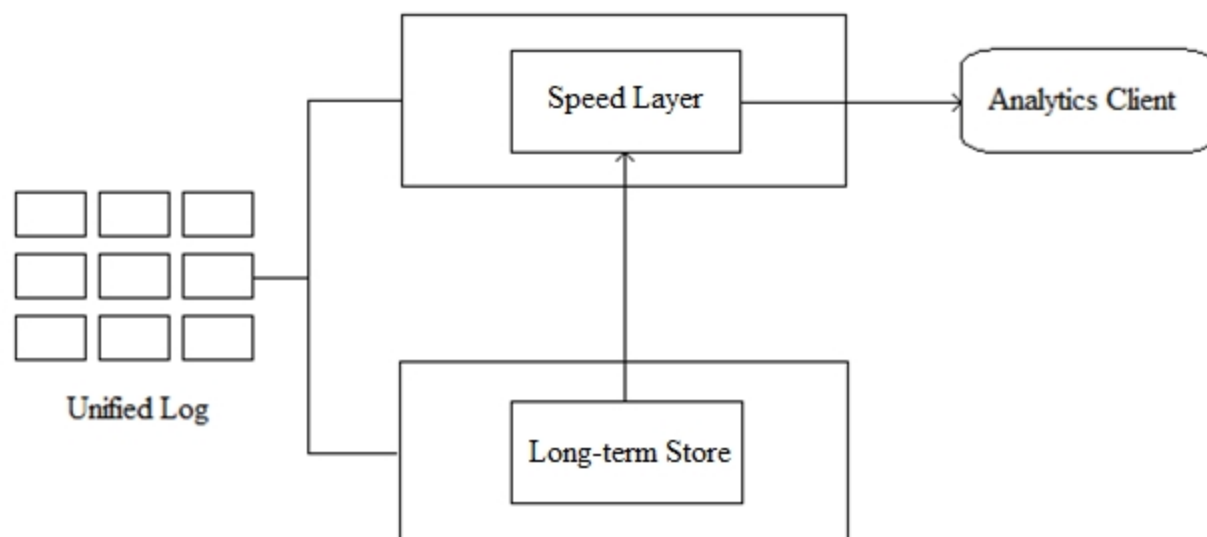
Question #: 2

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

You are planning a design pattern based on the Kappa architecture as shown in the exhibit.



Which Azure service should you use for each layer? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Speed Layer:

	▼
Azure Cosmos DB	
Azure Data Catalog	
Azure Data Factory	
Azure Synapse Analytics	

Long-term Store:

	▼
Azure Data Factory	
Azure Databricks	
Azure Synapse Analytics	
Azure Stream Analytics	

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 3

Topic #: 2

[\[All DP-201 Questions\]](#)

You need to design a telemetry data solution that supports the analysis of log files in real time.

Which two Azure services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Azure Databricks
- B. Azure Data Factory
- C. Azure Event Hubs
- D. Azure Data Lake Storage Gen2

Show Suggested Answer




Actual exam question from Microsoft's DP-201

Question #: 4

Topic #: 2

[\[All DP-201 Questions\]](#)

You are planning a design pattern based on the Lambda architecture as shown in the exhibit.



Which Azure service should you use for the hot path?

- A. Azure Databricks
- B. Azure Data Lake Storage Gen2
- C. Azure Data Factory
- D. Azure Synapse Analytics

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 5

Topic #: 2

[\[All DP-201 Questions\]](#)

You are designing an audit strategy for an Azure SQL Database environment.

You need to recommend a solution to provide real-time notifications for potential security breaches. The solution must minimize development effort.

Which destination should you include in the recommendation?

- A. Azure Blob storage
- B. Azure Synapse Analytics
- C. Azure Event Hubs
- D. Azure Log Analytics

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 6

Topic #: 2

[\[All DP-201 Questions\]](#)

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:

Support up to 1 million blobs.

-
- ⇒ Scaling must occur automatically.
- ⇒ 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.

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 7

Topic #: 2

[\[All DP-201 Questions\]](#)

A company purchases IoT devices to monitor manufacturing machinery. The company uses an IoT appliance to communicate with the IoT devices.

The company must be able to monitor the devices in real-time.

You need to design the solution.

What should you recommend?

- A. Azure Data Factory instance using Azure PowerShell
- B. Azure Analysis Services using Microsoft Visual Studio
- C. Azure Stream Analytics cloud job using Azure PowerShell
- D. Azure Data Factory instance using Microsoft Visual Studio

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 8

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

You plan to create a real-time monitoring app that alerts users when a device travels more than 200 meters away from a designated location.

You need to design an Azure Stream Analytics job to process the data for the planned app. The solution must minimize the amount of code developed and the number of technologies used.

What should you include in the Stream Analytics job? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Input type:
Stream
Reference

Input source:
Azure IoT Hub
Azure Event Hubs
Azure Blob storage

Function:
Aggregate
Geospatial
Windowing

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 9

Topic #: 2

[\[All DP-201 Questions\]](#)

A company purchases IoT devices to monitor manufacturing machinery. The company uses an IoT appliance to communicate with the IoT devices.

The company must be able to monitor the devices in real-time.

You need to design the solution.

What should you recommend?

- A. Azure Data Factory instance using the Azure portal
- B. Azure Analysis Services using Microsoft Visual Studio
- C. Azure Stream Analytics Edge application using Microsoft Visual Studio
- D. Azure Analysis Services using the Azure portal

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 10

Topic #: 2

[\[All DP-201 Questions\]](#)

You plan to ingest streaming social media data by using Azure Stream Analytics. The data will be stored in files in Azure Data Lake Storage, and then consumed by using Azure Databricks and PolyBase in Azure Synapse Analytics.

You need to recommend a Stream Analytics data output format to ensure that the queries from Databricks and PolyBase against the files encounter the fewest possible errors. The solution must ensure that the files can be queried quickly and that the data type information is retained.

What should you recommend?

- A. Avro
- B. CSV
- C. Parquet
- D. JSON

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 11

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

The following code segment is used to create an Azure Databricks cluster.

```
{
  "num_workers": null,
  "autoscale": {
    "min_workers": 2,
    "max_workers": 8
  },
  "cluster_name": "MyCluster",
  "spark_version": "latest-stable-scala2.11",
  "spark_conf": {
    "spark.databricks.cluster.profile": "serverless",
    "spark.databricks.repl.allowedLanguages": "sql,python,r"
  },
  "node_type_id": "Standard_DS13_v2",
  "ssh_public_keys": [],
  "custom_tags": {
    "ResourceClass": "Serverless"
  },
  "spark_env_vars": {
    "PYSPARK_PYTHON": "/databricks/python3/bin/python3"
  },
  "autotermination_minutes": 90,
  "enable_elastic_disk": true,
  "init_scripts": []
}
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
The Databricks cluster supports multiple concurrent users.	<input type="radio"/>	<input type="radio"/>
The Databricks cluster minimizes costs when running scheduled jobs that execute notebooks.	<input type="radio"/>	<input type="radio"/>
The Databricks cluster supports the creation of a Delta Lake table.	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 12

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

A company stores large datasets in Azure, including sales transactions and customer account information.

You must design a solution to analyze the data. You plan to create the following HDInsight clusters:

Cluster	Requirement
Sales	This cluster must be optimized for ad hoc HIVE queries.
Accounts	This cluster must be optimized for HIVE queries that are used in batch processes.

You need to ensure that the clusters support the query requirements.

Which cluster types 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

Cluster	Cluster type
Sales	<input type="text"/> ▼ Storm Hadoop Interactive Query Kafka
Accounts	<input type="text"/> ▼ Spark Hadoop Interactive Query Kafka

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 13

Topic #: 2

[\[All DP-201 Questions\]](#)

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 have streaming data that is received by Azure Event Hubs and stored in Azure Blob storage. The data contains social media posts that relate to a keyword of Contoso.

You need to count how many times the Contoso keyword and a keyword of Litware appear in the same post every 30 seconds. The data must be available to Microsoft Power BI in near real-time.

Solution: You use Azure Data Factory and an event trigger to detect when new blobs are created. You use mapping data flows in Azure Data Factory to aggregate and filter the data, and then send the data to an Azure SQL database. You consume the data in Power BI by using DirectQuery mode.

Does the solution meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 14

Topic #: 2

[\[All DP-201 Questions\]](#)

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 have streaming data that is received by Azure Event Hubs and stored in Azure Blob storage. The data contains social media posts that relate to a keyword of Contoso.

You need to count how many times the Contoso keyword and a keyword of Litware appear in the same post every 30 seconds. The data must be available to Microsoft Power BI in near real-time.

Solution: You create an Azure Stream Analytics job that uses an input from Event Hubs to count the posts that have the specified keywords, and then send the data to an Azure SQL database. You consume the data in Power BI by using DirectQuery mode.

Does the solution meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 15

Topic #: 2

[\[All DP-201 Questions\]](#)

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 have streaming data that is received by Azure Event Hubs and stored in Azure Blob storage. The data contains social media posts that relate to a keyword of Contoso.

You need to count how many times the Contoso keyword and a keyword of Litware appear in the same post every 30 seconds. The data must be available to Microsoft Power BI in near real-time.

Solution: You use Azure Databricks to create a Scala notebook. You use a Structured Streaming job to connect to the event hub that counts the posts that have the specified keywords, and then writes the data to a Delta table. You consume the data in Power BI by using DirectQuery mode.

Does the solution meet the goal?

A. Yes

B. No

Show Suggested Answer




Actual exam question from Microsoft's DP-201

Question #: 16

Topic #: 2

[\[All DP-201 Questions\]](#)

You are planning a design pattern based on the Lambda architecture as shown in the exhibit.



Which Azure service should you use for the hot path?

- A. Azure Synapse Analytics
- B. Azure SQL Database
- C. Azure Cosmos DB
- D. Azure Data Catalog

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 17

Topic #: 2

[\[All DP-201 Questions\]](#)

You are designing an enterprise data warehouse in Azure Synapse Analytics. You plan to load millions of rows of data into the data warehouse each day.

You must ensure that staging tables are optimized for data loading.

You need to design the staging tables.

What type of tables should you recommend?

- A. Round-robin distributed table
- B. Hash-distributed table
- C. Replicated table
- D. External table

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 18

Topic #: 2

[\[All DP-201 Questions\]](#)

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 have streaming data that is received by Azure Event Hubs and stored in Azure Blob storage. The data contains social media posts that relate to a keyword of Contoso.

You need to count how many times the Contoso keyword and a keyword of Litware appear in the same post every 30 seconds. The data must be available to Microsoft Power BI in near real-time.

Solution: You create an Azure Stream Analytics job that uses an input from Event Hubs to count the posts that have the specified keywords, and then send the data directly to Power BI.

Does the solution meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 19

Topic #: 2

[\[All DP-201 Questions\]](#)

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.

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 20

Topic #: 2

[\[All DP-201 Questions\]](#)

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 have an Azure Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that copies the data to a staging table in the data warehouse, and then uses a stored procedure to execute the R script.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 21

Topic #: 2

[\[All DP-201 Questions\]](#)

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 have an Azure Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 22

Topic #: 2

[\[All DP-201 Questions\]](#)

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 have an Azure Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes an Azure Databricks notebook, and then inserts the data into the data warehouse.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 23

Topic #: 2

[\[All DP-201 Questions\]](#)

A company purchases IoT devices to monitor manufacturing machinery. The company uses an Azure IoT Hub to communicate with the IoT devices.

The company must be able to monitor the devices in real-time.

You need to design the solution.

What should you recommend?

- A. Azure Data Factory instance using Azure Portal
- B. Azure Analysis Services using Microsoft Visual Studio
- C. Azure Stream Analytics Edge application using Microsoft Visual Studio
- D. Azure Data Factory instance using Microsoft Visual Studio

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 24

Topic #: 2

[\[All DP-201 Questions\]](#)

A company has a real-time data analysis solution that is hosted on Microsoft Azure. The solution uses Azure Event Hub to ingest data and an Azure Stream Analytics cloud job to analyze the data. The cloud job is configured to use 120 Streaming Units (SU).

You need to optimize performance for the Azure Stream Analytics job.

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

NOTE: Each correct selection is worth one point.

- A. Implement event ordering
- B. Scale the SU count for the job up
- C. Implement Azure Stream Analytics user-defined functions (UDF)
- D. Scale the SU count for the job down
- E. Implement query parallelization by partitioning the data output
- F. Implement query parallelization by partitioning the data output

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 25

Topic #: 2

[\[All DP-201 Questions\]](#)

You manage a process that performs analysis of daily web traffic logs on an HDInsight cluster. Each of the 250 web servers generates approximately 10 megabytes (MB) of log data each day. All log data is stored in a single folder in Microsoft Azure Data Lake Storage Gen 2.

You need to improve the performance of the process.

Which two changes should you make? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Combine the daily log files for all servers into one file
- B. Increase the value of the `mapreduce.map.memory` parameter
- C. Move the log files into folders so that each day's logs are in their own folder
- D. Increase the number of worker nodes
- E. Increase the value of the `hive.tez.container.size` parameter

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 26

Topic #: 2

[\[All DP-201 Questions\]](#)

A company purchases IoT devices to monitor manufacturing machinery. The company uses an IoT appliance to communicate with the IoT devices.

The company must be able to monitor the devices in real-time.

You need to design the solution.

What should you recommend?

- A. Azure Analysis Services using Azure Portal
- B. Azure Analysis Services using Microsoft Visual Studio
- C. Azure Stream Analytics Edge application using Microsoft Visual Studio
- D. Azure Data Factory instance using Microsoft Visual Studio

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 27

Topic #: 2

[\[All DP-201 Questions\]](#)

You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:

- ⇒ Send the output to Azure Synapse.
- ⇒ Identify spikes and dips in time series data.
- ⇒ Minimize development and configuration effort

Which should you include in the solution?

- A. Azure Databricks
- B. Azure Stream Analytics
- C. Azure SQL Database

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 28

Topic #: 2

[\[All DP-201 Questions\]](#)

You are designing an Azure Data Factory pipeline for processing data. The pipeline will process data that is stored in general-purpose standard Azure storage.

You need to ensure that the compute environment is created on-demand and removed when the process is completed.

Which type of activity should you recommend?

- A. Databricks Python activity
- B. Data Lake Analytics U-SQL activity
- C. HDInsight Pig activity
- D. Databricks Jar activity

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 29

Topic #: 2

[\[All DP-201 Questions\]](#)

A company installs IoT devices to monitor its fleet of delivery vehicles. Data from devices is collected from Azure Event Hub.

The data must be transmitted to Power BI for real-time data visualizations.

You need to recommend a solution.

What should you recommend?

- A. Azure HDInsight with Spark Streaming
- B. Apache Spark in Azure Databricks
- C. Azure Stream Analytics
- D. Azure HDInsight with Storm

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 30

Topic #: 2

[\[All DP-201 Questions\]](#)

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.

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 31

Topic #: 2

[\[All DP-201 Questions\]](#)

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:

- ⇒ 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.

Select and Place:

Answer Area

Technologies

Livy

Beeline

Azure Logic App

Azure API App

Requirement

Triggering of jobs

Monitoring of jobs

Technology

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 32

Topic #: 2

[\[All DP-201 Questions\]](#)

You have an Azure Databricks workspace named workspace1 in the Standard pricing tier. Workspace1 contains an all-purpose cluster named cluster1.

You need to reduce the time it takes for cluster1 to start and scale up. The solution must minimize costs.

What should you do first?

- A. Upgrade workspace1 to the Premium pricing tier.
- B. Create a pool in workspace1.
- C. Configure a global init script for workspace1.
- D. Create a cluster policy in workspace1.

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 33

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

You are designing a solution to process data from multiple Azure event hubs in near real-time.

Once processed, the data will be written to an Azure SQL database.

The solution must meet the following requirements:

- ⇒ Support the auditing of resource and data changes.
- ⇒ Support data versioning and rollback.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Azure service to use:

Azure Databricks
Azure Stream Analytics
Azure Analysis Services

Feature to use:

Delta
Replay
Persistence point

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 34

Topic #: 2

[\[All DP-201 Questions\]](#)

DRAG DROP -

You are designing a real-time processing solution for maintenance work requests that are received via email. The solution will perform the following actions:

- ⇒ Store all email messages in an archive.
- ⇒ Access weather forecast data by using the Python SDK for Azure Open Datasets.
- ⇒ Identify high priority requests that will be affected by poor weather conditions and store the requests in an Azure SQL database.

The solution must minimize costs.

How should you complete the solution? 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.

Select and Place:

Services

Azure Databricks

Azure Event Hubs

Azure IoT Hub

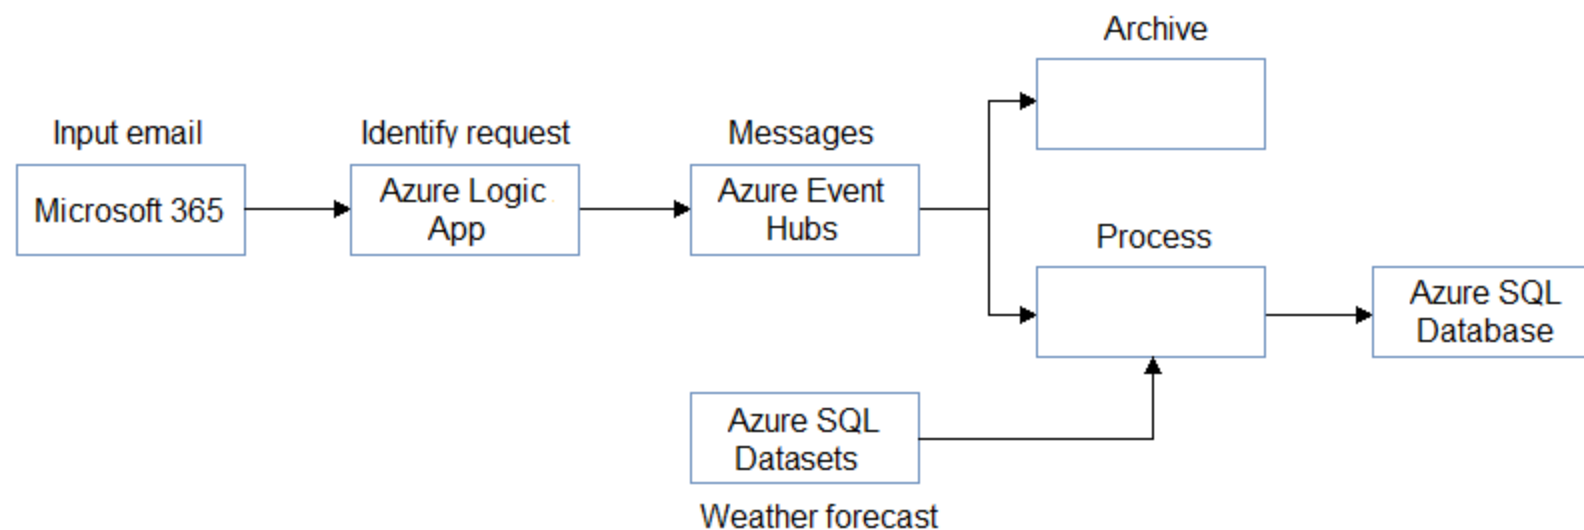
Azure Logic Apps

Azure SQL Database

Azure Storage

Azure Stream Analytics

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 35

Topic #: 2

[\[All DP-201 Questions\]](#)

You have a large amount of sensor data stored in an Azure Data Lake Storage Gen2 account. The files are in the Parquet file format.

New sensor data will be published to Azure Event Hubs.

You need to recommend a solution to add the new sensor data to the existing sensor data in real-time. The solution must support the interactive querying of the entire dataset.

Which type of server should you include in the recommendation?

- A. Azure SQL Database
- B. Azure Cosmos DB
- C. Azure Stream Analytics
- D. Azure Databricks

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 36

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

You have an Azure Storage account that generates 200,000 new files daily. The file names have a format of {YYYY}/{MM}/{DD}/{HH}/{CustomerID}.csv.

You need to design an Azure Data Factory solution that will load new data from the storage account to an Azure Data Lake once hourly. The solution must minimize load times and costs.

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

NOTE: Each correct selection is worth one point

Hot Area:

Answer Area

Load methodology:

Full Load
Incremental load
Load individual files as they arrive

Trigger:

Fixed schedule
New file
Tumbling window

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 37

Topic #: 2

[\[All DP-201 Questions\]](#)

You are designing a solution that will copy Parquet files stored in an Azure Blob storage account to an Azure Data Lake Storage Gen2 account.

The data will be loaded daily to the data lake and will use a folder structure of {Year}/{Month}/{Day}/.

You need to design a daily Azure Data Factory data load to minimize the data transfer between the two accounts.

Which two configurations should you include in the design? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Delete the files in the destination before loading new data.
- B. Filter by the last modified date of the source files.
- C. Delete the source files after they are copied.
- D. Specify a file naming pattern for the destination.

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 38

Topic #: 2

[\[All DP-201 Questions\]](#)

You have a C# application that process data from an Azure IoT hub and performs complex transformations.

You need to replace the application with a real-time solution. The solution must reuse as much code as possible from the existing application.

- A. Azure Databricks
- B. Azure Event Grid
- C. Azure Stream Analytics
- D. Azure Data Factory

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 39

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

You are designing an Azure Data Factory solution that will download up to 5 TB of data from several REST APIs.

The solution must meet the following staging requirements:

- ⇒ Ensure that the data can be landed quickly and in parallel to a staging area.
- ⇒ Minimize the need to return to the API sources to retrieve the data again should a later activity in the pipeline fail.

The solution must meet the following analysis requirements:

- ⇒ Ensure that the data can be loaded in parallel.
- ⇒ Ensure that users and applications can query the data without requiring an additional compute engine.

What should you include in the solution to meet the requirements? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Staging requirements:

Azure Blob storage

Azure SQL Database

Azure Synapse Analytics

Analysis requirements:

Azure Blob storage

Azure SQL Database

Azure Synapse Analytics

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 40

Topic #: 2

[\[All DP-201 Questions\]](#)

A company purchases IoT devices to monitor manufacturing machinery. The company uses an IoT appliance to communicate with the IoT devices.

The company must be able to monitor the devices in real-time.

You need to design the solution.

What should you recommend?

- A. Azure Data Factory instance using Azure Portal
- B. Azure Analysis Services using Microsoft Visual Studio
- C. Azure Stream Analytics cloud job using Azure Portal
- D. Azure Data Factory instance using Azure Portal

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 41

Topic #: 2

[\[All DP-201 Questions\]](#)

You are designing a real-time stream solution based on Azure Functions. The solution will process data uploaded to Azure Blob Storage.

The solution requirements are as follows:

- ⇒ Support up to 1 million blobs.
- ⇒ Scaling must occur automatically.
- ⇒ 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.

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 42

Topic #: 2

[\[All DP-201 Questions\]](#)

You plan to migrate data to Azure SQL Database.

The database must remain synchronized with updates to Microsoft Azure and SQL Server.

You need to set up the database as a subscriber.

What should you recommend?

- A. Azure Data Factory
- B. SQL Server Data Tools
- C. Data Migration Assistant
- D. SQL Server Agent for SQL Server 2017 or later
- E. SQL Server Management Studio 17.9.1 or later

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 43

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

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.

Hot Area:

Answer Area

Requirement	Workload	Tier
Provide managed clusters for running production jobs.	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Data Engineering only</p> <p>Data Analytics only</p> <p>Data Engineering and Data Analytics</p> </div> </div>	Standard
Provide persistent clusters that support auto-scaling for analytics processes.	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Data Engineering only</p> <p>Data Analytics only</p> <p>Data Engineering and Data Analytics</p> </div> </div>	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Standard</p> <p>Premium</p> </div> </div>
Provide role-based access control (RBAC) support for Notebooks.	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Data Engineering only</p> <p>Data Analytics only</p> <p>Data Engineering and Data Analytics</p> </div> </div>	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Standard</p> <p>Premium</p> </div> </div>

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 44

Topic #: 2

[\[All DP-201 Questions\]](#)

You design data engineering solutions for a company.

A project requires analytics and visualization of large set of data. The project has the following requirements:

- ⇒ Notebook scheduling
- ⇒ Cluster automation
- ⇒ Power BI Visualization

You need to recommend the appropriate Azure service. Your solution must minimize the number of services required.

Which Azure service should you recommend?

- A. Azure Batch
- B. Azure Stream Analytics
- C. Azure Databricks
- D. Azure HDInsight

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 45

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

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.

Hot Area:

Answer Area

Requirement	Service
Integrate the on-premises data into the cloud.	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Azure Databricks</p> <p>Azure Data Factory</p> <p>Azure Synapse Analytics</p> <p>Azure Batch</p> </div> </div>
Develop notebooks to transform the data.	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Azure Databricks</p> <p>Azure Data Factory</p> <p>Azure Synapse Analytics</p> <p>Azure Batch</p> </div> </div>
Run notebooks.	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Azure Databricks</p> <p>Azure Data Factory</p> <p>Azure Synapse Analytics</p> <p>Azure Batch</p> </div> </div>
Load the data.	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Azure Databricks</p> <p>Azure Data Factory</p> <p>Azure Synapse Analytics</p> <p>Azure Batch</p> </div> </div>
Store the transformed data.	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Azure Databricks</p> <p>Azure Data Factory</p> <p>Azure Synapse Analytics</p> <p>Azure Batch</p> </div> </div>

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 46

Topic #: 2

[\[All DP-201 Questions\]](#)

A company plans to use Apache Spark Analytics to analyze intrusion detection data.

You need to recommend a solution to analyze network and system activities for malicious activities and policy violations. The solution must minimize administrative efforts.

What should you recommend?

- A. Azure Data Factory
- B. Azure Data Lake Storage
- C. Azure Databricks
- D. Azure HDInsight

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 47

Topic #: 2

[\[All DP-201 Questions\]](#)

You are planning a streaming data solution that will use Azure Databricks. The solution will stream sales transaction data from an online store. The solution has the following specifications:

- ⇒ The output data will contain items purchased, quantity, line total sales amount, and line total tax amount.
- ⇒ Line total sales amount and line total tax amount will be aggregated in Databricks.
- ⇒ Sales transactions will never be updated. Instead, new rows will be added to adjust a sale.

You need to recommend an output mode for the dataset that will be processed by using Structured Streaming. The solution must minimize duplicate data.

What should you recommend?

- A. Append
- B. Complete
- C. Update

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 48

Topic #: 2

[\[All DP-201 Questions\]](#)

HOTSPOT -

You have an Azure event hub named retailhub that has 16 partitions. Transactions are posted to retailhub. Each transaction includes the transaction ID, the individual line items, and the payment details. The transaction ID is used as the partition key.

You are designing an Azure Stream Analytics job to identify potentially fraudulent transactions at a retail store. The job will use retailhub as the input. The job will output the transaction ID, the individual line items, the payment details, a fraud score, and a fraud indicator.

You plan to send the output to an Azure event hub named fraudhub.

You need to ensure that the fraud detection solution is highly scalable and processes transactions as quickly as possible.

How should you structure the output of the Stream Analytics job? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Number of partitions:

	▼
1	
8	
16	
32	

Partition key:

	▼
Fraud indicator	
Fraud score	
Individual line items	
Payment details	
Transaction ID	

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 49

Topic #: 2

[\[All DP-201 Questions\]](#)

DRAG DROP -

You have a CSV file in Azure Blob storage. The file does NOT have a header row.

You need to use Azure Data Factory to copy the file to an Azure SQL database. The solution must minimize how long it takes to copy the file.

How should you configure the copy process? To answer, drag the appropriate components to the correct locations. Each component 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:

Components

A copy activity that has an explicit schema mapping

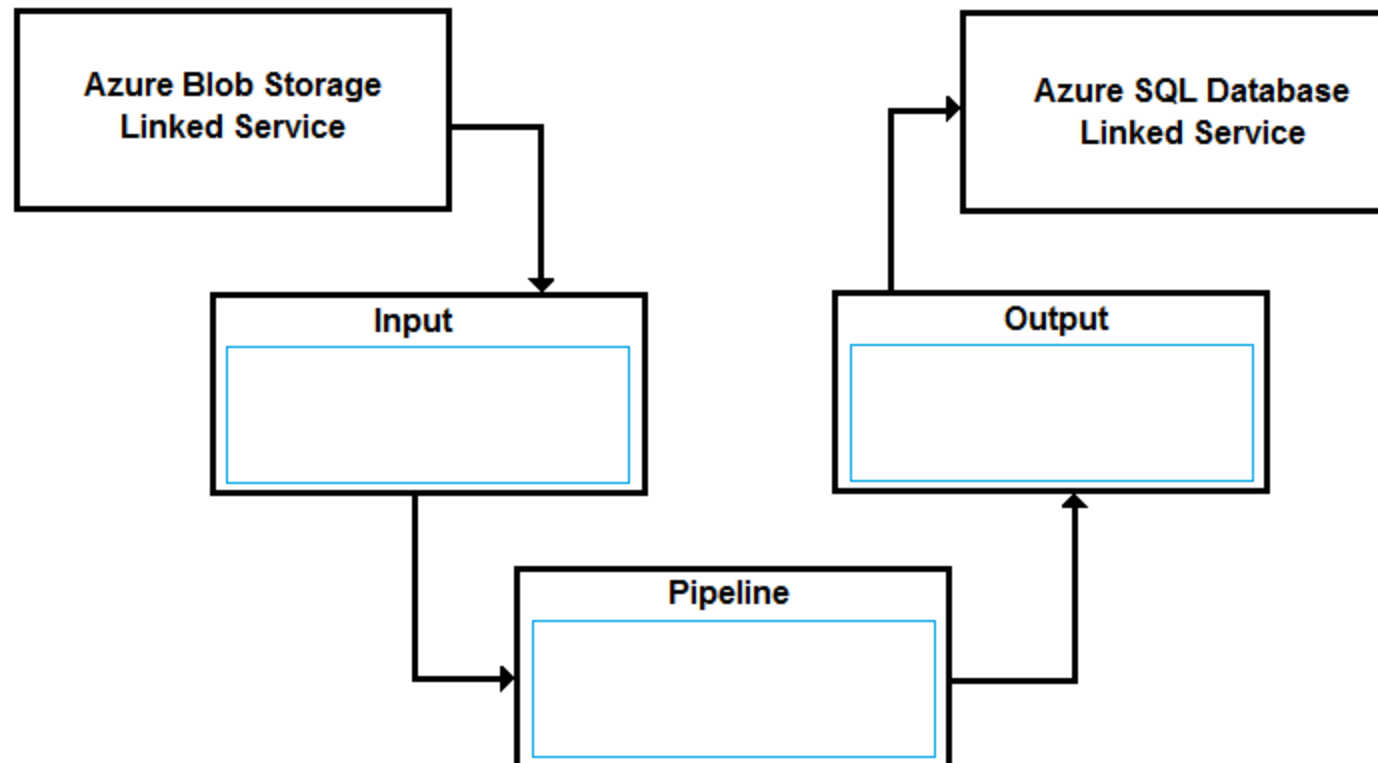
A data flow activity that has a general purpose compute type

A delimited text dataset that has a comma as a column delimiter

A Microsoft Excel dataset that has a specified sheet name

An Azure SQL Database dataset that has a specified schema and table

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's DP-201


Question #: 50

Topic #: 2

[\[All DP-201 Questions\]](#)

DRAG DROP -

You are planning a design pattern based on the Lambda architecture as shown in the exhibit.



Which Azure services should you use for the cold path? To answer, drag the appropriate services to the correct layers. 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.

Select and Place:

Services

Azure Data Lake Storage Gen2

Azure Event Hubs

Azure Log Analytics

Azure Synapse Analytics

Azure Stream Analytics

Answer Area

Layer2:

Layer3:

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 51

Topic #: 2

[\[All DP-201 Questions\]](#)

You need to recommend an Azure Cosmos DB solution that meets the following requirements:

- ⇒ All data that was NOT modified during the last 30 days must be purged automatically.
- ⇒ The solution must NOT affect ongoing user requests.

What should you recommend using to purge the data?

- A. an Azure Cosmos DB stored procedure executed by an Azure logic app
- B. an Azure Cosmos DB REST API Delete Document operation called by an Azure function
- C. Time To Live (TTL) setting in Azure Cosmos DB
- D. an Azure Cosmos DB change feed queried by an Azure function

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 52

Topic #: 2

[\[All DP-201 Questions\]](#)

You are planning a solution to aggregate streaming data that originates in Apache Kafka and is output to Azure Data Lake Storage Gen2. The developers who will implement the stream processing solution use Java.

Which service should you recommend using to process the streaming data?

- A. Azure Data Factory
- B. Azure Databricks
- C. Azure Event Hubs
- D. Azure Stream Analytics

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 53

Topic #: 2

[\[All DP-201 Questions\]](#)

DRAG DROP -

Your company has a custom human resources (HR) app named App1 that is deployed to mobile devices.

You are designing a solution that will use real-time metrics to indicate how employees use App1. The solution must meet the following requirements:

- ⇒ Use hierarchy data exported monthly from the company's HR enterprise application.
- ⇒ Process approximately 1 GB of telemetry data per day.
- ⇒ Minimize costs.

You need to recommend which Azure services to use to meet the requirements.

Which two services should you recommend? To answer, drag the appropriate services to the correct targets. 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.

Select and Place:

Services

Azure Event Hubs

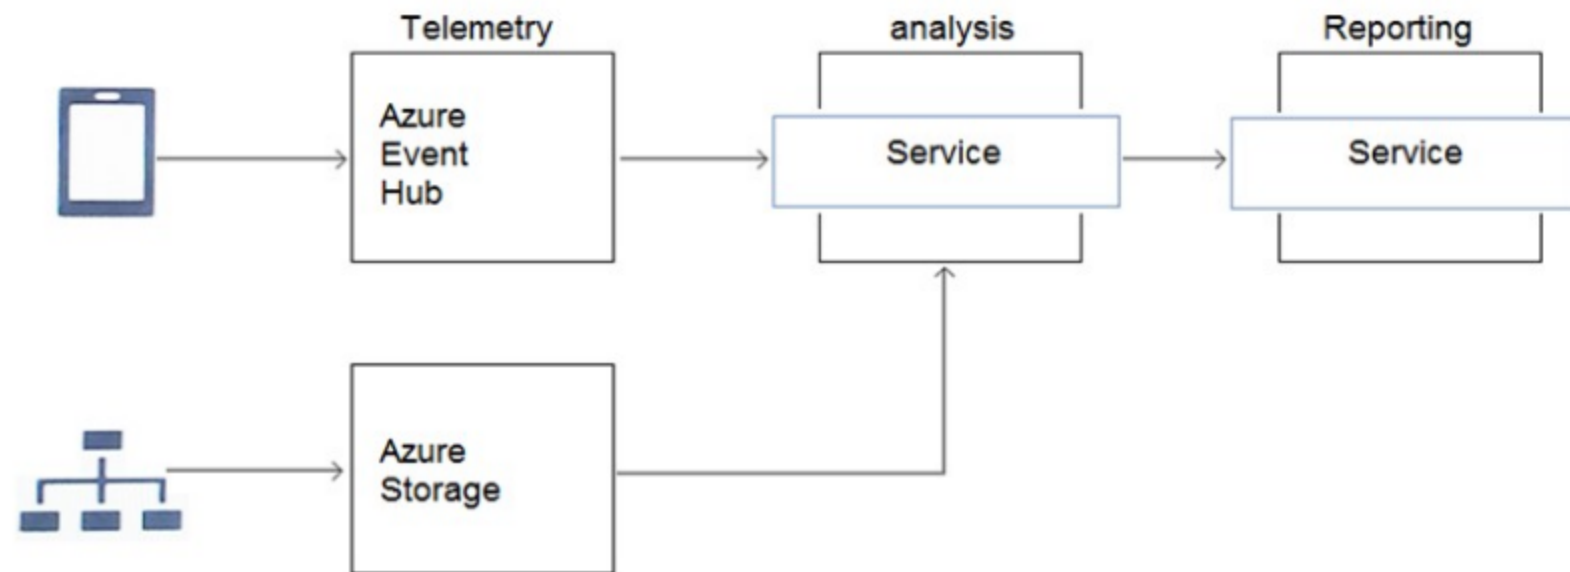
Azure Stream Analytics

Azure Synapse

Azure Data Share

Microsoft Power BI

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 54

Topic #: 2

[\[All DP-201 Questions\]](#)

You are planning an Azure solution that will aggregate streaming data.

The input data will be retrieved from tab-separated values (TSV) files in Azure Blob storage.

You need to output the maximum value from a specific column for every two-minute period in near real-time. The output must be written to Blob storage as a Parquet file.

What should you use?

- A. Azure Data Factory and mapping data flows
- B. Azure Data Factory and wrangling data flows
- C. Azure Stream Analytics window functions
- D. Azure Databricks and Apache Spark SQL window functions

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 55

Topic #: 2

[\[All DP-201 Questions\]](#)

You are designing a statistical analysis solution that will use custom proprietary Python functions on near real-time data from Azure Event Hubs.

You need to recommend which Azure service to use to perform the statistical analysis. The solution must minimize latency.

What should you recommend?

- A. Azure Synapse Analytics
- B. Azure Stream Analytics
- C. Azure Databricks
- D. Azure SQL Database

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 1

Topic #: 3

[\[All DP-201 Questions\]](#)

You are planning a big data solution in Azure.

You need to recommend a technology that meets the following requirements:

- ⇒ Be optimized for batch processing.
- ⇒ Support autoscaling.
- ⇒ Support per-cluster scaling.

Which technology should you recommend?

- A. Azure Synapse Analytics
- B. Azure HDInsight with Spark
- C. Azure Analysis Services
- D. Azure Databricks

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 2

Topic #: 3

[\[All DP-201 Questions\]](#)

You are designing an enterprise data warehouse in Azure Synapse Analytics that will contain a table named Customers. Customers will contain credit card information.

You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers. The solution must prevent all the salespeople from viewing or inferring the credit card information.

What should you include in the recommendation?

- A. row-level security
- B. data masking
- C. column-level security
- D. Always Encrypted

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 3

Topic #: 3

[\[All DP-201 Questions\]](#)

HOTSPOT -

A company plans to use Azure SQL Database to support a line of business application. The application will manage sensitive employee data.

The solution must meet the following requirements:

- ⇒ Encryption must be performed by the application.
- ⇒ Only the client application must have access keys for encrypting and decrypting data.
- ⇒ Data must never appear as plain text in the database.
- ⇒ The strongest possible encryption method must be used.
- ⇒ Grouping must be possible on selected data.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Data type	Encryption method
Grouping data	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <ul style="list-style-type: none"> Always Encrypted with randomized encryption Always Encrypted with deterministic encryption CREATE SYMMETRIC KEY statement CREATE CERTIFICATE statement </div> </div>
Non-Grouping data	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <ul style="list-style-type: none"> Always Encrypted with randomized encryption Always Encrypted with deterministic encryption CREATE SYMMETRIC KEY statement CREATE CERTIFICATE statement </div> </div>

Show Suggested Answer

Actual exam question from Microsoft's DP-201

Question #: 4

Topic #: 3

[\[All DP-201 Questions\]](#)

HOTSPOT -

You are designing the security for a mission critical Azure SQL database named DB1. DB1 contains several columns that store Personally Identifiable Information (PII) data

You need to recommend a security solution that meets the following requirements:

- ⇒ Ensures that DB1 is encrypted at rest
- ⇒ Ensures that data from the columns containing PII data is encrypted in transit

Which security solution should you recommend for DB1 and the columns? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

DB1:

	▼
Always Encrypted	
Dynamic data masking	
Row-level security	
Transparent Data Encryption (TDE)	

Columns:

	▼
Always Encrypted	
Dynamic data masking	
Row-level security	
Transparent Data Encryption (TDE)	

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 5

Topic #: 3

[\[All DP-201 Questions\]](#)

You are developing an application that uses Azure Data Lake Storage Gen 2.

You need to recommend a solution to grant permissions to a specific application for a limited time period.

What should you include in the recommendation?

- A. Azure Active Directory (Azure AD) identities
- B. shared access signatures (SAS)
- C. account keys
- D. role assignments

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 6

Topic #: 3

[\[All DP-201 Questions\]](#)

You need to recommend a security solution to grant anonymous users permission to access the blobs in a specific container only.

What should you include in the recommendation?

- A. access keys for the storage account
- B. shared access signatures (SAS)
- C. Role assignments
- D. the public access level for the container

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 7

Topic #: 3

[\[All DP-201 Questions\]](#)

You are designing a solution that will use Azure Databricks and Azure Data Lake Storage Gen2.

From Databricks, you need to access Data Lake Storage directly by using a service principal.

What should you include in the solution?

- A. shared access signatures (SAS) in Data Lake Storage
- B. access keys in Data Lake Storage
- C. an organizational relationship in Azure Active Directory (Azure AD)
- D. an application registration in Azure Active Directory (Azure AD)

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 8

Topic #: 3

[\[All DP-201 Questions\]](#)

You are designing a storage solution to store CSV files.

You need to grant a data scientist access to read all the files in a single container of an Azure Storage account. The solution must use the principle of least privilege and provide the highest level of security.

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

NOTE: Each correct selection is worth one point.

- A. Provide an access key.
- B. Assign the Storage Blob Data Reader role at the container level.
- C. Assign the Reader role to the storage account.
- D. Provide an account shared access signature (SAS).
- E. Provide a user delegation shared access signature (SAS).

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 9

Topic #: 3

[\[All DP-201 Questions\]](#)

You are designing an Azure Synapse solution that will provide a query interface for the data stored in an Azure Storage account. The storage account is only accessible from a virtual network.

You need to recommend an authentication mechanism to ensure that the solution can access the source data.

What should you recommend?

- A. a shared key
- B. an Azure Active Directory (Azure AD) service principal
- C. a shared access signature (SAS)
- D. anonymous public read access

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 10

Topic #: 3

[\[All DP-201 Questions\]](#)

HOTSPOT -

You have an Azure subscription that contains an Azure Data Lake Storage account. The storage account contains a data lake named DataLake1.

You plan to use an Azure data factory to ingest data from a folder in DataLake1, transform the data, and land the data in another folder.

You need to ensure that the data factory can read and write data from any folder in the DataLake1 file system. The solution must meet the following requirements:

- ⇒ Minimize the risk of unauthorized user access.
- ⇒ Use the principle of least privilege.
- ⇒ Minimize maintenance effort.

How should you configure access to the storage account for the data factory? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point

Hot Area:

Answer Area

Use to authenticate by using

Azure Active Directory (Azure AD)	a managed identity
a shared access signature (SAS)	a stored access policy
a shared key	an Authorization header

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 11

Topic #: 3

[\[All DP-201 Questions\]](#)

You are designing a real-time stream processing solution in Azure Stream Analytics. The solution must read data from a blob container in an Azure Storage account via a service endpoint.

You need to recommend an authentication mechanism for the solution.

What should you recommend?

- A. a managed identity
- B. a storage access signature (SAS)
- C. a user-assigned managed identity
- D. an account key

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 12

Topic #: 3

[\[All DP-201 Questions\]](#)

You are designing security for administrative access to Azure Synapse Analytics.

You need to recommend a solution to ensure that administrators use two-factor authentication when accessing the data warehouse from Microsoft SQL Server Management Studio (SSMS).

What should you include in the recommendation?

- A. Azure conditional access policies
- B. Azure Active Directory (Azure AD) Privileged Identity Management (PIM)
- C. Azure Key Vault secrets
- D. Azure Active Directory (Azure AD) Identity Protection

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 13

Topic #: 3

[\[All DP-201 Questions\]](#)

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 have an Azure SQL database that has columns. The columns contain sensitive Personally Identifiable Information (PII) data.

You need to design a solution that tracks and stores all the queries executed against the PII data. You must be able to review the data in Azure Monitor, and the data must be available for at least 45 days.

Solution: You add classifications to the columns that contain sensitive data. You turn on Auditing and set the audit log destination to use Azure Log Analytics.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 14

Topic #: 3

[\[All DP-201 Questions\]](#)

You have an Azure Storage account.

You plan to copy one million image files to the storage account.

You plan to share the files with an external partner organization. The partner organization will analyze the files during the next year.

You need to recommend an external access solution for the storage account. The solution must meet the following requirements:

⇒ Ensure that only the partner organization can access the storage account.

Ensure that access of the partner organization is removed automatically after 365 days.

▪

What should you include in the recommendation?

- A. shared keys
- B. Azure Blob storage lifecycle management policies
- C. Azure policies
- D. shared access signature (SAS)

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 15

Topic #: 3

[\[All DP-201 Questions\]](#)

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 have an Azure SQL database that has columns. The columns contain sensitive Personally Identifiable Information (PII) data.

You need to design a solution that tracks and stores all the queries executed against the PII data. You must be able to review the data in Azure Monitor, and the data must be available for at least 45 days.

Solution: You execute a daily stored procedure that retrieves queries from Query Store, looks up the column classifications, and stores the results in a new table in the database.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 16

Topic #: 3

[\[All DP-201 Questions\]](#)

A company stores sensitive information about customers and employees in Azure SQL Database.

You need to ensure that the sensitive data remains encrypted in transit and at rest.

What should you recommend?

- A. Transparent Data Encryption
- B. Always Encrypted with secure enclaves
- C. Azure Disk Encryption
- D. SQL Server AlwaysOn

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 17

Topic #: 3

[\[All DP-201 Questions\]](#)

DRAG DROP -

You are designing a data warehouse in Azure Synapse Analytics 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.

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Roles

Group

Department managers

IT

Role

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 18

Topic #: 3

[\[All DP-201 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 19

Topic #: 3

[\[All DP-201 Questions\]](#)

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 have an Azure SQL database that has columns. The columns contain sensitive Personally Identifiable Information (PII) data.

You need to design a solution that tracks and stores all the queries executed against the PII data. You must be able to review the data in Azure Monitor, and the data must be available for at least 45 days.

Solution: You create a SELECT trigger on the table in SQL Database that writes the query to a new table in the database, and then executes a stored procedure that looks up the column classifications and joins to the query text.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 20

Topic #: 3

[\[All DP-201 Questions\]](#)

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 have an Azure SQL database that has columns. The columns contain sensitive Personally Identifiable Information (PII) data.

You need to design a solution that tracks and stores all the queries executed against the PII data. You must be able to review the data in Azure Monitor, and the data must be available for at least 45 days.

Solution: You add classifications to the columns that contain sensitive data. You turn on Auditing and set the audit log destination to use Azure Blob storage.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 21

Topic #: 3

[\[All DP-201 Questions\]](#)

You have an Azure subscription that contains an Azure virtual machine and an Azure Storage account. The virtual machine will access the storage account.

You are planning the security design for the storage account.

You need to ensure that only the virtual machine can access the storage account.

Which two actions should you include in the design? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Select **Allow** trusted Microsoft services to access this storage account.
- B. Select **Allow** read access to storage logging from any network.
- C. Enable a virtual network service endpoint.
- D. Set the **Allow** access from setting to Selected networks.

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 22

Topic #: 3

[\[All DP-201 Questions\]](#)

You are designing an app that will provide a data cleaning and supplementing service for customers. The app will use Azure Data Factory to run a daily process to read and write data from Azure Storage blob containers.

You need to recommend an access mechanism for the customers to grant the app access to their data. The solution must meet the following requirements:

- ⇒ Provide access for a period of three months.
- ⇒ Restrict the app's access to specific containers.
- ⇒ Minimize administrative effort.
- ⇒ Minimize changes to the existing access controls of the customer's Azure Storage accounts.

What should you recommend?

- A. a shared key
- B. anonymous public read access
- C. a managed identity
- D. a shared access signature (SAS)

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 23

Topic #: 3

[\[All DP-201 Questions\]](#)

You need to implement an Azure Storage account that will use a Blob service endpoint that uses zone-redundant storage (ZRS).

The storage account must only accept connections from a virtual network over Azure Private Link.

What should you include in the implementation?

- A. a private endpoint for Azure Blob storage
- B. a customer-managed key
- C. a shared access signature (SAS)
- D. a firewall rule to allow traffic from the virtual network

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 24

Topic #: 3

[\[All DP-201 Questions\]](#)

You have an Azure Data Lake Storage account that has a virtual network service endpoint configured.

You plan to use Azure Data Factory to extract data from the Data Lake Storage account. The data will then be loaded to a data warehouse in Azure Synapse Analytics by using PolyBase.

Which authentication method should you use to access Data Lake Storage?

- A. shared access key authentication
- B. managed identity authentication
- C. account key authentication
- D. service principal authentication

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 25

Topic #: 3

[\[All DP-201 Questions\]](#)

HOTSPOT -

A company plans to use Azure SQL Database to support a line of business application. The application will manage sensitive employee data.

The solution must meet the following requirements:

- ⇒ Encryption must be performed by the application.
- ⇒ Only the client application must have access keys for encrypting and decrypting data.
- ⇒ Data must never appear as plain text in the database.
- ⇒ The strongest possible encryption method must be used.
- ⇒ Searching must be possible on selected data.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Data type	Encryption method
Searchable data	<div style="border: 1px solid gray; padding: 2px;"><div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;">▼</div><div style="padding: 2px;"><ul style="list-style-type: none">Always Encrypted with randomized encryptionAlways Encrypted with deterministic encryptionCREATE SYMMETRIC KEY statementCREATE CERTIFICATE statement</div></div>
Non-searchable data	<div style="border: 1px solid gray; padding: 2px;"><div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;">▼</div><div style="padding: 2px;"><ul style="list-style-type: none">Always Encrypted with randomized encryptionAlways Encrypted with deterministic encryptionCREATE SYMMETRIC KEY statementCREATE CERTIFICATE statement</div></div>

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 26

Topic #: 3

[\[All DP-201 Questions\]](#)

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

Show Suggested Answer





Actual exam question from Microsoft's DP-201

Question #: 27

Topic #: 3

[\[All DP-201 Questions\]](#)

You are designing the security for an Azure SQL database.

You have an Azure Active Directory (Azure AD) group named Group1.

You need to recommend a solution to provide Group1 with read access to the database only.

What should you include in the recommendation?

- A. a contained database user
- B. a SQL login
- C. an RBAC role
- D. a shared access signature (SAS)

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 28

Topic #: 3

[\[All DP-201 Questions\]](#)

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	

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 29

Topic #: 3

[\[All DP-201 Questions\]](#)

You store data in a data warehouse in Azure Synapse Analytics.

You need to design a solution to ensure that the data warehouse and the most current data is available within one hour of a datacenter failure.

Which three actions should you include in the design? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Each day, restore the data warehouse from a geo-redundant backup to an available Azure region.
- B. If a failure occurs, update the connection strings to point to the recovered data warehouse.
- C. If a failure occurs, modify the Azure Firewall rules of the data warehouse.
- D. Each day, create Azure Firewall rules that allow access to the restored data warehouse.
- E. Each day, restore the data warehouse from a user-defined restore point to an available Azure region.

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 30

Topic #: 3

[\[All DP-201 Questions\]](#)

HOTSPOT -

You are designing an Azure data factory that will copy data from Azure Blob storage to a data warehouse in Azure Synapse Analytics.

You need to recommend an authentication mechanism that meet the following requirements:

- ⇒ Identities must be validated by using Azure Active Directory (Azure AD).
- ⇒ Development and maintenance effort must be minimized.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Blob storage:

	▼
Account key	
Managed identity	
Service principal	
Shared access signature (SAS)	

Synapse Analytics:

	▼
Managed identity	
Service principal	
SQL authentication	

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 31

Topic #: 3

[\[All DP-201 Questions\]](#)

You have an Azure Data Lake Storage Gen2 account named adls2 that is protected by a virtual network.

You are designing a SQL pool in Azure Synapse that will use adls2 as a source.

What should you use to authenticate to adls2?

- A. a shared access signature (SAS)
- B. a managed identity
- C. a shared key
- D. an Azure Active Directory (Azure AD) user

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 32

Topic #: 3

[\[All DP-201 Questions\]](#)

You plan to create an Azure Synapse Analytics dedicated SQL pool.

You need to minimize the time it takes to identify queries that return confidential information as defined by the company's data privacy regulations and the users who executed the queries.

Which two components should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. dynamic data masking for columns that contain confidential information
- B. sensitivity-classification labels applied to columns that contain confidential information
- C. resource tags for databases that contain confidential information
- D. audit logs sent to a Log Analytics workspace

Show Suggested Answer



Actual exam question from Microsoft's DP-201

Question #: 1

Topic #: 4

[\[All DP-201 Questions\]](#)

HOTSPOT -

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	

Show Suggested Answer