

**100%** Money Back  
**Guarantee**

**Vendor:**Microsoft

**Exam Code:**DP-201

**Exam Name:**Designing an Azure Data Solution

**Version:**Demo

## QUESTION 1

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

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

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

Shops will upload data every 10 days.

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

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

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

---

## QUESTION 2

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

Correct Answer: C

Gremlin features fast queries and traversals with the most widely adopted graph query standard.

References: <https://docs.microsoft.com/th-th/azure/cosmos-db/graph-introduction?view=azurermps-5.7.0>

---

### QUESTION 3

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

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

Correct Answer: B

Azure resource costs must be minimized where possible.

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

The SLA for Planning Assistance is 70 percent, and multiday outages are permitted.

---

### QUESTION 4

#### HOTSPOT

Which Azure Data Factory components should you recommend using together to import the customer data from Salesforce to Data Lake Storage? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Integration runtime type:

	▼
Azure integration runtime	
Azure-SSIS integration runtime	
Self-hosted integration runtime	

Trigger type:

	▼
Event-based trigger	
Schedule trigger	
Tumbling window trigger	

Activity type:

	▼
Copy activity	
Lookup activity	
Stored procedure activity	

Correct Answer:

## Answer Area

Integration runtime type:

Azure integration runtime
Azure-SSIS integration runtime
Self-hosted integration runtime

Trigger type:

Event-based trigger
Schedule trigger
Tumbling window trigger

Activity type:

Copy activity
Lookup activity
Stored procedure activity

Box 1: Self-hosted integration runtime A self-hosted IR is capable of running copy activity between a cloud data store and a data store in private network.

Box 2: Schedule trigger Schedule every 8 hours Box 3: Copy activity Scenario: Customer data, including name, contact information, and loyalty number, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the

---

### QUESTION 5

#### HOTSPOT

You need to design storage for the solution.

Which storage services 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

**Solution component**

**Storage service**

**Images**

	▼
Azure Blob Storage	
Azure Data Lake Storage	
Azure SQL Database	
Azure SQL Data Warehouse	

**Customer data**

	▼
Azure Blob Storage	
Azure Cosmos DB	
Azure SQL Database	
Azure SQL Data Warehouse	

Correct Answer:

# Answer Area

**Solution component**

**Storage service**

**Images**

	▼
Azure Blob Storage	
Azure Data Lake Storage	
Azure SQL Database	
Azure SQL Data Warehouse	

**Customer data**

	▼
Azure Blob Storage	
Azure Cosmos DB	
Azure SQL Database	
Azure SQL Data Warehouse	

Images: Azure Data Lake Storage

Scenario: Image data must be stored in a single data store at minimum cost.

Customer data: Azure Blob Storage

Scenario: Customer data must be analyzed using managed Spark clusters.

Spark clusters in HDInsight are compatible with Azure Storage and Azure Data Lake Storage.

Azure Storage includes these data services: Azure Blob, Azure Files, Azure Queues, and Azure Tables.

References:

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

---

## QUESTION 6

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.

Correct Answer: B

When you create an Azure virtual machine (VM), you must create a virtual network (VNet) or use an existing VNet. You also need to decide how your VMs are intended to be accessed on the VNet. Incorrect Answers:

C: Azure Application Gateway is a web traffic load balancer that enables you to manage traffic to your web applications.

D: A VPN gateway is a specific type of virtual network gateway that is used to send encrypted traffic between an Azure virtual network and an on-premises location over the public Internet.

References: <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/network-overview>

---

## QUESTION 7

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

Correct Answer: C

Incorrect Answers:

D: Azure Database Migration Service is used to migrate on-premises SQL Server databases to the cloud.

References:

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

<https://azure.microsoft.com/en-us/blog/operationalize-azure-databricks-notebooks-using-data-factory/>

<https://azure.microsoft.com/en-us/blog/data-ingestion-into-azure-at-scale-made-easier-with-latest-enhancements-to-adf-copy-data-tool/>

---



**QUESTION 8**

You need to design the image processing and storage solutions.

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Image processing**

	▼
<b>Azure HDInsight</b>	
<b>Azure Databricks</b>	
<b>Azure Batch</b>	
<b>Azure Cognitive Services</b>	

**data storage for tagging data**

	▼
<b>Azure Blob Storage</b>	
<b>Azure Table Storage</b>	
<b>Azure Cosmos DB</b>	
<b>Azure SQL Database</b>	

Correct Answer:

**Image processing**

<b>Azure HDInsight</b>
<b>Azure Databricks</b>
<b>Azure Batch</b>
<b>Azure Cognitive Services</b>

**data storage for tagging data**

<b>Azure Blob Storage</b>
<b>Azure Table Storage</b>
<b>Azure Cosmos DB</b>
<b>Azure SQL Database</b>

From the scenario:

The company identifies the following business requirements:

1.

You must transfer all images and customer data to cloud storage and remove on-premises servers.

2.

You must develop an image object and color tagging solution.

The solution has the following technical requirements:

1.

Image data must be stored in a single data store at minimum cost.

2.

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

All cloud data must be encrypted at rest and in transit. The solution must support: hyper-scale storage of images

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

---

**QUESTION 9**

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

1.

Store standardized data that data scientists will explore in a curated folder.

2.

Ensure that applications cannot access the curated folder.

3.

Store staged data for import to applications in a raw folder.

4.

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 SQL Data Warehouse

B. Azure Blob storage

C. Azure Data Lake Storage Gen2

D. Azure SQL Database

Correct Answer: B

Azure Blob Storage containers is a general purpose object store for a wide variety of storage scenarios. Blobs are stored in containers, which are similar to folders. Incorrect Answers:

C: Azure Data Lake Storage is an optimized storage for big data analytics workloads.

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

---

## QUESTION 10

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

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

You 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

Correct Answer: A

Depending on what services and workloads are using the data, a good size to consider for files is 256 MB or greater. If the file sizes cannot be batched when landing in Data Lake Storage Gen1, you can have a separate compaction job that combines these files into larger ones.

Note: POSIX permissions and auditing in Data Lake Storage Gen1 comes with an overhead that becomes apparent when working with numerous small files. As a best practice, you must batch your data into larger files versus writing thousands or millions of small files to Data Lake Storage Gen1. Avoiding small file sizes can have multiple benefits, such as:

1.

Lowering the authentication checks across multiple files

2.

Reduced open file connections

3.

Faster copying/replication

4.

Fewer files to process when updating Data Lake Storage Gen1 POSIX permissions

References: <https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-best-practices>

---

## QUESTION 11

**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:

1. Ensures that DB1 is encrypted at rest

2. 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)	

Correct Answer:

## 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)	

## DB1: Transparent Data Encryption

Azure SQL Database currently supports encryption at rest for Microsoft-managed service side and client-side encryption scenarios.

Support for server encryption is currently provided through the SQL feature called Transparent Data Encryption.

Columns: Always encrypted

Always Encrypted is a feature designed to protect sensitive data stored in Azure SQL Database or SQL Server databases. Always Encrypted allows clients to encrypt sensitive data inside client applications and never reveal the encryption

keys to the database engine (SQL Database or SQL Server).

Note: Most data breaches involve the theft of critical data such as credit card numbers or personally identifiable information. Databases can be treasure troves of sensitive information. They can contain customers\' personal data (like national

identification numbers), confidential competitive information, and intellectual property. Lost or stolen data, especially customer data, can result in brand damage, competitive disadvantage, and serious fines--even lawsuits.

References:

<https://docs.microsoft.com/en-us/azure/security/fundamentals/encryption-atrest>

<https://docs.microsoft.com/en-us/azure/security/fundamentals/database-security-overview>

---

## QUESTION 12

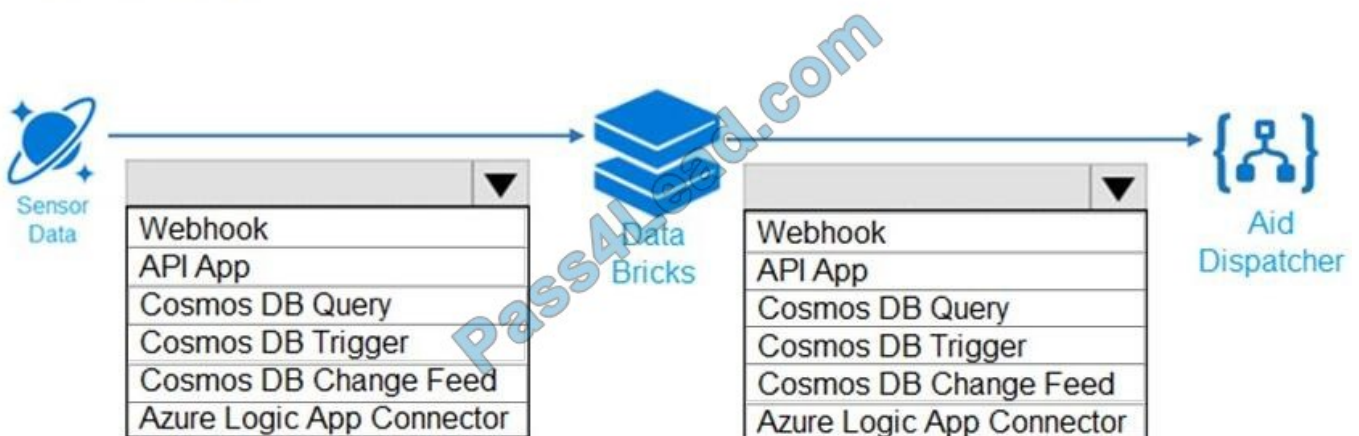
You need to ensure that emergency road response vehicles are dispatched automatically.

How should you design the processing system? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

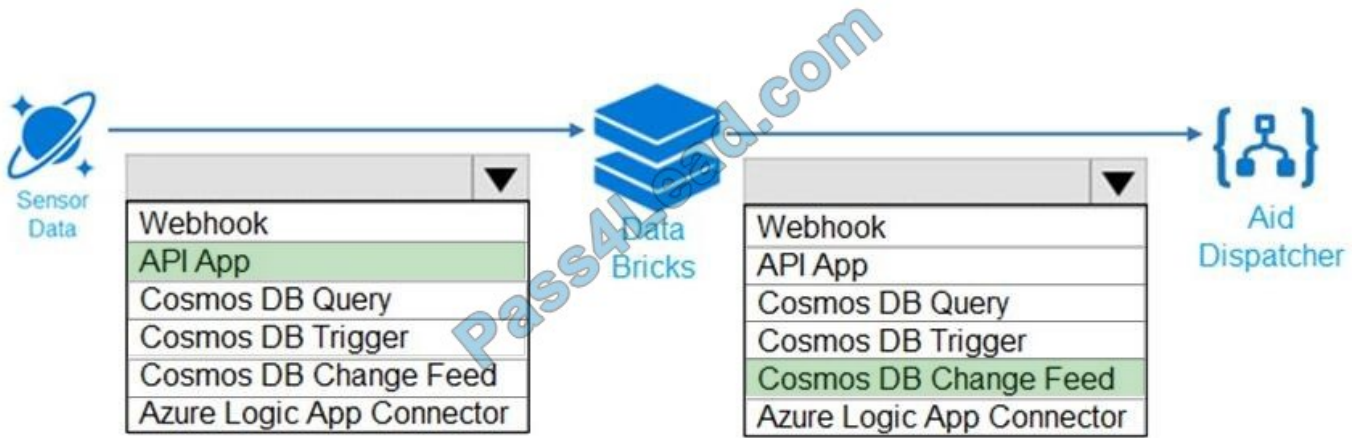
Hot Area:

### Answer Area

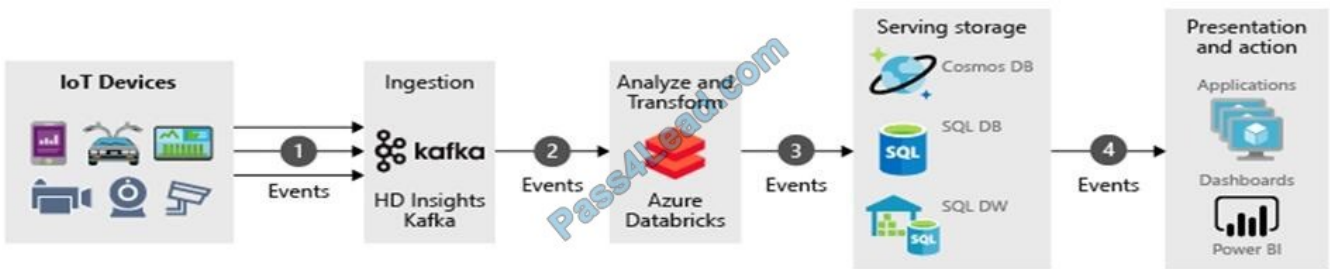


Correct Answer:

### Answer Area



Explanation: Box1: API App



Events generated from the IoT data sources are sent to the stream ingestion layer through Azure HDInsight Kafka as a stream of messages. HDInsight Kafka stores streams of data in topics for a configurable of time.

Kafka consumer, Azure Databricks, picks up the message in real time from the Kafka topic, to process the data based on the business logic and can then send to Serving layer for storage.


Downstream storage services, like Azure Cosmos DB, Azure SQL Data warehouse, or Azure SQL DB, will then be a data source for presentation and action layer.

Business analysts can use Microsoft Power BI to analyze warehoused data. Other applications can be built upon the serving layer as well. For example, we can expose APIs based on the service layer data for third party uses.

Box 2: Cosmos DB Change Feed

Change feed support in Azure Cosmos DB works by listening to an Azure Cosmos DB container for any changes. It then outputs the sorted list of documents that were changed in the order in which they were modified.

The change feed in Azure Cosmos DB enables you to build efficient and scalable solutions for each of these patterns, as shown in the following image:



Trigger call to an API when a document is inserted or modified

Real-time (stream) processing of data

Zero-downtime migrations

### Event-Computing and Notifications

Retail, Gaming, Content management

Azure Functions    Azure Notification Hubs    Azure App Service

### Stream Processing

IoT processing, Data Science & analytics

Azure Stream Analytics    Azure HDInsight    Apache Spark    Apache Storm

### Data movement

Enterprise data management

Azure Storage Blob    Azure Storage Table    Azure Data Lake    Azure Cosmos DB

Pass4Lead.com



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

## Try our product !

**100%** Guaranteed Success

**100%** Money Back Guarantee

**365** Days Free Update

**Instant Download** After Purchase

**24x7** Customer Support

Average **99.9%** Success Rate

More than **800,000** Satisfied Customers Worldwide

Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

## Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p><b>One Year Free Update</b> Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p><b>Money Back Guarantee</b> To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p><b>Security &amp; Privacy</b> We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information &amp; peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.