



SUCCEED AT THE
FIRST ATTEMPT!

<https://developer certification.com>



Questions and Answers Demo PDF

Microsoft

DP-201 Exam

Microsoft Designing an Azure Data Solution Exam

**Questions & Answers
Demo**



Version: 16.0

Question: 1

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

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

Answer: B

Explanation:

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

Scenario: Traffic sensors will occasionally capture an image of a vehicle for debugging purposes.

You must optimize performance of saving/storing vehicle images.

The impact of vehicle images on sensor data throughout must be minimized.

References:

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

Question: 2

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

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

Answer: A

Explanation:

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



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

References:

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

Question: 3

HOTSPOT

You need to design the SensorData collection.

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

NOTE: Each correct selection is worth one point.

Setting	Value
Default consistency level	<div><div></div><div><div>strong</div><div>session</div><div>eventual</div><div>consistent prefix</div><div>bounded staleness</div></div></div>
Partition key property	<div><div></div><div><div>Time</div><div>Location</div><div>Speed</div><div>License plate</div><div>Vehicle length</div></div></div>

Answer:

Setting	Value
Default consistency level	<div><div>▼</div><div><div>strong</div><div>session</div><div>eventual</div><div>consistent prefix</div><div>bounded staleness</div></div></div>
Partition key property	<div><div>▼</div><div><div>Time</div><div>Location</div><div>Speed</div><div>License plate</div><div>Vehicle length</div></div></div>

Explanation:

Box 1: Eventual

Traffic data insertion rate must be maximized.

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

With Azure Cosmos DB, developers can choose from five well-defined consistency models on the consistency spectrum. From strongest to more relaxed, the models include strong, bounded staleness, session, consistent prefix, and eventual consistency.

Box 2: License plate

This solution reports on all data related to a specific vehicle license plate. The report must use data from the SensorData collection.

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

Question: 4

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

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

Answer: B

Explanation:

Azure resource costs must be minimized where possible.



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

Question: 5

DRAG DROP

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation:



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

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

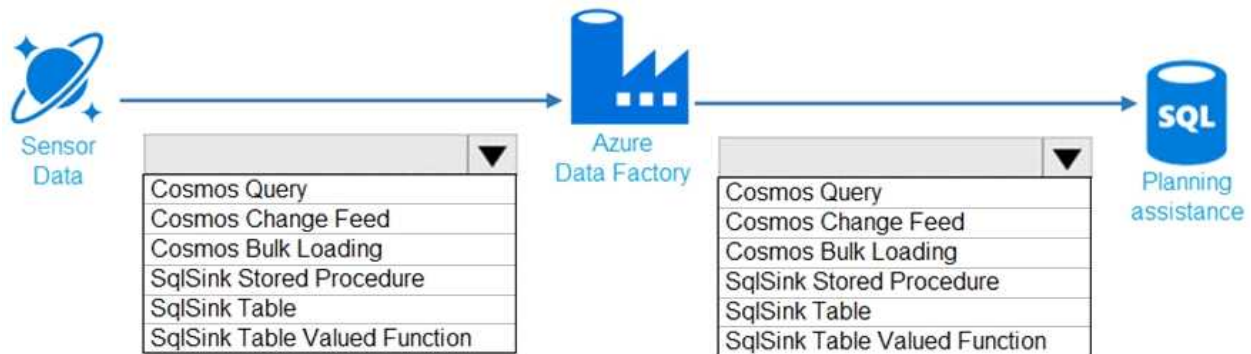
Question: 6

HOTSPOT

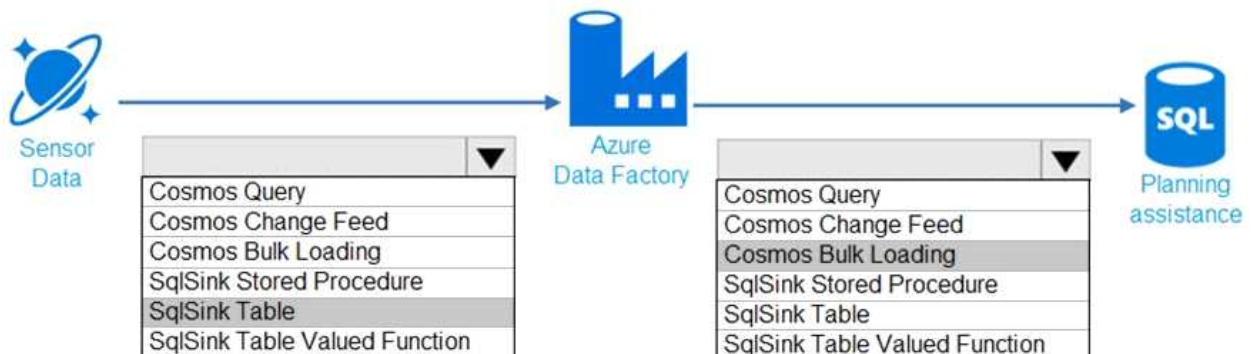
You need to design the data loading pipeline for Planning Assistance.

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

NOTE: Each correct selection is worth one point.



Answer:



Explanation:

Box 1: SqlSink Table

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

Box 2: Cosmos Bulk Loading

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

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

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

References:

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



Question: 7

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

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

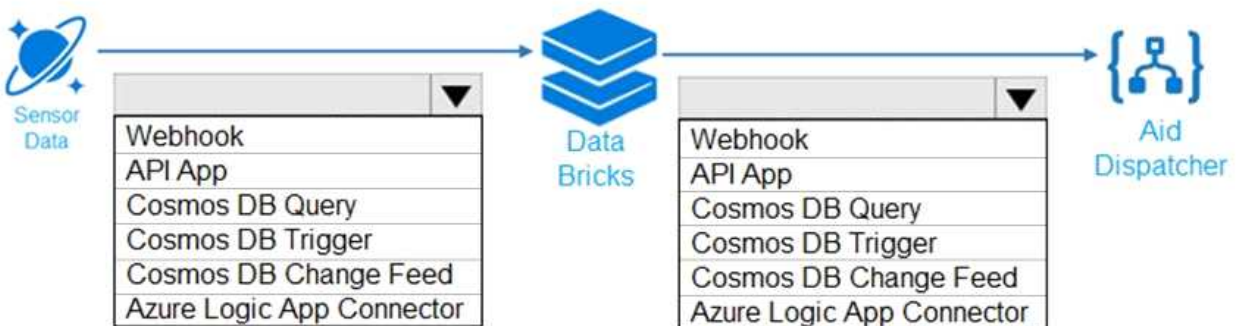
Answer: B

Question: 8

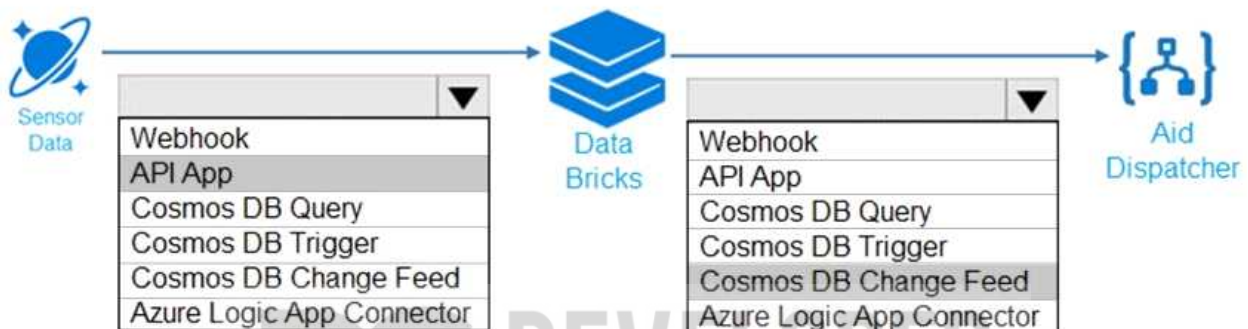
HOTSPOT

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.




Answer:



Explanation:



Box1: API App

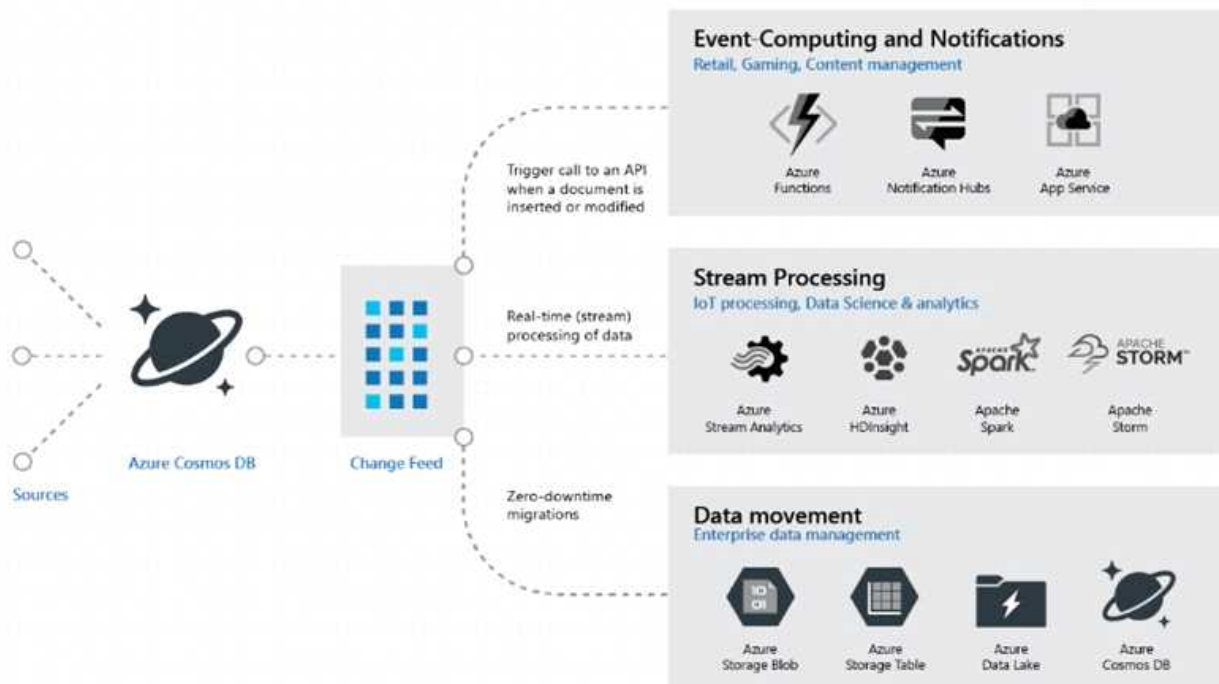


- Events generated from the IoT data sources are sent to the stream ingestion layer through Azure HDInsight Kafka as a stream of messages. HDInsight Kafka stores streams of data in topics for a configurable of time.
- Kafka consumer, Azure Databricks, picks up the message in real time from the Kafka 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:



References:

<https://docs.microsoft.com/bs-cyrl-ba/azure/architecture/example-scenario/data/realtime-analytics-vehicle-iot?view=azurermps-4.4.1>



Question: 9

DRAG DROP

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

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

NOTE: Each correct selection is worth one point.

Technologies	Answer Area	
	Requirement	Technology
Cosmos DB TTL	Backtrack reporting	
Cosmos DB indexes	Privacy and security policy	
Cosmos DB transactions		
Cosmos DB change feed		
Cosmos DB stored procedures		

Answer:

Requirement	Technology
Backtrack reporting	Cosmos DB indexes
Privacy and security policy	Cosmos DB TTL

Explanation:

Box 1: Cosmos DB indexes

The report for Backtrack must execute as quickly as possible.

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

Box 2: Cosmos DB TTL

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



vehicles on a specific road

vehicles driving above the speed limit

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

Incorrect Answers:

Cosmos DB stored procedures: Stored procedures are best suited for operations that are write heavy. When deciding where to use stored procedures, optimize around encapsulating the maximum amount of writes possible. Generally speaking, stored procedures are not the most efficient means for doing large numbers of read operations so using stored procedures to batch large numbers of reads to return to the client will not yield the desired benefit.

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy>

<https://docs.microsoft.com/en-us/azure/cosmos-db/time-to-live>

Question: 10

You need to recommend a solution for storing the image tagging data.
What should you recommend?

- A. Azure File Storage
- B. Azure Cosmos DB
- C. Azure Blob Storage
- D. Azure SQL Database
- E. Azure SQL Data Warehouse

Answer: C

Explanation:

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

Note: Azure Blob storage is Microsoft's object storage solution for the cloud. Blob storage is optimized for storing massive amounts of unstructured data. Unstructured data is data that does not adhere to a particular data model or definition, such as text or binary data.

Blob storage is designed for:

- Serving images or documents directly to a browser.
- Storing files for distributed access.
- Streaming video and audio.
- Writing to log files.
- Storing data for backup and restore, disaster recovery, and archiving.
- Storing data for analysis by an on-premises or Azure-hosted service.

References:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blobs-introduction>





<https://developer certification.com>