

IT 认证电子书



质 量 更 高 服 务 更 好

半年免费升级服务

<http://www.itrenzheng.com>

Exam : DP-201

**Title : Designing an Azure Data
Solution**

Version : DEMO

1. Topic 1, Trey Research

Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the **Question** button to return to the question.

Background

Trey Research is a technology innovator. The company partners with regional transportation department office to build solutions that improve traffic flow and safety.

The company is developing the following solutions:

| Solution | Comments |
|---------------------|--|
| Real Time Response | This solution will detect sudden changes in traffic flow including slow downs and stops that persist for more than one minute. The system will automatically dispatch emergency response vehicles to investigate issues. The solution will use a PySpark script to detect traffic flow changes. Script performance will be limited by available memory. |
| Backtrack | This solution will allow public safety officials to locate vehicles on roadways that implement traffic sensors. The solution must report changes in real time. |
| Planning Assistance | Transportation organizations will use Planning Assistance to analyze traffic data. The solution will allow users to define reports based on queries of the traffic data. The reports can be used for the following analyses: <ul style="list-style-type: none"> • current traffic load • correlation with recent local events such as sporting events • historical traffic • tracking the travel of a single vehicle |

Regional transportation departments installed traffic sensor systems on major highways across North America.

Sensors record the following information each time a vehicle passes in front of a sensor:

- Time
- Location in latitude and longitude
- Speed in kilometers per second (kmps)
- License plate number
- Length of vehicle in meters

Sensors provide data by using the following structure:

```
{
  "time" : "2014-09-15T23:14:25.72511732",
  "location" : {
    "type": "Point",
    "coordinates": [
      31.9,
      -4.8
    ]
  },
  "speed": 66.2,
  "license_plate": "WA-AJ0072W",
  "vehicle_length": 4.5
}
```

Traffic sensors will occasionally capture an image of a vehicle for debugging purposes. You must optimize performance of saving/storing vehicle images.

Traffic sensor data

- Sensors must have permission only to add items to the SensorData collection.
- Traffic data insertion rate must be maximized.
- Once every three months all traffic sensor data must be analyzed to look for data patterns that indicate sensor malfunctions.
- Sensor data must be stored in a Cosmos DB named treydata in a collection named SensorData
- The impact of vehicle images on sensor data throughout must be minimized.

Backtrack

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

Planning Assistance

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

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.

Privacy and security policy

- Azure Active Directory must be used for all services where it is available.
 - For privacy reasons, license plate number information must not be accessible in Planning Assistance.
 - Unauthorized usage of the Planning Assistance data must be detected as quickly as possible.
- Unauthorized usage is determined by looking for an unusual pattern of usage.
- Data must only be stored for seven years.

Performance and availability

- The report for Backtrack must execute as quickly as possible.
- The SLA for Planning Assistance is 70 percent, and multiday outages are permitted.
- All data must be replicated to multiple geographic regions to prevent data loss.
- You must maximize the performance of the Real Time Response system.

Financial requirements

Azure resource costs must be minimized where possible.

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>

2.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/>

3.HOTSPOT

You need to ensure that security policies for the unauthorized detection system are met.

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

| Setting | Value |
|-----------------------|--|
| Audit log destination | <div><div></div><div>Storage queue</div><div>Event Hub</div><div>Event Grid</div><div>Blob storage</div></div> |
| Detection app service | <div><div></div><div>Function App</div><div>Web App</div><div>API App</div></div> |

Answer:

| Setting | Value |
|-----------------------|--|
| Audit log destination | <div><div></div><div>Storage queue</div><div>Event Hub</div><div>Event Grid</div><div>Blob storage</div></div> |
| Detection app service | <div><div></div><div>Function App</div><div>Web App</div><div>API App</div></div> |

Explanation:

Box 1: Blob storage

Configure blob storage for audit logs.

Scenario: Unauthorized usage of the Planning Assistance data must be detected as quickly as possible.

Unauthorized usage is determined by looking for an unusual pattern of usage.

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

Box 2: Web Apps

SQL Advanced Threat Protection (ATP) is to be used.

One of Azure's most popular service is App Service which enables customers to build and host web applications in the programming language of their choice without managing infrastructure. App Service offers auto-scaling and high availability, supports both Windows and Linux. It also supports automated

deployments from GitHub, Visual Studio Team Services or any Git repository. At RSA, we announced that Azure Security Center leverages the scale of the cloud to identify attacks targeting App Service applications.

Reference:

<https://azure.microsoft.com/sv-se/blog/azure-security-center-can-identify-attacks-targeting-azure-app-service-applications/>

4.HOTSPOT

You need to design the authentication and authorization methods for sensors.

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

| Requirement | Method |
|----------------|---|
| Authentication | <div><div></div><div><div>▼</div><div>HMAC header</div><div>Resource Token</div><div>Azure Managed Identity</div><div>Storage account connection string</div></div></div> |
| Authorization | <div><div></div><div><div>▼</div><div>Custom RBAC role</div><div>Cosmos DB user</div><div>Azure Active Directory user</div><div>IoT device identity</div></div></div> |

Answer:

Requirement**Method****Authentication**

| |
|--|
| <div>▼</div> <div>HMAC header</div> <div>Resource Token</div> <div>Azure Managed Identity</div> <div>Storage account connection string</div> |
|--|

Authorization

| |
|--|
| <div>▼</div> <div>Custom RBAC role</div> <div>Cosmos DB user</div> <div>Azure Active Directory user</div> <div>IoT device identity</div> |
|--|

Explanation:

Sensor data must be stored in a Cosmos DB named treydata in a collection named SensorData Sensors must have permission only to add items to the SensorData collection

Box 1: Resource Token

Resource tokens provide access to the application resources within a Cosmos DB database. Enable clients to read, write, and delete resources in the Cosmos DB account according to the permissions they've been granted.

Box 2: Cosmos DB user

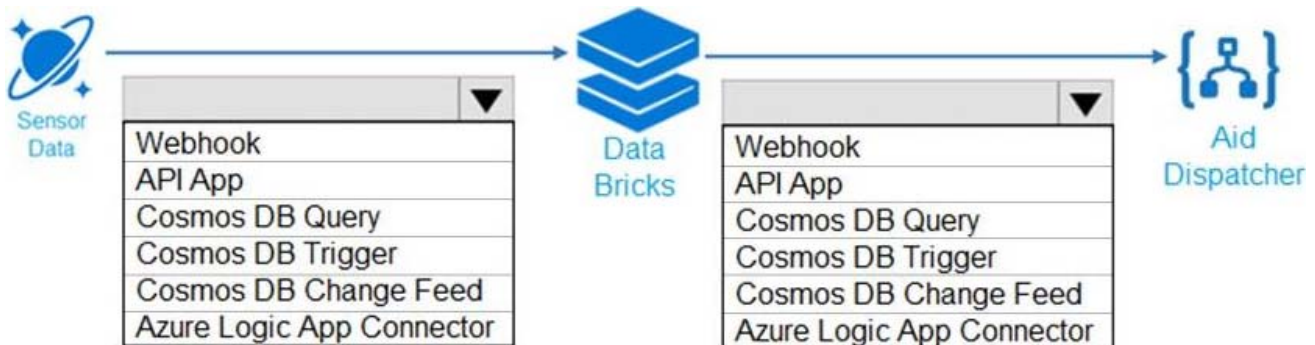
You can use a resource token (by creating Cosmos DB users and permissions) when you want to provide access to resources in your Cosmos DB account to a client that cannot be trusted with the master key.

Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/secure-access-to-data>

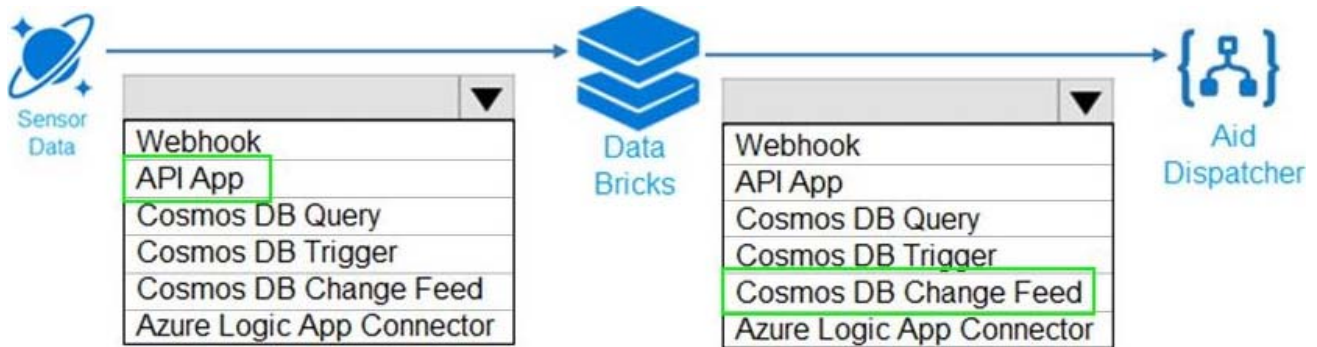
5.HOTSPOT

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.

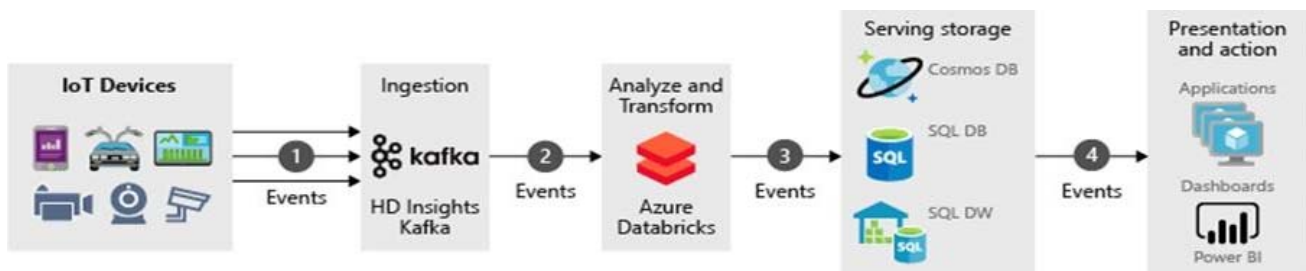


Answer:



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.

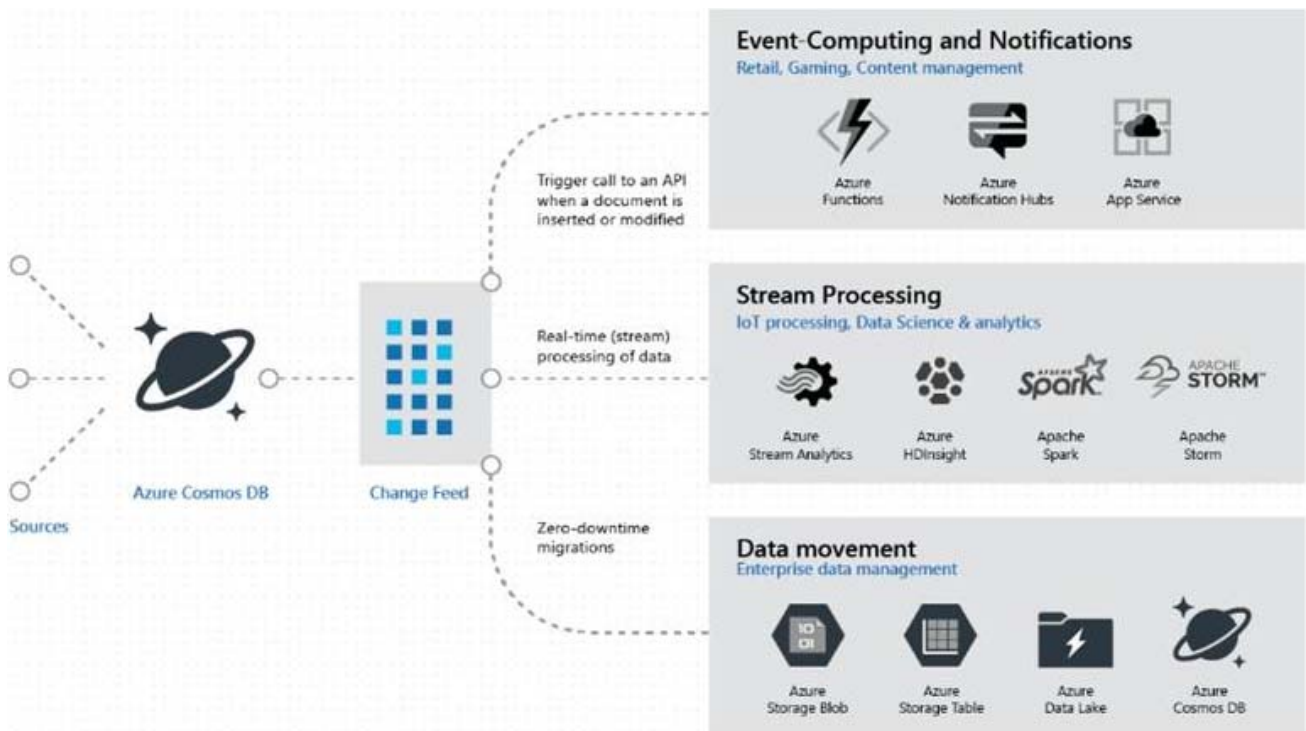
⇒ Downstream storage services, like Azure Cosmos DB, Azure SQL Data warehouse, or Azure SQL DB, will then be a data source for presentation and action layer.

⇒ 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=azuremps-4.4.1>