

Microsoft

DP-201 Exam

Designing an Azure Data Solution

Questions & Answers

Demo

Version: 17.0

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 question on 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 sections 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 on 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

a. 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 (km/s)

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.

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:

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 style="border: 1px solid black; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> strong session eventual consistent prefix bounded staleness </div>
Partition key property	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> Time Location Speed License plate Vehicle length </div>

Answer:

Setting	Value
Default consistency level	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> strong session <li style="background-color: #d0d0d0;">eventual consistent prefix bounded staleness </div>
Partition key property	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> Time Location Speed <li style="background-color: #d0d0d0;">License plate Vehicle length </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

HOTSPOT

You need to design the Planning Assistance database.

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation:

Box 1: No

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

Box 2: Yes

Box 3: Yes

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