

DP-201 Dumps

Designing an Azure Data Solution

<https://www.certleader.com/DP-201-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

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"/>

- A. Mastered
- B. Not Mastered

Answer: A

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

NEW QUESTION 2

- (Exam Topic 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>

NEW QUESTION 3

- (Exam Topic 1)

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

NEW QUESTION 4

- (Exam Topic 1)

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: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</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: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <ul style="list-style-type: none"> Time Location Speed License plate Vehicle length </div>

- A. Mastered
- B. Not Mastered

Answer: A

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>

NEW QUESTION 5

- (Exam Topic 1)

STION NO: 5 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 style="border: 1px solid black; padding: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <ul style="list-style-type: none"> HMAC header Resource Token Azure Managed Identity Storage account connection string </div>
Authorization	<div style="border: 1px solid black; padding: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <ul style="list-style-type: none"> Custom RBAC role Cosmos DB user Azure Active Directory user IoT device identity </div>

- A. Mastered
- B. Not Mastered

Answer: A

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.

References:

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

NEW QUESTION 6

- (Exam Topic 1)

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>

NEW QUESTION 7

- (Exam Topic 2)

You need to recommend a solution for storing customer data. What should you recommend?

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

Answer: C

Explanation:

From the scenario:

Customer data must be analyzed using managed Spark clusters.

All cloud data must be encrypted at rest and in transit. The solution must support: parallel processing of customer data.

References:

<https://www.microsoft.com/developerblog/2019/01/18/running-parallel-apache-spark-notebook-workloads-on-a>

NEW QUESTION 8

- (Exam Topic 2)

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>

NEW QUESTION 9

- (Exam Topic 2)

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.

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

A. Mastered

B. Not Mastered

Answer: A

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.

NEW QUESTION 10

- (Exam Topic 2)
You need to design the solution for analyzing customer data. What should you recommend?

- A. Azure Databricks
- B. Azure Data Lake Storage
- C. Azure SQL Data Warehouse
- D. Azure Cognitive Services
- E. Azure Batch

Answer: A

Explanation:

Customer data must be analyzed using managed Spark clusters. You create spark clusters through Azure Databricks. References:
<https://docs.microsoft.com/en-us/azure/azure-databricks/quickstart-create-databricks-workspace-portal>

NEW QUESTION 10

- (Exam Topic 2)
You need to design the encryption strategy for the tagging data and customer data.
What should you recommend? 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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

All cloud data must be encrypted at rest and in transit. Box 1: Transparent data encryption
Encryption of the database file is performed at the page level. The pages in an encrypted database are encrypted before they are written to disk and decrypted when read into memory.
Box 2: Encryption at rest
Encryption at Rest is the encoding (encryption) of data when it is persisted. References:
<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption?view=> <https://docs.microsoft.com/en-us/azure/security/azure-security-encryption-atrest>

NEW QUESTION 11

- (Exam Topic 2)
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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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>

NEW QUESTION 15

- (Exam Topic 3)

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

Answer: CDE

NEW QUESTION 16

- (Exam Topic 3)

You need to design a solution to meet the SQL Server storage requirements for CONT_SQL3. Which type of disk should you recommend?

- A. Standard SSD Managed Disk
- B. Premium SSD Managed Disk
- C. Ultra SSD Managed Disk

Answer: C

Explanation:

CONT_SQL3 requires an initial scale of 35000 IOPS.

Disk size (GiB)	4	8	16	32	64	128	256	512	1,024-65,536 (in increments of 1 TiB)
IOPS range	100-1,200	100-2,400	100-4,800	100-9,600	100-19,200	100-38,400	100-76,800	100-153,600	100-160,000
Throughput Cap (MBps)	300	600	1,200	2,000	2,000	2,000	2,000	2,000	2,000

The following table provides a comparison of ultra solid-state-drives (SSD) (preview), premium SSD, standard SSD, and standard hard disk drives (HDD) for managed disks to help you decide what to use.

	Ultra SSD (preview)	Premium SSD	Standard SSD	Standard HDD
Disk type	SSD	SSD	SSD	HDD
Scenario	IO-intensive workloads such as SAP HANA, top tier databases (for example, SQL Oracle), and other transaction-heavy workloads.	Production and performance sensitive workloads	Web servers, lightly used enterprise applications and dev/test	Backup, non-critical, infrequent access
Disk size	65,536 gibibyte (GiB) (Preview)	32,767 GiB	32,767 GiB	32,767 GiB
Max throughput	2,000 MiB/s (Preview)	900 MiB/s	750 MiB/s	500 MiB/s
Max IOPS	160,000 (Preview)	20,000	6,000	2,000

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-types>

NEW QUESTION 20

- (Exam Topic 3)

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

Answer: B

Explanation:

References:

<https://cloudblogs.microsoft.com/sqlserver/2018/12/17/confidential-computing-using-always-encrypted-withsec>

NEW QUESTION 24

- (Exam Topic 3)

You need to optimize storage for CONT_SQL3. What should you recommend?

- A. AlwaysOn
- B. Transactional processing
- C. General
- D. Data warehousing

Answer: B

Explanation:

CONT_SQL3 with the SQL Server role, 100 GB database size, Hyper-VM to be migrated to Azure VM. The storage should be configured to optimized storage for database OLTP workloads.

Azure SQL Database provides three basic in-memory based capabilities (built into the underlying database engine) that can contribute in a meaningful way to performance improvements:

In-Memory Online Transactional Processing (OLTP)

Clustered columnstore indexes intended primarily for Online Analytical Processing (OLAP) workloads Nonclustered columnstore indexes geared towards Hybrid

Transactional/Analytical Processing (HTAP) workloads

References:

<https://www.databasejournal.com/features/mssql/overview-of-in-memory-technologies-of-azure-sqldatabase.htm>

NEW QUESTION 28

- (Exam Topic 3)

You need to recommend a backup strategy for CONT_SQL1 and CONT_SQL2. What should you recommend?

- A. Use AzCopy and store the data in Azure.
- B. Configure Azure SQL Database long-term retention for all databases.
- C. Configure Accelerated Database Recovery.
- D. Use Dataloader.

Answer: B

Explanation:

Scenario: The database backups have regulatory purposes and must be retained for seven years.

NEW QUESTION 31

- (Exam Topic 3)

You need to recommend an Azure SQL Database service tier. What should you recommend?

- A. Business Critical
- B. General Purpose
- C. Premium
- D. Standard
- E. Basic

Answer: C

Explanation:

The data engineers must set the SQL Data Warehouse compute resources to consume 300 DWUs. Note: There are three architectural models that are used in Azure SQL Database:

- General Purpose/Standard
- Business Critical/Premium
- Hyperscale

NEW QUESTION 35

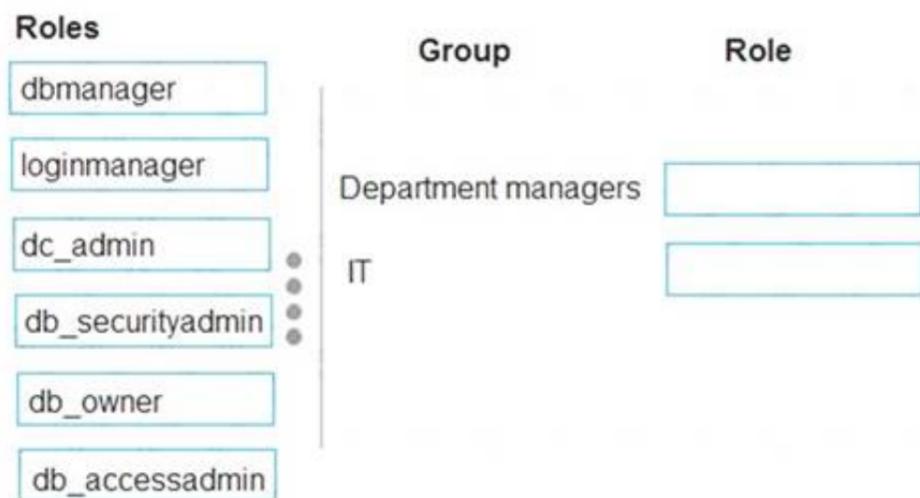
- (Exam Topic 3)

You are designing an Azure SQL Data Warehouse 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.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: dbmanager

Members of the dbmanager role can create new databases. Box 2: db_accessadmin

Members of the db_accessadmin fixed database role can add or remove access to the database for Windows logins, Windows groups, and SQL Server logins.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-manage-logins>

NEW QUESTION 37

- (Exam Topic 4)

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

Answer: C

Explanation:

<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>

NEW QUESTION 38

- (Exam Topic 4)

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.

Requirement	Service
Integrate the on-premises data into the cloud.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch
Develop notebooks to transform the data.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch
Run notebooks.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch
Load the data.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch
Store the transformed data.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Requirement	Service
Integrate the on-premises data into the cloud.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch
Develop notebooks to transform the data.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch
Run notebooks.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch
Load the data.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch
Store the transformed data.	<ul style="list-style-type: none"> Azure Databricks Azure Data Factory Azure SQL Data Warehouse Azure Batch

NEW QUESTION 39

- (Exam Topic 4)

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:

New blobs must be processed with a little delay as possible. 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.

Answer: C

Explanation:

Create a function, with the help of a blob trigger template, which is triggered when files are uploaded to or updated in Azure Blob storage.

You use a consumption plan, which is a hosting plan that defines how resources are allocated to your function app. In the default Consumption Plan, resources are added dynamically as required by your functions. In this serverless hosting, you only pay for the time your functions run. When you run in an App Service plan, you must manage the scaling of your function app.

References:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function>

NEW QUESTION 43

- (Exam Topic 4)

You need to design the unauthorized data usage detection system. What Azure service should you include in the design?

- A. Azure Databricks
- B. Azure SQL Data Warehouse
- C. Azure Analysis Services
- D. Azure Data Factory

Answer: B

NEW QUESTION 47

- (Exam Topic 4)

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

Answer: AE

Explanation:

To implement a lambda architecture on Azure, you can combine the following technologies to accelerate realtime big data analytics:

Azure Cosmos DB, the industry's first globally distributed, multi-model database service.

Apache Spark for Azure HDInsight, a processing framework that runs large-scale data analytics applications

Azure Cosmos DB change feed, which streams new data to the batch layer for HDInsight to process The Spark to Azure Cosmos DB Connector

E: You can use Apache Spark to stream data into or out of Apache Kafka on HDInsight using DStreams. References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/lambda-architecture>

NEW QUESTION 48

- (Exam Topic 4)

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: Create a user-defined restore point before data is uploaded. Delete the restore point after data corruption checks complete.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

User-Defined Restore Points

This feature enables you to manually trigger snapshots to create restore points of your data warehouse before and after large modifications. This capability ensures that restore points are logically consistent, which provides additional data protection in case of any workload interruptions or user errors for quick recovery time.

Note: A data warehouse restore is a new data warehouse that is created from a restore point of an existing or deleted data warehouse. Restoring your data warehouse is an essential part of any business continuity and disaster recovery strategy because it re-creates your data after accidental corruption or deletion.

References:

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

NEW QUESTION 52

- (Exam Topic 4)

You plan to deploy an Azure SQL Database instance to support an application. You plan to use the DTUbased purchasing model. Backups of the database must be available for 30 days and point-in-time restoration must be possible. You need to recommend a backup and recovery policy. What are two possible ways to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

- A. Use the Premium tier and the default backup retention policy.
- B. Use the Basic tier and the default backup retention policy.
- C. Use the Standard tier and the default backup retention policy.
- D. Use the Standard tier and configure a long-term backup retention policy.
- E. Use the Premium tier and configure a long-term backup retention policy.

Answer: DE

Explanation:

The default retention period for a database created using the DTU-based purchasing model depends on the service tier:

- ▶ Basic service tier is 1 week.
- ▶ Standard service tier is 5 weeks.
- ▶ Premium service tier is 5 weeks.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-long-term-retention>

NEW QUESTION 54

- (Exam Topic 4)

You have an on-premises MySQL database that is 800 GB in size. You need to migrate a MySQL database to Azure Database for MySQL. You must minimize service interruption to live sites or applications that use the database. What should you recommend?

- A. Azure Database Migration Service
- B. Dump and restore
- C. Import and export
- D. MySQL Workbench

Answer: A

Explanation:

You can perform MySQL migrations to Azure Database for MySQL with minimal downtime by using the newly introduced continuous sync capability for the Azure Database Migration Service (DMS). This functionality limits the amount of downtime that is incurred by the application. References:
<https://docs.microsoft.com/en-us/azure/mysql/howto-migrate-online>

NEW QUESTION 58

- (Exam Topic 4)

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.

Use case	Technology
Data abstraction	<div style="border: 1px solid gray; padding: 2px;"> ▼ Resilient Distributed Dataset (RDD) Dataset DataFrame </div>
Data format	<div style="border: 1px solid gray; padding: 2px;"> ▼ Avro parquet </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: DataFrame
DataFrames
Best choice in most situations.
Provides query optimization through Catalyst. Whole-stage code generation.
Direct memory access.
Low garbage collection (GC) overhead.

Not as developer-friendly as DataSets, as there are no compile-time checks or domain object programming. Box 2: parquet
The best format for performance is parquet with snappy compression, which is the default in Spark 2.x. Parquet stores data in columnar format, and is highly optimized in Spark.

NEW QUESTION 60

- (Exam Topic 4)

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

Answer: A

Explanation:

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:

Lowering the authentication checks across multiple files
Reduced open file connections

Faster copying/replication

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>

NEW QUESTION 63

- (Exam Topic 4)

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

Answer: B

Explanation:

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>

NEW QUESTION 68

- (Exam Topic 4)

A company is designing a solution that uses Azure Databricks.

The solution must be resilient to regional Azure datacenter outages. You need to recommend the redundancy type for the solution. What should you recommend?

- A. Read-access geo-redundant storage
- B. Locally-redundant storage
- C. Geo-redundant storage
- D. Zone-redundant storage

Answer: C

Explanation:

If your storage account has GRS enabled, then your data is durable even in the case of a complete regional outage or a disaster in which the primary region isn't recoverable.

References:

<https://medium.com/microsoftazure/data-durability-fault-tolerance-resilience-in-azure-databricks-95392982bac7>

NEW QUESTION 69

- (Exam Topic 4)

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. Which Azure service should you recommend?

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

Answer: D

Explanation:

A databrick job is a way of running a notebook or JAR either immediately or on a scheduled basis.

Azure Databricks has two types of clusters: interactive and job. Interactive clusters are used to analyze data collaboratively with interactive notebooks. Job clusters are used to run fast and robust automated workloads using the UI or API.

You can visualize Data with Azure Databricks and Power BI Desktop.

References:

<https://docs.azuredatabricks.net/user-guide/clusters/index.html> <https://docs.azuredatabricks.net/user-guide/jobs.html>

NEW QUESTION 71

- (Exam Topic 4)

A company has locations in North America and Europe. The company uses Azure SQL Database to support business apps.

Employees must be able to access the app data in case of a region-wide outage. A multi-region availability solution is needed with the following requirements:

- Read-access to data in a secondary region must be available only in case of an outage of the primary region.
- The Azure SQL Database compute and storage layers must be integrated and replicated together.

You need to design the multi-region high availability solution.

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

NOTE: Each correct selection is worth one point.

Option	Value
Service tier	<div style="border: 1px solid gray; padding: 5px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> Basic Standard General Premium </div>
Redundancy type	<div style="border: 1px solid gray; padding: 5px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> SQL Sync Zone-redundancy Geo-redundant storage </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Standard

The following table describes the types of storage accounts and their capabilities:

Storage account type	Supported services	Supported performance tiers	Supported access tiers	Replication options	Deployment model ¹
General-purpose V2	Blob, File, Queue, Table, and Disk	Standard, Premium ⁵	Hot, Cool, Archive ³	LRS, ZRS ⁴ , GRS, RA-GRS	Resource Manager
General-purpose V1	Blob, File, Queue, Table, and Disk	Standard, Premium ⁵	N/A	LRS, GRS, RA-GRS	Resource Manager, Classic

Box 2: Geo-redundant storage

If your storage account has GRS enabled, then your data is durable even in the case of a complete regional outage or a disaster in which the primary region isn't recoverable.

Note: If you opt for GRS, you have two related options to choose from:

GRS replicates your data to another data center in a secondary region, but that data is available to be read only if Microsoft initiates a failover from the primary to secondary region.

Read-access geo-redundant storage (RA-GRS) is based on GRS. RA-GRS replicates your data to another data center in a secondary region, and also provides you with the option to read from the secondary region. With RA-GRS, you can read from the secondary region regardless of whether Microsoft initiates a failover from the primary to secondary region.

Scenario	LRS	ZRS	GRS	RA-GRS
Node unavailability within a data center	Yes	Yes	Yes	Yes
An entire data center (zonal or non-zonal) becomes unavailable	No	Yes	Yes	Yes
A region-wide outage	No	No	Yes	Yes

References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-introduction> <https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs>

NEW QUESTION 72

- (Exam Topic 4)

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.

Option	Value
File type for exporting the on-premises database	BACPAC DACPAC VHDX
Azure storage type for exported data	Blob Disk Table File

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-import>

NEW QUESTION 76

- (Exam Topic 4)

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

Answer: B

Explanation:

Ensure that files stored are larger, not smaller than 250MB.

You can have a separate compaction job that combines these files into larger ones.

Note: The file 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:

Lowering the authentication checks across multiple files
Reduced open file connections

Faster copying/replication

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>

NEW QUESTION 77

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your DP-201 Exam with Our Prep Materials Via below:

<https://www.certleader.com/DP-201-dumps.html>