Microsoft

Exam Questions DP-201

Designing an Azure Data Solution

NEW QUESTION 1
- (Exam Topic 1)
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.

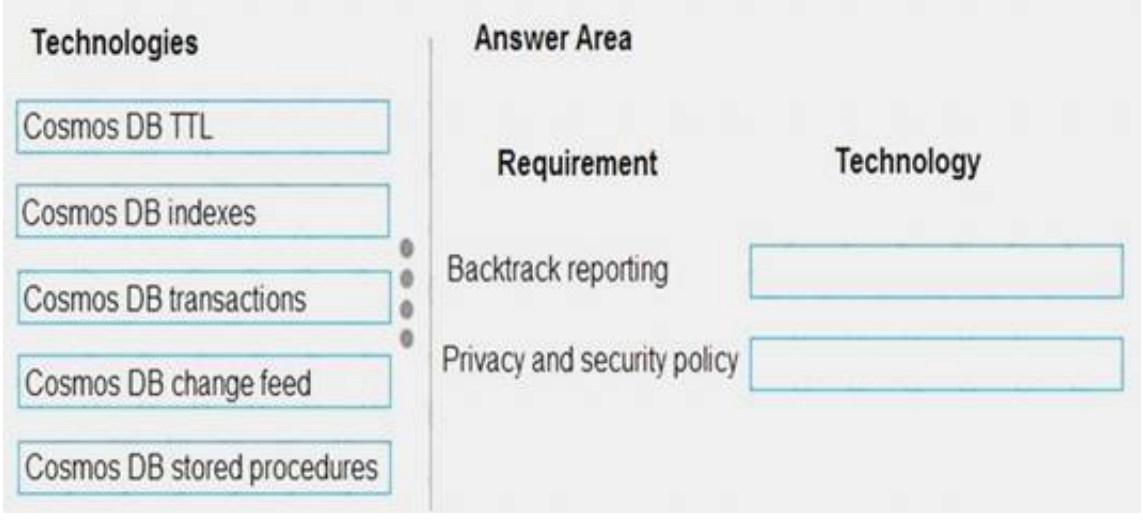

A. Mastered
B. Not Mastered

Answer: A

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.

NEW QUESTION 2
- (Exam Topic 1)
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.

NEW QUESTION 3
- (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.

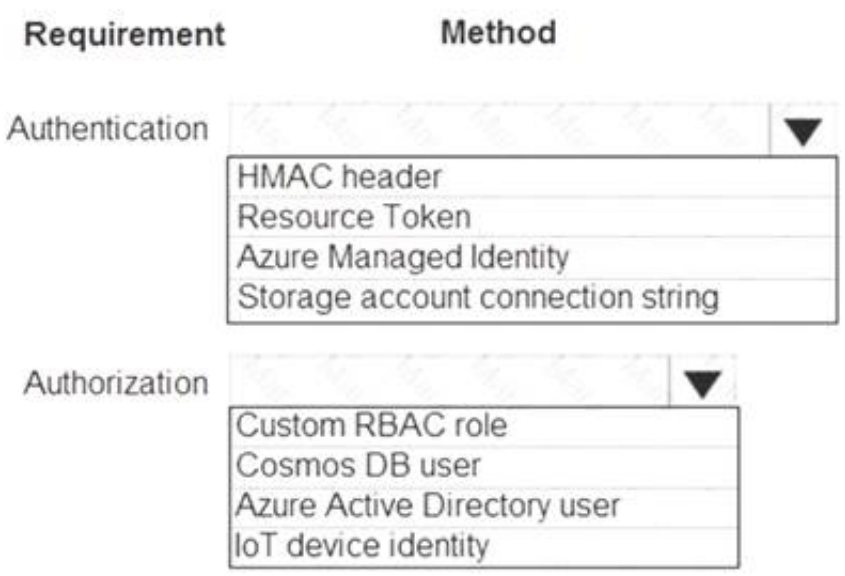

A. Mastered
B. Not Mastered

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 4
- (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 5
- (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 6
- (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 7
- (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.

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 8
- (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 9
- (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 10
- (Exam Topic 3)
You need to recommend the appropriate storage and processing solution? What should you recommend?

A. Enable auto-shrink on the database.
B. Flush the blob cache using Windows PowerShell.
C. Enable Apache Spark RDD (RDD) caching.
D. Enable Databricks IO (DBIO) caching.
E. Configure the reading speed using Azure Data Studio.

Answer: C

Explanation:
Scenario: You must be able to use a file system view of data stored in a blob. You must build an architecture that will allow Contoso to use the DB FS filesystem layer over a blob store.
Databricks File System (DBFS) is a distributed file system installed on Azure Databricks clusters. Files in DBFS persist to Azure Blob storage, so you won't lose data even after you terminate a cluster.
The Databricks Delta cache, previously named Databricks IO (DBIO) caching, accelerates data reads by creating copies of remote files in nodes' local storage using a fast intermediate data format. The data is cached automatically whenever a file has to be fetched from a remote location. Successive reads of the same data are then performed locally, which results in significantly improved reading speed.

NEW QUESTION 10
- (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 12

- (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.

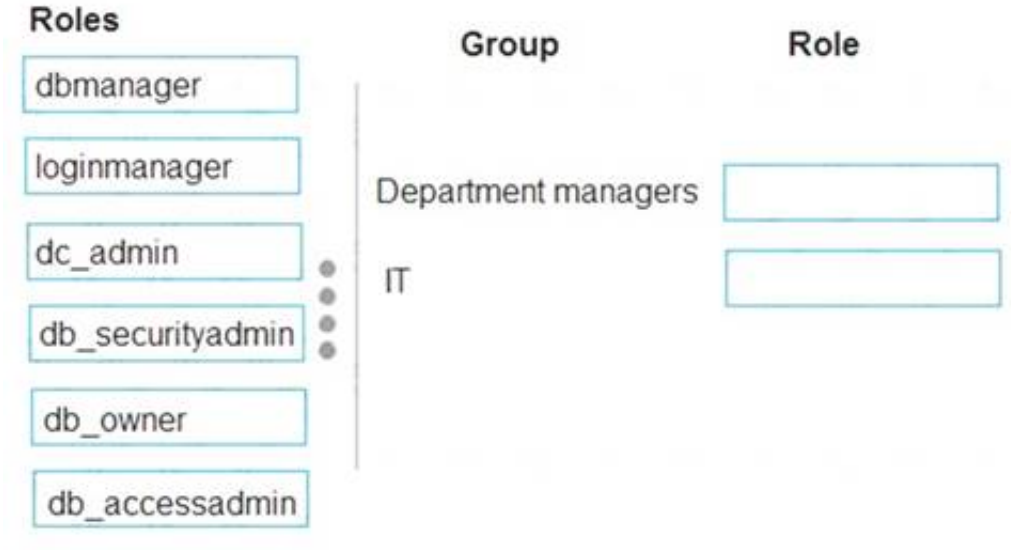

A. Mastered
B. Not Mastered

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 16
- (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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 17
- (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 18
- (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: Configure database-level auditing in Azure SQL Data Warehouse and set retention to 10 days.
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 22

- (Exam Topic 4)
You are designing a solution for a company. You plan to use Azure Databricks. You need to recommend workloads and tiers to meet the following requirements:
- Provide managed clusters for running production jobs.
- Provide persistent clusters that support auto-scaling for analytics processes.
- Provide role-based access control (RBAC) support for Notebooks.

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

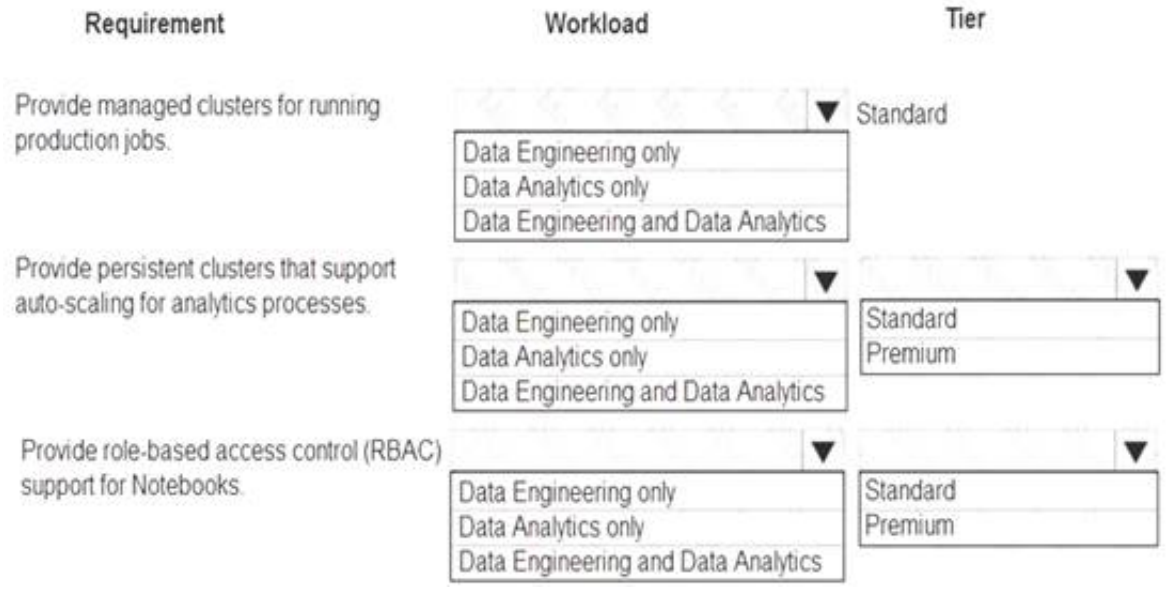

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Data Engineering Only
Box 2: Data Engineering and Data Analytics Box 3: Standard
Box 4: Data Analytics only Box 5: Premium
Premium required for RBAC. Data Analytics Premium Tier provide interactive workloads to analyze data collaboratively with notebooks
References:
https://azure.microsoft.com/en-us/pricing/details/databricks/

NEW QUESTION 27
- (Exam Topic 4)
A company is evaluating data storage solutions.
You need to recommend a data storage solution that meets the following requirements: Minimize costs for storing blob objects.
Optimize access for data that is infrequently accessed. Data must be stored for at least 30 days.
Data availability must be at least 99 percent. What should you recommend?

A. Premium
B. Cold
C. Hot
D. Archive

Answer: B

Explanation:
Azure's cool storage tier, also known as Azure cool Blob storage, is for infrequently-accessed data that needs to be stored for a minimum of 30 days. Typical use cases include backing up data before tiering to archival systems, legal data, media files, system audit information, datasets used for big data analysis and more.
The storage cost for this Azure cold storage tier is lower than that of hot storage tier. Since it is expected that the data stored in this tier will be accessed less frequently, the data access charges are high when compared to hot tier. There are no additional changes required in your applications as these tiers can be accessed using
APIs in the same manner that you access Azure storage. References:
https://cloud.netapp.com/blog/low-cost-storage-options-on-azure

NEW QUESTION 28
- (Exam Topic 4)
A company installs IoT devices to monitor its fleet of delivery vehicles. Data from devices is collected from Azure Event Hub.
The data must be transmitted to Power BI for real-time data visualizations. You need to recommend a solution.
What should you recommend?

A. Azure HDInsight with Spark Streaming
B. Apache Spark in Azure Databricks
C. Azure Stream Analytics
D. Azure HDInsight with Storm

Answer: C

Explanation:
Step 1: Get your IoT hub ready for data access by adding a consumer group.
Step 2: Create, configure, and run a Stream Analytics job for data transfer from your IoT hub to your Power BI account.
Step 3: Create and publish a Power BI report to visualize the data. References:
https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-live-data-visualization-in-power-bi

NEW QUESTION 29
- (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 34
- (Exam Topic 4)
You are designing a recovery strategy for your Azure SQL Databases.
The recovery strategy must use default automated backup settings. The solution must include a Point-in time restore recovery strategy.
You need to recommend which backups to use and the order in which to restore backups.
What should you recommend? To answer, select the appropriate configuration in the answer area.
NOTE: Each correct selection is worth one point.

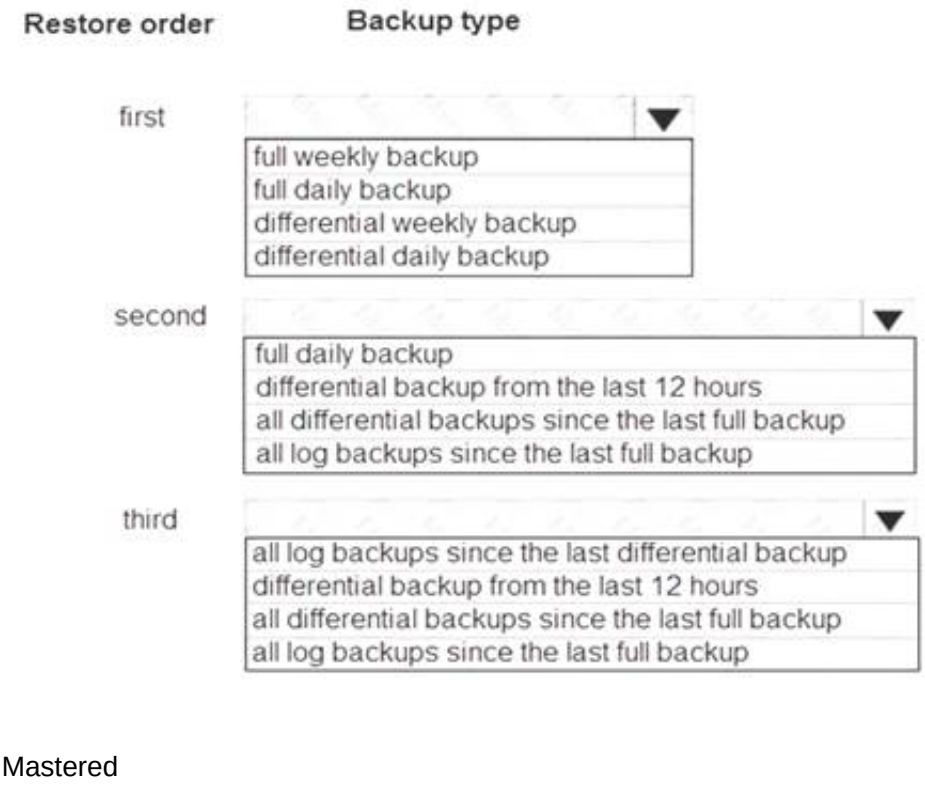

A. Mastered
B. Not Mastered

Answer: A

Explanation:
All Basic, Standard, and Premium databases are protected by automatic backups. Full backups are taken every week, differential backups every day, and log backups every 5 minutes.
References:
https://azure.microsoft.com/sv-se/blog/azure-sql-database-point-in-time-restore/

NEW QUESTION 35
- (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 39
- (Exam Topic 4)
You are designing a solution for a company. The solution will use model training for objective classification. You need to design the solution.
What should you recommend?

A. an Azure Cognitive Services application
B. a Spark Streaming job
C. interactive Spark queries
D. Power BI models
E. a Spark application that uses Spark MLib.

Answer: E

Explanation:
Spark in SQL Server big data cluster enables AI and machine learning.
You can use Apache Spark MLlib to create a machine learning application to do simple predictive analysis on an open dataset.
MLlib is a core Spark library that provides many utilities useful for machine learning tasks, including utilities that are suitable for:
- Classification
- Regression
- Clustering
- Topic modeling
- Singular value decomposition (SVD) and principal component analysis (PCA)
- Hypothesis testing and calculating sample statistics

References:
https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-machine-learning-mllib-ipython

NEW QUESTION 44
- (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 46
- (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 50
- (Exam Topic 4)
You need to design the storage for the telemetry capture system. What storage solution should you use in the design?

A. Azure Databricks

B. Azure SQL Data Warehouse
C. Azure Cosmos DB

Answer: C

NEW QUESTION 52
- (Exam Topic 4)
You are evaluating data storage solutions to support a new application.
You need to recommend a data storage solution that represents data by using nodes and relationships in graph structures.
Which data storage solution should you recommend?

A. Blob Storage
B. Cosmos DB
C. Data Lake Store
D. HDInsight

Answer: B

Explanation:
For large graphs with lots of entities and relationships, you can perform very complex analyses very quickly. Many graph databases provide a query language that you can use to traverse a network of relationships efficiently.
Relevant Azure service: Cosmos DB
References:
https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview

NEW QUESTION 53
- (Exam Topic 4)
You are designing an application. You plan to use Azure SQL Database to support the application.
The application will extract data from the Azure SQL Database and create text documents. The text documents will be placed into a cloud-based storage solution.
The text storage solution must be accessible from an SMB network share.
You need to recommend a data storage solution for the text documents. Which Azure data storage type should you recommend?

A. Queue
B. Files
C. Blob
D. Table

Answer: B

Explanation:
Azure Files enables you to set up highly available network file shares that can be accessed by using the standard Server Message Block (SMB) protocol.
References:
https://docs.microsoft.com/en-us/azure/storage/common/storage-introduction https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-overview

NEW QUESTION 56
- (Exam Topic 4)
You are designing an Azure Data Factory pipeline for processing data. The pipeline will process data that is stored in general-purpose standard Azure storage.
You need to ensure that the compute environment is created on-demand and removed when the process is completed.
Which type of activity should you recommend?

A. Databricks Python activity
B. Data Lake Analytics U-SQL activity
C. HDInsight Pig activity
D. Databricks Jar activity

Answer: C

Explanation:
The HDInsight Pig activity in a Data Factory pipeline executes Pig queries on your own or on-demand HDInsight cluster.
References:
https://docs.microsoft.com/en-us/azure/data-factory/transform-data-using-hadoop-pig

NEW QUESTION 60
- (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.

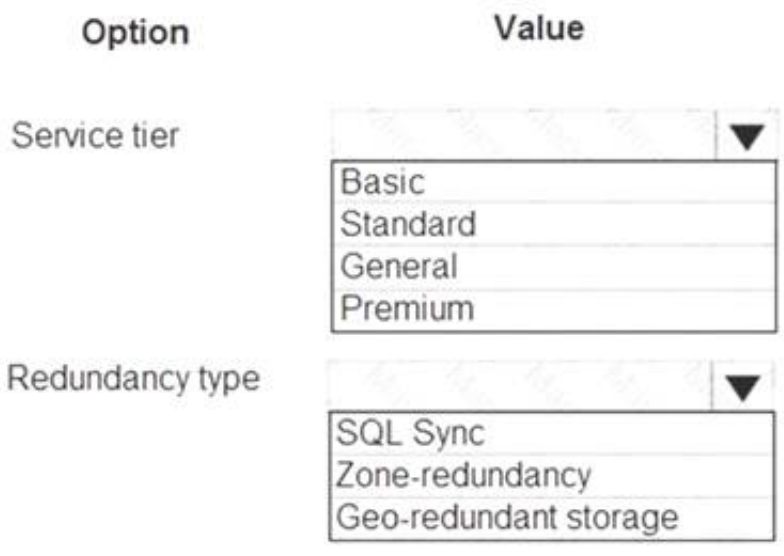

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Standard
The following table describes the types of storage accounts and their capabilities:

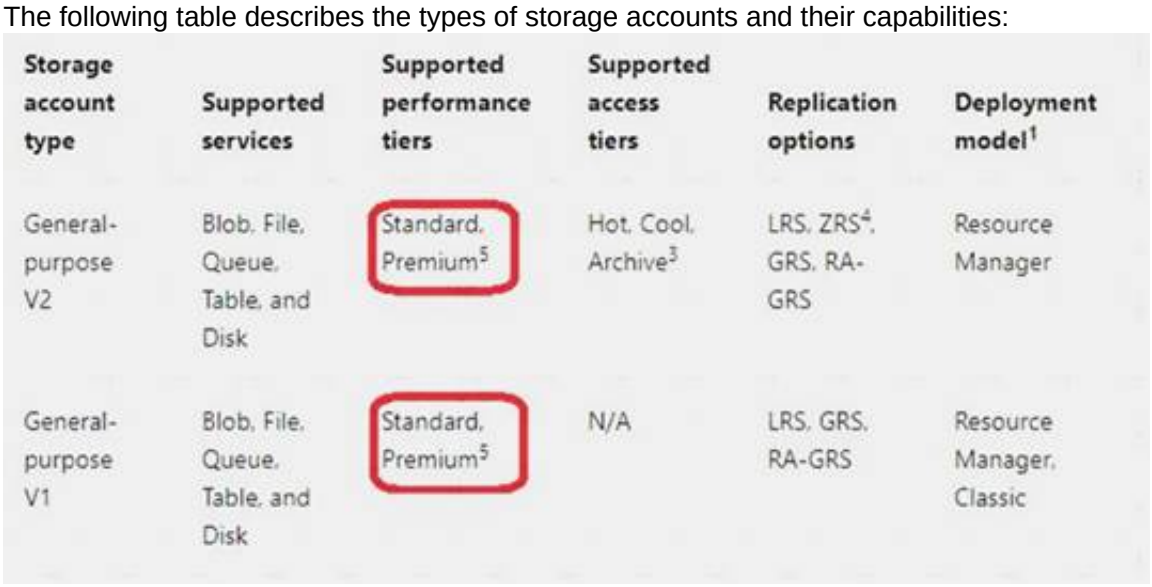

Storage account type	Supported services	Supported performance tiers	Supported access tiers	Replication options	Deployment model[1]
General-purpose V2	Blob, File, Queue, Table, and Disk	Standard, Premium[5]	Hot, Cool, Archive[3]	LRS, ZRS[4], GRS, RA-GRS	Resource Manager
General-purpose V1	Blob, File,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

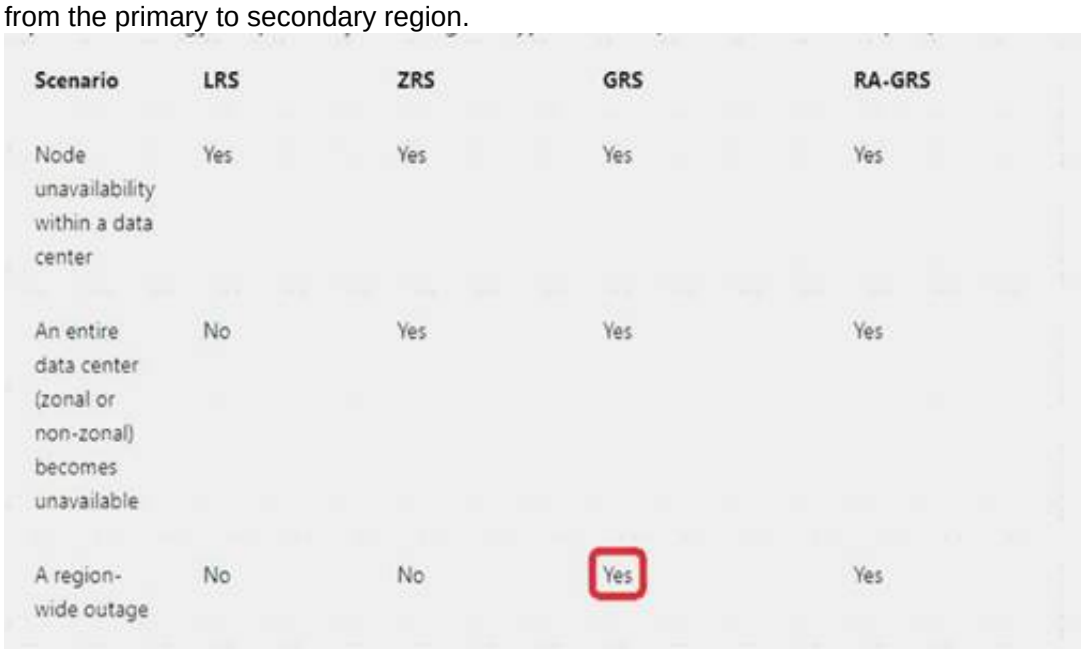

Scenario	LRS	ZRS	GRS	RA-GRS
Node unavailability within a data center	Yes	Yes	Yes	Yes
An entire data center (zonal or non-zonal) becomes unavailable	No	Yes	Yes	Yes
A region-wide outage	No	No	Yes	Yes

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 64
- (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 Azure SQL Database that will use elastic pools. You plan to store data about customers in a table. Each record uses a value for CustomerID.
You need to recommend a strategy to partition data based on values in CustomerID. Proposed Solution: Separate data into customer regions by using horizontal partitioning. Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
We should use Horizontal Partitioning through Sharding, not divide through regions.
Note: Horizontal Partitioning - Sharding: Data is partitioned horizontally to distribute rows across a scaled out data tier. With this approach, the schema is identical on all participating databases. This approach is also called “sharding”. Sharding can be performed and managed using (1) the elastic database tools libraries or (2)
self-sharding. An elastic query is used to query or compile reports across many shards.
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-query-overview

NEW QUESTION 69
- (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 Azure SQL Database that will use elastic pools. You plan to store data about customers in a table. Each record uses a value for CustomerID.
You need to recommend a strategy to partition data based on values in CustomerID. Proposed Solution: Separate data into customer regions by using vertical partitioning. Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Vertical partitioning is used for cross-database queries. Instead we should use Horizontal Partitioning, which also is called charding.
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-query-overview

NEW QUESTION 70
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

DP-201 Practice Exam Features:

* DP-201 Questions and Answers Updated Frequently

* DP-201 Practice Questions Verified by Expert Senior Certified Staff

* DP-201 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DP-201 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year