DP-201 Dumps

Designing an Azure Data Solution

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

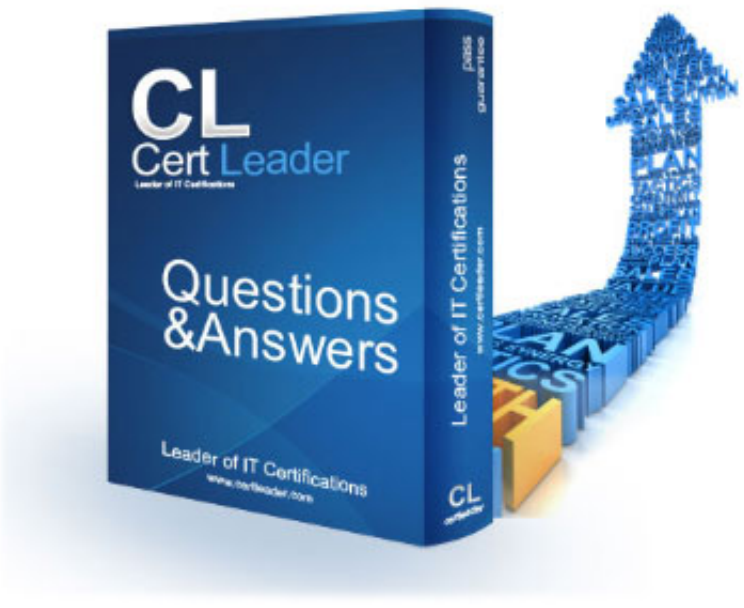

NEW QUESTION 1
- (Exam Topic 1)
You need to design the runtime environment for the Real Time Response system. What should you recommend?

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 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)
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.

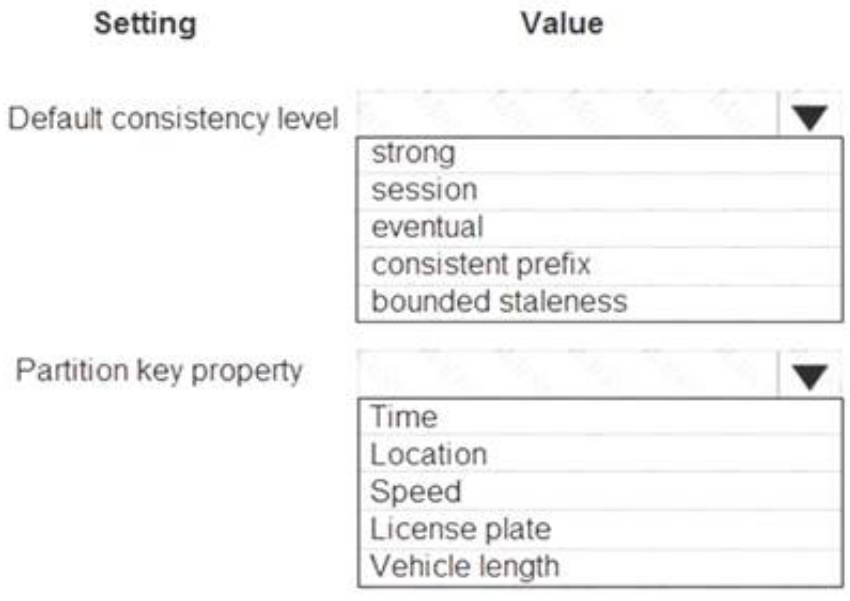

A. Mastered
B. Not Mastered

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 4
- (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 5
- (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 6
- (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 7
- (Exam Topic 2)
You need to design a backup solution for the processed customer data. What should you include in the design?

A. AzCopy
B. AdlCopy
C. Geo-Redundancy
D. Geo-Replication

Answer: C

Explanation:
Scenario: All data must be backed up in case disaster recovery is required.
Geo-redundant storage (GRS) is designed to provide at least 99.99999999999999% (16 9's) durability of objects over a given year by replicating your data to a secondary region that is hundreds of miles away from
the primary region. 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://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs

NEW QUESTION 8
- (Exam Topic 3)
You need to design the disaster recovery solution for customer sales data analytics.
Which three actions should you recommend? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Provision multiple Azure Databricks workspaces in separate Azure regions.
B. Migrate users, notebooks, and cluster configurations from one workspace to another in the same region.
C. Use zone redundant storage.
D. Migrate users, notebooks, and cluster configurations from one region to another.
E. Use Geo-redundant storage.
F. Provision a second Azure Databricks workspace in the same region.

Answer: ADE

Explanation:
Scenario: The analytics solution for customer sales data must be available during a regional outage. To create your own regional disaster recovery topology for databricks, follow these requirements:
1. Provision multiple Azure Databricks workspaces in separate Azure regions
2. Use Geo-redundant storage.
3. Once the secondary region is created, you must migrate the users, user folders, notebooks, cluster configuration, jobs configuration, libraries, storage, init scripts, and reconfigure access control.
Note: Geo-redundant storage (GRS) is designed to provide at least 99.99999999999999% (16 9's) durability of objects over a given year by replicating your data to a secondary region that is hundreds of miles away from the primary region. 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://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs

NEW QUESTION 9
- (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 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 10
- (Exam Topic 3)
You need to design network access to the SQL Server data.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

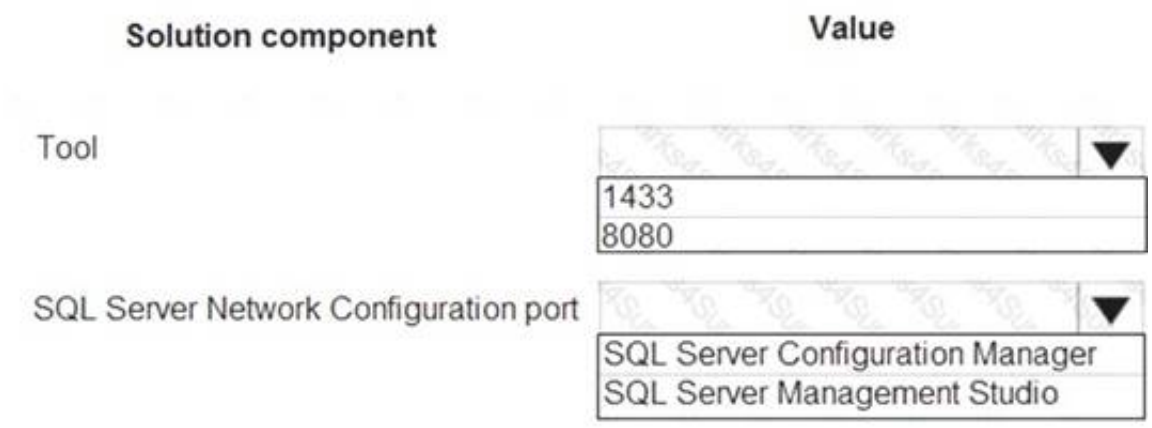

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: 8080
1433 is the default port, but we must change it as CONT_SQL3 must not communicate over the default ports. Because port 1433 is the known standard for SQL Server, some organizations specify that the SQL Server port number should be changed to enhance security.
Box 2: SQL Server Configuration Manager
You can configure an instance of the SQL Server Database Engine to listen on a specific fixed port by using the SQL Server Configuration Manager.
References:
https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-a-server-to-listen-on-a-speci

NEW QUESTION 13
- (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 15
- (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:

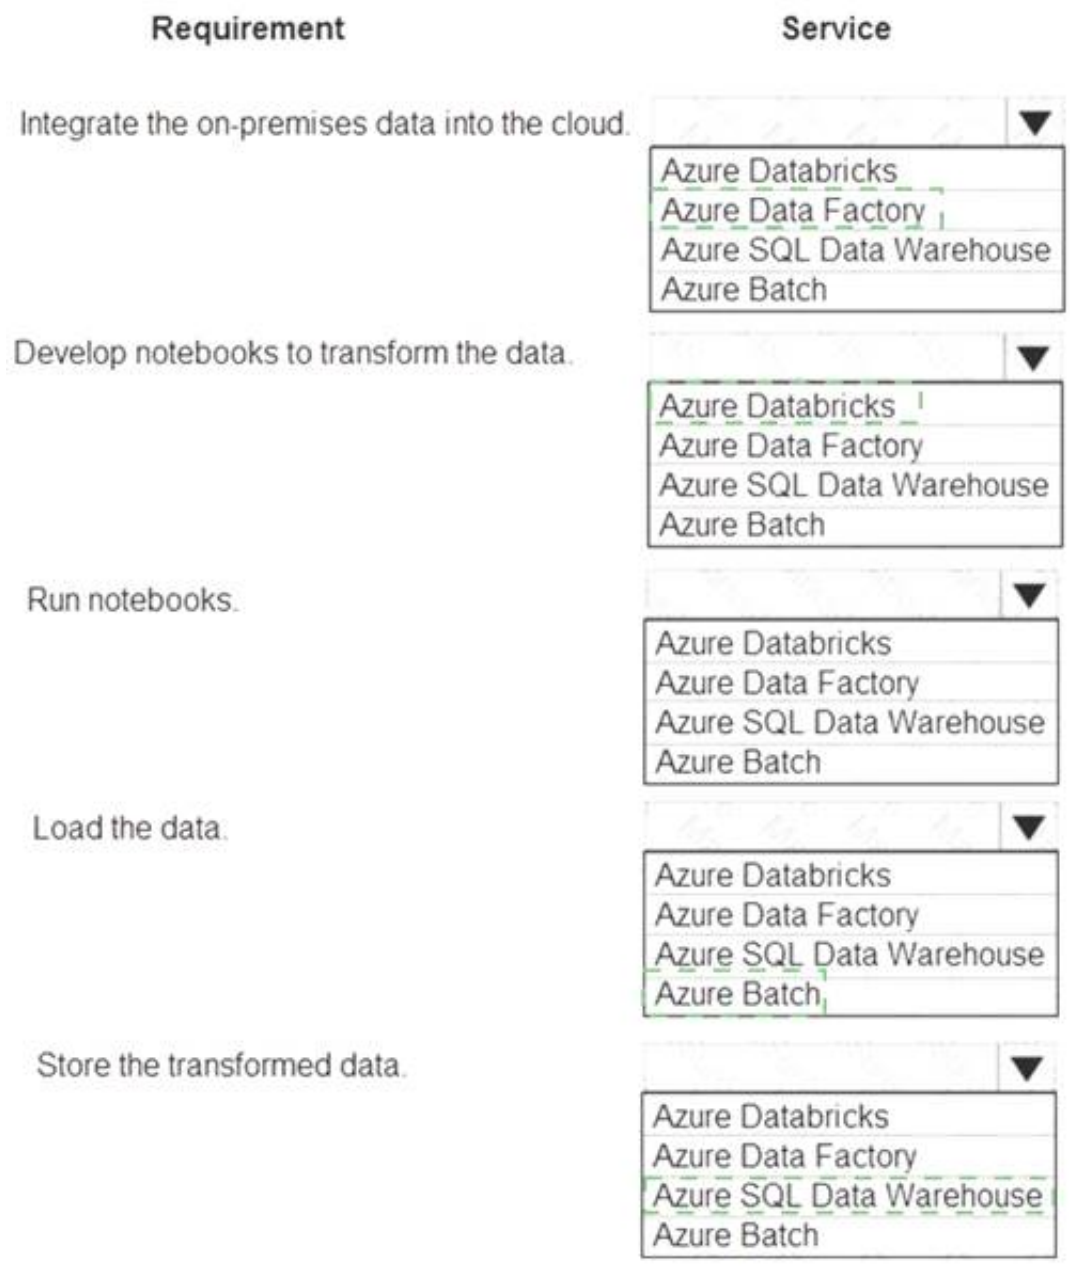

NEW QUESTION 19
- (Exam Topic 4)
You are designing a real-time stream solution based on Azure Functions. The solution will process data uploaded to Azure Blob Storage.

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 23
- (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)
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 larger than 250MB. 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 32
- (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 36
- (Exam Topic 4)
You are designing an Azure Databricks interactive cluster.
You need to ensure that the cluster meets the following requirements: Enable auto-termination
Retain cluster configuration indefinitely after cluster termination. What should you recommend?

A. Start the cluster after it is terminated.
B. Pin the cluster
C. Clone the cluster after it is terminated.
D. Terminate the cluster manually at process completion.

Answer: B

Explanation:
To keep an interactive cluster configuration even after it has been terminated for more than 30 days, an administrator can pin a cluster to the cluster list.
References:
https://docs.azuredatabricks.net/user-guide/clusters/terminate.html

NEW QUESTION 39
- (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 40
- (Exam Topic 4)
You are designing an Azure Databricks cluster that runs user-defined local processes. You need to recommend a cluster configuration that meets the following requirements:
• Minimize query latency.
• Reduce overall costs.
• Maximize the number of users that can run queries on the cluster at the same time. Which cluster type should you recommend?

A. Standard with Autoscaling
B. High Concurrency with Auto Termination
C. High Concurrency with Autoscaling
D. Standard with Auto Termination

Answer: C

Explanation:
High Concurrency clusters allow multiple users to run queries on the cluster at the same time, while minimizing query latency. Autoscaling clusters can reduce overall costs compared to a statically-sized cluster.
References:
https://docs.azuredatabricks.net/user-guide/clusters/create.html https://docs.azuredatabricks.net/user-guide/clusters/high-concurrency.html#high-concurrency https://docs.azuredatabricks.net/user-guide/clusters/terminate.html https://docs.azuredatabricks.net/user-guide/clusters/sizing.html#enable-and-configure-autoscaling

NEW QUESTION 42
- (Exam Topic 4)
A company has an application that uses Azure SQL Database as the data store.
The application experiences a large increase in activity during the last month of each year.
You need to manually scale the Azure SQL Database instance to account for the increase in data write operations.
Which scaling method should you recommend?

A. Scale up by using elastic pools to distribute resources.
B. Scale out by sharding the data across databases.
C. Scale up by increasing the database throughput units.

Answer: C

Explanation:
As of now, the cost of running an Azure SQL database instance is based on the number of Database Throughput Units (DTUs) allocated for the database. When determining the number of units to allocate for the
solution, a major contributing factor is to identify what processing power is needed to handle the volume of expected requests.
Running the statement to upgrade/downgrade your database takes a matter of seconds.

NEW QUESTION 44
- (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 48
- (Exam Topic 4)
A company is developing a mission-critical line of business app that uses Azure SQL Database Managed Instance. You must design a disaster recovery strategy for the solution.
You need to ensure that the database automatically recovers when full or partial loss of the Azure SQL Database service occurs in the primary region.
What should you recommend?

A. Failover-group
B. Azure SQL Data Sync
C. SQL Replication
D. Active geo-replication

Answer: A

Explanation:
Auto-failover groups is a SQL Database feature that allows you to manage replication and failover of a group of databases on a SQL Database server or all databases in a Managed Instance to another region (currently in public preview for Managed Instance). It uses the same underlying technology as active geo-replication. You can initiate failover manually or you can delegate it to the SQL Database service based on a user-defined policy.
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-auto-failover-group

NEW QUESTION 50
- (Exam Topic 4)
You have a Windows-based solution that analyzes scientific data. You are designing a cloud-based solution that performs real-time analysis of the data.
You need to design the logical flow for the solution.
Which two actions should you recommend? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Send data from the application to an Azure Stream Analytics job.
B. Use an Azure Stream Analytics job on an edge devic
C. Ingress data from an Azure Data Factory instance and build queries that output to Power BI.
D. Use an Azure Stream Analytics job in the clou
E. Ingress data from the Azure Event Hub instance and build queries that output to Power BI.
F. Use an Azure Stream Analytics job in the clou
G. Ingress data from an Azure Event Hub instance and build queries that output to Azure Data Lake Storage.
H. Send data from the application to Azure Data Lake Storage.
I. Send data from the application to an Azure Event Hub instance.

Answer: CF

Explanation:
Stream Analytics has first-class integration with Azure data streams as inputs from three kinds of resources: Azure Event Hubs
Azure IoT Hub Azure Blob storage References:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-define-inputs

NEW QUESTION 51
- (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.

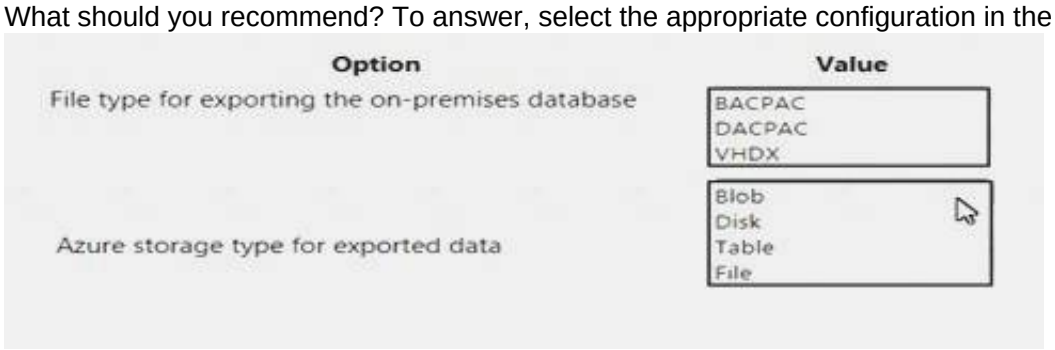

A. Mastered
B. Not Mastered

Answer: A

Explanation:
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-import

NEW QUESTION 55
- (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 shards by using horizontal partitioning.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
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) selfsharding.
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 56
- (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 57
- (Exam Topic 4)
You are developing a solution that performs real-time analysis of IoT data in the cloud. The solution must remain available during Azure service updates.
You need to recommend a solution.
Which two actions should you recommend? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Deploy an Azure Stream Analytics job to two separate regions that are not in a pair.
B. Deploy an Azure Stream Analytics job to each region in a paired region.
C. Monitor jobs in both regions for failure.
D. Monitor jobs in the primary region for failure.
E. Deploy an Azure Stream Analytics job to one region in a paired region.

Answer: BC

Explanation:

Stream Analytics guarantees jobs in paired regions are updated in separate batches. As a result there is a sufficient time gap between the updates to identify potential breaking bugs and remediate them.
Customers are advised to deploy identical jobs to both paired regions.
In addition to Stream Analytics internal monitoring capabilities, customers are also advised to monitor the jobs as if both are production jobs. If a break is identified to be a result of the Stream Analytics service update, escalate appropriately and fail over any downstream consumers to the healthy job output. Escalation to support will prevent the paired region from being affected by the new deployment and maintain the integrity of the paired jobs.
References:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-job-reliability

NEW QUESTION 59
- (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 61
.....

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