# DP-201 Dumps

# Designing an Azure Data Solution

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

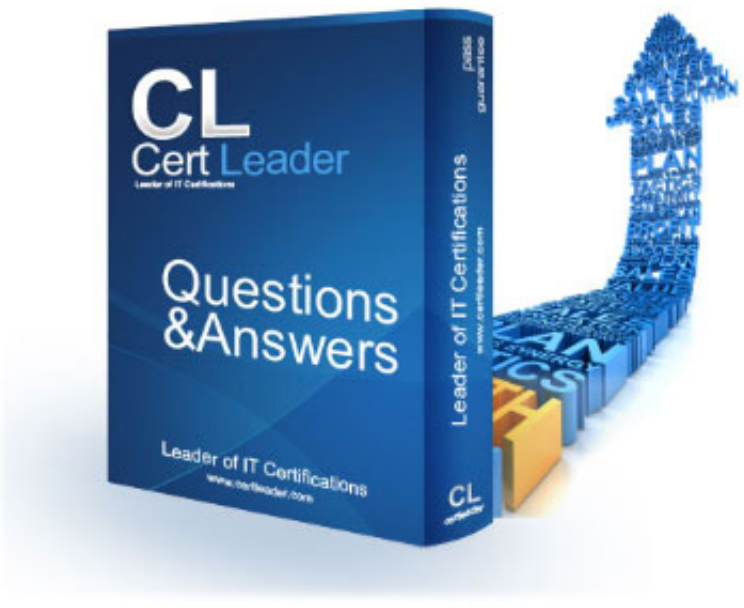

**NEW QUESTION 1**
- (Exam Topic 2)
You need to recommend a solution for storing customer data. What should you recommend?

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 2**
- (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 3**
- (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.
**Answer Area**

| Solution component | Tool |
|---|---|
| Image processing | Azure HDInsight<br>Azure Databricks<br>Azure Batch<br>Azure Cognitive Services |
| data storage for tagging data | Azure Blob Storage<br>Azure Table Storage<br>Azure Cosmos DB<br>Azure SQL Database |

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 4**
- (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 5**
- (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 6**
- (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 7**
- (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 8**
- (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.

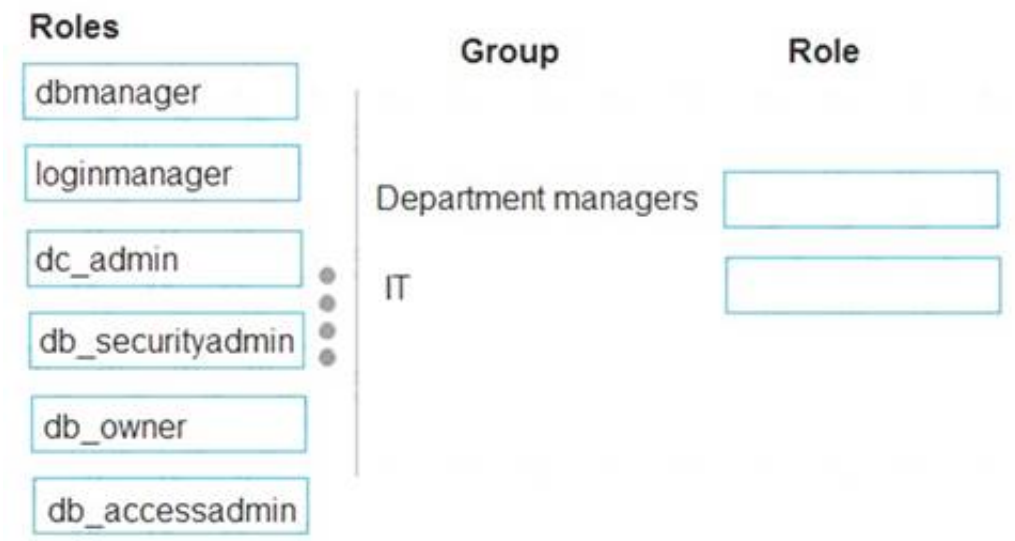


A. Mastered
B. Not Mastered

**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 9**
- (Exam Topic 4)
A company manufactures automobile parts. The company installs IoT sensors on manufacturing machinery. You must design a solution that analyzes data from the sensors.
You need to recommend a solution that meets the following requirements: Data must be analyzed in real-time.
Data queries must be deployed using continuous integration. Data must be visualized by using charts and graphs.
Data must be available for ETL operations in the future. The solution must support high-volume data ingestion.
Which three actions should you recommend? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Use Azure Analysis Services to query the dat
B. Output query results to Power BI.
C. Configure an Azure Event Hub to capture data to Azure Data Lake Storage.
D. Develop an Azure Stream Analytics application that queries the data and outputs to Power B
E. Use AzureData Factory to deploy the Azure Stream Analytics application.
F. Develop an application that sends the IoT data to an Azure Event Hub.
G. Develop an Azure Stream Analytics application that queries the data and outputs to Power B
H. Use AzurePipelines to deploy the Azure Stream Analytics application.
I. Develop an application that sends the IoT data to an Azure Data Lake Storage container.

**Answer:** BCD

**NEW QUESTION 10**
- (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 10**
- (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. | Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch |
| Develop notebooks to transform the data. | Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch |
| Run notebooks. | Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch |
| Load the data. | Azure Databricks / Azure Data Factory / Azure SQL Data Warehouse / Azure Batch |
| Store the transformed data. | 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. | Azure Databricks<br>Azure Data Factory<br>Azure SQL Data Warehouse<br>Azure Batch |
| Develop notebooks to transform the data. | Azure Databricks<br>Azure Data Factory<br>Azure SQL Data Warehouse<br>Azure Batch |
| Run notebooks. | Azure Databricks<br>Azure Data Factory<br>Azure SQL Data Warehouse<br>Azure Batch |
| Load the data. | Azure Databricks<br>Azure Data Factory<br>Azure SQL Data Warehouse<br>Azure Batch |
| Store the transformed data. | Azure Databricks<br>Azure Data Factory<br>Azure SQL Data Warehouse<br>Azure Batch |

**NEW QUESTION 12**
- (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 14**
- (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 17**
- (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 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.
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 23**
- (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 27**
- (Exam Topic 4)
You are designing an Azure SQL Data Warehouse. You plan to load millions of rows of data into the data warehouse each day.
You must ensure that staging tables are optimized for data loading. You need to design the staging tables.
What type of tables should you recommend?

A. Round-robin distributed table
B. Hash-distributed table
C. Replicated table
D. External table

**Answer:** A

**Explanation:**
To achieve the fastest loading speed for moving data into a data warehouse table, load data into a staging table. Define the staging table as a heap and use round-robin for the distribution option.
References:
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data

**NEW QUESTION 32**
- (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 36**
- (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 40**
- (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 41**

- (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 42**
- (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 43**
- (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 46**
- (Exam Topic 4)
A company stores large datasets in Azure, including sales transactions and customer account information. You must design a solution to analyze the data. You plan to create the following HDInsight clusters:
You need to ensure that the clusters support the query requirements.
Which cluster types should you recommend? To answer, select the appropriate configuration in the answer area.
NOTE: Each correct selection is worth one point.

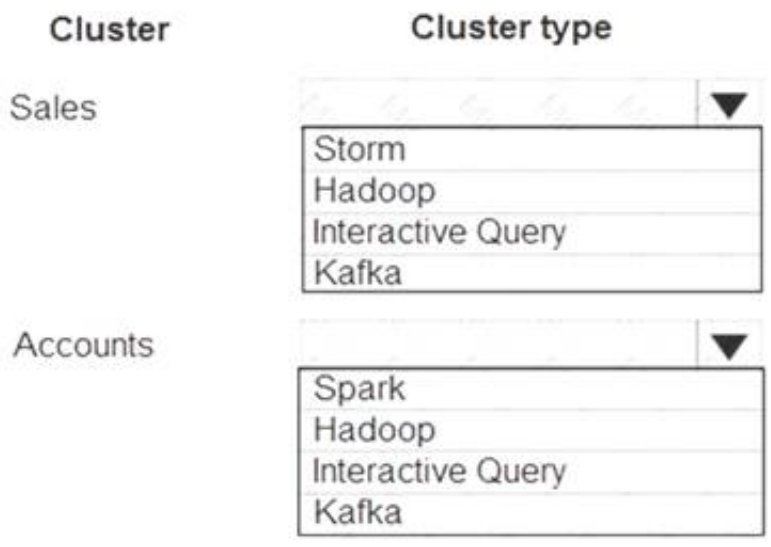


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Interactive Query
Choose Interactive Query cluster type to optimize for ad hoc, interactive queries. Box 2: Hadoop
Choose Apache Hadoop cluster type to optimize for Hive queries used as a batch process.
Note: In Azure HDInsight, there are several cluster types and technologies that can run Apache Hive queries. When you create your HDInsight cluster, choose the appropriate cluster type to help optimize performance for your workload needs.
For example, choose Interactive Query cluster type to optimize for ad hoc, interactive queries. Choose Apache Hadoop cluster type to optimize for Hive queries used as a batch process. Spark and HBase cluster types can also run Hive queries.
References:
https://docs.microsoft.com/bs-latn-ba/azure/hdinsight/hdinsight-hadoop-optimize-hive-query?toc=%2Fko-kr%2

**NEW QUESTION 50**
- (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.

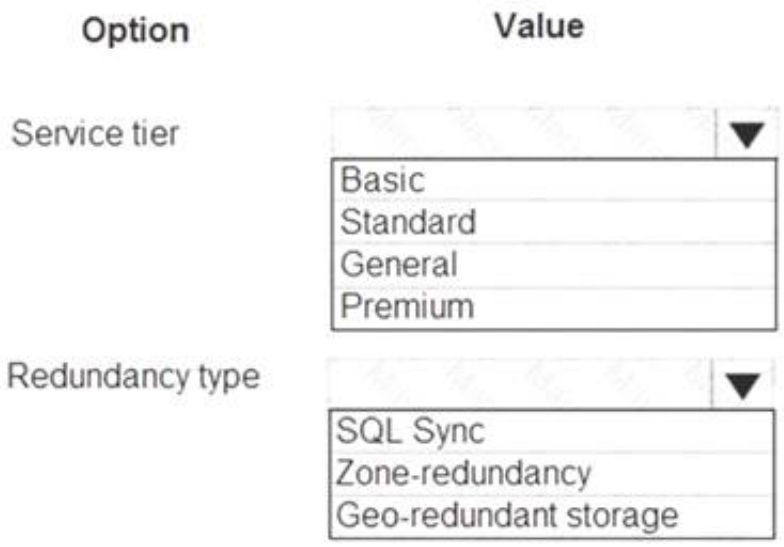


A. Mastered
B. Not Mastered

**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 51**
- (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 54**
- (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 59**
- (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 60**
- (Exam Topic 4)
You need to design the system for notifying law enforcement officers about speeding vehicles.
How should you design the pipeline? To answer, drag the appropriate services to the correct locations. Each service 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.

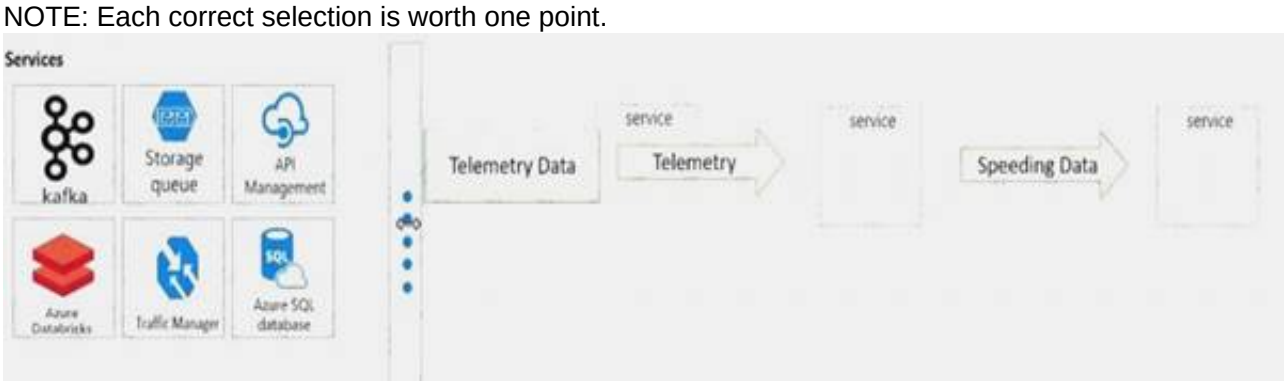


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**


**NEW QUESTION 61**
- (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 65**
.....

# 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