

workloads:

- A workload for data engineers who will use Python and SQL.
- A workload for jobs that will run notebooks that use Python, Spark, Scala, and SQL.
- A workload that data scientists will use to perform ad hoc analysis in Scala and R.

The enterprise architecture team at your company identifies the following standards for Databricks environments:

- The data engineers must share a cluster.
- The job cluster will be managed by using a request process whereby data scientists and data engineers provide packaged notebooks for deployment to the cluster.
- All the data scientists must be assigned their own cluster that terminates automatically after 120 minutes of inactivity. Currently, there are three data scientists.

You need to create the Databrick clusters for the workloads.

Solution: You create a Standard cluster for each data scientist, a High Concurrency cluster for the data engineers, and a Standard cluster for the jobs.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

We would need a High Concurrency cluster for the jobs. Note: Standard clusters are recommended for a single user. Standard can run workloads developed in any language: Python, R, Scala, and SQL. A high concurrency cluster is a managed cloud resource. The key benefits of high concurrency clusters are that they provide Apache Spark-native fine-grained sharing for maximum resource utilization and minimum query latencies.

<https://docs.azuredatabricks.net/clusters/configure.html>

NEW QUESTION 141

You plan to create an Azure Databricks workspace that has a tiered structure. The workspace will contain the following three workloads:

- A workload for data engineers who will use Python and SQL.
- A workload for jobs that will run notebooks that use Python, Spark, Scala, and SQL.
- A workload that data scientists will use to perform ad hoc analysis in Scala and R.

The enterprise architecture team at your company identifies the following standards for Databricks environments:

- The data engineers must share a cluster.
- The job cluster will be managed by using a request process whereby data scientists and data engineers provide packaged notebooks for deployment to the cluster.
- All the data scientists must be assigned their own cluster that terminates automatically after 120 minutes of inactivity. Currently, there are three data scientists.

You need to create the Databrick clusters for the workloads.

Solution: You create a Standard cluster for each data scientist, a High Concurrency cluster for the data engineers, and a High Concurrency cluster for the jobs.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

We need a High Concurrency cluster for the data engineers and the jobs. Note: Standard clusters are recommended for a single user. Standard can run workloads developed in any language: Python, R, Scala, and SQL. A high concurrency cluster is a managed cloud resource. The key benefits of high concurrency clusters are that they provide Apache Spark-native fine-grained sharing for maximum resource utilization and minimum query latencies.

<https://docs.azuredatabricks.net/clusters/configure.html>

NEW QUESTION 142

You have an Azure Stream Analytics query. The query returns a result set that contains 10,000 distinct values for a column named

clusterID. You monitor the Stream Analytics job and discover high latency. You need to reduce the latency. Which two actions should you perform? (Each correct answer presents a complete solution. Choose two.)

- A. Add a pass-through query.
- B. Add a temporal analytic function.
- C. Scale out the query by using PARTITION BY.
- D. Convert the query to a reference query.
- E. Increase the number of streaming units.

Answer: CE

Explanation:

C: Scaling a Stream Analytics job takes advantage of partitions in the input or output. Partitioning lets you divide data into subsets based on a partition key. A process that consumes the data (such as a Streaming Analytics job) can consume and write different partitions in parallel, which increases throughput.

E: Streaming Units (SUs) represents the computing resources that are allocated to execute a Stream Analytics job. The higher the number of SUs, the more CPU and memory resources are allocated for your job. This capacity lets you focus on the query logic and abstracts the need to manage the hardware to run your Stream Analytics job in a timely manner.

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-parallelization>

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-streaming-unit-consumption>

NEW QUESTION 143

A company uses Azure Data Lake Gen 1 Storage to store big data related to consumer behavior. You need to implement logging.

Solution: Create an Azure Automation runbook to copy events.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead configure Azure Data Lake Storage diagnostics to store logs and metrics in a storage account.

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-diagnostic-logs>

NEW QUESTION 144

You have an Azure data solution that contains an Azure SQL data warehouse named DW1. Several users execute adhoc queries to DW1 concurrently. You regularly perform automated data loads to DW1. You need to ensure that the automated data loads have enough memory available to complete quickly and successfully when the adhoc queries run. What should you do?

- A. Hash distribute the large fact tables in DW1 before performing the automated data loads.
- B. Assign a larger resource class to the automated data load queries.
- C. Create sampled statistics for every column in each table of DW1.
- D. Assign a smaller resource class to the automated data load queries.

Answer: B

Explanation:

To ensure the loading user has enough memory to achieve maximum compression rates, use loading users that are a member of a medium or large resource class.

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data>

NEW QUESTION 145

Drag and Drop

You deploy an Azure SQL database named DB1 to an Azure SQL server named SQL1. Currently, only the server admin has access to DB1. An Azure Active Directory (Azure AD) group named Analysts contains all the users who must have access to DB1. You have the following data security requirements:


- The Analysts group must have read-only access to all the views and tables in the Sales schema of DB1.
- A manager will decide who can access DB1. The manager will not interact directly with DB1.
- Users must not have to manage a separate password solely to access DB1.

Which four actions should you perform in sequence to meet the data security requirements? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Actions

Answer Area

- From the Azure portal, set the Active Directory admin for SQL1.
- From Microsoft SQL Server Management Studio (SSMS), set the Active Directory admin by using Transact-SQL.
- On DB1, grant the SELECT and VIEW DEFINITION permissions on the Sales schema to the Analysts group user.
- Create a login for each Azure AD user in the Analysts group by using Transact-SQL.
- From Microsoft SQL Server Management Studio (SSMS), sign in to SQL1 by using the account set as the Active Directory admin.
- On DB1, create a contained user for the Analysts group by using Transact-SQL.



www.pastleader.com


Answer:

Actions

Answer Area

- From Microsoft SQL Server Management Studio (SSMS), set the Active Directory admin by using Transact-SQL.
- Create a login for each Azure AD user in the Analysts group by using Transact-SQL.

- From the Azure portal, set the Active Directory admin for SQL1.
- On DB1, create a contained user for the Analysts group by using Transact-SQL.
- From Microsoft SQL Server Management Studio (SSMS), sign in to SQL1 by using the account set as the Active Directory admin.
- On DB1, grant the SELECT and VIEW DEFINITION permissions on the Sales schema to the Analysts group user.



www.pastleader.com

Explanation:

Step 1: From the Azure Portal, set the Active Directory admin for SQL1. Provision an Azure Active Directory administrator for your Azure SQL Database server. You can provision an Azure Active Directory administrator for your Azure SQL server in the Azure portal and by using PowerShell.

Step 2: On DB1, create a contained user for the Analysts group by using Transact-SQL Create contained database users in your database mapped to Azure AD identities. To create an Azure AD-based contained database user (other than the server administrator that owns the database), connect to the database with an Azure AD identity, as a user with at least the ALTER ANY USER

permission. Then use the following Transact-SQL syntax: CREATE USER <Azure_AD_principal_name> FROM EXTERNAL PROVIDER;

Step 3: From Microsoft SQL Server Management Studio (SSMS), sign in to SQL1 by using the account set as the Active Directory admin. Connect to the user database or data warehouse by using SSMS or SSDT To confirm the Azure AD administrator is properly set up, connect to the master database using the Azure AD administrator account. To provision an Azure AD-based contained database user (other than the server administrator that owns the database), connect to the database with an Azure AD identity that has access to the database.

Step 4: On DB1, grant the VIEW and SELECT DEFINITION.

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-aad-authentication-configure>

NEW QUESTION 146

Drag and Drop

You have an Azure subscription that contains an Azure Databricks environment and an Azure Storage account. You need to implement secure communication between Databricks and the storage account. You create an Azure key vault. Which four actions should you perform in sequence? (To answer, move the actions from the list of actions to the answer area and arrange them in the correct order.)

Actions	Answer Area
Mount the storage account.	
Add a secret scope to the Databricks environment.	⤵
Add a secret to the key vault.	⤴
Pause the Databricks cluster.	⤵
Retrieve an access key from the storage account.	

www.passleader.com

Answer:

Actions	Answer Area
	Mount the storage account.
	Retrieve an access key from the storage account.
	Add a secret to the key vault.
Pause the Databricks cluster.	Add a secret scope to the Databricks environment.

www.passleader.com

Explanation:

Managing secrets begins with creating a secret scope. To reference secrets stored in an Azure Key Vault, you can create a secret scope backed by Azure Key Vault.

<https://docs.microsoft.com/en-us/azure/azure-databricks/store-secrets-azure-key-vault>

NEW QUESTION 147

.....

Get the newest PassLeader DP-200 VCE dumps here: <https://www.passleader.com/dp-200.html> (148 Q&As Dumps)

And, DOWNLOAD the newest PassLeader DP-200 PDF dumps from Cloud Storage for free:

https://drive.google.com/open?id=1CTHwJ44u5IT4tsb2qo8oThaQ5c_ywun1