


Microsoft Data Certification DP-200 Exam



- **Vendor: Microsoft**
- **Exam Code: DP-200**
- **Exam Name: Implementing an Azure Data Solution**

Get Complete Version Exam DP-200 Dumps with VCE and PDF Here



<https://www.passleader.com/dp-200.html>

NEW QUESTION 137

You have an Azure Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB. You plan to copy the data from the storage account to an Azure SQL data warehouse. You need to prepare the files to ensure that the data copies quickly.

Solution: You modify the files to ensure that each row is less than 1 MB.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

When exporting data into an ORC File Format, you might get Java out-of-memory errors when there are large text columns. To work around this limitation, export only a subset of the columns. <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data>

NEW QUESTION 138

You have an Azure Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB. You plan to copy the data from the storage account to an Azure SQL data warehouse. You need to prepare the files to ensure that the data copies quickly.

Solution: You modify the files to ensure that each row is more than 1 MB.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead modify the files to ensure that each row is less than 1 MB.

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

NEW QUESTION 139

You plan to deploy an Azure Cosmos DB database that supports multi-master replication. You need to select a consistency level for the database to meet the following requirements:

- Provide a recovery point objective (RPO) of less than 15 minutes.
- Provide a recovery time objective (RTO) of zero minutes.

What are three possible consistency levels that you can select? (Each correct answer presents a complete solution. Choose three.)

- A. Strong
- B. Bounded Staleness
- C. Eventual
- D. Session
- E. Consistent Prefix

Answer: CDE

Explanation:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels-choosing>

NEW QUESTION 140

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 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.

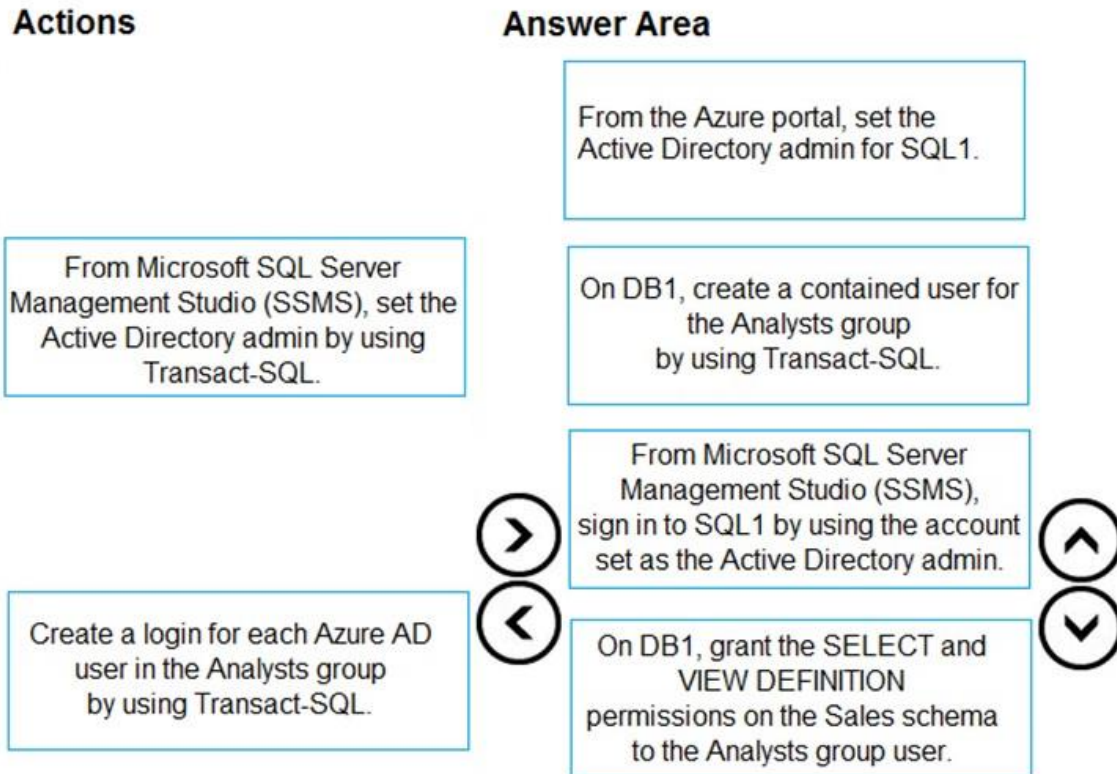
Create a login for each Azure AD user in the Analysts group by using Transact-SQL.

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.



Answer:



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
<input type="text" value="Mount the storage account."/>	
<input type="text" value="Add a secret scope to the Databricks environment."/>	⤵
<input type="text" value="Add a secret to the key vault."/>	⤴
<input type="text" value="Pause the Databricks cluster."/>	⤴
<input type="text" value="Retrieve an access key from the storage account."/>	⤴

Answer:

Actions	Answer Area
	<input type="text" value="Mount the storage account."/>
	<input type="text" value="Retrieve an access key from the storage account."/>
	<input type="text" value="Add a secret to the key vault."/>
<input type="text" value="Pause the Databricks cluster."/>	<input type="text" value="Add a secret scope to the Databricks environment."/>

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

.....

NEW QUESTION 198

You are migrating a corporate research analytical solution from an internal datacenter to Azure. 200 TB of research data is currently stored in an on-premises Hadoop cluster. You plan to copy it to Azure Storage. Your internal datacenter is connected to your Azure Virtual Network (VNet) with Express Route private peering. The Azure Storage service endpoint is accessible from the same VNet. Corporate policy dictates that the research data cannot be transferred over public internet. You need to securely migrate the research data online. What should you do?

- A. Transfer the data using Azure Data Box Disk devices.
- B. Transfer the data using Azure Data Factory in distributed copy (DistCopy) mode, with an Azure Data Factory self-hosted Integration Runtime (IR) machine installed in the on-premises datacenter.
- C. Transfer the data using Azure Data Factory in native Integration Runtime (IR) mode, with an Azure Data Factory self-hosted IR machine installed on the Azure VNet.
- D. Transfer the data using Azure Data Box Heavy devices.

Answer: C

Explanation:

You should transfer the data using Azure Data Factory in native IR mode, with an Azure Data Factory selfhosted IR machine installed on the Azure VNet. This approach supports data transfer via Express Route and uses Azure Data Factory IR as an engine to copy the data.

Incorrect:

Not A: You should not transfer the data using Azure Data Box Disk devices. This approach is for offline transfer scenarios. Also, an Azure Data Box Disk device has a total capacity of only 40 TB.

Not B: You should not transfer the data using Azure Data Factory in DistCopy mode, with an Azure Data Factory self-hosted IR machine installed in the on-premises datacenter. The DistCp tool does not support Express Route private peering with an Azure Storage VNet endpoint.

Not D: You should not transfer the data using Azure Data Box Heavy devices. While these devices have a capacity of 1 PB, they are intended for offline transfer scenarios only.

NEW QUESTION 199

You are a data engineer for your company. Your company has an on-premises SQL Server instance that contains 16 databases. Four of the databases require Common Language Runtime (CLR) features. You must be able to manage each database separately because each database has its own resource needs. You plan to migrate these databases to Azure. You want to migrate the databases by using a backup and restore process by using SQL commands. You need to choose the most appropriate deployment option to migrate the databases. What should you use?

- A. Azure SQL Database with an elastic pool.
- B. Azure Cosmos DB with the SQL (DocumentDB) API.
- C. Azure SQL Database managed instance.
- D. Azure Cosmos DB with the Table API.

Answer: C

Explanation:

You should use an Azure SQL Database managed instance deployment. This deployment option is almost 100% compatible with an on-premises instance, including the ability to use CLR features. When you back up the databases on-premises, you can execute a restore command to migrate the databases in Azure. This is referred to as lift and shift.

Incorrect:

Not A: You should not use an Azure SQL Database with an elastic pool deployment. An elastic pool allows you to deploy multiple databases to a single logical instance and have all databases share a pool of resources. You configure the resource usage up front by choosing a purchasing model. You cannot take advantage of CLR features with an elastic pool.

Not B: You should not use an Azure Cosmos DB with the SQL (DocumentDB) API deployment. Cosmos DB is a multimodel database that supports five APIs for storage and queries, including SQL, Table, Cassandra, Gremlin, and MongoDB. The SQL API allows you to access data by using SQL-like queries. You cannot restore SQL Server databases to Cosmos DB by using SQL commands.

Not D: You should not use an Azure Cosmos DB with the Table API deployment. The Table API is similar to Azure Tables. This deployment is useful if you are migrating an application from Azure Tables to Cosmos DB. With Azure Tables, you can access data by using Language Integrated Query (LINQ) and OData. You cannot restore SQL Server databases to Cosmos DB by using SQL commands.

NEW QUESTION 200

SIMULATION

.....

Answer:

Step 1: Create a new SQL database named DB3:

1. Select SQL in the left-hand menu of the Azure portal. If SQL is not in the list, select All services, then type SQL in the search box.
2. Select + Add to open the Select SQL deployment option page. Select Single Database. You can view additional information about the different databases by selecting Show details on the Databases tile.
3. Select Create.
4. Enter the required fields if necessary.
5. Leave the rest of the values as default and select Review + Create at the bottom of the form.
6. Review the final settings and select Create. Use DB3 as database name.

On the SQL Database form, select Create to deploy and provision the resource group, server, and database.

Step 2: Create your elastic pool using the Azure portal:

1. Select Azure SQL in the left-hand menu of the Azure portal. If Azure SQL is not in the list, select All services, then type Azure SQL in the search box.
2. Select + Add to open the Select SQL deployment option page.
3. Select Elastic pool from the Resource type drop-down in the SQL Databases tile. Select Create to create your elastic pool.
4. Configure your elastic pool with the following values:
 - Name: Provide a unique name for your elastic pool, such as myElasticPool.
 - Subscription: Select your subscription from the drop-down.
 - ResourceGroup: Select the resource group.
 - Server: Select the server.
5. Select Configure elastic pool.
6. On the Configure page, select the Databases tab, and then choose to Add database.
7. Add the Azure SQL database named DB2, and the new SQL database named DB3 that you created in Step 1.
8. Select Review + create to review your elastic pool settings and then select Create to create your elastic pool.

Explanation:

<https://docs.microsoft.com/bs-latn-ba/azure/sql-database/sql-database-elastic-pool-failover-group-tutorial>

NEW QUESTION 201 SIMULATION

.....

Answer:

Create a general-purpose v2 storage account, which provides access to all of the Azure Storage services: blobs, files, queues, tables, and disks:

1. On the Azure portal menu, select All services. In the list of resources, type Storage Accounts. As you begin typing, the list filters based on your input. Select Storage Accounts.
2. On the Storage Accounts window that appears, choose Add.
3. Select the subscription in which to create the storage account.
4. Under the Resource group field, select Create new. Enter the name for your new resource group, as shown in the following image.
5. Next, enter the name account10543936 for your storage account.
6. Select a location for your storage account, or use the default location.
7. Leave these fields set to their default values:
 - Deployment model: Resource Manager
 - Performance: Standard
 - Account kind: StorageV2 (general-purpose v2)
 - Replication: Read-access geo-redundant storage (RA-GRS)
 - Access tier: Hot
8. Select Review + Create to review your storage account settings and create the account.

9. Select Create.

Explanation:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-create>

NEW QUESTION 202

SIMULATION

.....

Answer:

You can enable Availability Zones by using Azure portal when creating an Azure Cosmos account.

You can enable Availability Zones by using Azure portal:

Step 1: enable the Geo-redundancy, Multi-region Writes:

1. In Azure Portal search for and select Azure Cosmos DB.
2. Locate the Cosmos DB database named cosmos10543936
3. Access the properties for cosmos10543936
4. Enable the Geo-redundancy, Multi-region Writes:
- Location: West US region

Step 2: Add region from your database account:

1. In to Azure portal, go to your Azure Cosmos account, and open the Replicate data globally menu.
2. To add regions, select the hexagons on the map with the + label that corresponds to your desired region(s). Alternatively, to add a region, select the + Add region option and choose a region from the drop-down menu:
- Add: West US region
3. To save your changes, select OK.

Explanation:

<https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability>

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-manage-database-account>

NEW QUESTION 203

SIMULATION

.....

Answer:

Provision an Azure Active Directory administrator for your managed instance. Each Azure SQL server (which hosts a SQL Database or SQL Data Warehouse) starts with a single server administrator account that is the administrator of the entire Azure SQL server. A second SQL Server administrator must be created, that is an Azure AD account. This principal is created as a contained database user in the master database:

1. In the Azure portal, in the upper-right corner, select your connection to drop down a list of possible Active Directories. Choose the correct Active Directory as the default Azure AD. This step links the subscription-associated Active Directory with Azure SQL server making sure that the same subscription is used for both Azure AD and SQL Server. (The Azure SQL server can be hosting either Azure SQL Database or Azure SQL Data Warehouse.)
2. Search for and select the SQL server SQL10543936.
3. In SQL Server page, select Active Directory admin.
4. In the Active Directory admin page, select Set admin.
5. In the Add admin page, search for user User1-10543936@ExamUsers.com, select it, and then select Select. (The Active Directory admin page shows all members and groups of your Active Directory. Users or groups that are grayed out cannot be selected because they are not supported as Azure AD administrators.)
6. At the top of the Active Directory admin page, select SAVE.

Explanation:

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

NEW QUESTION 204
SIMULATION

.....

Answer:

You can use Private Endpoints for your Azure Storage accounts to allow clients on a virtual network (VNet) to securely access data over a Private Link. Create your Private Endpoint:

1. On the upper-left side of the screen in the Azure portal, Storage > Storage account, and select your storage account storage10543936.
2. Select Networking.
3. Select Add Private Endpoint.
4. In Create Private Endpoint, enter or select this information:
- Virtual network: Select VNET1 from the resource group.
5. Select OK.
6. Select Review + create. You're taken to the Review + create page where Azure validates your configuration.

Explanation:

<https://docs.microsoft.com/en-us/azure/private-link/create-private-endpoint-storage-portal>

NEW QUESTION 205
SIMULATION

.....

Answer:

1. In the Azure portal, browse to the database db1-copy10543936 that you want to set up for geo-replication.
2. On the SQL database page, select geo-replication, and then select the region to create the secondary database: US West region.
3. Select or configure the server and pricing tier for the secondary database.
4. Click Create to add the secondary.
5. The secondary database is created and the seeding process begins.
6. When the seeding process is complete, the secondary database displays its status.

Explanation:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-active-geo-replication-portal>

NEW QUESTION 206
SIMULATION

.....

Answer:

Enable soft delete for blobs on your storage account by using Azure portal:

1. In the Azure portal, select your storage account.
2. Navigate to the Data Protection option under Blob Service.
3. Click Enabled under Blob soft delete.
4. Enter the number of days you want to retain for under Retention policies. Here enter 10.
5. Choose the Save button to confirm your Data Protection settings.

Note: Azure Storage now offers soft delete for blob objects so that you can more easily recover your data when it is erroneously modified or deleted by an application or other storage account user. Currently you can retain soft deleted data for between 1 and 365 days.

Explanation:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-soft-delete>

NEW QUESTION 207

You are monitoring an Azure Stream Analytics job. You discover that the Backlogged Input Events metric is increasing slowly and is consistently non-zero. You need to ensure that the job can handle

all the events. What should you do?

- A. Change the compatibility level of the Stream Analytics job.
- B. Increase the number of streaming units (SUs).
- C. Create an additional output stream for the existing input stream.
- D. Remove any named consumer groups from the connection and use \$default.

Answer: B

Explanation:

Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently non-zero, you should scale out your job. You should increase the Streaming Units. Note: 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.

<https://docs.microsoft.com/bs-cyrl-ba/azure/stream-analytics/stream-analytics-monitoring>

NEW QUESTION 208

SIMULATION

.....

Answer:

Set up auditing for your database. The following section describes the configuration of auditing using the Azure portal:

1. Go to the Azure portal.
2. Navigate to Auditing under the Security heading in your SQL database db2/server pane.
3. If you prefer to enable auditing on the database level, switch Auditing to ON.

Note: By default the audit database data retention period is set to 100 days.

Explanation:

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

NEW QUESTION 209

SIMULATION

.....

Answer:

You can configure your storage account to accept requests from secure connections only by setting the Secure transfer required property for the storage account. Require secure transfer for an existing storage account:

1. Select the existing storage account storage10543936 in the Azure portal.
2. In the storage account menu pane, under SETTINGS, select Configuration.
3. Under Secure transfer required, select Enabled.

Explanation:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-require-secure-transfer>

NEW QUESTION 210

SIMULATION

.....

Answer:

SQL Data Warehouse compute resources can be scaled by increasing or decreasing data warehouse units:

1. Click SQL data warehouses in the left page of the Azure portal.
2. Select datawarehouse from the SQL data warehouses page. The data warehouse opens.
3. Click Scale.

4. In the Scale panel, move the slider left or right to change the DWU setting. Double the DWU setting.

5. Click Save. A confirmation message appears. Click yes to confirm or no to cancel.

Explanation:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/quickstart-scale-compute-portal>

NEW QUESTION 211

.....

Get Complete Version Exam DP-200 Dumps with VCE and PDF Here



<https://www.passleader.com/dp-200.html>