

Google.ADP.v2026-09-11.q22

Exam Code:	ADP
Exam Name:	Associate Data Practitioner
Certification Provider:	Google
Free Question Number:	22
Version:	v2026-09-11
# of views:	105
# of Questions views:	220
https://www.exam-tests.com/ADP-exam/Google.ADP.v2026-09-11.q22.html	

NEW QUESTION: 1

Your company stores historical data in Cloud Storage. You need to ensure that all data is saved in a bucket for at least three years. What should you do?

- A. Set temporary object holds.
- B. Enable Object Versioning.
- C. Change the bucket storage class to Archive.
- D. Set a bucket retention policy.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 2

You need to design a data pipeline to process large volumes of raw server log data stored in Cloud Storage.

The data needs to be cleaned, transformed, and aggregated before being loaded into BigQuery for analysis.

The transformation involves complex data manipulation using Spark scripts that your team developed. You need to implement a solution that leverages your team's existing skillset, processes data at scale, and minimizes cost. What should you do?

- A. Use Cloud Data Fusion to visually design and manage the pipeline.
- B. Use Dataproc to run the transformations on a cluster.
- C. Use Dataform to define the transformations in SQLX.
- D. Use Dataflow with a custom template for the transformation logic.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 3

Your company uses Looker as its primary business intelligence platform. You want to use LookML to visualize the profit margin for each of your company's products in your Looker Explores and dashboards.

You need to implement a solution quickly and efficiently. What should you do?

- A.** Apply a filter to only show products with a positive profit margin.
- B.** Create a new dimension that categorizes products based on their profit margin ranges (e.g., high, medium, low).
- C.** Define a new measure that calculates the profit margin by using the existing revenue and cost fields.
- D.** Create a derived table that pre-calculates the profit margin for each product, and include it in the Looker model.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 4

Your data science team needs to collaboratively analyze a 25 TB BigQuery dataset to support the development of a machine learning model. You want to use Colab Enterprise notebooks while ensuring efficient data access and minimizing cost. What should you do?

- A.** Copy the BigQuery dataset to the local storage of the Colab Enterprise runtime, and analyze the data using Pandas.
- B.** Create a Dataproc cluster connected to a Colab Enterprise notebook, and use Spark to process the data in BigQuery.
- C.** Use BigQuery magic commands within a Colab Enterprise notebook to query and analyze the data.
- D.** Export the BigQuery dataset to Google Drive. Load the dataset into the Colab Enterprise notebook using Pandas.

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 5

You are migrating data from a legacy on-premises MySQL database to Google Cloud. The database contains various tables with different data types and sizes, including large tables with millions of rows and transactional data. You need to migrate this data while maintaining data integrity, and minimizing downtime and cost. What should you do?

- A.** Export the MySQL database to CSV files, transfer the files to Cloud Storage by using Storage Transfer Service, and load the files into a Cloud SQL for MySQL instance.
- B.** Use Cloud Data Fusion to migrate the MySQL database to MySQL on Compute Engine.
- C.** Use Database Migration Service to replicate the MySQL database to a Cloud SQL for MySQL instance.
- D.** Set up a Cloud Composer environment to orchestrate a custom data pipeline. Use a Python script to extract data from the MySQL database and load it to MySQL on Compute Engine.

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 6

You have a Dataproc cluster that performs batch processing on data stored in Cloud Storage. You need to schedule a daily Spark job to generate a report that will be emailed to stakeholders. You need a fully-managed solution that is easy to implement and minimizes complexity. What should you do?

- A.** Use Cloud Run functions to trigger the Spark job and email the report.
- B.** Use Dataproc workflow templates to define and schedule the Spark job, and to email the report.
- C.** Use Cloud Composer to orchestrate the Spark job and email the report.
- D.** Use Cloud Scheduler to trigger the Spark job. and use Cloud Run functions to email the report.

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 7

You manage data at an ecommerce company. You have a Dataflow pipeline that processes order data from Pub/Sub, enriches the data with product information from Bigtable, and writes the processed data to BigQuery for analysis. The pipeline runs continuously and processes thousands of orders every minute. You need to monitor the pipeline's performance and be alerted if errors occur. What should you do?

- A.** Use the Dataflow job monitoring interface to visually inspect the pipeline graph, check for errors, and configure notifications when critical errors occur.
- B.** Use Cloud Logging to view the pipeline logs and check for errors. Set up alerts based on specific keywords in the logs.
- C.** Use BigQuery to analyze the processed data in Cloud Storage and identify anomalies or inconsistencies. Set up scheduled alerts based when anomalies or inconsistencies occur.
- D.** Use Cloud Monitoring to track key metrics. Create alerting policies in Cloud Monitoring to trigger notifications when metrics exceed thresholds or when errors occur.

Answer: **D** [\(LEAVE A REPLY\)](#)

NEW QUESTION: 8

Your organization has several datasets in BigQuery. The datasets need to be shared with your external partners so that they can run SQL queries without needing to copy the data to their own projects. You have organized each partner's data in its own BigQuery dataset. Each partner should be able to access only their data. You want to share the data while following Google-recommended practices. What should you do?

- A.** Use Analytics Hub to create a listing on a private data exchange for each partner dataset. Allow each partner to subscribe to their respective listings.
- B.** Grant the partners the bigquery.user IAM role on the BigQuery project.
- C.** Export the BigQuery data to a Cloud Storage bucket. Grant the partners the storage.objectUser IAM role on the bucket.

D. Create a Dataflow job that reads from each BigQuery dataset and pushes the data into a dedicated Pub /Sub topic for each partner. Grant each partner the pubsub. subscriber IAM role.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 9

Your team is building several data pipelines that contain a collection of complex tasks and dependencies that you want to execute on a schedule, in a specific order. The tasks and dependencies consist of files in Cloud Storage, Apache Spark jobs, and data in BigQuery. You need to design a system that can schedule and automate these data processing tasks using a fully managed approach. What should you do?

A. Use Cloud Tasks to schedule and run the jobs asynchronously.

B. Use Cloud Scheduler to schedule the jobs to run.

C. Create directed acyclic graphs (DAGS) in Cloud Composer. Use the appropriate operators to connect to Cloud Storage, Spark, and BigQuery.

D. Create directed acyclic graphs (DAGS) in Apache Airflow deployed on Google Kubernetes Engine. Use the appropriate operators to connect to Cloud Storage, Spark, and BigQuery.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 10

You are developing a data ingestion pipeline to load small CSV files into BigQuery from Cloud Storage. You want to load these files upon arrival to minimize data latency. You want to accomplish this with minimal cost and maintenance. What should you do?

A. Create a Cloud Composer pipeline to load new files from Cloud Storage to BigQuery and schedule it to run every 10 minutes.

B. Use the bq command-line tool within a Cloud Shell instance to load the data into BigQuery.

C. Create a Dataproc cluster to pull CSV files from Cloud Storage, process them using Spark, and write the results to BigQuery.

D. Create a Cloud Run function to load the data into BigQuery that is triggered when data arrives in Cloud Storage.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 11

You have a BigQuery dataset containing sales data

a. This data is actively queried for the first 6 months. After that, the data is not queried but needs to be retained for 3 years for compliance reasons. You need to implement a data management strategy that meets access and compliance requirements, while keeping cost and administrative overhead to a minimum. What should you do?

- A.** Partition a BigQuery table by month. After 6 months, export the data to Coldline storage. Implement a lifecycle policy to delete the data from Cloud Storage after 3 years.
- B.** Store all data in a single BigQuery table without partitioning or lifecycle policies.
- C.** Set up a scheduled query to export the data to Cloud Storage after 6 months. Write a stored procedure to delete the data from BigQuery after 3 years.
- D.** Use BigQuery long-term storage for the entire dataset. Set up a Cloud Run function to delete the data from BigQuery after 3 years.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 12

Your company is migrating their batch transformation pipelines to Google Cloud. You need to choose a solution that supports programmatic transformations using only SQL. You also want the technology to support Git integration for version control of your pipelines. What should you do?

- A.** Use Dataform workflows.
- B.** Use Dataflow pipelines.
- C.** Use Cloud Composer operators.
- D.** Use Cloud Data Fusion pipelines.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 13

Your company is building a near real-time streaming pipeline to process JSON telemetry data from small appliances. You need to process messages arriving at a Pub/Sub topic, capitalize letters in the serial number field, and write results to BigQuery. You want to use a managed service and write a minimal amount of code for underlying transformations. What should you do?

- A.** Use the "Pub/Sub to BigQuery" Dataflow template with a UDF, and write the results to BigQuery.
- B.** Use a Pub/Sub to Cloud Storage subscription, write a Cloud Run service that is triggered when objects arrive in the bucket, performs the transformations, and writes the results to BigQuery.
- C.** Use a Pub/Sub to BigQuery subscription, write results directly to BigQuery, and schedule a transformation query to run every five minutes.
- D.** Use a Pub/Sub push subscription, write a Cloud Run service that accepts the messages, performs the transformations, and writes the results to BigQuery.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 14

You are storing data in Cloud Storage for a machine learning project. The data is frequently accessed during the model training phase, minimally accessed after 30 days,

and unlikely to be accessed after 90 days. You need to choose the appropriate storage class for the different stages of the project to minimize cost. What should you do?

- A.** Store the data in Nearline storage during the model training phase. Transition the data to Archive storage 30 days after model deployment, and to Coldline storage 90 days after model deployment.
- B.** Store the data in Standard storage during the model training phase. Transition the data to Durable Reduced Availability (DRA) storage 30 days after model deployment, and to Coldline storage 90 days after model deployment.
- C.** Store the data in Standard storage during the model training phase. Transition the data to Nearline storage 30 days after model deployment, and to Coldline storage 90 days after model deployment.
- D.** Store the data in Nearline storage during the model training phase. Transition the data to Coldline storage 30 days after model deployment, and to Archive storage 90 days after model deployment.

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 15

Your company currently uses an on-premises network file system (NFS) and is migrating data to Google Cloud. You want to be able to control how much bandwidth is used by the data migration while capturing detailed reporting on the migration status. What should you do?

- A.** Use Cloud Storage FUSE.
- B.** Use Storage Transfer Service.
- C.** Use a Transfer Appliance.
- D.** Use gcloud storage commands.

Answer: **B** [\(LEAVE A REPLY\)](#)

NEW QUESTION: 16

Your organization is conducting analysis on regional sales metrics. Data from each regional sales team is stored as separate tables in BigQuery and updated monthly. You need to create a solution that identifies the top three regions with the highest monthly sales for the next three months. You want the solution to automatically provide up-to-date results. What should you do?

- A.** Create a BigQuery table that performs a union across all of the regional sales tables. Use the `row_number()` window function to query the new table.
- B.** Create a BigQuery materialized view that performs a union across all of the regional sales tables. Use the `rank()` window function to query the new materialized view.
- C.** Create a BigQuery materialized view that performs a cross join across all of the regional sales tables. Use the `row_number()` window function to query the new materialized view.
- D.** Create a BigQuery table that performs a cross join across all of the regional sales tables. Use the `rank()` window function to query the new table.

Answer: ([**SHOW ANSWER**](#)**)**

Valid ADP Dumps shared by BraindumpsPass.com for Helping Passing ADP Exam! BraindumpsPass.com now offer the **newest ADP exam dumps**, the BraindumpsPass.com ADP exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com ADP dumps with Test Engine here: <https://www.braindumpsPass.com/Google/ADP-practice-exam-dumps.html> (66 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 17

You need to design a data pipeline that ingests data from CSV, Avro, and Parquet files into Cloud Storage.

The data includes raw user input. You need to remove all malicious SQL injections before storing the data in BigQuery. Which data manipulation methodology should you choose?

- A. ETLT
- B. ELT
- C. ETL
- D. EL

Answer: C ([**LEAVE A REPLY**](#)**)**

NEW QUESTION: 18

Your company uses Looker to visualize and analyze sales data

a. You need to create a dashboard that displays sales metrics, such as sales by region, product category, and time period. Each metric relies on its own set of attributes distributed across several tables. You need to provide users the ability to filter the data by specific sales representatives and view individual transactions. You want to follow the Google-recommended approach. What should you do?

- A. Use BigQuery to create multiple materialized views, each focusing on a specific sales metric. Build the dashboard using these views.
- B. Create a single Explore with all sales metrics. Build the dashboard using this Explore.
- C. Use Looker's custom visualization capabilities to create a single visualization that displays all the sales metrics with filtering and drill-down functionality.
- D. Create multiple Explores, each focusing on each sales metric. Link the Explores together in a dashboard using drill-down functionality.

Answer: B ([**LEAVE A REPLY**](#)**)**

NEW QUESTION: 19

You created a customer support application that sends several forms of data to Google Cloud. Your application is sending:

1. Audio files from phone interactions with support agents that will be accessed during trainings.
2. CSV files of users' personally identifiable information (PII) that will be analyzed with SQL.
3. A large volume of small document files that will power other applications.

You need to select the appropriate tool for each data type given the required use case, while following Google- recommended practices. Which should you choose?

- A. Cloud Storage BigQuery Firestore
- B. Filestore Cloud SQL for PostgreSQL Datastore
- C. Filestore Bigtable BigQuery
- D. Cloud Storage CloudSQL for PostgreSQL Bigtable

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 20

You created a curated dataset of market trends in BigQuery that you want to share with multiple external partners. You want to control the rows and columns that each partner has access to. You want to follow Google-recommended practices. What should you do?

- A. Create a separate Cloud Storage bucket for each partner. Export the dataset to each bucket and assign each partner to their respective bucket. Grant bucket-level access by using IAM roles.
- B. Grant each partner read access to the BigQuery dataset by using IAM roles.
- C. Create a separate project for each partner and copy the dataset into each project. Publish each dataset in Analytics Hub. Grant dataset-level access to each partner by using subscriptions.
- D. Publish the dataset in Analytics Hub. Grant dataset-level access to each partner by using subscriptions.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 21

You are designing a BigQuery data warehouse with a team of experienced SQL developers. You need to recommend a cost- effective, fully-managed, serverless solution to build ELT processes with SQL pipelines.

Your solution must include source code control, environment parameterization, and data quality checks. What should you do?

- A. Use Dataform to build, orchestrate, and monitor the pipelines.
- B. Use Cloud Composer to orchestrate and run data workflows.
- C. Use Dataproc to run MapReduce jobs for distributed data processing.
- D. Use Cloud Data Fusion to visually design and manage the pipelines.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 22

Your organization plans to move their on-premises environment to Google Cloud. Your organization's network bandwidth is less than 1 Gbps. You need to move over 500 ## of data to Cloud Storage securely, and only have a few days to move the data. What should you do?

- A.** Request multiple Transfer Appliances, copy the data to the appliances, and ship the appliances back to Google Cloud to upload the data to Cloud Storage.
- B.** Connect to Google Cloud using VPN. Use the gcloud storage command to move the data to Cloud Storage.
- C.** Connect to Google Cloud using Dedicated Interconnect. Use the gcloud storage command to move the data to Cloud Storage.
- D.** Connect to Google Cloud using VPN. Use Storage Transfer Service to move the data to Cloud Storage.

Answer: ([SHOW ANSWER](#))

Valid ADP Dumps shared by BraindumpsPass.com for Helping Passing ADP Exam! BraindumpsPass.com now offer the **newest ADP exam dumps**, the BraindumpsPass.com ADP exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com ADP dumps with Test Engine here: <https://www.braindumpspass.com/Google/ADP-practice-exam-dumps.html> (66 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)