

## [28/July/2018 Updated PassLeader Free 80q Professional Data Engineer VCE Dumps Cover All New Questions (Part A)]

New Updated Professional Data Engineer Exam Questions from PassLeader Professional Data Engineer PDF dumps! Welcome to download the newest PassLeader Professional Data Engineer VCE

dumps: <https://www.passleader.com/professional-data-engineer.html> (80 Q&As) Keywords: Professional Data Engineer exam dumps, Professional Data Engineer exam questions, Professional Data Engineer VCE dumps, Professional Data Engineer PDF dumps, Professional Data Engineer practice tests, Professional Data Engineer study guide, Professional Data Engineer braindumps, **Google Cloud Certifications: Professional Data Engineer** Exam P.S. New Professional Data Engineer dumps PDF:

<https://drive.google.com/open?id=1m882ngsiRO1BOHineV4lQUv9jgF5Lpue> P.S. New Professional Cloud Architect dumps PDF: <https://drive.google.com/open?id=19jt3GbCmVz-pmGbZv8zjAu0NH7423IQ2> NEW QUESTION 1 Suppose you have a table that includes a nested column called "city" inside a column called "person", but when you try to submit the following query in BigQuery, it gives you an error: `SELECT person FROM `project1.example.table1` WHERE city = "London"` How would you correct the error? A. Add `", UNNEST(person)"` before the WHERE clause. B. Change "person" to "person.city". C. Change "person" to "city.person". D. Add `", UNNEST(city)"` before the WHERE clause. Answer: A Explanation: To access the person.city column, you need to "UNNEST(person)" and JOIN it to table1 using a comma.

[https://cloud.google.com/bigquery/docs/reference/standard-sql/migrating-from-legacy-sql#nested\\_repeated\\_results](https://cloud.google.com/bigquery/docs/reference/standard-sql/migrating-from-legacy-sql#nested_repeated_results) NEW QUESTION 2 What are two of the benefits of using denormalized data structures in BigQuery? A. Reduces the amount of data processed, reduces the amount of storage required. B. Increases query speed, makes queries simpler. C. Reduces the amount of storage required, increases query speed. D. Reduces the amount of data processed, increases query speed. Answer: B Explanation: Denormalization increases query speed for tables with billions of rows because BigQuery's performance degrades when doing JOINS on large tables, but with a denormalized data structure, you don't have to use JOINS, since all of the data has been combined into one table. Denormalization also makes queries simpler because you do not have to use JOIN clauses. Denormalization increases the amount of data processed and the amount of storage required because it creates redundant data.

[https://cloud.google.com/solutions/bigquery-data-warehouse#denormalizing\\_data](https://cloud.google.com/solutions/bigquery-data-warehouse#denormalizing_data) NEW QUESTION 3 Which of these statements about exporting data from BigQuery is false? A. To export more than 1 GB of data, you need to put a wildcard in the destination filename. B. The only supported export destination is Google Cloud Storage. C. Data can only be exported in JSON or Avro format. D. The only compression option available is GZIP. Answer: C Explanation: Data can be exported in CSV, JSON, or Avro format. If you are exporting nested or repeated data, then CSV format is not supported. <https://cloud.google.com/bigquery/docs/exporting-data> NEW QUESTION 4 What are all of the BigQuery operations that Google charges for? A. Storage, queries, and streaming inserts. B. Storage, queries, and loading data from a file. C. Storage, queries, and exporting data. D. Queries and streaming inserts. Answer: A Explanation: Google charges for storage, queries, and streaming inserts. Loading data from a file and exporting data are free operations.

<https://cloud.google.com/bigquery/pricing> NEW QUESTION 5 Which of the following is not possible using primitive roles? A. Give a user viewer access to BigQuery and owner access to Google Compute Engine instances. B. Give UserA owner access and UserB editor access for all datasets in a project. C. Give a user access to view all datasets in a project, but not run queries on them. D. Give GroupA owner access and GroupB editor access for all datasets in a project. Answer: C Explanation: Primitive roles can be used to give owner, editor, or viewer access to a user or group, but they can't be used to separate data access permissions from job-running permissions. [https://cloud.google.com/bigquery/docs/access-control#primitive\\_iam\\_roles](https://cloud.google.com/bigquery/docs/access-control#primitive_iam_roles) NEW QUESTION 6 Which of these statements about BigQuery caching is true? A. By default, a query's results are not cached. B. BigQuery caches query results for 48 hours. C. Query results are cached even if you specify a destination table. D. There is no charge for a query that retrieves its results from cache. Answer: D Explanation: When query results are retrieved from a cached results table, you are not charged for the query. BigQuery caches query results for 24 hours, not 48 hours. Query results are not cached if you specify a destination table. A query's results are always cached except under certain conditions, such as if you specify a destination table. <https://cloud.google.com/bigquery/querying-data#query-caching> NEW QUESTION 7 Which of these sources can you not load data into BigQuery from? A. File upload

B. Google Drive C. Google Cloud Storage D. Google Cloud SQL Answer: D Explanation: You can load data into BigQuery from a file upload, Google Cloud Storage, Google Drive, or Google Cloud Bigtable. It is not possible to load data into BigQuery directly from Google Cloud SQL. One way to get data from Cloud SQL to BigQuery would be to export data from Cloud SQL to Cloud Storage and then load it from there.

<https://cloud.google.com/bigquery/loading-data> NEW QUESTION 8 Which of the following statements about Legacy SQL and Standard SQL is not true? A. Standard SQL is the preferred query language for BigQuery.

B. If you write a query in Legacy SQL, it might generate an error if you try to run it with Standard SQL.

C. One difference between the two query languages is how you specify fully-qualified table names (i.e. table names that include their associated project name). D. You need to set a query language for each dataset and the default is Standard SQL. Answer: D Explanation: You do not set a query language for each dataset. It is set each time you run a query and the default query language is Legacy SQL. Standard SQL has been the preferred query language since BigQuery 2.0 was released. In legacy SQL, to query a table with a project-qualified name, you use a colon (:), as a separator. In standard SQL, you use a period, instead. Due to the differences in syntax between the two query languages (such as with project-qualified table names), if you write a query in Legacy SQL, it might generate an error if you try to run it with Standard SQL.

<https://cloud.google.com/bigquery/docs/reference/standard-sql/migrating-from-legacy-sql> NEW QUESTION 9 How would you query specific partitions in a BigQuery table? A. Use the DAY column in the WHERE clause

B. Use the EXTRACT(DAY) clause C. Use the PARTITIONTIME pseudo-column in the WHERE clause D. Use DATE BETWEEN in the WHERE clause Answer: C Explanation:

Partitioned tables include a pseudo column named \_PARTITIONTIME that contains a date-based timestamp for data loaded into the table. To limit a query to particular partitions (such as Jan 1st and 2nd of 2017), use a clause similar to this: WHERE \_PARTITIONTIME BETWEEN TIMESTAMP('2017-01-01') AND TIMESTAMP('2017-01-02')

[https://cloud.google.com/bigquery/docs/partitioned-tables#the\\_partitiontime\\_pseudo\\_column](https://cloud.google.com/bigquery/docs/partitioned-tables#the_partitiontime_pseudo_column) NEW QUESTION 10 Which SQL keyword can be used to reduce the number of columns processed by BigQuery? A. BETWEEN

B. WHERE C. SELECT D. LIMIT Answer: C

Explanation: SELECT allows you to query specific columns rather than the whole table. LIMIT, BETWEEN, and WHERE clauses will not reduce the number of columns processed by BigQuery.

[https://cloud.google.com/bigquery/launch-checklist#architecture\\_design\\_and\\_development\\_checklist](https://cloud.google.com/bigquery/launch-checklist#architecture_design_and_development_checklist) NEW QUESTION 11 To give a user read permission for only the first three columns of a table, which access control method would you use?

A. Primitive role. B. Predefined role. C. Authorized view.

D. It's not possible to give access to only the first three columns of a table. Answer: C Explanation: An authorized view allows you to share query results with particular users and groups without giving them read access to the underlying tables. Authorized views can only be created in a dataset that does not contain the tables queried by the view. When you create an authorized view, you use the view's SQL query to restrict access to only the rows and columns you want the users to see.

<https://cloud.google.com/bigquery/docs/views#authorized-views> NEW QUESTION 12 What are two methods that can be used to denormalize tables in BigQuery? A. 1. Split table into multiple tables 2. Use a partitioned table

B. 1. Join tables into one table 2. Use nested repeated fields C. 1. Use a partitioned table 2. Join tables into one table D. 1. Use nested repeated fields 2. Use a partitioned table Answer: B

Explanation: The conventional method of denormalizing data involves simply writing a fact, along with all its dimensions, into a flat table structure. For example, if you are dealing with sales transactions, you would write each individual fact to a record, along with the accompanying dimensions such as order and customer information. The other method for denormalizing data takes advantage of BigQuery's native support for nested and repeated structures in JSON or Avro input data. Expressing records using nested and repeated structures can provide a more natural representation of the underlying data. In the case of the sales order, the outer part of a JSON structure would contain the order and customer information, and the inner part of the structure would contain the individual line items of the order, which would be represented as nested, repeated elements.

[https://cloud.google.com/solutions/bigquery-data-warehouse#denormalizing\\_data](https://cloud.google.com/solutions/bigquery-data-warehouse#denormalizing_data) NEW QUESTION 13 Which of these is not a supported method of putting data into a partitioned table? A. If you have existing data in a separate file for

each day, then create a partitioned table and upload each file into the appropriate partition. B. Run a query to get the records for a specific day from an existing table and for the destination table, specify a partitioned table ending with the day in the format "\$YYYYMMDD". C. Create a partitioned table and stream new records to it every day.

D. Use ORDER BY to put a table's rows into chronological order and then change the table's type to

"Partitioned". Answer: D Explanation: You cannot change an existing table into a partitioned table. You must create a partitioned table from scratch. Then you can either stream data into it every day and the data will automatically be put in the right partition, or you can load data into a specific partition by using "\$YYYYMMDD" at the end of the table name.

<https://cloud.google.com/bigquery/docs/partitioned-tables> NEW QUESTION 14 Which of these operations can you perform from the BigQuery Web UI? A.&#160;&#160;&#160; Upload a file in SQL format. B.&#160;&#160;&#160; Load data with nested and repeated fields. C.&#160;&#160;&#160; Upload a 20 MB file. D.&#160;&#160;&#160; Upload multiple files using a wildcard. Answer: B Explanation: You can load data with nested and repeated fields using the Web UI. You cannot use the Web UI to: - Upload a file greater than 10 MB in size - Upload multiple files at the same time - Upload a file in SQL format All three of the above operations can be performed using the "bq" command. <https://cloud.google.com/bigquery/loading-data> NEW

QUESTION 15 Which methods can be used to reduce the number of rows processed by BigQuery? A.&#160;&#160;&#160; Splitting tables into multiple tables; putting data in partitions. B.&#160;&#160;&#160; Splitting tables into multiple tables; putting data in partitions; using the LIMIT clause. C.&#160;&#160;&#160; Putting data in partitions; using the LIMIT clause. D.&#160;&#160;&#160; Splitting tables into multiple tables; using the LIMIT clause. Answer: A Explanation: If you split a

table into multiple tables (such as one table for each day), then you can limit your query to the data in specific tables (such as for particular days). A better method is to use a partitioned table, as long as your data can be separated by the day. If you use the LIMIT clause, BigQuery will still process the entire table. <https://cloud.google.com/bigquery/docs/partitioned-tables> Download the

newest PassLeader Professional Data Engineer dumps from passleader.com now! 100% Pass Guarantee! Professional Data Engineer PDF dumps & Professional Data Engineer VCE

dumps:&#160;<https://www.passleader.com/professional-data-engineer.html> (80 Q&As) (New Questions Are 100% Available and Wrong Answers Have Been Corrected! Free VCE simulator!) P.S. New Professional Data Engineer dumps PDF:

<https://drive.google.com/open?id=1m882ngsiRO1BOHineV4IQUv9jgF5Lpue> P.S. New Professional Cloud Architect dumps PDF:

<https://drive.google.com/open?id=19jt3GbCmVz-pmGbZv8zjAu0NH7423IQ2>