Snowflake ARA-C01 Latest Study Notes | ARA-C01 Reliable Test Prep

What's more, part of that Actual4Labs ARA-C01 dumps now are free: https://drive.google.com/open?id=1QFv4L_9UKzjpOowEZWrJtX-MmOEX5H8O

Actual4Labs is benefiting more and more candidates for our excellent ARA-C01 exam torrent which is compiled by the professional experts accurately and skillfully. We are called the best friend on the way with our customers to help pass their ARA-C01 exam and help achieve their dreaming certification. The reason is that we not only provide our customers with valid and Reliable ARA-C01 Exam Materials, but also offer best service online since we uphold the professional ethical. So you can feel relax to have our ARA-C01 exam guide for we are a company with credibility.

Snowflake ARA-C01 Certification Exam has a reputation for being one of the most challenging Snowflake certification exams. It requires a deep understanding of Snowflake architecture and features, as well as a strong ability to apply that knowledge to complex real-world scenarios. However, earning this certification can be a significant boost to a professional's career, as it demonstrates their expertise in architecting Snowflake solutions at an advanced level.

Snowflake ARA-C01 Certification Exam is computer-based and consists of 60 multiple-choice questions. ARA-C01 exam is timed, and candidates have 90 minutes to complete it. To pass the exam, candidates must score at least 70%. ARA-C01 exam is administered by Pearson VUE and can be taken at any of their authorized testing centers.

>> Snowflake ARA-C01 Latest Study Notes <<

Pass Guaranteed 2026 ARA-C01: SnowPro Advanced Architect Certification Unparalleled Latest Study Notes

In the same way, IE, Firefox, Opera and Safari, and all the major browsers support the web-based Snowflake ARA-C01 practice test. So it requires no special plugins. The web-based SnowPro Advanced Architect Certification (ARA-C01) practice exam software is genuine, authentic, and real so feel free to start your practice instantly with SnowPro Advanced Architect Certification (ARA-C01) practice test.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q140-Q145):

NEW QUESTION # 140
You are creating a TASK to query a table streams created on the raw table and insert subsets of rows into multiple tables. You are following the below steps, but when you reached the step to resume the task, you received an error message as below.
Why is this error thrown and who can give you the required privilege?

Steps to be followed to get this error
-- Create a landing table to store raw JSON data.

-- Snowpipe could load data into this table. create or replace table raw (var variant);
-- Create a stream to capture inserts to the landing table.
-- A task will consume a set of columns from this stream. create or replace stream rawstream1 on table raw;
-- Create a second stream to capture inserts to the landing table.
-- A second task will consume another set of columns from this stream. create or replace stream rawstream2 on table raw;
-- Create a table that stores the names of office visitors identified in the raw data. create or replace table names (id int, first_name string, last_name string);
-- Create a table that stores the visitation dates of office visitors identified in the raw data.
create or replace table visits (id int, dt date);
-- Create a task that inserts new name records from the rawstream1 stream into the names table
-- every minute when the stream contains records.
-- Replace the 'mywh' warehouse with a warehouse that your role has USAGE privilege on. create or replace task raw_to_names
warehouse = etl_wh
schedule = '1 minute'
when
system$stream_has_data('rawstream1')
as
merge into names n
using (select var:id id, var:fname fname, var:lname lname from rawstream1) r1 on n.id = to_number(r1.id)
when matched then update set n.first_name = r1.fname, n.last_name = r1.lname
when not matched then insert (id, first_name, last_name) values (r1.id, r1.fname, r1.lname)
;
-- Create another task that merges visitation records from the rawstream1 stream into the visits table
-- every minute when the stream contains records.
-- Records with new IDs are inserted into the visits table;
-- Records with IDs that exist in the visits table update the DT column in the table.
-- Replace the 'mywh' warehouse with a warehouse that your role has USAGE privilege on. create or replace task raw_to_visits
warehouse = etl_wh schedule = '1 minute' when
system$stream_has_data('rawstream2') as
merge into visits v
using (select var:id id, var:visit_dt visit_dt from rawstream2) r2 on v.id = to_number(r2.id) when matched then update set v.dt = r2.visit_dt
when not matched then insert (id, dt) values (r2.id, r2.visit_dt);
-- Resume both tasks.
alter task raw_to_names resume;

- A. The role used to resume the task does not have EXECUTE TASK privilege. Only ACCOUNTADMIN can provide that privilege to the role.
- B. The role used to resume the task does not have EXECUTE TASK privilege. Only TASK OWNER can provide that privilege to the role.
- C. The role used to resume the task does not have EXECUTE TASK privilege. Both SECURITYADMIN and ACCOUNTADMIN can provide that privilege to the role.

Answer: A

NEW QUESTION # 141
Which feature provides the capability to define an alternate cluster key for a table with an existing cluster key?

- A. Search optimization
- B. Materialized view
- C. Result cache
- D. External table

Answer: B

Explanation:
A materialized view is a feature that provides the capability to define an alternate cluster key for a table with an existing cluster key. A materialized view is a pre-computed result set that is stored in Snowflake and can be queried like a regular table. A materialized view can have a different cluster key than the base table, which can improve the performance and efficiency of queries on the materialized view. A materialized view can also support aggregations, joins, and filters on the base table data. A materialized view is

automatically refreshed when the underlying data in the base table changes, as long as the AUTO_REFRESH parameter is set to true1.
Reference:
Materialized Views | Snowflake Documentation

NEW QUESTION # 142
Which of the below commands will use warehouse credits?

- A. SELECT COUNT(*) FROM SNOWFLAKE;
- B. SELECT MAX(FLAKE_ID) FROM SNOWFLAKE;
- C. SELECT COUNT(FLAKE_ID) FROM SNOWFLAKE GROUP BY FLAKE_ID;
- D. SHOW TABLES LIKE 'SNOWFL%';

Answer: A,B,C

Explanation:
* Warehouse credits are used to pay for the processing time used by each virtual warehouse in Snowflake.
A virtual warehouse is a cluster of compute resources that enables executing queries, loading data, and performing other DML operations. Warehouse credits are charged based on the number of virtual warehouses you use, how long they run, and their size1.
* Among the commands listed in the question, the following ones will use warehouse credits:
* SELECT MAX(FLAKE_ID) FROM SNOWFLAKE: This command will use warehouse credits because it is a query that requires a virtual warehouse to execute. The query will scan the
* SNOWFLAKE table and return the maximum value of the FLAKE_ID column2. Therefore, option B is correct.
* SELECT COUNT(*) FROM SNOWFLAKE: This command will also use warehouse credits because it is a query that requires a virtual warehouse to execute. The query will scan the SNOWFLAKE table and return the number of rows in the table3. Therefore, option C is correct.
* SELECT COUNT(FLAKE_ID) FROM SNOWFLAKE GROUP BY FLAKE_ID: This command will also use warehouse credits because it is a query that requires a virtual warehouse to execute. The query will scan the SNOWFLAKE table and return the number of rows for each distinct value of the FLAKE_ID column4. Therefore, option D is correct.
* The command that will not use warehouse credits is:
* SHOW TABLES LIKE 'SNOWFL%': This command will not use warehouse credits because it is a metadata operation that does not require a virtual warehouse to execute. The command will return the names of the tables that match the pattern 'SNOWFL%' in the current database and schema5. Therefore, option A is incorrect.
References: : Understanding Compute Cost : MAX Function : COUNT Function : GROUP BY Clause : SHOW TABLES

NEW QUESTION # 143
What Snowflake features should be leveraged when modeling using Data Vault?

- A. Snowflake's ability to hash keys so that hash key joins can run faster than integer joins
- B. Snowflake's support of multi-table inserts into the data model's Data Vault tables
- C. Data needs to be pre-partitioned to obtain a superior data access performance
- D. Scaling up the virtual warehouses will support parallel processing of new source loads

Answer: B

Explanation:
These two features are relevant for modeling using Data Vault on Snowflake. Data Vault is a data modeling approach that organizes data into hubs, links, and satellites. Data Vault is designed to enable high scalability, flexibility, and performance for data integration and analytics. Snowflake is a cloud data platform that supports various data modeling techniques, including Data Vault. Snowflake provides some features that can enhance the Data Vault modeling, such as:
Snowflake's support of multi-table inserts into the data model's Data Vault tables. Multi-table inserts (MTI) are a feature that allows inserting data from a single query into multiple tables in a single DML statement. MTI can improve the performance and efficiency of loading data into Data Vault tables, especially for real-time or near-real-time data integration. MTI can also reduce the complexity and maintenance of the loading code, as well as the data duplication and latency12.
Scaling up the virtual warehouses will support parallel processing of new source loads. Virtual warehouses are a feature that allows provisioning compute resources on demand for data processing. Virtual warehouses can be scaled up or down by changing the size of the warehouse, which determines the number of servers in the warehouse. Scaling up the virtual warehouses can improve the performance and concurrency of processing new source loads into Data Vault tables, especially for large or complex data sets. Scaling up the virtual warehouses can also leverage the parallelism and distribution of Snowflake's architecture, which can optimize

the data loading and querying34.
Reference:
Snowflake Documentation: Multi-table Inserts
Snowflake Blog: Tips for Optimizing the Data Vault Architecture on Snowflake Snowflake Documentation: Virtual Warehouses
Snowflake Blog: Building a Real-Time Data Vault in Snowflake

NEW QUESTION # 144
Two queries are run on the customer_address table:
create or replace TABLE CUSTOMER_ADDRESS (CA_ADDRESS_SK NUMBER(38,0), CA_ADDRESS_ID VARCHAR(16), CA_STREET_NUMBER VARCHAR(IO) CA_STREET_NAME VARCHAR(60), CA_STREET_TYPE VARCHAR(15), CA_SUITE_NUMBER VARCHAR(10), CA_CITY VARCHAR(60), CA_COUNTY VARCHAR(30), CA_STATE VARCHAR(2), CA_ZIP VARCHAR(10), CA_COUNTRY VARCHAR(20), CA_GMT_OFFSET NUMBER(5,2), CA_LOCATION_TYPE VARCHAR(20)); ALTER TABLE DEMO_DB.DEMO_SCH.CUSTOMER_ADDRESS ADD SEARCH OPTIMIZATION ON SUBSTRING(CA_ADDRESS_ID); Which queries will benefit from the use of the search optimization service? (Select TWO).

- A. select*fromDEMO_DB.DEMO_SCH.CUSTOMER_ADDRESSWhereCA_ADDRESS_IDNOT LIKE '%AAAAAAAAPHPPL%';
- B. select * from DEMO_DB.DEMO_SCH.CUSTOMER_ADDRESS Where CA_ADDRESS_ID= substring('AAAAAAAAPHPPLBAAASKDJHASLKDJHASKJD',1,16);
- C. select * from DEMO_DB.DEMO_SCH.CUSTOMER_ADDRESS Where substring(CA_ADDRESS_ID,1,8)= substring('AAAAAAAAPHPPLBAAASKDJHASLKDJHASKJD',1,8);
- D. select*fromDEMO_DB.DEMO_SCH.CUSTOMER_ADDRESSWhereCA_ADDRESS_IDLIKE '%PHPP%';
- E. select*fromDEMO_DB.DEMO_SCH.CUSTOMER_ADDRESSWhereCA_ADDRESS_IDLIKE '%BAAASKD%';

Answer: B,C

Explanation:
The use of the search optimization service in Snowflake is particularly effective when queries involve operations that match exact substrings or start from the beginning of a string. The ALTER TABLE command adding search optimization specifically for substrings on the CA_ADDRESS_ID field allows the service to create an optimized search path for queries using substring matches.
* Option A benefits because it directly matches a substring from the start of the CA_ADDRESS_ID, aligning with the optimization's capability to quickly locate records based on the beginning segments of strings.
* Option B also benefits, despite performing a full equality check, because it essentially compares the full length of CA_ADDRESS_ID to a substring, which can leverage the substring index for efficient
* retrieval.Options C, D, and E involve patterns that do not start from the beginning of the string or use negations, which are not optimized by the search optimization service configured for starting substring matches.References: Snowflake's documentation on the use of search optimization for substring matching in SQL queries.

NEW QUESTION # 145
.....

- 2026 Updated 100% Free ARA-C01 – 100% Free Latest Study Notes | ARA-C01 Reliable Test Prep □ Download 《 ARA-C01 》 for free by simply searching on □ www.pdfvce.com □ □Flexible ARA-C01 Learning Mode
- ARA-C01 Materials □ Reliable ARA-C01 Exam Topics □ Flexible ARA-C01 Learning Mode ☃ Search on □ www.exam4labs.com □ for { ARA-C01 } to obtain exam materials for free download □ARA-C01 Learning Materials
- Flexible ARA-C01 Learning Mode □ ARA-C01 Test Answers □ Flexible ARA-C01 Learning Mode □ Open [www.pdfvce.com] enter 《 ARA-C01 》 and obtain a free download □Vce ARA-C01 Exam
- ARA-C01 latest study torrent - ARA-C01 practice download pdf □ Open website 《 www.easy4engine.com 》 and search for ➽ ARA-C01 □ for free download □ARA-C01 PDF VCE
- ARA-C01 Test Answers □ ARA-C01 Learning Materials □ ARA-C01 Materials □ Search for [ARA-C01] on 「 www.pdfvce.com 」 immediately to obtain a free download □Flexible ARA-C01 Learning Mode
- Test ARA-C01 Book □ ARA-C01 Latest Exam Registration □ Reliable ARA-C01 Exam Topics □ Download " ARA-C01 " for free by simply entering { www.verifieddumps.com } website □ARA-C01 Test Answers
- ARA-C01 Test Free □ ARA-C01 New Test Camp □ Reliable ARA-C01 Exam Topics □ Open 【 www.pdfvce.com 】 and search for ➠ ARA-C01 □ to download exam materials for free □ARA-C01 Reliable Test Sample
- Free PDF Quiz Snowflake - Updated ARA-C01 Latest Study Notes □ Search for ☀ ARA-C01 □☀□ and download it for free on □ www.prepawaypdf.com □ website □ARA-C01 Latest Exam Registration
- www.cncircus.com.cn, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, letterboxd.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, pixabay.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, Disposable vapes

BONUS!!! Download part of Actual4Labs ARA-C01 dumps for free: https://drive.google.com/open?id=1QFv4L_9UKzjpOowEZWrJtX-MmOEX5H8O