



Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.



ARA-C01 MCQs
ARA-C01 TestPrep
ARA-C01 Study Guide
ARA-C01 Practice Test
ARA-C01 Exam Questions



killexams.com

SnowFlake

ARA-C01

SnowPro Advanced Architect Certification

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/ARA-C01>



Question: 191

What conditions should be true for a table to consider search optimization

- A. The table size is at least 100 GB
- B. The table is not clustered OR The table is frequently queried on columns other than the primary cluster key
- C. The table can be of any size

Answer: A,B

Explanation:

Search optimization works best to improve the performance of a query when the following conditions are true:

For the table being queried:

Question: 192

One of your query is taking a long time to finish, when you open the query profiler you see that lot of data is spilling to the remote disk(Bytes spilled to remote storage).

What may be the cause of this?

- A. The amount of memory available for the servers used to execute the operation might not be sufficient to hold intermediate results
- B. The size of the AWS bucket used to hold the data is not sufficient for the query
- C. Number of disks attached to the virtual warehouse is not enough for the processing

Answer: A

Explanation:

This is again a question based on work experience. One variation of this may be, you will be given a query profile snapshot which will be having a huge number against Bytes spilled to remote storage. You will be asked to find the possible cause

Queries Too Large to Fit in Memory

For some operations (e.g. duplicate elimination for a huge data set), the amount of memory available for the servers used to execute the operation might not be sufficient to hold intermediate results. As a result, the query processing engine will start spilling the data to local disk. If the local disk space is not sufficient, the spilled data is then saved to remote disks.

This spilling can have a profound effect on query performance (especially if remote disk is used for spilling). To alleviate this, we recommend:

Question: 193

While loading data into a table from stage, which are the valid copyOptions

- A. CONTINUE
- B. SKIP_FILE
- C. SKIP_FILE_<NUM>
- D. SKIP_FILE_<NUM>%
- E. ABORT_STATEMENT
- F. ERROR_STATEMENT

Answer: A,B,C,D,E

Explanation:

Copy Options (copyOptions)%

You can specify one or more of the following copy options (separated by blank spaces, commas, or new lines):

ON_ERROR = CONTINUE | SKIP_FILE | SKIP_FILE_<num> | SKIP_FILE_<num>% | ABORT_STATEMENT

String (constant) that specifies the action to perform when an error is encountered while loading data from a file:

Supported Values	Notes
CONTINUE	Continue loading the file. The COPY statement returns an error message for a maximum of one error encountered per data file. Note that the difference between the ROWS_PARSED and ROWS_LOADED column values represents the number of rows that include detected errors. However, each of these rows could include multiple errors. To view all errors in the data files, use the VALIDATION_MODE parameter or query the VALIDATE function.
SKIP_FILE	Skip file if any errors encountered in the file.
SKIP_FILE_<num> (e.g. SKIP_FILE_10)	Skip file when the number of errors in the file is equal to or exceeds the specified number.
SKIP_FILE_<num>% (e.g. SKIP_FILE_10%)	Skip file when the percentage of errors in the file exceeds the specified percentage.
ABORT_STATEMENT	Abort the COPY statement if any error is encountered.

Question: 194

For this object, Snowflake executes code outside Snowflake; the executed code is known as remote service.

What is this object called?

- A. External procedure
- B. External function
- C. External Script
- D. External job

Answer: B

Explanation:

An external function calls code that executes outside Snowflake; the executed code is known as a remote service.

Users can write and call their own remote services, or call remote services written by third parties. These remote services can be written using any HTTP server stack, including cloud serverless compute services such as AWS Lambda.

From the perspective of a user running a SQL statement, an external function behaves like any other scalar function. A SQL statement performs the following actions: Calls the function, optionally passing parameters.

Receives a value back from the function.

In SQL statements, external functions generally behave like UDFs (user-defined functions). For example, external functions follow these rules:

Inside Snowflake, an external function is represented as a database object. That object is created in a specific database and schema, and can be referenced using dot notation (e.g. MY_DATABASE.MY_SCHEMA.MY_EXTERNAL_FUNCTION()).

An external function can appear in any clause of a SQL statement in which other types of functions can appear (e.g. the WHERE clause).

The returned value can be a compound value, such as a VARIANT that contains JSON.

External functions can be overloaded; two different functions can have the same name but different signatures (different numbers or data types of input parameters).

An external function can be part of a more complex expression: select upper(zipcode_to_city_external_function(zipcode)) from address_table;

<https://docs.snowflake.com/en/sql-reference/external-functions-introduction.html#what-is-an-external-function>

Question: 195

Validation mode can take the below options

- A. RETURN_<n>_ROWS
- B. RETURN_ERRORS
- C. RETURN_ALL_ERRORS
- D. RETURN_SEVERE_ERRORS_ONLY

Answer: A,B,C

Explanation:

VALIDATION_MODE = RETURN_n_ROWS | RETURN_ERRORS | RETURN_ALL_ERRORS

String (constant) that instructs the COPY command to validate the data files instead of loading them into the specified table; i.e. the COPY command tests the files for errors but does not load them. The command validates the data to be loaded and returns results based on the validation option specified:

Supported Values	Notes
<code>RETURN_<n>_ROWS</code> (e.g. <code>RETURN_10_ROWS</code>)	Validates the specified number of rows, if no errors are encountered; otherwise, fails at the first error encountered in the rows.
<code>RETURN_ERRORS</code>	Returns all errors (parsing, conversion, etc.) across all files specified in the COPY statement.
<code>RETURN_ALL_ERRORS</code>	Returns all errors across all files specified in the COPY statement, including files with errors that were partially loaded during an earlier load because the <code>ON_ERROR</code> copy option was set to <code>CONTINUE</code> during the load.

Question: 196

Which semi structured data function interprets an input string as a JSON document, producing a VARIANT value.

- A. PARSE_JSON
- B. PARSE_XML
- C. STRIP_JSON

Answer: A

Explanation:

Try a hands-on exercise to understand this

create or replace table vartab (n number(2), v variant); insert into vartab

select column1 as n, parse_json(column2) as v

from values (1, 'null'),

(2, null),

(3, 'true'),

(4, '-17'),

(7, '"Om ara pa ca na dhih" '),

(8, '[-1, 12, 289, 2188, false,]'),

(9, '{ "x" : "abc", "y" : false, "z": 10} ') as vals;

select n, v, typeof(v) from vartab;

Question: 197

Remote service in external function can be an AWS Lambda function

- A. TRUE
- B. FALSE

Answer: A

Explanation:

remote service

A remote service is stored and executed outside Snowflake, and returns a value. For example, remote services can be implemented as:

An AWS Lambda function.

An HTTPS server (e.g. Node.js) running on an EC2 instance.

To be called by the Snowflake external function feature, the remote service must:

Expose an HTTPS endpoint.

Accept JSON inputs and return JSON outputs.

Question: 198

Bytes spilled to remote storage in query profile indicates volume of data spilled to remote disk

A. TRUE

B. FALSE

Answer: A

Explanation:

This question may come in various format in the exam, so let us not mug it up. Let us understand what it means.

When you run large aggregations, sorts in snowflake the processing usually happens in memory of the virtual warehouse. But if the virtual warehouse is not properly sized and if it does not have enough memory, the intermediate results starts spilling to remote disk(in AWS, it will be S3). When this happens, it impacts the query performance because now you are retrieving your results from remote disk instead of memory. In most of these cases, if it is a regular occurrence you may need to move to a bigger warehouse.

Also read this section referred in the link

<https://docs.snowflake.com/en/user-guide/ui-query-profile.html#queries-too-large-to-fit-in-memory>

Question: 199

```
{"stuId":2000,"stuCourse":"Snowflake"}
```

How will you write a query that will check if stuId in JSON in #1 is also there in JSON in#2

A. with stu_demography as (select parse_json(column1) as src, src:stuId as ID from values('{"stuId":2000, "stuName":"Amy"}')),

B. stu_course as (select parse_json(column1) as src, src:stuId as ID from

values('{"stuId":2000,"stuCourse":"Snowflake"}')) select case when stdemo.ID in(select ID from stu_course) then 'True' else 'False' end as result from stu_demography stdemo;

C. with stu_demography as (select parse_json(column1) as src, src['stuId'] as ID from values('{"stuId":2000,"stuName":"Amy"}')), stu_course as (select parse_json(column1) as src, src['stuId'] as ID from values('{"stuId":2000,"stuCourse":"Snowflake"}')) select case when stdemo.ID in(select ID from stu_course) then 'True' else 'False' end as result from stu_demography stdemo;

D. SELECT CONTAINS('{"stuId":2000, "stuName":"Amy"}', '{"stuId":2000,"stuCourse":"Snowflake"}');

E. with stu_demography as (select parse_json(column1) as src, src['STUID'] as ID from values('{"stuId":2000,"stuName":"Amy"}')), stu_course as (select parse_json(column1) as src, src['stuId'] as ID from values('{"stuId":2000,"stuCourse":"Snowflake"}')) select case when stdemo.ID in(select ID from stu_course) then 'True' else 'False' end as result from stu_demography stdemo;

Answer: B,C

Explanation:

I would like you to try this out in your snowflake instance and find that out

Please note that this may not be the way the question will appear in the certification exam, but why we are still learning this?

Question: 200

In the default access control hierarchy, both securityadmin and sysadmin are owned by accountadmin

- A. TRUE
- B. FALSE

Answer: A


Explanation:

Role hierarchy is an important concept that you should read thoroughly. More than one question may appear in the exam on this topic. Please remember in snowflake you cannot assign a privilege to a user directly. You need to create role and grant privileges to the role and then assign users to the role. As role can be assigned to another role also.

<https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#role-hierarchy-and-privilege-inheritance>

Role Hierarchy and Privilege Inheritance

The following diagram illustrates the hierarchy for the system-defined roles along with the recommended structure for additional, user-defined custom roles:



Question: 201

You are running a large join on snowflake. You ran it on a medium warehouse and it took almost an hour to run. You then tried to run the join on a large warehouse but still the performance did not improve.

What may be the most possible cause of this.

- A. There may be a symptom on skew in your data where one of the value of the column is significantly more than rest of the values in the column
- B. Your warehouses do not have enough memory
- C. Since you have configured an warehouse with a low auto-suspend value, your warehouse is going down frequently

Answer: A

Explanation:

In the snowflake advanced architect exam, 40% of the questions will be based on work experience and this is one such question. You need to have a very good hold on the concepts of Snowflake. So, what may be happening here?

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including MCQs, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.



Exam Questions:

Killexams.com provides exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these questions, candidates can familiarize themselves with the content and format of the real exam.

Exam MCQs:

Killexams.com offers exam MCQs in PDF format. These questions contain a comprehensive collection of questions and answers that cover the exam topics. By using these MCQs, candidate can enhance their knowledge and improve their chances of success in the certification exam.

Practice Test:

Killexams.com provides practice test through their desktop test engine and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice test cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success:

Killexams.com offers a success guarantee with the exam MCQs. Killexams claim that by using this materials, candidates will pass their exams on the first attempt or they will get refund for the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exam.

Updated Contents:

Killexams.com regularly updates its question bank of MCQs to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.