

[Q97-Q118 Pass Your COF-C02 Exam Easily With 100% Exam Passing Guarantee [2024]



[Q97-Q118] Pass Your COF-C02 Exam Easily With 100% Exam Passing Guarantee [2024]

Pass Your COF-C02 Exam Easily With 100% Exam Passing Guarantee [2024]
COF-C02 Dumps are Available for Instant Access from DumpsMaterials

NO.97 Which account__usage views are used to evaluate the details of dynamic data masking? (Select TWO)

- * ROLES
- * POLICY_REFERENCES
- * QUERY_HISTORY
- * RESOURCE_MONITORS
- * ACCESS_HISTORY

Explanation

<https://docs.snowflake.com/en/user-guide/security-column-ddm-intro.html#:~:text=DESCRIBE%20MASKING%>

NO.98 Two users share a virtual warehouse named wh dev 01. When one of the users loads data, the other one experiences performance issues while querying data.

How does Snowflake recommend resolving this issue?

- * Scale up the existing warehouse.
- * Create separate warehouses for each user.
- * Create separate warehouses for each workload.
- * Stop loading and querying data at the same time.

NO.99 How many resource monitors can be assigned at the account level?

- * 1
- * 2
- * 3
- * 4

NO.100 What are the main differences between the account usage views and the information schema views? (Select TWO).

- * No active warehouse is needed to query account usage views but one is needed to query information schema views.
- * Account usage views do not contain data about tables but information schema views do.
- * Account usage views contain dropped objects but information schema views do not.
- * Data retention for account usage views is 1 year but is 7 days to 6 months for information schema views, depending on the view.
- * Information schema views are read-only but account usage views are not.

NO.101 A user is preparing to load data from an external stage

Which practice will provide the MOST efficient loading performance?

- * Organize files into logical paths
- * Store the files on the external stage to ensure caching is maintained
- * Use pattern matching for regular expression execution
- * Load the data in one large file

NO.102 How long can a data consumer who has a pre-signed URL access data files using Snowflake?

- * Indefinitely
- * Until the result_cache expires
- * Until the retention_time is met
- * Until the expiration time is exceeded

NO.103 When should you consider disabling auto-suspend for a Virtual Warehouse? Choose 2 answers

- * When users will be using compute at different times throughout a 24/7 period
- * When managing a steady workload
- * When the compute must be available with no delay or lag time
- * When you don't want to have to manually turn on the Warehouse each time a user needs it

NO.104 Which of the following Snowflake features provide continuous data protection automatically? (Select TWO).

- * Internal stages
- * Incremental backups
- * Time Travel
- * Zero-copy clones
- * Fail-safe

Time travel and fail safe are the two continuous data protection features that support the recovery of data automatically.

Snowflake provides powerful CDP features for ensuring the maintenance and availability of your historical data (i.e. data that has been changed or deleted):

Querying, cloning, and restoring historical data in tables, schemas, and databases for up to 90 days through Snowflake Time Travel.

Disaster recovery of historical data (by Snowflake) through Snowflake Fail-safe.

<https://docs.snowflake.com/en/user-guide/data-availability.html>

NO.105 How many days is load history for Snowpipe retained?

- * 1 day
- * 7 days
- * 14 days
- * 64 days

NO.106 How would a user run a multi-cluster warehouse in maximized mode?

- * Configure the maximum clusters setting to `“Maximum.”`
- * Turn on the additional clusters manually after starting the warehouse.
- * Set the minimum Clusters and maximum Clusters settings to the same value.
- * Set the minimum clusters and maximum clusters settings to different values.

NO.107 Which Snowflake objects can be restored using Time Travel? (Select VNO).

- * Roles
- * Users
- * Databases
- * Schemas
- * Virtual warehouses

NO.108 What feature can be used to reorganize a very large table on one or more columns?

- * Micro-partitions
- * Clustering keys
- * Key partitions
- * Clustered partitions

Explanation

<https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html>

NO.109 Which database objects can be shared with the Snowflake secure data sharing feature? (Choose two.)

- * Files
- * External tables
- * Secure User-Defined Functions (UDFs)
- * Sequences
- * Streams

NO.110 What will happen if a Snowflake user increases the size of a suspended virtual warehouse?

- * The provisioning of new compute resources for the warehouse will begin immediately.
- * The warehouse will remain suspended but new resources will be added to the query acceleration service.
- * The provisioning of additional compute resources will be in effect when the warehouse is next resumed.
- * The warehouse will resume immediately and start to share the compute load with other running virtual warehouses.

NO.111 How long is the Fail-safe period for temporary and transient tables?

- * There is no Fail-safe period for these tables.
- * 1 day
- * 7 days

- * 31 days
- * 90 days

NO.112 What happens to the underlying table data when a CLUSTER BY clause is added to a Snowflake table?

- * Data is hashed by the cluster key to facilitate fast searches for common data values
- * Larger micro-partitions are created for common data values to reduce the number of partitions that must be scanned
- * Smaller micro-partitions are created for common data values to allow for more parallelism
- * Data may be colocated by the cluster key within the micro-partitions to improve pruning performance

NO.113 Which service or feature in Snowflake is used to improve the performance of certain types of lookup and analytical queries that use an extensive set of WHERE conditions?

- * Data classification
- * Query acceleration service
- * Search optimization service
- * Tagging

NO.114 What is cached during a query on a virtual warehouse?

- * All columns in a micro-partition
- * Any columns accessed during the query
- * The columns in the result set of the query
- * All rows accessed during the query

NO.115 A virtual warehouse is created using the following command:

Create warehouse my_WH with

warehouse_size = MEDIUM

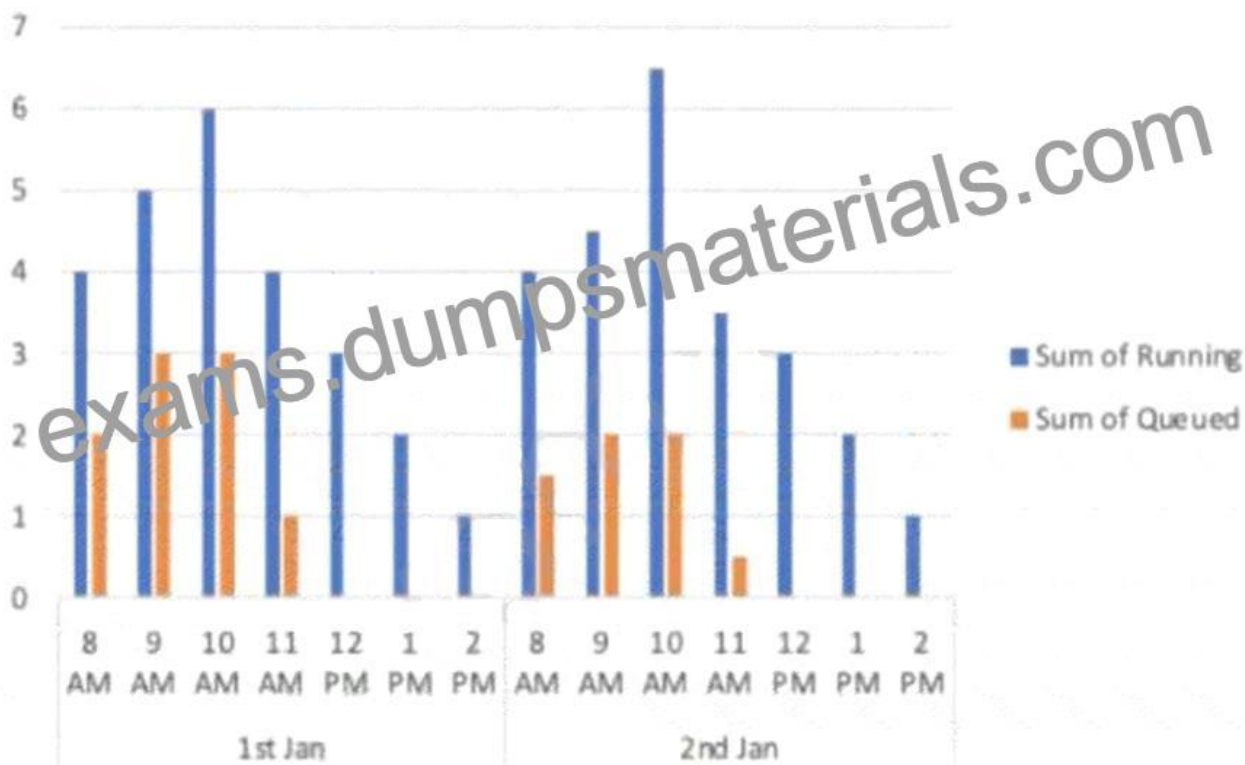
min_cluster_count = 1

max_cluster_count = 1

auto_suspend = 60

auto_resume = true;

The image below is a graphical representation of the warehouse utilization across two days.



What action should be taken to address this situation?

- * Increase the warehouse size from Medium to 2XL.
- * Increase the value for the parameter MAX_CONCURRENCY_LEVEL.
- * Configure the warehouse to a multi-cluster warehouse.
- * Lower the value of the parameter STATEMENT_QUEUED_TIMEOUT_IN_SECONDS.

NO.116 Which Snowflake table supports unstructured data?

- * Directory
- * Transient
- * Temporary
- * Permanent

NO.117 A user is loading JSON documents composed of a huge array containing multiple records into Snowflake. The user enables the strip_outer_array file format option. What does the STRIP_OUTER_ARRAY file format do?

- * It removes the last element of the outer array.
- * It removes the outer array structure and loads the records into separate table rows,
- * It removes the trailing spaces in the last element of the outer array and loads the records into separate table columns
- * It removes the NULL elements from the JSON object eliminating invalid data and enables the ability to load the records

Explanation

Data Size Limitations

The VARIANT data type imposes a 16 MB size limit on individual rows.

For some semi-structured data formats (e.g. JSON), data sets are frequently a simple concatenation of multiple documents. The

JSON output from some software is composed of a single huge array containing multiple records. There is no need to separate the documents with line breaks or commas, though both are supported.

If the data exceeds 16 MB, enable the STRIP_OUTER_ARRAY file format option for the COPY INTO

<table> command to remove the outer array structure and load the records into separate table rows:

copy into <table>

from @~/<file>.json

file_format = (type = '‘JSON’ strip_outer_array = true);

<https://docs.snowflake.com/en/user-guide/semistructured-considerations.html>

NO.118 From what stage can a Snowflake user omit the FROM clause while loading data into a table?

- * The user stage
- * The table stage
- * The internal named stage
- * The external named stage

Study resources for the Valid COF-C02 Braindumps: <https://www.dumpsmaterials.com/COF-C02-real-torrent.html>