ARA-C01 Q&As

SnowPro Advanced: Architect Certification Exam

Pass Snowflake ARA-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.geekcert.com/ara-c01.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Snowflake Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A large manufacturing company runs a dozen individual Snowflake accounts across its business divisions. The company wants to increase the level of data sharing to support supply chain optimizations and increase its purchasing leverage with multiple vendors.

The company\\'s Snowflake Architects need to design a solution that would allow the business divisions to decide what to share, while minimizing the level of effort spent on configuration and management. Most of the company divisions use Snowflake accounts in the same cloud deployments with a few exceptions for European-based divisions.

According to Snowflake recommended best practice, how should these requirements be met?

A. Migrate the European accounts in the global region and manage shares in a connected graph architecture. Deploy a Data Exchange.

B. Deploy a Private Data Exchange in combination with data shares for the European accounts.

C. Deploy to the Snowflake Marketplace making sure that invoker_share() is used in all secure views.

D. Deploy a Private Data Exchange and use replication to allow European data shares in the Exchange.

Correct Answer: B

Explanation: According to Snowflake recommended best practice, the requirements of the large manufacturing company should be met by deploying a Private Data Exchange in combination with data shares for the European accounts. A Private Data Exchange is a feature of the Snowflake Data Cloud platform that enables secure and governed sharing of data between organizations. It allows Snowflake customers to create their own data hub and invite other parts of their organization or external partners to access and contribute data sets. A Private Data Exchange provides centralized management, granular access control, and data usage metrics for the data shared in the exchange1. A data share is a secure and direct way of sharing data between Snowflake accounts without having to copy or move the data. A data share allows the data provider to grant privileges on selected objects in their account to one or more data consumers in other accounts2. By using a Private Data Exchange in combination with data shares, the company can achieve the following benefits: The business divisions can decide what data to share and publish it to the Private Data Exchange, where it can be discovered and accessed by other members of the exchange. This reduces the effort and complexity of managing multiple data sharing relationships and configurations. The company can leverage the existing Snowflake accounts in the same cloud deployments to create the Private Data Exchange and invite the members to join. This minimizes the migration and setup costs and leverages the existing Snowflake features and security. The company can use data shares to share data with the European accounts that are in different regions or cloud platforms. This allows the company to comply with the regional and regulatory requirements for data sovereignty and privacy, while still enabling data collaboration across the organization. The company can use the Snowflake Data Cloud platform to perform data analysis and transformation on the shared data, as well as integrate with other data sources and applications. This enables the company to optimize its supply chain and increase its purchasing leverage with multiple vendors. The other options are incorrect because they do not meet the requirements or follow the best practices. Option A is incorrect because migrating the European accounts to the global region may violate the data sovereignty and privacy regulations, and deploying a Data Exchange may not provide the level of control and management that the company needs. Option C is incorrect because deploying to the Snowflake Marketplace may expose the company\\'s data to unwanted consumers, and using invoker_share() in secure views may not provide the desired level of security and governance. Option D is incorrect because using replication to allow European data shares in the Exchange may incur additional costs and complexity, and may not be necessary if data shares can be used instead. References: Private Data Exchange | Snowflake Documentation, Introduction to Secure Data Sharing | Snowflake Documentation

QUESTION 2

Which system functions does Snowflake provide to monitor clustering information within a table (Choose two.)

A. SYSTEM$CLUSTERING_INFORMATION

B. SYSTEM$CLUSTERING_USAGE

C. SYSTEM$CLUSTERING_DEPTH

D. SYSTEM$CLUSTERING_KEYS

E. SYSTEM$CLUSTERING_PERCENT

Correct Answer: AC

Explanation: According to the Snowflake documentation, these two system functions are provided by Snowflake to monitor clustering information within a table. A system function is a type of function that allows executing actions or returning information about the system. A clustering key is a feature that allows organizing data across micro-partitions based on one or more columns in the table. Clustering can improve query performance by reducing the number of files to scan. SYSTEM$CLUSTERING_INFORMATION is a system function that returns clustering information, including average clustering depth, for a table based on one or more columns in the table. The function takes a table name and an optional column name or expression as arguments, and returns a JSON string with the clustering information. The clustering information includes the cluster by keys, the total partition count, the total constant partition count, the average overlaps, and the average depth1. SYSTEM$CLUSTERING_DEPTH is a system function that returns the clustering depth for a table based on one or more columns in the table. The function takes a table name and an optional column name or expression as arguments, and returns an integer value with the clustering depth. The clustering depth is the maximum number of overlapping micro-partitions for any micro-partition in the table. A lower clustering depth indicates a better clustering2. References: SYSTEM$CLUSTERING_INFORMATION | Snowflake Documentation SYSTEM$CLUSTERING_DEPTH | Snowflake Documentation

QUESTION 3

A Snowflake Architect is designing a multi-tenant application strategy for an organization in the Snowflake Data Cloud and is considering using an Account Per Tenant strategy.

Which requirements will be addressed with this approach? (Choose two.)

A. There needs to be fewer objects per tenant.

B. Security and Role-Based Access Control (RBAC) policies must be simple to configure.

C. Compute costs must be optimized.

D. Tenant data shape may be unique per tenant.

E. Storage costs must be optimized.

Correct Answer: DE

An Account Per Tenant strategy means creating a separate Snowflake account for each tenant (customer or business unit) of the multi-tenant application. This approach has some advantages and disadvantages compared to other strategies,

such as Database Per Tenant or Schema Per Tenant. One advantage is that each tenant can have a unique data shape, meaning they can define their own tables, views, and other objects without affecting other tenants. This allows for more

flexibility and customization for each tenant.

Therefore, option D is correct.

Another advantage is that storage costs can be optimized, because each tenant can use their own storage credits and manage their own data retention policies. This also reduces the risk of data spillover or cross-tenant access. Therefore,

option E is correct.

However, this approach also has some drawbacks, such as:

References: : Multi-Tenant Application Strategies

QUESTION 4

When loading data from stage using COPY INTO, what options can you specify for the ON_ERROR clause? (Choose three.)

A. CONTINUE

B. SKIP_FILE

C. ABORT_STATEMENT

D. FAIL

Correct Answer: ABC

The ON_ERROR clause is an optional parameter for the COPY INTO command that specifies the behavior of the command when it encounters errors in the files. The ON_ERROR clause can have one of the following values1:

Therefore, options A, B, and C are correct.

References: : COPY INTO

QUESTION 5

What are some of the characteristics of result set caches? (Choose three.)

A. Time Travel queries can be executed against the result set cache.

B. Snowflake persists the data results for 24 hours.

C. Each time persisted results for a query are used, a 24-hour retention period is reset.

D. The data stored in the result cache will contribute to storage costs.

E. The retention period can be reset for a maximum of 31 days.

F. The result set cache is not shared between warehouses.

Correct Answer: BCE

Explanation: Comprehensive and Detailed Explanation: According to the SnowPro Advanced: Architect documents and learning resources, some of the characteristics of result set caches are: Snowflake persists the data results for 24 hours. This means that the result set cache holds the results of every query executed in the past 24 hours, and can be reused if the same query is submitted again and the underlying data has not changed1. Each time persisted results for a query are used, a 24-hour retention period is reset. This means that the result set cache extends the lifetime of the results every time they are reused, up to a maximum of 31 days from the date and time that the query was first executed1. The retention period can be reset for a maximum of 31 days. This means that the result set cache will purge the results after 31 days, regardless of whether they are reused or not. After 31 days, the next time the query is submitted, a new result is generated and persisted1. The other options are incorrect because they are not characteristics of result set caches. Option A is incorrect because Time Travel queries cannot be executed against the result set cache. Time Travel queries use the AS OF clause to access historical data that is stored in the storage layer, not the result set cache2. Option D is incorrect because the data stored in the result set cache does not contribute to storage costs. The result set cache is maintained by the service layer, and does not incur any additional charges1. Option F is incorrect because the result set cache is shared between warehouses. The result set cache is available across virtual warehouses, so query results returned to one user are available to any other user on the system who executes the same query, provided the underlying data has not changed1. References: Using Persisted Query Results | Snowflake Documentation, Time Travel | Snowflake Documentation