TestkingIT

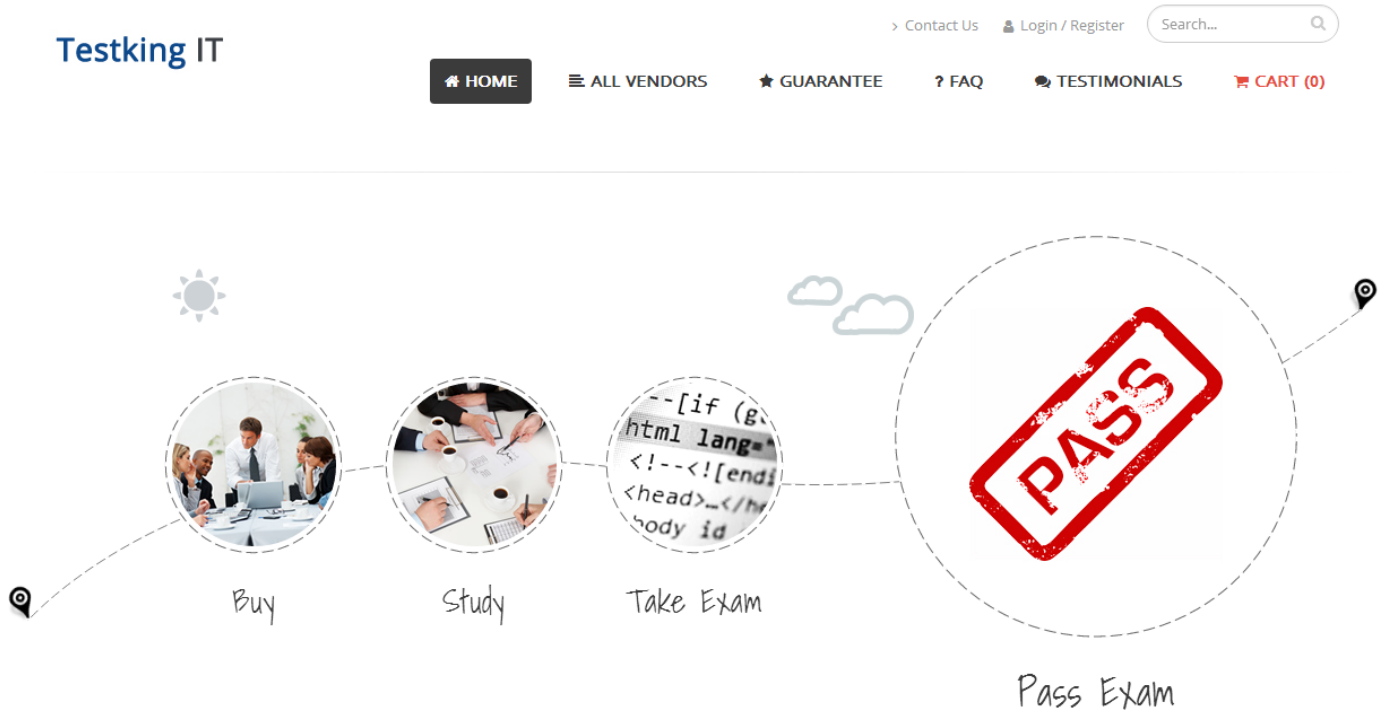

Try **Desktop Test Engine** before you buy

We're not the only ones **happy** about TestKingsIT Practice Material ...

48236+ customers in 100+ countries use TestKingsIT Test Engine. Meet our customers.

iMessenger VOREED GetCustom JET ORANGE iCompany Paradoxx

http://www.testkingit.com/

Latest practice material - Exam Cram - TestKingIT

Exam : **ARA-C01**

Title : SnowPro Advanced Architect Certification

Vendor : Snowflake

Version : DEMO

NO.1 An Architect Is designing a data lake with Snowflake. The company has structured, semi-structured, and unstructured data. The company wants to save the data inside the data lake within the Snowflake system. The company is planning on sharing data among Its corporate branches using Snowflake data sharing.
What should be considered when sharing the unstructured data within Snowflake?
A. A pre-signed URL should be used to save the unstructured data into Snowflake in order to share data over secure views, with no time limit for the URL.
B. A scoped URL should be used to save the unstructured data into Snowflake in order to share data over secure views, with a 24-hour time limit for the URL.
C. A file URL should be used to save the unstructured data into Snowflake in order to share data over secure views, with a 7-day time limit for the URL.
D. A file URL should be used to save the unstructured data into Snowflake in order to share data over secure views, with the "expiration_time" argument defined for the URL time limit.
Answer: D
Explanation:
According to the Snowflake documentation, unstructured data files can be shared by using a secure view and Secure Data Sharing. A secure view allows the result of a query to be accessed like a table, and a secure view is specifically designated for data privacy. A scoped URL is an encoded URL that permits temporary access to a staged file without granting privileges to the stage. The URL expires when the persisted query result period ends, which is currently 24 hours. A scoped URL is recommended for file administrators to give scoped access to data files to specific roles in the same account. Snowflake records information in the query history about who uses a scoped URL to access a file, and when. Therefore, a scoped URL is the best option to share unstructured data within Snowflake, as it provides security, accountability, and control over the data access. References:
Sharing unstructured Data with a secure view
Introduction to Loading Unstructured Data

NO.2 An Architect executes the following statements in order:
CREATE TABLE emp (id INTEGER);
INSERT INTO emp VALUES (1),(2);
CREATE TEMPORARY TABLE emp (id INTEGER);
INSERT INTO emp VALUES (1);
Then executes:
SELECT COUNT(*) FROM emp;
DROP TABLE emp;
SELECT COUNT(*) FROM emp;
What will be the result?
A. COUNT() = 2
COUNT() = 1
B. COUNT() = 1
COUNT() = 2
C. COUNT() = 2
COUNT() = 2
D. The final query results in an error.
Answer: B

Explanation:
In Snowflake, temporary tables take precedence over permanent tables when they share the same name and exist in the same session. After creating the permanent emp table and inserting two rows, a temporary table with the same name is created. From that point forward in the session, all references to emp resolve to the temporary table, not the permanent one.
The insert following the creation of the temporary table adds one row to the temporary table. Therefore, the first SELECT COUNT(*) FROM emp returns 1, reflecting the single row in the temporary table.
When DROP TABLE emp is executed, it drops the temporary table first because it shadows the permanent table. After the temporary table is dropped, the permanent table named emp becomes visible again within the session. The final SELECT COUNT(*) FROM emp then queries the permanent table, which still contains the original two rows inserted earlier.
This behavior is frequently tested in SnowPro Architect exams to validate understanding of object resolution precedence, session scope, and temporary object behavior.
=========

NO.3 A Snowflake account has the following parameters:
* MIN_DATA_RETENTION_TIME_IN_DAYS is set to 5 at the account level.
* DATA_RETENTION_TIME_IN_DAYS is set to 4 on database DB1, and 6 on Schema1 in DB1.
* DATA_RETENTION_TIME_IN_DAYS is set to 5 on database DB2, and 8 on Schema2 in DB2.What will be the result?

A. DB1 and Schema1 retained 5 days; DB2 and Schema2 retained 5 days.

B. DB1 and Schema1 retained 4 days; DB2 and Schema2 retained 5 days.

C. DB1 retained 4 days and Schema1 6 days; DB2 retained 5 days and Schema2 8 days.

D. DB1 retained 5 days and Schema1 6 days; DB2 retained 5 days and Schema2 8 days.

Answer: D
Comprehensive and Detailed 150 to 250 words of Explanation From Snowflake SnowPro Architect exam scope and all publicly documented material:
Snowflake Time Travel retention is governed by a parameter hierarchy and by edition-dependent limits, but the key behavior in this scenario is the minimum retention guardrail.
MIN_DATA_RETENTION_TIME_IN_DAYS sets a floor: objects cannot effectively use a DATA_RETENTION_TIME_IN_DAYS value below that minimum. With an account-level minimum of 5 days, a database-level setting of 4 days is below the permitted minimum, so the effective retention for DB1 cannot be 4 days; it must be at least 5 days. Meanwhile, Schema1 is set to 6 days, which is above the minimum and should apply at the schema level (schemas can override their parent database settings within allowed bounds). For DB2, the database setting is 5 days, which matches the minimum and is valid; Schema2 at 8 days is also valid and applies to objects under that schema where applicable. From an architectural and operational perspective, this highlights two exam-relevant points: (1) retention is controlled via hierarchical parameters (account # database # schema # table), and (2) minimum retention settings enforce governance constraints across environments, preventing overly aggressive reductions that might violate recovery requirements.
=========

NO.4 A new table and streams are created with the following commands:
CREATE OR REPLACE TABLE LETTERS (ID INT, LETTER STRING) ;
CREATE OR REPLACE STREAM STREAM_1 ON TABLE LETTERS;

CREATE OR REPLACE STREAM STREAM_2 ON TABLE LETTERS APPEND_ONLY = TRUE;
The following operations are processed on the newly created table:
INSERT INTO LETTERS VALUES (1, 'A');
INSERT INTO LETTERS VALUES (2, 'B');
INSERT INTO LETTERS VALUES (3, 'C');
TRUNCATE TABLE LETTERS;
INSERT INTO LETTERS VALUES (4, 'D');
INSERT INTO LETTERS VALUES (5, 'E');
INSERT INTO LETTERS VALUES (6, 'F');
DELETE FROM LETTERS WHERE ID = 6;
What would be the output of the following SQL commands, in order?
SELECT COUNT (*) FROM STREAM_1;
SELECT COUNT (*) FROM STREAM_2;

A. 2 & 6

B. 2 & 3

C. 4 & 3

D. 4 & 6

Answer: C

Explanation:
In Snowflake, a stream records data manipulation language (DML) changes to its base table since the stream was created or last consumed. STREAM_1 will show all changes including the TRUNCATE operation, while STREAM_2, being APPEND_ONLY, will not show deletions like TRUNCATE. Therefore, STREAM_1 will count the three inserts, the TRUNCATE (counted as a single operation), and the subsequent two inserts before the delete, totaling 4. STREAM_2 will only count the three initial inserts and the two after the TRUNCATE, totaling 3, as it does not count the TRUNCATE or the delete operation.
References: The explanation is based on the Snowflake documentation on streams, which details how streams track changes and the difference between standard and APPEND_ONLY streams12.

NO.5 A Developer is having a performance issue with a Snowflake query. The query receives up to 10 different values for one parameter and then performs an aggregation over the majority of a fact table. It then joins against a smaller dimension table. This parameter value is selected by the different query users when they execute it during business hours. Both the fact and dimension tables are loaded with new data in an overnight import process.
On a Small or Medium-sized virtual warehouse, the query performs slowly. Performance is acceptable on a size Large or bigger warehouse. However, there is no budget to increase costs. The Developer needs a recommendation that does not increase compute costs to run this query.
What should the Architect recommend?

A. Create a task that will run the 10 different variations of the query corresponding to the 10 different parameters before the users come in to work. The query results will then be cached and ready to respond quickly when the users re-issue the query.

B. Create a task that will run the 10 different variations of the query corresponding to the 10 different parameters before the users come in to work. The task will be scheduled to align with the users' working hours in order to allow the warehouse cache to be used.

C. Enable the search optimization service on the table. When the users execute the query, the search

optimization service will automatically adjust the query execution plan based on the frequently-used parameters.

D. Create a dedicated size Large warehouse for this particular set of queries. Create a new role that has USAGE permission on this warehouse and has the appropriate read permissions over the fact and dimension tables. Have users switch to this role and use this warehouse when they want to access this data.

Answer: C

Explanation:

Enabling the search optimization service on the table can improve the performance of queries that have selective filtering criteria, which seems to be the case here. This service optimizes the execution of queries by creating a persistent data structure called a search access path, which allows some micro-partitions to be skipped during the scanning process. This can significantly speed up query performance without increasing compute costs1.

References

*Snowflake Documentation on Search Optimization Service1.

NO.6 Assuming all Snowflake accounts are using an Enterprise edition or higher, in which development and testing scenarios would be copying of data be required, and zero-copy cloning not be suitable? (Select TWO).

A. Developers create their own datasets to work against transformed versions of the live data.

B. Production and development run in different databases in the same account, and Developers need to see production-like data but with specific columns masked.

C. Data is in a production Snowflake account that needs to be provided to Developers in a separate development/testing Snowflake account in the same cloud region.

D. Developers create their own copies of a standard test database previously created for them in the development account, for their initial development and unit testing.

E. The release process requires pre-production testing of changes with data of production scale and complexity. For security reasons, pre-production also runs in the production account.

Answer: B C

Explanation:

https://docs.snowflake.com/en/user-guide/tag-based-masking-policies#considerations

NO.7 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 * from DEMO_DB.DEMO_SCH.CUSTOMER_ADDRESS Where substring (CA_ADDRESS_ID,1,8)= substring('AAAAAAAAPHPPLBAAASKDJHASLKDJHASKJD',1,8);

B. select * from DEMO_DB.DEMO_SCH.CUSTOMER_ADDRESS Where CA_ADDRESS_ID= substring('AAAAAAAAPHPPLBAAASKDJHASLKDJHASKJD',1,16);

C. select*fromDEMO_DB.DEMO_SCH.CUSTOMER_ADDRESSWhereCA_ADDRESS_IDLIKE '% BAAASKD%';
D. select*fromDEMO_DB.DEMO_SCH.CUSTOMER_ADDRESSWhereCA_ADDRESS_IDLIKE '% PHPP%';
E. select*fromDEMO_DB.DEMO_SCH.CUSTOMER_ADDRESSWhereCA_ADDRESS_IDNOT LIKE '%AAAAAAAAPHPPL%';
Answer: A B
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 theCA_ADDRESS_IDfield allows the service to create an optimized search path for queries using substring matches.
Option Abenefits because it directly matches a substring from the start of theCA_ADDRESS_ID, aligning with the optimization's capability to quickly locate records based on the beginning segments of strings.
References:Snowflake's documentation on the use of search optimization for substring matching in SQL queries.

NO.8 What built-in Snowflake features make use of the change tracking metadata for a table? (Choose two.)
A. The MERGE command
B. The UPSERT command
C. The CHANGES clause
D. A STREAM object
E. The CHANGE_DATA_CAPTURE command
Answer: A D
Explanation:
In Snowflake, the change tracking metadata for a table is utilized by the MERGE command and the STREAM object. The MERGE command uses change tracking to determine how to apply updates and inserts efficiently based on differences between source and target tables. STREAM objects, on the other hand, specifically capture and store change data, enabling incremental processing based on changes made to a table since the last stream offset was committed.
References:Snowflake Documentation on MERGE and STREAM Objects.

NO.9 When using the Snowflake Connector for Kafka, what data formats are supported for the messages? (Choose two.)
A. CSV
B. XML
C. Avro
D. JSON
E. Parquet
Answer: C D
Explanation:
The data formats that are supported for the messages when using the Snowflake Connector for Kafka are Avro and JSON. These are the two formats that the connector can parse and convert into

Snowflake table rows. The connector supports both schemaless and schematized JSON, as well as Avro with or without a schema registry1. The other options are incorrect because they are not supported data formats for the messages. CSV, XML, and Parquet are not formats that the connector can parse and convert into Snowflake table rows. If the messages are in these formats, the connector will load them as VARIANT data type and store them as raw strings in the table2. References: Snowflake Connector for Kafka | Snowflake Documentation, Loading Protobuf Data using the Snowflake Connector for Kafka | Snowflake Documentation

NO.10 What considerations need to be taken when using database cloning as a tool for data lifecycle management in a development environment? (Select TWO).

A. Any pipes in the source are not cloned.

B. Any pipes in the source referring to internal stages are not cloned.

C. Any pipes in the source referring to external stages are not cloned.

D. The clone inherits all granted privileges of all child objects in the source object, including the database.

E. The clone inherits all granted privileges of all child objects in the source object, excluding the database.

Answer: A C

NO.11 An Architect is designing a disaster recovery plan for a global fraud reporting system. The plan must support near real-time systems using Snowflake data, operate near regional centers with fully redundant failover, and must not be publicly accessible.

Which steps must the Architect take? (Select THREE).

A. Create multiple replicating Snowflake Standard edition accounts.

B. Establish one Snowflake account using a Business Critical edition or higher.

C. Establish multiple Snowflake accounts in each required region with independent data sets.

D. Set up Secure Data Sharing among all Snowflake accounts in the organization.

E. Create a Snowflake connection object.

F. Create a failover group for the fraud data for each regional account.

Answer: B, C, F

Mission-critical, near real-time systems with strict availability and security requirements require advanced Snowflake features. Business Critical edition (or higher) is required to support failover groups and cross- region replication with higher SLA guarantees and compliance capabilities (Answer B). To meet regional proximity and redundancy requirements, multiple Snowflake accounts must be deployed in each required region, ensuring independence and isolation between regional environments (Answer C).

Failover groups are the core Snowflake mechanism for disaster recovery. They replicate selected databases, schemas, and roles across accounts and allow controlled promotion of secondary accounts to primary during failover events (Answer F). Secure Data Sharing alone does not provide DR or replication, and connection objects are unrelated to availability or redundancy.

This design aligns with SnowPro Architect best practices for multi-region disaster recovery, enabling low- latency regional access, controlled failover, and strong isolation without exposing systems to the public internet.

=========

NO.12 An Architect wants to build an automated ETL pipeline that reads data from an external stage using an external table, performs transformations on changed data, joins with dimension tables, and loads results into a target table.
What should the Architect use?
A. Streams and tasks
B. Dynamic tables
C. Materialized views
D. Snowpipe with auto-ingest
Answer: A
Streams and tasks provide Snowflake's native framework for building automated, incremental ETL pipelines.
A stream can track changes in the external table (or staging table), while tasks orchestrate the execution of transformation logic and loading into target tables (Answer A).
This approach supports complex transformations, joins with dimension tables, and controlled scheduling or event-driven execution. Materialized views cannot perform joins with arbitrary tables and are not suitable for complex ETL. Dynamic tables simplify transformations but are not designed to consume change data directly from external tables. Snowpipe focuses on ingestion only and does not support downstream transformations.
SnowPro Architect exams frequently test understanding of when to use streams and tasks versus newer abstractions like dynamic tables.
=========

NO.13 A company needs to have the following features available in its Snowflake account:
1. Support for Multi-Factor Authentication (MFA)
2. A minimum of 2 months of Time Travel availability
3. Database replication in between different regions
4. Native support for JDBC and ODBC
5. Customer-managed encryption keys using Tri-Secret Secure
6. Support for Payment Card Industry Data Security Standards (PCI DSS)
In order to provide all the listed services, what is the MINIMUM Snowflake edition that should be selected during account creation?
A. Standard
B. Enterprise
C. Business Critical
D. Virtual Private Snowflake (VPS)
Answer: C
Explanation:
According to the Snowflake documentation1, the Business Critical edition offers the following features that are relevant to the question:
Support for Multi-Factor Authentication (MFA): This is a standard feature available in all Snowflake editions1.
A minimum of 2 months of Time Travel availability: This is an enterprise feature that allows users to access historical data for up to 90 days1.
Database replication in between different regions: This is an enterprise feature that enables users to replicate databases across different regions or cloud platforms1.

Native support for JDBC and ODBC: This is a standard feature available in all Snowflake editions1.
Customer-managed encryption keys using Tri-Secret Secure: This is a business critical feature that provides enhanced security and data protection by allowing customers to manage their own encryption keys1.
Support for Payment Card Industry Data Security Standards (PCI DSS): This is a business critical feature that ensures compliance with PCI DSS regulations for handling sensitive cardholder data1.
Therefore, the minimum Snowflake edition that should be selected during account creation to provide all the listed services is the Business Critical edition.
References:
Snowflake Editions | Snowflake Documentation

NO.14 What is the MOST efficient way to design an environment where data retention is not considered critical, and customization needs are to be kept to a minimum?

A. Use a transient database.

B. Use a transient schema.

C. Use a transient table.

D. Use a temporary table.

Answer: A

Explanation:
Transient databases in Snowflake are designed for situations where data retention is not critical, and they do not have the fail-safe period that regular databases have. This means that data in a transient database is not recoverable after the Time Travel retention period. Using a transient database is efficient because it minimizes storage costs while still providing most functionalities of a standard database without the overhead of data protection features that are not needed when data retention is not a concern.

NO.15 An Architect executes the following query:

```
SELECT query_hash,
COUNT(*) AS query_count,
SUM(QH.EXECUTION_TIME) AS total_execution_time,
SUM((QH.EXECUTION_TIME / (1000 * 60 * 60)) * 8) AS c
FROM SNOWFLAKE.ACCOUNT_USAGE.QUERY_HISTORY QH
WHERE warehouse_name = 'WH_L'
AND DATE_TRUNC('day', start_time) >= CURRENT_DATE() - 3
GROUP BY query_hash
ORDER BY c DESC
LIMIT 10;
```

What information does this query provide? (Select TWO).

A. It shows the total execution time and credit estimates for the 10 most expensive individual queries executed on WH_L over the last 3 days.

B. It shows the total execution time and credit estimates for the 10 most expensive query groups (identical or similar queries) executed on WH_L over the last 3 days.

C. It shows the total execution time and credit estimates for the 10 most frequently run query groups executed on WH_L over the last 3 days.

D. It calculates relative cost by converting execution time to minutes and multiplying by credits used.

E. It calculates relative cost by converting execution time to hours and multiplying by credits used.

Answer: B, E

This query groups results by QUERY_HASH, which represents logically identical SQL statements. As a result, the aggregation is performed at the query group level, not at the individual execution level. This allows architects to identify patterns where the same query (or same logical SQL) repeatedly consumes a large amount of compute (Answer B).

The cost calculation converts execution time from milliseconds to hours by dividing by (1000 * 60 * 60) and then multiplies the result by 8, which represents the hourly credit consumption of the WH_L warehouse size.

This provides a relative estimate of credit usage per query group, not an exact billing value but a useful approximation for cost analysis (Answer E).

The query does not identify the most frequently executed queries; although COUNT(*) is included, the ordering is done by calculated cost (c), not by frequency. This type of analysis is directly aligned with SnowPro Architect responsibilities, helping architects optimize workloads, refactor expensive query patterns, and right-size warehouses to control costs.

=========

NO.16 When loading data from stage using COPY INTO, what options can you specify for the ON_ERROR clause?

A. CONTINUE

B. SKIP_FILE

C. ABORT_STATEMENT

D. FAIL

Answer: A B C

Explanation:

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:

CONTINUE: This value instructs the command to continue loading the file and return an error message for a maximum of one error encountered per data file. The difference between the ROWS_PARSED and ROWS_LOADED column values represents the number of rows that include detected errors. To view all errors in the data files, use the VALIDATION_MODE parameter or query the VALIDATE function1.

SKIP_FILE: This value instructs the command to skip the file when it encounters a data error on any of the records in the file. The command moves on to the next file in the stage and continues loading. The skipped file is not loaded and no error message is returned for the file1.

ABORT_STATEMENT: This value instructs the command to stop loading data when the first error is encountered. The command returns an error message for the file and aborts the load operation. This is the default value for the ON_ERROR clause1.

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

COPY INTO <table>

NO.17 A company has a Snowflake environment running in AWS us-west-2 (Oregon). The company needs to share data privately with a customer who is running their Snowflake environment in Azure East US 2 (Virginia).

What is the recommended sequence of operations that must be followed to meet this requirement?

A. 1. Create a share and add the database privileges to the share2. Create a new listing on the Snowflake Marketplace3. Alter the listing and add the share4. Instruct the customer to subscribe to the listing on the Snowflake Marketplace

B. 1. Ask the customer to create a new Snowflake account in Azure EAST US 2 (Virginia)2. Create a share and add the database privileges to the share3. Alter the share and add the customer's Snowflake account to the share

C. 1. Create a new Snowflake account in Azure East US 2 (Virginia)2. Set up replication between AWS us- west-2 (Oregon) and Azure East US 2 (Virginia) for the database objects to be shared3. Create a share and add the database privileges to the share4. Alter the share and add the customer's Snowflake account to the share

D. 1. Create a reader account in Azure East US 2 (Virginia)2. Create a share and add the database privileges to the share3. Add the reader account to the share4. Share the reader account's URL and credentials with the customer

Answer: C

Explanation:

Option C is the correct answer because it allows the company to share data privately with the customer across different cloud platforms and regions. The company can create a new Snowflake account in Azure East US 2 (Virginia) and set up replication between AWS us-west-2 (Oregon) and Azure East US 2 (Virginia) for the database objects to be shared. This way, the company can ensure that the data is always up to date and consistent in both accounts. The company can then create a share and add the database privileges to the share, and alter the share and add the customer's Snowflake account to the share. The customer can then access the shared data from their own Snowflake account in Azure East US 2 (Virginia).

Option A is incorrect because the Snowflake Marketplace is not a private way of sharing data. The Snowflake Marketplace is a public data exchange platform that allows anyone to browse and subscribe to data sets from various providers. The company would not be able to control who can access their data if they use the Snowflake Marketplace.

Option B is incorrect because it requires the customer to create a new Snowflake account in Azure East US 2 (Virginia), which may not be feasible or desirable for the customer. The customer may already have an existing Snowflake account in a different cloud platform or region, and may not want to incur additional costs or complexity by creating a new account.

Option D is incorrect because it involves creating a reader account in Azure East US 2 (Virginia), which is a limited and temporary way of sharing data. A reader account is a special type of Snowflake account that can only access data from a single share, and has a fixed duration of 30 days. The company would have to manage the reader account's URL and credentials, and renew the account every 30 days. The customer would not be able to use their own Snowflake account to access the shared data, and would have to rely on the company's reader account.

References:

Snowflake Replication

Secure Data Sharing Overview

Snowflake Marketplace Overview

Reader Account Overview

NO.18 A table, EMP_ TBL has three records as shown:

```
create or replace TABLE EMP_TBL (
ID NUMBER(38,0),
NAME VARCHAR(16777216)
);
```

ID	NAME
1	Name1
2	Name2
3	Name3

The following variables are set for the session:

```
set tbl_ref = 'EMP_TBL';
set col_ref = 'NAME';
set (var1, var2, var3) = (select 'Name1', 'Name2', 'Name3');
```

Which SELECT statements will retrieve all three records? (Select TWO).

A. Select * FROM Stbl_ref WHERE Scol_ref IN ('Name1','Nam2','Name3');

B. SELECT * FROM EMP_TBL WHERE identifier(Scol_ref) IN ('Namel','Name2', 'Name3');

C. SELECT * FROM identifier<Stbl_ref> WHERE NAME IN ($var1, $var2, $var3);

D. SELECT * FROM identifier($tbl_ref) WHERE ID IN Cvarl','var2','var3');

E. SELECT * FROM $tb1_ref WHERE $col_ref IN ($var1, Svar2, Svar3);

Answer: B E

Explanation:

The correct answer is B and E because they use the correct syntax and values for the identifier function and the session variables.

The identifier function allows you to use a variable or expression as an identifier (such as a table name or column name) in a SQL statement. It takes a single argument and returns it as an identifier. For example, identifier($tbl_ref) returns EMP_TBL as an identifier.

The session variables are set using the SET command and can be referenced using the $ sign. For example,

$var1 returns Name1 as a value.

Option A is incorrect because it uses Stbl_ref and Scol_ref, which are not valid session variables or identifiers. They should be $tbl_ref and $col_ref instead.

Option C is incorrect because it uses identifier<Stbl_ref>, which is not a valid syntax for the identifier function. It should be identifier($tbl_ref) instead.

Option D is incorrect because it uses Cvarl, var2, and var3, which are not valid session variables or values.

They should be $var1, $var2, and $var3 instead. References:

Snowflake Documentation: Identifier Function

Snowflake Documentation: Session Variables

Snowflake Learning: SnowPro Advanced: Architect Exam Study Guide

NO.19 What are purposes for creating a storage integration? (Choose three.)

A. Control access to Snowflake data using a master encryption key that is maintained in the cloud provider' s key management service.

B. Store a generated identity and access management (IAM) entity for an external cloud provider regardless of the cloud provider that hosts the Snowflake account.

C. Support multiple external stages using one single Snowflake object.

D. Avoid supplying credentials when creating a stage or when loading or unloading data.

E. Create private VPC endpoints that allow direct, secure connectivity between VPCs without traversing the public internet.

F. Manage credentials from multiple cloud providers in one single Snowflake object.

Answer: B C D

Explanation:

The purpose of creating a storage integration in Snowflake includes:

B).Store a generated identity and access management (IAM) entity for an external cloud provider- This helps in managing authentication and authorization with external cloud storage without embedding credentials in Snowflake. It supports various cloud providers like AWS, Azure, or GCP, ensuring that the identity management is streamlined across platforms.

C).Support multiple external stages using one single Snowflake object- Storage integrations allow you to set up access configurations that can be reused across multiple external stages, simplifying the management of external data integrations.

D).Avoid supplying credentials when creating a stage or when loading or unloading data- By using a storage integration, Snowflake can interact with external storage without the need to continuously manage or expose sensitive credentials, enhancing security and ease of operations.

References:Snowflake documentation on storage integrations, found within the SnowPro Advanced: Architect course materials.

NO.20 An Architect needs to ensure that users can upload data from Snowsight into an existing table.

What privileges must be granted? (Select THREE).

A. Database: USAGE

B. Database: OWNERSHIP

C. Schema: CREATE TABLE

D. Schema: USAGE

E. Table: SELECT

F. Table: OWNERSHIP

Answer: A, D, E

Uploading data into an existing table via Snowsight requires sufficient privileges to access the database and schema and to interact with the target table. Database-level USAGE is required to access objects within the database (Answer A). Schema-level USAGE is required to access the schema containing the table (Answer D).

Table-level SELECT is required by Snowsight to validate and preview the data and table structure during the upload process (Answer E). OWNERSHIP privileges are not required and would grant excessive control.

CREATE TABLE is unnecessary when uploading into an existing table.

This reflects Snowflake's least-privilege security model and is a common SnowPro Architect exam topic when designing user self-service data ingestion workflows.
=========

NO.21 A company has a table with that has corrupted data, named Data. The company wants to recover the data as it was 5 minutes ago using cloning and Time Travel.
What command will accomplish this?
A. CREATE CLONE TABLE Recover_Data FROM Data AT(OFFSET => -60*5);
B. CREATE CLONE Recover_Data FROM Data AT(OFFSET => -60*5);
C. CREATE TABLE Recover_Data CLONE Data AT(OFFSET => -60*5);
D. CREATE TABLE Recover Data CLONE Data AT(TIME => -60*5);
Answer: C
Explanation:
This is the correct command to create a clone of the table Data as it was 5 minutes ago using cloning and Time Travel. Cloning is a feature that allows creating a copy of a database, schema, table, or view without duplicating the data or metadata. Time Travel is a feature that enables accessing historical data (i.e. data that has been changed or deleted) at any point within a defined period. To create a clone of a table at a point in time in the past, the syntax is:
CREATE TABLE <clone_name> CLONE <source_table> AT (OFFSET => <offset_in_seconds>); The OFFSET parameter specifies the time difference in seconds from the present time. A negative value indicates a point in the past. For example, -60*5 means 5 minutes ago. Alternatively, the TIMESTAMP parameter can be used to specify an exact timestamp in the past. The clone will contain the data as it existed in the source table at the specified point in time12.
Snowflake Documentation: Cloning Objects
Snowflake Documentation: Cloning Objects at a Point in Time in the Past

NO.22 An Architect with the ORGADMIN role wants to change a Snowflake account from an Enterprise edition to a Business Critical edition.
How should this be accomplished?
A. Run an ALTER ACCOUNT command and create a tag of EDITION and set the tag to Business Critical.
B. Use the account's ACCOUNTADMIN role to change the edition.
C. Failover to a new account in the same region and specify the new account's edition upon creation.
D. Contact Snowflake Support and request that the account's edition be changed.
Answer: D
Explanation:
To change the edition of a Snowflake account, an organization administrator (ORGADMIN) cannot directly alter the account settings through SQL commands or the Snowflake interface. The proper procedure is to contact Snowflake Support to request an edition change for the account. This ensures that the change is managed correctly and aligns with Snowflake's operational protocols.
References: This process is outlined in the Snowflake documentation, which specifies that changes to an account's edition should be facilitated through Snowflake Support1.

NO.23 Which feature provides the capability to define an alternate cluster key for a table with an existing cluster key?
A. External table

B. Materialized view

C. Search optimization

D. Result cache

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.

Materialized Views | Snowflake Documentation