Snowflake ADA-C02 Dumps

Shared by Sparks on 25-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A virtual warehouse report_wh is configured with AUTO_RESUME=TRUE and AUTO_SUSPEND=300. A user has been granted the accountant role.

An application using the accountant role must use this warehouse to run financial reports and track the compute credits consumed by the warehouse. What minimum warehouse privileges should be granted to the role to satisfy these requirements?

Choose two

P2P exams

Options:

A- OPERATE
B- MODIFY
C- MONITOR
D- USAGE
E- OWNERSHIP

Answer:

C, D

Explanation:

Snowflake separates warehouse execution, monitoring, and administration through distinct privileges. The accountant role requires USAGE to execute financial-reporting queries and MONITOR to inspect warehouse activity and consumption statistics.

The USAGE privilege authorizes a role to submit queries to the specified virtual warehouse. Because report_wh has AUTO_RESUME=TRUE, Snowflake automatically resumes the suspended warehouse when a query is submitted. Consequently, the accountant role does not require OPERATE merely to execute reports.

The MONITOR privilege enables inspection of current and historical warehouse queries and usage statistics. This provides the necessary visibility into warehouse activity and compute consumption without allowing modification of warehouse configuration.

AUTO_SUSPEND=300 directs Snowflake to suspend the warehouse following approximately five minutes of inactivity, reducing unnecessary compute consumption.

Option A grants additional authority to suspend and resume the warehouse manually. Option B permits administrative configuration changes, including resizing. Option E grants full ownership-

level control and violates the principle of least privilege.

Therefore, granting MONITOR and USAGE satisfies the specified operational requirements without introducing unnecessary administrative permissions.

For account-wide billing and consumption analysis, administrators should distinguish warehouse-level MONITOR from the separate global MONITOR USAGE privilege.

Official Snowflake references: Access Control Privileges: Virtual Warehouse Privileges; Configuring Access Control; Virtual Warehouse Credit Usage.

Question 2

Question Type: MultipleChoice

An Administrator has a warehouse that is intended to have a credit quota of 3000 for every calendar year. The Administrator must create a resource monitor that performs these tasks:

At 80% usage, notify the account Administrators.

At 100% usage, suspend the warehouse and notify the account Administrators.

At 120% usage, stop all running executions, suspend the warehouse, and notify the account Administrators. Which SQL command meets these requirements?

Options:

A- CREATE OR REPLACE RESOURCE MONITOR RM1 WITH CREDIT_QUOTA = 3000
START_TIMESTAMP = '2022-01-01 00:00 CET' TRIGGERS
ON 80 PERCENT DO NOTIFY
ON 100 PERCENT DO SUSPEND
ON 120 PERCENT DO SUSPEND_IMMEDIATE;
ALTER WAREHOUSE WH1 SET RESOURCE_MONITOR = RM1;
B- CREATE OR REPLACE RESOURCE MONITOR RM1 WITH CREDIT_QUOTA = 3000
FREQUENCY = YEARLY START_TIMESTAMP = '2022-01-01 00:00 CET' TRIGGERS
ON 80 PERCENT DO NOTIFY ON 100 PERCENT DO SUSPEND
ON 120 PERCENT DO SUSPEND_IMMEDIATE;
ALTER WAREHOUSE WH1 SET RESOURCE_MONITOR = RM1;
C- CREATE OR REPLACE RESOURCE MONITOR RM1 WITH CREDIT_QUOTA = 3000
START_TIMESTAMP = '2022-01-01 00:00 CET' TRIGGERS
ON 80 PERCENT DO NOTIFY
ON 100 PERCENT DO NOTIFY AND SUSPEND
ON 120 PERCENT DO NOTIFY AND SUSPEND_IMMEDIATE; ALTER WAREHOUSE WH1 SET

RESOURCE_MONITOR = RM1;
D- CREATE OR REPLACE RESOURCE MONITOR RM1 WITH CREDIT_QUOTA = 3000
FREQUENCY = YEARLY TRIGGERS
ON 80 PERCENT DO NOTIFY ON 100 PERCENT DO SUSPEND
ON 120 PERCENT DO SUSPEND_IMMEDIATE;
ALTER WAREHOUSE WH1 SET RESOURCE_MONITOR = RM1;

Answer:

B

Explanation:

Snowflake resource monitors enforce credit-consumption controls through configurable quotas, reset intervals, start timestamps, and threshold-triggered actions.

Option B is the intended configuration because it specifies a 3,000-credit quota, a YEARLY reset frequency, an explicit January 1 start timestamp, and the three required actions.

At 80%, the NOTIFY action generates alerts for administrators who have enabled resource-monitor notifications without suspending the warehouse.

At 100%, SUSPEND prevents new query execution while allowing statements already running to complete. Snowflake also generates the associated suspension notifications.

At 120%, SUSPEND_IMMEDIATE cancels executing statements and suspends the assigned warehouse immediately, generating additional notifications.

Option A omits FREQUENCY despite defining START_TIMESTAMP. Snowflake requires the two scheduling parameters to be specified together for a customized schedule.

Option C omits FREQUENCY and uses unsupported combined actions such as DO NOTIFY AND SUSPEND.

Option D specifies YEARLY but omits the required START_TIMESTAMP.

The ALTER WAREHOUSE statement is necessary because creating a resource monitor does not automatically assign it to a warehouse.

Outdated exhibit warning: The January 1, 2022 timestamp is historical. For deployment in a current environment, the administrator should use an appropriate valid calendar-year start date rather than execute the sample unchanged. Notification delivery must also be enabled for the relevant administrators.

Official Snowflake references: CREATE RESOURCE MONITOR; Working with Resource Monitors --- Schedules, Actions, and Notifications.

Question 3

Question Type: MultipleChoice

Which tasks are available to the ORGADMIN role?

Choose three

Options:

A- Create one or more accounts in the organization.
B- View a list of all regions enabled for the organization.
C- Create secure views on application tables within the organization.
D- View usage information for all accounts in the organization.
E- Perform zero-copy cloning on account data.
F- Create a reader account to share data with another organization.

Answer:

A, B, D

Explanation:

The ORGADMIN role provides organization-level administrative capabilities across Snowflake accounts. Its responsibilities are distinct from database administration and object-level security management within individual accounts.

Option A is correct because organization administrators can create and manage Snowflake accounts across supported cloud platforms and regions. This enables centralized account provisioning without requiring independent account-creation procedures for each deployment.

Option B is correct because organization administrators can execute SHOW REGIONS to identify the regions available for account creation. The returned information includes Snowflake region identifiers and their corresponding cloud providers.

Option D is correct because organization administrators can access organization-wide consumption and administrative information through the SNOWFLAKE.ORGANIZATION_USAGE schema, subject to the applicable organization-account access model.

Option C is incorrect because creating secure views requires the appropriate privileges within the account containing the database objects.

Option E is incorrect because zero-copy cloning is controlled by database and object-level

privileges rather than organization-administrator authority.

Option F is incorrect because reader-account creation is an account-level data-sharing operation.

Important current documentation update: Snowflake is phasing out ORGADMIN administration in regular accounts for multi-account organizations. The recommended approach is to use GLOBALORGADMIN in a dedicated organization account. The original answer remains valid for the ORGADMIN role described in the question.

Official Snowflake references: Organization Administrators; SHOW REGIONS; Organization Usage.

Question 4

Question Type: MultipleChoice

A Snowflake customer is incurring higher-than-expected costs while migrating data-warehouse workloads from on premises to Snowflake. The migration workloads run on one warehouse and consist of many small INSERT statements rather than bulk-loading large extracts. That warehouse is configured as a single 2XL cluster because many parallel INSERT statements are scheduled during nightly loads.

How can the Administrator lower costs while minimizing the overall time required to load data-warehouse history?

Options:

A- There should be another 2XL warehouse deployed to handle a portion of the load queries.
B- The 2XL warehouse should be changed to 4XL to increase the number of threads available for parallel load queries.
C- The warehouse should be kept as a SMALL or XSMALL and configured as a multi-cluster warehouse to handle the parallel load queries.
D- The INSERTs should be converted to several tables to avoid contention on large tables that slows down query processing.

Answer:

C

Explanation:

Snowflake distinguishes between scaling up a warehouse to improve the resources available to individual statements and scaling out through multiple clusters to accommodate concurrent

workloads.

Option C is the intended answer because a multi-cluster warehouse can increase parallel query-processing capacity without requiring a continuously running 2X-Large warehouse. Smaller clusters reduce the cost of each running cluster, while auto-scaling provides additional clusters when concurrency increases.

Option A introduces another expensive warehouse without addressing inefficient workload execution. Option B increases compute capacity but does not necessarily improve the throughput of numerous small INSERT statements. Option D introduces additional tables without establishing that table-level contention is the primary bottleneck.

However, Snowflake's official cost-optimization guidance explicitly identifies frequent single-row INSERT operations as inefficient. These operations can generate substantial cloud-services overhead because Snowflake is designed primarily for analytical workloads rather than traditional OLTP processing.

Technical qualification: Option C is the strongest available answer, but it does not address the principal ingestion-design problem. Snowflake recommends batching INSERT operations or using bulk-loading mechanisms such as COPY INTO whenever appropriate. Smaller multi-cluster warehouses are not guaranteed to outperform a larger single warehouse for every DML workload.

For a production-quality examination item, the scenario should distinguish concurrency bottlenecks from inefficient ingestion architecture.

Official Snowflake references: Optimizing Cloud Services for Cost; Planning a Data Load; Multi-cluster Warehouses.

Question 5

Question Type: MultipleChoice

Which Snowflake capabilities are typically used for rollback scenarios?

Choose two

Options:

A- SELECT SYSTEM$CANCEL_QUERY('problematic_query_id');
B- CREATE TABLE prd_table_bkp CLONE prd_table BEFORE (STATEMENT = > 'problematic_query_id');
C- CREATE TABLE prd_table_bkp as SELECT * from table(result_scan('problematic_query_id'()));
D- ALTER TABLE prd_table SWAP WITH prd_table_bkp;
E- Contact Snowflake Support to retrieve Fail-safe data.

Answer:

B, D

Explanation:

Snowflake Time Travel and zero-copy cloning provide an efficient mechanism for recovering from accidental data modifications without requiring a conventional full-table backup.

Option B uses Time Travel cloning with the BEFORE clause to create a table representing the historical state immediately preceding a specified statement. The STATEMENT parameter identifies the relevant query ID, and Snowflake reconstructs the qualifying historical version from retained data.

Option D uses ALTER TABLE ... SWAP WITH to exchange the original and restored table objects atomically. This enables the recovered data to become available under the production table name while retaining the displaced table as another object.

Option A cancels an executing query but does not reverse changes that have already committed.

Option C queries a previous statement's result set rather than reconstructing the table's historical state. RESULT_SCAN is not a substitute for Time Travel recovery.

Option E involves Snowflake Fail-safe, which is a last-resort disaster-recovery capability, not the normal administrative mechanism for reversing an erroneous statement.

Technical qualification: The source's problematic_query_id is a placeholder, and the SQL would require an actual qualifying statement ID. Successful recovery also requires the historical version to remain within the configured Time Travel retention period. The administrator must review dependencies, object privileges, and downstream effects before swapping production tables.

Official Snowflake references: Understanding & Using Time Travel; AT | BEFORE; ALTER TABLE; Cloning Considerations.

Question 6

Question Type: MultipleChoice

An Administrator must create a sample from the LINEITEM table. The sample must not be repeatable, and the sampling function must select data in blocks of rows.

Which SELECT command generates a sample of 20% of the table?

Options:

A- select * from LINEITEM sample bernoulli (20);
B- select * from LINEITEM sample system (20);
C- select * from LINEITEM tablesample block (20 rows);
D- select * from LINEITEM tablesample system (20) seed (1);

Answer:

B

Explanation:

Snowflake supports statistical sampling through the SAMPLE and TABLESAMPLE SQL clauses. The two principal sampling methods are BERNOULLI, also called ROW sampling, and SYSTEM, also called BLOCK sampling.

SYSTEM sampling evaluates blocks of rows rather than independently selecting individual rows. The specified percentage represents the probability that each block will be included in the resulting sample.

Option B is correct because: SELECT *

FROM LINEITEM

SAMPLE SYSTEM (20);

selects blocks using a 20% inclusion probability without specifying a deterministic random seed. Consequently, subsequent executions may select different blocks and produce different samples. Option A uses BERNOULLI sampling, which evaluates rows individually rather than selecting blocks.

Option C is incorrect because SYSTEM or BLOCK sampling does not support fixed-row-count sampling through the proposed syntax.

Option D uses the correct block-based sampling method but supplies SEED (1), requesting reproducible sampling behavior. This contradicts the requirement that the sample not be repeatable.

An important distinction is that SYSTEM (20) does not guarantee exactly 20% of the table's rows. Because Snowflake selects entire blocks probabilistically, the actual number of sampled rows can vary depending on block composition and selection.

Official Snowflake references: SAMPLE / TABLESAMPLE: Sampling Methods; Probability; SEED / REPEATABLE.

Question 7

Question Type: MultipleChoice

A Snowflake organization, MYORG, contains two Snowflake accounts:

Account Name	Snowflake Region	Snowflake Edition
ACCOUNT1	AWS_EU_WEST_2	ENTERPRISE
ACCOUNT2	AZURE_WESTEUROPE	STANDARD

ACCOUNT1 has a database named PROD_DB and has the ORGADMIN role enabled. Management wants PROD_DB replicated to ACCOUNT2.

Options:

A- re any configuration steps required in ACCOUNT1 before database replication can be configured and started in ACCOUNT2?

B- USE ROLE ORGADMIN;
SELECT SYSTEM$GLOBAL_ACCOUNT_SET_PARAMETER('MYORG.ACCOUNT1',
'ENABLE_ACCOUNT_DATABASE_REPLICATION',
'TRUE'
);
SELECT SYSTEM$GLOBAL_ACCOUNT_SET_PARAMETER('MYORG.ACCOUNT2',
'ENABLE_ACCOUNT_DATABASE_REPLICATION',
'TRUE'
);
USE ROLE ACCOUNTADMIN; ALTER DATABASE PROD_DB
ENABLE REPLICATION TO ACCOUNTS MYORG.ACCOUNT2;

C- USE ROLE ORGADMIN;
SELECT SYSTEM$GLOBAL_ACCOUNT_SET_PARAMETER('MYORG.ACCOUNT1',
'ENABLE_ACCOUNT_DATABASE_REPLICATION',
'TRUE'
);
USE ROLE ACCOUNTADMIN; ALTER DATABASE PROD_DB
ENABLE REPLICATION TO ACCOUNTS MYORG.ACCOUNT2 IGNORE EDITION CHECK;

D- No configuration steps are necessary in ACCOUNT1. Replicating databases across accounts within the same Snowflake organization is enabled by default.

E- It is not possible to replicate a database from an Enterprise edition Snowflake account to a Standard edition Snowflake account.

Answer:

A

Explanation:

Snowflake supports replication of databases between accounts belonging to the same organization, including accounts in different cloud providers and regions.

Option A correctly performs two prerequisites. First, an organization administrator enables replication for both the source and destination accounts by calling SYSTEM$GLOBAL_ACCOUNT_SET_PARAMETER with ENABLE_ACCOUNT_DATABASE_REPLICATION set to TRUE.

Second, an appropriately authorized account administrator enables replication on PROD_DB and explicitly identifies ACCOUNT2 as an approved destination.

Option B is incorrect because it does not enable replication for ACCOUNT2. Additionally, IGNORE EDITION CHECK is not necessary when the primary account uses Enterprise Edition and the destination uses Standard Edition.

Snowflake reserves that override for specified circumstances involving Business Critical or higher editions and lower-edition destination accounts or related compliance differences.

Option C is incorrect because the required replication features must be enabled for participating accounts.

Option D is incorrect because Snowflake supports database replication from Enterprise Edition to Standard Edition.

Current-version qualification: These commands represent the legacy database-replication workflow. Snowflake now supports replication groups and recommends group-based replication for broader business-continuity configurations. The ORGADMIN role is also being phased out in favor of GLOBALORGADMIN in a dedicated organization account.

Official Snowflake references: SYSTEM$GLOBAL_ACCOUNT_SET_PARAMETER; ALTER DATABASE; Introduction to Replication and Failover Across Multiple Accounts.

Question 8

Question Type: MultipleChoice

Which roles may be used to create network policies in Snowflake accounts?

Choose three

Options:

A- SYSADMIN
B- SECURITYADMIN
C- ACCOUNTADMIN
D- ORGADMIN
E- Any role with the global permission of CREATE NETWORK POLICY
F- Any rote that owns the database where the network policy is created

Answer:

B, C, E

Explanation:

Snowflake network policies are account-level security objects that control inbound connectivity through permitted and blocked network locations.

Creating a network policy requires the account-level CREATE NETWORK POLICY privilege. Snowflake assigns this privilege to SECURITYADMIN and higher roles by default.

Therefore, SECURITYADMIN can create network policies through its predefined security-administration capabilities. ACCOUNTADMIN can also create them because its role hierarchy includes the necessary administrative privileges.

Additionally, Snowflake permits authorized administrators to delegate the global CREATE NETWORK POLICY privilege to custom roles. Such a custom role can create network policies without requiring the complete privileges associated with SECURITYADMIN or ACCOUNTADMIN.

An example is:

GRANT CREATE NETWORK POLICY ON ACCOUNT

TO ROLE NETWORK_SECURITY_ADMIN;

Option A is incorrect because SYSADMIN does not possess this privilege by default. Option D is incorrect because ORGADMIN focuses on organization-level administration rather than automatically inheriting account-level network-policy privileges.

Option F is incorrect because network policies are not database objects. Ownership of an individual database does not confer authority to create account-level network policies.

Following creation, policy assignment and modification remain subject to the appropriate access-control requirements. Administrators should delegate only those privileges necessary to support the intended security operations.

Official Snowflake references: CREATE NETWORK POLICY; Controlling Network Traffic with Network Policies; Access Control Privileges.

Question 9

Question Type: MultipleChoice

In which scenario does using an external table simplify a data pipeline?

Options:

A- When accessing a Snowflake table from a relational database
B- When accessing a Snowflake table from an external database within the same region
C- When continuously writing data from a Snowflake table to external storage
D- When accessing a Snowflake table that references data files located in cloud storage

Answer:

D

Explanation:

Snowflake external tables provide a relational interface for querying structured and semi-structured data files stored outside Snowflake, typically in cloud object storage such as Amazon S3, Azure Blob Storage, or Google Cloud Storage.

An external table references files through an external stage. Snowflake maintains file-level metadata and exposes the external data through a table-like abstraction, allowing SQL queries without requiring an initial COPY INTO operation to populate an internal Snowflake table.

This simplifies data pipelines when source files must remain in cloud storage, particularly for data lakes, historical archives, and externally managed datasets.

Option A describes accessing Snowflake from another relational database, which is not the purpose of an external table.

Option B similarly concerns external database connectivity rather than querying files stored in an external stage.

Option C describes exporting data from Snowflake to external storage, a process generally implemented through unloading operations such as COPY INTO a storage location.

Option D correctly identifies the external-table use case.

External tables are read-only and do not support conventional INSERT, UPDATE, or DELETE operations. Query performance may also be lower than native Snowflake tables because the underlying data remains externally stored.

Official Snowflake references: Introduction to External Tables; Creating External Tables; Querying External Tables.

Question 10

Question Type: MultipleChoice

A user has enrolled in multi-factor authentication (MFA) to connect to Snowflake. The user tells the Snowflake Administrator that they lost their mobile phone the prior evening.

Which action should the Administrator take to enable the user to sign in to the system without revoking the user's MFA enrollment?

Options:

A- Alter the user and set MINS_TO_BYPASS_MFA to a value that will disable MFA long enough for the user to log in.
B- Alter the user and set DISABLE_MFA to true, which will suspend the MFA requirement for 24 hours.
C- Instruct the user to connect to Snowflake using SnowSQL, which does not support MFA authentication.
D- Instruct the user to append the normal URL with /?mode=mfa_bypass & code= to log on.

Answer:

A

Explanation:

Snowflake provides temporary MFA bypass functionality to recover user access when an enrolled authentication device becomes unavailable. The appropriate administrative action is to configure the MINS_TO_BYPASS_MFA property on the affected user.

For example:

ALTER USER user1

SET MINS_TO_BYPASS_MFA = 30;

This command permits the user to authenticate temporarily without completing MFA for the configured 30-minute period. Importantly, the operation does not revoke the user's existing MFA enrollment or permanently disable the authentication requirement.

Option B is unsuitable because DISABLE_MFA is a Boolean setting and does not automatically provide a 24-hour expiration period. It is not equivalent to a controlled, time-limited bypass.

Option C is incorrect because Snowflake supports MFA authentication through SnowSQL and other supported clients. Changing the connection client does not inherently eliminate MFA enforcement.

Option D proposes an unsupported URL-based authentication bypass.

From a security-administration perspective, the administrator should first validate the user's identity through an approved recovery procedure, establish the shortest practical bypass interval, and guide the user through registration of a replacement MFA method.

Current Snowflake documentation also describes ALTER USER ... ENROLL MFA as a recovery mechanism for users who must register a new authenticator.

Official Snowflake references: Multi-factor Authentication; Recovering a User Who Is Locked Out; ALTER USER.

To Get Premium Files for ADA-C02 Visit
https://www.p2pexams.com/products/ada-c02

For More Free Questions Visit
https://www.p2pexams.com/snowflake/pdf/ada-c02

