



Up-to-date Practice Test with Latest Questions and Answers covering latest syllabus and topics of the exam. Makes you ready to face actual exam.



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



killexams.com

SnowFlake

ARA-C01

SnowPro Advanced Architect Certification

ORDER FULL VERSION

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



Question: 1326

In a Snowflake account, a fact table experiences daily UPDATE statements affecting 10% of its 1 TB active storage. With Enterprise Edition and default retention, calculate the approximate additional monthly storage cost attributable to one week's worth of changes entering Fail-safe (assume average pricing of \$23/TB/month and compression factors are neutral for estimation).

- A. Approximately \$16 per month
- B. Approximately \$53 per month
- C. Approximately \$161 per month
- D. Approximately \$322 per month

Answer: C

Explanation: Daily changes create ~100 GB of new micro-partitions (10% of 1 TB). Over 7 days, this accumulates ~700 GB historical data. When Time Travel expires, this volume enters the 7-day Fail-safe, incurring costs for the full week at ~\$23/TB/month, equating to roughly $(0.7 \text{ TB} * \$23) \approx \16 per week or ~\$161 over a month when considering ongoing churn patterns. Lower options underestimate the weekly accumulation impact.

Question: 1327

External OAuth token includes custom claim 'department'. To map sessions to departmental roles dynamically, use:

- A. Policy on token validation
- B. Custom scope mapping to roles
- C. EXTERNAL_OAUTH_ANY_ROLE_MODE with claim array
- D. EXTERNAL_OAUTH_TOKEN_USER_MAPPING_CLAIM = 'department' for role

Answer: B

Explanation: Scopes in tokens map directly to Snowflake roles; custom claims require standard scope usage for role assignment.

Question: 1328

A query with WHERE date_col BETWEEN '2025-01-01' AND '2025-01-31' AND active_flag = TRUE on a well-clustered table on date_col scans 10% of micro-partitions despite flag filter. Why limited additional pruning on active_flag?

- A. Pruning requires clustering on both columns
- B. Active_flag not in clustering key prevents any pruning
- C. BETWEEN predicates disable constant pruning
- D. Flag is boolean (low cardinality); metadata overlaps extensively

Answer: D

Explanation: Low-cardinality columns like booleans have values in nearly all micro-partitions, yielding high overlaps and minimal pruning even with constants. Clustering improves high/medium-cardinality columns for effective skipping. Multi-column clustering could help if flag correlates, but low cardinality limits gains.

Question: 1329

Scenario: ORGADMIN creates account; sets ADMIN_NAME='admin' PASSWORD='pass'. Admin logs in, creates users. Best practice for ACCOUNTADMIN? (Select one)

- A. Use ORGADMIN directly
- B. Grant to two non-default users
- C. Set as default
- D. Single admin

Answer: B

Explanation: Assign ACCOUNTADMIN to at least two, never default; use custom lower roles daily.

Question: 1330

A company is looking to share its marketing data with external partners while ensuring that no sensitive information is exposed. What is the most effective way to achieve this using Snowflake's features?

- A. Share the entire dataset with external partners and rely on them to handle data security.
- B. Use a reader account to limit access to the marketing data.
- C. Create secure views that filter out sensitive information before sharing.
- D. Replicate the marketing data to a separate account for sharing.

Answer: C

Explanation: Creating secure views that filter out sensitive information allows the company to share only the necessary marketing data with external partners, effectively protecting sensitive information while

facilitating collaboration.

Question: 1331

A company notices that their query performance varies significantly between similar queries. They suspect that workload management might be an issue. What feature should they investigate?

- A. Auto-resume settings
- B. Warehouse size adjustments
- C. Concurrency scaling
- D. Query acceleration service

Answer: D

Explanation: Investigating the query acceleration service can help the company understand how resource allocation is managed and whether enabling it can provide more consistent performance across similar queries.

Question: 1332

A large enterprise wants to enforce strict separation of duties in Snowflake. The ACCOUNTADMIN role must not be used for daily operations, and the RBAC design should ensure that: (1) only one role is responsible for managing users and roles, (2) another role manages objects and warehouses, and (3) a third role is used for data access. Which Snowflake role design best meets these requirements following Snowflake's recommended practices?

- A. Use SECURITYADMIN to manage users and roles, SYSADMIN to manage object and warehouse creation, and create custom data access roles granted to end users
- B. Use ACCOUNTADMIN for all management tasks, and grant it to all senior DBAs
- C. Use SYSADMIN for all tasks and grant users direct access to SYSADMIN, using custom roles only for BI users
- D. Create a single custom role that has USERADMIN, SECURITYADMIN, and SYSADMIN privileges and assign it to the platform team

Answer: A

Explanation: Snowflake recommends that the ACCOUNTADMIN role be reserved for account-level tasks such as billing and global configuration, while daily operations are handled by specialized roles to enforce separation of duties. The SECURITYADMIN role is designed to administer users and roles, including granting and revoking privileges, making it the appropriate role for identity and access management. The SYSADMIN role is intended to own warehouses, databases, and schemas, making it the correct role for managing the creation and lifecycle of objects and compute resources. Data access should be mediated by custom roles that are granted object privileges and then assigned to end users, thereby providing finer-

grained, business-oriented access control. Consolidating all powerful privileges into a single role, or using ACCOUNTADMIN or SYSADMIN for everything, violates the principle of least privilege and weakens governance.

Question: 1333

A data engineer is tasked with generating a report that aggregates sales data from multiple regions. They need to calculate the total sales and average sales per region for the last quarter. Which SQL query correctly implements this requirement, assuming the sales data is stored in a table named `sales_data` with columns `region`, `sales`, and `sale_date`?

A.

```
```sql
SELECT region, SUM(sales) AS total_sales, AVG(sales) AS average_sales
FROM sales_data
WHERE sale_date >= DATEADD(quarter, -1, CURRENT_DATE)
GROUP BY region;
```
```

B.

```
```sql
SELECT region, COUNT(sales) AS total_sales, AVG(sales) AS average_sales
FROM sales_data
WHERE sale_date >= DATEADD(month, -3, CURRENT_DATE)
GROUP BY region;
```
```

C.

```
```sql
SELECT region, SUM(sales) AS average_sales, COUNT(sales) AS total_sales
FROM sales_data
WHERE sale_date >= DATEADD(quarter, -1, CURRENT_DATE)
GROUP BY region;
```
```

D.

```
```sql
SELECT region, AVG(sales) AS total_sales, SUM(sales) AS average_sales
FROM sales_data
WHERE sale_date >= DATEADD(month, -3, CURRENT_DATE)
GROUP BY region;
```
```

Answer: A

Explanation: The correct query aggregates sales data by region for the last quarter, using `SUM` to calculate total sales and `AVG` for average sales. The `WHERE` clause filters the records to only include

those from the last quarter.

Question: 1334

(Select one) Optimize ELT cost: 100 TB/month from Kafka via Spark connector writing Parquet stages, Snowpipe loads, matview refreshes. Cuts 50% compute?

- A. All above
- B. Switch to Kafka Snowpipe connector direct, auto-cluster stages
- C. Use dbt incremental with unique_key + is_incremental() block
- D. Spark write optimized Parquet (zstd, partition by date), cluster load table

Answer: A

Explanation: Kafka connector bypasses Spark for native ingestion. dbt incremental skips full reprocess. Optimized Parquet + clustering prunes loads.

Question: 1335

A data lake ingestion project requires a service application to load files from cloud storage into Snowflake using a non-interactive mechanism. The security team states that: (1) no static password should be used, (2) credentials must be rotated periodically without changing the application code, and (3) the application must not rely on any user's personal SSO session. Which Snowflake authentication method best satisfies these constraints?

- A. Username and password authentication stored as environment variables inside the application container
- B. SAML SSO authentication using a technical user login initiated via a browserless client
- C. Key-pair authentication with a dedicated service user whose private key is managed and rotated by a secrets manager
- D. Native Snowflake MFA enforced for the service account with a fixed device token

Answer: C

Explanation: Key-pair authentication allows a Snowflake user to authenticate using a public/private key pair instead of a password, eliminating the need for static passwords while still enabling strongly authenticated, non-interactive access for services. Storing and rotating the private key in a secrets manager supports the requirement for periodic rotation without changing the application code, as the application retrieves the current key dynamically from the central store. Because key-pair authentication does not rely on a browser-based SSO flow or a human's session, it cleanly separates the service account from personal SSO. SAML SSO flows generally assume user interaction and are not ideal for non-interactive service access. Storing a static password as an environment variable directly contradicts the no-static-password requirement, and enforcing MFA for an automated process is impractical and often unsupported in headless environments.

Question: 1336

An enterprise streaming platform uses the Snowflake Kafka Connector to ingest from 100 topics with varying schemas, requiring schema evolution and data quality checks.

Configuration considerations include:

- A. Enable Snowflake's native schema detection per topic
- B. Use private key authentication for connector security
- C. One pipe per topic with separate file formats
- D. Single connector instance managing multiple topics via include list

Answer: D

Explanation: The connector supports multiple topics per instance, using Schema Registry for evolution across topics efficiently. Separate pipes per topic scale poorly; native detection is supplemented by registry; key-based auth enhances security.

Question: 1337

JavaScript procedure calling external API via HTTP. How?

- A. External access integration
- B. Use external functions
- C. Not possible directly
- D. Python procedure with requests

Answer: A

Explanation: Procedures support outbound HTTP with external access integrations.

Question: 1338

A multi-cluster warehouse (Large, min=3, max=8, Standard policy) supports critical production dashboards. To ensure no queuing during unpredictable spikes while controlling runaway costs.

- A. Set `min_cluster_count=1`
- B. Change to Economy policy
- C. Keep Standard and monitor for adjustments
- D. Switch to Maximized mode

Answer: C

Explanation: Standard policy aggressively scales to prevent queuing by starting clusters immediately upon load detection, ideal for critical dashboards requiring responsiveness. Maximized runs all clusters constantly, wasting credits. Economy risks queuing. Monitoring allows fine-tuning `max_cluster_count` based on historical utilization.

Question: 1339

Scenario: A global company has Snowflake accounts in AWS (primary) and Azure (DR). During a planned failover test, promoting the Azure replica fails due to object dependencies. What must be included in the replication group?

- A. User sessions and query history
- B. All databases containing objects referenced by shared views
- C. Only the primary database
- D. Virtual warehouses and roles

Answer: B

Explanation: For cross-object dependencies like views spanning databases or shares, all related databases must be in the same replication/failover group to ensure consistent promotion of secondaries to primary. Warehouses, roles, and sessions are account-level and replicate separately or not required for data failover. Single database suffices only for isolated objects.

Question: 1340

A resource monitor is set with `FREQUENCY = WEEKLY`, `START_TIMESTAMP = '2026-01-01'`, and `CREDIT_QUOTA = 5000`. The monitor has thresholds for 75% notify and 100% suspend. If 4000 credits are consumed by mid-week, then an additional 2000 early next week, what sequence of actions occurs?

- A. Notify mid-this week; quota resets at week end, new consumption starts fresh
- B. No notify until 3750; suspend after 5000 cumulative
- C. Suspend immediately after 5000 total, ignoring weekly reset
- D. Notify at 4000 this week; suspend early next week after 1000 more

Answer: A

Explanation: With `FREQUENCY = WEEKLY`, credit usage resets to zero at the defined interval boundary. Consumption reaching 4000 triggers notify (80% > 75%). At week reset, usage resets, so additional consumption starts anew against the 5000 quota. Actions are evaluated per interval independently, preventing carry-over enforcement from prior periods.

Question: 1341

Advanced row policy uses multiple mapping tables with joins for hierarchical entitlements.

- A. Performance issues unavoidable.
- B. Not allowed; single table.
- C. Limit subqueries; use functions.
- D. Supported in policy body with JOINS.

Answer: D

Explanation: Policy bodies support complex SQL including multi-table joins and subqueries, though optimization via clustering and memoizable functions recommended for performance.

Question: 1342

Poor clustering most directly impacts:

- A. Partition pruning ratio
- B. Compression ratio
- C. Warehouse resume time
- D. Storage costs

Answer: A

Explanation: Increases scanned partitions for same predicates.

Question: 1343

In a cross-cloud failover scenario from GCP to AWS, post-failover queries experience initial slow metadata resolution.

- A. Pre-warm metadata cache via dummy queries
- B. Replication automatically synchronizes metadata, initial queries populate cache normally
- C. Increase cloud services credits quota
- D. Use larger warehouses for metadata

Answer: B

Explanation: Cloud services layer metadata (including pruning stats) replicates with databases. Post-failover, cache rebuilds on first access naturally, with minimal delay. No pre-warming needed, quotas unrelated, warehouses not for metadata.

Question: 1344

Architect designs global data mesh. Domains share data cross-cloud. Key Snowgrid benefit? (Select one)

- A. Zero-ETL live sharing via replication
- B. Built-in data catalog
- C. Serverless compute sharing
- D. Automatic query routing

Answer: A

Explanation: Snowgrid enables live, governed data access across boundaries without ETL.

Question: 1345

A privacy team requires that personal data subject to GDPR can be deleted within 30 days of a valid erasure request, including in Time Travel. The Snowflake account currently uses the maximum Time Travel retention for key databases. Which configuration change is necessary to support this requirement?

- A. Reduce Time Travel retention for relevant databases or tables to align with the 30-day deletion SLA and adjust operational processes accordingly
- B. Enable Fail-safe for all data indefinitely
- C. Increase Time Travel to the maximum retention window for audit purposes
- D. Disable all logical deletion mechanisms and retain only physical backups

Answer: A

Explanation: Time Travel retains historical data for a configurable window; to meet a 30-day deletion SLA, retention for GDPR-covered data should be tuned so that no Time Travel copy persists beyond the required period, and processes adjusted to honor erasure requests. Fail-safe is intended for disaster recovery and extends retention beyond Time Travel but should be considered carefully for GDPR. Increasing retention conflicts with erasure requirements.

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including Exam Questions, 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.



Practice Exam Questions Based on Current Exam Objectives

Killexams.com provides practice exam questions aligned with the latest official exam objectives and latest syllabus. Our content is reviewed and updated regularly to reflect recent changes announced by certification vendors. By studying these practice questions, candidates will cover the structure, difficulty level, and topics of the actual exam, helping them prepare more effectively and efficiently.

Comprehensive Practice Exams (PDF Format)

Killexams.com offers multiple-choice questions (MCQs) in easy-to-read PDF format, covering all major domains of the exam. Each PDF contains a structured collection of practice questions and verified answers designed to support focused study. These MCQs help candidates reinforce key concepts, identify knowledge gaps, and improve exam readiness through consistent practice.

Realistic Practice Tests (Online Test Engine & Desktop Test Engine)

To support hands-on preparation, Killexams.com provides practice tests through both an Online Test Engine and a Desktop Test Engine. These tools are designed to simulate a real exam environment, allowing candidates to practice under exam-like conditions, with latest syllabus and topics of the exam. Performance tracking, test history, and result analysis help users evaluate their progress and focus on areas that need improvement.

Risk-Free Purchase Policy

Killexams.com follows a transparent and customer-friendly purchase policy. If users are not satisfied with the study materials, they may request assistance or a refund in accordance with our published terms and conditions. This policy reflects our commitment to customer satisfaction, fairness, and confidence in our preparation resources.

Regularly Updated Content

Our practice question bank is reviewed and updated on an ongoing basis to stay aligned with the latest exam outlines and vendor updates. This ensures candidates are studying up-to-date, relevant material, and preparing with content that reflects current exam expectations, helping them stay confident and well-prepared.