PRACTICE QUESTIONS

Snowflake

ADS-B01

SnowPro Specialty: AI Governance & Security

PUBLISHED BY aiexamcerts

EDITION	Free Set 2
QUESTIONS	10
ISSUED	November 8, 2026
EXAM	Snowflake ADS-B01, SnowPro Specialty: AI Governance & Security

Shared by **aiexamcerts** on aiexamcerts

10 practice questions for the Snowflake ADS-B01 exam, SnowPro Specialty: AI Governance & Security, each followed by its answer and a short explanation. More free sets and the full practice exam: **aiexamcerts.com/snowflake/ads-b01**. Spotted a mistake? Email support@aiexamcerts.com with the exam code ADS-B01 and the question number.

---

1. An account review finds the SNOWFLAKE.CORTEX_USER database role still granted to PUBLIC. What is the accurate assessment of this default?

A Every user in the account can call the Cortex AI SQL functions until it is revoked

B Only users with the ACCOUNTADMIN role can call the functions regardless of this grant

C The grant only permits viewing function documentation, not calling the functions

D The grant lets Snowflake use the account's prompts to train models for other customers

Answer: A

By default CORTEX_USER is granted to PUBLIC, so any user can invoke the functions. Security-conscious deployments revoke it from PUBLIC and grant it to specific roles. Snowflake does not use customer prompts to train models for others, and the grant is not limited to ACCOUNTADMIN.

---

2. An analyst sees clear text in an AI-generated summary table although the source columns are masked. Which TWO explanations are plausible? (Choose two.)

A Masking policies expire after 24 hours

B The table was created from a query run under a role that the policy allows to see clear text

C Cortex functions always disable masking in output tables

D A policy condition lists a role the analyst holds

E Time Travel disables masking policies

Answer: B, D

Derived tables store whatever the creating query returned, so a privileged role can persist clear text, and permissive policy logic can also reveal data. Policies do not expire and functions do not disable them.

---

3. A schema created WITH _____ ACCESS lets only the schema owner or a role with MANAGE GRANTS issue grants on its objects.

Answer: MANAGED

Managed access schemas centralize privilege management so object owners cannot grant on their own.

---

4. Which TWO statements about Trust Center are correct? (Choose two.)

A It automatically remediates every finding without administrator action

B It scans table contents for toxic LLM output

C Scanner packages, such as the CIS Benchmarks package, group related security checks

D It replaces ACCOUNT_USAGE views for auditing

E Access to view and manage findings is governed by Trust Center application roles

Answer: C, E

Trust Center organizes checks into scanner packages and gates access through application roles. It reports rather than auto-fixes, and it is complementary to `ACCOUNT_USAGE` rather than a replacement.

---

5. An investigator needs to know which users read a specific column through an AI-triggered query. Which `ACCOUNT_USAGE` data best answers this?

A The `ACCESS_HISTORY` view, using its accessed-objects columns for the table and column

B The `WAREHOUSE_METERING_HISTORY` view

C The `STORAGE_USAGE` view

D The `LOGIN_HISTORY` view

Answer: A

`ACCESS_HISTORY` records, per query, which objects and columns were read, and by which user. Metering and storage views report credits and bytes, and login history records authentication only.

---

6. Which TWO measures reduce the impact of prompt injection in a RAG application built on Cortex? (Choose two.)

A Treat retrieved document text as untrusted data rather than instructions

B Raise the temperature so the model varies its answers

C Run the application under a role limited to the objects it truly needs

D Run all queries as ACCOUNTADMIN to avoid permission errors

E Disable query tags so injected text is not logged

Answer: A, C

Retrieved content may carry hidden instructions, so it must be handled as data, and least privilege bounds what a hijacked model can reach. Temperature does not block injection, ACCOUNTADMIN maximizes damage, and disabling tags removes auditing value.

---

7. Which TWO actions follow least privilege for a service user that runs an AI workload? (Choose two.)

A Create a dedicated custom role holding only the required object privileges

B Grant OWNERSHIP of the schema so grants cannot fail

C Reuse one shared credential across every application

D Grant the SECURITYADMIN role for convenience

E Set the user's default role to that custom role, not ACCOUNTADMIN

Answer: A, E

A narrow dedicated role and a restricted default role bound what a compromised service can do. Ownership, shared credentials and SECURITYADMIN each expand privilege or blur accountability.

---

8. Masking policies change what a query returns at run time without altering the stored data.

A True

B False

Answer: A (True)

The underlying values remain unchanged; the policy decides per query what the role sees, which is why role logic must be verified.

---

9. SNOWFLAKE.CORTEX_USER is not an account role but a database _____ that is granted with GRANT DATABASE ROLE.

Answer: role

`CORTEX_USER` is a database role in the SNOWFLAKE database, granted to account roles.

---

10. Which `ACCOUNT_USAGE` view reports credit consumption attributed to Cortex AI functions for cost and usage review?

A LOGIN_HISTORY, which records user sign-in attempts

B CORTEX_AI_FUNCTIONS_USAGE_HISTORY

C ROW_ACCESS_POLICIES, which lists the policies defined

D SESSIONS

Answer: B

`CORTEX_AI_FUNCTIONS_USAGE_HISTORY` reports the tokens and credits billed for each AI function call, with the model, user and roles, so unusual AI consumption can be spotted. The older `CORTEX_FUNCTIONS_USAGE_HISTORY` view is no longer updated. Login, policy definition and session views do not record Cortex credits.