



---

**TRUSTED, ACCURATE AND  
RELIABLE!**

---

**The most comprehensive IT certification  
preparation materials in the industry!**

All rights reserved. No part of this document may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other non-commercial uses permitted by copyright law. Unauthorized copying, reselling, or distribution of this document is strictly prohibited and may result in legal action.

<https://www.virtulearner.com>  
[support@virtulearner.com](mailto:support@virtulearner.com)

**Snowflake**

**ADA-C01**

SnowPro Advanced:

Administrator Certification

Exam

**QUESTION: 1**

When a role is dropped, which role inherits ownership of objects owned by the dropped role?

- A. The SYSADMIN role
- B. The role above the dropped role in the RBAC hierarchy
- C. The role executing the command
- D. The SECURITYADMIN role

**Answer(s):** B

**Explanation:**

According to the Snowflake documentation<sup>1</sup>, when a role is dropped, ownership of all objects owned by the dropped role is transferred to the role that is directly above the dropped role in the role hierarchy. This is to ensure that there is always a single owner for each object in the system.

1: Drop Role | Snowflake Documentation

**QUESTION: 2**

Company A uses Snowflake to manage audio files of call recordings. Company A hired Company B, who also uses Snowflake, to transcribe the audio files for further analysis.

Company A's Administrator created a share.

What object should be added to the share to allow Company B access to the files?

- A. A secure view with a column for file URLs.
- B. A secure view with a column for pre-signed URLs.
- C. A secure view with a column for METADATA\$FILENAME.
- D. A secure view with a column for the stage name and a column for the file path.

**Answer(s):** B

**Explanation:**

According to the Snowflake documentation<sup>1</sup>, pre-signed URLs are required to access external files in a share. A secure view can be used to generate pre-signed URLs for the audio files stored in an external stage and expose them to the consumer account. Option A is incorrect because file URLs alone are not sufficient to access external files in a share. Option C is incorrect because METADATA\$FILENAME only returns the file name, not the full path or URL. Option D is incorrect because the stage name and file path are not enough to generate pre-signed URLs.

**QUESTION: 3**

A retailer uses a TRANSACTIONS table (100M rows, 1.2 TB) that has been clustered by the STORE\_ID column (varchar(50)). The vast majority of analyses on this table are grouped by STORE\_ID to look at store performance.

There are 1000 stores operated by the retailer but most sales come from only 20 stores. The

Administrator notes that most queries are currently experiencing poor pruning, with large amounts of bytes processed by even simple queries.

Why is this occurring?

- A. The STORE\_ID should be numeric.
- B. The table is not big enough to take advantage of the clustering key.
- C. Sales across stores are not uniformly distributed.
- D. The cardinality of the stores to transaction count ratio is too low to use the STORE\_ID as a clustering key.

**Answer(s): C**

**Explanation:**

According to the Snowflake documentation<sup>1</sup>, clustering keys are most effective when the data is evenly distributed across the key values. If the data is skewed, such as in this case where most sales come from only 20 stores out of 1000, then the micro-partitions will not be well-clustered and the pruning will be poor. This means that more bytes will be scanned by queries, even if they filter by STORE\_ID. Option A is incorrect because the data type of the clustering key does not affect the pruning. Option B is incorrect because the table is large enough to benefit from clustering, if the data was more balanced. Option D is incorrect because the cardinality of the clustering key is not relevant for pruning, as long as the key values are distinct.

1: Considerations for Choosing Clustering for a Table | Snowflake Documentation

**QUESTION: 4**

A team is provisioning new lower environments from the production database using cloning. All production objects and references reside in the database, and do not have external references.

What set of object references needs to be re-pointed before granting access for usage?

- A. Sequences, views, and secure views
- B. Sequences, views, secure views, and materialized views
- C. Sequences, storage integrations, views, secure views, and materialized views
- D. There are no object references that need to be re-pointed

**Answer(s): C**

**Explanation:**

According to the Snowflake documentation<sup>1</sup>, when an object in a schema is cloned, any future grants defined for this object type in the schema are applied to the cloned object unless the COPY GRANTS option is specified in the CREATE statement for the clone operation. However, some objects may still reference the source object or external objects after cloning, which may cause issues with access or functionality. These objects include:

- Sequences: If a table column references a sequence that generates default values, the cloned table may reference the source or cloned sequence, depending on where the sequence is defined. To avoid conflicts, the sequence reference should be re-pointed to the desired sequence using the ALTER TABLE command<sup>2</sup>.
- Storage integrations: If a stage or a table references a storage integration, the cloned object may still reference the source storage integration, which may not be accessible or valid in the