


SNOWPRO ADVANCED ARCHITECT PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://snowproadvancedarchitect.examzify.com>



Copyright © 2026 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain accurate, complete, and timely information about this product from reliable sources.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	5
Answers	8
Explanations	10
Next Steps	15

SAMPLE

Introduction

Preparing for a certification exam can feel overwhelming, but with the right tools, it becomes an opportunity to build confidence, sharpen your skills, and move one step closer to your goals. At Examzify, we believe that effective exam preparation isn't just about memorization, it's about understanding the material, identifying knowledge gaps, and building the test-taking strategies that lead to success.

This guide was designed to help you do exactly that.

Whether you're preparing for a licensing exam, professional certification, or entry-level qualification, this book offers structured practice to reinforce key concepts. You'll find a wide range of multiple-choice questions, each followed by clear explanations to help you understand not just the right answer, but why it's correct.

The content in this guide is based on real-world exam objectives and aligned with the types of questions and topics commonly found on official tests. It's ideal for learners who want to:

- Practice answering questions under realistic conditions,
- Improve accuracy and speed,
- Review explanations to strengthen weak areas, and
- Approach the exam with greater confidence.

We recommend using this book not as a stand-alone study tool, but alongside other resources like flashcards, textbooks, or hands-on training. For best results, we recommend working through each question, reflecting on the explanation provided, and revisiting the topics that challenge you most.

Remember: successful test preparation isn't about getting every question right the first time, it's about learning from your mistakes and improving over time. Stay focused, trust the process, and know that every page you turn brings you closer to success.

Let's begin.

How to Use This Guide

This guide is designed to help you study more effectively and approach your exam with confidence. Whether you're reviewing for the first time or doing a final refresh, here's how to get the most out of your Examzify study guide:

1. Start with a Diagnostic Review

Skim through the questions to get a sense of what you know and what you need to focus on. Your goal is to identify knowledge gaps early.

2. Study in Short, Focused Sessions

Break your study time into manageable blocks (e.g. 30 – 45 minutes). Review a handful of questions, reflect on the explanations.

3. Learn from the Explanations

After answering a question, always read the explanation, even if you got it right. It reinforces key points, corrects misunderstandings, and teaches subtle distinctions between similar answers.

4. Track Your Progress

Use bookmarks or notes (if reading digitally) to mark difficult questions. Revisit these regularly and track improvements over time.

5. Simulate the Real Exam

Once you're comfortable, try taking a full set of questions without pausing. Set a timer and simulate test-day conditions to build confidence and time management skills.

6. Repeat and Review

Don't just study once, repetition builds retention. Re-attempt questions after a few days and revisit explanations to reinforce learning. Pair this guide with other Examzify tools like flashcards, and digital practice tests to strengthen your preparation across formats.

There's no single right way to study, but consistent, thoughtful effort always wins. Use this guide flexibly, adapt the tips above to fit your pace and learning style. You've got this!

Questions

SAMPLE

1. What is the result of the following query on the table null_count_test?
`select count(*), count(nct.*), count(col1), count(distinct col1), count(distinct col1, col2), approx_count_distinct(*) from null_count_test nct;`
 - A. Syntax error: Invalid function approx_count_distinct(*)
 - B. 8, 5, 6, 6, 6, 5
 - C. 9, 5, 6, 6, 6, 7
 - D. 9, 5, 6, 5, 5, 5
2. Which of the following is not a type of Snowflake product release?
 - A. Full
 - B. Patch release
 - C. Behaviour change release
 - D. Hotfix release
3. Why is it not recommended to use `SELECT *` in the definition of a materialized view?
 - A. Because it's always better to specify the columns you want to show in the materialized view
 - B. It guarantees real-time update of new base columns
 - C. If new columns are added to the base table, those new columns are not automatically propagated to the materialized views
 - D. It prevents any confusion about the underlying schema
4. When resizing a multi-cluster warehouse, which clusters are resized?
 - A. Clusters that started after the warehouse is resized
 - B. All of them
 - C. Clusters that are running
 - D. None
5. How many days of query history can you view in the History tab?
 - A. 7 days
 - B. 14 days
 - C. 30 days
 - D. 60 days

- 6. Which of the following objects can be included in a Snowflake share?**
- A. Regular views
 - B. Stored Procedures
 - C. External table
 - D. Tasks
- 7. Which statement accurately describes Snowflake network policies?**
- A. They optimize query performance
 - B. They provide options for managing network configurations for the Snowflake service
 - C. They enforce data retention policies
 - D. They manage user authentication methods
- 8. Quiesce mode refers to which of the following descriptions?**
- A. Active processing
 - B. Waiting to shut down
 - C. Fully shut down
 - D. Paused but ready
- 9. Why does Snowflake advise against adding more than 3–4 columns to a cluster key?**
- A. Because data is not going to be clustered correctly
 - B. Because it tends to increase costs more than benefits
 - C. Because cluster keys must be exactly 3-4 columns
 - D. Because more columns have no effect on performance
- 10. What is the most likely cause if a large join query takes hours to complete even after increasing the warehouse size?**
- A. The query was still running before the new size of the warehouse was applied
 - B. The warehouse had auto-suspend turned on and restarted mid-query
 - C. The warehouse doesn't have enough power still
 - D. One of the column's values is significantly more than the rest of the values in the column; that's why it produces a skew in your data.

Answers

SAMPLE

1. D
2. D
3. C
4. B
5. B
6. C
7. B
8. B
9. B
10. D

SAMPLE

Explanations

SAMPLE

1. What is the result of the following query on the table `null_count_test`?
- ```
select count(*), count(nct.*), count(col1), count(distinct col1), count(distinct col1, col2), approx_count_distinct(*) from null_count_test nct;
```
- A. Syntax error: Invalid function `approx_count_distinct(*)`
  - B. 8, 5, 6, 6, 5
  - C. 9, 5, 6, 6, 6, 7
  - D. 9, 5, 6, 5, 5, 5**

Understanding how NULLs influence different aggregate functions helps you read these results correctly. The query mixes several counting methods: a row-level count, standard counts by column, and counts of distinct values, including an approximate distinct. First, `count(*)` returns the total number of rows in the table, regardless of NULLs in any column. So you get all rows, which is 9 here. Second, `count(nct.*)` uses a row expression. It counts rows where the entire row value is not NULL. If some rows have every column NULL, they don't contribute to this count. In this case, there are 4 such all-NULL rows, leaving 5 rows that have at least one non-NULL column for this measure. Third, `count(col1)` tallies non-NULL values in a single column. That column has 6 non-NULL entries. Fourth, `count(distinct col1)` counts how many distinct non-NULL values appear in that column. There are 5 distinct non-NULL values. Fifth, `count(distinct col1, col2)` counts distinct combinations of those two columns, ignoring any rows where either component is NULL. The result here is 5 distinct non-NULL pairs. Finally, `approx_count_distinct(*)` estimates the number of distinct rows (treated as a row value) using a probabilistic algorithm. In this dataset, the estimate is 5, matching the distinct-row intuition from the exact counts. Putting it together, the resulting sequence is 9, 5, 6, 5, 5, 5.

2. Which of the following is not a type of Snowflake product release?

- A. Full
- B. Patch release
- C. Behaviour change release
- D. Hotfix release**

Snowflake organizes updates into planned, labeled releases to communicate precisely what changes to expect. A full release adds new features and broader platform changes. A patch release contains bug fixes and stability improvements, usually without changing how you use existing functionality. A behaviour change release is intended for updates that may alter how features behave or interact, giving users notice to adjust as needed. A hotfix release isn't part of Snowflake's formal release taxonomy. In the broader software world, a hotfix is an urgent, isolated fix outside the normal update cadence. Snowflake handles urgent issues by incorporating fixes into the ongoing patch or maintenance updates, rather than issuing a distinct hotfix release. Because of that, hotfix release isn't a recognized Snowflake product release type.

### 3. Why is it not recommended to use `SELECT *` in the definition of a materialized view?

- A. Because it's always better to specify the columns you want to show in the materialized view
- B. It guarantees real-time update of new base columns
- C. If new columns are added to the base table, those new columns are not automatically propagated to the materialized views**
- D. It prevents any confusion about the underlying schema

*The idea being tested is how a materialized view derives its structure from its defining `SELECT`. When you use `SELECT *`, the materialized view takes all columns that exist in the base table at the moment the MV is created, and that column list becomes fixed for the MV. If you later add new columns to the base table, those new columns don't automatically become part of the materialized view. The MV's schema won't evolve with the base table—you'd have to redefining the MV with an explicit column list (or recreate it) to include or adjust columns. Explicitly listing columns gives you deliberate control over exactly what is materialized, making maintenance predictable. It avoids schema drift where changes to the base table aren't reflected in the MV, which is why the recommended practice is to avoid `SELECT *` in MV definitions.*

### 4. When resizing a multi-cluster warehouse, which clusters are resized?

- A. Clusters that started after the warehouse is resized
- B. All of them**
- C. Clusters that are running
- D. None

*In a multi-cluster warehouse, the size you set is applied uniformly to every cluster that belongs to that warehouse. A multi-cluster warehouse is designed to provide high concurrency by spinning up multiple clusters, all of which share the same size setting. When you resize, Snowflake updates the compute power of each cluster to the new size so performance remains consistent across the entire warehouse. The min/max cluster controls are about how many clusters can run, not about which clusters get resized, and resizing doesn't target only clusters that are running or those created later. So the change applies to all clusters in that warehouse.*

### 5. How many days of query history can you view in the History tab?

- A. 7 days
- B. 14 days**
- C. 30 days
- D. 60 days

*The History tab is designed to show a practical window of recent activity, specifically the last 14 days. This two-week view gives you quick access to rising issues or long-running queries without overwhelming you with too much data. In this window you can see each query's start time, duration, status, user, and the SQL text, which helps you diagnose failures or performance problems in the near past. If you need access to longer-term history beyond 14 days, Snowflake provides longer-retention options through `ACCOUNT_USAGE` views, such as `QUERY_HISTORY`, which retain data for a much longer period (up to about a year in many setups).*

## 6. Which of the following objects can be included in a Snowflake share?

- A. Regular views
- B. Stored Procedures
- C. External table**
- D. Tasks

Snowflake data sharing publishes data objects to other accounts, exposing objects that actually provide data rather than executable code. An external table fits this model because it represents data located outside Snowflake (in external storage) that you can surface through the share along with the necessary access metadata. This lets the consumer query the external data via Snowflake without duplicating it. The other items are not data-access objects published by a share: stored procedures are executable code, tasks are scheduled jobs, and regular views depend on underlying data within your account and aren't standalone data sources exposed through a share.

## 7. Which statement accurately describes Snowflake network policies?

- A. They optimize query performance
- B. They provide options for managing network configurations for the Snowflake service**
- C. They enforce data retention policies
- D. They manage user authentication methods

Snowflake network policies control where and how users can connect to Snowflake by defining allowed networks and connectivity rules. This means you can specify IP addresses or ranges that are permitted, enforce secure transport, and attach these rules to specific users or roles. The reason this description fits best is that network policies are about the network access layer for the Snowflake service, shaping who can reach the account rather than how data is processed or stored. They aren't about query performance, which depends on warehouse sizing and query optimization; they don't set data retention policies, which govern how long data is kept; and they aren't the method for choosing authentication approaches, though they work alongside authentication to restrict login origins. In practice you create a network policy, define the allowed networks, and apply it to users; login attempts from unapproved networks are blocked, enhancing security while permitting access from trusted locations. They can also support private connectivity options to further control how traffic reaches Snowflake.

## 8. Quiesce mode refers to which of the following descriptions?

- A. Active processing
- B. Waiting to shut down**
- C. Fully shut down
- D. Paused but ready

Quiesce mode is the state where a Snowflake resource is winding down its activity. In this mode, the system stops accepting new work while allowing currently running tasks to complete, so it's effectively waiting for the workload to drain before shutdown. That's why the description "Waiting to shut down" fits best: no new processing begins, but the system hasn't fully shut down yet. This differs from active processing (where work is ongoing), fully shut down (no compute active at all), or paused but ready (could resume quickly but isn't in the shutdown-wind-down state). Use quiesce mode when you plan to decommission or maintenance a warehouse while preserving in-flight queries.

**9. Why does Snowflake advise against adding more than 3–4 columns to a cluster key?**

- A. Because data is not going to be clustered correctly
- B. Because it tends to increase costs more than benefits**
- C. Because cluster keys must be exactly 3-4 columns
- D. Because more columns have no effect on performance

*Clustering keys influence how Snowflake organizes data into micro-partitions so that queries can prune what they don't need to scan. Each additional column in the cluster key adds more criteria that must be maintained and tracked, which increases the overhead of DML operations and metadata management. The real benefit comes only when queries frequently filter on those columns; beyond a few well-chosen columns, the extra pruning you gain tends to be small while the maintenance costs grow. That mismatch between limited extra benefit and rising cost is why keeping cluster keys to around 3–4 columns is recommended. Snowflake isn't saying you must have exactly that many, and it isn't saying more columns will never help, but the typical outcome is higher costs with diminishing returns.*

**10. What is the most likely cause if a large join query takes hours to complete even after increasing the warehouse size?**

- A. The query was still running before the new size of the warehouse was applied
- B. The warehouse had auto-suspend turned on and restarted mid-query
- C. The warehouse doesn't have enough power still
- D. One of the column's values is significantly more than the rest of the values in the column; that's why it produces a skew in your data.**

*When a join processes a lot of data, uneven work distribution across compute resources can cause hours-long runtimes. If one value in a join key is Extremely common, that value ends up matching a large portion of the data, so one or a few nodes have to do most of the heavy lifting while the rest sit idle. Even larger warehouses increase parallelism, but the overall time is still dominated by that slowest worker handling the skewed key. That's why data skew in the join key is the most plausible reason for the persistent long runtime. To verify, check the distribution of the join key values in both tables and identify any values that occur far more often than the rest. If skew is present, you can address it by filtering out the outliers early, restructuring the query to reduce the amount of data joined (perhaps by aggregating first or performing the join in stages), or adjusting clustering/partitioning strategies to improve how data is pruned before the join.*

## Next Steps

Congratulations on reaching the final section of this guide. You've taken a meaningful step toward passing your certification exam and advancing your career.

As you continue preparing, remember that consistent practice, review, and self-reflection are key to success. Make time to revisit difficult topics, simulate exam conditions, and track your progress along the way.

If you need help, have suggestions, or want to share feedback, we'd love to hear from you. Reach out to our team at [hello@examzify.com](mailto:hello@examzify.com).

Or visit your dedicated course page for more study tools and resources:

<https://snowproadvancedarchitect.examzify.com>

We wish you the very best on your exam journey. You've got this!

SAMPLE