

# SnowPro Advanced Architect Practice Test (Sample)

## Study Guide



**Everything you need from our exam experts!**

**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** ..... 16

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. In the command referencing an internal stage `@xf_tuts.public.%emp_raw` from a Put command, which internal stage name is indicated by the percent sign?
  - A. SCHEMA
  - B. DATABASE
  - C. TABLE
  - D. TABLE
  
2. To optimize parallel loads, what compressed data file size is recommended by Snowflake?
  - A. 1-2 MB
  - B. 10-50 MB
  - C. 100-250 MB (or larger) in size compressed
  - D. 500-1000 MB
  
3. Which statement is true about Snowflake materialized views?
  - A. ORDER BY clause is allowed in a materialized view.
  - B. A materialized view can query only a single table.
  - C. Materialized views support UDFs.
  - D. GROUP BY keys in a materialized view can include columns not present in the SELECT list.
  
4. How does Snowflake deploy a full release to accounts?
  - A. It moves all accounts at once
  - B. It never moves accounts
  - C. It uses a staged approach
  - D. It uses a random selection
  
5. Data Load metadata expires after how many days?
  - A. 64 Days
  - B. 30 Days
  - C. 90 Days
  - D. 120 Days

- 6. What is the default clustering state for a Snowflake table with no clustering key defined?**
- A. It is automatically clustered by most recently used column.**
  - B. It uses automatic micro-partitioning with random clustering.**
  - C. It has no clustering key defined by default.**
  - D. It uses a hidden primary key for clustering.**
- 7. What is the recommended pattern for calling the Snowpipe loadHistoryScan endpoint to avoid rate limits?**
- A. Reading the last 10 minutes of history every 8 minutes.**
  - B. Reading the last 24 hours of history every minute.**
  - C. Reading the last seven days of history every hour.**
  - D. Reading the last hour of history every 15 minutes.**
- 8. How can you determine the last refresh time of a materialized view?**
- A. Check REFRESHED\_ON and BEHIND\_BY columns in the output of SHOW MATERIALIZED VIEWS**
  - B. Check CREATED\_ON in the output of SHOW MATERIALIZED VIEWS**
  - C. Query INFORMATION\_SCHEMA.MATERIALIZED\_VIEWS**
  - D. Describe MATERIALIZED\_VIEW**
- 9. Which of the following are types of Snowflake product releases?**
- A. Full**
  - B. Behaviour change release**
  - C. All of the above**
  - D. Patch release**
- 10. True or False: A single clustering key can contain one or more columns or expressions**
- A. False**
  - B. Depends on data**
  - C. True**
  - D. Cannot be**

## Answers

SAMPLE

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

SAMPLE

## **Explanations**

SAMPLE

**1. In the command referencing an internal stage @xf\_tuts.public.%emp\_raw from a Put command, which internal stage name is indicated by the percent sign?**

- A. SCHEMA**
- B. DATABASE**
- C. TABLE**
- D. TABLE**

The percent sign in a stage reference marks a per-table stage, meaning a table-level internal stage for that specific table. In @xf\_tuts.public.%emp\_raw, xf\_tuts is the database, public is the schema, and %emp\_raw designates the table's own internal stage for the table named emp\_raw. So, the internal stage indicated by the percent sign is a TABLE (table-level) stage. The database and schema parts come before the percent sign, while the percent sign itself signals the table-level scope.

**2. To optimize parallel loads, what compressed data file size is recommended by Snowflake?**

- A. 1-2 MB**
- B. 10-50 MB**
- C. 100-250 MB (or larger) in size compressed**
- D. 500-1000 MB**

Maximizing parallel loads relies on having enough independent files for Snowflake to process at the same time. Each file can be loaded by a separate loader thread, so you want enough work units to keep the compute warehouse busy without creating excessive per-file overhead. The recommended compressed file size is 100-250 MB because it strikes a balance: large enough to amortize startup and I/O costs per file, yet small enough to generate many concurrent loaders. If files are too small, overhead dominates and throughput drops; if files are too large, you get fewer parallel tasks and underutilize the warehouse. So aiming for around 100-250 MB compressed per file helps maximize parallelism and overall load performance.

### 3. Which statement is true about Snowflake materialized views?

- A. ORDER BY clause is allowed in a materialized view.
- B. A materialized view can query only a single table.**
- C. Materialized views support UDFs.
- D. GROUP BY keys in a materialized view can include columns not present in the SELECT list.

Materialized views in Snowflake are precomputed results based on a query that references a single base table. This design keeps the refresh and maintenance straightforward, which is why the defining query is restricted to one table. That makes the statement about querying only a single table the correct one. Consider why the others aren't true: an ORDER BY clause isn't meaningful in a materialized view's stored results, since Snowflake doesn't guarantee a particular storage order or retrieval order for MV data. So including ORDER BY in the MV definition isn't supported. User-defined functions aren't part of the standard MV definition behavior either, as MVs rely on deterministic, reproducible results and have restrictions around nonstandard functions. Finally, when you use GROUP BY, all non-aggregated columns typically need to appear in the SELECT list; you can't group by a column that isn't included in the output.

### 4. How does Snowflake deploy a full release to accounts?

- A. It moves all accounts at once
- B. It never moves accounts
- C. It uses a staged approach**
- D. It uses a random selection

A staged rollout approach is used when deploying a full release to Snowflake accounts. The update is rolled out to a subset of accounts first, so engineers can monitor performance, compatibility with existing workloads, and catch any issues early. If all looks good, the release is gradually extended to more accounts until everyone is updated. This strategy minimizes disruption, reduces risk of widespread outages, and allows quick rollback or fixes if unexpected problems arise. Rolling out to all accounts at once would be risky, and applying changes randomly or not moving accounts at all wouldn't provide the controlled, observable path needed for a safe, smooth release.

### 5. Data Load metadata expires after how many days?

- A. 64 Days**
- B. 30 Days
- C. 90 Days
- D. 120 Days

Data Load metadata retention defines how long Snowflake keeps details about each load, so you can review what happened during data loads (file names, load IDs, row counts, errors, etc.) without letting metadata grow without bound. The default period is 64 days, meaning after that time the load metadata is purged. This provides a practical window for troubleshooting and audits while steering storage usage to a reasonable level. If you need longer history, you can adjust the retention setting or export the metadata before it expires. Other durations like 30 days or 90/120 days aren't the standard default, so they aren't the correct value in this context.

6. What is the default clustering state for a Snowflake table with no clustering key defined?
- A. It is automatically clustered by most recently used column.
  - B. It uses automatic micro-partitioning with random clustering.
  - C. It has no clustering key defined by default.**
  - D. It uses a hidden primary key for clustering.

In Snowflake, clustering keys are optional. If you don't define a clustering key for a table, there is no clustering key associated with that table by default. Snowflake still stores data in automatic micro-partitions, but without a defined clustering key there's no explicit guidance for how data is grouped for clustering. The other options misrepresent how Snowflake handles partitioning: it doesn't cluster by the most recently used column, it doesn't use random clustering, and there's no hidden primary key used for clustering.

7. What is the recommended pattern for calling the Snowpipe loadHistoryScan endpoint to avoid rate limits?
- A. Reading the last 10 minutes of history every 8 minutes.**
  - B. Reading the last 24 hours of history every minute.
  - C. Reading the last seven days of history every hour.
  - D. Reading the last hour of history every 15 minutes.

Polling strategy should balance data freshness with rate-limit safety. The best approach uses a small, sliding window and a cadence that keeps the window overlapping with the next poll. Reading the last 10 minutes of history every 8 minutes does this: each call grabs a short slice of the most recent history, and the 8-minute cadence ensures continuous coverage with overlap. If some history arrives late, it's still captured on the next poll within a short grace period, while keeping the number of calls and the data per call moderate to avoid hitting rate limits. Other patterns are less suitable: querying a 24-hour window every minute creates excessive traffic and data volume, quickly exhausting quotas. Querying the last seven days every hour is too slow to surface timely events and risks missing late-arriving history. Querying the last hour every 15 minutes is better than the extremes but still yields more latency and larger per-call payload than a small, overlapping window with a tighter cadence.

**8. How can you determine the last refresh time of a materialized view?**

- A. Check REFRESHED\_ON and BEHIND\_BY columns in the output of SHOW MATERIALIZED VIEWS**
- B. Check CREATED\_ON in the output of SHOW MATERIALIZED VIEWS**
- C. Query INFORMATION\_SCHEMA.MATERIALIZED\_VIEWS**
- D. Describe MATERIALIZED\_VIEW**

The last refresh time is provided directly in the SHOW MATERIALIZED VIEWS output. Look at the REFRESHED\_ON timestamp—the value there is the exact time the materialized view was last refreshed. The BEHIND\_BY column is also useful because it shows how far behind the latest data the view is, giving you a sense of currentness. This approach is preferred because CREATED\_ON only tells when the object was created, not when it was last refreshed. Querying INFORMATION\_SCHEMA.MATERIALIZED\_VIEWS or DESCRIBE MATERIALIZED\_VIEW typically won't yield the last refresh timestamp or the current staleness, so they don't directly answer when the view was last refreshed.

**9. Which of the following are types of Snowflake product releases?**

- A. Full**
- B. Behaviour change release**
- C. All of the above**
- D. Patch release**

Snowflake releases come in three distinct forms: full releases, patch releases, and behavior-change releases. A full release introduces new features and broader updates across the platform. A patch release focuses on bug fixes and minor improvements without changing how things work. A behavior-change release is specifically about changes that alter how Snowflake behaves, which may require customers to adjust configurations or workloads. Because all three types exist, the option that includes all of them is the correct one. This categorization helps you plan testing and upgrade steps: explore new capabilities with full releases, stabilize systems with patch releases, and review notes for behavior changes to avoid surprises in production.

**10. True or False: A single clustering key can contain one or more columns or expressions**

**A. False**

**B. Depends on data**

**C. True**

**D. Cannot be**

Clustering keys can be composite, meaning a single clustering key can span multiple components—columns or expressions. When you define a clustering key with `CLUSTER BY`, you can include several columns (for example, `CLUSTER BY (region, order_date)`) or even expressions (such as `CLUSTER BY (DATE_TRUNC('MONTH', order_date))`). This single key then guides how Snowflake organizes data into micro-partitions to improve pruning for queries that filter on those components. That's why the statement is true: a single clustering key can contain one or more columns or expressions. The other options don't fit because the ability to compose a clustering key is not dependent on the data, nor is it an impossibility to have multiple components in a clustering key.

SAMPLE

## 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