

# DEA-C01 Practice Test Online, DEA-C01 Valid Test Questions



The advertisement is a rectangular box with a blue border. At the top, the text 'DEA-C01' is written in large, bold, blue letters. Below it, 'SnowPro' is in a smaller blue font, followed by 'Advanced' and 'Data Engineer' in the same blue font. To the right of this text is an illustration of a graduation cap on top of a stack of books, with a laptop and a smartphone below. At the bottom of the box, the text 'Certification Questions & Exams Dumps' is written in white on a blue background, and the website 'www.edurely.com' is written in blue on a white background.

What's more, part of that FreeCram DEA-C01 dumps now are free: <https://drive.google.com/open?id=12gyJUcY0mhP0ySNkHadimmMMFyBcLmdZ>

With the development of information and communications technology, we are now living in a globalized world. DEA-C01 information technology learning is correspondingly popular all over the world. Modern technology has changed the way how we live and work. When it comes to the study materials selling in the market, qualities are patchy. But our DEA-C01 test material has been recognized by multitude of customers, which possess of the top-class quality, can help you pass exam successfully. On the other hand, our DEA-C01 Latest Dumps are designed by the most experienced experts, thus it can not only teach you knowledge, but also show you the method of learning in the most brief and efficient ways.

## Snowflake DEA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"><li>• Performance Optimization: This topic assesses the ability to optimize and troubleshoot underperforming queries in Snowflake. Candidates must demonstrate knowledge in configuring optimal solutions, utilizing caching, and monitoring data pipelines. It focuses on ensuring engineers can enhance performance based on specific scenarios, crucial for Snowflake Data Engineers and Software Engineers.</li></ul>
Topic 2	<ul style="list-style-type: none"><li>• Security: The Security topic of the DEA-C01 test covers the principles of Snowflake security, including the management of system roles and data governance. It measures the ability to secure data and ensure compliance with policies, crucial for maintaining secure data environments for Snowflake Data Engineers and Software Engineers.</li></ul>
Topic 3	<ul style="list-style-type: none"><li>• Storage and Data Protection: The topic tests the implementation of data recovery features and the understanding of Snowflake's Time Travel and micro-partitions. Engineers are evaluated on their ability to create new environments through cloning and ensure data protection, highlighting essential skills for maintaining Snowflake data integrity and accessibility.</li></ul>

Topic 4	<ul style="list-style-type: none"> <li>• <b>Data Transformation:</b> The SnowPro Advanced: Data Engineer exam evaluates skills in using User-Defined Functions (UDFs), external functions, and stored procedures. It assesses the ability to handle semi-structured data and utilize Snowpark for transformations. This section ensures Snowflake engineers can effectively transform data within Snowflake environments, critical for data manipulation tasks.</li> </ul>
Topic 5	<ul style="list-style-type: none"> <li>• <b>Data Movement:</b> Snowflake Data Engineers and Software Engineers are assessed on their proficiency to load, ingest, and troubleshoot data in Snowflake. It evaluates skills in building continuous data pipelines, configuring connectors, and designing data sharing solutions.</li> </ul>

>> DEA-C01 Practice Test Online <<

## Pass Guaranteed Quiz Snowflake - The Best DEA-C01 Practice Test Online

If you want to find the best DEA-C01 study materials, the first thing you need to do is to find a bank of questions that suits you. Our DEA-C01 learning material is prepared by experts in strict accordance with the exam outline of the DEA-C01 certification exam, whose main purpose is to help students to pass the exam with the least amount of time and effort. We can claim that if you study with our DEA-C01 Practice Engine for 20 to 30 hours, then you will be sure to pass the exam.

### Snowflake SnowPro Advanced: Data Engineer Certification Exam Sample Questions (Q215-Q220):

#### NEW QUESTION # 215

Jonas, a Lead Performance Engineer, identified that some of the operation of his query which functionally remove the duplicates from huge data set is spilling the data to remote disk. How can he alleviate spilling to a remote disk for better query performance?

- A. Jonas can recommend using a large warehouse which effectively increase the available memory/local disk space for the operations.
- B. He can Process data in smaller batches to manage workload.
- C. Spilling do not have a profound effect on query performance (especially if remote disk is used for spilling).
- D. Data Sharing can be helpful to improve query performance.

**Answer: A,B**

Explanation:

Explanation

For some operations (e.g. duplicate elimination for a huge data set), the amount of memory available for the compute resources used to execute the operation might not be sufficient to hold intermediate results. As a result, the query processing engine will start spilling the data to local disk. If the local disk space is not sufficient, the spilled data is then saved to remote disks.

This spilling can have a profound effect on query performance (especially if remote disk is used for spilling).

To alleviate this, It is recommend that:

Using a larger warehouse (effectively increasing the available memory/local disk space for the operation), and/or Processing data in smaller batches.

#### NEW QUESTION # 216

```

1.+-----+
2.| SYSTEM$CLUSTERING_INFORMATION('SF_DATA', '(COL1, COL3')) |
3.|-----|
4.| { |
5.| "cluster_by_keys": "(COL1, COL3)", |
6.| "total_partition_count": 1156, |
7.| "total_constant_partition_count": 0, |
8.| "average_overlaps": 117.5484, |
9.| "average_depth": 64.0701, |
10.| "partition_depth_histogram": { |
11.| "00000": 0, |
12.| "00001": 0, |

```

```

13. | "00002" : 3, |
14. | "00003" : 3, |
15. | "00004" : 4, |
16. | "00005" : 6, |
17. | "00006" : 3, |
18. | "00007" : 5, |
19. | "00008" : 10, |
20. | "00009" : 5, |
21. | "00010" : 7, |
22. | "00011" : 6, |
23. | "00012" : 8, |
24. | "00013" : 8, |
25. | "00014" : 9, |
26. | "00015" : 8, |
27. | "00016" : 6, |
28. | "00032" : 98, |
29. | "00064" : 269, |
30. | "00128" : 698 |
31. | } |
32. | } |
33. +-----+

```

The Above example indicates that the SF\_DATA table is not well-clustered for which of following valid reasons?

- A. Most of the micro-partitions are grouped at the lower-end of the histogram, with the majority of micro-partitions having an overlap depth between 64 and 128.
- B. High average of overlap depth across micro-partitions.
- C. ALL of the above
- D. Zero (0) constant micro-partitions out of 1156 total micro-partitions.
- E. High average of overlapping micro-partitions.

**Answer: C**

#### NEW QUESTION # 217

What are Common Query Problems a Data Engineer can identified using Query Profiler?

- A. Inefficient Pruning
- B. Ineffective Data Sharing
- C. Queries Too Large to Fit in Memory
- D. "Exploding" Joins i.e Joins resulting due to a "Cartesian product"

**Answer: A,C,D**

Explanation:

Explanation

"Exploding" Joins

One of the common mistakes SQL users make is joining tables without providing a join condition (resulting in a "Cartesian product"), or providing a condition where records from one table match multiple records from another table. For such queries, the Join operator produces significantly (often by orders of magnitude) more tuples than it consumes.

This can be observed by looking at the number of records produced by a Join operator in the profile interface, and typically is also reflected in Join operator consuming a lot of time.

Queries Too Large to Fit in Memory

For some operations (e.g. duplicate elimination for a huge data set), the amount of memory available for the compute resources used to execute the operation might not be sufficient to hold intermediate results. As a result, the query processing engine will start spilling the data to local disk. If the local disk space is not sufficient, the spilled data is then saved to remote disks.

This spilling can have a profound effect on query performance (especially if remote disk is used for spilling).

Spilling statistics can be checked in Query Profile Interface.

Inefficient Pruning

Snowflake collects rich statistics on data allowing it not to read unnecessary parts of a table based on the query filters. However, for this to have an effect, the data storage order needs to be correlated with the query filter attributes.

The efficiency of pruning can be observed by comparing Partitions scanned and Partitions total statistics in the TableScan

operators. If the former is a small fraction of the latter, pruning is efficient. If not, the pruning did not have an effect. Of course, pruning can only help for queries that actually filter out a significant amount of data. If the pruning statistics do not show data reduction, but there is a Filter operator above TableScan which filters out a number of records, this might signal that a different data organization might be beneficial for this query.

### NEW QUESTION # 218

A company is building a data lake for a new analytics team. The company is using Amazon S3 for storage and Amazon Athena for query analysis. All data that is in Amazon S3 is in Apache Parquet format.

The company is running a new Oracle database as a source system in the company's data center. The company has 70 tables in the Oracle database. All the tables have primary keys.

Data can occasionally change in the source system. The company wants to ingest the tables every day into the data lake.

Which solution will meet this requirement with the LEAST effort?

- A. Create an Apache Sqoop job in Amazon EMR to read the data from the Oracle database. Configure the Sqoop job to write the data to Amazon S3 in Parquet format.
- **B. Create an AWS Database Migration Service (AWS DMS) task for ongoing replication. Set the Oracle database as the source. Set Amazon S3 as the target. Configure the task to write the data in Parquet format.**
- C. Create an Oracle database in Amazon RDS. Use AWS Database Migration Service (AWS DMS) to migrate the on-premises Oracle database to Amazon RDS. Configure triggers on the tables to invoke AWS Lambda functions to write changed records to Amazon S3 in Parquet format.
- D. Create an AWS Glue connection to the Oracle database. Create an AWS Glue bookmark job to ingest the data incrementally and to write the data to Amazon S3 in Parquet format.

**Answer: B**

Explanation:

[https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Target.S3.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html)

### NEW QUESTION # 219

In one of your created Schema, you have been required to create Internal Stages, what are the In-correct considerations you can noticed from the below options? [Select All that Apply]

- A. Table stage type is designed to store files that are staged and managed by one or more users but only loaded into a single table.
- B. A table stage is available for each table created in Snowflake.
- **C. User stages can be altered or dropped just like Table Stage.**
- D. A named internal stage type can store files that are staged and managed by one or more users and loaded into one or more tables.

**Answer: C**

Explanation:

Explanation

A stage specifies where data files are stored (i.e. "staged") so that the data in the files can be loaded into a table.

Types of Internal Stages

User Stages

Table Stages

Named Stages

By default, each user and table in Snowflake is automatically allocated an internal stage for staging data files to be loaded. In addition, you can create named internal stages.

File staging information is required during both steps in the data loading process:

You must specify an internal stage in the PUT command when uploading files to Snowflake.

You must specify the same stage in the COPY INTO <table> command when loading data into a table from the staged files.

Consider the best type of stage for specific data files. Each option provides benefits and potential drawbacks.

User Stages

Each user has a Snowflake stage allocated to them by default for storing files. This stage is a convenient option if your files will only be accessed by a single user, but need to be copied into multiple tables.

User stages have the following characteristics and limitations:

User stages are referenced using @~; e.g. use LIST @~ to list the files in a user stage.

Unlike named stages, user stages cannot be altered or dropped.

User stages do not support setting file format options. Instead, you must specify file format and copy options as part of the COPY INTO <table> command.

This option is not appropriate if

Multiple users require access to the files.

The current user does not have INSERT privileges on the tables the data will be loaded into.

Table Stages

Each table has a Snowflake stage allocated to it by default for storing files. This stage is a convenient option if your files need to be accessible to multiple users and only need to be copied into a single table.

Table stages have the following characteristics and limitations:

Table stages have the same name as the table; e.g. a table named mytable has a stage referenced as @%mytable.

Unlike named stages, table stages cannot be altered or dropped.

Table stages do not support transforming data while loading it (i.e. using a query as the source for the COPY command).

Note that a table stage is not a separate database object; rather, it is an implicit stage tied to the table itself. A table stage has no grantable privileges of its own. To stage files to a table stage, list the files, query them on the stage, or drop them, you must be the table owner (have the role with the OWNERSHIP privilege on the table).

This option is not appropriate if you need to copy the data in the files into multiple tables.

Named Stages

Named stages are database objects that provide the greatest degree of flexibility for data loading:

Users with the appropriate privileges on the stage can load data into any table.

Because the stage is a database object, the security/access rules that apply to all objects apply. The privileges to use a stage can be granted or revoked from roles. In addition, ownership of the stage can be transferred to another role.

If you plan to stage data files that will be loaded only by you, or will be loaded only into a single table, then you may prefer to simply use either your user stage or the stage for the table into which you will be loading data.

Named stages are optional but recommended when you plan regular data loads that could involve multiple users and/or tables.

## NEW QUESTION # 220

.....

Our DEA-C01 training materials have won great success in the market. Tens of thousands of the candidates are learning on our DEA-C01 practice engine. First of all, our DEA-C01 study dumps cover all related tests about computers. It will be easy for you to find your prepared learning material. If you are suspicious of our DEA-C01 Exam Questions, you can download the free demo from our official websites.

**DEA-C01 Valid Test Questions:** <https://www.freecram.com/Snowflake-certification/DEA-C01-exam-dumps.html>

- Frequent DEA-C01 Updates  DEA-C01 New Study Questions  Test DEA-C01 Pattern  Download  DEA-C01  for free by simply entering { [www.prepawayexam.com](http://www.prepawayexam.com) } website  Exam Vce DEA-C01 Free
- Get Help From Top Pdfvce DEA-C01 Exam Practice Questions  The page for free download of  DEA-C01  on > [www.pdfvce.com](http://www.pdfvce.com) < will open immediately  DEA-C01 Reliable Test Cram
- SnowPro Advanced: Data Engineer Certification Exam exam questions - DEA-C01 torrent pdf - SnowPro Advanced: Data Engineer Certification Exam actual dumps  Search for  DEA-C01  and download it for free immediately on { [www.examdiscuss.com](http://www.examdiscuss.com) }  Real DEA-C01 Torrent
- Latest DEA-C01 Test Objectives  DEA-C01 Sample Exam  Vce DEA-C01 Download  Easily obtain free download of ( DEA-C01 ) by searching on  [www.pdfvce.com](http://www.pdfvce.com)   Study DEA-C01 Demo
- Get Help From Top [www.examcollectionpass.com](http://www.examcollectionpass.com) DEA-C01 Exam Practice Questions  Immediately open  [www.examcollectionpass.com](http://www.examcollectionpass.com)  and search for  DEA-C01  to obtain a free download  New DEA-C01 Braindumps Ebook
- DEA-C01 Exam Course  DEA-C01 Exams Dumps  DEA-C01 Sample Exam  Simply search for  DEA-C01  for free download on 《 [www.pdfvce.com](http://www.pdfvce.com) 》  Test DEA-C01 Pattern
- Free PDF Quiz 2026 Snowflake DEA-C01: SnowPro Advanced: Data Engineer Certification Exam Pass-Sure Practice Test Online  Download [ DEA-C01 ] for free by simply entering { [www.examcollectionpass.com](http://www.examcollectionpass.com) } website  DEA-C01 Exam Course
- DEA-C01 Test Book  Latest DEA-C01 Test Objectives  Real DEA-C01 Torrent  Enter  [www.pdfvce.com](http://www.pdfvce.com)   and search for “DEA-C01 ” to download for free  DEA-C01 Authorized Exam Dumps
- Unparalleled DEA-C01 Practice Test Online | Easy To Study and Pass Exam at first attempt - Fantastic DEA-C01: SnowPro Advanced: Data Engineer Certification Exam  Open  [www.practicevce.com](http://www.practicevce.com)   and search for [ DEA-C01 ] to download exam materials for free  DEA-C01 Test Book
- SnowPro Advanced: Data Engineer Certification Exam exam questions - DEA-C01 torrent vce - SnowPro Advanced: Data

Engineer Certification Exam pdf dumps ☐ Search for ☐ DEA-C01 ☐ and download it for free on ☼ www.pdfvce.com  
☐☼☐ website ☐Test DEA-C01 Pattern

- Get Help From Top www.prepawayete.com DEA-C01 Exam Practice Questions ☐ Search for ✓ DEA-C01 ☐✓☐ and obtain a free download on ☐ www.prepawayete.com ☐ ☐DEA-C01 Exams Dumps
- privatebookmark.com, isaiahxynq279856.blogripley.com, icelisting.com, amiecznut437867.dekaronwiki.com, iwandctf656490.ssnblog.com, berthawawf946828.wikinewspaper.com, emiliahncm593821.nizarblog.com, zaynwjca698126.qodsblog.com, iwanreka378261.theisblog.com, socialwebconsult.com, Disposable vapes

BTW, DOWNLOAD part of FreeCram DEA-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=12gyJUcY0mhP0ySNkHadinnMMFyBcLmdZ>