

Cloudera CDP-3002

CLOUDERA DATA ENGINEER CERTIFICATION QUESTIONS &
ANSWERS

**Get Instant Access to Vital Exam Acing
Materials | Study Guide | Sample
Questions | Practice Test**

CDP-3002

[Cloudera Data Engineer \(CDP-3002\)](#)

50 Questions Exam – 55% Cut Score – Duration of 90 minutes

Table of Contents

Discover More about the CDP-3002 Certification	2
Cloudera CDP-3002 Data Engineer Certification Details:	2
CDP-3002 Syllabus:	3
Broaden Your Knowledge with Cloudera CDP-3002 Sample Questions:	4
Avail the Study Guide to Pass Cloudera CDP-3002 Data Engineer Exam:	6
Career Benefits:	7

Discover More about the CDP-3002 Certification

Are you interested in passing the Cloudera CDP-3002 exam? First discover, who benefits from the CDP-3002 certification. The CDP-3002 is suitable for a candidate if he wants to learn about Engineering. Passing the CDP-3002 exam earns you the Cloudera Data Engineer (CDP-3002) title.

While preparing for the CDP-3002 exam, many candidates struggle to get the necessary materials. But do not worry; your struggling days are over. The CDP-3002 PDF contains some of the most valuable preparation tips and the details and instant access to useful [**CDP-3002 study materials just at one click.**](#)

Cloudera CDP-3002 Data Engineer Certification Details:

Exam Name	Cloudera Data Engineer
Exam Code	CDP-3002
Exam Duration	90 minutes
Exam Questions	50
Passing Score	55%
Exam Price	\$330 (USD)
Books / Training	Preparing with Cloudera Data Engineering Advanced Spark Application Performance Tuning DENG-152 Cloudera Iceberg Fundamentals (FREE OnDemand) Data Engineering Orchestrating workflows and pipelines Apache Airflow Iceberg
Exam Registration	Cloudera
Sample Questions	Cloudera Data Engineer Certification Sample Question
Practice Exam	Cloudera Data Engineer Certification Practice Exam

CDP-3002 Syllabus:

Objective	Details	Weight
Spark	<ul style="list-style-type: none">- Fundamentals on Spark over Kubernetes- Work with DataFrames- Understand Distribute Processing- Implement Hive and Spark Integration- Understand Distributed Persistence	48%
Airflow	<ul style="list-style-type: none">- Implement incremental extraction in Apache Airflow from source system- Use Apache Airflow to schedule ETL pipelines- Use Apache Airflow to schedule quality checks- Work with DAGs	10%
Performance Tuning	<ul style="list-style-type: none">- Know Basic tools in (Spark) Performance Tuning- Understand Optimization Framework and Explain plans- Understand Inferring Schemas- Work with Improving Join Performance Leverage Caching Data for Reuse- Work with Partitioned and Bucketed Tables	22%
Deployment	<ul style="list-style-type: none">- Use the API and CLI- Work in the Data Engineering Service	10%
Iceberg	<ul style="list-style-type: none">- Understand Iceberg	10%

Broaden Your Knowledge with Cloudera CDP-3002 Sample Questions:

Question: 1

A Spark SQL query on a partitioned table scans all partitions. How can you limit the scan to only required partitions?

- a) Lower Kerberos ticket timeout
- b) Enable speculative execution
- c) Increase replication factor
- d) Add partition filter in the WHERE clause

Answer: d

Question: 2

A Spark query is taking too long. You run `df.explain("formatted")` and see multiple Exchange stages. What does this indicate?

- a) Data locality achieved
- b) Caching in memory
- c) Shuffle operations between stages
- d) UDF optimization applied

Answer: c

Question: 3

A Spark app should run driver pods in a specific Kubernetes node pool with SSDs. Which property enables this?

- a) `spark.executor.memory`
- b) `spark.kubernetes.node.selector`
- c) `spark.sql.broadcastTimeout`
- d) `spark.dynamicAllocation.initialExecutors`

Answer: b

Question: 4

You want to see executor memory usage during a job. Which Spark UI tab provides this?

- a) Storage
- b) SQL
- c) DAG
- d) Executors

Answer: d

Question: 5

A Spark job writing into Hive fails with “Permission denied” error. Ranger is enabled. What’s the first step?

- a) Check Ranger policy granting write access to target table/database
- b) Restart Hive Metastore service
- c) Lower spark.sql.shuffle.partitions
- d) Enable speculative execution

Answer: a

Question: 6

Your Spark executors fail due to insufficient memory in Kubernetes pods. Which property should be tuned?

- a) spark.sql.shuffle.partitions
- b) spark.hadoop.dfs.blocksize
- c) spark.executor.memory
- d) spark.driver.extraJavaOptions

Answer: c

Question: 7

You need to add a new column email to an Iceberg table without rewriting existing data. Which statement is correct?

- a) You must rebuild the entire table
- b) This requires Ranger catalog integration
- c) Only Hive supports altering schemas
- d) Iceberg supports column addition via schema evolution

Answer: d

Question: 8

A DataFrame is used in multiple downstream aggregations. What should you do to reduce recomputation?

- a) Add more shuffle partitions
- b) Cache or persist the DataFrame
- c) Increase HDFS replication
- d) Enable speculative execution

Answer: b

Question: 9

Which method removes a cached DataFrame from memory/disk?

- a) unpersist()
- b) clear()
- c) stop()
- d) drop()

Answer: a

Question: 10

A data engineer creates a Cloudera Iceberg table for analytical queries. What is Iceberg primarily designed to solve compared to Hive tables?

- a) Schema evolution and hidden partitioning
- b) Faster Kerberos ticket renewal
- c) Improved Yarn node scheduling
- d) Automated replication to Ranger

Answer: a

Avail the Study Guide to Pass Cloudera CDP-3002 Data Engineer Exam:

- Find out about the CDP-3002 syllabus topics. Visiting the official site offers an idea about the exam structure and other important study resources. Going through the syllabus topics help to plan the exam in an organized manner.
- Once you are done exploring the [**CDP-3002 syllabus**](#), it is time to plan for studying and covering the syllabus topics from the core. Chalk out the best plan for yourself to cover each part of the syllabus in a hassle-free manner.
- A study schedule helps you to stay calm throughout your exam preparation. It should contain your materials and thoughts like study hours, number of topics for daily studying mentioned on it. The best bet to clear the exam is to follow your schedule rigorously.
- The candidate should not miss out on the scope to learn from the CDP-3002 training. Joining the Cloudera provided training for CDP-3002 exam helps a candidate to strengthen his practical knowledge base from the certification.
- Learning about the probable questions and gaining knowledge regarding the exam structure helps a lot. Go through the [**CDP-3002 sample questions**](#) and boost your knowledge

- Make yourself a pro through online practicing the syllabus topics. CDP-3002 practice tests would guide you on your strengths and weaknesses regarding the syllabus topics. Through rigorous practicing, you can improve the weaker sections too. Learn well about time management during exam and become confident gradually with practice tests.

Career Benefits:

Passing the CDP-3002 exam, helps a candidate to prosper highly in his career. Having the certification on the resume adds to the candidate's benefit and helps to get the best opportunities.

Here Is the Trusted Practice Test for the CDP-3002 Certification

VMExam.Com is here with all the necessary details regarding the CDP-3002 exam. We provide authentic practice tests for the CDP-3002 exam. What do you gain from these practice tests? You get to experience the real exam-like questions made by industry experts and get a scope to improve your performance in the actual exam. Rely on VMExam.Com for rigorous, unlimited two-month attempts on the CDP-3002 practice tests, and gradually build your confidence. Rigorous practice made many aspirants successful and made their journey easy towards grabbing the Cloudera Data Engineer (CDP-3002).

Start Online Practice of CDP-3002 Exam by Visiting URL

<https://www.analyticsexam.com/cloudera/cdp-3002-cloudera-data-engineer>