



Cloudera CDP-3002

CLOUDERA DATA ENGINEER CERTIFICATION QUESTIONS & ANSWERS

Exam Summary – Syllabus – Questions

CDP-3002

Cloudera Data Engineer (CDP-3002)

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

www.AnalyticsExam.Com

Table of Contents

Know Your CDP-3002 Certification Well:	2
Cloudera CDP-3002 Data Engineer Certification Details:	2
CDP-3002 Syllabus:	3
Cloudera CDP-3002 Sample Questions:	3
Study Guide to Crack Cloudera Data Engineer CDP-3002 Exam:	6

Know Your CDP-3002 Certification Well:

The CDP-3002 is best suitable for candidates who want to gain knowledge in the Cloudera Engineering. Before you start your CDP-3002 preparation you may struggle to get all the crucial Cloudera Data Engineer materials like CDP-3002 syllabus, sample questions, study guide.

But don't worry the CDP-3002 PDF is here to help you prepare in a stress free manner.

The PDF is a combination of all your queries like-

- What is in the CDP-3002 syllabus?
- How many questions are there in the CDP-3002 exam?
- Which Practice test would help me to pass the CDP-3002 exam at the first attempt?

Passing the CDP-3002 exam makes you Cloudera Data Engineer (CDP-3002). Having the Cloudera Data Engineer certification opens multiple opportunities for you. You can grab a new job, get a higher salary or simply get recognition within your current organization.

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%

Cloudera CDP-3002 Sample Questions:

Question: 1

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: 2

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: 3

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

Question: 4

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: 5

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: 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

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: 8

Which method removes a cached DataFrame from memory/disk?

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

Answer: a

Question: 9

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: 10

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

Study Guide to Crack Cloudera Data Engineer CDP-3002 Exam:

- Getting details of the CDP-3002 syllabus, is the first step of a study plan. This pdf is going to be of ultimate help. Completion of the syllabus is must to pass the CDP-3002 exam.
- Making a schedule is vital. A structured method of preparation leads to success. A candidate must plan his schedule and follow it rigorously to attain success.
- Joining the Cloudera provided training for CDP-3002 exam could be of much help. If there is specific training for the exam, you can discover it from the link above.
- Read from the CDP-3002 sample questions to gain your idea about the actual exam questions. In this PDF useful sample questions are provided to make your exam preparation easy.
- Practicing on CDP-3002 practice tests is must. Continuous practice will make you an expert in all syllabus areas.

Reliable Online Practice Test for CDP-3002 Certification

Make AnalyticsExam.Com your best friend during your Cloudera Data Engineer exam preparation. We provide authentic practice tests for the CDP-3002 exam. Experts design these online practice tests, so we can offer you an exclusive experience of taking the actual CDP-3002 exam. We guarantee you 100% success in your first exam attempt if you continue practicing regularly. Don't bother if you don't get 100% marks in initial practice exam attempts. Just utilize the result section to know your strengths and weaknesses and prepare according to that until you get 100% with our practice tests. Our evaluation makes you confident, and you can score high in the CDP-3002 exam.

Start Online Practice of CDP-3002 Exam by Visiting URL

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