

DUMPS ARENA

ACA Big Data Certification Exam

Alibaba Cloud ACA-BigData1

Version Demo

Total Demo Questions: 10

Total Premium Questions: 78

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

The data development mode in DataWorks has been upgraded to the three-level structure comprising of _____, _____, and _____. (Number of correct answers: 3)

Score 2

- A. Project
- B. Solution
- C. Business flow
- D. Directory

ANSWER: A B

QUESTION NO: 2

In a scenario where a large enterprise plans to use MaxCompute to process and analyze its data, tens of thousands of tables and thousands of tasks are expected for this project, and a project team of 40 members is responsible for the project construction and O&M. From the perspective of engineering, which of the following can considerably reduce the cost of project construction and management?

Score 2

- A. Develop directly on MaxCompute and use script-timed scheduling tasks
- B. Use DataWorks
- C. Use Eclipse
- D. Use a private platform specially developed for this project

ANSWER: B

QUESTION NO: 3

The data development mode in DataWorks has been upgraded to the three-level structure

comprising of _____, _____, and _____. (Number of correct answers: 3)

Score 2

- A. Project
- B. Solution
- C. Business flow
- D. Directory

ANSWER: A B D

QUESTION NO: 4

A Log table named log in MaxCompute is a partition table, and the partition key is dt. A new partition is created daily to store the new data of that day. Now we have one

month's data, starting from dt='20180101' to dt='20180131', and we may use _____
to delete the data on 20180101.

- A. delete from log where dt='20180101'
- B. truncate table where dt='20180101'
- C. drop partition log (dt='20180101')
- D. alter table log drop partition(dt='20180101')

ANSWER: D

QUESTION NO: 5

A company originally handled the local data services through the Java programs. The local data have been migrated to MaxCompute on the cloud, now the data can be accessed through modifying the Java code and using the Java APIs provided by MaxCompute.

Score 1

- A. True
- B. False

ANSWER: A**QUESTION NO: 6**

In DataWorks, a task should be instantiated first before a scheduled task is running every time, that is, generating a corresponding instance which is executed for running the scheduled task. The status is different in each phase of the scheduling process, including _____. (Number of correct answers: 3)

- A. Not running
- B. Running
- C. Running Successfully

ANSWER: A B C**QUESTION NO: 7**

Your company stores user profile records in an OLTP databases. You want to join the serecords with web server logs you have already ingested into the Hadoop file system.

What is the best way to obtain and ingest these user records?

- A. Ingest with Hadoop streaming
- B. Ingest using Hive
- C. Ingest with sqoop import
- D. Ingest with Pig's LOAD command

ANSWER: C**QUESTION NO: 8**

Which of the following Hadoop ecosystem componets can you choose to setup a streaming log analysis system?(Number of correct answers: 3)

Score 2

- A. Apache Flume
- B. Apache Kafka
- C. Apache Spark
- D. Apache Lucene

ANSWER: A C D

QUESTION NO: 9

MaxCompute SQL is suitable for processing less real-time massive data, and employs a syntax similar to that of SQL. The efficiency of data query can be improved through creating proper indexes in the table.

- A. True
- B. False

ANSWER: B

QUESTION NO: 10

A distributed file system like GFS and Hadoop are design to have much larger block(or chunk) size like 64MB or 128MB, which of the following descriptions are correct? (Number of correct answers: 4)

Score 2

- A. It reduces clients' need to interact with the master because reads and writes on the same block(or chunk) require only one initial request to the master for block location information
- B. Since on a large block(or chunk), a client is more likely to perform many operations on a given block, it can reduce network overhead by keeping a persistent TCP connection to the metadata server over an extended period of time
- C. It reduces the size of the metadata stored on the master
- D. The servers storing those blocks may become hot spots if many clients are accessing the same small files
- E. If necessary to support even larger file systems, the cost of adding extra memory to the meta data server is a big price

ANSWER: A B C D