

DUMPS ARENA

Oracle Big Data 2016 Implementation Essentials

Oracle 1z0-449

Version Demo

Total Demo Questions: 10

Total Premium Questions: 72

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

The NoSQL KVStore experiences a node failure. One of the replicas is promoted to primary.

How will the NoSQL client that accesses the store know that there has been a change in the architecture?

- A. The KVLite utility updates the NoSQL client with the status of the master and replica.
- B. KVStoreConfig sends the status of the master and replica to the NoSQL client.
- C. The NoSQL admin agent updates the NoSQL client with the status of the master and replica.
- D. The Shard State Table (SST) contains information about each shard and the master and replica status for the shard.

ANSWER: D**QUESTION NO: 2**

Which three pieces of hardware are present on each node of the Big Data Appliance? (Choose three.)

- A. high capacity SAS disks
- B. memory
- C. redundant Power Delivery Units
- D. InfiniBand ports
- E. InfiniBand leaf switches

ANSWER: A B D**QUESTION NO: 3**

What are the two roles performed by the Big Data Appliance and the Exadata Database Machine in an Oracle Big Data Management solution? (Choose two.)

- A. Data Warehouse
- B. Data Definer
- C. Data Analyzer
- D. Data Reservoir
- E. Data Connector
- F. Data Integrator

ANSWER: E F

QUESTION NO: 4

The `hdfs_stream` script is used by the Oracle SQL Connector for HDFS to perform a specific task to access data.

What is the purpose of this script?

- A. It is the preprocessor script for the Impala table.
- B. It is the preprocessor script for the HDFS external table.
- C. It is the streaming script that creates a database directory.
- D. It is the preprocessor script for the Oracle partitioned table.
- E. It defines the jar file that points to the directory where Hive is installed.

ANSWER: B

QUESTION NO: 5

Your customer has three XML files in HDFS with the following contents. Each XML file contains comments made by users on a specific day. Each comment can have zero or more "likes" from other users. The customer wants you to query this data and load it into the Oracle Database on Exadata.

How should you parse this data?

```
mydata/comments1.xml
```

```
<comments date="2013-12-30">  
<comment id="12345" user="john" text="It is raining :("/>  
<comment id="56789" user="kelly" text="I won the lottery!">  
<like user="john"/>  
<like user="mike"/>  
</comment>  
</comments>
```

```
mydata/comments2.xml
```

```
<comments date="2013-12-31">  
<comment id="54321" user="mike" text="Happy New Year!">  
<like use="laura"/>  
</comment>  
</comments>
```

```
mydata/comments3.xml
```

```
<comments date="2014-01-01">  
<comment id="87654" user="mike" text="I don't feel so good."/>  
<comment id="23456" user="john" text="what a beautiful day!">  
<like user="kelly"/>  
<like user="phil"/>  
</comment>  
</comments>
```

- A. by creating a table in Hive and using MapReduce to parse the XML data by column
- B. by configuring the Oracle SQL Connector for HDFS and parsing by using SerDe
- C. by using the XML file module in the Oracle XQuery for Hadoop Connector
- D. by using the built-in functions for reading JSON in the Oracle XQuery for Hadoop Connector

ANSWER: B

QUESTION NO: 6

What are two reasons that a MapReduce job is not run when the following command is executed in Hive? (Choose two.)

```
hive> create view v_consolidated_credit_products as select
```

```
hcd.customer_id, hcd.credit_card_limits, hcd.credit_balance, hmda.n_mortgages, hmda.mortgage_amount  
from mortgages_department_agg hmda join credit_department hcd on  
hcd.customer_id=hmda.customer_id;
```

OK

Time taken: 0.316 seconds

- A. The MapReduce job is run when the view is accessed.
- B. MapReduce is run; the command output is incorrect.
- C. MapReduce is not run with Hive. Hive bypasses the MapReduce layer.
- D. A view only defines the metadata in the RDBMS store.
- E. MapReduce is run in the background.

ANSWER: A D

QUESTION NO: 7

Your customer needs to analyze large numbers of log files after combining them with customer profile and interaction data. The customer wants to be able to analyze these files back to their go-live date, which was last October.

Which solution should you choose to accomplish this?

- A. Oracle Database
- B. Hadoop
- C. Berkeley DB
- D. Oracle NoSQL

ANSWER: B

QUESTION NO: 8

Identify two valid steps for setting up a Hive-to-Hive transformation by using Oracle Data Integrator. (Choose two.)

- A. Ensure that the Hive server is up and running.
- B. Create a Logical Schema object.
- C. Configure ODI by using the mammoth utility.
- D. Ensure that Apache Sentry is configured.

ANSWER: A B

QUESTION NO: 9

What is the output of the following six commands when they are executed by using the Oracle XML Extensions for Hive in the Oracle XQuery for Hadoop Connector?

1. `$ echo "xxx" > src.txt`
2. `$ hive --auxpath $OXH_HOME/hive/lib -i $OXH_HOME/hive/init.sql`
3. `hive> CREATE TABLE src (dummy STRING);`
4. `hive> LOAD DATA LOCAL INPATH 'src.txt' OVERWRITE INTO TABLE src;`
5. `hive> SELECT * FROM src;`
OK
xxx
6. `hive> SELECT xml_query ("x/y", "123456") FROM src;`

- A. xyz
- B. 123
- C. 456
- D. xxx
- E. x/y

ANSWER: B

QUESTION NO: 10

What are two of the main steps for setting up Oracle XQuery for Hadoop? (Choose two.)

- A. unpacking the contents of oXH-version.zip into the installation directory
- B. installing the Oracle SQL Connector for Hadoop
- C. configuring an Oracle wallet
- D. installing the Oracle Loader for Hadoop

ANSWER: A D