DUMPSBOSS.

Oracle Big Data 2016 Implementation Essentials

Oracle 1z0-449

Version Demo

Total Demo Questions: 10

Total Premium Questions: 72

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

Your customer has an older starter rack Big Data Appliance (BDA) that was purchased in 2013. The customer would like to know what the options are for growing the storage footprint of its server.

Which two options are valid for expanding the customer's BDA footprint? (Choose two.)

A. Elastically expand the footprint by adding additional high capacity nodes.

B. Elastically expand the footprint by adding additional Big Data Oracle Database Servers.

C. Elastically expand the footprint by adding additional Big Data Storage Servers.

D. Racks manufactured before 2014 are no longer eligible for expansion.

E. Upgrade to a full 18-node Big Data Appliance.

ANSWER: D E

QUESTION NO: 2

Your customer uses Active Directory to manage user accounts. You are setting up Hadoop Security for the customer's Big Data Appliance.

How will you integrate Hadoop and Active Directory?

A. Set up Kerberos' Key Distribution Center to be the Active Directory keystore.

B. Configure Active Directory to use Kerberos' Key Distribution Center.

C. Set up a one-way cross-realm trust from the Kerberos realm to the Active Directory realm.

D. Set up a one-way cross-realm trust from the Active Directory realm to the Kerberos realm.

ANSWER: C

QUESTION NO: 3

What two actions do the following commands perform in the Oracle R Advanced Analytics for Hadoop Connector? (Choose two.)

ore.connect (type="HIVE")

ore.attach ()

A. Connect to Hive.

B. Attach the Hadoop libraries to R.

C. Attach the current environment to the search path of R.

D. Connect to NoSQL via Hive.

ANSWER: A C

QUESTION NO: 4

Identify two features of the Hadoop Distributed File System (HDFS). (Choose two.)

A. It is written to store large amounts of data.

B. The file system is written in C#.

C. It consists of Mappers, Reducers, and Combiners.

D. The file system is written in Java.

ANSWER: A D

QUESTION NO: 5

Your customer has a Big Data Appliance and an Exadata Database Machine and would like to extend security. Select two ways that security works in Big Data SQL. (Choose two.)

A. On the Big Data Appliance, Hadoop's native security is used.

B. On the Exadata Database Machine, Oracle Advanced Security is used for fine-grained access control.

C. On the Big Data Appliance, Oracle Advanced Hadoop Security is used for fine grained access control.

D. On the Big Data Appliance, Oracle Identity Management is used.

E. On the Big Data Appliance, data is encrypted by using Oracle Transparent Data Encryption (TDE).

ANSWER: B E

QUESTION NO: 6

What kind of workload is MapReduce designed to handle?

A. batch processing

B. interactive

C. computational

D. real time

E. commodity

ANSWER: A

QUESTION NO: 7

Your customer's Hadoop cluster displays an error while running a MapReduce task by using Yarn. The error states that a container could not be created.

Which resource in Yarn creates containers?

A. Task Tracker

B. Resource Manager

C. NameNode

D. MasterNode

E. Node Manager

ANSWER: E

QUESTION NO: 8

Your customer needs to access Hive tables, HDFS, and Data Pump format files as the data source.

When installing and configuring the Oracle SQL Connector for HDFS, which two locations must be installed or configured to fulfill this requirement? (Choose two.)

A. Hadoop cluster

B. HiveQL

C. Impala

D. Oracle database

E. Oracle NoSQL

ANSWER: A D

QUESTION NO: 9

Your customer needs the data that is generated from social media such as Facebook and Twitter, and the customer's website to be consumed and sent to an HDFS directory for analysis by the marketing team.

Identify the architecture that you should configure.

A. multiple flume agents with collectors that output to a logger that writes to the Oracle Loader for Hadoop agent

B. multiple flume agents with sinks that write to a consolidated source with a sink to the customer's HDFS directory

C. a single flume agent that collects data from the customer's website, which is connected to both Facebook and Twitter, and writes via the collector to the customer's HDFS directory

D. multiple HDFS agents that write to a consolidated HDFS directory

E. a single HDFS agent that collects data from the customer's website, which Is connected to both Facebook and Twitter, and writes via the Hive to the customer's HDFS directory

ANSWER: B

QUESTION NO: 10

Your customer has 10 web servers that generate logs at any given time. The customer would like to consolidate and load this data as it is generated into HDFS on the Big Data Appliance.

Which option should the customer use?

A. Set up a zookeeper agent to capture the transactions and write them to HDFS.

B. Write a hive query to listen for new logs and save them in a Hive table.

C. Set up a flume agent to capture the transactions and write them to HDFS.

D. Set up an hbase agent to capture the transactions and write them to HDFS.

E. Set up a web server agent in Apache Oozie to write the data to HDFS.

ANSWER: C