

PDFBraindumps



Latest Pdf Braindumps	Top Certifications	Top Vendors		
<ul style="list-style-type: none">▶ LRP-614▶ BCABA▶ JN0-740▶ 250-405▶ DS-200▶ SDM_2002001040▶ ST0-250▶ H12-221▶ M2180-716	<ul style="list-style-type: none">▶ ISEB Certification▶ OCE▶ NVIDIA Certifications▶ Network+▶ IBM Certified Integrat▶ CCDH▶ IBM Certified Advanc▶ eserver Certified Spe▶ SAP-Certifications▶ Network Appliance N	<ul style="list-style-type: none">▶ HCNP▶ IFPUG Certifications▶ dotMobi Certification▶ SCMA▶ MCSD▶ NCLP▶ XMLMaster Certificat▶ CS5▶ CHA	<ul style="list-style-type: none">▶ ISEB▶ ASTQB▶ Aruba▶ Data Center Universit▶ HRCI▶ CIW▶ Patchlink▶ International Consorti▶ Acme-Packet	<ul style="list-style-type: none">▶ Fortinet▶ Ericsson▶ Liferay▶ Novell▶ Huawei▶ RSA▶ MYSQL▶ ISM▶ CheckPoint

<http://www.pdfbraindumps.com>

Latest pdf braindumps provider, high pass rate

Exam : **1z1-449**

Title : Oracle Big Data 2017
Implementation Essentials

Vendor : Oracle

Version : DEMO

NO.1 What access driver does the Oracle SQL Connector for HDFS use when reading HDFS data by using external tables?

- A. ORACLE_DATA_PUMP
- B. ORACLE_LOADER
- C. ORACLE_HDP
- D. ORACLE_BDSQL
- E. HADOOP_LOADER
- F. ORACLE_HIVE_LOADER

Answer: B

Explanation:

Oracle SQL Connector for HDFS creates the external table definition for Data Pump files by using the metadata from the Data Pump file header. It uses the ORACLE_LOADER access driver with the preprocessor access parameter. It also uses a special access parameter named EXTERNAL VARIABLE DATA, which enables ORACLE_LOADER to read the Data Pump format files generated by Oracle Loader for Hadoop.

References: https://docs.oracle.com/cd/E37231_01/doc.20/e36961/sqlch.htm#BDCUG356

NO.2 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

Explanation:

The hdfs_stream script is the preprocessor for the Oracle Database external table created by Oracle SQL Connector for HDFS.

References: https://docs.oracle.com/cd/E37231_01/doc.20/e36961/start.htm#BDCUG107

NO.3 What does the flume sink do in a flume configuration?



- A. sinks the log file that is transmitted into Hadoop
- B. hosts the components through which events flow from an external source to the next destination
- C. forwards events to the source

- D. consumes events delivered to it by an external source such as a web server
- E. removes events from the channel and puts them into an external repository

Answer: D

Explanation:

A Flume source consumes events delivered to it by an external source like a web server. The external source sends events to Flume in a format that is recognized by the target Flume source. When a Flume source receives an event, it stores it into one or more channels. The channel is a passive store that keeps the event until it's consumed by a Flume sink.

References: <https://flume.apache.org/FlumeUserGuide.html>

NO.4 What does the following line do in Apache Pig?

```
products = LOAD '/user/oracle/products' AS (prod_id, item);
```

- A. The products table is loaded by using data pump with prod_id and item.
- B. The LOAD table is populated with prod_id and item.
- C. The contents of /user/oracle/products are loaded as tuples and aliased to products.
- D. The contents of /user/oracle/products are dumped to the screen.

Answer: C

Explanation:

The LOAD function loads data from the file system.

Syntax: LOAD 'data' [USING function] [AS schema];

Terms: 'data'

The name of the file or directory, in single quote

References: <https://pig.apache.org/docs/r0.11.1/basic.html#load>

NO.5 How should you encrypt the Hadoop data that sits on disk?

- A. Enable Transparent Data Encryption by using the Mammoth utility.
- B. Enable HDFS Transparent Encryption by using bdacli on a Kerberos-secured cluster.
- C. Enable HDFS Transparent Encryption on a non-Kerberos secured cluster.
- D. Enable Audit Vault and Database Firewall for Hadoop by using the Mammoth utility.

Answer: B

Explanation:

HDFS Transparent Encryption protects Hadoop data that's at rest on disk.

When the encryption is enabled for a cluster, data write and read operations on encrypted zones (HDFS directories) on the disk are automatically encrypted and decrypted.

This process is "transparent" because it's invisible to the application working with the data.

The cluster where you want to use HDFS Transparent Encryption must have Kerberos enabled.

NO.6 Which user in the following set of ACL entries has read, write and execute permissions?

```
user::rwx group::r-x other::r-x
```

```
default:user::rwx
```

```
default:user:bruce:rwx #effective:r-x
```

```
default:group::r-x
```

```
default:group:sales:rwx #effective:r-x
```

```
default:mask::r-x
```

default:other::r-x

- A. sales
- B. bruce
- C. all users
- D. hdfs
- E. user

Answer: E

Explanation:

References: <https://askubuntu.com/questions/257896/what-is-meant-by-mask-and-effective-in-the-output-from-getfacl>

NO.7 Your customer keeps getting an error when writing a key/value pair to a NoSQL replica.

What is causing the error?

- A. The master may be in read-only mode and as result, writes to replicas are not being allowed.
- B. The replica may be out of sync with the master and is not able to maintain consistency.
- C. The writes must be done to the master.
- D. The replica is in read-only mode.
- E. The data file for the replica is corrupt.

Answer: C

Explanation:

Replication Nodes are organized into shards.

A shard contains a single Replication Node which is responsible for performing database writes, and which copies those writes to the other Replication Nodes in the shard. This is called the master node.

All other Replication

Nodes in the shard are used to service read-only operations.

Note: Oracle NoSQL Database provides multi-terabyte distributed key/value pair storage that offers scalable throughput and performance. That is, it services network requests to store and retrieve data which is organized into key-value pairs.

References:

https://docs.oracle.com/cd/E26161_02/html/GettingStartedGuide/introduction.html