ISLEVER

P2090-038

IBM InfoSphere BigInsights Technical Mastery Test v2

DEMO

https://www.islever.com/p2090-038.html

https://www.islever.com/ibm.html

For the most up-to-date exam questions and materials, we recommend visiting our website, where you can access the latest content and resources.

QUESTION NO: 1

Which of the following options best describes the differences between a traditional data warehouse environment and a Hadoop environment?

A. Traditional data warehousing environments are mostly ideal for analyzing structured data from various systems, while a Hadoop environment is well suited to deal with structured, semi-structured, and unstructured data, as well as when a data discovery process is needed.
B. Hadoop environments are mostly ideal for analyzing structured and semi-structured data from a single system, while traditional data warehousing environment is well suited to deal with unstructured data, as well as when a data discovery process is needed.
C. Typically, data stored in Hadoop environments is cleaned up before storing in the distributed file-system.
D. Typically, data stored in data warehousing environments is rarely filtered and pre-processed. On the other hand, data injected into Hadoop environments is always pre-processed and filtered.

Answer: A
Explanation:

QUESTION NO: 2

What is Big SQL?

A. Big SQL is a feature in Data Explorer that allows for indexing of data from SQL sources such as data warehouses.
B. Big SQL is a feature in BigInsights that allows for native SQL query access for Hadoop, providing full ANSI SQL 92 compliance and standardSQL syntax such as joins, for data contained in a variety of formats such as structured Hive tables, Hbase tables, or csv and other delimitedfiles in HDFS.
C. Big SQL is a feature in Streams that allows for real time analysis of data via standard SQL syntax.
D. Big SQL is a feature in BigInsights that provides a SQL like interface to data contained in Hbase tables only. Other data sources in HDFS mustbe accessed via other means such as HiveQL.

Answer: B
Explanation:

QUESTION NO: 3

What is the InfoSphere BigInsights Credential Store?

A. The InfoSphere BigInsights credentials store is a table stored in the HBase relational database that stores passwords, tokens, and other potentially sensitive information.

B. The InfoSphere BigInsights credentials store is a designated folder on the distributed file system (DFS) that stores passwords, tokens, and other potentially sensitive information.
C. The InfoSphere BigInsights credentials store is a designated folder in the local file system (not HDFS) that stores the authorities and privileges for all users in the BigInsights environment.
D. The InfoSphere BigInsights credentials store is a designated file defined by an environment variable that stores the authorities and privileges for all users in the BigInsights environment.

Answer: B
Explanation:

QUESTION NO: 4

What does ig Data?represent?What does ?ig Data?represent?

A. A Hadoop feature capable of processing vast amounts of data in-parallel on large clusters of commodity hardware in a reliable, fault-tolerant manner.
B. A concept and platform of technologies with the characteristics of the ? Vs? that is able to handle large amounts of unstructured,semi-structured, and structured raw data unlike traditional systems.
C. A database feature capable of converting pre-existing structured data into unstructured raw data.
D. Only data stored in the BIGDATA table in any relational database.

Answer: B
Explanation:

QUESTION NO: 5

How do existing applications usually connect to InfoSphere BigInsights using the Big SQL feature?

A. Applications will connect using custom made connectors programmed in SPL.
B. Applications will connect using standard JDBC and ODBC drivers that come with InfoSphere BigInsights.
C. Applications will connect using the JAQL programming language.
D. Applications will connect using either HiveQL or Pig programming languages.

Answer: B
Explanation:

QUESTION NO: 6

Which of the following components is NOT included in the BigInsights Basic Edition distribution?

A. Hadoop Distributed File System.

B. Hive.
C. Pig.
D. BigSheets.

Answer: D
Explanation:

QUESTION NO: 7

Hadoop environments are optimized for:

A. Processing transactions (random access).
B. Low latency data access.
C. Batch processing on large files.
D. Intensive calculation with little data.

Answer: C
Explanation:

QUESTION NO: 8

Which of the following options DOES NOT describe an advantage of InfoSphere BigInsights on top of regular standalone Hadoop distributions?

A. Integrated install
B. BigSheets
C. Adaptive MapReduce
D. Support to HiveQL and PigLatin languages.

Answer: D
Explanation:

QUESTION NO: 9

For maximum stability, the recommended number of dedicated hosts in order to implement the high availability configuration for BigInsights (consisting of the HA manager, the NameNode, and the JobTracker) is:

A. 1
B. 2
C. 3
D. 4

Answer: C