P2090-038^{Q&As}

IBM InfoSphere BigInsights Technical Mastery Test v2

Pass IBM P2090-038 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/p2090-038.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by IBM Official Exam Center

⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

QUESTION 1

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.

Correct Answer: B

QUESTION 2

Which of the following options is CORRECT?

A. InfoSphere BigInsights is based on a branched Hadoop distribution, and therefore backwards compatibility is not guaranteed.

B. InfoSphere BigInsights is a distributed file system used as base for Hadoop distributions.

C. InfoSphere BigInsights is based on the non forked core Hadoop distribution, but backwards compatibility with the Apache Hadoop project is not guaranteed, therefore applications written for Hadoop might not run on BigInsights.

D. InfoSphere BigInsights is based on the non forked core Hadoop distribution, and backwards compatibility with the Apache Hadoop project will always be maintained. Therefore, all applications written for Hadoop will run on BigInsights.

Correct Answer: D

QUESTION 3

Which of the following options best describes the proper usage of MapReduce jobs in Hadoop environments?

A. MapReduce jobs are used to process vast amounts of data in-parallel on large clusters of commodity hardware in a reliable, fault-tolerant manner.

B. MapReduce jobs are used to process small amounts of data in-parallel on expensive hardware, without fault-tolerance.

C. MapReduce jobs are used to process structured data in sequence, with fault-tolerance.

D. MapReduce jobs are used to execute sequential search outside the Hadoop environment using a built-in UDF to access information stored in non-relational databases.

Correct Answer: A

QUESTION 4

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.

Correct Answer: C

QUESTION 5

What is a major difference between HDFS and GPFS-FPO?

A. There is no difference, GPFS-FPO is just IBM version of HDFS in BigInsights.

B. There is no difference, as GPFS-FPO is the new open source standard version for the Hadoop file system.

C. Unlike HDFS, GPFS-FPO is a kernel level, POSIX compliant, and highly available file system that allows for standard applications to also use the storage area marked for Hadoop use.

D. Unlike HDFS, GPFS-FPO is a kernel level, POSIX compliant file system that allows for standard applications to also use the storage area marked for Hadoop use. Like HDFS, GPFS- FPO also has an issue with the meta data name node service which may be a single point of failure.

Correct Answer: C

P2090-038 PDF Dumps   P2090-038 Exam Questions   P2090-038 Braindumps