



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.geekcert.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

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

QUESTION 2

InfoSphere BigInsights offers the following benefits to your organization, EXCEPT:

- A. It cuts costs by providing an efficient compression algorithm based on row level compression using a dictionary to store repetitive patterns.
- B. It complements your existing infrastructure by extending your data collection and analysis capabilities.
- C. It integrates data from a variety of structured and unstructured sources.
- D. It enables analysis at scale.

Correct Answer: A

QUESTION 3

Which of the following InfoSphere BigInsights editions includes the Adaptive MapReduce feature?

- A. InfoSphere BigInsights Express Edition.
- B. InfoSphere BigInsights Enterprise Edition.
- C. InfoSphere BigInsights Extended Developer Edition.
- D. InfoSphere BigInsights Advanced Enterprise Edition.

Correct Answer: B

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

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

[Latest P2090-038 Dumps](#)

[P2090-038 Practice Test](#)

[P2090-038 Braindumps](#)