



# P2090-038<sup>Q&As</sup>

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.passapply.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

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 2

How is data stored in a Hadoop cluster?

- A. The data is divided into blocks, and copies of these blocks are replicated across multiple servers in the Hadoop cluster.
- B. The data converted into a single block, and the block is stored in just one of the servers in the Hadoop cluster.
- C. The data is divided into blocks, each block is stored in a different server in the Hadoop cluster, but the blocks are not replicated.
- D. The data is converted into a single block, and copies of this block are replicated across multiple servers in the Hadoop cluster.

Correct Answer: A

---

### QUESTION 3

Which of the following InfoSphere BigInsights features provides a vast library of extractors enabling actionable insights from large amounts of native textual data?

- A. Text Analytics.
- B. Adaptive MapReduce.
- C. General Parallel File System.
- D. BigSheets.

Correct Answer: A

---



#### QUESTION 4

Which of the following statements should be used to import data from DB2 into InfoSphere BigInsights?

- A. You can use the Big SQL LOAD FROM statement to load data into a Big SQL table.
- B. You can use the Big SQL IMPORT FROM statement to load data into a HBase table.
- C. You can use the Big SQL LOAD FROM statement to load data into a Hive table.
- D. You can use the Big SQL IMPORT FROM statement to load data into a Big SQL table.

Correct Answer: A

---

#### QUESTION 5

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.

Correct Answer: A

[Latest P2090-038 Dumps](#)

[P2090-038 Study Guide](#)

[P2090-038 Exam Questions](#)