

**100%** Money Back  
**Guarantee**

**Vendor:**IBM

**Exam Code:**P2090-038

**Exam Name:**IBM InfoSphere BigInsights Technical  
Mastery Test v2

**Version:**Demo

### QUESTION 1

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.

Correct Answer: B

---

### QUESTION 2

Which of the following options is NOT CORRECT?

- A. Big data solutions are ideal for analyzing not only raw structured data, but semi structured and unstructured data from a wide variety of sources.
- B. Big data solutions are ideal when all, or most, of the data needs to be analyzed versus a sample of the data; or a sampling of data isn't nearly as effective as a larger set of data from which to derive analysis.
- C. Big data solutions are ideal for Online Transaction Processing (OLTP) environments.
- D. Big data solutions are ideal for iterative and exploratory analysis when business measures on data are not predetermined.

Correct Answer: C

---

### 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

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

Correct Answer: C

---

#### QUESTION 5

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 6

Which of the following integrated development environment (IDE) can be used to develop programs to be used with InfoSphere BigInsights?

- A. InfoSphere Data Explorer.
- B. IBM Data Studio with InfoSphere Data Explorer plug-in.
- C. InfoSphere Streams Studio
- D. Eclipse with InfoSphere BigInsights plug-in.

Correct Answer: D

---

#### QUESTION 7

Which of the following options is CORRECT in regards to integrating InfoSphere BigInsights and warehousing environments?

- A. There is no direct interaction between InfoSphere BigInsights and data warehousing environments. However, InfoSphere Streams can be used to stream data between the two environments.

B. Frequently accessed data should be maintained in the BigInsights environment while old?or outdated information can be offloaded to the Frequently accessed data should be maintained in the BigInsights environment while ?old?or outdated information can be offloaded to the warehouse environment.

C. InfoSphere BigInsights Enterprise Edition also comes with a standard Java Database Connectivity (JDBC) connector, making it possible for organizations to quickly integrate with a wide variety of data and information systems, including Oracle, Microsoft SQL Server, MySQL and Teradata.

D. There is no integration between InfoSphere BigInsights and warehousing environments since they deal with different type of data.

Correct Answer: C

---

### QUESTION 8

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.

Correct Answer: D

---

### QUESTION 9

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 10

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 11

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

---

### QUESTION 12

What is the difference between Hadoop MapReduce and IBM Adaptive MapReduce feature available in InfoSphere BigInsights?

- A. Hadoop Map Reduce is optimized for operating on small files or splits, while IBM Adaptive MapReduce is optimized for operating on largepartitioned files.
- B. Hadoop MapReduce is optimized for operating on large files, while IBM Adaptive MapReduce is configurable to operate optimized on large or small files or splits
- C. Hadoop MapReduce is optimized for operating on small files or splits, while IBM Adaptive MapReduce is optimized for operating on large files stored in individual blocks.
- D. Hadoop MapReduce is optimized for operating on small partitioned tables stored in the HBase component, while IBM? Adaptive MapReduceis optimized for operating on large partitioned files.

Correct Answer: B