

IBM

Exam C2030-102

IBM BigData Architect

Version: 6.0

[Total Questions: 55]

Question No : 1

An IBM Big Data platform is well suited to deal which of the following kinds of data types?

- A. Structured data in row format only
- B. Semi-structured and unstructured data only
- C. Text data, sensor data, and audio data only
- D. Semi-structured, unstructured, and structured data

Answer: D

Reference:

<https://www-304.ibm.com/industries/publicsector/fileservice?contentid=239170>

Question No : 2

A bank wants to build a system that tracks all ATM and online transactions in real-time. They want to build a personalized model of their customer's financial activity by incorporating enterprise data as well as social media data. The system must be able to learn and adapt over a period of time. These personalized models will be used for real time promotions as well as for any fraud or crime detections. Given these requirements, which of the following would recommend?

- A. Spark
- B. Hadoop
- C. Cloudand
- D. Netezza

Answer: D

Question No : 3

In a typical Hadoop HA cluster, two separate machines are configured as which of the following?

- A. Data Nodes
- B. Edge Nodes

- C. Name Nodes
- D. None of the Above

Answer: A

Reference:

<http://hadoop.apache.org/docs/current/hadoop-project-dist/hadoop-hdfs/HDFSHighAvailabilityWithQJM.html>

Question No : 4

You have a need for Storm real time processing and you realize that your Storm processing is detrimental to the timely execution of your MapReduce batch jobs. Which of the following would be your best course of action?

- A. Implement a Storm-YARN integration to facilitate the management of elastic workloads
- B. Implement the Oozie 2.0 framework optimized for elastic workload management
- C. Use a combination of Flume and Oozie 2.0 to enable the Flume built-in elastic flow automation
- D. Implement Apache ACE 2.0 for Storm

Answer: C

Reference:

<http://www.ibm.com/developerworks/library/bd-ooziehadoop/>

Question No : 5

Which of the following statements is TRUE regarding cloud applications?

- A. Migrating a legacy application to the cloud is a simple solution to drive down cost
- B. Architecting and deploying a scalable cloud application requires a private cloud implementation
- C. To be truly elastic, a cloud application must be tightly coupled to the physical cloud environment
- D. Leveraging a private vs. public cloud may result in sacrificing some of the core

advantages of cloud computing

Answer: B

Question No : 6

What is the most important aspect of data center disaster recovery?

- A. A complete damage assessment
- B. Control of critical assets
- C. Restoration of business functions
- D. Protection of individual life

Answer: D

Question No : 7

A large application vendor wants to port their existing distributed applications to run on Hadoop. In order to be competitive they need to provide monitoring and keep the size of the monitored applications consistent with the configuration. This implies the ability to deploy a replacement, for example, for any failed components. Which of the following would be a workable solution?

- A. Nagios with YARN
- B. Slider with YARN
- C. Oozie with Lucene
- D. OPTIM Performance Manager

Answer: B

Question No : 8

The NameNode determines the rack id to which each DataNode belongs via the process outlined in which of the following?

- A. HDFS APIs
- B. Hadoop Rack Awareness