


Ultimate Guide to Prepare H13-711_V3.0 Certification Exam for HCIA-Big Data in 2023 [Q53-Q75]



Ultimate Guide to Prepare H13-711_V3.0 Certification Exam for HCIA-Big Data in 2023
Use Real H13-711_V3.0 Dumps - Huawei Correct Answers updated on 2023

The certification exam includes multiple-choice questions and simulations that test the candidate's knowledge and skills in big data technologies. The exam is designed to be challenging, but it is also designed to be fair and objective. Candidates who pass the Huawei H13-711_V3.0 certification exam can demonstrate their skills and expertise in big data technologies and gain a competitive edge in the job market.

The HCIA-Big Data certification is a professional certification targeted at candidates who have basic knowledge of the data center, as well as the basics of big data, data analysis, and data mining. This certification is best suited for mid-level IT professionals who want to enhance their skills in the data science domain. The exam tests candidates' ability to design, deploy, maintain a big data platform, and manipulate big data analytically using Hadoop, Kafka, and Spark.

NEW QUESTION 53

In Flink, the checkpoint mechanism can ensure that when an application fails during operation, it can be restored from a certain checkpoint. In this process, stream snapshots are created in sequence according to the order of data inflow.

- * True
- * False

NEW QUESTION 54

The following about Hive Component capabilities in the architecture. Which is the correct description?

- * ThriftServer for thriftcatchmouth, as JDBCservices and willHive and other applicationsintegrated
- * Compiler pressAccording to missionaccording to Dependent relationship is executed separately Map/Reduce task
- * Executor is responsible for editing Translate voQL and will which translates into a series of mutual rely of Map/Reduce Task
- * Optimizer is an optimizer, divided into logical Jiyou optimizer and physical optimizer. HiveQL Generated execution plan MapReduce Task Enterrow optimization

NEW QUESTION 55

The following about Which Kafka feature description is incorrect??

- * Support looking for horizontal expansion
- * Support multi-level cascade
- * Support message partitioning
- * High phagocytosis rate

NEW QUESTION 56

What data does Fusion Insight Manager back up regularly? (Multiple choice)

- * NameNode
- * LDAPServer
- * OMS
- * DBService

NEW QUESTION 57

Regarding Hive log collection on the Fusioninsight HD Manager interface, which option is incorrect?

- * You can specify an instance for log collection, for example, specify a separate collection of MetaStore logs
- * Log collection can be performed in a specified time period, for example, only logs from 2016-1-1 to 2016-1-40 are collected
- * You can specify the node IP for log collection, for example, only download the log of a certain IP
- * Specific users can be designated for log collection, for example, only logs generated by userA can be downloaded

NEW QUESTION 58

Regarding the Supervisor of Fusioninsight HD Streaming, which of the following is correct?

- * Supervisor is responsible for resource allocation and task scheduling
- * Supervisor is responsible for accepting the tasks assigned by Nimbus, starting and stopping the Worker process under its own management
- * Supervisor is a process that runs specific processing logic
- * Supervisor is a component that receives data in Topology and then performs processing

NEW QUESTION 59

To set the maximum resource usage of queue QueueA in YARN, which parameter needs to be configured?

- * yarn.scheduler.capacity.root.QueueA.user-limit-factor
- * yarn.scheduler.capacity.root.QueueA.minimum-user-limit-percent
- * yarn.scheduler.capacity.root.QueueA.state
- * yarn.scheduler.capacity.root.QueueA.maximum-capacity

NEW QUESTION 60

In the processing node Bolt of Streaming. Which of the following operations can be done

- * join operation
- * Connect to the database
- * filter
- * business processing

NEW QUESTION 61

Which of the following statements about Kafka Partition are correct? (Multiple choice)

- * The Partition mechanism is introduced to ensure the high throughput of Kafka
- * Each Partition is an ordered and immutable message queue
- * The number of Partitions determines the maximum number of concurrent consumers in each consumer group
- * Each Partition corresponds to a log file at the storage level

NEW QUESTION 62

HDFS supports large file storage, and supports multiple users to write to the same file, as well as modify any location of the file.

- * True
- * False

NEW QUESTION 63

Which parts of the data need to be read to execute the HBase data reading business?

- * HLog
- * MemStore
- * HFile
- * HMaster

NEW QUESTION 64

When executing the HBase data reading business, which parts of the data need to be read?

- * HFile
- * HLog
- * MemStore
- * HMaster

NEW QUESTION 65

In the replica placement strategy of HDFS, what is the distance between different servers in the same rack?

2

NEW QUESTION 66

Which of the following HBase components is correct? (Multiple choice)

- * Column-oriented
- * Type of NoSQL data
- * Distributed
- * The stored data is in the form of K-V

NEW QUESTION 67

In Hive architecture, () groupPiece is responsible for table, column and Partition etc. yuandataRead, write and update operations

- * Metastore

NEW QUESTION 68

The following pair in LdapServer groupweaving model Which statement is correct?

- * The root of the Ldap Server directory tree generally defines the domain name the C?cin Compnitment
- * Ldap Server directory information is organized and stored based on a tree structure.
- * Each node in the Ldap Server directory tree is called an entry, and I have its own uniquely distinguishable name? ?(Plantnad? ?)
- * Specific objects cannot be defined under organizational units.

NEW QUESTION 69

The following descriptions about FusionInsight multi-tenant management are correct

- * One-stop management of tenant resources, computing resources (CPU/memory/IO), storage resources (HDFS), service resources (HBase…)
- * Visual multi-level tenant management, matching with the enterprise organizational structure, simplifies system resource allocation and management
- * Tenant resource isolation based on container mechanism, escorting tenant SLA
- * A multi-level tenant model that matches the organizational structure of the enterprise, different departments correspond to different tenants, and tenants can be dynamically added or deleted as needed

NEW QUESTION 70

HBase quickly determines that user data does not exist. (fill in the blank)
describe

NEW QUESTION 71

Which module is responsible for Fusioninsight Manager user data storage?

- * CAS
- * AOS
- * Kerberos
- * LDAP

NEW QUESTION 72

In the WebHCat architecture, which of the following operations can users perform through the secure HTTPS protocol?

- * Perform Hive DDL operations
- * Run MapReduce tasks
- * Run Hive HQL tasks
- * All the above are correct

NEW QUESTION 73

When the number of nodes in the Zookeeper cluster is 5 nodes, how many nodes are equivalent to the disaster tolerance of the cluster?

- * 3
- * 4

- * 6
- * None of the above

NEW QUESTION 74

When using the configuration planning tool to plan the Fusioninsight HD V100R002C50 cluster, you can use the configuration planning tool of the V100R2002C60 version to plan.

- * True
- * False

NEW QUESTION 75

Which instance must the Loader instance be deployed with in FusionInsightHD?

- * DataNode
- * RegionServer
- * ResourceManager
- * NodeManager

HCIA-Big Data -H13-711_V3.0 Exam-Practice-Dumps: https://www.validbraindumps.com/H13-711_V3.0-exam-prep.html