

Huawei.H13-711-ENU.v2025-02-12.q181

Exam Code:	H13-711-ENU
Exam Name:	HCIA-Big Data(Huawei Certified ICT Associate-Big Data)
Certification Provider:	Huawei
Free Question Number:	181
Version:	v2025-02-12
# of views:	1361
# of Questions views:	1810
https://www.exam-tests.com/H13-711-ENU-exam/Huawei.H13-711-ENU.v2025-02-12.q181.html	

NEW QUESTION: 1

Is the description of the role of the standby NameNode in the HDFS system correct?

(Multiple choice)

- A. The standby NameNode should be deployed on the same node as the primary NameNode
- B. The standby NameNode has no requirement for memory
- C. Hot Standby of Primary NameNode
- D. Help the main NameNode to merge edit logs and reduce the startup time of the main NameNode

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 2

Which components in Fusioninsight HD can Zookeeper provide distributed management support for? (Multiple choice)

- A. Spark
- B. HBase
- C. Hive
- D. Loader

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 3

Under the HDFS federation mechanism, metadata between NameNodes is not shared.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 4

Which of the following descriptions about the function of HMaster in HBase is correct?
(Multiple choice)

- A. Responsible for load balancing of RegionServer
- B. Responsible for creating tables, modifying tables/deleting tables
- C. Region load balancing, Region splitting and Region allocation after splitting
- D. Region migration after RegionServer failure

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 5

In Fusioninsight HD, which of the following is not a Hive flow control feature? (Multiple choice)

- A. Support threshold control on the number of connections established by each user
- B. Support threshold control on the number of connections established by a specific user
- C. Support threshold control on the number of connections established per unit time
- D. Support threshold control on the total number of established connections

Answer: A,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 6

Fusioninsight Manager's statement about the configuration function of the service is wrong?

- A. Instance-level configuration takes effect only for this instance
- B. Instance-level configurations also take effect on other instances
- C. After the configuration is saved, the service needs to be restarted to take effect
- D. Service-level configurations can take effect for all instances

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 7

The following provides a variety of Redis optimization methods, which option is wrong?

- A. Turn off Transparent Huge Pages
- B. Modify the minimum number of tcp connections in Linux
- C. Simplified key names and key values
- D. Modify the Linux kernel memory allocation strategy or execute sysctl

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 8

Which of the following options is an advantage of Kerberos? (Multiple choice)

- A. Prevent replay attacks
- B. Can prevent brute force cracking
- C. Support two-way authentication
- D. have higher performance

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 9

What query types does Elasticsearch have? (Multiple choice)

- A. Retrieval based on score
- B. Term-Based Search
- C. Based on full-text search
- D. Retrieval based on metadata

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 10

Which of the following descriptions about Fusioninsight CTBase is wrong?

- A. The read-write data interface of CTBase uniformly encapsulates the interface defined by the row, automatically merges and parses cold fields, and does not need to be merged and interpreted in the application
- B. The java API of CTBase provides a set of interfaces for HBase connection pool management, and internal connection sharing reduces the difficulty of client application development.
- C. CTBase provides a set of WebUI for metadata definition, provides a table design tool only for medical use, and reduces the difficulty of table design
- D. CTBase is a cluster table development framework based on HBase

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 11

Which of the following descriptions of cloud search service CSS is correct? (Multiple choice)

- A. Compatible with Elasticsearch APIs
- B. TB-level data retrieval is returned in seconds, and billion-level 100-dimensional star retrieval is returned in seconds
- C. Does not support the fusion search of rich text and text such as sound
- D. The business of expanding nodes and disks is not interrupted, and the business of updating the dictionary is not interrupted

Answer: A,B,D ([LEAVE A REPLY](#))

NEW QUESTION: 12

Which items can FusionInsight Manager perform health check on? (Multiple choice)

- A. Example
- B. service
- C. role
- D. Host

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 13

In Huawei Fusioninsight, the HBase table design tool, connection pool management and enhanced SDK can simplify the business development of complex data tables.

A. True

B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 14

Regarding the difference and connection between Spark SQL & Hive, which of the following statements is correct? (Multiple choice)

A. Spark SQL is compatible with most Hive syntax and functions

B. The execution engine of Spark SQL is Spark core, and the default execution engine of Hive is MapReduce

C. Spark SQL cannot use Hive's custom functions

D. Spark SQL relies on Hive metadata

Answer: A,B,D ([LEAVE A REPLY](#))

NEW QUESTION: 15

The flume architecture shown in the figure below is mainly used to import data from outside the cluster into the cluster.

A. True

B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 16

Which of the following statements about Fusioninsight HBase visual modeling is correct? (Multiple choice)

A. Visual modeling helps DBAs to model and design, lowering the threshold for using HBase

B. Qualifier HBase column: each column represents a certain attribute of business data

C. Realize the division of labor: DBAs focus on data table modeling, developers focus on user table names and columns used

D. Column user table column: each column represents a KeyValue

Answer: A,B,C ([LEAVE A REPLY](#))

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated**

and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPass.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 17

Which of the following descriptions about the key features of Flink is wrong?

- A. Compared with Flink, SparkStreaming has lower latency
- B. Compared with Streaming in FusionInght HD, Flink has higher throughput
- C. The Flink stream processing engine can provide support for both stream processing and batch processing applications
- D. checkpoint implements Flink's fault tolerance

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 18

Regarding RDD, which of the following statements is wrong?

- A. RDD has a lineage mechanism (Lineage)
- B. RDD is Spark's abstraction of basic data
- C. RDD is a read-only, partitionable distributed dataset
- D. RDD is stored on disk by default

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 19

"Group by" in Hive refers to dividing each data set into several small data sets through certain rules, and then grouping data for several small data sets

- A. False
- B. True

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 20

What is the resource management framework that comes with Spark?

- A. Mesos
- B. Docker
- C. Standalone
- D. YARN

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 21

Which of the following descriptions about Zookeeper features is wrong?

- A. Message update can only succeed or fail, there is no intermediate state.
- B. The number of Zookeeper nodes must be an odd number.

- C. Updates sent by the client are applied in the order in which they were sent.
- D. A message must be received by more than half of the servers. It will be successfully written to disk

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 22

Which command in Redis is to view the remaining lifetime of the key?

- A. persist
- B. expired
- C. expire
- D. ttl

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 23

If some information is allowed to be lost during message processing, which of the following methods can be used to close the message reliability processing mechanism? (Multiple choice)

- A. Set the parameter Config.Topology_ACKKRS to False
- B. When a spout sends a message, it uses an interface that does not specify a message ID for sending.
- C. Bolt uses Unanchor to send messages
- D. Set the parameter Config.Topology_ACKKRS to 0

Answer: B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 24

Which of the following types of data is not semi-structured data?

- A. JSON
- B. XML
- C. HTML
- D. Two-dimensional table

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 25

Which of the following descriptions about Hive features is wrong?

- A. Flexible and convenient ETL
- B. Only supports MapReduce computing engine
- C. Easy to use and easy to program
- D. Can directly access HDFS files and HBase

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 26

HiveServer compiles and parses HQL statements submitted by users into corresponding Yarn tasks, Spark tasks, or HDPS operations, thereby completing data extraction, conversion, and analysis.

A. False

B. True

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 27

When the client writes data to any Zookeeper node, Zookeeper needs to fully synchronize the data to other nodes before it can return to the client that the write is successful.

A. False

B. True

Answer: A [\(LEAVE A REPLY\)](#)

NEW QUESTION: 28

Hive supports normal view and materialized view.

A. True

B. False

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 29

The data flow of Flume can be sent to different Channels according to the information of headers.

A. True

B. False

Answer: A [\(LEAVE A REPLY\)](#)

NEW QUESTION: 30

Which of the following descriptions about the data management and architecture faced by traditional enterprises is correct? (Multiple choice)

A. Traditional databases do not consider the diversity of data, especially compatible with structured data, unstructured data and semi-structured data.

B. Traditional databases do not require high data processing time, while big data requires real-time data processing

C. Traditional databases are not suitable for processing PB-level data.

D. The operation and maintenance of massive data needs to ensure data stability, support high concurrency and increase server load.

Answer: A,B,C,D [\(LEAVE A REPLY\)](#)

NEW QUESTION: 31

In MapReduce programming, the code should not only describe what to do, but also how to do it.

A. True

B. False

Answer: ([SHOW ANSWER](#))

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPass.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 32

In FusionInsight HD products, a typical Kafka cluster includes several Producers, several Borkers, several Consumers and a Zookeeper cluster.

A. True

B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 33

Which of the following does not include the scheme authentication method of Zookeeper?

A. auth

B. world

C. sasl

D. digest

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 34

Hive in Fusioninsight includes two roles, HiveServer and MetaStore.

A. False

B. True

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 35

Regarding DataSet, which of the following statements is wrong?

A. DataSet can perform most operations without deserialization

B. DataSet performs sort, filter, shuffle and other operations that need to be deserialized

- C. DataSet is highly similar to RDD, and its performance is better than RDD
- D. DataSet is a strongly typed collection of domain-specific objects

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 36

Structured Streaming can define different data writing methods in the Output stage, including which of the following methods? (Multiple choice)

- A. Complete Mode
- B. Update Mode
- C. Append Mode
- D. General Mode

Answer: A,B,C ([LEAVE A REPLY](#))

NEW QUESTION: 37

Which of the following scenarios is HDFS not suitable for? (Multiple choice)

- A. Large number of small file storage
- B. Large file storage and access
- C. Random write
- D. Streaming data access

Answer: A,C ([LEAVE A REPLY](#))

NEW QUESTION: 38

The storage formats supported by Hive include? (Multiple choice)

- A. Parquet
- B. RCFile
- C. SequenceFile
- D. TextFile

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 39

Which of the following options belong to Fusioninsight data security? (Multiple choice)

- A. Operating system security hardening
- B. User rights authentication management
- C. Component Data Encryption
- D. Data integrity check

Answer: C,D ([LEAVE A REPLY](#))

NEW QUESTION: 40

Which is not the channel type of Flume?

- A. Memory Channel
- B. File Channel

C. JDBC Channel

D. HDFS Channel

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 41

Which statement about HIVE is wrong?

A. The best usage scenario of Hive is the batch processing job of large data sets

B. Hive is built on Hadoop based on static batch processing. Hadoop usually has high latency and requires a lot of overhead when submitting and scheduling

C. The Hive query operation process strictly follows the Hadoop MapReduce role execution model. Hive converts the user's HiveQL statement to the MapReduce Hadoop cluster through the interpreter.

D. Hive can realize low-latency and fast query on large-scale data sets

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 42

Which service process is used to manage the Region of HBase in the Hadoop platform?

A. DataNodes

B. HMaster

C. RegionServer

D. Zookeeper

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 43

Which application scenarios are HBase not suitable for?

A. Massive data (TB, PB) application scenarios

B. Semi-structured data application scenarios

C. High throughput application scenarios

D. Application scenarios that require full ACID characteristics

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 44

Which of the following options are Huawei Cloud MRS advantages? (Multiple choice)

A. Ease of use

B. Scalability

C. High reliability

D. high performance

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 45

In the FusionInsight HD cluster planning, what kind of scenarios is suitable for the unified deployment of management nodes, control nodes and data nodes?

- A. More than 200 nodes
- B. 100-200 nodes
- C. Below 30 nodes
- D. More than 100 nodes

Answer: (SHOW ANSWER)

NEW QUESTION: 46

Synchronize data between partition replication in Kafka. Copying data from the leader of the partition to the follower requires a thread (replicationFetcherThread), and the follower (a follower is equivalent to a consumer) actively pulls messages from the leader in batches, which greatly improves the throughput.

- A. False
- B. True

Answer: B (LEAVE A REPLY)

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPass.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 47

Which of the following nodes are included in the basic system architecture of HDFS?
(Multiple choice)

- A. NameNode
- B. Resource Manager
- C. Data Node
- D. Node Manager

Answer: (SHOW ANSWER)

NEW QUESTION: 48

All Zookeeper nodes can process requests

- A. True
- B. False

Answer: A (LEAVE A REPLY)

NEW QUESTION: 49

In the Loader of Fusioninsight HD, one connector can only be assigned to one job.

A. False

B. True

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 50

Which of the following options in the HBase command has the worst performance

A. list

B. put

C. Scan

D. get

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 51

The applicationMaster uses polling to apply for and receive resources from the ResourceManager through the RPC protocol.

A. True

B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 52

The fault-tolerant mechanism of Spark Streaming means that any partition error in RDD can be recalculated and generated according to its parent RDD. If the parent RDD is lost, you need to find the original data in the disk.

A. False

B. True

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 53

Which of the following descriptions about HBase's file storage module (HBase FileStream, HFS for short) is correct? (Multiple choice)

A. HFS encapsulates the interface between HBase and HDFS

B. Provide functions such as file storage, reading, and deletion for upper-layer applications

C. Applied to the upper layer of Fusioninsight HD

D. HFS is an independent module of HBase

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 54

Which of the following scenarios is Spark suitable for? (Multiple choice)

- A. Graph computing
- B. Interactive query
- C. Real-time stream processing
- D. Batch processing

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 55

Which instance must be deployed with the Loader instance in FusionInsight HD?

- A. DataNodes
- B. Resource Manager
- C. Node Manager
- D. RegionServer

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 56

What does Fusioninsight HD HBase use by default as its underlying file storage system?

- A. MapReduce
- B. Memory
- C. HDFS
- D. Hadoop

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 57

Flink is a computing framework combining batch processing and stream processing. Its core is a stream data processing engine for data classification and parallel computing.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 58

The data source types of Flink stream processing include? (Multiple choice)

- A. Collections
- B. Files
- C. Socket streams
- D. JDBC

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 59

In the Fusioninsight HD system, if a Loader node in the cluster is abnormal, and other services are normal, the normal use of the Loader service function will not be affected.

- A. True

B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 60

When deploying HBase in a Fusioninsight HD cluster with N nodes, it is recommended to deploy _____ HMaster processes and _ _____ RegionServer processes.

A. N, N

B. 2,2

C. 2, N

D. 3, N

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 61

What is the default resource scheduler for queues in YARN?

A. All of the above are correct

B. Capacity Scheduler

C. Fair Scheduler

D. FIFO scheduler

Answer: B ([LEAVE A REPLY](#))

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPass.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 62

The graphical health inspection tool Fusioninsight Tool is composed of FusionCare and SysChecker.

A. False

B. True

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 63

The RowKey of a table in HBase is divided into SplitKey 9, e, a, z. How many Regions does this table have?

A. 4

B. 6

C. 3

D. 5

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 64

What are the characteristics of Spark? (Multiple choice)

A. fast

B. Clever

C. spirit

D. light

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 65

Which components in the Fusioninsight HD platform support table and column encryption?
(Multiple choice)

A. HDFS

B. HBase

C. Flink

D. Hive

Answer: B,D ([LEAVE A REPLY](#))

NEW QUESTION: 66

The driven source is that flume periodically and actively acquires data.

A. True

B. False

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 67

The combine of MapReduce in the Map stage is a pre-grouping process, which is optional.

A. True

B. False

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 68

The index data of ElasticSearch can only be saved in the HDFS system.

A. True

B. False

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 69

Which of the following programming languages is implemented in Spark?

- A. C
- B. Scala
- C. JAVA
- D. C++

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 70

Which of the following is not a state store supported by Flink?

- A. FileStateBackend
- B. MemoryStateBackend
- C. RocksDBStateBackend
- D. FsStateBackend

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 71

RDD has Transformation and Action operators. Which of the following belongs to the Action operator?

- A. filter
- B. map
- C. reduceByKey
- D. saveAsTextFile

Answer: (SHOW ANSWER)

NEW QUESTION: 72

In Spark on YARN mode, the node without NodeManager cannot start the executor to execute the Task

- A. True
- B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 73

After using the Streaming client shell command to submit the topology in the Fusioninsight HD system, I use the Storm UI to view and find that the topology has not processed data for a long time. What may be the reason? (Multiple choice)

- A. Supervisor is a component that receives data in topology and then performs processing
- B. There is a logic error in the topology business, and it cannot run normally after submission
- C. The slots resources of the Supervisor are exhausted. After the topology is submitted, there are no slots to start the worker process.

D. The topology is too complex or the concurrency is too large, causing the worker startup time to be too long, exceeding the waiting time of the Supervisor

Answer: B,C,D (LEAVE A REPLY)

NEW QUESTION: 74

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

A. Each Partition corresponds to a log file at the storage level

B. Each Partition is an ordered and immutable message queue

C. Introduce the Partition mechanism to ensure the high throughput of Kafka

D. The number of Partitions determines the maximum number of concurrent consumers in each consumer group

Answer: A,B,C,D (LEAVE A REPLY)

NEW QUESTION: 75

Hive is a data warehouse infrastructure built on Hadoop. It provides a set of tools that can be used to extract, transform and load (ETL), which is a mechanism for storing, querying and analyzing large-scale data stored in Hadoop.

A. True

B. False

Answer: A (LEAVE A REPLY)

NEW QUESTION: 76

What is the default Block Size of HDFS in the Fusioninsight HD system?

A. 64M

B. 128M

C. 256M

D. 32M

Answer: B (LEAVE A REPLY)

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPASS.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 77

Regarding the Controller and NodeAgent in FusionInsightManager, which statement is correct?

- A. Each node must deploy a Controller
- B. NodeAgent accepts the commands issued by the Controller and executes specific actions
- C. Controller sends heartbeat to NodeAgent every 3 seconds
- D. NodeAgent is open source enhanced

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 78

What are the successful cases of Huawei Fusioninsight HD in the industry? (Multiple choice)

- A. Smart Park
- B. Digital government
- C. Finance
- D. Intelligent transportation

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 79

The functions of Nimbus in the Streaming architecture include? (Multiple choice)

- A. Task scheduling
- B. Start/shutdown worker process
- C. Monitor task execution status
- D. Resource allocation

Answer: A,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 80

Which of the following descriptions about the basic operations of Hive SQL is correct?

- A. You can specify the column separator when creating the table
- B. Create an external table using the external keyword. To create an ordinary table, you need to specify the internal keyword
- C. Creating an external table must specify location information
- D. When loading data to Hive, the source data must be a path of HDFS

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 81

Spark Streaming computing is based on DStream, which decomposes streaming computing into a series of short batch jobs.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 82

The data node is the working node of HDFS. What are its functions described below?
(Multiple choice)

- A. Records all operations on file creation, deletion, renaming, etc.
- B. Store and retrieve data according to the scheduling of the client or the name node
- C. Responsible for data storage and reading
- D. Periodically send the name node a list of the blocks it stores.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 83

The FusionInsight HD management node provides a dual-node HA mechanism. After the primary node fails, the primary and backup nodes need to manually implement Failover.

- A. False
- B. True

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 84

What configuration files can the Fusioninsight HD LLD configuration planning tool generate? (Multiple choice)

- A. The installation template file for the cluster
- B. Configuration files for HDFS and YARN
- C. Monitoring Alert Threshold Profiles
- D. The configuration file CheckNodes.Config required to execute Precheck

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 85

Which of the following sub-products does the Fusioninsight family include? (Multiple choice)

- A. Fusioninsight Miner
- B. Fusioninsight Farmer
- C. Fusioninsight HD
- D. GaussDB 200

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 86

In the Fusioninsight Hadoop cluster, through df-hT query on a certain node, the partitions seen include the following:

/var/log

/srv/BigData

/srv/BigData/hadoop/data5

/srv/BigData/solr/solrserver3

/srv/BigData/dbdataom

What is the planning combination of the optimal Raid level for the disks corresponding to these partitions?

- A. Raid0 Raid0 Raid0 Raid0 Raid0
- B. Non-Raid Non-Raid Non-Raid Non-Raid Raid1
- C. Raid1 Raid1 Non-Raid Non-Raid Raid1
- D. Raid0 Raid1 Raid0 Non-Raid1

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 87

ElasticSearch is strongly dependent on Zookeeper.

- A. False
- B. True

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 88

What physical resources does Redis mainly consume?

- A. Bandwidth
- B. Disk I/O
- C. Hard disk space
- D. memory

Answer: (SHOW ANSWER)

NEW QUESTION: 89

Which operations in Hive can be merged? (Multiple choice)

- A. JOIN
- B. SELECT
- C. UNION ALL
- D. GROUP BY

Answer: C,D ([LEAVE A REPLY](#))

NEW QUESTION: 90

In the Fusioninsight HD system, which LDAP data synchronization method is used?

- A. Data cross synchronization
- B. One-way synchronization
- C. Two-way synchronization
- D. Isolation is out of sync

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 91

Which of the following measures are used to protect the integrity of data? (Multiple choice)

- A. Rebuild the copy data of the failed data disk
- B. Safe Mode
- C. Cluster data balance
- D. Metadata Reliability Assurance

Answer: A,B,C,D ([LEAVE A REPLY](#))

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPass.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 92

Multiple channels can be configured in the properties.protertises configuration file of Flume to transmit data.

- A. False
- B. True

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 93

The "UNION ALL" operator in Hive is used to combine the result sets of two or more SELECT statements, and duplicate values are not allowed in the result set.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 94

In the flume architecture, a Source can connect to multiple channels.

- A. True
- B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 95

Which of the following is a node of ElasticSearch?

- A. Index
- B. Cluster

C. EsMaster

D. EsNode

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 96

Which of the following are the actual application scenarios of Kafka? (Multiple choice)

A. Log Collection

B. Resource management

C. Aggregated Statistical System Operational Data

D. Site Activity Tracking

Answer: A,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 97

Which of the following statements about Zookeeper's Leader election is correct? (Multiple choice)

A. When the number of instances is 8, 5 votes are required to become the leader node, and the disaster recovery capacity is 4

B. When the number of instances n is an odd number, assuming $n=2x+1$, x votes are required to become a leader node

C. When Zookeeper elects the leader, more than half of the votes are required

D. When the number of instances n is an odd number, assuming $n=2x+1$, then $x+1$ votes are required to become a leader node

Answer: C,D ([LEAVE A REPLY](#))

NEW QUESTION: 98

How many shards does an index library of Elasticsearch have by default?

A. 5

B. 4

C. 6

D. 3

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 99

Which of the following options is an important role for Spark? (Multiple choice)

A. Node manager

B. Resource Manager

C. DateNode

D. Driver

Answer: A,B,D ([LEAVE A REPLY](#))

NEW QUESTION: 100

A typical Kafka cluster includes several Producers, several Brokers, several Consumers and a ZooKeeper cluster.

A. True

B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 101

Fusioninsight HD preparations before installation, what steps are included? (Multiple choice)

A. Complete the hardware installation

B. Complete node host operating system installation

C. Prepare tools and software. For example, Putty, LLD, Fusioninsight HD software installation package, etc.

D. Prepare planning data. Such as network parameters and role deployment locations

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 102

What is the file format of data storage in HBase?

A. SequenceFile

B. HLog

C. HFile

D. TextFile

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 103

In a point-to-point message system, one or more consumers can consume the data in the queue, but a message can only be consumed once.

A. True

B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 104

The driver in Spark on YARN mode can only run on the client.

A. True

B. False

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 105

Hardware failure is considered normal. In order to solve this problem, HDFS designed a copy mechanism. By default, for a file, HDFS will save () copies?

A. 1

B. 4

C. 3

D. 2

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 106

The HDFS data reading process includes the following steps, please select the correct order.

A. If the end of the data block has been reached, then close the connection to the DataNode, and then search for the next data block again until all the data is read

B. After obtaining this input stream, the client calls the read method to read the data. The input stream selects the nearest DataNode to establish a connection and read data

C. Obtain the storage location of the data block corresponding to this file in NameNode by calling NameNode remotely through RPC

D. The client calls close to close the input stream

E. The client calls the open method of the FileSystem instance to obtain the input stream corresponding to the file

Answer: A,B,C,D,E ([LEAVE A REPLY](#))

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpspass.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 107

Which of the following sub-products does the Fusioninsight family include? (Multiple choice)

A. Miner

B. Manager

C. Farmer

D. HD

E. Libra

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 108

In the Hadoop system, which statement about the client uploading files to the HDFS file system is correct?

- A.** The client divides the file into multiple blocks, and writes them to each DataNode in sequence according to the address information of the DataNode
- B.** The file data of the client is passed to the DataNode through the NameNode
- C.** The client writes the entire file to each DataNode in sequence according to the address information of the DataNode, and then the DataNode divides the file into multiple blocks
- D.** The client only uploads data to one DataNode, and then the NameNode is responsible for block replication

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 109

Which of the following options are suitable for MapReduce?

- A.** Iterative calculation
- B.** Stream computing
- C.** Offline Computing
- D.** Real-time interactive computing

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 110

Flume supports monitoring and transferring new files in the directory, which can realize data transfer. Which type of source is described above?

- A.** spooling directory source
- B.** exec source
- C.** syslog source
- D.** http source

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 111

Fusioninsight is Huawei's unified platform for enterprise-level big data storage, query, and analysis. It can help enterprises quickly build massive data information processing systems, and discover new value points and business opportunities through real-time and non-real-time analysis and mining of massive information data.

- A.** False
- B.** True

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 112

Compared with open source Sqoop, what are the enhanced features of Loader? (Multiple choice)

- A.** high performance

- B. High reliability
- C. Security
- D. Graphical

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 113

What is the unit of data scale for traditional data processing?

- A. GB
- B. EB
- C. PB
- D. TB

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 114

Which two roles are responsible for the active and standby state management of the NameNode node and the merging of metadata files in HDFS?

- A. ZKFC and Standby NameNode
- B. ZKFC and master NameNode
- C. Primary NameNode and JournalNode
- D. Active NameNode and Standby NameNode

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 115

When the Flume process is cascaded, which of the following sink types are used to receive the messages sent by the previous hop Flume? (Multiple choice)

- A. Null sink
- B. Hive sink
- C. Avro sink
- D. Thrift sink

Answer: C,D ([LEAVE A REPLY](#))

NEW QUESTION: 116

The memory size allocated by YARN to Container in the Hadoop system can be set by the parameter yarn.app.mapreduce am.resource.mb

- A. True
- B. False

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 117

YARN's label-based scheduling, which of the following options is used for labeling?

- A. Node Manager

- B. Container
- C. APPMaster
- D. Resource Manager

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 118

Assume that HDFS only saves 2 copies when writing data. Then during the writing process, HDFS Client first writes data to DataNode1 and then writes data to DataNode2. ()

- A. False
- B. True

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 119

Which of the following are the characteristics of Elasticsearch? (Multiple choice)

- A. High fault tolerance
- B. high performance
- C. High scalability
- D. High reliability

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 120

Regarding the task selection of the capacity scheduler, which of the following statements is correct?

- A. Minimum queue level first
- B. Maximum Queue Level Priority
- C. Resource reclamation request queue priority
- D. The queue with the lowest resource utilization is given priority

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 121

Spark on YARN-client is suitable for production environments because you can see the output of APP faster.

- A. False
- B. True

Answer: A ([LEAVE A REPLY](#))

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated**

and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPass.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 122

Which of the following scenarios are not good for Flink components? (Multiple choice)

- A. Batch processing
- B. Iterative calculation
- C. Data storage
- D. Stream processing

Answer: B,C ([LEAVE A REPLY](#))

NEW QUESTION: 123

What are the performance bottlenecks of traditional data processing? (Multiple choice)

- A. Data storage costs are high
- B. Limited scalability
- C. Insufficient streaming data processing performance
- D. Missing batch data processing

Answer: A,B ([LEAVE A REPLY](#))

NEW QUESTION: 124

What services can Huawei MRS provide to customers? (Multiple choice)

- A. Statistical Analysis and Data Mining
- B. Security Control Based on Kerberos Certificate
- C. Offline and real-time data processing
- D. Multi-node deployment

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 125

Which service process is the Region of HBase managed by?

- A. HRegionServer
- B. ZooKeeper
- C. DataNodes
- D. HMaster

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 126

Which of the following statements about the segment file in Kafka Logs is correct? (Multiple choice)

- A. The message can be quickly located through the index information

- B.** Map all the index metadata to memory, which can avoid the IO disk operation of the index data of the segment file
- C.** The index file is stored sparsely, which can greatly reduce the space occupied by the metadata of the index file
- D.** Sparse storage, that is, the original complete data, only select multiple pieces at intervals for storage

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 127

What are the main features of HBase? (Multiple choice)

- A.** scalable
- B.** high performance
- C.** Column-oriented
- D.** High reliability

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 128

When the Fusioninsight Manager interface shows that the Hive service status is Bad, what are the possible reasons? (Multiple choice)

- A.** The MetaStore instance is unavailable
- B.** HDFS service is unavailable
- C.** The HBase service is unavailable
- D.** DBService is unavailable

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 129

In the MRS platform, which component does the Flume data flow not need to pass through in the node?

- A.** Source
- B.** Sink
- C.** Topic
- D.** Channel

Answer: **C** ([LEAVE A REPLY](#))

NEW QUESTION: 130

Where is the HBase metadata Meta Region routing information stored?

- A.** HMaster
- B.** Meta table
- C.** Zookeeper
- D.** Root table

Answer: **C** ([LEAVE A REPLY](#))

NEW QUESTION: 131

The following figure shows the label storage strategy of HDFS. Observe the following figure, which data nodes will the HBase data be stored on? (Multiple choice)

- A. DataNode E
- B. DataNode B
- C. DataNode F
- D. DataNode A

Answer: B,D ([LEAVE A REPLY](#))

NEW QUESTION: 132

Which of the following statements about the atomicity in the key features of ZooKeeper is correct?

- A. A message received by one server will be received by all servers
- B. The update can only complete or fail in full, not in part
- C. No matter which server in the cluster, the same view is displayed externally
- D. Updates sent by the client will be applied in the order they were sent

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 133

TaskSlot in Flink is mainly used for resource isolation, including memory resources and CPU resources.

- A. True
- B. False

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 134

After the Loader job in Fusioninsight HD is submitted to YARN, the job cannot be stopped manually.

- A. True
- B. False

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 135

Hadoop's HDFS is a distributed file system, which of the following application scenarios is suitable for data storage and management? (Multiple choice)

- A. Streaming data access
- B. Low latency read
- C. Large number of small file storage
- D. High capacity, high throughput

Answer: A,D ([LEAVE A REPLY](#))

NEW QUESTION: 136

The nodes in the ElasticSearch cluster are divided into master and slave.

- A. True
- B. False

Answer: A ([LEAVE A REPLY](#))

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpspass.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (**500** Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 137

Which of the following descriptions about Kafka partition offset is wrong?

- A. offset is a string of type String
- B. The position of each message in the file is called offset (offset)
- C. Consumers track records through (offset, partition, topic)
- D. Uniquely mark a message

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 138

During operation, the Kafka cluster directly depends on which of the following components?

- A. Zookeeper
- B. Spark
- C. HDFS
- D. HBase

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 139

Kerberos can only provide security authentication for services in the cluster.

- A. True
- B. False

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 140

Flume is used to collect data, and the basic unit of data transmitted by the device is?

- A. Event
- B. Block
- C. Packet
- D. Split

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 141

Spark Streaming has higher real-time performance than Storm.

- A. False
- B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 142

Which of the following descriptions about Hive log collection on the Fusioninsight Manager interface is wrong?

- A. You can specify a time period for log collection, for example, only collect logs from 2016-1-1 to 2016-1-10.
- B. You can specify the node IP for log collection, for example, only download the logs of a certain IP.
- C. You can specify an instance for log collection, such as specifying to collect metastore logs.
- D. Specific users can be designated for log collection, for example, only the logs generated by UserA are downloaded.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 143

In Kafka, which of the following statement about Producer sending data is wrong? (Multiple choice)

- A. Proucer, as a message producer, can directly write data to Zookeeper
- B. Producer is responsible for producing data, Consumer is responsible for consuming data, and a connection needs to be established between Producer and Consumer.
- C. Proucer connects to any normal Broker instance to generate data
- D. Proucer can send data to Broker and send data to Consumer

Answer: A,B,D ([LEAVE A REPLY](#))

NEW QUESTION: 144

In the task scheduling process of YARN, which of the following tasks is the responsibility of the Application Master?

- A. Start the Map or Reduce task
- B. Allocate Containers

- C. Applying for and receiving resources
- D. Set up the running environment for the task

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 145

When the Client of HDFS writes a file, the writing position of the first copy of the data is determined by the NameNode, and the writing position of other copies is determined by the Data Node.

- A. True
- B. False

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 146

In Fusioninsight HD, if you need to view the users and permission groups currently logged into HBase, what commands can you execute in the HBase shell?

- A. who
- B. use_permission
- C. get_user
- D. whoami

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 147

When MRS Loader creates a job, what is the role of the connector?

- A. Configure the connection method between the job and the external data source
- B. Make sure there is a conversion step
- C. Configure the connection method between the job and the internal data source
- D. Provide optimization parameters to improve data import and export performance

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 148

In the era of big data, which of the following challenges are faced by enterprises? (Multiple choice)

- A. There are problems such as noise, missing, and non-standard storage types in scattered data, which requires a lot of data preprocessing.
- B. Technological advances of competitors.
- C. The data is scattered among the various departments of the enterprise, and the same data is stored in different formats within each department.
- D. Data structure diversification.

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 149

What is the core module of spark?

- A. spark sql
- B. spark core
- C. mapreduce
- D. spark streaming

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 150

Which of the following description about HBase's Region Split process is wrong?

- A. The table will suspend service during the Split process
- B. During the Split process, the file is not really split, only the reference file is created
- C. Split splits a Region into two Regions in order to reduce the data size in the Region
- D. The Region that is split during the Split process will suspend the service

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 151

Which of the following statements about the interaction between Flink and other components is correct? (Multiple choice)

- A. Flink can send received components to Kafka
- B. The operation of Flink tasks depends on YARN for resource scheduling management
- C. The implementation of Flink's checkpoint depends on Zookeeper
- D. Flink reads and writes data in the HDFS file system

Answer: A,B,C,D ([LEAVE A REPLY](#))

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPASS.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (**500** Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 152

In YARN, which of the following tasks is the "slave" node responsible for? (Multiple choice)

- A. Supervise the life cycle management of the Container
- B. Unified management and allocation of all resources in the cluster
- C. Manage logs and ancillary services used by different applications
- D. Monitor the resource usage (memory, CPU, etc.) of each Container

Answer: A,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 153

What kind of computing tasks is the MapReduce component in Hadoop good at handling?

- A. Iterative calculation
- B. Real-time interactive computing
- C. Stream computing
- D. Offline Computing

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 154

Which of the following components must depend on Zookeeper to run? (Multiple choice)

- A. Spark
- B. HDFS
- C. YARN
- D. HBase

Answer: B,D [\(LEAVE A REPLY\)](#)

NEW QUESTION: 155

Fusioninsight tool is a set of health detection tools provided for technical support engineers and maintenance engineers. It can check the health status of cluster-related nodes and services, discover potential problems in the cluster in advance, and generate health check reports. It is convenient for technical support engineers and maintenance engineers to quickly understand the health status of the system.

- A. False
- B. True

Answer: B [\(LEAVE A REPLY\)](#)

NEW QUESTION: 156

Which of the following application scenarios can HBase be used for? (Multiple choices)

- A. Message/Order Storage
- B. User portrait
- C. Time series data storage
- D. Recommended by a friend

Answer: A,B [\(LEAVE A REPLY\)](#)

NEW QUESTION: 157

Which of the following statements about ZKFC is wrong?

- A. The Standby NameNode perceives the status of the Active NameNode through Zookeeper. Once the Active NameNode goes down, the Standby NameNode will perform the master promotion operation

B. ZKFC (ZKFailoverController) is used as a client of Zookeeper cluster to monitor the status information of NameNode

C. The ZKFC of HDFS NameNode connects to Zookeeper, and saves information such as the host name in Zookeeper

D. The ZKFC process needs to be deployed in the NameNode node and Zookeeper's Leader node

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 158

topology will automatically end running after the task is completed.

A. False

B. True

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 159

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

A. MemStore

B. HLog

C. HMaster

D. HFile

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 160

Which of the following options are the key features of Zookeeper? (Multiple choice)

A. Wait for irrelevance

B. Atomicity

C. Reliability

D. Eventual consistency

Answer: A,B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 161

When viewing the partition details of a topic in Kafka, which of the following commands should be used?

A. bin/kafka-topic.sh -create

B. bin/kafka-topic.sh -describe

C. bin/kafka-topic.sh -delete

D. bin/kafka-topic.sh -list

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 162

Which of the following statements about Transformation in Flink is correct? (Multiple choice)

- A. KeyBy is to group the source data according to the key to ensure that the metadata of the same key are divided into the same group
- B. FlatMap can segment text
- C. The Filter operation is to perform a Boolean function on each element
- D. You can set the time window through the window

Answer: A,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 163

The Fusioninsight HD cluster planned by a bank has 90 nodes. If three control nodes are planned, it is recommended to plan () data nodes in the cluster.

- A. 90
- B. 86
- C. 85
- D. 87

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 164

Which of the following descriptions about the characteristics of the Kafka Partition copy is correct?

- A. Replicas are in units of partitions. Each partition has its own slave copy of the master copy
- B. Follower synchronizes data from Leader by pulling
- C. Both consumers and producers read and write data from the Leader, and can also directly interact with the Follower
- D. The master copy is called Leader, and the slave copy is called Follower

Answer: A,B,D ([LEAVE A REPLY](#))

NEW QUESTION: 165

The main role of Zookeeper in distributed applications does not include which of the following options?

- A. Election of Master nodes
- B. Allocate cluster resources
- C. Store server information in the cluster
- D. Ensure data consistency on each node

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 166

Which of the following options does the type of time operation supported by Flink include? (Multiple choice)

- A. End time
- B. Acquisition time
- C. Event time
- D. Processing time

Answer: C,D ([LEAVE A REPLY](#))

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPass.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (**500 Q&As Dumps, 40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 167

When a RegionServer fails, the Region on the failed RegionServer is migrated by ().

Answer:

HMaster

NEW QUESTION: 168

Flink is a unified computing framework combining batch processing and stream processing. Its core is a stream data processing engine for data distribution and parallel computing.

- A. True
- B. wrong

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 169

Which of the following scenarios is Hive not suitable for?

- A. Real-time online data analysis
- B. Data mining, such as user behavior analysis, interest segmentation, and regional display
- C. Non-real-time analysis, such as log analysis, statistical analysis
- D. Data summary, such as daily and weekly user clicks, click ranking

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 170

FusionInsight HD cluster upgrade, which of the following descriptions is correct? (Multiple choice)

- A. During the upgrade process, manual operation of active and standby OMS switching is not allowed
- B. The root account passwords of all hosts in the cluster are consistent
- C. Capacity expansion is not allowed during the observation period
- D. Keep the network open to avoid abnormal upgrades due to network problems

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 171

The figure below shows the storage locations of files A, B, C, and D, where A and B are related, and their storage locations conform to the Colocation same distribution strategy.

- A. False
- B. True

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 172

Which of the following components does the platform architecture of Huawei's big data solution include? (Multiple choice)

- A. GaussDB 200
- B. Hadoop layer
- C. Datafarm layer
- D. FusionInsight Manager

Answer: B,C,D ([LEAVE A REPLY](#))

NEW QUESTION: 173

What log download methods does FusionInsight HD support? (Multiple choice)

- A. Download logs for all components of installed components
- B. Download the log of the specified host
- C. Download the log of a module of a single component
- D. Download logs for a specified period of time

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 174

The emergence of HFS solves the need to store a large number of small files (below 10MB) in HDFS. At the same time, it is also necessary to store some mixed scenes with large files (10MB or more)

- A. True
- B. False

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 175

In Hive, which of the following descriptions about bucketing is correct? (Multiple choice)

- A. Data can put different data into different buckets according to the way of buckets
- B. The advantage of bucketing is that it can obtain higher query processing efficiency and make sampling more efficient
- C. You can specify the number of buckets when creating a table
- D. Bucket cannot be sorted

Answer: A,B,C ([LEAVE A REPLY](#))

NEW QUESTION: 176

Which of the following is not a mandatory option when creating a Loader job?

- A. type
- B. name
- C. Priority
- D. connection

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 177

Among the following options, which one is correct about the meaning of Zookeeper reliability?

- A. Reliability means that a message received by a Server will be accepted by all Servers
- B. Reliability means that no matter which server displays the same view to the outside world
- C. Reliability is achieved through active and standby deployment
- D. Reliability means that the update can only succeed or fail, there is no intermediate state

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 178

Which of the following statements about the Ldap Server organizational model is correct? (Multiple choice)

- A. Each node in the LDAP Server directory tree is called an entry and has its own unique and distinguishable name DN (Distinguished Name).
- B. LDAP Server directory information is organized and stored based on a tree structure.
- C. The root of the LDAP Server directory tree generally defines the domain name DC (Domain Component).
- D. Specific objects cannot be defined under organizational units.

Answer: A,B,C ([LEAVE A REPLY](#))

NEW QUESTION: 179

Normally, tables stored in text files in Hive use carriage returns as their line separators. In Huawei Fusioninsight Hive, you can specify the input and output format processing classes for table data.

- A. True

B. wrong

Answer: (SHOW ANSWER)

NEW QUESTION: 180

The overall process of kafka Consumer writing data is that the Consumer connects to the Leader Broker where the specified Topic Partition is located, and uses it to actively obtain messages from Kafka.

A. False

B. True

Answer: B (LEAVE A REPLY)

NEW QUESTION: 181

The MOB data in HBase is directly stored in HFile format on HDFS, and then the address information and size information of this file are stored as value on the store that manages HBase, and these files are managed centrally through tools. In this way, the compaction and split frequency of HBase can be greatly reduced and the performance can be improved.

A. False

B. True

Answer: B (LEAVE A REPLY)

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPASS.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

Valid H13-711-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711-ENU Exam! BraindumpsPass.com now offer the **newest H13-711-ENU exam dumps**, the BraindumpsPass.com H13-711-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711-ENU dumps with Test Engine here: <https://www.braindumpsPASS.com/Huawei/H13-711-ENU-practice-exam-dumps.html> (500 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)