

Huawei.H13-711_V3.0-ENU.v2022-10-26.q316

Exam Code:	H13-711_V3.0-ENU
Exam Name:	HCIA-Big Data V3.0
Certification Provider:	Huawei
Free Question Number:	316
Version:	v2022-10-26
# of views:	4708
# of Questions views:	3160
https://www.exam-tests.com/H13-711_V3.0-ENU-exam/Huawei.H13-711_V3.0-ENU.v2022-10-26.q316.html	

NEW QUESTION: 1

In the FusiinsightHD platform, which components support list encryption?

- A. HBase
- B. HDFS
- C. Flink
- D. Hive

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

NEW QUESTION: 2

What modules does the Spark framework include?

- A. Spark GraphX
- B. Spark Mllib
- C. Spark Streaming
- D. Spark SQL

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

NEW QUESTION: 3

To enable the log aggregation function of the Yam component in the Hadoop platform, which parameter needs to be configured?

- A. yarn.nodemanager.local-dirs
- B. yarn.nodemanager.log-dirs
- C. yarn.log-aggregation-enable
- D. yarn.acl.enable

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

NEW QUESTION: 4

Which of the following options belong to Hive's data storage model?

- A. All the above are correct
- B. Bucket
- C. Partition
- D. Table
- E. Database

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

NEW QUESTION: 5

Spark and Hadoop are not suitable for iterative calculation computing.

- A. False
- B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 6

Which of the following descriptions of Zookeeper features is wrong?

- A. Message update can only succeed or fail, there is no intermediate state.
- B. A message needs to be received by more than half of the servers. It can be successfully written to disk
- C. Updates sent by the client will be applied in the order in which they were sent.
- D. The number of Zookeeper nodes must be an odd number.

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

NEW QUESTION: 7

server.channels.ch1.transactioncapacity represents the transaction size, that is, the number of events that the current channel supports for transaction processing. It can be set to the same size as the batchsize of the source, but not less than the batchsize.

- A. True
- B. False

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

NEW QUESTION: 8

Which of the following commands can be used to create node data?

- A. get /node
- B. create /node
- C. set/node data
- D. ls/node

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 9

Regarding the comparison between Hive and traditional data warehouse, which of the following descriptions is wrong?

- A. Hive is based on HDFS storage. Theoretically, the storage capacity can be expanded infinitely. The storage capacity of traditional data warehouses will have an upper limit
- B. Since Hive is based on a big data platform, query efficiency is faster than traditional data warehouses
- C. Since Hive data is stored in HDFS, it can ensure high fault tolerance and high reliability of data
- D. Hive metadata storage is independent of data storage, thereby decoupling metadata and data. High flexibility, while traditional data warehouse data application is single, low flexibility

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

NEW QUESTION: 10

Which of the following descriptions about bucketing is correct in Hive? (Multiple choice)

- A. Unsortable in the bucket
- B. Different data can be put into different buckets according to the way of buckets
- C. The advantage of bucketing is to obtain higher query processing efficiency and make sampling more efficient
- D. You can specify the number of buckets when creating a table

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 11

Which of the following languages are supported by the Flink streaming data processing interface Datastream API? (Multiple choice)

- A. Scala
- B. C Language
- C. Python
- D. Java

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 12

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

- A. True
- B. False

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

NEW QUESTION: 13

In data during stream processing, each thing piece Time Which of the following three can be divided into?

- A. ingestion time, the time the event arrives at the stream processing system
- B. event time, the time when the event occurred

- C. Processing time,eventquiltssystem processing time
- D. delay time,Right nowthingincident to processingextendlate time

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 14

When a Region in HBase performs a Split operation, at what stage does the process of separating an HFile file into two regions occur?

- A. Split process
- B. Flush process
- C. Compaction process
- D. HFile separation process

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

NEW QUESTION: 15

Which of the following extensions belong to Elasticsearch? (Multiple choice)

- A. head
- B. bigdesk
- C. IKAAnalyzer
- D. hadoop

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

NEW QUESTION: 16

Which of the description of HIVE is wrong?

- A. The best use scenario for Hive is batch processing 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 for submission and scheduling.
- C. The Hive query operation process strictly follows the Hadoop MapReduce execution model. Hive converts the user's HiveQL statements into 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](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

NEW QUESTION: 17

What log download methods does FusionInsight HD support?

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

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

NEW QUESTION: 18

What is the core module of spark?

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

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

NEW QUESTION: 19

In the Fusioninsight Manager interface, which of the following options is not included in the operation of the loader?

- A. Configure loader parameters
- B. Switch the active and standby Loader nodes
- C. Start the loader instance
- D. View loader service status

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

NEW QUESTION: 20

On the Fusioninsight Manager interface, when an alarm of insufficient disk capacity of Kafka is received, and the cause of the alarm has been eliminated from the hard disk hardware failure, the system administrator needs to consider expanding the capacity to solve this problem.

- A. False
- B. True

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

NEW QUESTION: 21

What is the default resource scheduler in YARN?()

- A. FIFO scheduler
- B. Fair scheduler none of the
- C. capacity scheduler

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

NEW QUESTION: 22

The following is for the LdapSearer organization?Type statement is correct

- A. Specific objects cannot be defined under organizational units.
- B. The tree of the LdapSearer directory tree generally defines the domain name dc (Dcmain Compen?)
- C. Each node in the LdapSearer directory is called?day, and?has its own uniquely distinguishable name?IN (Distingarl shed Mama)
- D. LdapSearer directory information is based on? ?structure to carry out?woven and stored.

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

NEW QUESTION: 23

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 24

Which is the correct description of the Loader job in FusionInsight HD?

- A. After Loader submits the job to Yam for execution, if a Mapper task fails to execute, it can automatically retry
- B. After the execution of Loader job fails, the dam data will be generated, which needs to be manually cleared by the user
- C. After Loader submits a job to Yam for execution, no other jobs can be submitted until the job is completed.
- D. After the Loader submits the job to Yam for execution, if the Loader service is abnormal at this time, the job execution fails.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 25

In the process of using Flume to transmit data, in order to prevent data loss due to the restart of the Flume process. Which of the following Channel types can be used

- A. HDES Channel
- B. File Channel
- C. JDBC Channel
- D. Memory Channel

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

NEW QUESTION: 26

When the Loader of Fusioninsight HD creates a job, what is the function of the connector?

- A. Provide optimization parameters to provide data import and export performance
- B. Configure how jobs connect to internal data sources
- C. Determine what conversion steps are in place
- D. Configure how jobs connect to external data sources

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

NEW QUESTION: 27

When importing data into a Hive table, the validity of the data will not be checked, but only when the data is read.

- A. True
- B. False

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

NEW QUESTION: 28

Sparki, like Hadoop-, is not suitable for generational computing.

- A. True
- B. False

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

NEW QUESTION: 29

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

- A. NameNode
- B. Data Node
- C. NodeManager
- D. ResourceManager

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

NEW QUESTION: 30

Compared with MapReduce computing, Spark uses _____ resources to obtain _____ times the performance?

	Madoop	Spark	Spark
Data size	102.5TB	102TB	1000TB
Time consuming	72 minutes	23 minutes	234 minutes
Number of nodes	2100	206	190
Cores	50400	6592	6080
Rate	1.4TB/min	4.27TB/min	4.27TB/min
Rate/node	0.67GB/min	20.7GB/min	22.5GB/min
Daytona gray Category Sorting Basis Rules	Yes	Yes	Yes

- A.** 1/10, 3
B. 1/10, 4
C. 1/8, 3
D. 1/8, 4

Answer: A (LEAVE A REPLY)

NEW QUESTION: 31

Regarding the alarm about insufficient disk capacity of Kafkat, which of the following analysis is incorrect for the possible reasons?

- A.** Unreasonable business planning results in uneven data distribution and makes some disks reach the upper limit of usage
B. Caused by the failure of the rocker node
C. The data storage time is configured too long, and the accumulated data reaches the upper limit of the disk usage.
D. The disk configuration used to store Kafka data (such as the number of disks, size, etC.) cannot meet the current business data stream declaration, resulting in the disk usage reaching the upper limit

Answer: B (LEAVE A REPLY)

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

NEW QUESTION: 32

There are two ways to clean up Kafka logs: delete and compact. The default value is delete.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 33

kafka logs segment file naming rule is that the first segment of the global partition starts from 0, and each subsequent segment file is named the maximum offset (offset message number) of the previous global partition.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 34

In FusionInsight HD cluster planning, manage nodes&control node&What kind of scenarios is the unified data node deployment solution suitable for?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 35

From the perspective of life cycle, which stages do data go through? (Multiple choice)

- A. Data presentation
- B. Data collection
- C. Data management
- D. Data storage
- E. Data analysis

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 36

How many nodes are required to install the Nimbus role of the Streaming component of FusionInsight HD?

- A. 1
- B. 3
- C. 2

D. 4

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

NEW QUESTION: 37

Which of the following options are advantages of FusioninsightSparkSQL?

- A. Support data insert, update and delete
- B. SQL Compatibility
- C. everything aboveA. True
- D. Stable and high performance
- E. long-term stability test

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 38

Which of the following options is a process that MapReduce must have? (Multiple choice)

- A. Partition
- B. Reduce
- C. Combine
- D. Map

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 39

Which of the following options are correctHBaseProcess description of reading and datastatedcorrect?

- A. The client can maintain the visited regionLocation information to speed up the query
- B. Clientcannot maintain visitedregionLocation information to speed up the query
- C. The client needs to read the datato HMaster Involvement
- D. customersNo need to read data on the sideHMaster participation

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

NEW QUESTION: 40

When FusionInsight Manager is connected to an external management platform, which interfaces are supported? (Multiple choice)

- A. Syslog
- B. SNMP
- C. FTP
- D. VPN

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

NEW QUESTION: 41

What information is included in calling the Zookeeper client command?

- A. server nickname

- B. username
- C. The port number
- D. IP address

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

NEW QUESTION: 42

In the MRS platform, which component does the F1ume data flow not need to pass through in the node?

- A. Source
- B. Sink
- C. LTopic
- D. Channel

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

NEW QUESTION: 43

What is the module used to manage the active and standby states of the Loader Server process in Loader?

- A. Job Scheduler
- B. ResourcesManager
- C. Job Manager
- D. HAManager

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

NEW QUESTION: 44

Which of the following are the characteristics of ElasticSearch?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 45

Fusioninsight Manger supports REST interface, SNMP interface and SYSLOG interface externally

- A. True
- B. False

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

NEW QUESTION: 46

What steps are included in the preparation for Fusioninsight HD installation?

- A. Complete the hardware installation

- B. Complete the node host OS installation
- C. Prepare planning data. such as network parameters and role deployment locations
- D. Prepare tools and software. Such as Putty, LLD, Fusioninsight HD software installation package, etC.

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

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpsPASS.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 47

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

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

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

NEW QUESTION: 48

The underlying data of HBase exists in the form of?

- A. row storage
- B. column store
- C. real-time storage
- D. keyvalue

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

NEW QUESTION: 49

MapReduce - When a task fails, the task can be recalculated through a retry mechanism.

- A. False
- B. True

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

NEW QUESTION: 50

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

- A. Real-time stream processing

- B. Graph calculation
- C. Batch processing
- D. Interactive query

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 51

Which of the following options is correct about the meaning of Zookeeper reliability?

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

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

NEW QUESTION: 52

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

- A. Data integrity check
- B. Component data encryption
- C. User authority authentication management
- D. Operating system security hardening

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 53

What is the resource management framework that comes with Spark?

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

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

NEW QUESTION: 54

When importing data to the Hive table, the validity of the data will not be checked, only when the data is read.

- A. False
- B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 55

Which of the following HDFS commands can be used to check the integrity of data blocks?

- A. HDFS balancer -threshold 1
- B. HDFS fsck -delete

- C. HDFS fsck /
- D. HDFS dfsadmin -report

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 56

When Fusioninsight HD deploys Solr, if you choose to store the index on a local disk, it is recommended to mount a disk separately to the data directory of each SolrServer instance, and configure the disk as RAID0 or RAID5.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 57

What are the characteristics of Spark? (Multiple choice)

- A. Light
- B. Spirit
- C. skillful
- D. Fast

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 58

Which of the following nodes in HDFS is ZKFC process deployed on? (Multiple choice)

- A. DataNode
- B. All of the above are wrong
- C. Standby NameNode
- D. Active NameNode

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

NEW QUESTION: 59

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

- A. Metastore

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 60

Which component controls the primary and secondary arbitration of the NameNode in HDFS?

- A. HDFS Client
- B. ResourceManager
- C. NodeManager
- D. ZooKeeper Failover Controller

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

NEW QUESTION: 61

In the Fusioninsight HD product, a typical Kafka cluster contains several producers, several consumers, and a zookeeper cluster.

- A. False
- B. True

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

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumps.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 62

Which of the following operations cannot be recorded in the Fusioninsight HD system audit log ()

- A. Start and stop the service instance
- B. delete service instance
- C. Query history monitoring
- D. Manually clear the camp

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

NEW QUESTION: 63

Which of the following contents can be viewed in the Loader historical job record?

- A. dirty data link
- B. Error lines/number of files
- C. Job start/run time
- D. job status

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

NEW QUESTION: 64

Which of the following options belong to the advantages of HUAWEI CLOUD MRS? (Multiple choice)

- A. High performance
- B. Scalability
- C. Ease of use
- D. High reliability

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

NEW QUESTION: 65

When installing a FusionInsight HD cluster in safe mode, which components must be installed?

- A. LDAPServer
- B. Zookeeper
- C. HDFS
- D. KrbServer

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

NEW QUESTION: 66

What is the meaning of this command "ALTER TABLE employee1 ADD columns (columnlstring);" in Hive?

- A. Create a table
- B. Delete table
- C. Modify the file format
- D. Add columns

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

NEW QUESTION: 67

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

- A. When the number of instances n is an odd number, assuming $n=2x+1$, x votes are required to become the leader node
- B. When Zookeeper elects a leader, more than half of the votes are required
- C. When the number of instances n is an odd number, assuming $n=2x+1$, $x+1$ votes are required to become the leader node
- D. When the number of instances is 8, 5 votes are required to become the leader node, and the disaster tolerance is 4

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

NEW QUESTION: 68

When creating a Loader job, in which of the following steps can the number of Maps be set?

- A. output
- B. input settings
- C. Basic Information
- D. convert

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

NEW QUESTION: 69

Sparki - like Hadoop - is not suitable for iterative computing.

- A. True

B. False

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

NEW QUESTION: 70

The figure below shows the configuration of HDFS tiered storage. If the number of copies of a block is 4. Which of the following statements is wrong?

A. All 4 blocks will be written to RAM_DISK

B. The first block copy will be written to RAM_DISK, the rest of the replicas are written to DISK

C. If it is a copy other than the first copy that fails to write, try to write to the storage type specified by "Alternative storage policy for copy"

D. If the first block copy is written to RAM_If DISK fails, try to write to the storage type specified by "Alternative Storage Policy"

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

NEW QUESTION: 71

In FusionInsight HD products, which statement about Kafka components are correct? (Multiple choice)

A. Unavailability of the ZooKeeper service will cause the Kafka service to be unavailable

B. When deleting a topic, you must ensure that the Kafka service configuration delete.topic.enable is set to true

C. Kafka installation and operation log save path is /srv/Bigdata/kafka

D. Must use admin user or Kafka admin group user to create Topic

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

NEW QUESTION: 72

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

A. filter

B. saveAsTextFile

C. reduceByKey

D. map

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

NEW QUESTION: 73

FusionInsight HD uses HBase client to write 10 pieces of data in batches.

A Regionserver node contains 2 Regions of the table, A and B respectively.

Two of the 10 pieces of data belong to A and 4 belong to B.

How many RPC requests need to be sent to the Regionserver to write these 10 pieces of data?

A. 3

B. 4

C. 2

D. 1

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 74

Fusioninsight Manger externally supports REST interface, SNMP interface, and SYSLOG interface.

A. False

B. True

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

NEW QUESTION: 75

Hive is not suitable for which of the following scenarios?

A. Data summary. For example, clicks per user per day, click ranking

B. Real-time online data analysis

C. Non-real-time analysis. Such as log analysis, statistical analysis

D. Data mining. Such as user behavior analysis, interest analysis, regional display

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

NEW QUESTION: 76

Comparing Spark and MapReduce, which of the following statements are correct?

A. Shuffle is not a necessary process in Spark tasks

B. Spark can cache intermediate results in memory reducing disk IO

C. Intermediate results of Spark tasks can only be stored in memory

D. Shuffle is a necessary process in Spark tasks

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

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpsPass.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 77

Which of the following scenarios in HBase will trigger the Flush operation?

A. When the number of files in WALs reaches a threshold

B. The total size of the MemStore in the Region has reached the preset Flush Size threshold

C. HBasePeriod refreshMemstore, silentthink period is 1Hour

D. The ratio of the total memory occupied by the MemStore to the total memory of the RegionServer exceeds the preset threshold size

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

NEW QUESTION: 78

If HDFS is deployed with a single name node, what limitations may exist? (multiple choice)

A. The amount of leaf swallowing of the entire distributed file system is limited by the throughput of a single name node

B. The number of objects (files, blocks) that a name node can hold is limited by the size of the memory space

C. Failure of this unique name node will render the entire cluster unavailable

D. It will affect the functional implementation of the upper-level components based on HDFS

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

NEW QUESTION: 79

Which of the following descriptions about the deployment of big data components on Kunpeng and x86 servers is correct? (Multiple choice)

A. Support mixed deployment of Kunpeng server and ordinary x86 server in a single cluster

B. No shortcomings in performance

C. Achieve autonomous controllability of some equipment

D. Single component (such as HDFS) supports mixed deployment of Kunpeng server and x86 server

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 80

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

A. False

B. True

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

NEW QUESTION: 81

If using Redis pairAn ordered collection is sorted, which data type?

A. sorted set

B. Sat

C. Hash

D. string

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

NEW QUESTION: 82

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

- A. True
- B. False

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

NEW QUESTION: 83

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

Answer:

priority

NEW QUESTION: 84

The FusionInsight HD cluster contains many kinds of services, and each service consists of thousands of roles. Which of the following are the roles of the service?()

- A. HDFS
- B. NameNode
- C. Hbase
- D. DataNode

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

NEW QUESTION: 85

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

- A. HDFS
- B. HBase
- C. Flink
- D. Hive

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

NEW QUESTION: 86

What are the key features of Streaming in Huawei's big data product Fusioninsight HD?

- A. message reliability
- B. flexibility
- C. Scalability
- D. Disaster recovery capability

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

NEW QUESTION: 87

Which of the following statements about ZKFC is wrong?

- A. The ZKFC process needs to be deployed in NameNode's node and Zookeeper's Leader node
- B. Standby NameNode perceives the status of Active NameNode through Zookeeper. Once the Active NameNode is down, the Standby NameNode will perform the main upgrade operation.

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 88

Which of the following optimization methods belong to Redis? (Multiple choice)

A. Turn off persistence

B. Reduced key value

C. Limit Redis memory size

D. slowlog configuration

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 89

Which of the following are the functions that Spark can provide?

A. Distributed file system

B. Distributed memory computing engine

C. Unified scheduling of cluster resources

D. Stream processing capabilities

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

NEW QUESTION: 90

Which of the following descriptions about the deployment of big data components on Kunpeng and X86 servers is correct?

A. No shortcomings in performance

B. Realize the autonomous control of some equipment

C. Single component (for example: HDFS) supports mixed deployment of Kunpeng server and X86 server

D. Supports mixed deployment of Kunpeng servers and ordinary X86 servers in a single cluster

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 91

The HDFS data reading process includes the following steps, please choose the correct order. (Drag title) Order: ECADB

A. 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.

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

C. The location where the data block corresponding to this file in NameNode is obtained by calling NameNode remotely through RPC

D. If the end of the data block has been reached. Then close the connection with this DataNodel, and then re-find the next data block. until all data is read.

E. The client calls close. to close the input stream.

Answer: ([SHOW ANSWER](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumps.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 92

What information does a Key Value format in the HBase data file HFile contain?

- A. Timestamp
- B. Key
- C. KeyType
- D. Value

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 93

Through the unified user management system in the big data platform, unified management of users, roles and organizations of various open-source component application systems in the platform can be realized, and cross-domain single sign-on logout and unified identity authentication function among various application systems can be realized.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 94

When a RegionServer fails, the Region on the failed RegionServer is migrated. (fill in the blank)

Answer:

HMaster

NEW QUESTION: 95

The following description about Flink barrier is wrong ()

- A. barrier) is periodically inserted into the data stream and flows with it as part of the data stream
- B. Barriers are at the heart of Flink snapshots

C. A barrier compares the snapshot data of this cycle with the next one. Periodic snapshot data is separated

D. When inserting a barrier, it will temporarily block the data flow

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

multiple choice

NEW QUESTION: 96

All Kafka messages will be persisted to the hard disk. At the same time, Kafka ensures data reliability by setting Replication on Topic Partitions.

A. True

B. False

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

NEW QUESTION: 97

Which of the following conversion rules can be implemented by Loader?

A. long time conversion

B. splice conversion

C. Incremental conversion

D. Null conversion

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

NEW QUESTION: 98

What kinds of Kafka message transmission guarantees are usually provided?

A. At least once (At Least once)

B. At Most Three Times

C. At Most once

D. Exactly Once

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

NEW QUESTION: 99

The following descriptions about FusionInsight multi-tenant management are correct:

A. 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

B. Tenant resource isolation based on container mechanism, ensuring tenant SLA

C. One-stop management of tenant resources, computing resources (CPU/memory/IO), storage resources (HDFS), service resources (HBase....)

D. Visual multi-level tenant management, matching with the enterprise organizational structure, simplifies system resource allocation and management

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

NEW QUESTION: 100

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

- A. False
- B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 101

How many copies of HDFS Block in FusionInsight HD system are saved by default?

- A. 2
- B. Not sure
- C. 1
- D. 3

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

NEW QUESTION: 102

FusionInsight Hadoop:In the cluster, the cluster size has 70 nodes. If the recommended deployment scheme is adopted, what partitions may exist on the management node?

- A. /srv/BigData/zookeeper
- B. /srv/BigData/dbdata-om
- C. /srv/BigData/hadoop/data5
- D. /srv/BigData
- E. /srv/BigData/jumalnode

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 103

Which of the following functions can be achieved by the kafka-cluster mirroring tool?

- A. None of the above is correct
- B. Kafka single cluster data recovery
- C. Kafka single cluster data backup
- D. Kafka cluster data synchronization solution

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 104

Which of the following descriptions of the characteristics of Kafka Partition copies are correct?
(Multiple choice)

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

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

NEW QUESTION: 105

Which of the following statements about the comparison of Sparkstreamingi and Streaming is incorrect?

- A. The execution logic of Streamingl is to start immediately, and then recycle after running.
- B. The throughput of SparkStreaming is about 2-5 times that of Streaming
- C. SparkStreaming:It is a micro-batch framework, and events need to be processed when they are exhausted to a certain amount
- D. SparkStreaming event processing latency is higher than Streaming

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 106

Which fields can ElasticSearchl's balancing algorithm be applied to?

- A. volume reduction
- B. Expansion
- C. Import Data
- D. export data

Answer: ([SHOW ANSWER](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:
https://www.braindumpsPASS.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html
(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 107

Flink is suitable for applications with high concurrency processing data and millisecond delay.

- A. False
- B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 108

In order to ensure the reliability of snapshot storage of streaming applications, where are the snapshots mainly stored?

- A. In a single-machine database with high reliability
- B. in the local file system

- C. In the memory of the JobManager
- D. in HDFS

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

NEW QUESTION: 109

One of the goals of the data center technology architecture is to be able to solve the disadvantages of traditional data warehouses that tend to become data islands and big data platforms that tend to become data swamps.

- A. False
- B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 110

In Huawei's big data solution b, which of the following components are included in the bottom layer of hadoop?

- A. miner
- B. fink
- C. spark
- D. hive

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

NEW QUESTION: 111

With the continuous increase in the volume of big data, the requirements for the physical security of data storage are getting higher and higher, and higher requirements are put forward for multiple copies of data and disaster tolerance mechanisms.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 112

Spark is a memory-based computing engine. All data during the running of Spark programs can only be stored in memory.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 113

In the Fusioninsight Hadoop cluster, through df-hT query on a 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 planned combination of the best Raid level for the disks corresponding to these partitions?

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

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

NEW QUESTION: 114

The smallest processing unit of HBase is Region, and the routing information between User Region and RegionServer is stored in Zookeeper.

- A. False
- B. True

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

NEW QUESTION: 115

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

- A. False
- B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 116

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

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

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

NEW QUESTION: 117

Regarding Flume, which of the following statements is false?

- A. Data transfer between Flume cascade nodes does not support compression
- B. There is a need for encryption inside processes such as Source to Channel to Sink
- C. Data transfer between Flume cascade nodes supports encryption
- D. Flume supports multi-cascading and multiplexing

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

NEW QUESTION: 118

HashThe value of type stores the mapping of fields and field values, and fields and field values can only be()

Answer:

string

NEW QUESTION: 119

Spark Streaming? The fault tolerance mechanism means that if any partition in the RDD fails, it can be recalculated and generated according to its parent RDD. If the parent RDD is lost, you need to go to the disk to find the original data.

A. False

B. True

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

NEW QUESTION: 120

The smallest processing unit of HBase is Region. Where is the routing information between User Region and Region Server stored in?

A. Master

B. HDFS

C. meta table

D. Zookeeper

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 121

In the Huawei big data solution, which of the following components are included in the hadoop layer? (Multiple choice)

A. Flink

B. Hive

C. Spark

D. Miner

Answer: ([SHOW ANSWER](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

NEW QUESTION: 122

Which statement is correct about worker (worker process), Executor (thread) and task (task)?

- A. Each worker can only run Executors (threads) for one topology
- B. Each worker can run multiple Executors (threads)
- C. Each Executor (thread) can run different components (spout or bolt) of the task (task)
- D. Each Executor (thread) can run a number of tasks (tasks)

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 123

Regarding the basic operation of Hive table building, which of the descriptions is correct?

- A. Once the table is created, you can no longer modify the column name
- B. Once the table is created, no more new columns can be added
- C. Once the table is created, the table name cannot be modified
- D. When creating an external table, you need to specify the external keyword

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

NEW QUESTION: 124

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

- A. False
- B. True

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

NEW QUESTION: 125

On the FusionInsight Manager interface, when a warning about insufficient kafka wipe disk capacity is received, and the cause of the warning has been eliminated from the hard disk hardware failure, the system administrator needs to consider expanding the capacity to solve this problem.

- A. False
- B. True

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

NEW QUESTION: 126

What data does FusionInsightManager regularly back up?

- A. DBService
- B. NameNode
- C. OMS
- D. LDAPServer

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

NEW QUESTION: 127

Hive does not check whether the data conforms to the schema during load. Hive follows the schema on read (read mode). Only when reading is Hive checks and parses specific data fields and schemas.

- A. True
- B. False

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

NEW QUESTION: 128

What is the physical storage unit of Region in HBase?

- A. ColumnFamily
- B. Column
- C. Region
- D. Row

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

NEW QUESTION: 129

BelowcloseAtWhat does Flick's fault tolerance mechanism describe errors?

- A. SaveThe point mechanism ensures thatAfter the program is upgraded or modified and developed, etC. Also able to restore from a saved state bit healthy boot
- B. SavepointmechanismGuaranteed in real timeprogram operationrun time, even ifunexpectedlycan alsoEnterRowsincelrecover
- C. Chackpointmechanism guaranteesAfter the program is upgraded or modified and developed, etC. Also able to restore from a saved state bit healthy boot
- D. Chackpoint mechanismGuaranteed in real timeprogram operationrun time, even ifunexpectedlycan alsoEnterRowsincelrecovercomplex

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

NEW QUESTION: 130

What are the main features of the YARN capacity scheduler

- A. flexibility
- B. Capacity Guarantee
- C. multiple tenancy
- D. Dynamically update configuration files

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 131

During the deployment of FusioninsightHD , the CheckNode. Config configuration file must be called when performing precheck to check each node.

A. False

B. True

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

NEW QUESTION: 132

Which option is wrong in the following description of Hive log collection on the Fusioninsight Manager interface?

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 the instance for log collection, for example, specify the log collection of metastore.

D. Specific users can be designated for log collection, for example, only logs generated by UserA can be downloaded.

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

NEW QUESTION: 133

Colocation file-level colocation realizes fast file access, avoiding a large amount of network overhead caused by data relocation.

A. True

B. False

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

NEW QUESTION: 134

In the era of big data, which of the following challenges companies face? (Multiple choice)

A. Scattered data has problems such as noise, missing, irregular storage type, etc., which requires a lot of data preprocessing.

B. The data is scattered among various departments of the enterprise, and the same data is inconsistent in the internal storage format of each department.

C. The technological advancement of competitors.

D. The data structure is diversified.

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

NEW QUESTION: 135

HBaseTo achieve the scale of bench drills with billions of accidents and millions of columns, which of the following options is required?

A. Redis

B. HDFS

C. ElasticSearch

D. kafka

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

NEW QUESTION: 136

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

- A. Name
- B. Connection
- C. Priority
- D. Type

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

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpsPASS.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 137

In Huawei Fusioninsight HD, which of the following components are Flink strongly dependent on? (Multiple choice)

- A. Kafka
- B. Zookeeper
- C. YARN
- D. HDFS

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

NEW QUESTION: 138

Which of the following options is a key feature of the HDFS architecture?

- A. Metadata Persistence Mechanism
- B. robust mechanism
- C. multiple access mechanism
- D. HA high reliability

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 139

In FusionInsight HD, which of the following components do Flink mainly interact with? (Multiple choice)

- A. zookeeper

- B. Kafka
- C. HDFS
- D. Yarn

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

NEW QUESTION: 140

Kafka is a distributed message publishing and subscribing system. It only forwards messages and does not save messages.

- A. True
- B. False

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

NEW QUESTION: 141

As shown in the figure, which of the following options are wide dependencies

- A. group by A->B
- B. union D&E->F
- C. map:C->D
- D. join B&F->G

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 142

In the FusionInsight Hadoop cluster, the cluster scale has 70 nodes. If the recommended deployment plan is adopted, what partitions might exist on the management node? (Multiple choice)

- A. /srv/BigData
- B. /srv/BigData/zookeeper
- C. /srv/BigData/dbdata-om
- D. /srv/BigData/hadoop/data5
- E. /srv/BigData/jumalnode

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 143

In the MRS service, the unavailability of the Zookeeper service will lead to the unavailability of the kafka service.

- A. False
- B. True

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

NEW QUESTION: 144

Layer 2 networking in the Fusioninsight cluster networking design means that the Layer 2 switching cluster nodes in the cluster are in a subnet. It is suitable for cluster scenarios with less than 200 nodes.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 145

Regarding the order of the YARN scheduler allocating resources, which one of the following descriptions is correct?

- A. Any machine -> same rack -> local resource
- B. Same rack -> any machine -> local resource
- C. Any machine -> local resource -> same rack
- D. Local resources -> same rack -> any machine

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

NEW QUESTION: 146

Which of the following options are suitable for MapReduce scenarios?

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

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

NEW QUESTION: 147

In Hadoop, if `yarn.scheduler.capacity.root.Queue A`.

`minimum-user-limitpercent` is set to 50, which of the following statements is wrong?

- A. Each user in QueueA can only get up to 50% of the resources.
- B. QueueA must ensure that each user gets at least 50% of the resources.
- C. A user submits a task and can use 100% of QueueA's resources.
- D. If there are already 2 user tasks running in QueueA, then the task submitted by the third user needs to wait for the release of resources.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 148

In the FusionInsight cluster, which of the following components does Spark mainly interact with?

- A. YARN
- B. Hive
- C. HDFS
- D. Zookeeper

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 149

When the loader in Fusioninsight HD imports files from the SFTP server, which of the following file types does not require encoding conversion and data conversion and is the fastest?

- A. binary_file
- B. graph_file
- C. sequence_file
- D. text_file

Answer: A (LEAVE A REPLY)

NEW QUESTION: 150

What physical resources does Redis consume mainly?

- A. Disk IO
- B. Memory
- C. Bandwidth
- D. Hard Disk Space

Answer: B (LEAVE A REPLY)

NEW QUESTION: 151

Hive is a data warehouse software based on Hadoop, which can query and manage PB-level distributed data. Which of the following descriptions about Hive features is incorrect? ()

- A. Only supports MapReduce computing engine
- B. Easy to use and easy to program
- C. Flexible and convenient ETL (extract/transform/load)
- D. Direct access to HDFS files and HBase

Answer: A (LEAVE A REPLY)

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpsPass.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 152

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

- A. You can specify a time period for log collection, for example, only collect logs from 2016-1-1 to 2016-1-10.
- B. The node IP can be specified for log collection, for example, only the logs of a certain IP can be downloaded.
- C. You can specify a specific user for log collection, for example, only download logs generated by UserA.
- D. You can specify an instance for log collection, for example, specify to collect metastore logs.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 153

Which of the following components is included in the platform architecture of Huawei's big data solution? (Multiple choice)

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

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

NEW QUESTION: 154

Regarding DataSet, which of the following statements is incorrect?

- A. DataSet needs to be deserialized to perform sort, filter, shuffle and other operations
- B. DataSet is a strongly typed collection composed of objects of a specific domain
- C. DataSet can perform most operations without deserialization
- D. DataSet is highly similar to RDD, and its performance is better than RDD

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

NEW QUESTION: 155

According to the way of data acquisition in Flue, Source can be divided into driven Sourcet and()Source.

- A. polling

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 156

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

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

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

NEW QUESTION: 157

Huawei has extensive experience in big data solutions in finance, public security, railways, government affairs and other industries.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 158

What processes are included in the HBase service of Fusioninsight HD? (Multiple choice)

- A. Data Node
- B. HRegionServer
- C. Slave
- D. HMaster

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

NEW QUESTION: 159

In many small file scenarios, Spark will start a lot of tasks. When there is a Shuffle operation in the SQL logic, it will greatly increase the number of hash buckets, which will seriously affect performance. In Fusioninsight, for small file scenarios, the () operator is usually used to merge the partitions generated by the small files in the Table to reduce the number of partitions, to avoid generating too many hash buckets during shuffle and improve performance?

- A. group by
- B. join
- C. connect
- D. coalosce

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

NEW QUESTION: 160

Which of the following OS Which version is recommended for building a Fusioninsight V1R2C60 cluster?

- A. CentOS6.6
- B. SUSE11 SP1/SP2/SP3 for AMD64&Inter64
- C. RedHat-6.7-x86_64F Ubuntu6.3
- D. redhat-6.4-x86_64
- E. RedHat-6.5-x86_64

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

NEW QUESTION: 161

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

- A. JOIN
- B. SELECT

C. UNION ALL

D. GROUP BY

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 162

Which one of the following descriptions of the characteristics of big data and traditional databases are incorrect?

A. Big data is the data processing of "fish in the pond" with a clear goal; the data processing of traditional database is to judge whether other kinds of "fish" exist or not through some "fish"

B. In big data, there is no unified data tool, that is, "No size fits all"; in traditional databases, in a specific business scenario, one tool can often solve a problem, that is, "One size fits all"

C. There are many types of data processed by big data, including structured, unstructured, and semi-structured data; the traditional database has a relatively single type of data, often structured data predominates.

D. The data scale of big data is very large, generally TB and PB are used as the data processing unit; the data size in traditional databases is generally small, and MB is often used as the data processing unit

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

NEW QUESTION: 163

Which of the following functions can Spark provide? (Multiple choice)

A. Stream processing function

B. Distributed File System

C. Distributed memory computing engine

D. Unified scheduling of cluster resources

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

NEW QUESTION: 164

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

A. Resource recovery request queue priority

B. The queue with the lowest resource utilization has priority

C. The smallest queue level priority

D. Maximum queue level priority

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

NEW QUESTION: 165

In Hive, which of the following description of the partition is wrong?

A. Each partition is a directory

B. There can be more partitions or buckets under the partition

C. The number of partitions is fixed

D. The data table can be partitioned according to the value of a field

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

NEW QUESTION: 166

Planning FusionInsight HD clusters, if customers have no performance requirements for functional testing, under the case of cost savings, they can deploy management nodes, control nodes, and data nodes in one. How many nodes are required at least?

- A. 2
- B. 5
- C. 4
- D. 3

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

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpspass.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 167

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

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

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

NEW QUESTION: 168

In the Kafka cluster, the role of the Kafka server is

- A. Consumer
- B. Producer
- C. Brocker
- D. Zookeeper

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

NEW QUESTION: 169

Which component controls the primary and secondary arbitration of NameNode in HDFS

- A. ZooKeeper Failover Controller

- B. HDFS Client
- C. ResourceManager
- D. NodeManager

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

NEW QUESTION: 170

carriage return as its row separator. In Huawei Fusioninsight Hiver, you can specify the input and output format processing classes of table data.

- A. False
- B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 171

The Streaming of Fusioninsight HD is developed based on the open source Apache Storm, which is a distributed offline computing framework.

- A. True
- B. False

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

NEW QUESTION: 172

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

- A. True
- B. False

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

NEW QUESTION: 173

Which of the following descriptions of Flink key features are wrong?

- A. The Flink streaming engine can provide functions that support stream processing and batch processing applications at the same time
- B. Compared with Flink, SparkStreaming has lower latency
- C. checkpoint realizes Flink's fault tolerance
- D. Compared with Streaming in FusionInght HD, Flink has higher throughput

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

NEW QUESTION: 174

A telecommunications company plan open exhibition Bigdata business, target business have customer grouping, calendar Historical bill analysis, real-time call charge analysis and other services. Which of the following options is the most appropriate in terms of functionality and cost to meet business needs?

- A. Deploy MapReduce with Stare

- B. Deploy Storm
- C. Deploy Map Rabace
- D. Deploy Soark separately

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 175

inWebHCatin the architecture. Users can use the secureHTTPSprotocolexecuteDo which of the following actions?

- A. runMap Reduce tasks
- B. runHiveHQL tasks
- C. implementHive DDLoperate
- D. all of the abovecorrect

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 176

When installing the Streaming component of FusionInsight HD, how many nodes are required to install by the Nimbus role?

- A. 4
- B. 2
- C. 1
- D. 3

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

NEW QUESTION: 177

In the task scheduling of YARN. Once the ApplicationMasterE applies for the resource, it communicates with the corresponding ResourceManagerj and asks it to start the task

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 178

The following aboutWhich Kafka feature description is incorrect??

- A. Support multi-level cascade
- B. Support looking for horizontal expansion
- C. High phagocytosis rate
- D. Support message partitioning

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

NEW QUESTION: 179

In Zookeeper's service model, the Leader node exists in the active-standby mode. All other nodes belong to the Follower node.

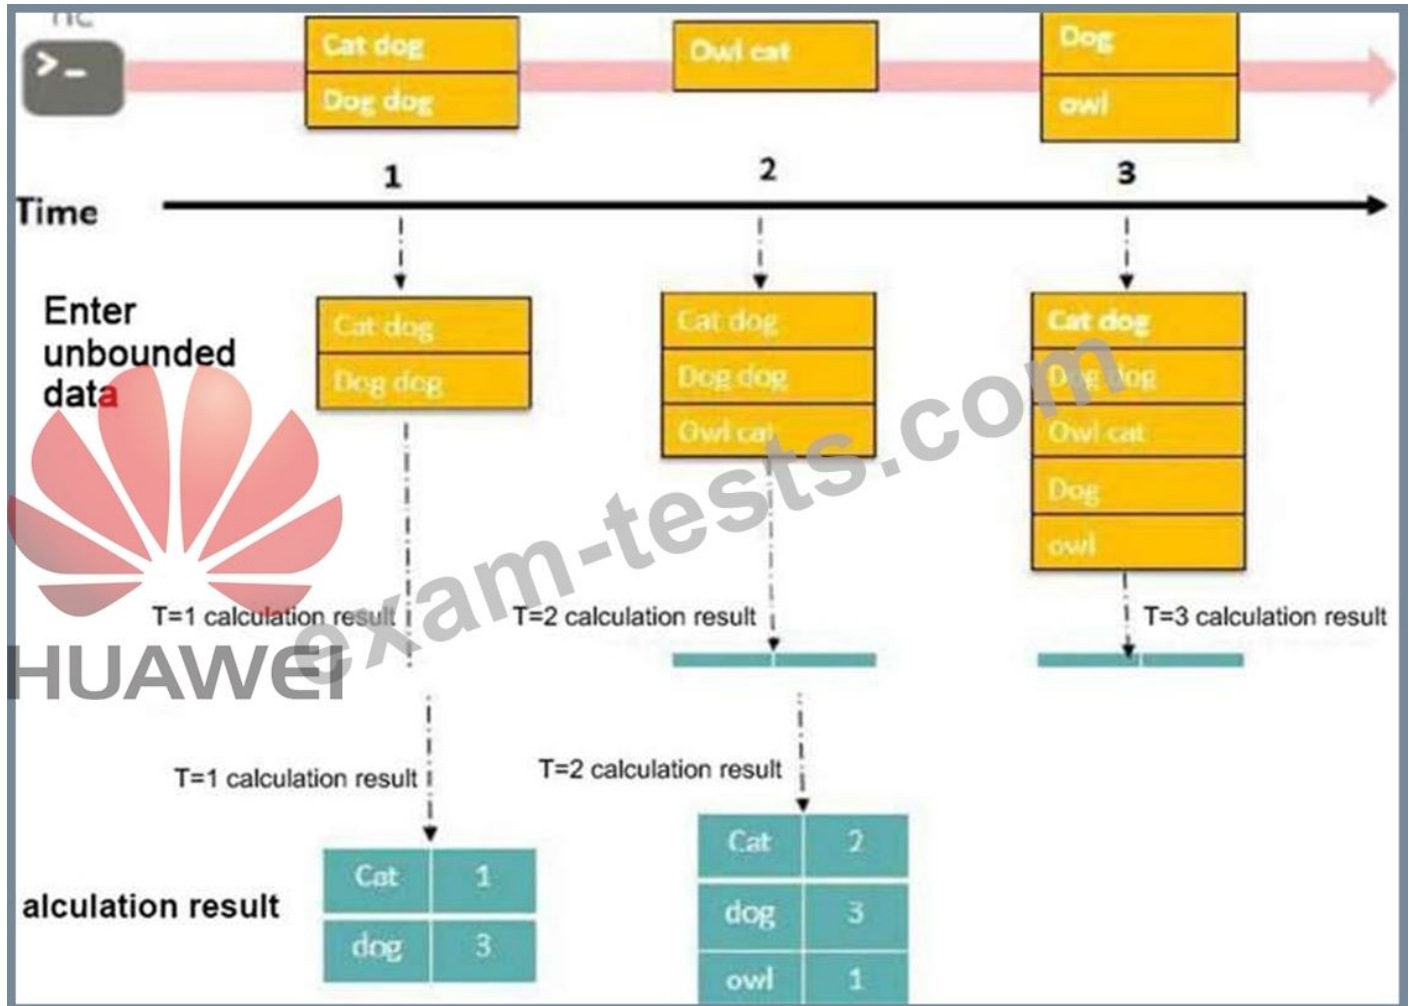
A. False

B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 180

Through observation, it can be concluded that the final T3 calculation result is?



A. Cat 2, dog 4, owl 2

B. Dog 1, owl 1

C. Cat 1, cat 1, dog 2, dog 2, owl 2

D. Cat 2, dog 3, owl 1

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

NEW QUESTION: 181

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

A. Consumers track records by (offset, partition, topic)

B. offset is a String type string

C. Uniquely tag a message

D. The position of each message in the file is called offset

Answer: ([SHOW ANSWER](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:
https://www.braindumpsPASS.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html
(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 182

Which of the descriptions of the Loader job in FusionInsightHD is correct?

- A. After the Loader submits a job to YARN for execution, other jobs cannot be submitted until the job is executed.
- B. After the Loader submits the job to YARN for execution, if the Loader phase is abnormal at this time, the job fails to execute.
- C. The Loader submits the job to the small YARN for execution. If a Mapper task fails to execute, it can automatically retry
- D. After the Loader job fails to execute, garbage data will be generated, which needs to be cleared manually by the user

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

NEW QUESTION: 183

Which of the following statements about worker, Executor and task are correct? (Multiple choice)

- A. Each worker can run multiple Executors
- B. Each worker can only run Executor for one topology
- C. Each Executor can run tasks of different components (spout or bolt)
- D. Each Executor can run a task

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

NEW QUESTION: 184

Starting from version 2.7.3 of HDFS, what is the default Block Size?

- A. 128MB
- B. 64MB
- C. 32MB
- D. 16MB

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

NEW QUESTION: 185

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

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

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

NEW QUESTION: 186

inDGGplatformshelfunder constructioncarryfor? ?Enterprise grouparity ofAccording to management. Data asset management is visualized and supported? ?,? ?Wait. through data? ?, the realization material? ?view visible. Provide data intelligent search and operation?The number of modules to be controlled is ().

- A. Data Development
- B. Data Integration
- C. Data Asset Management
- D. Specification design

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

NEW QUESTION: 187

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

- A. 1
- B. 4
- C. 3
- D. 2

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

NEW QUESTION: 188

Neither Spark nor Hadoopi are suitable for iterative computing scenarios.

- A. False
- B. True

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

NEW QUESTION: 189

Which of the following description about FusionInsight multi-tenant management is correct?
(Multiple choice)

- A. Visualized multi-level tenant management, matching the corporate organizational structure, simplifying system resource allocation and management
- B. Isolation of tenant resources based on the container mechanism, escorting tenant SLA
- C. One-stop management of tenant resources, computing resources (CPU/memory/IO), storage resources (HDFS), service resources (HBase...)

D. A multi-level tenant model that matches the organizational structure of the enterprise, different departments correspond to different tenants, and tenants are dynamically added or deleted on demand

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

NEW QUESTION: 190

In Hive, which of the following statements about partitions is incorrect?

- A. The number of partitions is fixed
- B. The data table can be partitioned by the value of a field
- C. There can be further partitions or buckets under the partition
- D. Each partition is a journal

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

NEW QUESTION: 191

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. This can greatly reduce HBase's compression and split frequency and improve performance.

- A. True
- B. False

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

NEW QUESTION: 192

About RDDs. Which of the following statements is false?

- A. RDD is a read-only, partitionable distributed dataset
- B. RDD is Spark's abstraction of underlying data
- C. RDDs have a lineage mechanism (Lineed)
- D. RDDs are stored on disk by default

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 193

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 necessary to store some mixed scenes of large files (above 10MB)

- A. True
- B. False

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

NEW QUESTION: 194

In the FusionInsight HD product, a typical Kafka cluster contains several Producers, several Borkers, several Consumers and a Zookeeper cluster

- A. False

B. True

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

NEW QUESTION: 195

How is the HBaseM master elected?

- A. HMaster is a dual-master mode and does not need to be adjudicated
- B. Adjudicated by RegionServer
- C. Adjudication via Zookeeper
- D. Randomly selected

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

NEW QUESTION: 196

Fusioninsight HD system audit log is unavailable to record which of the following operations?

- A. Manually clear warnings
- B. Check history monitoring
- C. Start and stop service instance
- D. Delete service instance

Answer: ([SHOW ANSWER](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpspass.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 197

What are the success stories of Huawei Fusioninsight HD industry? (Multiple choice)

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

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

NEW QUESTION: 198

Regarding the basic operation of Hive table creation, which description is correct?

- A. Once the table is built, the table name cannot be modified
- B. Once the table is built, no new columns can be added

- C. Once the table is created, you can no longer modify the column name
- D. The external keyword needs to be formulated when creating an external table

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

NEW QUESTION: 199

Which of the following statements about the read/write process of the leader node of Zookeeper after receiving the data change request is correct?

- A. Write to disk first, then write to memory
- B. Simultaneous writes to disk and memory
- C. write to memory only
- D. Write to memory first, then write to disk

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

NEW QUESTION: 200

In the Zookeeper service model, the Leader node exists in the active/standby mode, and other nodes belong to the Follower node.

- A. False
- B. True

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

NEW QUESTION: 201

The following offon HDFSarity ofholding a spoontransformedProcess description,pleaseArrange from top to bottom in the correct order.

- A. put frimage?combine into new? ?document.
- B. Saccn? ?node periodically fromMa? ?download?and?files
- C. in HDFS Sectionformat onceRear, ?will generate fs?me and?twodelivery
- D. to avoid?Increasing, ?will be merged periodically?become new?.
- E. to replace?file, and will?coincident name?make up new?and and.

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

NEW QUESTION: 202

Which of the following descriptions of HBase FileStream (HFS) are correct? (Multiple choice)

- A. HFS is an independent module of HBase
- B. Provide file storage, read, and delete functions for upper-level applications
- C. HFS encapsulates the interface between HBase and HDFS
- D. Applied in the upper application of Fusioninsight HD

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 203

Which of the following does not belong to the characteristics of the MapReduce component in Hadoop?

- A. Good scalability
- B. High fault tolerance
- C. Real-time calculation
- D. Easy to program

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

NEW QUESTION: 204

In the FusionInsight HD cluster, there are currently 300 nodes in the cluster. If the recommended deployment scheme is adopted, which partitions will definitely not exist on the control node?

(Multiple choice)

- A. /srv/BigData
- B. /srv/Bigdata/hadoop/data5
- C. /srv/Bigdata/umatnode
- D. /srv/BigData/dbdata_om
- E. /

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

NEW QUESTION: 205

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

- A. False
- B. True

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

NEW QUESTION: 206

The correct steps for data transformation in Fusioninsight HD are

- A. Extract, convert, output
- B. load, transform, extract
- C. load, convert, export
- D. Input settings, conversions, outputs

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

NEW QUESTION: 207

Which of the following are used to configure the number of memory and CPU of the NodeManager? (Multiple choice)

- A. YARN.modemanager.resource.memory-mb
- B. YARN.scheduler.capacity.root.QueueA.maximum-capacity
- C. YARN.nodemanager.vmem-pmom-ratio
- D. YARN.nodemanager.resource.cpu-vcore

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

NEW QUESTION: 208

Which service process manages the HBase Region in the Hadoop platform?

- A. DataNode
- B. HMaster
- C. RegionServer
- D. Zookeeper

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 209

Which of the following RESTful request methods belong to Elasticsearch? (Multiple choice)

- A. Get
- B. Post
- C. Delete
- D. Update

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

NEW QUESTION: 210

Which of the following are the main functions of FusionInsight Manager? (Multiple choice)

- A. System Management
- B. Security Management
- C. Data integration
- D. Service governance

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 211

In the FusionInsight cluster, which of the following components does Spark mainly interact with? (Multiple choice)

- A. HDFS
- B. Zookeeper
- C. Hive
- D. YARN

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

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

NEW QUESTION: 212

As shown in the figure, the following description of the message read by the Kafka message consumer Consumer1 is wrong?

- A. Consumers read messages sequentially from right to left.
- B. The messages generated by the producer are placed at the end of the topic one by one.
- C. Consumer uses offset to record the position of the read
- D. The blue in the picture is a topic of Kafka, which can be understood as a queue, and each grid represents a message.

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

NEW QUESTION: 213

Which of the following targets can Fusioninsight HD Loader export HDFS data to? (Multiple choice)

- A. SFTP server
- B. FTP server
- C. Oracle database
- D. DB2 database

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

NEW QUESTION: 214

When deploying Fusioninsight HD, how many FlumeServer nodes in the same cluster are recommended to deploy at least?

- A. 2
- B. 3
- C. 4
- D. 1

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

NEW QUESTION: 215

In the Fusioninsight HD system, which of the following methods cannot view the results of Loader job execution?

- A. View through NodeManager
- B. View through Manager's alarm
- C. View the YARN Task Management
- D. View through the Loader UI interface

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

NEW QUESTION: 216

The ResourceManager adopts a high-availability solution. When the Active ResourceManager finds a fault, it can only start the standby ResourceManager through the built-in ZooKeeper and switch its state to active.

A. False

B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 217

Redis uses the `exists` command to determine whether a key exists.

A. True

B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 218

Which is not the channel type of Flume?

A. Memory Channel

B. File Channel

C. HDFS Channel

D. JDBC Channel

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

NEW QUESTION: 219

In the FusionInsight HD system, if dirty data is generated while the Loader job is running, the status of the Loader job execution result must be failed.

A. True

B. False

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

NEW QUESTION: 220

What is the default resource scheduler for queues in YARN?

Answer:

capacityscheduler

NEW QUESTION: 221

Which of the following descriptions of Kafka partition offset is incorrect?

A. Uniquely mark a message

B. Consumers track records through (offset, partition, topic)

C. The position of each message in the file is called offset (offset)

D. offset is a String type string

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

NEW QUESTION: 222

As an authentication server center, Kerberos1 can provide unified authentication services to all services in the cluster and secondary development applications of customers.

- A. True
- B. False

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

NEW QUESTION: 223

The index data of Elasticsearch can only be stored in the HDFS system.

- A. True
- B. False

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

fill in the blank

NEW QUESTION: 224

In the Mapreduce process, by default, a shard is a block and a mapTask.

- A. False
- B. True

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

NEW QUESTION: 225

In Streaming, one accurate message reliability level is achieved through the ACK mechanism.

- A. True
- B. False

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

NEW QUESTION: 226

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

- A. True
- B. False

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

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

NEW QUESTION: 227

When Fusioninsight Manager is connected to an external management platform, which of the following interfaces does it support? (Multiple choice)

- A. FTP
- B. VPN
- C. SNMP
- D. Syslog

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

NEW QUESTION: 228

ApplicationMasters apply for and receive resources from ResourceManager1 through the RPC protocol in a polling manner.

- A. True
- B. False

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

NEW QUESTION: 229

A user information table meg_table is saved in the HBase of Fusioninsight HD, Rowkey is the user id, and one of the columns is the user nickname. Now write three KeyValues in this column in order: 001:Li, 001:Mary, 001:Lily, May I ask scan 'meg_table', (VERSION=>2) which data will be returned?

- A. 001: Mary, 001: Lily
- B. 001: Lily
- C. 001: Li
- D. 001: Li, 001: Mary, 001: Lily

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

NEW QUESTION: 230

Which of the following are not part of the core elements of KrbServer?

- A. Kerberos Client
- B. KDC(Key Distribution Center)
- C. Kerberos KDC Client
- D. Kerberos KDC Server

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

NEW QUESTION: 231

Which of the following HDFS commands can be used to check the integrity of data blocks?

- A. HDFS fsck-delete

- B. HDES balancer -threshold 1
- C. HDFS fsck
- D. HDFS dfsadmin -report

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

NEW QUESTION: 232

Which of the following options are included in the key features of ZooKeeper?

- A. Reliablesex
- B. wait for irrelevance
- C. atomicity
- D. eventual consistency

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

NEW QUESTION: 233

Which of the following operations cannot be recorded in the Fusioninsight HD system audit log?

Answer:

Queryhistorymonitoring

NEW QUESTION: 234

Which of the following processing methods is unreasonable when the window has been closed and the calculation result has been produced when the delayed event occurs?

- A. Reactivate closed windows and recalculate to correct results
- B. Treat delayed events as error messages and discard
- C. Collect delayed events for additional processing
- D. Delayed events - must not ignore calculations

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

NEW QUESTION: 235

The order in which the YARN scheduler allocates resources, which one of the following descriptions is correct?

- A. Any machine -> same rack -> local resources
- B. Local resources -> same rack->Any machine
- C. Same rack->any machine->local resources
- D. Any machine -> local resources -> same rack

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 236

Which of the following statements about the LdapServer organizational model is correct?

- A. DAPServer directory information is organized and stored based on a tree structure.
- B. The root of the LDAPServer directory tree generally defines the domain name DC (Domain Component).

- C. Concrete objects cannot be defined under an organizational unit.
- D. Each node in the LDAPServer directory tree is called an entry and has its own unique distinguished name DN (Distinguished Name).

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

NEW QUESTION: 237

Which of the following functions can be achieved by Loader's Shell script? (Multiple choice)

- A. Data source query
- B. Start the job
- C. View job status
- D. Stop the job

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

NEW QUESTION: 238

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. True
- B. False

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

NEW QUESTION: 239

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

- A. True
- B. False

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

NEW QUESTION: 240

Which of the following options is correct about the meaning of Zookeeper reliability?

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

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

NEW QUESTION: 241

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

- A. Expansion of nodes and disk services will not be interrupted, and the thesaurus will be updated without interruption
- B. TB-level data retrieval is returned in seconds
- C. Does not support the fusion search of text such as sound
- D. Compatible with Elasticsearch APIs

Answer: ([SHOW ANSWER](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpsPASS.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 242

In order to improve the fault tolerance of kafka, kafka supports partition replication strategy. What is wrong with the following description of Leader partition and Follower partition?

- A. If the leader fails, then another follower will take over (becoming the new leader)
- B. Kafka needs to select a leader for partition replication, and the leader is responsible for the read and write operations of the partition, and other replica nodes are only responsible for data synchronization
- C. It is impossible for each node of a Kafka cluster to be the leader and follower of each other
- D. Since the leader Server carries all the request pressure, from the overall consideration of the cluster, Kafka will distribute the leader on each instance in a balanced manner to ensure stable overall performance

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

NEW QUESTION: 243

In the Output stage, Structured Streaming can define different data writing methods, including which of the following methods?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 244

In the HBase component, data read and write services need to be connected to the Master for execution.

- A. False
- B. True

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

NEW QUESTION: 245

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

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

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

NEW QUESTION: 246

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

- A. Traditional databases do not require high data processing time, while big data requires real-time data processing
- B. Mass data operation and maintenance needs to ensure data stability and support high concurrency while increasing server burden.
- C. Traditional databases do not consider the diversity of data, and are especially compatible with structured data, unstructured data and semi-structured data.
- D. Traditional databases are not suitable for processing PB of data.

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

NEW QUESTION: 247

The label-based scheduling of YARN is to label which of the following options?

- A. Applaster
- B. ResourceNanager
- C. Container
- D. NodeManager

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

NEW QUESTION: 248

SELECT aa.salary**B**. address FROM employee aa **JoIN** SELECT adress FROM employee info where provine='zhejiang') What types of operations does **bb ON**aa.nanme=**bB**. name contain?

- A. subquery
- B. JOIN Cha Xun
- C. Import Data
- D. create table

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

NEW QUESTION: 249

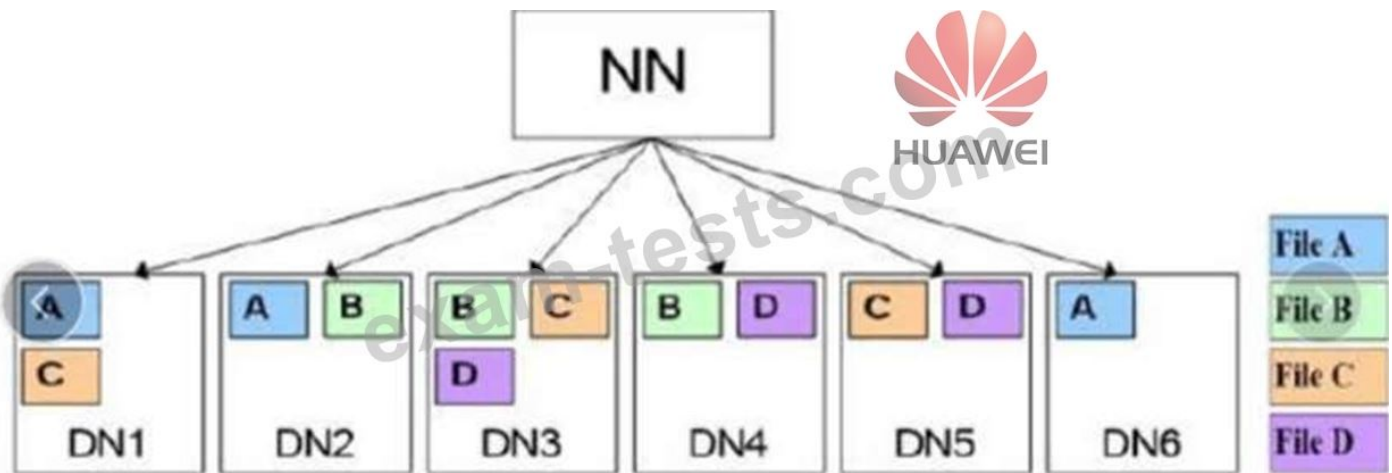
The checkpoint mechanism in Flink continuously draws snapshots of streaming applications, and the state snapshots of streaming applications can only be saved in the HDFS file system.

- A. False
- B. True

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

NEW QUESTION: 250

The following figure shows the storage locations of files A, B, C, and D. Among them, A and B are related, and their storage locations conform to the colocation strategy of Colocation.

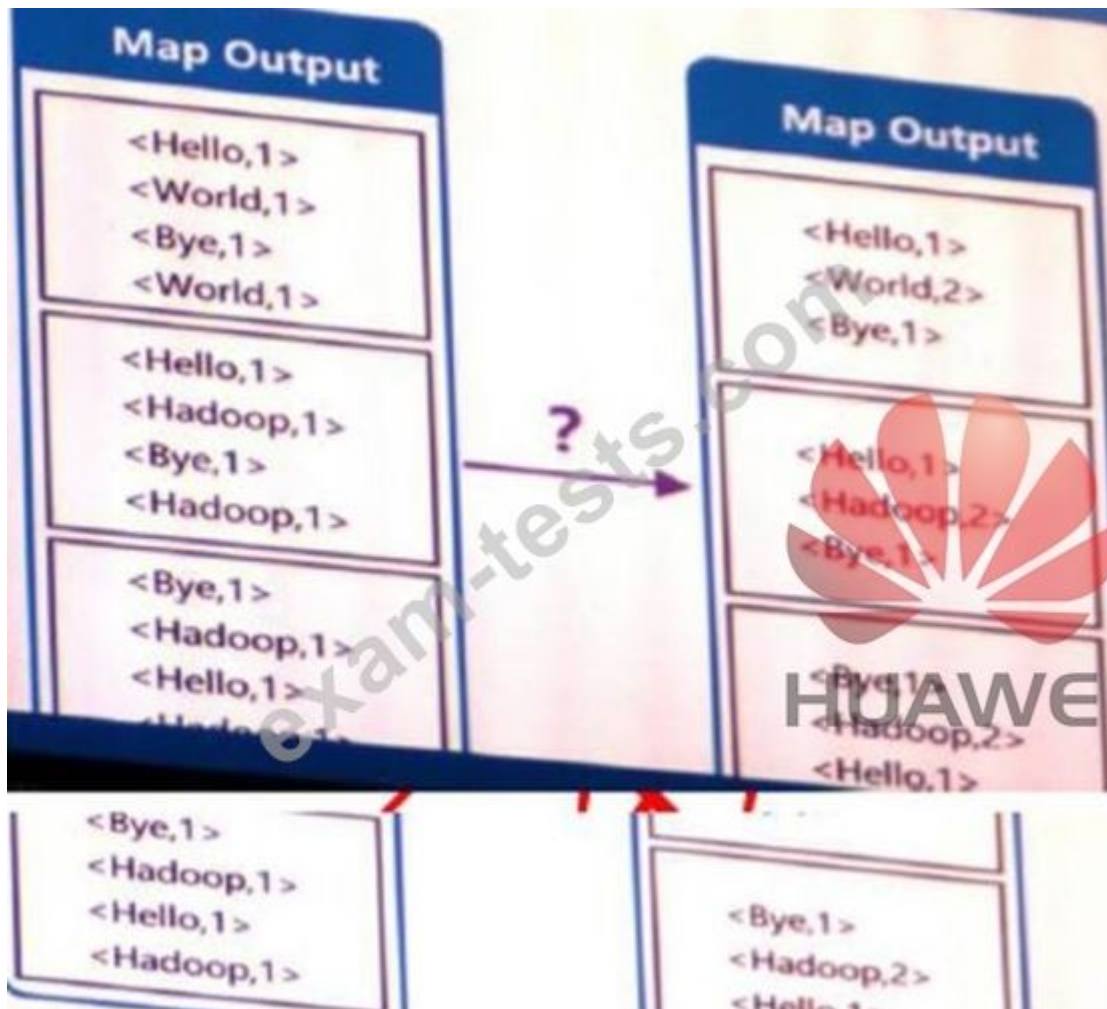


- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 251

Which process does the Map stage describe?



- A. Combine
- B. Spill/Merge
- C. Partition
- D. Sort

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 252

In Fusioninsight HD, which of the following is not a flow control feature of Hive

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

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

NEW QUESTION: 253

Which of the following scenarios are not what Flink components good at?

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

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

NEW QUESTION: 254

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

- A. True
- B. False

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

NEW QUESTION: 255

LdapServerff? ?(group) is to unify usersgroup managementreason. likefruituserAdd to? ? middle.? ?middle? ?record.

- A. FALSE
- B. TRUE

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

NEW QUESTION: 256

When installing the Streaming component of Fusioninsight HD, how many nodes does the Nimbus role require to installed?

- A. 3
- B. 4
- C. 1
- D. 2

Answer: ([SHOW ANSWER](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpsPass.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 257

A user needs to build a FusionInsightHD cluster with 350 nodes. Which planning solution is the best?

- A. The management node and data node are deployed together, the control node is independently deployed, and the two-layer networking

B. Management node, control node, and data node are deployed in one, layer-2 networking

C. Management nodes, control nodes, and data nodes are independently deployed, with three-layer networking

D. Integrated deployment of management nodes and control nodes, independent deployment of data nodes, and Layer 2 networking

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

NEW QUESTION: 258

Which of the following options belong to Hive's data storage model?

A. all of the above

B. database

C. surface

D. partition

E. bucket

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

NEW QUESTION: 259

Kafka, as a distributed messaging system, supports online and offline message processing, and provides javaAPI for the docking of other components. In the Fusioninsight solution, Kafka belongs to the FusionInsight HD module.

A. False

B. True

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

NEW QUESTION: 260

When planning and deploying a FusionInsight cluster, it is recommended that the management node be best deployed(), the control node needs to be deployed at least()Piece,Data nodes need to be deployed at least()Piece.

A. 1,2, 2

B. 1, 3, 2

C. 2, 3, 1

D. 2, 3, 3

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 261

In the Fusioninsight product, which statement is correct about the Kafka component?

A. After Kafka is installed, the sensitive data storage directory cannot be configured.

B. When the Producer of Kafkal sends a message, it can specify which Consumer consumes the message

C. When creating a topic, the number of replicas must not be greater than the number of currently surviving Broker instances, otherwise the topic creation will fail

D. Kafka will store metadata information in Zookeeper for

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

NEW QUESTION: 262

File is a kind of distributed, high reliable and high available of cloud service. used to be effective compare?, polymerization? Move a lot of logs? ?.

A. FALSE

B. TRUE

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

NEW QUESTION: 263

Which of the following options is correct for the description of Compaction?

A. Minor trigger frequency is higher than Major

B. There are two types of Minor and Major

C. Compaction reduces the number of HFile files by merging files

D. Compaction reduces the number of HFile files by eliminating expired data files

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 264

RDD has Transformation and Action operators. Which of the following are Action operators?

A. saveAsTextFile

B. filter

C. map

D. reduceByKey

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

NEW QUESTION: 265

The driving source is that flume periodically actively acquires data.

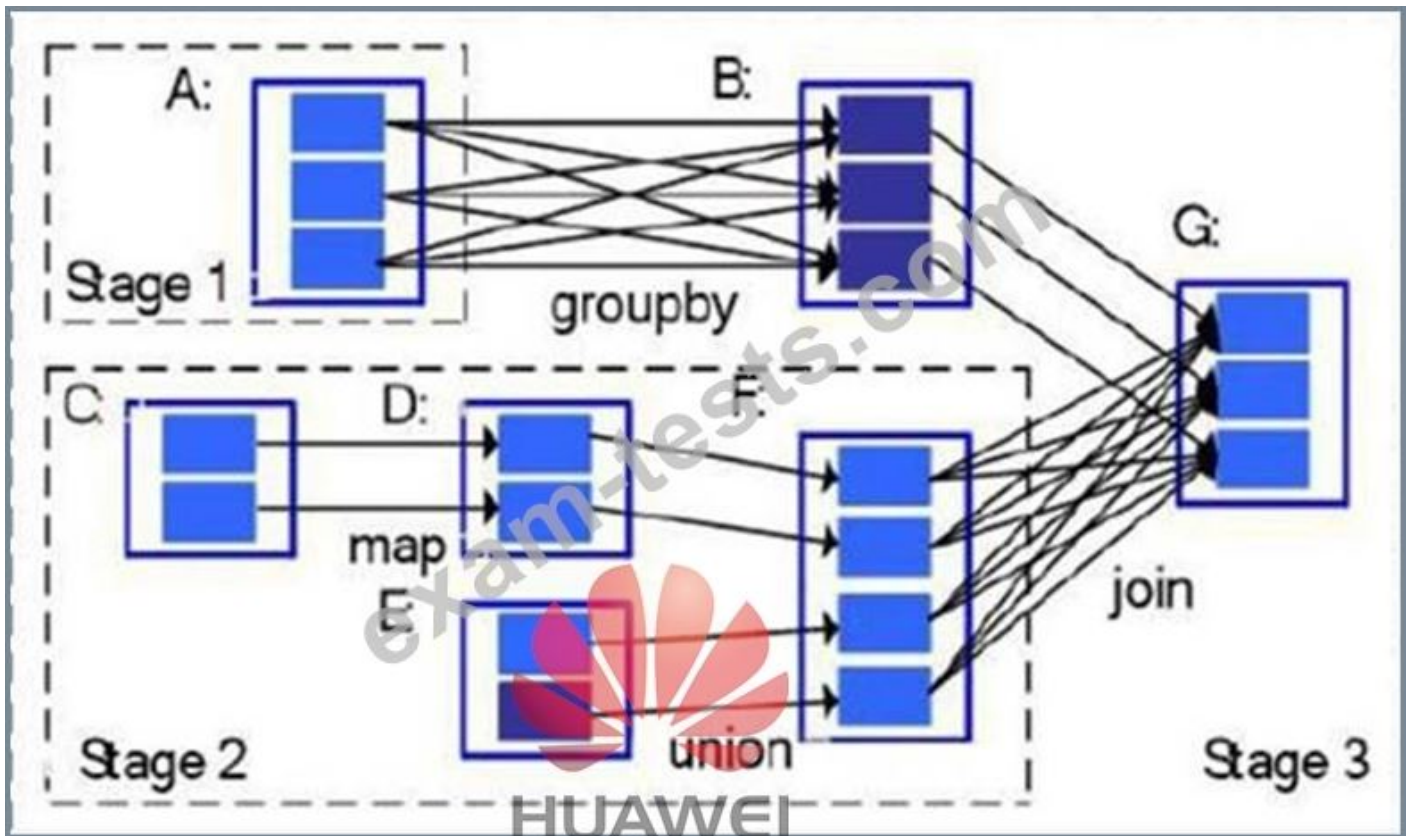
A. True

B. False

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

NEW QUESTION: 266

Which of the following options belong to broadly dependent? (Multiple choice)



A. map: C->D

B. union: D&E->F

C. group by: A->B

D. join: B&F->G

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

NEW QUESTION: 267

Spark, like Hadoop, is not suitable for iterative computing.

A. True

B. False

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

NEW QUESTION: 268

There are currently 3 racks, and there is a file that needs to be stored in 3 copies, of which copy 1 and copy 2 are stored in the same rack as the client but on different servers. According to HDFS's copy placement strategy, it must be stored in other racks.

A. False

B. True

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

NEW QUESTION: 269

Regarding RDD, which of the following statements is wrong?

A. RDD is stored on disk by default

- B. RDD is an abstraction of Spark's basic data
- C. RDD is a read-only, partitionable distributed data set
- D. RDD has a lineage mechanism

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

NEW QUESTION: 270

Which of the descriptions of the Loader job in FusionInsightHD is correct?

- A. Loader submits a job to Yam for execution. No other jobs can be submitted until the job is executed.
- B. After the Loader submits the job to Yam for execution, if the Loader service is abnormal at this time, the job execution fails.
- C. After the Loader job fails to execute, garbage data will be generated, which needs to be cleared manually by the user
- D. After the Loader submits the job to Yam for execution, if a Mapper task fails to execute, it can automatically retry

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

NEW QUESTION: 271

There is no single point of failure in the Discovery module of elasticSearch.

- A. FALSE
- B. TRUE

Answer: ([SHOW ANSWER](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpspass.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 272

In Spark on YARN mode, nodes without NodeManager deployed cannot start executor to execute tasks.

- A. True
- B. False

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

NEW QUESTION: 273

When deploying Fusioninsight HD, how many FlumeServer nodes are recommended to be deployed in the same cluster?

- A. 3
- B. 4
- C. 2
- D. 1

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

NEW QUESTION: 274

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

- A. 1
- B. 3
- C. 4
- D. 2

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

NEW QUESTION: 275

Among the Kerberos service roles, KerberosServer mainly provides authentication functions, and KerberosAdmin mainly provides user management functions.

- A. True
- B. False

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

NEW QUESTION: 276

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

- A. False
- B. True

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

NEW QUESTION: 277

The following figure shows the label storage strategy of HDFS. Observe the figure below, which data nodes will HBase data be stored on

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

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

NEW QUESTION: 278

User Rights Management Role-Based Access Control (RBAC), which provides visual multi-group unified user rights management in clusters.

A. False

B. True

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

NEW QUESTION: 279

Which of the following scenarios is not applicable to Hive?

A. Non-real-time analysis, such as log analysis, statistical analysis

B. Data summary, such as daily and weekly user clicks, click ranking

C. Real-time online data analysis

D. Data mining, such as user behavior analysis, interest division, regional display.

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

NEW QUESTION: 280

Each stage of a Spark task can be divided into jobs, and the division mark is shuffle.

A. True

B. False

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

NEW QUESTION: 281

Fusioninsight Spark SQL, like the community Spark JDBCServer, only supports single-tenant binding to a YARN resource queue and multi-tenancy, and does not support multi-tenant parallel execution.

A. False

B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 282

ElasticSearch expansion fieldKing Packageincluding whichsome?

A. The index data of a single instance of ElasticSearch is too large

B. The node needs to reinstall the operating system

C. Excessive consumption of physical resources

D. The amount of cluster data is reduced

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

NEW QUESTION: 283

When installing the Streaming component of Fusioninsight HD, the Nimbus role requires several nodes to be installed

A. 2

B. 3

C. 1

D. 4

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

NEW QUESTION: 284

The OpenScanner process in HBase creates two different Scanners to read HFile and MemStore data. The Scanner corresponding to HFile is StoreFileScanner, and the Scanner corresponding to MemStore is MemStoreScanner.

A. False

B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 285

Regarding the Kafka message consumer reading the message, which of the following is wrong?

A. The blue is a topic of Kafka, which can be understood as a queue, and each grid represents a message.

B. The messages generated by the producer are placed at the end of the topic one by one.

C. Consumer uses offset to record the read position

D. The consumer reads the messages sequentially from right to left.

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

NEW QUESTION: 286

Hadoop's NameNode is used to store the metadata of the file system.

A. False

B. True

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

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpspass.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 287

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

A. False

B. True

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

NEW QUESTION: 288

Which of the following descriptions is correct about the deployment of big data components on Kunpeng and X86 servers? (Multiple choice)

- A. Support mixed deployment of Kunpeng server and ordinary X86 server in a single cluster
- B. No shortcomings in performance
- C. Single component (for example: HDFS) supports mixed deployment of Kunpeng server and X86 server
- D. Achieve autonomous controllability of some equipment

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

NEW QUESTION: 289

In the cooperative work of Zookeeper and YARN, when Active ResourceManagerj fails, which directory will Standby ResourceManager obtain Application-related information from?

- A. Metastore
- B. Storage
- C. State store
- D. warehouse

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

NEW QUESTION: 290

In the Fusioninsight product, which of the following descriptions about kafka topic is wrong?

- A. Each topic can only be divided into one partition area
- B. Each partition corresponds to a log file at the storage level, and all message data is recorded in the log file
- C. Every message published to Kafka has a category, which is called topic and can also be understood as a queue for storing messages
- D. The number of Topic partitions can be configured when created

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

NEW QUESTION: 291

There are 230 nodes in the FusionInsight HD cluster of a university. When clustering is performed, which of the following processes should be deployed on the data nodes? (Multiple choice)

- A. SolrServer
- B. NodeManager
- C. NameNode
- D. DBServer
- E. RegionServer

F. DataNode

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 292

All messages in Kafka will be persisted to the hard disk. At the same time, Kafka ensures data reliability by setting Replication to Topic Partitio.

- A. True
- B. False

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

NEW QUESTION: 293

The Fusioninsight HD cluster planned by a bank has 90 nodes. If 3 control nodes are planned, how many data nodes in the cluster is recommended to plan?

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

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

NEW QUESTION: 294

Which way does Elasticsearch index data use?

- A. Keyword Index
- B. Forward Index
- C. Inverted index
- D. Full-text index

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

NEW QUESTION: 295

ElasticSearch is strongly dependent on Zookeeper.

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 296

The size of the memory allocated by YARN to Containerb in the Hadoop system, which can be set by the parameter yarn.app.mapreduceam.resource.mb

- A. False
- B. True

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 297

The memory size allocated by YARN to the 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: 298

Which of the following are the optimization methods of Redis?

- A. Limit Redis memory size
- B. Turn off persistence
- C. Reduced key-values
- D. slowlog configuration

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

NEW QUESTION: 299

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

- A. Proucer connects to any normal Broker instance to generate data
- B. As a message producer, Proucer can write data directly to Zookeeper
- C. Proucer is responsible for producing data, Consumer is responsible for consuming data, and a connection needs to be established between Producer and Consumer.
- D. Proucer can not only send data to Broker, but also send data to Consumer

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

NEW QUESTION: 300

Which statement about the DataNodel of HDFS in Huawei Fusioninsight HD system is correct?

- A. Does not check the validity of the data
- B. Blocks on a DataNode. can be the same
- C. Blocks stored in different DataNodes must be different
- D. Periodically send the block-related information of this node to the NameNode

Answer: (SHOW ANSWER)

NEW QUESTION: 301

Which of the following commands can be used to create node data?

- A. create /node
- B. set/node data
- C. get/node
- D. 1s/node

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

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpspass.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 302

The data node is the working node of HDFS. Which of the following are its functions? (multiple choice)

- A. Data storage and retrieval are performed according to the scheduling of the client or the name node
- B. Responsible for storing and reading data
- C. Periodically send the namenode a list of its own stored blocks.
- D. Records all operations for file creation, deletion, renaming, etC.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 303

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

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

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

NEW QUESTION: 304

What data types of application scenarios is HBase of Hadoop not suitable for?

- A. Mass data application scenarios
- B. High throughput application scenarios
- C. Semi-structured data application scenarios
- D. Large file application scenarios

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 305

Regarding Flume, which of the following statement is wrong? (Multiple choice)

- A. Flume supports multi-cascade and multiplexing
- B. Data transmission between Flume cascade nodes does not support compression
- C. Data transmission between Flume cascade nodes supports encryption

D. There is a need for encryption in processes such as Source to Channel to Sink

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

NEW QUESTION: 306

Which of the following statements about CarbonData in Fusioninsight is correct?

- A. The purpose of using carbon is to provide ultra-fast responses to ad-hoc queries on big data.
- B. cArbon is also a high-performance analytics engine that integrates data sources with spark.
- C. arbon is a new Apache Hadoop native file format that uses advanced columnar storage, indexing, compression, and encoding techniques to improve computational efficiency to help accelerate data queries over petabytes of magnitude, and can be used for faster interactive queries.
- D. cArbon uses a combination of lightweight compression and heavyweight compression to compress data, which can reduce data storage space by 60%-80% and greatly save hardware storage costs.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 307

In Hadoop, if yarn.scheduler.capacity.root.Queue A.

minimum-user-limit-percent is set to 50, which of the following statements is wrong?

- A. A user-submitted task can use 100% of QueueA's resources
- B. Each user in QueueA can only get up to 50% of the resources
- C. QueueA must ensure that each user gets at least 50% of the resources
- D. If there are already 2 user tasks running in QueueA, then the task submitted by the third user needs to wait for the release of resources

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 308

When the scale of the Hadoop cluster is large, the distribution of data will be very critical. Users need to decide whether the cluster is to be expanded and whether the data needs to be balanced according to the data distribution. Which of the following statements about Fusioninsight resource distribution monitoring are correct? (Multiple choice)

- A. Through the DashBoard interface, you can view the distribution of host resources. For example, a list of hosts with a memory occupancy rate of 50-75%, and link redirection is provided
- B. Help users quickly find the node with the highest resource consumption and adopt appropriate measures
- C. Through the resource distribution viewing interface of each service homepage, you can view the distribution of key resources
- D. Can help users quickly gather on the most critical resource consumption

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 309

Which of the following options cannot be achieved through big data technology?

Answer:

businessmodeldiscovery

NEW QUESTION: 310

Which of the following options belong to Fusioninsight data security?

- A. Operating system security hardening
- B. User authority authentication management
- C. Component data encryption
- D. data integrity check

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 311

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

- A. True
- B. False

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 312

Which of the following is not a role or service involved in the HBase read data process?

- A. Zookeeper
- B. HDFS
- C. HMaster
- D. HRegionServer

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 313

What is the data scale unit of traditional data processing?

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

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

NEW QUESTION: 314

Which of the following statements about Fusioninsight HBase visual modeling are correct?

(Multiple choice)

- A. Visual modeling helps DBA perform modeling and design, lowering the threshold for using HBase

- B. Column user table column: each column represents a KeyValue
- C. Realize the division of labor: DBA pays attention to data table modeling, and developers pay attention to user table names and columns used
- D. Qualifier HBase column: each column represents a certain attribute of business data

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 315

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

- A. All the above are correct
- B. Perform Hive DDL operations
- C. Run MapReduce tasks
- D. Run Hive HQL tasks

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 316

Which of the following are the key features of the HDFS architecture? (Multiple choice)

- A. Metadata persistence mechanism
- B. HA high reliability
- C. Multi-mode access mechanism
- D. Robust mechanism

Answer: ([SHOW ANSWER](#))

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumps.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

Valid H13-711_V3.0-ENU Dumps shared by BraindumpsPass.com for Helping Passing H13-711_V3.0-ENU Exam! BraindumpsPass.com now offer the **newest H13-711_V3.0-ENU exam dumps**, the BraindumpsPass.com H13-711_V3.0-ENU exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com H13-711_V3.0-ENU dumps with Test Engine here:

https://www.braindumpspass.com/Huawei/H13-711_V3.0-ENU-practice-exam-dumps.html

(646 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)