

Examcollection

<http://www.ipass4sure.com/examcollection.htm>

CCA-470

Cloudera

*Cloudera Certified Administrator for Apache Hadoop
CDH4 Upgrade (CCA4H)*

The CCA-470 practice exam is written and formatted by Certified Senior IT Professionals working in today's prospering companies and data centers all over the world! The CCA-470 Practice Test covers all the exam topics and objectives and will prepare you for success quickly and efficiently. The CCA-470 exam is very challenging, but with our CCA-470 questions and answers practice exam, you can feel confident in obtaining your success on the CCA-470 exam on your FIRST TRY!

Cloudera CCA-470 Exam Features

- Detailed questions and answers for CCA-470 exam
- Try a demo before buying any Cloudera exam
- CCA-470 questions and answers, updated regularly
- Verified CCA-470 answers by Experts and bear almost 100% accuracy
- CCA-470 tested and verified before publishing
- CCA-470 examcollection vce questions with exhibits
- CCA-470 same questions as real exam with multiple choice options

Acquiring Cloudera certifications are becoming a huge task in the field of I.T. More over these exams like CCA-470 exam are now continuously updating and accepting this challenge is itself a task. This CCA-470 test is an important part of Cloudera certifications. We have the resources to prepare you for this. The CCA-470 exam is essential and core part of Cloudera certifications and once you clear the exam you will be able to solve the real life problems yourself. Want to take advantage of the Real CCA-470 Test and save time and money while developing your skills to pass your Cloudera CCA-470 Exam? Let us help you climb that ladder of success and pass your CCA-470 now!

DEMO EXAM

For Full Version visit

<http://www.ipass4sure.com/allexams.asp>

QUESTION: 1

Which two updates occur when a client application opens a stream to begin a file write on a cluster running MapReduce v1 (MRv1)?

- A. Once the write stream closes on the DataNode, the DataNode immediately initiates a block report to the NameNode.
- B. The change is written to the NameNode disk.
- C. The metadata in the RAM on the NameNode is flushed to disk.
- D. The metadata in RAM on the NameNode is flushed disk.
- E. The metadata in RAM on the NameNode is updated.
- F. The change is written to the edits file.

Answer: D, F

Explanation:

Note: Namenode stores modifications to the filesystem as a log appended to a native filesystem file (edits). When a Namenode starts up, it reads HDFS state from an image file (fsimage) and then applies edits from edits log file. It then writes new HDFS state to (fsimage) and starts normal operation with an empty edits file. Since namenode merges fsimage and edits files only during start up, edits file could get very large over time on a large cluster. Another side effect of larger edits file is that next restart of Namenode takes longer. The secondary namenode merges fsimage and edits log periodically and keeps edits log size within a limit. It is usually run on a different machine than the primary Namenode since its memory requirements are on the same order as the primary namenode. The secondary namenode is started by bin/start-dfs.sh on the nodes specified in conf/masters file.

QUESTION: 2

For a MapReduce job, on a cluster running MapReduce v1 (MRv1), what's the relationship between tasks and task templates?

- A. There are always at least as many task attempts as there are tasks.
- B. There are always at most as many tasks attempts as there are tasks.
- C. There are always exactly as many task attempts as there are tasks.
- D. The developer sets the number of task attempts on job submission.

Answer: C

QUESTION: 3

What action occurs automatically on a cluster when a DataNode is marked as dead?

- A. The NameNode forces re-replication of all the blocks which were stored on the dead DataNode.
- B. The next time a client submits job that requires blocks from the dead DataNode, the JobTracker receives no heart beats from the DataNode. The JobTracker tells the NameNode that the DataNode is dead, which triggers block re-replication on the cluster.
- C. The replication factor of the files which had blocks stored on the dead DataNode is temporarily reduced, until the dead DataNode is recovered and returned to the cluster.
- D. The NameNode informs the client which write the blocks that are no longer available; the client then re-writes the blocks to a different DataNode.

Answer: A

Explanation:

How NameNode Handles data node failures? NameNode periodically receives a Heartbeat and a Blockreport from each of the DataNodes in the cluster. Receipt of a Heartbeat implies that the DataNode is functioning properly. A Blockreport contains a list of all blocks on a DataNode. When NameNode notices that it has not received a heartbeat message from a data node after a certain amount of time, the data node is marked as dead. Since blocks will be under replicated the system begins replicating the blocks that were stored on the dead datanode. The NameNode Orchestrates the replication of data blocks from one datanode to another. The replication data transfer happens directly between datanodes and the data never passes through the namenode. Note:If the Name Node stops receiving heartbeats from a Data Node it presumes it to be dead and any data it had to be gone as well.Based on the block reports it had been receiving from the dead node, the Name Node knows which copies of blocks died along with the node and can make the decision to re-replicate those blocks to other Data Nodes.It will also consult the Rack Awareness data in order to maintain the two copies in one rack, one copy in another rack replica rule when deciding which Data Node should receive a new copy of the blocks.

Reference:

24 Interview Questions & Answers for Hadoop MapReduce developers, How NameNode Handles data node failures'

QUESTION: 4

How does the NameNode know DataNodes are available on a cluster running MapReduce v1 (MRv1)

- A. DataNodes listed in the dfs.hosts file. The NameNode uses as the definitive list of available DataNodes.
- B. DataNodes heartbeat in the master on a regular basis.
- C. The NameNode broadcasts a heartbeat on the network on a regular basis, and DataNodes respond.
- D. The NameNode send a broadcast across the network when it first starts, and DataNodes respond.

Answer: B

Explanation:

How NameNode Handles data node failures? NameNode periodically receives a Heartbeat and a Blockreport from each of the DataNodes in the cluster. Receipt of a Heartbeat implies that the DataNode is functioning properly. A Blockreport contains a list of all blocks on a DataNode. When NameNode notices that it has not received a heartbeat message from a data node after a certain amount of time, the data node is marked as dead. Since blocks will be under replicated the system begins replicating the blocks that were stored on the dead datanode. The NameNode Orchestrates the replication of data blocks from one datanode to another. The replication data transfer happens directly between datanodes and the data never passes through the namenode.

Reference:

24 Interview Questions & Answers for Hadoop MapReduce developers, How NameNode Handles data node failures?

QUESTION: 5

Which three distcp features can you utilize on a Hadoop cluster?

- A. Use distcp to copy files only between two clusters or more. You cannot use distcp to copy data between directories inside the same cluster.
- B. Use distcp to copy HBase table files.
- C. Use distcp to copy physical blocks from the source to the target destination in your cluster.
- D. Use distcp to copy data between directories inside the same cluster.
- E. Use distcp to run an internal MapReduce job to copy files.

Answer: B, D, E

Explanation:

DistCp (distributed copy) is a tool used for large inter/intra-cluster copying. It uses Map/Reduce to effect its distribution, error handling and recovery, and reporting. It expands a list of files and directories into input to map tasks, each of which will

copy a partition of the files specified in the source list. Its Map/Reduce pedigree has endowed it with some quirks in both its semantics and execution.

Reference:

Hadoop DistCp Guide

QUESTION: 6

How does HDFS Federation help HDFS Scale horizontally?

- A. HDFS Federation improves the resiliency of HDFS in the face of network issues by removing the NameNode as a single-point-of-failure.
- B. HDFS Federation allows the Standby NameNode to automatically resume the services of an active NameNode.
- C. HDFS Federation provides cross-data center (non-local) support for HDFS, allowing a cluster administrator to split the Block Storage outside the local cluster.
- D. HDFS Federation reduces the load on any single NameNode by using the multiple, independent NameNode to manage individual parts of the filesystem namespace.

Answer: D

Explanation:

HDFS FederationIn order to scale the name service horizontally, federation uses multiple independent Namenodes/Namespaces. The Namenodes are federated, that is, the Namenodes are independent and don't require coordination with each other. The datanodes are used as common storage for blocks by all the Namenodes. Each datanode registers with all the Namenodes in the cluster. Datanodes send periodic heartbeats and block reports and handles commands from the Namenodes.

Reference:

Apache Hadoop 2.0.2-alpha
<http://hadoop.apache.org/docs/current/>

QUESTION: 7

Choose which best describe a Hadoop cluster's block size storage parameters once you set the HDFS default block size to 64MB?

- A. The block size of files in the cluster can be determined as the block is written.
- B. The block size of files in the Cluster will all be multiples of 64MB.
- C. The block size of files in the duster will all at least be 64MB.
- D. The block size of files in the cluster will all be the exactly 64MB.



Pass4sure Certification Exam Features;

- Pass4sure offers over **4500** Certification exams for professionals.
- More than **98,800** Satisfied Customers Worldwide.
- Average **99.8%** Success Rate.
- Over **150** Global Certification Vendors Covered.
- Services of **Professional & Certified Experts** available via support.
- Free **90 days** updates to match real exam scenarios.
- **Instant Download Access!** No Setup required.
- Price as low as **\$19**, which is 80% more **cost effective** than others.
- **Verified answers** researched by industry experts.
- Study Material **updated** on regular basis.
- Questions / Answers are downloadable in **PDF** format.
- Mobile Device Supported (**Android, iPhone, iPod, iPad**)
- **No authorization** code required to open exam.
- **Portable** anywhere.
- **Guaranteed Success.**
- **Fast**, helpful support **24x7**.



View list of All certification exams offered;
<http://www.ipass4sure.com/allexams.asp>

View list of All Study Guides (SG);
<http://www.ipass4sure.com/study-guides.asp>

View list of All Audio Exams (AE);
<http://www.ipass4sure.com/audio-exams.asp>

Download Any Certification Exam DEMO.
<http://www.ipass4sure.com/samples.asp>

To purchase Full version of exam click below;
<http://www.ipass4sure.com/allexams.asp>

3COM	CompTIA	Filemaker	IBM	LPI	OMG	Sun
ADOBE	ComputerAssociates	Fortinet	IISFA	McAfee	Oracle	Sybase
APC	CWNP	Foundry	Intel	McData	PMI	Symantec
Apple	DELL	Fujitsu	ISACA	Microsoft	Polycom	TeraData
BEA	ECCouncil	GuidanceSoftware	ISC2	Mile2	RedHat	TIA
BICSI	EMC	HDI	ISEB	NetworkAppliance	Sair	Tibco
CheckPoint	Enterasys	Hitachi	ISM	Network-General	SASInstitute	TruSecure
Cisco	ExamExpress	HP	Juniper	Nokia	SCP	Veritas
Citrix	Exin	Huawei	Legato	Nortel	See-Beyond	Vmware
CIW	ExtremeNetworks	Hyperion	Lotus	Novell	Google	

and many others.. See complete list [Here](#)

