

Rebalance an HDFS Cluster

Introduction/Objective

This document describes a design for redistributing data blocks when an HDFS cluster becomes imbalanced.

Motivation

A HDFS cluster can easily become imbalanced, for example, when a new data node joins the cluster. Because it does not hold much data, any map task assigned to the machine most likely does not access local data, thus increasing the use of network bandwidth. On the other hand, when some data nodes become full, new data blocks are placed on only non-full data nodes, thus reducing their read parallelism. It is important to redistribute data blocks when imbalance occurs.

Requirements

1. Rebalancing maintains data availability guarantees in the sense that it does not cause any block loss, change the number of replicas that a block has, or reduce the number of racks that a block resides.
2. An administrator should be able to invoke and interrupt rebalancing from a command line.
3. Block moving should be throttled so that it does not saturate the network.
4. Rebalancing should not cause a name node to be too busy to serve any incoming request.

Architecture Overview

1. What's balance?

A cluster is balanced if there is no under-utilized or above-utilized data node in the cluster. A data node's utilization is defined as the percentage of its used space.

An under-utilized data node is a node whose utilization is less than (average utilization – threshold).

An over-utilized data node is a node whose utilization is greater than (average utilization + threshold).

A threshold is user configurable. It can be set or changed from a command line. A default value is a utilization of 10%.

2. When to rebalance?

An administrator issues a command to trigger rebalancing. Rebalancing automatically shuts off when the cluster is balanced or when it does not make any progress. The command is defined as below:

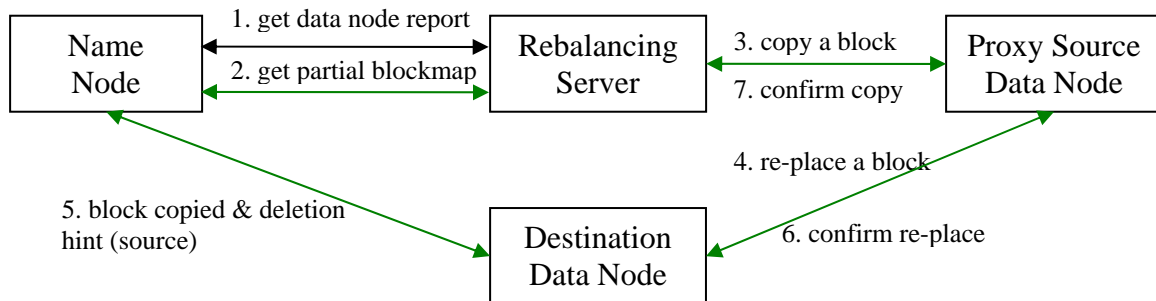
```
Hadoop balance [-t <threshold>]
```

Start rebalancing and change the default threshold if the option `-t` is set.

An administrator can simply press Ctrl-C to stop the rebalancing process.

3. How to rebalance?

To prevent the rebalancing from making the name node too busy to serve incoming requests, all the rebalancing decisions are made at a separate process, a rebalancing server at the machine where the command is issued. This implies that a rebalancing command should run at a machine that has fast connection to the DFS cluster.



Picture 1: Interactions among rebalancing server, name node, and data nodes

The rebalancing server makes rebalancing decision iteration by iteration. At each iteration the major goal is to make every over/under-utilized data node less imbalanced rather than reducing the number of imbalanced data nodes. In this way rebalancing is able to maximize the use of the network bandwidth. Picture 1 shows the interactions among a rebalancing server, a name node, and data nodes. Green lines represent new protocols for rebalancing. The numbers show the communication order. The rebalancing server first asks the name node for a data node report (step 1). After choosing all source and destination data nodes, it asks the name node for each source's partial blocks map (step 2). It then decides the blocks to move from sources to destinations. For each block, it picks up a proxy source which also contains the block but closer to the destination or less loaded than the source. It then instructs the proxy source to copy the block to the destination together with a hint that the block should be deleted from the source(step 3).

The proxy source then requests the destination to re-place the block on the source to its local disk together with the block data (step 4) and the destination copies the block to its disk. After the copy is done, the destination notifies the name node with a hint that a replica of the block should be deleted from the source (step 5). The name node then chooses a replica of the block to delete in favor of the hint if it does not violate requirement 1. The destination node then notifies the proxy source of the re-place status (step 6) and the proxy source notifies the rebalancing server of the operation status (step 7).

Architecture Detail

Rebalancing Server Design

1. A rebalancing command starts a Rebalancing Server at the user's machine.
2. It asks the name node for a data node report.
 - construct the network topology and calculate average utilization
 - construct an over-utilized data node list, an under-utilized data node list, an above-average utilized data node list, and a below-average utilized data node list
 - print the utilization variance on the screen as a progress report.
 - exit if the cluster is balanced or no progress has been made for three iterations; otherwise do the following:
 - go through the over/under-utilized data node list and make rebalancing decisions
 - a. choose source data nodes & destination data nodes. The selection criteria are:
 - For an under-utilized data node, source data nodes are chosen randomly from data nodes in the following priority:
 - (a). Over-utilized nodes in the same rack
 - (b). Above-average-utilized nodes in the same rack
 - (c). Over-utilized nodes in the other racks
 - (d). Above-average-utilized nodes in the other racks
 - For an over-utilized data node, destination data nodes are randomly chosen from data nodes in the following priority:
 - (a). Under-utilized nodes in the same rack
 - (b). Below-average-utilized data nodes in the same rack
 - (c). Under-utilized nodes in other racks
 - (d). Below-average-utilized nodes in the other racks
 - b. decide the number of bytes (not the blocks) to move from a source to a destination.
 - if the source is over-utilized, the size is $\text{MIN}(\text{size of the threshold}, 1\text{GB})$; if the source is above-average-utilized, the size is equal to $\text{MIN}(\text{its capacity} - \text{average capacity}, 2\text{GB})$
 - if the destination is under-utilized, the size is equal to $\text{MIN}(\text{the threshold size}, 1\text{GB})$; if the destination is below-average-utilized, the size is equal to $\text{MIN}(\text{average capacity} - \text{its capacity}, 2\text{GB})$

3. The rebalancing server asks the name node for a partial blocks map for each source. Each request is in the form of <source, size>, where size is 1GB.
4. The name node randomly chooses a certain number of blocks so that the total size of the chosen blocks is equal to the requested size. It sends all blocks back to the rebalancing server.
5. The rebalancing server keeps on choosing blocks to move for each source data node until the total number of chosen blocks is equal to the size decided previously or all the received blocks have been examined.
 - a. The block selection should not violate requirement 1. The selection criteria is that:
 - The block should not have other replicas on the rack that the destination resides if the source & the destination are on different racks;
 - The destination should not hold a replica of the block;
 - Disallow more than one replicas of the block to move concurrently;
 - b. Block move tasks are put to a pendingMove queue and copying requests are sent to source data nodes.
 - c. The pendingMove queue is indexed by sources, destinations, and blocks. It makes sure that both a source and a destination's queue size is at most 5, thus limiting the number of concurrent block transfers to be 5 per node. It also makes sure that at any given time, only one replica of a block is transferring.
 - d. A task is removed from pendingMove when it receives a confirmation from the source or when the timer is expired.
6. When all blocks are scanned, repeat from 3.
7. When all moves scheduled in step 2 are done and the pendingMove queue becomes empty, repeat from 2.

Data Node Support for Rebalancing

1. When a data node receives an OP_COPY command, it sends an OP_REPLACE command to the destination and transfer the block to the destination.
2. After the destination data node copied the block, it informs the name node that a block is copied with a hint that a replica at the source should be removed. The name node updates its block map. If the block is over replicated, choose a replica to remove in favor of the hint if removing the replica does not violate requirement 1.
3. The destination then informs the proxy server of the operation status and the proxy source notifies the rebalancing server of the status.
4. Each data node has a limited bandwidth for rebalancing. The default value for the bandwidth is 5MB/s.
5. Throttling is done at both source & destination sides. Each data node limits maximum number of concurrent data transfers (including both sending and receiving) for the rebalancing purpose to be 5. In the worst case, each data transfer has a limited bandwidth of 1MB/s.
6. Each sender & receiver has a Throttler. The primary method of the class is throttle(int numOfBytes). The parameter numOfBytes indicates the total number of bytes that the caller has sent or received since the last throttle is called. The method calculates the caller's I/O rate. If the rate is faster than the bandwidth limit, it sleeps to slow down the data transfer. After it wakes up, it adjusts its bandwidth limit if the

number of concurrent data transfers is changed. Throttle() is called periodically whenever a chunk is sent or received.

Rebalancing protocols

All new protocols to support rebalancing are marked as green in picture 1.

1. Name node protocol

A rebalancing server shares part of the name node work load, acting as a secondary name node. So a new interface, NamenodeProtocol, is created to support any secondary name node's communications with the name node.

```
/*
*****
* Protocol that a secondary NameNode uses to communicate with the NameNode.
* It's used to get part of the name node state
*****
interface NamenodeProtocol extends VersionedProtocol {
    public static final long versionID = 0L;

    /** Get a list of blocks belonged to <code>datanode</code>
     * whose total size is equal to <code>size</code>
     * @param datanode a data node
     * @param size requested size
     * @return a list of blocks
     */
    public LocatedBlocks getBlocks(DatanodeInfo datanode, long size)
        throws IOException; // Step 2 in Picture 1
}

```

2. Change to DatanodeProtocol

```
/**
* Allows a DataNode to tell the NameNode about blocks recently copied to this data node,
* each with a hint that which replica to remove if the block becomes over replicated.
*/
public void blockCopied(DatanodeRegistration registration, Block blocks[],
    DatanodeRegistration delHints[]) throws IOException; // Step 5 in Picture 1

```

3. Communications with Data nodes

Two rebalancing related operations will be added to communication protocols that define communications to data nodes. OP_COPY is to send to a proxy source that requests the proxy source to copy a block to another data node, while OP_REPLACE is to send to a destination that requests the destination to re-place a block at the source to the destination. Data nodes also send reply when operations are finished. The message formats are defined below:

OP_COPY request:	OP_COPY	BkID (long)	Source (UTF8)	Destination (DatanodeInfo)	// step 3
OP_REPLACE request:	OP_REPLACE	BkID (long)	Source (UTF8)	Data_With_Crc	//step 4
OP reply:	Operation status (short)				// steps 6,7

Table 1 Formats of Rebalancing Messages sent to Data Nodes

Discussion of possible race conditions

Race Conditions	How to handle it
Proxy Source data node does not have the block when OP_COPY is received	Return OP_FAILURE to the rebalancing server
Destination data node already has the block when OP_REPLACE is received.	Return OP_FAILURE to the proxy source
Proxy source receives an INVALIDATE on the same block that is being read.	
Destination data node receives an OP_WRITE or another OP_REPLACE on the same block when OP_REPLACE is in progress.	Reject the write command/the late coming OP_REPLACE.
Rebalancing server's request to name node failed	The rebalancing server exits
destination's confirmation to name node failed	block move continues, and block report will send the received block to the namenode later.
Any request to data node failed	Abort the block move

Table 2 Race Conditions and their solutions

Test Plan

Protocol test

1. Block Replacement Protocol Test

First set up a 3 node cluster with nodes NA, NB, and NC, which are on different racks. Then create a file with one block B with a replication factor 3. Finally add a new node ND to the cluster on the same rack as NC.

ID	Test Case	Expected Behavior
1	copy block B from ND to NA with del hint NC	Fail because proxy source ND does not have the block
2	Copy block B from NA to NB with del hint N	Fail because the destination NB contains the block
3	Copy block B from NA to ND with del hint NB	Succeed now block B is on NA, NC, and ND
4	Copy block B from NB to NC with del hint NA	Succeed but NA is not a valid del hint. So block B is on NA and NB. The third replica is either on NC or ND

2. Throttling test

ID	Test case	Expected Behavior
1	Create a throttler with 1MB/s bandwidth. Send 6MB data, and throttle at 0.5MB, 0.75MB, and in the end.	Actual bandwidth should be less than or equal to the expected bandwidth 1MB/s

3. Namenode Protocol Test: getBlocks

Set up a 2-node cluster and create a file with a length of 2 blocks and a replication factor of 2.

ID	Test case	Expected Behavior
1	Get blocks from datanode 0 with a size of 2 blocks	Return 2 blocks
2	Get blocks from datanode 0 with a size of 1 block	Return 1 block
3	Get blocks from datanode 0 with a size of 1	Return 1 block
4	Get blocks from datanode 0 with a size of 0	Receive an IOException
5	Get blocks from datanode 0 with a size of -1	Receive an IOException
6	Get blocks from a non-existent datanode	Receive an IOException

4. Balancer Test

All balancer tests are run with a threshold of 10% unless otherwise noted. All nodes have the same capacity unless otherwise noted. Test cases 3, 4, and 5 are automatic while all the other cases are manual. All tests are expected to meet the following requirements unless otherwise noted.

REQ1. Rebalancing does not cause the loss of a block;

REQ2. Rebalancing does not change the number of replicas that a block had;

REQ3. Rebalancing does not decrease the number of racks that a block had.

Besides the algorithm should try best to satisfy the following but provides no guarantee.

REQ4. Rebalancing process makes the cluster to be less and less imbalanced.

Id	Test Case	Expected Behavior
1	Bring up a one-node dfs cluster. Set files' replication factor to be 1 and fill up the node to 30% full. Then add an empty data node	Old node is 25% utilized and the new node is 5% utilized.

	with same capacity on the same rack.	
2	The same as 1 except that the empty new data node is on a different rack.	The same as above.
3	The same as 1 except that the empty new data node is half of the capacity as the old one.	Old one is 25% utilized and the new one is 10% utilized
4	Bring up a 2-node cluster and fill one node to be 60% and the other to be 10% full. All nodes are on different racks.	One node is 40% utilized and the other one is 30% utilized
5	Bring up a dfs cluster with nodes A and B. Set files' replication factor to be 2 and fill up the cluster to 30% full. Then add an empty data node C. All three nodes are on the same rack.	Old ones are 25% utilized and the new one is 10%
6	The same as test case 4 except that A, B, and C are on different racks.	The same as above
7	a. The same as test case 5 except that interrupt rebalancing.	The cluster is less imbalanced
	b. Restart rebalancing until it is done.	The same as 5
8	The same as test case 5 except that shut down namenode while rebalancing.	Rebalancing is interrupted
9	The same as test case 5 except that writing while rebalancing.	The cluster most likely bomes balanced, but may fluctuate
10	The same as test case 5 except that deleting while rebalancing	The same as above
11	The same as test case 5 except that writing & deleting while rebalancing	The same as above
12	Bring up a 3-node cluster. All nodes are on different racks. Set Files' replication factor to be 2. Fill 2 nodes to its 30% capacity and fill the other to its 60% capacity.	The first 2 nodes are 35% utilized & the 3 rd one is 50% utilized
13	The same as test case 11 except that fill 1 node to its 30% capacity, 1 node to 50%, the other to 80%.	Nodes become 43.33%, 50%, and 56.67% utilized respectively
14	Bring up a dfs cluster with 3 nodes, 2 of which are on rack A and another node is on rack B. Set replication factor to be 2 and fill up the cluster to its 20% capacity and add one empty node to rack A.	No blocks can be moved
15	The same as test case 12 except that the new node is added to rack B	Nodes become 15%, 15%, 20%, and 10% utilized respectively
16	Bring up a 4-node cluster. Nodes A, B are on rack 1 and Nodes C, D are on rack 2. Set replication factor to be 3 and fill Node A to 60%, B to 30%, C to 50%, and D to 20%.	Nodes become 50%, 40%, 20%, and 30% utilized respectively
17	Scalability test: populate a 750-node cluster	Cluster becomes balanced; File I/O's performance should not be noticeable slower.
	a. Run rebalancing after 3 node are added. b. Run rebalancing after 2 racks of nodes	

	(60 nodes) are added. c. Run rebalancing after 2 racks of nodes are added and run file writing/deleting at the same time	
--	---	--

5 Balancer Parameter Test

Id	Test Case	Expected Behavior
1	-threshold 10000000	exit with err msg
2	-threshold -10000000	exit with err msg
3	-threshold a	exit with err msg
4	-threshold -1	exit with err msg
5	-threshold 0	ok
6	-threshold -1.5	exit with err msg
7	-threshold 10.5	ok
8	-threshold 10.5a	exit with err msg
9	-thresholdd 10	exit with err msg
10	-threshold 10 10	exit with err msg