

CouldOnlyBeReplicatedTo

Could Only Be Replicated To ...

A common message people see is "could only be replicated to 0 nodes, instead of ...".

What does this mean? It means that the [Block Replication](#) mechanism of HDFS could not make any copies of a file (more specifically a block within a file) it wanted to create.

This can be caused by various conditions, including:

- No [DataNode](#) instances being up and running.
Action: look at the servers, see if the processes are running.
- The [DataNode](#) instances cannot talk to the server, through networking or Hadoop configuration problems.
Action: look at the logs of one of the [DataNodes](#).
- Your [DataNode](#) instances have no hard disk space in their configured data directories.
Action: look at the `dfs.data.dir` list in the node configurations, verify that at least one of the directories exists, and is writeable by the user running the Hadoop processes. Then look at the logs.
- Your [DataNode](#) instances have run out of space. Look at the disk capacity via the Namenode web pages.
Action: delete old files. Compress under-used files. Buy more disks for existing servers (if there is room), upgrade the existing servers to bigger drives, or add some more servers.
- The reserved space for a DN (as set in `dfs.datanode.du.reserved`) is greater than the remaining free space, so the DN thinks it has no free space.
Action: look at the value of this option and compare it with the amount of available space in your datanodes.
- There's not enough threads in the datanodes, and requests are being rejected.
Action: Look in the datanode logs, and the value of `dfs.datanode.handler.count`
- Some configuration problem is preventing effective two-way communication. In particular, we have seen that the combination of settings below to trigger the connectivity problem:

```
dfs.data.transfer.protection = authentication
dfs.encrypt.data.transfer = false
```

Action: check to see if this combination is set. If so, either disable protection or enable encryption.

- You may also get this message due to permissions.

This is not a problem in Hadoop, it is a problem (possibly configuration) in your cluster that you are going to have to fix on your own. Sorry.