

I'm new to Talend and while following steps in [Getting Started Guide](#), the tHDFSput failed with below error creating a ZERO byte file on my HDFS(Cloudera VM). After reading thro' forums I enabled dfs.client.use.datanode.hostname as "true". It would be helpful if someone can explain the root cause and the fix.

Starting job write_to_hdfs at 23:40 03/09/2016.

[statistics] connecting to socket on port 4033

[statistics] connected

[WARN]: org.apache.hadoop.util.NativeCodeLoader - Unable to load native-hadoop library for your platform... using builtin-java classes where appl

[ERROR]: org.apache.hadoop.util.Shell - Failed to locate the winutils binary in the hadoop binary path

java.io.IOException: Could not locate executable null\bin\winutils.exe in the Hadoop binaries.

at org.apache.hadoop.util.Shell.getQualifiedBinPath(Shell.java:378)

at org.apache.hadoop.util.Shell.getWinUtilsPath(Shell.java:393)

at org.apache.hadoop.util.Shell.<clinit>(Shell.java:386)

at org.apache.hadoop.util.StringUtils.<clinit>(StringUtils.java:79)

at org.apache.hadoop.security.Groups.parseStaticMapping(Groups.java:130)

at org.apache.hadoop.security.Groups.<init>(Groups.java:94)

at org.apache.hadoop.security.Groups.<init>(Groups.java:74)

at org.apache.hadoop.security.Groups.getUserToGroupsMappingService(Groups.java:303)

at org.apache.hadoop.security.UserGroupInformation.initialize(UserGroupInformation.java:283)

at org.apache.hadoop.security.UserGroupInformation.setConfiguration(UserGroupInformation.java:311)

at bigdata_demos.write_to_hdfs_0_1.write_to_hdfs.tHDFSput_1Process(write_to_hdfs.java:303)

at bigdata_demos.write_to_hdfs_0_1.write_to_hdfs.runJobInTOS(write_to_hdfs.java:685)

at bigdata_demos.write_to_hdfs_0_1.write_to_hdfs.main(write_to_hdfs.java:542)

File /user/cloudera/talend/directors.txt could only be replicated to 0 nodes instead of minReplication (=1). There are 1 datanode(s) running and

at org.apache.hadoop.hdfs.server.blockmanagement.BlockManager.chooseTarget4NewBlock(BlockManager.java:1595)

at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.getAdditionalBlock(FSNamesystem.java:3287)

at org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.addBlock(NameNodeRpcServer.java:677)

at org.apache.hadoop.hdfs.server.namenode.AuthorizationProviderProxyClientProtocol.addBlock(AuthorizationProviderProxyClientProtocol.java:213)

at org.apache.hadoop.hdfs.protocolPB.ClientNamenodeProtocolServerSideTranslatorPB.addBlock(ClientNamenodeProtocolServerSideTranslatorPB.java:485)

at org.apache.hadoop.hdfs.protocol.proto.ClientNamenodeProtocolProtos\$ClientNamenodeProtocol\$2.callBlockingMethod(ClientNamenodeProtocolProtos.java:1000)

at org.apache.hadoop.ipc.ProtobufRpcEngine\$Server\$ProtoBufRpcInvoker.call(ProtobufRpcEngine.java:617)

[WARN]: org.apache.hadoop.hdfs.DFSClient - DataStreamer Exception

org.apache.hadoop.ipc.RemoteException(java.io.IOException): File /user/cloudera/talend/directors.txt could only be replicated to 0 nodes instead of

at org.apache.hadoop.hdfs.server.blockmanagement.BlockManager.chooseTarget4NewBlock(BlockManager.java:1595)

at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.getAdditionalBlock(FSNamesystem.java:3287)

at org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.addBlock(NameNodeRpcServer.java:677)

at org.apache.hadoop.hdfs.server.namenode.AuthorizationProviderProxyClientProtocol.addBlock(AuthorizationProviderProxyClientProtocol.java:213)

at org.apache.hadoop.hdfs.protocolPB.ClientNamenodeProtocolServerSideTranslatorPB.addBlock(ClientNamenodeProtocolServerSideTranslatorPB.java:485)

at org.apache.hadoop.hdfs.protocol.proto.ClientNamenodeProtocolProtos\$ClientNamenodeProtocol\$2.callBlockingMethod(ClientNamenodeProtocolProtos.java:1000)

at org.apache.hadoop.ipc.RPC\$Server.call(RPC.java:1073)

at org.apache.hadoop.ipc.Server\$Handler\$1.run(Server.java:2086)

at org.apache.hadoop.ipc.Server\$Handler\$1.run(Server.java:2082)

at java.security.AccessController.doPrivileged(Native Method)

at javax.security.auth.Subject.doAs(Subject.java:415)

```
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1693)
at org.apache.hadoop.ipc.Server$Handler.run(Server.java:2080)
at org.apache.hadoop.ipc.ProtobufRpcEngine$Server$ProtoBufRpcInvoker.call(ProtobufRpcEngine.java:617)
at org.apache.hadoop.ipc.RPC$Server.call(RPC.java:1073)
at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2086)
at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2082)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:415)
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1693)
at org.apache.hadoop.ipc.Server$Handler.run(Server.java:2080)

at org.apache.hadoop.ipc.Client.call(Client.java:1471)
at org.apache.hadoop.ipc.Client.call(Client.java:1408)
at org.apache.hadoop.ipc.ProtobufRpcEngine$Invoker.invoke(ProtobufRpcEngine.java:230)
at com.sun.proxy.$Proxy9.addBlock(Unknown Source)
at org.apache.hadoop.hdfs.protocolPB.ClientNamenodeProtocolTranslatorPB.addBlock(ClientNamenodeProtocolTranslatorPB.java:404)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(Unknown Source)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(Unknown Source)
at java.lang.reflect.Method.invoke(Unknown Source)
at org.apache.hadoop.io.retry.RetryInvocationHandler.invokeMethod(RetryInvocationHandler.java:256)
at org.apache.hadoop.io.retry.RetryInvocationHandler.invoke(RetryInvocationHandler.java:104)
at com.sun.proxy.$Proxy10.addBlock(Unknown Source)
at org.apache.hadoop.hdfs.DFSOutputStream$DataStreamer.locateFollowingBlock(DFSOutputStream.java:1704)
at org.apache.hadoop.hdfs.DFSOutputStream$DataStreamer.nextBlockOutputStream(DFSOutputStream.java:1500)
at org.apache.hadoop.hdfs.DFSOutputStream$DataStreamer.run(DFSOutputStream.java:668)
File /user/cloudera/talend/movies.csv could only be replicated to 0 nodes instead of minReplication (=1). There are 1 datanode(s) running and 1 r

at org.apache.hadoop.hdfs.server.blockmanagement.BlockManager.chooseTarget4NewBlock(BlockManager.java:1595)
at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.getAdditionalBlock(FSNamesystem.java:3287)
at org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.addBlock(NameNodeRpcServer.java:677)
at org.apache.hadoop.hdfs.server.namenode.AuthorizationProviderProxyClientProtocol.addBlock(AuthorizationProviderProxyClientProtocol.java:213)

at org.apache.hadoop.hdfs.protocolPB.ClientNamenodeProtocolServerSideTranslatorPB.addBlock(ClientNamenodeProtocolServerSideTranslatorPB.java:485)

at org.apache.hadoop.hdfs.protocol.proto.ClientNamenodeProtocolProtos$ClientNamenodeProtocol$2.callBlockingMethod(ClientNamenodeProtocolProtos.ja

at org.apache.hadoop.ipc.ProtobufRpcEngine$Server$ProtoBufRpcInvoker.call(ProtobufRpcEngine.java:617)
[WARN ]: org.apache.hadoop.hdfs.DFSClient - DataStreamer Exception
org.apache.hadoop.ipc.RemoteException(java.io.IOException): File /user/cloudera/talend/movies.csv could only be replicated to 0 nodes instead of n

at org.apache.hadoop.hdfs.server.blockmanagement.BlockManager.chooseTarget4NewBlock(BlockManager.java:1595)
at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.getAdditionalBlock(FSNamesystem.java:3287)
at org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.addBlock(NameNodeRpcServer.java:677)
at org.apache.hadoop.hdfs.server.namenode.AuthorizationProviderProxyClientProtocol.addBlock(AuthorizationProviderProxyClientProtocol.java:213)

at org.apache.hadoop.hdfs.protocolPB.ClientNamenodeProtocolServerSideTranslatorPB.addBlock(ClientNamenodeProtocolServerSideTranslatorPB.java:485)

at org.apache.hadoop.hdfs.protocol.proto.ClientNamenodeProtocolProtos$ClientNamenodeProtocol$2.callBlockingMethod(ClientNamenodeProtocolProtos.ja

at org.apache.hadoop.ipc.RPC$Server.call(RPC.java:1073)
at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2086)
at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2082)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:415)
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1693)
```

```
at org.apache.hadoop.ipc.Server$Handler.run(Server.java:2080)
at org.apache.hadoop.ipc.ProtobufRpcEngine$Server$ProtoBufRpcInvoker.call(ProtobufRpcEngine.java:617)
at org.apache.hadoop.ipc.RPC$Server.call(RPC.java:1073)
at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2086)
at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2082)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:415)
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1693)
at org.apache.hadoop.ipc.Server$Handler.run(Server.java:2080)

at org.apache.hadoop.ipc.Client.call(Client.java:1471)
at org.apache.hadoop.ipc.Client.call(Client.java:1408)
at org.apache.hadoop.ipc.ProtobufRpcEngine$Invoker.invoke(ProtobufRpcEngine.java:230)
at com.sun.proxy.$Proxy9.addBlock(Unknown Source)
at org.apache.hadoop.hdfs.protocolPB.ClientNamenodeProtocolTranslatorPB.addBlock(ClientNamenodeProtocolTranslatorPB.java:404)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(Unknown Source)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(Unknown Source)
at java.lang.reflect.Method.invoke(Unknown Source)
at org.apache.hadoop.io.retry.RetryInvocationHandler.invokeMethod(RetryInvocationHandler.java:256)
at org.apache.hadoop.io.retry.RetryInvocationHandler.invoke(RetryInvocationHandler.java:104)
at com.sun.proxy.$Proxy10.addBlock(Unknown Source)
at org.apache.hadoop.hdfs.DFSOutputStream$DataStreamer.locateFollowingBlock(DFSOutputStream.java:1704)
at org.apache.hadoop.hdfs.DFSOutputStream$DataStreamer.nextBlockOutputStream(DFSOutputStream.java:1500)
at org.apache.hadoop.hdfs.DFSOutputStream$DataStreamer.run(DFSOutputStream.java:668)
44084 milliseconds
[statistics] disconnected
Job write_to_hdfs ended at 23:41 03/09/2016. [exit code=0]
```

Files created in HDFS:

Thanks !!