

I am getting below error intermittently in Spark job which loads data from one hawq table into another.
org.postgresql.util.PSQLException: Connection to ServerName:7432 refused. Check that the hostname and port are correct and that the postmaster is accepting TCP/IP connections.
Caused by: java.net.ConnectException: Connection timed out: connect
We have another utility written in java which also does the same, but this error has not appeared even once. We have been using the same jdbc driver and same connection details.
The same job has been executed successfully with same configuration couple of times. We tested the job with both dns name and ip address in the db configuration.
The pg_hba.conf file has been verified and are correct, the "TRUST" access has been given.
Appreciate your help in providing some pointers.

System Details-

1. The HortonWorks Data Platform V2.4.0 has been used and the job is running in YARN mode.

2. Talend - 6.2.1

I have created a simple spark job to load the data from one HAWQ table to another. Used following components in the job -

1. tHDFCConfiguration
2. tJDBCConfiguration
3. tJDBCInput
4. tJDBCOutput

JDBC has been used to connect to HAWQ database and details are given below -

Driver Class - "org.postgresql.Driver"

JAR - postgresql-9.4-1206-jdbc41.jar

JDBC URL - "jdbc: postgresql://servername:7432/databasename"

Log Output ->

Starting job HawqLoad at 05:09 16/08/2016.

connecting to socket on port 3467

connected

SLF4J: Class path contains multiple SLF4J bindings.

SLF4J: Found binding in

SLF4J: Found binding in

SLF4J: See for an explanation.

SLF4J: Actual binding is of type

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

: org.apache.spark.SparkConf - In Spark 1.0 and later spark.local.dir will be overridden by the value set by the cluster manager (via SPARK_LOCAL_DIRS in mesos/standalone and LOCAL_DIRS in YARN).

: org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory - The short-circuit local reads feature cannot be used because UNIX Domain sockets are not available on Windows.

: org.apache.spark.deploy.yarn.Client -

hdp.version is not found,

Please set HDP_VERSION=xxx in spark-env.sh,

or set -Dhdp.version=xxx in spark.{driver|yarn.am}.extraJavaOptions

or set SPARK_JAVA_OPTS="-Dhdp.verion=xxx" in spark-env.sh

If you're running Spark under HDP.

org.postgresql.util.PSQLException: Connection to **ServerName**:7432 refused. Check that the hostname and port are correct and that the postmaster is accepting TCP/IP connections.

at org.postgresql.core.v3.ConnectionFactoryImpl.openConnectionImpl(ConnectionFactoryImpl.java:247)

at org.postgresql.core.ConnectionFactory.openConnection(ConnectionFactory.java:65)

at org.postgresql.jdbc2.AbstractJdbc2Connection.<init>(AbstractJdbc2Connection.java:149)

at org.postgresql.jdbc3.AbstractJdbc3Connection.<init>(AbstractJdbc3Connection.java:35)

at org.postgresql.jdbc3g.AbstractJdbc3gConnection.<init>(AbstractJdbc3gConnection.java:22)

at org.postgresql.jdbc4.AbstractJdbc4Connection.<init>(AbstractJdbc4Connection.java:47)

at org.postgresql.jdbc4.Jdbc4Connection.<init>(Jdbc4Connection.java:30)

at org.postgresql.Driver.makeConnection(Driver.java:414)

at org.postgresql.Driver.connect(Driver.java:282)

at java.sql.DriverManager.getConnection(DriverManager.java:664)

at java.sql.DriverManager.getConnection(DriverManager.java:208)

at org.apache.spark.sql.execution.datasources.jdbc.JDBCRDD\$\$anonfun\$getConnector\$1.apply(JDBCRDD.scala:188)

at org.apache.spark.sql.execution.datasources.jdbc.JDBCRDD\$\$anonfun\$getConnector\$1.apply(JDBCRDD.scala:181)

```
at org.apache.spark.sql.execution.datasources.jdbc.JDBCRDD$.resolveTable(JDBCRDD.scala:121)
at org.apache.spark.sql.execution.datasources.jdbc.JDBCRelation.<init>(JDBCRelation.scala:91)
at org.apache.spark.sql.execution.datasources.jdbc.DefaultSource.createRelation(DefaultSource.scala:60)
at org.apache.spark.sql.execution.datasources.ResolvedDataSource$.apply(ResolvedDataSource.scala:158)
at org.apache.spark.sql.DataFrameReader.load(DataFrameReader.scala:119)
at local_project.hawqload_0_1.HawqLoad.tJDBCInput_1Process(HawqLoad.java:1489)
at local_project.hawqload_0_1.HawqLoad.run(HawqLoad.java:1950)
at local_project.hawqload_0_1.HawqLoad.runJobInTOS(HawqLoad.java:1754)
at local_project.hawqload_0_1.HawqLoad.main(HawqLoad.java:1641)
Caused by: java.net.ConnectException: Connection timed out: connect
at java.net.DualStackPlainSocketImpl.connect0(Native Method)
at java.net.DualStackPlainSocketImpl.socketConnect(DualStackPlainSocketImpl.java:79)
at java.net.AbstractPlainSocketImpl.doConnect(AbstractPlainSocketImpl.java:350)
at java.net.AbstractPlainSocketImpl.connectToAddress(AbstractPlainSocketImpl.java:206)
at java.net.AbstractPlainSocketImpl.connect(AbstractPlainSocketImpl.java:188)
at java.net.PlainSocketImpl.connect(PlainSocketImpl.java:172)
at java.net.SocksSocketImpl.connect(SocksSocketImpl.java:392)
at java.net.Socket.connect(Socket.java:589)
at org.postgresql.core.PGStream.<init>(PGStream.java:61)
at org.postgresql.core.v3.ConnectionFactoryImpl.openConnectionImpl(ConnectionFactoryImpl.java:129)
... 21 more
org.postgresql.util.PSQLException: Connection to ServerName:7432 refused. Check that the hostname and port are correct and that the postmaster is accepting TCP/IP connections.
at org.postgresql.core.v3.ConnectionFactoryImpl.openConnectionImpl(ConnectionFactoryImpl.java:247)
at org.postgresql.core.ConnectionFactory.openConnection(ConnectionFactory.java:65)
at org.postgresql.jdbc2.AbstractJdbc2Connection.<init>(AbstractJdbc2Connection.java:149)
at org.postgresql.jdbc3.AbstractJdbc3Connection.<init>(AbstractJdbc3Connection.java:35)
at org.postgresql.jdbc3g.AbstractJdbc3gConnection.<init>(AbstractJdbc3gConnection.java:22)
at org.postgresql.jdbc4.AbstractJdbc4Connection.<init>(AbstractJdbc4Connection.java:47)
at org.postgresql.jdbc4.Jdbc4Connection.<init>(Jdbc4Connection.java:30)
at org.postgresql.Driver.makeConnection(Driver.java:414)
at org.postgresql.Driver.connect(Driver.java:282)
at java.sql.DriverManager.getConnection(DriverManager.java:664)
at java.sql.DriverManager.getConnection(DriverManager.java:208)
disconnected
: org.apache.spark.rpc.netty.Dispatcher - Message RemoteProcessDisconnected( ServerName:60427) dropped.
java.lang.IllegalStateException: RpcEnv already stopped.
at org.apache.spark.rpc.netty.Dispatcher.postMessage(Dispatcher.scala:159)
at org.apache.spark.rpc.netty.Dispatcher.postToAll(Dispatcher.scala:109)
at org.apache.spark.rpc.netty.NettyRpcHandler.connectionTerminated(NettyRpcEnv.scala:630)
at org.apache.spark.network.server.TransportRequestHandler.channelUnregistered(TransportRequestHandler.java:94)
at org.apache.spark.network.server.TransportChannelHandler.channelUnregistered(TransportChannelHandler.java:89)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
at io.netty.channel.ChannelInboundHandlerAdapter.channelUnregistered(ChannelInboundHandlerAdapter.java:53)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
at org.apache.spark.sql.execution.datasources.jdbc.JDBCRDD$$anonfun$getConnector$1.apply(JDBCRDD.scala:188)
at org.apache.spark.sql.execution.datasources.jdbc.JDBCRDD$$anonfun$getConnector$1.apply(JDBCRDD.scala:181)
at org.apache.spark.sql.execution.datasources.jdbc.JDBCRDD$.resolveTable(JDBCRDD.scala:121)
at org.apache.spark.sql.execution.datasources.jdbc.JDBCRelation.<init>(JDBCRelation.scala:91)
at org.apache.spark.sql.execution.datasources.jdbc.DefaultSource.createRelation(DefaultSource.scala:60)
at org.apache.spark.sql.execution.datasources.ResolvedDataSource$.apply(ResolvedDataSource.scala:158)
at org.apache.spark.sql.DataFrameReader.load(DataFrameReader.scala:119)
at local_project.hawqload_0_1.HawqLoad.tJDBCInput_1Process(HawqLoad.java:1489)
```

at local_project.hawqload_0_1.HawqLoad.run(HawqLoad.java:1950)
at local_project.hawqload_0_1.HawqLoad.runJobInTOS(HawqLoad.java:1754)
at local_project.hawqload_0_1.HawqLoad.main(HawqLoad.java:1641)
Caused by: java.net.ConnectException: Connection timed out: connect
at io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
at io.netty.channel.ChannelInboundHandlerAdapter.channelUnregistered(ChannelInboundHandlerAdapter.java:53)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
at io.netty.channel.ChannelInboundHandlerAdapter.channelUnregistered(ChannelInboundHandlerAdapter.java:53)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
at io.netty.channel.DefaultChannelPipeline.fireChannelUnregistered(DefaultChannelPipeline.java:739)
at io.netty.channel.AbstractChannel\$AbstractUnsafe\$8.run(AbstractChannel.java:659)
at io.netty.util.concurrent.SingleThreadEventExecutor.runAllTasks(SingleThreadEventExecutor.java:328)
at java.net.DualStackPlainSocketImpl.connect0(Native Method)
at java.net.DualStackPlainSocketImpl.socketConnect(DualStackPlainSocketImpl.java:79)
at java.net.AbstractPlainSocketImpl.doConnect(AbstractPlainSocketImpl.java:350)
at java.net.AbstractPlainSocketImpl.connectToAddress(AbstractPlainSocketImpl.java:206)
at java.net.AbstractPlainSocketImpl.connect(AbstractPlainSocketImpl.java:188)
at java.net.PlainSocketImpl.connect(PlainSocketImpl.java:172)
at java.net.SocksSocketImpl.connect(SocksSocketImpl.java:392)
at java.net.Socket.connect(Socket.java:589)
at org.postgresql.core.PGStream.<init>(PGStream.java:61)
at org.postgresql.core.v3.ConnectionFactoryImpl.openConnectionImpl(ConnectionFactoryImpl.java:129)
... 21 more
at io.netty.util.concurrent.SingleThreadEventExecutor.confirmShutdown(SingleThreadEventExecutor.java:627)
at io.netty.channel.nio.NioEventLoop.run(NioEventLoop.java:362)
at io.netty.util.concurrent.SingleThreadEventExecutor\$2.run(SingleThreadEventExecutor.java:111)
at java.lang.Thread.run(Thread.java:745)
: org.apache.spark.rpc.netty.Dispatcher - Message RemoteProcessDisconnected(**ServerName**:60427) dropped.
java.lang.IllegalStateException: RpcEnv already stopped.
at org.apache.spark.rpc.netty.Dispatcher.postMessage(Dispatcher.scala:159)
at org.apache.spark.rpc.netty.Dispatcher.postToAll(Dispatcher.scala:109)
at org.apache.spark.rpc.netty.NettyRpcHandler.connectionTerminated(NettyRpcEnv.scala:630)
at org.apache.spark.network.server.TransportRequestHandler.channelUnregistered(TransportRequestHandler.java:94)
at org.apache.spark.network.server.TransportChannelHandler.channelUnregistered(TransportChannelHandler.java:89)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
at io.netty.channel.ChannelInboundHandlerAdapter.channelUnregistered(ChannelInboundHandlerAdapter.java:53)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
at io.netty.channel.ChannelInboundHandlerAdapter.channelUnregistered(ChannelInboundHandlerAdapter.java:53)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
at io.netty.channel.ChannelInboundHandlerAdapter.channelUnregistered(ChannelInboundHandlerAdapter.java:53)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
at io.netty.channel.DefaultChannelPipeline.fireChannelUnregistered(DefaultChannelPipeline.java:739)
at io.netty.channel.AbstractChannel\$AbstractUnsafe\$8.run(AbstractChannel.java:659)
at io.netty.util.concurrent.SingleThreadEventExecutor.runAllTasks(SingleThreadEventExecutor.java:328)
at io.netty.util.concurrent.SingleThreadEventExecutor.confirmShutdown(SingleThreadEventExecutor.java:627)
at io.netty.channel.nio.NioEventLoop.run(NioEventLoop.java:362)
at io.netty.util.concurrent.SingleThreadEventExecutor\$2.run(SingleThreadEventExecutor.java:111)

```
at java.lang.Thread.run(Thread.java:745)
Exception in thread "main" java.lang.RuntimeException: TalendJob: 'HawqLoad' - Failed with exit code: 1.
at local_project.hawqload_0_1.HawqLoad.main(HawqLoad.java:1651)
: local_project.hawqload_0_1.HawqLoad - TalendJob: 'HawqLoad' - Failed with exit code: 1.
Job HawqLoad ended at 05:10 16/08/2016.
```