

I had a issue when i m using tPigLoad component.
When i use hCatLoader option, i have a error, but when i use PigStorage option it is ok.
My envt is sandbox hortonwork 1.3 and TOS Big Data 5.4
thks for your help.

error :

Starting job Copy_of_Copy_of_pig_hcatLoader at 11:39 11/03/2014.

connecting to socket on port 3994

connected

: org.apache.pig.backend.hadoop.executionengine.HExecutionEngine - Connecting to hadoop file system at: hdfs://192.168.235.130:8020

: org.apache.pig.backend.hadoop.executionengine.HExecutionEngine - Connecting to map-reduce job tracker at: 192.168.235.130:50300

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

SLF4J: Failed to load class "org.slf4j.impl.StaticLoggerBinder".

SLF4J: Defaulting to no-operation (NOP) logger implementation

SLF4J: See <http://www.slf4j.org/codes.html#StaticLoggerBinder> for further details.

: hive.metastore - Trying to connect to metastore with URI thrift://192.168.235.130:9083

: hive.metastore - Waiting 1 seconds before next connection attempt.

: hive.metastore - Connected to metastore.

: org.apache.pig.tools.pigstats.ScriptState - Pig features used in the script: GROUP_BY

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MRCompiler - File concatenation threshold: 100 optimistic? false

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.CombinerOptimizer - Choosing to move algebraic foreach to combiner

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MultiQueryOptimizer - MR plan size before optimization: 1

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MultiQueryOptimizer - MR plan size after optimization: 1

: org.apache.pig.tools.pigstats.ScriptState - Pig script settings are added to the job

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - mapred.job.reduce.markreset.buffer.percent is not set, set to default 0.3

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - Using reducer estimator:

org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.InputSizeReducerEstimator

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.InputSizeReducerEstimator - BytesPerReducer=1000000000 maxReducers=999 totalInputFileSize=0

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - Setting Parallelism to 1

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - creating jar file Job6144726076440503266.jar

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - jar file Job6144726076440503266.jar created

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - Setting up single store job

: org.apache.pig.data.SchemaTupleFrontend - Key is false, will not generate code.

: org.apache.pig.data.SchemaTupleFrontend - Starting process to move generated code to distributed cacche

: org.apache.pig.data.SchemaTupleFrontend - Setting key with classes to deserialize []

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - 1 map-reduce job(s) waiting for submission.

: org.apache.hadoop.mapred.JobClient - Use GenericOptionsParser for parsing the arguments. Applications should implement Tool for the same.

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - 0% complete

: org.apache.hadoop.io.compress.snappy.LoadSnappy - Snappy native library not loaded

: org.apache.hadoop.mapred.FileInputFormat - Total input paths to process : 1

: org.apache.pig.backend.hadoop.executionengine.util.MapRedUtil - Total input paths (combined) to process : 1

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - HadoopJobId: job_201403110429_0002

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - Processing aliases

tPigAggregate_1_GROUP,tPigAggregate_1_row3_RESULT,tPigLoad_1_row2_RESULT

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - detailed locations: M:

tPigLoad_1_row2_RESULT,tPigAggregate_1_row3_RESULT,tPigAggregate_1_GROUP C: tPigAggregate_1_row3_RESULT,tPigAggregate_1_GROUP R: tPigAggregate_1_row3_RESULT

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - More information at: http://192.168.235.130:50030/jobdetails.jsp?jobid=job_201403110429_0002

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - job job_201403110429_0002 has failed! Stop running all dependent jobs

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - 100% complete

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.Launcher - There is no log file to write to.

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.Launcher - Backend error message

java.io.IOException: Deserialization error: org.apache.hadoop.hive.metastore.api.Table; local class incompatible: stream classdesc serialVersionUID = -946662244473213550, local class serialVersionUID = 3777104583087858118

at org.apache.hcatalog.common.HCatUtil.deserialize(HCatUtil.java:117)

```
at org.apache.hcatalog.mapreduce.HCatSplit.readFields(HCatSplit.java:144)
at org.apache.hadoop.io.serializer.WritableSerialization$WritableDeserialzier.deserialize(WritableSerialization.java:67)
at org.apache.hadoop.io.serializer.WritableSerialization$WritableDeserialzier.deserialize(WritableSerialization.java:40)
at org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.PigSplit.readFields(PigSplit.java:233)
at org.apache.hadoop.io.serializer.WritableSerialization$WritableDeserialzier.deserialize(WritableSerialization.java:67)
at org.apache.hadoop.io.serializer.WritableSerialization$WritableDeserialzier.deserialize(WritableSerialization.java:40)
at org.apache.hadoop.mapred.MapTask.getSplitDetails(MapTask.java:389)
at org.apache.hadoop.mapred.MapTask.runNewMapper(MapTask.java:726)
at org.apache.hadoop.mapred.MapTask.run(MapTask.java:363)
at org.apache.hadoop.mapred.Child$4.run(Child.java:255)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:396)
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1232)
at org.apache.hadoop.mapred.Child.main(Child.java:249)
Caused by: java.io.InvalidClassException: org.apache.hadoop.hive.metastore.api.Table; local class incompatible: stream classdesc serialVersionUID = -946662244473213550, local class
serialVersionUID = 3777104583087858118
at java.io.ObjectStreamClass.initNonProxy(ObjectStreamClass.java:560)
at java.io.ObjectInputStream.readNonProxyDesc(ObjectInputStream.java:1582)
at java.io.ObjectInputStream.readClassDesc(ObjectInputStream.java:1495)
at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1731)
at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1328)
at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1946)
at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1870)
at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1752)
at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1328)
at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1946)
at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1870)
at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1752)
at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1328)
at java.io.ObjectInputStream.readObject(ObjectInputStream.java:350)
at org.apache.hcatalog.common.HCatUtil.deserialize(HCatUtil.java:115)
... 14 more
: org.apache.pig.tools.pigstats.SimplePigStats - ERROR 2997: Unable to recreate exception from backed error: java.io.IOException: Deserialization error:
org.apache.hadoop.hive.metastore.api.Table; local class incompatible: stream classdesc serialVersionUID = -946662244473213550, local class serialVersionUID = 3777104583087858118
at org.apache.hcatalog.common.HCatUtil.deserialize(HCatUtil.java:117)
at org.apache.hcatalog.mapreduce.HCatSplit.readFields(HCatSplit.java:144)
at org.apache.hadoop.io.serializer.WritableSerialization$WritableDeserialzier.deserialize(WritableSerialization.java:67)
at org.apache.hadoop.io.serializer.WritableSerialization$WritableDeserialzier.deserialize(WritableSerialization.java:40)
at org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.PigSplit.readFields(PigSplit.java:233)
at org.apache.hadoop.io.serializer.WritableSerialization$WritableDeserialzier.deserialize(WritableSerialization.java:67)
at org.apache.hadoop.io.serializer.WritableSerialization$WritableDeserialzier.deserialize(WritableSerialization.java:40)
at org.apache.hadoop.mapred.MapTask.getSplitDetails(MapTask.java:389)
at org.apache.hadoop.mapred.MapTask.runNewMapper(MapTask.java:726)
at org.apache.hadoop.mapred.MapTask.run(MapTask.java:363)
at org.apache.hadoop.mapred.Child$4.run(Child.java:255)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:396)
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1232)
at org.apache.hadoop.mapred.Child.main(Child.java:249)
Caused by: java.io.InvalidClassException: org.apache.hadoop.hive.metastore.api.Table; local class incompatible: stream classdesc serialVersionUID = -946662244473213550, local class
serialVersionUID = 3777104583087858118
at java.io.ObjectStreamClass.initNonProxy(ObjectStreamClass.java:560)
at java.io.ObjectInputStream.readNonProxyDesc(ObjectInputStream.java:1582)
```

```
at java.io.ObjectInputStream.readClassDesc(ObjectInputStream.java:1495)
at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1731)
at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1328)
at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1946)
at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1870)
at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1752)
at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1328)
at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1946)
at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1870)
at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1752)
at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1328)
at java.io.ObjectInputStream.readObject(ObjectInputStream.java:350)
at org.apache.hcatalog.common.HCatUtil.deserialize(HCatUtil.java:115)
... 14 more
: org.apache.pig.tools.pigstats.PigStatsUtil - 1 map reduce job(s) failed!
: org.apache.pig.tools.pigstats.SimplePigStats - Script Statistics:
HadoopVersion PigVersion UserId StartedAt FinishedAt Features
1.2.0.1.3.0.0-107 0.11.1.1.3.0.0-107 rui.martins 2014-03-11 11:39:49 2014-03-11 11:41:03 GROUP_BY
Failed!
Failed Jobs:
JobId Alias Feature Message Outputs
job_201403110429_0002 tPigAggregate_1_GROUP,tPigAggregate_1_row3_RESULT,tPigLoad_1_row2_RESULT GROUP_BY,COMBINER Message: Job failed! Error - # of failed Map Tasks
exceeded allowed limit. FailedCount: 1. LastFailedTask: task_201403110429_0002_m_000000 /user/rui.martins/data/out,
Input(s):
Failed to read data from "talend.users"
Output(s):
Failed to produce result in "/user/rui.martins/data/out"
Counters:
Total records written : 0
Total bytes written : 0
Spillable Memory Manager spill count : 0
Total bags proactively spilled: 0
Total records proactively spilled: 0
Job DAG:
job_201403110429_0002

: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - Failed!
disconnected
Job Copy_of_Copy_of_pig_hcatLoader ended at 11:41 11/03/2014.
```