

Hi

i m try to user TPigStoreResult to store file in hdfs. But i had an error :

Thks for your response :

I use Hortonwork sandbox 2.0 / Talend TOS BigData 5.4.1

connecting to socket on port 3862

connected

: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS

: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.tracker is deprecated. Instead, use mapreduce.jobtracker.address

: org.apache.pig.backend.hadoop.executionengine.HExecutionEngine - Connecting to hadoop file system at: hdfs://192.168.235.129:8020

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

: 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:278)

at org.apache.hadoop.util.Shell.getWinUtilsPath(Shell.java:300)

at org.apache.hadoop.util.Shell.<clinit>(Shell.java:293)

at org.apache.hadoop.util.StringUtils.<clinit>(StringUtils.java:76)

at org.apache.hadoop.conf.Configuration.getTrimmedStrings(Configuration.java:1546)

at org.apache.hadoop.hdfs.DFSCClient.<init>(DFSCClient.java:519)

at org.apache.hadoop.hdfs.DFSCClient.<init>(DFSCClient.java:453)

at org.apache.hadoop.hdfs.DistributedFileSystem.initialize(DistributedFileSystem.java:136)

at org.apache.hadoop.fs.FileSystem.createFileSystem(FileSystem.java:2433)

at org.apache.hadoop.fs.FileSystem.access\$200(FileSystem.java:88)

at org.apache.hadoop.fs.FileSystem\$Cache.getInternal(FileSystem.java:2467)

at org.apache.hadoop.fs.FileSystem\$Cache.get(FileSystem.java:2449)

at org.apache.hadoop.fs.FileSystem.get(FileSystem.java:367)

at org.apache.hadoop.fs.FileSystem.get(FileSystem.java:166)

at org.apache.pig.backend.hadoop.datastorage.HDataStorage.init(HDataStorage.java:72)

at org.apache.pig.backend.hadoop.datastorage.HDataStorage.<init>(HDataStorage.java:58)

at org.apache.pig.backend.hadoop.executionengine.HExecutionEngine.init(HExecutionEngine.java:215)

at org.apache.pig.backend.hadoop.executionengine.HExecutionEngine.init(HExecutionEngine.java:122)

at org.apache.pig.impl.PigContext.connect(PigContext.java:301)

at org.apache.pig.PigServer.<init>(PigServer.java:220)

at org.apache.pig.PigServer.<init>(PigServer.java:205)

at bigdatademo.copy\_of\_step\_4\_pig\_count\_code\_0\_1.Copy\_of\_Step\_4\_Pig\_Count\_Code.tPigLoad\_2Process(Copy\_of\_Step\_4\_Pig\_Count\_Code.java:1991)

at bigdatademo.copy\_of\_step\_4\_pig\_count\_code\_0\_1.Copy\_of\_Step\_4\_Pig\_Count\_Code.runJobInTOS(Copy\_of\_Step\_4\_Pig\_Count\_Code.java:2691)

at bigdatademo.copy\_of\_step\_4\_pig\_count\_code\_0\_1.Copy\_of\_Step\_4\_Pig\_Count\_Code.main(Copy\_of\_Step\_4\_Pig\_Count\_Code.java:2514)

: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS

: org.apache.hadoop.conf.Configuration.deprecation - dfs.df.interval is deprecated. Instead, use fs.df.interval

: org.apache.hadoop.conf.Configuration.deprecation - mapred.task.tracker.http.address is deprecated. Instead, use mapreduce.tasktracker.http.address

: org.apache.hadoop.conf.Configuration.deprecation - dfs.max.objects is deprecated. Instead, use dfs.namenode.max.objects

: org.apache.hadoop.conf.Configuration.deprecation - mapred.userlog.retain.hours is deprecated. Instead, use mapreduce.job.userlog.retain.hours

: org.apache.hadoop.conf.Configuration.deprecation - hadoop.native.lib is deprecated. Instead, use io.native.lib.available

: org.apache.hadoop.conf.Configuration.deprecation - mapred.local.dir.minspacestart is deprecated. Instead, use mapreduce.tasktracker.local.dir.minspacestart

: org.apache.hadoop.conf.Configuration.deprecation - mapred.shuffle.read.timeout is deprecated. Instead, use mapreduce.reduce.shuffle.read.timeout

: org.apache.hadoop.conf.Configuration.deprecation - io.sort.spill.percent is deprecated. Instead, use mapreduce.map.sort.spill.percent

: org.apache.hadoop.conf.Configuration.deprecation - mapred.reduce.parallel.copies is deprecated. Instead, use mapreduce.reduce.shuffle.parallelcopies

: org.apache.hadoop.conf.Configuration.deprecation - mapred.submit.replication is deprecated. Instead, use mapreduce.client.submit.file.replication

: org.apache.hadoop.conf.Configuration.deprecation - mapred.local.dir.minspacekill is deprecated. Instead, use mapreduce.tasktracker.local.dir.minspacekill

: org.apache.hadoop.conf.Configuration.deprecation - mapred.task.profile is deprecated. Instead, use mapreduce.task.profile

: org.apache.hadoop.conf.Configuration.deprecation - dfs.data.dir is deprecated. Instead, use dfs.datanode.data.dir

: org.apache.hadoop.conf.Configuration.deprecation - mapred.heartbeats.in.second is deprecated. Instead, use mapreduce.jobtracker.heartbeats.in.second

: org.apache.hadoop.conf.Configuration.deprecation - mapred.output.compress is deprecated. Instead, use mapreduce.output.fileoutputformat.compress

: org.apache.hadoop.conf.Configuration.deprecation - mapred.healthChecker.interval is deprecated. Instead, use mapreduce.tasktracker.healthchecker.interval

: org.apache.hadoop.conf.Configuration.deprecation - mapred.task.timeout is deprecated. Instead, use mapreduce.task.timeout  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.temp.dir is deprecated. Instead, use mapreduce.cluster.temp.dir  
: org.apache.hadoop.conf.Configuration.deprecation - jobclient.completion.poll.interval is deprecated. Instead, use mapreduce.client.completion.pollinterval  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.tracker.persist.jobstatus.active is deprecated. Instead, use mapreduce.jobtracker.persist.jobstatus.active  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.name.dir is deprecated. Instead, use dfs.namenode.name.dir  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.output.compression.codec is deprecated. Instead, use mapreduce.output.fileoutputformat.compress.codec  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.shuffle.merge.percent is deprecated. Instead, use mapreduce.reduce.shuffle.merge.percent  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.map.max.attempts is deprecated. Instead, use mapreduce.map.maxattempts  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.reduce.input.buffer.percent is deprecated. Instead, use mapreduce.reduce.input.buffer.percent  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.task.cache.levels is deprecated. Instead, use mapreduce.jobtracker.taskcache.levels  
: org.apache.hadoop.conf.Configuration.deprecation - io.sort.factor is deprecated. Instead, use mapreduce.task.io.sort.factor  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.jobtracker.instrumentation is deprecated. Instead, use mapreduce.jobtracker.instrumentation  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.userlog.limit.kb is deprecated. Instead, use mapreduce.task.userlog.limit.kb  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.speculative.execution.slowNodeThreshold is deprecated. Instead, use mapreduce.job.speculative.slownodethreshold  
: org.apache.hadoop.conf.Configuration.deprecation - fs.checkpoint.dir is deprecated. Instead, use dfs.namenode.checkpoint.dir  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.skip.map.max.skip.records is deprecated. Instead, use mapreduce.map.skip.maxrecords  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.block.size is deprecated. Instead, use dfs.blocksize  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.access.time.precision is deprecated. Instead, use dfs.namenode.accesstime.precision  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.tracker.jobhistory.lru.cache.size is deprecated. Instead, use mapreduce.jobtracker.jobhistory.lru.cache.size  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.tracker.persist.jobstatus.hours is deprecated. Instead, use mapreduce.jobtracker.persist.jobstatus.hours  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.tracker.handler.count is deprecated. Instead, use mapreduce.jobtracker.handler.count  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.reduce.markreset.buffer.percent is deprecated. Instead, use mapreduce.reduce.markreset.buffer.percent  
: org.apache.hadoop.conf.Configuration.deprecation - io.sort.mb is deprecated. Instead, use mapreduce.task.io.sort.mb  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.task.profile.maps is deprecated. Instead, use mapreduce.task.profile.maps  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.map.tasks.speculative.execution is deprecated. Instead, use mapreduce.map.speculative  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.replication.min is deprecated. Instead, use dfs.namenode.replication.min  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.reduce.tasks is deprecated. Instead, use mapreduce.job.reduces  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.min.split.size is deprecated. Instead, use mapreduce.input.fileinputformat.split.minsize  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.name.edits.dir is deprecated. Instead, use dfs.namenode.edits.dir  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.replication.considerLoad is deprecated. Instead, use dfs.namenode.replication.considerLoad  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.tasktracker.dns.nameserver is deprecated. Instead, use mapreduce.tasktracker.dns.nameserver  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.tasktracker.taskmemorymanager.monitoring-interval is deprecated. Instead, use  
mapreduce.tasktracker.taskmemorymanager.monitoringinterval  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.tasktracker.expiry.interval is deprecated. Instead, use mapreduce.jobtracker.expire.trackers.interval  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.balance.bandwidthPerSec is deprecated. Instead, use dfs.datanode.balance.bandwidthPerSec  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.max.tracker.failures is deprecated. Instead, use mapreduce.job.maxtaskfailures.per.tracker  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.jobtracker.split.metainfo.maxsize is deprecated. Instead, use mapreduce.job.split.metainfo.maxsize  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.tracker.persist.jobstatus.dir is deprecated. Instead, use mapreduce.jobtracker.persist.jobstatus.dir  
: org.apache.hadoop.conf.Configuration.deprecation - job.end.retry.attempts is deprecated. Instead, use mapreduce.job.end-notification.retry.attempts  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.reduce.tasks.speculative.execution is deprecated. Instead, use mapreduce.reduce.speculative  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.job.counters.limit is deprecated. Instead, use mapreduce.job.counters.max  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.safemode.threshold.pct is deprecated. Instead, use dfs.namenode.safemode.threshold-pct  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.task.tracker.task-controller is deprecated. Instead, use mapreduce.tasktracker.taskcontroller  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.jobtracker.maxtasks.per.job is deprecated. Instead, use mapreduce.jobtracker.maxtasks.perjob  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.reduce.child.log.level is deprecated. Instead, use mapreduce.reduce.log.level  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.reduce.max.attempts is deprecated. Instead, use mapreduce.reduce.maxattempts  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.http.address is deprecated. Instead, use dfs.namenode.http-address  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.map.output.compression.codec is deprecated. Instead, use mapreduce.map.output.compress.codec  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.shuffle.input.buffer.percent is deprecated. Instead, use mapreduce.reduce.shuffle.input.buffer.percent  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.task.tracker.report.address is deprecated. Instead, use mapreduce.tasktracker.report.address  
: org.apache.hadoop.conf.Configuration.deprecation - keep.failed.task.files is deprecated. Instead, use mapreduce.task.files.preserve.failedtasks  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.name.dir.restore is deprecated. Instead, use dfs.namenode.name.dir.restore

: org.apache.hadoop.conf.Configuration.deprecation - dfs.https.client.keystore.resource is deprecated. Instead, use dfs.client.https.keystore.resource  
: org.apache.hadoop.conf.Configuration.deprecation - tasktracker.http.threads is deprecated. Instead, use mapreduce.tasktracker.http.threads  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.speculative.execution.slowTaskThreshold is deprecated. Instead, use mapreduce.job.speculative.slowtaskthreshold  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.backup.address is deprecated. Instead, use dfs.namenode.backup.address  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.backup.http.address is deprecated. Instead, use dfs.namenode.backup.http-address  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.acls.enabled is deprecated. Instead, use mapreduce.cluster.acls.enabled  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.max.tracker.blacklists is deprecated. Instead, use mapreduce.jobtracker.tasktracker.maxblacklists  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.tasktracker.indexcache.mb is deprecated. Instead, use mapreduce.tasktracker.indexcache.mb  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.skip.attempts.to.start.skipping is deprecated. Instead, use mapreduce.task.skip.start.attempts  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.tasktracker.reduce.tasks.maximum is deprecated. Instead, use mapreduce.tasktracker.reduce.tasks.maximum  
: org.apache.hadoop.conf.Configuration.deprecation - jobclient.output.filter is deprecated. Instead, use mapreduce.client.output.filter  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.permissions is deprecated. Instead, use dfs.permissions.enabled  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.jobtracker.restart.recover is deprecated. Instead, use mapreduce.jobtracker.restart.recover  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.local.dir is deprecated. Instead, use mapreduce.cluster.local.dir  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.speculative.execution.speculativeCap is deprecated. Instead, use mapreduce.job.speculative.speculativecap  
: org.apache.hadoop.conf.Configuration.deprecation - jobclient.progress.monitor.poll.interval is deprecated. Instead, use mapreduce.client.progressmonitor.pollinterval  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.safemode.extension is deprecated. Instead, use dfs.namenode.safemode.extension  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.datanode.max.xcievers is deprecated. Instead, use dfs.datanode.max.transfer.threads  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.map.child.log.level is deprecated. Instead, use mapreduce.map.log.level  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.output.compression.type is deprecated. Instead, use mapreduce.output.fileoutputformat.compress.type  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.tracker.retiredjobs.cache.size is deprecated. Instead, use mapreduce.jobtracker.retiredjobs.cache.size  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.https.need.client.auth is deprecated. Instead, use dfs.client.https.need-auth  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.tasktracker.dns.interface is deprecated. Instead, use mapreduce.tasktracker.dns.interface  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.task.profile.reduces is deprecated. Instead, use mapreduce.task.profile.reduces  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.https.address is deprecated. Instead, use dfs.namenode.https-address  
: org.apache.hadoop.conf.Configuration.deprecation - job.end.retry.interval is deprecated. Instead, use mapreduce.job.end-notification.retry.interval  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.jobtracker.job.history.block.size is deprecated. Instead, use mapreduce.jobtracker.jobhistory.block.size  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.child.tmp is deprecated. Instead, use mapreduce.task.tmp.dir  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.replication.interval is deprecated. Instead, use dfs.namenode.replication.interval  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.map.tasks is deprecated. Instead, use mapreduce.job.maps  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.committer.job.setup.cleanup.needed is deprecated. Instead, use mapreduce.job.committer.setup.cleanup.needed  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.tasktracker.map.tasks.maximum is deprecated. Instead, use mapreduce.tasktracker.map.tasks.maximum  
: org.apache.hadoop.conf.Configuration.deprecation - fs.checkpoint.edits.dir is deprecated. Instead, use dfs.namenode.checkpoint.edits.dir  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.queue.name is deprecated. Instead, use mapreduce.job.queue.name  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.write.packet.size is deprecated. Instead, use dfs.client-write-packet-size  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.jobtracker.taskScheduler is deprecated. Instead, use mapreduce.jobtracker.taskscheduler  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.skip.reduce.max.skip.groups is deprecated. Instead, use mapreduce.reduce.skip.maxgroups  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.permissions.supergroup is deprecated. Instead, use dfs.permissions.superusergroup  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.tracker.http.address is deprecated. Instead, use mapreduce.jobtracker.http.address  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.healthChecker.script.timeout is deprecated. Instead, use mapreduce.tasktracker.healthchecker.script.timeout  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.tasktracker.instrumentation is deprecated. Instead, use mapreduce.tasktracker.instrumentation  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.system.dir is deprecated. Instead, use mapreduce.jobtracker.system.dir  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.reuse.jvm.num.tasks is deprecated. Instead, use mapreduce.job.jvm.numtasks  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.inmem.merge.threshold is deprecated. Instead, use mapreduce.reduce.merge.inmem.threshold  
: org.apache.hadoop.conf.Configuration.deprecation - topology.script.number.args is deprecated. Instead, use net.topology.script.number.args  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.reduce.slowstart.completed.maps is deprecated. Instead, use mapreduce.job.reduce.slowstart.completedmaps  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.umaskmode is deprecated. Instead, use fs.permissions.umask-mode  
: org.apache.hadoop.conf.Configuration.deprecation - dfs.secondary.http.address is deprecated. Instead, use dfs.namenode.secondary.http-address  
: org.apache.hadoop.conf.Configuration.deprecation - fs.checkpoint.period is deprecated. Instead, use dfs.namenode.checkpoint.period  
: org.apache.hadoop.conf.Configuration.deprecation - topology.node.switch.mapping.impl is deprecated. Instead, use net.topology.node.switch.mapping.impl  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.tasktracker.tasks.sleeptime-before-sigkill is deprecated. Instead, use mapreduce.tasktracker.tasks.sleeptimebefore sigkill  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.compress.map.output is deprecated. Instead, use mapreduce.map.output.compress  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.merge.recordsBeforeProgress is deprecated. Instead, use mapreduce.task.merge.progress.records

: org.apache.hadoop.conf.Configuration.deprecation - mapred.shuffle.connect.timeout is deprecated. Instead, use mapreduce.reduce.shuffle.connect.timeout  
: org.apache.hadoop.conf.Configuration.deprecation - io.bytes.per.checksum is deprecated. Instead, use dfs.bytes-per-checksum  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.job.counters.limit is deprecated. Instead, use mapreduce.job.counters.max  
: org.apache.hadoop.conf.Configuration.deprecation - io.bytes.per.checksum is deprecated. Instead, use dfs.bytes-per-checksum  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.job.counters.limit is deprecated. Instead, use mapreduce.job.counters.max  
: org.apache.hadoop.conf.Configuration.deprecation - io.bytes.per.checksum is deprecated. Instead, use dfs.bytes-per-checksum  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.job.counters.limit is deprecated. Instead, use mapreduce.job.counters.max  
: org.apache.hadoop.conf.Configuration.deprecation - io.bytes.per.checksum is deprecated. Instead, use dfs.bytes-per-checksum  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.job.counters.limit is deprecated. Instead, use mapreduce.job.counters.max  
: org.apache.hadoop.conf.Configuration.deprecation - io.bytes.per.checksum is deprecated. Instead, use dfs.bytes-per-checksum  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.input.dir.recursive is deprecated. Instead, use mapreduce.input.fileinputformat.input.dir.recursive  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.max.split.size is deprecated. Instead, use mapreduce.input.fileinputformat.split.maxsize  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.min.split.size.per.rack is deprecated. Instead, use mapreduce.input.fileinputformat.split.minsize.per.rack  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.min.split.size.per.node is deprecated. Instead, use mapreduce.input.fileinputformat.split.minsize.per.node  
: hive.metastore - Trying to connect to metastore with URI thrift://192.168.235.129:9083  
: hive.metastore - Waiting 1 seconds before next connection attempt.  
: hive.metastore - Connected to metastore.  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.pig.tools.pigstats.ScriptState - Pig features used in the script: GROUP\_BY,ORDER\_BY,FILTER  
: org.apache.pig.newplan.logical.optimizer.LogicalPlanOptimizer - {RULES\_ENABLED=}  
: org.apache.pig.newplan.logical.rules.ColumnPruneVisitor - Columns pruned for tPigLoad\_2\_row1\_RESULT: \$0, \$1, \$2, \$3, \$4, \$5, \$6, \$7, \$8, \$10, \$11, \$12, \$13, \$14  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.textoutputformat.separator is deprecated. Instead, use mapreduce.output.textoutputformat.separator  
: 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: 3  
: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MultiQueryOptimizer - MR plan size after optimization: 3  
: org.apache.hadoop.yarn.client.RMProxy - Connecting to ResourceManager at /192.168.235.129:8050  
: 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.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.reduce.tasks is deprecated. Instead, use mapreduce.job.reduces  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.min.split.size is deprecated. Instead, use mapreduce.input.fileinputformat.split.minsize  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.reduce.tasks.speculative.execution is deprecated. Instead, use mapreduce.reduce.speculative  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.job.counters.limit is deprecated. Instead, use mapreduce.job.counters.max  
: org.apache.hadoop.conf.Configuration.deprecation - io.bytes.per.checksum is deprecated. Instead, use dfs.bytes-per-checksum  
: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - Reduce phase detected, estimating # of required reducers.  
: 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 Job1200056317904842081.jar  
: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - jar file Job1200056317904842081.jar created  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.jar is deprecated. Instead, use mapreduce.job.jar  
: 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 cache  
: 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.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - 0% complete  
: org.apache.hadoop.yarn.client.RMProxy - Connecting to ResourceManager at /192.168.235.129:8050  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.classpath.files is deprecated. Instead, use mapreduce.job.classpath.files  
: org.apache.hadoop.conf.Configuration.deprecation - fs.default.name is deprecated. Instead, use fs.defaultFS  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.cache.files is deprecated. Instead, use mapreduce.job.cache.files  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.output.value.class is deprecated. Instead, use mapreduce.job.output.value.class  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.job.counters.limit is deprecated. Instead, use mapreduce.job.counters.max  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.output.key.comparator.class is deprecated. Instead, use mapreduce.job.output.key.comparator.class  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.combine.class is deprecated. Instead, use mapreduce.job.combine.class  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.map.class is deprecated. Instead, use mapreduce.job.map.class  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.job.name is deprecated. Instead, use mapreduce.job.name  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.output.value.groupfn.class is deprecated. Instead, use mapreduce.job.output.group.comparator.class  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.reduce.class is deprecated. Instead, use mapreduce.job.reduce.class  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.inputformat.class is deprecated. Instead, use mapreduce.job.inputformat.class  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.output.dir is deprecated. Instead, use mapreduce.output.fileoutputformat.outputdir  
: org.apache.hadoop.conf.Configuration.deprecation - mapreduce.outputformat.class is deprecated. Instead, use mapreduce.job.outputformat.class  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.output.key.class is deprecated. Instead, use mapreduce.job.output.key.class  
: org.apache.hadoop.conf.Configuration.deprecation - io.bytes.per.checksum is deprecated. Instead, use dfs.bytes-per-checksum  
: org.apache.hadoop.mapreduce.JobSubmitter - Hadoop command-line option parsing not performed. Implement the Tool interface and execute your application with ToolRunner to remedy this.  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.task.id is deprecated. Instead, use mapreduce.task.attempt.id  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.input.dir is deprecated. Instead, use mapreduce.input.fileinputformat.inputdir  
: 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.hadoop.mapreduce.JobSubmitter - number of splits:1  
: org.apache.hadoop.conf.Configuration.deprecation - user.name is deprecated. Instead, use mapreduce.job.user.name  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.cache.files.filesizes is deprecated. Instead, use mapreduce.job.cache.files.filesizes  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.cache.files.timestamps is deprecated. Instead, use mapreduce.job.cache.files.timestamps  
: org.apache.hadoop.conf.Configuration.deprecation - mapred.working.dir is deprecated. Instead, use mapreduce.job.working.dir  
: org.apache.hadoop.mapreduce.JobSubmitter - Submitting tokens for job: job\_1393941587298\_0001  
: org.apache.hadoop.yarn.client.api.impl.YarnClientImpl - Submitted application application\_1393941587298\_0001 to ResourceManager at /192.168.235.129:8050  
: org.apache.hadoop.mapreduce.Job - The url to track the job: [http://sandbox.hortonworks.com:8088/proxy/application\\_1393941587298\\_0001/](http://sandbox.hortonworks.com:8088/proxy/application_1393941587298_0001/)  
: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - HadoopJobId: job\_1393941587298\_0001  
: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - Processing aliases  
tPigAggregate\_1\_GROUP,tPigAggregate\_1\_row3\_RESULT,tPigFilterColumns\_1\_row2\_RESULT,tPigFilterRow\_1\_row5\_RESULT,tPigLoad\_2\_row1\_RESULT  
: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - detailed locations: M:  
tPigLoad\_2\_row1\_RESULT,tPigFilterRow\_1\_row5\_RESULT,tPigFilterColumns\_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 - Oops! Some job has failed! Specify -stop\_on\_failure if you want Pig to stop immediately on failure.  
: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - job job\_1393941587298\_0001 has failed! Stop running all dependent jobs  
: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - 100% complete  
: 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  
2.2.0.2.0.6.0-76 0.12.0.2.0.6.0-76 rui.martins 2014-03-04 14:31:33 2014-03-04 14:34:21 GROUP\_BY,ORDER\_BY,FILTER  
Failed!  
Failed Jobs:  
JobId Alias Feature Message Outputs  
job\_1393941587298\_0001 tPigAggregate\_1\_GROUP,tPigAggregate\_1\_row3\_RESULT,tPigFilterColumns\_1\_row2\_RESULT,tPigFilterRow\_1\_row5\_RESULT,tPigLoad\_2\_row1\_RESULT  
GROUP\_BY,COMBINER Message: Job failed!  
Input(s):  
Failed to read data from "talend.weblog"  
Output(s):  
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\_1393941587298\_0001 -> null,  
null -> null,  
null  
  
: org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - Failed!