



链滴

Spark SQL 操作 hive 过程 rename 过程时间 长

作者: [ludengke95](#)

原文链接: <https://ld246.com/article/1568951362396>

来源网站: 链滴

许可协议: [署名-相同方式共享 4.0 国际 \(CC BY-SA 4.0\)](#)



Spark SQL 操作hive过程rename过程时间长

情况简介

hive版本: 1.2.1, spark版本: 2.3.0

2亿数据去重spark 任务时间: 12.5h (4h (去重) + 2.5h (不知道spark在干嘛, driver端没有日志, xecutor也没有日志) + 6h (Rname操作))

部分Rename日志。

```
2019-09-19 22:34:22,097 [Driver] INFO hive ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00002-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00002-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true
2019-09-19 22:34:22,111 [Driver] INFO hive ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00003-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00003-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true
2019-09-19 22:34:22,128 [Driver] INFO hive ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00004-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00004-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true
2019-09-19 22:34:22,143 [Driver] INFO hive ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00005-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00005-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true
```

2019-09-19 22:34:22,160 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00006-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00006-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,175 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00007-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00007-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,192 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00008-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00008-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,207 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00009-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00009-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,223 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00010-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00010-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,238 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00011-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00011-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,253 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00012-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00012-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,267 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00013-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00013-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,281 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00014-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00014-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,296 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00015-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00015-5d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true

2019-09-19 22:34:22,315 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_393364298521072924-1/-ext-10000/date=2018-02-24/part-00016-50d4ca62-4853-46d4-a1c9-2c155442904.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00016-5

```
d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true
2019-09-19 22:34:22,331 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/
pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_39336429852
1072924-1/-ext-10000/date=2018-02-24/part-00017-50d4ca62-4853-46d4-a1c9-2c15544290
4.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00017-5
d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true
2019-09-19 22:34:22,345 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/
pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_39336429852
1072924-1/-ext-10000/date=2018-02-24/part-00018-50d4ca62-4853-46d4-a1c9-2c15544290
4.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00018-5
d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true
2019-09-19 22:34:22,361 [Driver] INFO hive.ql.metadata.Hive - Renaming src: hdfs://cluster/
pps/hive/warehouse/partitioned/.hive-staging_hive_2019-09-19_19-06-31_561_39336429852
1072924-1/-ext-10000/date=2018-02-24/part-00019-50d4ca62-4853-46d4-a1c9-2c15544290
4.c000, dest: hdfs://cluster/apps/hive/warehouse/partitioned/date=2018-02-24/part-00019-5
d4ca62-4853-46d4-a1c9-2c15544290f4.c000, Status:true
```

spark sql执行hive sql任务

1. 会现在目标表中（1.21版本之后是默认位置目标表的文件夹）生成一个以hive-staging开头的临时文件夹，结果会在临时文件夹存放
2. 执行完成后会，将临时文件夹rename，放到对应的目标表文件下。

```

// 10/ Move individual files from src dir to dest dir.
boolean destIsSubDir = isSubDir(srcf, destf, fs, isSrcLocal);
try {
    if (inheritPerms || replace) {
        try {
            destStatus = shims.getFullFileStatus(conf, fs, destf.getParent());
            //if destf is an existing directory:
            //if replace is true, delete followed by rename(mv) is equivalent to replace
            //if replace is false, rename (mv) actually move the src under dest dir
            //if destf is an existing file, rename is actually a replace, and do not need
            // to delete the file first
            if (replace && !destIsSubDir) {
                LOG.debug("The path " + destf.toString() + " is deleted");
                fs.delete(destf, true);
            }
        } catch (FileNotFoundException ignore) {
            //if dest dir does not exist, any re
            if (inheritPerms) {
                destStatus = shims.getFullFileStatus(conf, fs, destf.getParent());
            }
        }
    }
    if (!isSrcLocal) {
        // For NOT local src file, rename the file

        if (hdfsEncryptionShim != null && (hdfsEncryptionShim.isPathEncrypted(srcf) || hdfsEncryptionShim.isPathEncrypted(destf) || !hdfsEncryptionShim.arePathsOnSameEncryptionZone(srcf, destf))) {
            LOG.info("Copying source " + srcf + " to " + destf + " because HDFS encryption zones
            success = FileUtils.copy(srcf.getFileSystem(conf), srcf, destf.getFileSystem(conf),
            true, // delete source
            replace, // overwrite destination
            conf);
        } else {
            if (destIsSubDir) {
                FileStatus[] srcs = fs.listStatus(srcf, FileUtils.HIDDEN_FILES_PATH_FILTER);
                if (srcs.length == 0) {
                    success = true; // Nothing to move.
                }
                for (FileStatus status : srcs) {
                    success = FileUtils.copy(srcf.getFileSystem(conf), status.getPath(), destf.getFileSystem(conf),
                    true, // delete source
                    replace, // overwrite destination
                    conf);

                    if (!success) {
                        throw new HiveException("Unable to move source " + status.getPath() + " to destination " + destf.getPath());
                    }
                }
            } else {
                success = fs.rename(srcf, destf);
            }
        }
    } else {
        // For local src file, copy to hdfs
        fs.copyFromLocalFile(srcf, destf);
        success = true;
    }
}

```

判断sparksql产生的中间文件是否和目标目录一致

如果一致，执行hadoop copy操作，删除源文件

如果不一致，执行hadoop rename操作。

无论一不一致，在spark日志里面看到的都是rename日志，但是实际上执行hadoop文件操作不一样。

从代码中可以看出，有两种策略：如果源目录和目标目录是同一个根目录，则会源目录下的每个文件行复制操作。反之，执行rename操作（只涉及namenode元数据，不会有额外数据操作）。

解决方案

修改hive-site.xml配置文件参数：

```

<property>
  <name>hive.exec.stagingdir</name>
  <value>/tmp/hive/.hive-staging</value>
  <description>hive任务生成临时文件夹地址</description>
</property>
</property>

```

```
<name>hive.insert.into.multilevel.dirs</name>
<value>true</value>
<description>hive.insert.into.mulltilevel.dirs设置成false的时候, insert 目标目录的上级目录必
存在; trued的时候允许不存在</description>
</property>
```

参考资料

1. [hadoop, hive中的mv\(rename\)操作](#)
2. [hive添加完hive.exec.stagingdir参数, 有的SQL报FileNotFoundException错](#)