

# Reprocessing runs from a RePiped job

This process assumes that there has been a successful run of the RePipe task for the Run in question. If the run is less than 30 days old, you might be able to do the repiping and reprocessing by L1 using the instructions on [Instructions for rolling back a halfPipe stream and subsequent L1 stream](#). Try that first.

The following instructions are adapted from instructions provided by Maria Elena (the original e-mail with her instructions is included at the end of the page). I've adapted it to use some environment variables to save a bunch of typing. After setting the environment variables in step 1, you can just copy and paste most of these commands (expect the ones in step 6).

These instructions assume you are logged in as glastraw

- 1) Set up some environment variables for the run and deliveries.

```
setenv rn <run number> e.g. 569213577
```

```
setenv dn1 <first delivery number> e.g. 190115006
```

```
setenv dn2 <second delivery number>
```

Note: Usually any given run will span two deliveries. If it only spans one, you only have to do the following for the single run. If the run spans three deliveries, just add a third one (i.e. dn3) and a third copy of each of the commands in the following steps, adjusting as needed.

- 2) Move to the staging area

```
cd /nfs/farm/g/glast/u28/stage
```

- 3) Create the staging directories that would have been created by the half-pipe.

```
mkdir -p $dn1/r0$rn $dn2/r0$rn
```

- 4) Copy over the event lists (Note: The following assumes that the repiped files are in the directory Repipe/\$rn/. If not change the command accordingly.)

```
cp ../RePipe/$rn/r0$rn/r0$rn-delivered.txt $dn1/r0$rn/r0${rn}_events_$dn1.txt
```

```
cp ../RePipe/$rn/r0$rn/r0$rn-delivered.txt $dn2/r0$rn/r0${rn}_events_$dn2.txt
```

- 5) Copy over the retired runs lists

```
cp ../RePipe/$rn/r0$rn/r0${rn}-retired.txt $dn1/retired_runs_$dn1.txt
```

```
cp ../RePipe/$rn/r0$rn/r0${rn}-retired.txt $dn2/retired_runs_$dn2.txt
```

- 6) Copy over the magic7 data

```
cp /nfs/farm/g/glast/u41/L1/runs/echo $rn | cut -c1-3`r0$rn/r0${rn}_${dn1}_v000_magic7Hp.txt $dn1/magic7_$dn1.txt
```

```
cp /nfs/farm/g/glast/u41/L1/runs/echo $rn | cut -c1-3`r0$rn/r0${rn}_${dn2}_v001_magic7Hp.txt $dn2/magic7_$dn2.txt
```

Note that the version number might be larger than v000 or v001, and there might be multiple versions for any given directory if the run had to rolled back, so you should check the run directory to see what the largest version number is for each delivery.

- 7) Copy over the event files

```
mv ../RePipe/$rn/r0$rn/r0${rn}-e000000000000000* $dn1/r0$rn/
```

```
mv ../RePipe/$rn/r0$rn/r0${rn}-e000000000000000* $dn2/r0$rn/
```

Just to be safe, do an `ls ../RePipe/$rn/r0$rn/` just to make sure no .evt files were left in the the RePipe directory. If they were, move them to the last downlink directory.

Note: if the run spanned three deliveries, you'll need to distribute some of the files from the second mv command to the third delivery directory. If there was only one, you can leave out all the 0's before the wildcard "\*" in the first command.

- 8) Then follow the instructions on this page - [Instructions for rolling back a halfPipe stream and subsequent L1 stream](#) - starting at the "Fix the Chunk Lists" section, with the following variations in the "Fix the Failed L1Proc Streams" section.

a) There may not be any run locks. If so, just roll back the stream you want to start first. (if there is a run lock, you don't need to make any changes to the instructions on the linked page)

b) Wait until the findChunk process is running (this will create a lock file).

c) Roll back all the other streams.

Note: Be sure to just rollback from findChunks. **Do not** roll back the entire stream from the button at the top of the page.

- 9) Once the streams are rolled back you need to create a dummy chunktoken as the RePipe process doesn't create them.

```
touch /nfs/farm/g/glast/u28/stage/chunktokens/r0$rn/pippo
```

This file and directory also need to be removed manually once the last of the deliveries has started processing. If you don't, the cleanupCompleteRun process will fail. If you forget, remove the chunk token and roll back cleanupCompleteRun.

Verify that the cleanupCompleteRun task actually ran for the last delivery. There have been cases where it hasn't run even when following these instructions. *Waiting on full details from Maria Elena on how to properly launch this job if it doesn't start automatically and what errors to look for.*

If the log of checkRun shows hpFinal=False and there are no other errors, you need to set the halfpipe status to "Complete" per the last step of #10 below and rollback the checkRun task.

10) Once all the L1Proc streams have finished run the following command

```
/afs/slac/g/glast/ground/bin/pipeline --mode PROD createStream --stream <run number> --define "runNumber=<run number>,  
l1RunStatus=Complete" setL1Status
```

to set the run completion status to "Complete".

Also follow the instruction on the [HTF Run completion status](#) page if necessary to set the status in the Runs section to complete as well.

11) Notify the data quality shifter that the run is ready to be reviewed.

## Some Issues to Watch Out For

### Overlapping runs in the same delivery

I've seen this happen exactly once but it caused some problems in the repipe. If you have two runs that need to be reprocessed, and they both have parts in the same delivery, completely process the first one before starting the second one. I had repiped both of the runs and copied all the event files over and then the processing of the first run removed the staging directory and the data for the second run was gone and I had to repipe it again.

### Issues with Magic7 data

Sometimes, especially with runs that span three deliveries, an arbitrary assignment of event files to each delivery as indicated by step 7 above results in the corresponding magic7 data files not covering the time range covered by the event files. This causes the doChunk substreams either at the beginning or end of the set to fail depending on where the data gap lies. The magic7 data typically extends quite a ways before and after the time of the delivery so this usually isn't a problem but it occasionally pops up. There are two solutions.

One is to move event files around so that they are covered by the magic7 data for each delivery. You could do this in advance by looking at roughly the chunk ranges in the originally processing and moving the appropriate files. Or if you don't discover it until after the fact, moving things then. However, in the latter case, many of the event files may have already been moved off the staging disk and that may require another repipe to set everything back up.

The other option is to merge in the magic7 data from the previous or later delivery (depending on where the problem is). This is what I typically do as it then just requires rolling back the failed doChunk runs. Simply grab the other magic7 data file from the run directory (the ones copied in step 6) into a temporary file, remove the overlapping data, and then concatenate the two files in the right order and rename the resultant file properly (i.e. *magic7\_<down load number>.txt*) in the proper directory. Then you can just roll back the failed processes and it will continue on.

## Original e-mail

This is the text of the original email from Maria Elena that started the process of creating this page:

Ok, here's how to create empty halfPipe folders after Re-Piping:

```
cd /nfs/farm/g/glast/u28/stage  
  
mkdir 181122013 181122014  
  
mkdir 181122013/r0564602913 181122014/r0564602913  
  
cp ../RePipe/564602913/r0564602913/r0564602913-delivered.txt 181122013/delivered_events_181122013.txt  
cp ../RePipe/564602913/r0564602913/r0564602913-delivered.txt 181122014/delivered_events_181122014.txt  
cp ../RePipe/564602913/r0564602913/r0564602913-retired.txt 181122013/retired_runs_181122013.txt  
cp ../RePipe/564602913/r0564602913/r0564602913-retired.txt 181122014/retired_runs_181122014.txt
```

```
cp /nfs/farm/g/glast/u41/L1/runs/564/r0564602913/r0564602913_181122013_v000_magic7Hp.txt 181122013
/magic7_181122013.txt

cp /nfs/farm/g/glast/u41/L1/runs/564/r0564602913/r0564602913_181122014_v001_magic7Hp.txt 181122014
/magic7_181122014.txt

mv ../RePipe/564602913/r0564602913/r0564602913-e000000000000000* 181122013/r0564602913/

mv ../RePipe/564602913/r0564602913/r0564602913-e000000000000000* 181122014/r0564602913/
```

Then proceed as usual with making new chunk lists etc.

Or course this specific run is the *\*wrong\** example because it was created with L1Proc/5.6, so in the end I didn't go

through with the actual L1 reprocessing. But this is how you would do the thing with the folders.