

Migrating HalfPipe to s3df

Directories:

- Incoming
 1. `/nfs/farm/g/glast/u23/ISOC-flight/Archive/fcopy -> /sdf/group/fermi/n/u23/ISOC-flight/Archive/fcopy` : incoming fastcopy
 2. `/nfs/farm/g/glast/u23/ISOC-flight/Archive/level0 -> /sdf/group/fermi/n/u23/ISOC-flight/Archive/level0` : incoming raw data. Both dirs were copied to s3df April 26, 2023. That's why tests on dev are done with runs of April 1, 2023.
- Still before HalfPipe
 1. RetDef files in `/sdf/group/fermi/a/isoc/flightOps/volumes/vol5/offline/halfPipe/v7r0p0_devel/u42/ISOC-flight/Downlinks/<dl>/`. Temporary, was `/nfs/farm/g/glast/u42/ISOC-flight/Downlinks/<dl>/`. **These files typically have time stamps of some dozens of seconds before the submission of the HalfPipe stream. Which process creates these files? We should ask Don!**
 2. `du of u42 (in MB)`:

722940	ISOC
713	ISOC-devel
2530798	ISOC-flight
538211	ISOC-flight/Downlinks
1974982	ISOC-flight/FswDumps
17726	ISOC-flight/Online
1075	ISOC-nightly
16641	ISOC-test
3272165	total

 - a. 2 TB of FswDumps. What's that? Does someone care?
 - b. **We would prefer to have u42/ISOC-flight/Downlinks copied to s3df.** It contains the downlink dirs since launch, a few transient dirs, and an evt dir. In evt are some evt files whose role isn't clear, we shouldn't lose them.
- During/after HalfPipe
 1. `/sdf/group/fermi/a/isoc/flightOps/volumes/vol5/offline/halfPipe/v7r0p0_devel/u28/stage` : staging directory, defined in config/stagedir

Processes:

- `launchChunks`: is just an embedded jython scriptlet launching a stream for every chunk: `for i in range(0, numChunks): pipeline.createSubstream("doChunk", i+1)`
- `doChunk.makeEvt.sh`
 - each `makeEvt` creates identical lock files `${HALFPIPE_OUTPUTBASE}/lock/halfpipe-${HALFPIPE_DOWNLINKID}`. Have to check when this file is being removed.
 - RetDef files `${HALFPIPE_OUTPUTBASE}/${HALFPIPE_DOWNLINKID}/RetDef-${HALFPIPE_DOWNLINKID}-${HALFPIPE_CHUNKID}.xml` are missing. u42 isn't on s3df. We copied these files from `rhel6-64:/nfs/farm/g/glast/u42/ISOC-flight/Downlinks/<dl>/`.
- `mergeIndices.sh`
 - `tokendir` is read dynamically: `tokendir=`cat ${taskBase}/config/stagedir/chunktokens`. For testing we changed `stagedir` from `/nfs/farm/g/glast/u28/stage` to `/sdf/group/fermi/a/isoc/flightOps/volumes/vol5/offline/halfPipe/v7r0p0_devel/u28/stage`.
 - at the end a `mergeEvt` stream for each run is created. This won't work in the container. This task is being taken by `launchMerge`, an embedded scriptlet in the xml.
- `launchMerge`: embedded scriptlet to launch `mergeEvt`
- `mergeEvt.sh`
 - should receive `${HALFPIPE_RUNSTART}` from `launchMerge`. Doesn't work right now. Workaround: read the decimal runid from `PIPELINE_STREAM` and convert to hex. Michael thinks this should be safe.
 - ~~not fault of mergeEvt, I guess: mergeEvt runs always on sdmlan14[56]. There is no /scratch but /scratch /lscratch needed --bind and /sdf/group/fermi/a/isoc/flightOps/volumes/vol5/offline/halfPipe/v7r0p0_devel/outputBase is mapped to uid 7234, while glastraw is 9503. Ticket open. Michael's fault. 7234 is Omodei, glastraw is 9503. All ok.~~
 - ~~Michael did something else yesterday to make mergeEvt finish. But probably the above should be understood and fixed first.~~
- `launchL1`
 - should work (I (Nicola) haven't added the "submit" script yet, but it gets to the point where it call the "submit stream" command successfully created.
 - `launchL1` creates streams of task into `OnlineAnalysis`. It seems to deal with calibration runs. The last time an `intOnlineAnalysis` stream finished successfully was 2013, the last partial 2017! We have to figure out if this task has any use, needs to be fixed, or simply can be skipped.
- `launchOnline`: **there is still the question about the "evt" directory...** we need to check on the rhel6 who is making this dir and how it is used? We also have **multiple directories called "stage"** (one in `/sdf/group/fermi/a/isoc/flightOps/volumes/vol5/offline/halfPipe/v7r0p0_devel/u28/stage/` one in `/sdf/group/fermi/a/isoc/flightOps/volumes/vol5/offline/halfPipe/v7r0p0_devel/outputBase/stage/` we probably need a bit of cleaning.