

L1 Data Processing Sequence

L1 currently uses NFS for long-term storage and AFS for short- and medium-term temporary storage. Both of these are likely to change, so think long-term where it says NFS below, and temporary when it says AFS.

Numbers below are for a single orbit, gleaned from examining OktoberTest data. They are approximate, and not always perfectly consistent.

There are a fair amount of parallelism and concurrency in the task, so the description below will not always be in order. Here's a flowchart for the



processing of one run:

There are 2 levels of parallel processing: Most work is done on chunks, which are created by upstreqam software (the halfpipe). Reconstruction breaks chunks up into crumbs. Each recon crumb job must copy the entire digi chunk file to scratch before processing its portion of the file.

.evt chunk files made by the halfpipe (~25 per run @ 125 MB) are copied from AFS to scratch. Each twice (digitization and fastmon). All read in parallel.

6.25 GB out (AFS)

Fastmon chunk results are moved from scratch to AFS, then from AFS to scratch for merging. Run-level fastmon file is moved from scratch to NFS, then copied from NFS to scratch for alarm scanning. These are small.

Digitized chunk files (150 MB) are moved from scratch to AFS, then copied from AFS to scratch for merging (1 copy of each, all read serially), digi and recon end-of-run and trending (4 copies), svac (1), and recon (usually 6 currently). That's 11 copies of each all read in parallel, plus 1 copy serial.

3.75 GB in (AFS)

45 GB out (AFS)

Run-level digi file (3.5 GB) is moved from scratch to NFS.

3.5 GB in (NFS)

Recon crumb files (~150 @ 125 MB recon, 29 MB cal, 8 MB merit, 0.1 MB gcr) are moved from scratch to AFS, then copied back to scratch for merging into chunks. Merges of each file type are performed in parallel. Crumbs are merged into chunks in groups of 6, the files going into a chunk are read serially. Chunk files (750 MB recon, 175 MB cal, 50 MB merit, .5 MB gcr) are moved from scratch to AFS. After all crumb-to-chunk merges are complete, crumb files are removed from AFS, they should usually stay on AFS for an hour or 2.

23 GB out (AFS)

23 GB in (AFS)

Recon-output chunk files are copied from AFS to scratch for merging into run-level files (18 GB recon, 4.1 GB cal, 1.2 GB merit, 12 MB gcr). The different types are merged in parallel, the chunks of each type are read serially. Run-level files are moved to NFS.

23 GB out (AFS)

23 GB in (NFS)

Recon and cal chunks are copied from AFS to scratch for recon end-of-run and trending (2 copies) and recon only for svac. The digi files are copied for all 3 also, but we already counted those. All parallel.

63 GB out (AFS)

Run-level merit file is copied from NFS to scratch for FT1 and FT2 generation. Also run-level digi file for FT2.

6 GB out (NFS)

FT1 and FT2 files get copied back & forth a couple of times. They're small.

Chunk-level monitoring (4 types, all fairly small <~1MB) and svac (190 MB) files are moved from scratch to AFS, then copied to scratch for merging. Run-level files are moved from scratch to NFS. svac run file is ~4.5 GB.

4.5 GB out (AFS)

4.5 GB in (NFS)

Chunk-level files stay on AFS until the run is completely processed. This is ~2hr in the best case, frequently ~6hr. Worst case is a week.

Every file that is written to NFS is read out more-or-less immediately by the data catalog crawler.

31 GB out (NFS)

----- total data moved -----

27 GB in (AFS)

165 GB out (AFS)

31 GB in (NFS)

37 GB out (NFS)

This is for one orbit, we expect 2 orbits per downlink. We're trying to do most of this in an hour. So that's ~400 GB/hr (~100MB/s) to and from temporary storage, and ~140 GB/hr (~40 MB/s) to/from permanent storage.

Worst-case requirement for volume of temporary storage (everything has to stay there for a week) is ~3.5 TB.