

Requirements document:

Inclusion of recocert in primary data processing with SAMGrid

Version 1.4

Preamble: As part of the ongoing 'gridification' of production activities at DØ primary processing is to be carried out with SAMGrid. To minimise the overall manpower effort this move has been made part of the SAM v7 transition. Data quality at DØ is monitored using the recocert package. Primary processing, reprocessing and MC generation are monitored using this package, though the actual procedure varies by activity. For primary processing, recocert is currently run immediately after the custom scripts which run reco on the d0farm. The output is stored on central sam owned discs (sam/durable1&2 on d0srv03). The verification of the recocert plots is the responsibility of the Data Quality (DQ) Group, as is the management of the durable locations (with the SAM team providing guidance wrt management scripts). This document details the requirements for including recocert in primary processing via SAMGrid and gives a possible implementation, following informal discussions between the SAMGrid team and DØ.

Requirements: Currently recocert is run on every unmerged thumbnail file, generating ~10MB file. There is currently no automated recovery of failed jobs as much data quality certification does not require all jobs to complete. However luminosity block by luminosity block certification of the calorimeter data quality (that runs in recocert) is carried out and this requires job recovery, so automated job recovery would be a medium term goal. Here automated job recovery means that it shall be possible to create in a single command a job that reruns recocert on all files missing a recocert output in from a given input dataset.

We intend to keep the output from recocert (once suitably compressed) on tape in SAM, however after initial creation the recocert output shall be stored to a durable location. This durable location shall be the same configurable location (at FNAL) independent of where the recocert jobs are run. The output from recocert is histograms, so data tier 'root-histogram'.

As the size of the histograms is largely independent of the number of entries we would prefer to run on merged thumbnails so reducing the output volume and the compression required before storage from the durable location into the tape system. The experiment would like to move to primary processing with SAMGrid as soon as possible and thus our immediate priority is for the inclusion of recocert without automated error recovery.

We request that the log files are available to the DQ team, i.e included in the logfiles of the grid job accessible from the monitoring pages. The current SamGrid policy of storing logfiles for at least 21 day is acceptable, but at the lower limit.

The need for automated data quality checks (all be it on fraction of the data) is particularly important as we emerge from shutdown.

Possible implementation: Following informal discussions between DØ and the SAMGrid team, and based around the need to expediently implement the existing functionality, without wasting manpower, we put forward the following for discussion.

- 1) Include running of recocert in one of the existing job types ignoring the need for recovery.
 - a. Choosing the merge jobs seems the best choice as it yields a compression factor of 5-10 wrt the current running.
 - b. The execution of recocert should be optional so that it can easily be turned off after the final solution is in place or in case of trouble. The default shall be to *not* run recocert.
 - c. Run recocert after the merged thumbnails have been stored to SAM, as part of the existing procedure. This will minimise the impact in case of failures as the normal merge step is complete before recocert is started.
 - d. Recocert output shall be stored to a durable location specified in the JDF.
 - e. To allow recovery from SAM information at a later stage, it is important that the parentage information for stored output is accurate.
 - f. The name of the output files should be constructed from the input filename as cert_<inputfilename>.root .
 - g. Standard error and output files should be included to the grids output tarballs.
- 2) Medium term: Consider creation of a new job type to allow job recovery.
 - a. A separate job type could be used to either only recover failure of a production using the above option or to take over the full production load, thereby possibly further reducing the number of output files.
(We will have to discuss whether we want to stick to one input to one output or whether it will be more efficient to have multiple inputs for one output. This may become a JDF options, like input_files_per_output_file)
 - b. Recovery should be implemented based on SAM as it is done for production jobs. Consider adapting of d0repro tools to submit and recover recocert jobs.
 - c. Responsibility for recovery needs to be assigned to central production or DQ. It seems reasonable the DQ group takes over recovery until the procedure has proven to be stable enough to become part of the central production.

Priorities: Inclusion of the basic functionality as is, i.e. without error recovery, is of the highest priority, at the same level as the completion of the MC migration to v7. The implementation of automated recovery has a lower priority as other, a priori higher priority goals, such as the inclusion of skimming are currently on hold and an overall priority needs to be set.