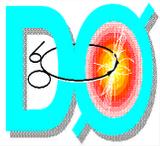


SAM Status and (short term) Plans

Lee Lueking

Dzero RACE

May 23, 2002



Current Status: V4_1_0_9 (bug fixes)

- 1) --impatient-end bug fixed.
- 2) --max-prefetched-files improved. Now it is supposed to limit SECOND and all subsequent tape mounts after project has started but it still allows all files on the FIRST(initial) tape to appear on the disc. Some of them will be locked some will not. Reason: As Per discussion with Igor. He reminded that project usually contains files that are not on single tape. We sacrifice disk usage to minimize tape mounts even if it violates maxPrefetchedFiles policy. maxPrefetchedFiles regulates how many files are pre-staged WHEN project finds out that no files are cached during its progress.
- 3) Dead project bug fix. Station now stops delivering files to a dead project.
- 4) Clearing of cached files on nodes not participating in delivery.
- 5) 2-line job submission bug in pbs adapter.
- 6) Fixed problem with arithmetic progression of in trasation requests that is related to --excess-satisfaction (for N=10, station would invoke 10 file transfers, then 9, 8, etc.).
- 7) Fixed problem with erasing files on non-local disks when looking for new space.



Current Status: V4_1_0_9 (bug fixes)

- 8) Fixed bug causing core dumps when file arrives at the node with no consumers.
- 9) Fixed bug when last file gets delivered to the file server cache, but intrastation transfer is never scheduled.
- 10) Fixed bug when file gets delivered to the worker node, but project is never notified.
- 11) Fixed problem with project in a distributed environment never getting end-of-file message if there are undeliverable files.
- 12) Fixed --file-cut bug. Undelivered files are counted against it.
- 13) Fixed bug which crashed station after project crashed.
- 14) Fixed project stalling bug with file delivery errors.



Current Status: V4_1_0_9 (New features)

- 1) Introduced start/stop, and the ability for station to act as its own client for stopping purposes.
- 2) Introduced --retry-attempts for failed file transfers.
- 3) Start/stop for fss.
- 4) Introduced --route flag that sets additional delivery rules for the file delivery (not intra-station transfers). Syntax is --route=<regexp>\<node1>[,node2]...
- 5) Station can get relationship between projects and jobs that are running in the batch system.



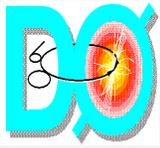
Current Status: V4_1_0_9 (Outstanding Issues)

- 1) Problem with small projects overlapping with big ones started first. The small one won't make any progress until the files are delivered for the big one.
 - 2) --route does not work with --constrain-delivery.
 - 3) Unlocked files change group ownership upon station restart. This has been fixed temporarily by hardcoding dzero group as the default for orphans.
 - 4) end of stream does not work if there is a file delivery error after consumer has been established
 - 5) station should not retrieve locations for all files when a project starts, but do it on a need-to-know basis
- 1) Introduced start/stop, and the ability for station to act as its own client for stopping purposes.



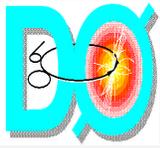
Current Status: V4_1_0_9 (Outstanding Issues)

- 1) Problem with small projects overlapping with big ones started first. The small one won't make any progress until the files are delivered for the big one.
 - 2) --route does not work with --constrain-delivery.
 - 3) Unlocked files change group ownership upon station restart. This has been fixed temporarily by hardcoding dzero group as the default for orphans.
 - 4) end of stream does not work if there is a file delivery error after consumer has been established
 - 5) station should not retrieve locations for all files when a project starts, but do it on a need-to-know basis
- 1) Introduced start/stop, and the ability for station to act as its own client for stopping purposes.



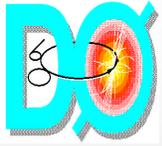
Current Status: V4_1_0_9 (Outstanding Minor Issues)

- 1) Possible problem with caching algorithm (not enough space for new files/encp's running out of space). (previous release)
- 2) Another possible problem with caching/uncaching algorithm (not uncaching the least-recently-used files). (previous release)
- 3) Multi-submission of jobs for distributed sam project. Code needs to be thoroughly tested. (previous release)
- 4) --impatiend-end improvement for the "Chris scenario".
- 5) The group disk allocation in the distributed environment needs to be designed/discussed, as well as oversubscribing of the available disk space in a single node environment.
- 6) Station tries to deliver a file to a node that is down (farm environment).



Current Status: V4_1_0_9 (Outstanding Minor Issues)

- 7) Problem with file transfers in error. Those can create situations in which project is not getting good files while bad files are retried many times. File transfers in error should not be retried immediately, but instead a new file transfer request should be created, put back into the queue, and optimizer should be asked for a new authorization (with reduced number of retrials for that location). In the short term, the number of retrials will be made configurable.
- 8) Intrastation transfers get behind encp's in the queue, causing slow file delivery.
- 9) get_enpc_priority.py script needs to be dealt with...
- 10) Inner workings of LRU need to be looked into.
- 11) endProcess needs to be implemented in the project master.
- 12) Station dump output needs to be fixed.



SAM Tasklist Highlites

- CRC transfer and verification
- Sam_batch wrapper
- Site optimizer
- Farm CP for intra-station transfers
- Reengineer cache management.
- File x-fer monitor
- Sam admin and bootstrap changes (and general distribution plan)
- Requirements for Lyon interfaces
- Nodes on VPN issue
- Decentralize name service
- D0mino backend and clued0 (almost done)
- Crummy file status
- Big ones like site autonomy, global data replica work, consolidation of station servers.



TNG DB Server Project

- True multi-threading supports multiple simultaneous clients. Number of active logins to database is configurable (database pool), drops idle connections.
- Memory (L1) and persistent (L2) server side object caching.
- Additional monitoring. Dynamic configuration.
- Simplification of D0om interface, pass data only (no CORBA objects). Provide user identification and tracking.
- More object oriented. More easily maintainable.
- Proxy feature will allow servers to be connected together to provide maximum reliability and scalability of the overall system.
- Jim Kowalkowski, and Steve White have done full evaluation and design. Now working on implementation. Expect to have first working version May 31 (no L2 cache or Proxy). All features in July.