Data Retention Policy Version 3

Policy by Folder

Red text indicates a change from the previous policy.

| Space | Quota | Backup | Lifetime | Comment |
|---------------------|-----------------|----------------|--------------------|--|
| xtc | None | Tape archive | 4 months | Raw data |
| usrdaq | None | Tape archive | 4 months | Raw data from users' DAQ systems |
| hdf5 | None | Tape archive | 4 months | Data translated to HDF5 |
| scratch | None | None | 4 months | Temporary data (lifetime not guaranteed) |
| xtc/hdf5 | 10TB | n/a | 2 years | Selected XTC and HDF5 runs |
| fte | 10TB | None | 2 years | Filtered, translated, compressed |
| results | 4TB | Tape backup | 2 years | Analysis results |
| calib | None | Tape backup | 2 years | Calibration data |
| User home | 20GB | Disk + tape | Indefinite | User code |
| Tape archive | - | - | 10 years | Raw data (xtc, hdf5, usrdaq) |
| Tape backup | - | - | Indefinite | User home, results and calib folder |
| Disk backup | - | - | Indefinite | Accessible under ~/.zfs/ |

Rationale for Proposed Policy

In the past couple of years we have observed some aspects of the LCLS data retention policy which are not ideal:

- 1. All experiments are treated equally even if a few institutions copy the data home and don't need to have the data on disk at SLAC: ideally we would reserve that disk space for other experiments that do rely on it.
- 2. Some folders and the different storage classes (short, medium and long term) were not always properly understood or used (e.g. ftc was often treated as scratch).
- 3. It's been hard to maintain the promise of preserving all the data on disk for their supposed lifetime: this has proved particularly tricky for scratch where the users can easily write tens of terabytes in a few hours.
- 4. Deleting data too early, i.e. when files are still being actively accesses, can cause large, and concurrent, restore operations from tape which affect negatively the performance of the system.

Also, we have studied the data usage over time and we have observed that:

• The rate at which data are accessed starts decreasing around 130 days after the experiment ends, both for raw and generated data.

Proposed Policy

Based on the observations above we propose the following modifications:

- Change the lifetime for raw and scratch data on disk to 4 months.
- After the initial 4 months period, the expiration status of a run is determined by the access pattern of its files.
 - This will be enforced by the constant monitoring of the file systems to determine which files were accessed and when. We will be also
 tracking if the files were actually analyzed or just "touched" to dodge the policy.
 - Analyzing the data extends the lifetime of the accessed data by 1 month from the access date.
 - Data restored from tape will stay on disk for 1 month, i.e. it's treated as a file access. (Users will be able to restore files by themselves through the web portal interface.)
- Eliminate the short and medium storage classes.
- Eliminate the /ftc folder.
 - Data currently under ftc will be moved to scratch and the scratch policy will apply.
- Rename the usr folder as usrdaq.
- Rename the res folder as results.
- Increase the quota of the results folder to 4TB.

Notes

- Please do not store under the scratch folder data that you cannot recreate because this directory is not backed up and the oldest files
 on scratch may be deleted at any time to make space for data from new experiments.
- The tape archive (xtc, hdf5, usrdaq) and the tape backup (results, home) are fundamentally different:
 - In the tape archive the folders are frozen after the end of the experiments and their contents are stored on tape once.
 - In the tape backup, the system takes snapshots of the folders as appear at a given time. this implies that files which are deleted from disk are eventually, i.e. after a long enough time, also deleted from tape.
- · For raw data the cleanup operations will affect all files, i.e. all streams and chunks, which make up one run, rather than individual files.

- After 2 years from the end of an experiment we'll remove the experiment from disk. At that point we'll take a snapshot of the results and calib folders and archive them to tape so that we can, upon request, restore an entire experiment back to disk. After 10 years we plan to remove the tapes with the archived raw data from the silos and store them in a safe environment.
- •
- The new policy will apply to all experiments, i.e. it will be retroactive, and its deployment date will coincide with the start of Run 14 (August 10th 2016). •

User Home

- Please do not store large files under your home, this space is meant for code/scripts, documents, etc, not science data.
 Users can check the used and available space under their home with a command like:

df -h ~<username>

• Users can automatically access snapshots of the backup of their home here:

~<username>/.zfs/snapshot/