

The Data Lifetime model

Simone Campana
CERN

The new data lifecycle model

- Every dataset will have a lifetime set at creation
 - The lifetime can be infinite (e.g. RAW data)
- The lifetime can be extended
 - E.g. if the dataset is recently accessed. Or if there is a known exception
- Every dataset will have a retention policy
 - E.g. RAW need at least 2 copies on tape. Need at least one copy of AODs on tape.
- Lifetime being agreed with ATLAS Computing Resources management

Effect of the data lifecycle model

- Datasets with expired lifetime can disappear at any time from (data)disk and datatape
 - groupdisk and localgroupdisk exempt
- “Organized” expiration lists will be distributed to groups
- ATLAS Distributed Computing will flexibly manage data replication and reduction
 - Within the boundaries of lifetime and retention
- For example
 - Increase/reduce the number of copies based on data popularity
 - Re-distribute data at T2s rather than T1s and viceversa
 - Move data to tape and free up disk space

Apply lifetime: what I would do in DQ2

- Write a script and get a list of all ATLAS datasets once a month.
 - With creationtime and atime.
- Apply the lifetime policy for each project/type as defined by CREM
- Get a list of datasets to delete, group them by project/type
- Email the lists to physics groups
- Delete after 4 weeks if no feedback (hide after 2 weeks)
 - If there is feedback, whitelist the complains

Flexibly manage replication/reduction

- I would keep little but some planned replication
- I would keep PD2P
- I would run periodic (monthly?) rearrangement scripts
 - Taking into account retention policy/lifetime/popularity
- Examples follow

Example: EVNTs

- Retention policy: do not lose
 - Keep at least 2 copies on disk (no tape)
- Lifetime: 4 years
- Pre-placement: 4 copies at T1s
- PD2P: none
- Reduction:
 - if not accessed in 1 year, reduce to 2 copies at T1s
 - If not accessed in 2 years, reduce to 2 copies (T1 and T2)

Example: data AODs

- Retention policy: do not lose
 - Keep at least 1 copy on tape
- Lifetime: 2 years
- Pre-placement: 1 copy on tape, 2 copies at T1s, 2 copies at T2s
- PD2P: replicate at T2s if used (tuning PD2P algorithm)
- Reduction:
 - if not accessed in 6 months, reduce to 1 tape + 1 T1 disk copy
 - if not accessed in 12months, reduce to 1 tape + 1 disk copy
 - if not accessed in 18months, reduce to 1 tape copy

Backup



Impact: gain of space

What happens if we set today the lifetime to the unused data?

Examples from my favorite types:

- ESDs: we have life w/o ESD, but some are immortal
- RDOs: produced on demand, deleted on demand (never)
- HITS: as they are on tape, why bother cleaning?

Datatype	Lifetime (months)	Expired data (TB)	Total data (TB)
ESD	12	17,327	23,262
RDO	12	942	2,223
HITS	24	6,163	14,938

As next step, we will do a full dry run based on the lifetimes discussed with Computing Resource Management

Impact: staging from tape

What happens if we remove all “unused” data from disk and keep it on tape?


“unused” here = not accessed in 9 months

Simulation based on last year’s data access

Tape access from Reconstruction and Reprocessing in 2014

We would have to restage from tape 20TB/week, compare with 1PB/week for reco/repro (2% increase). In terms of number of files, it is a 10% increase

Data staged per week (TB)



dashbord Staging Volume 2014-01-01 00:00 to 2014-08-31 00:00 UTC

