

Librarians workshop 2018-05-30 – **AFS phaseout**

- Background, AFS phaseout in general
 - Slow progress: FUSE / FUSEX, \$HOME only in LS2
- Good:
 - SW distribution is assumed to be a solved issue (*i.e* “use CVMFS”, “use RPM” or “use containers”)
 - Reassuring few mentions of AFS today
 - (massive) compilation is assumed to be a solved issue (*i.e do it locally*)
 - See EP-SFT status report from March 2017 – lots of ‘✓’

However, ...

- Several areas under */afs/cern.ch/sw/lcg/* see ongoing updates
- Some areas are even occasionally overloaded ... (ROOT releases)
- No(?) agreed(?) plan to eventually remove **all** areas under */afs/cern.ch/sw/lcg/*

= Agreement needed

- On **goals**
- On **strategy**

Proto-Goals

- “No AFS for large-scale software distribution”
 - Individual user and group programs can still be shipped to LXBATCH via AFS/EOS
- “No AFS for large-scale software building”
 - If for “technical reasons”¹ we need a staging area between build and CVMFS, it should not be user-accessible.



¹ such as EOS-FUSE being too slow for this

Proto-Strategy

- Per-product (/version?)
 - Discover residual usage
 - Negotiate end dates
 - Hide
 - Remove



- Overall
 - Communicate
 - Overall timelines
 - Track



- Discover:
 - By asking experiments and users upfront (?)
 - *And: agree that this usage is currently OK*
 - Residual unknown use: via **restrictive ACLs**
 - Easy to apply and revert, easy to tailor to agreed usage
 - Discover: which version of what software is being accessed by whom?
 - As high-level as possible, fine-grained as needed
 - Note: also should catch EP-SFT / LCG-internal access

- Negotiate:
 - Re-opening some area happens against agreed deadline
 - Per-experiment policy:
 - who can request re-opening (individual user, SW coordinator)?
 - Some (buggy) SW releases perhaps *should* not be used?
 - Track (JIRA?)

- Hide:
 - Unmount the volume
 - Easy, fast (minutes) to revert
 - Note: might not catch multi-mounted volumes. Tant pis.
- Remove:
 - Delete the volume
 - AFS backups are kept for 6months
 - Would take hours to revert
 - Most data is already on CVMFS, else could backup in CASTOR

- Communicate:
 - Discuss (&agree) here
 - users will need to know that access errors are intentional, and whom to talk to - otherwise will waste time debugging.
 - At least, need “post-facto” place (experiment statusboard, ITSSB)

(if OK so far...)

- Who does what, until when?
 - Announcements:
 - SFT-wide:
 - Per-experiment users:
 - “existing (known, agreed) usage” collection:
.....
 - Remove access:
 - Status tracking: