



CERN AFS phaseout: motivation, status & plans

Jan Iven, Alberto Pace

2021-05-19

What & Why & How

AFS: secure global shared (old) filesystem.
CMU → IBM → TransArc → OpenAFS

2012: OpenAFS community split

..2015: dev effort decline + longstanding issues (IPv6/crypto →? US defund?)

- ‘classic’ CERN computing services built around AFS, since >25years.
- Opensource - but “fragile” kernel client
- upcoming LHC “RUN3”.. ⇒ **F.U.D**

2016: risk-based decision: phaseout **CERN AFS** during LHC “long shutdown 2”

- no drop-in identified - migrate usecase-by-usecase
 - But: EOS played a central role
- (also: hoped for behaviour shift to web-based / mobile / BYOD computing...)

Formal project: representation from experiment & user groups

- Identify use cases, identify migration target, copy data; assess criticality
- IT: coordinate, but much work to be done by users.

Things done

AFS project spaces
One-by-one review

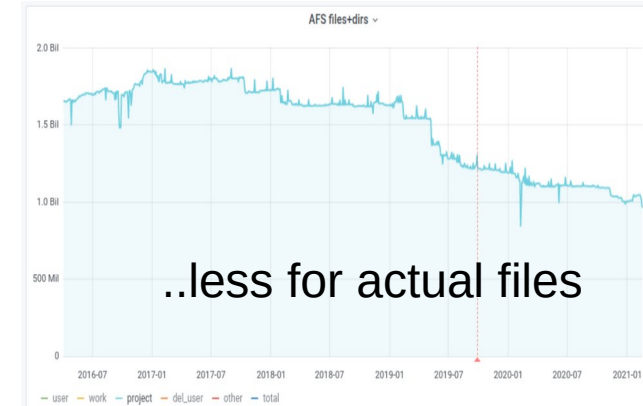
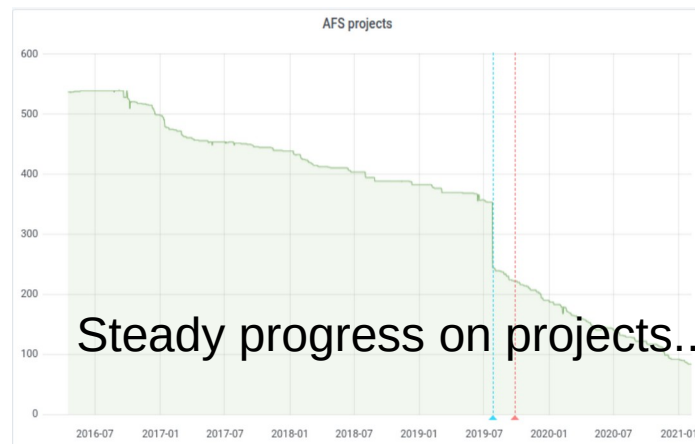
Archival

Non-controversial, but often no 'owner', still 'minimally used', 'historical' web sites...

Software distribution: → CVMFS (mostly..).

Investigations

use case discovery, also “what is left in /work?” (50% ROOT data), “have users moved on and forgot their data” (no), “dependencies from outside” (yes: “disco” tests)




Issues seen

EOS-FUSEX

- Initially: “convenience” access
 - Slow, did not cover many usecases
 - User expectations: 100% drop-in
- ⇒ fixed: rewritten
 - (caching+POSIX, complex)
- Slow roll-out
 - mixed with remote datacenter

EOS namespace

- Fast=in-memory - getting full..
 - ..even without +4B AFS files
- ⇒ fixed: rewritten
- Slow roll-out
 - Now: split & sharded



\$HOME

- semantics \neq sync & share
 - Integration & authentication issues
 - Data conflicts (DFS migration)
- ⇒ still open

BATCH


- High potential for DDoS
- conflicting use cases

Communication & confusion

- Same bugs, different instances.
- Stop/Pause/Start
- project structure vs lambda users & coffee tables

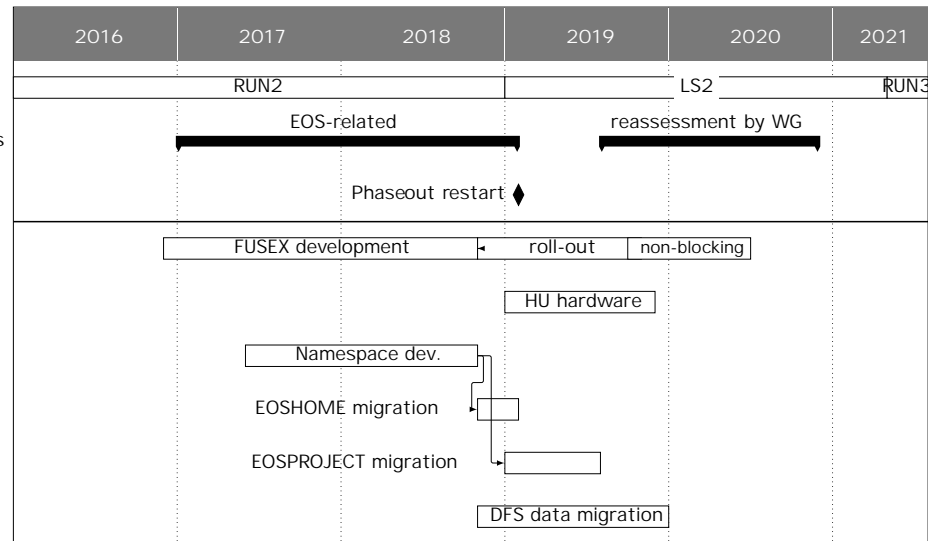
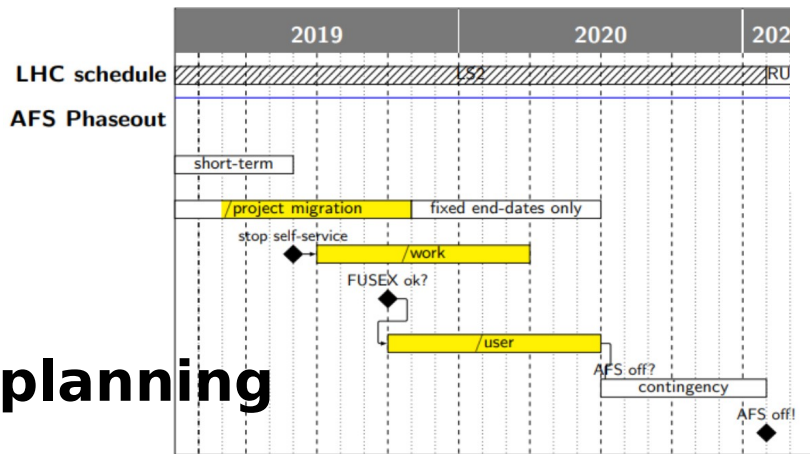
Plans.. and their evolution

2017 planning



AFS

EOS



“veto” periods
(unplanned)

reassessment

2019:

- 1) Time running out.. CERN AFS still used
 - Would have required major effort - by users
- 2) Upstream OpenAFS “still not dead yet”
 - Risks not realized ⇒ effort not justified

} ⇒ IT study group

Study group verdict (2020):

- no urgent need to change the current home directory
- continue to support AFS at an appropriate level
 - .. but also
- reduce the dependency on and utilization of AFS
- Long-term reducing the usage of AFS is crucial

Status & next?

AFS Service is stable & maintained

- smaller, and much better understood
- Experiments have minimized their exposure

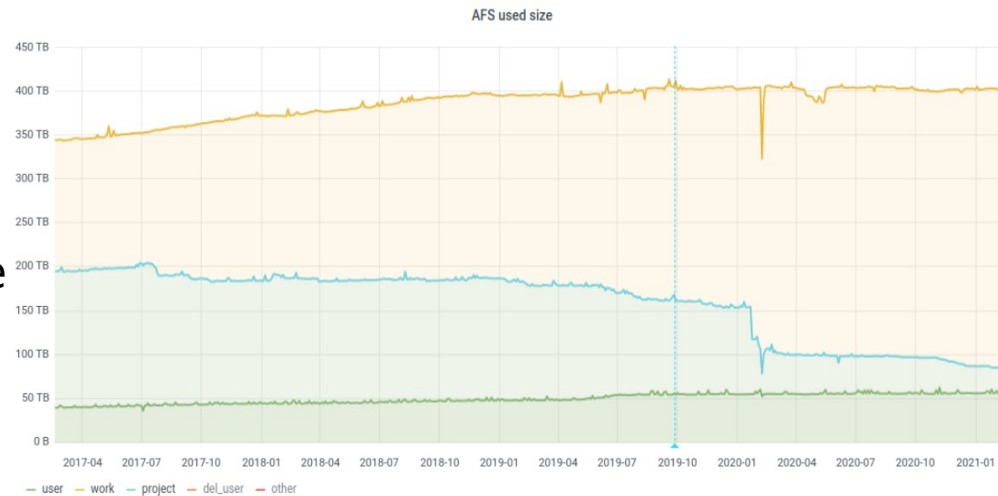
Usage has stabilized

- vs fast growth on all other storage services
- New use cases go “elsewhere”

Archival & cleanup: ongoing..

CERN IT will run AFS during LHC RUN3

- For homedirectories, existing personal workspace, and where EOS is not suitable (yet)
- Longer-term:
 - Keep running - while it still works for the current use cases
 - Batch (continued scaling) may increasingly affect stability



(no criticism of the OpenAFS project is implied – opensource, dedicated developers, reactive, architecture is astonishingly resilient...)



Questions?

home.cern