



DSS

Data & Storage Services

CERN
IT
Department

OpenAFS and IPv6 (at CERN)

Arne Wiebalck

HEPiX IPv6 WG F2F Meeting
CERN, Jan 24/25 2013



The information presented here is extracted from recent presentations of the OpenAFS community, discussions with the main OpenAFS developers and corresponding postings to the OpenAFS mailing lists (mostly from Jeffrey Altman which is one of the OpenAFS lead developers).



As of now:

- There is **no IPv6 support** in OpenAFS
- There is **no activity to add IPv6** to OpenAFS



- AFS is a complex distributed system
 - not: p2p with new address type
- IPv4 is embedded everywhere
 - DB schemes, RPC message formats, UBIK voting algorithm, command line parsers, ...
- Backwards-compatibility high priority
 - all combinations of IPv4 and IPv6 servers and clients
 - mixed server setups
- "Multi-home" support missing
 - everywhere except client-fs interaction



DSS

OK, so when will it come? :-)

- Estimate of development effort: 18 months
 - AFS standardization?
- On the AFS wishlist since 10+ years
 - maybe even from HEPiX
- But: there has never been a "serious" request
 - "serious" = "with funding"
- **IPv6 for OpenAFS won't come by itself**
 - funding situation needs to change
 - don't assume OpenAFS is for free!



DSS What to do when IPv6 happens?

- Option 1: Do nothing
- Option 2: Fund IPv6 for OpenAFS
- Option 3: Implement it ourselves
- Option 4: Look into AFS alternatives



DSS Option 1: Do nothing

- Dual-stack with private IPv4 addresses could "solve" the problem (at CERN)
 - no address translation needed, i.e. no performance hit
- Servers need to stay on public IPv4 addresses
 - $O(100)$ addresses
- But: no access for external IPv6-only clients
 - ... no access to any AFS cell!
- And: no access to external AFS cells from CERN AFS clients
 - could be solved with NAT
 - demand unclear



- Two commercial companies with lead OpenAFS developers
 - YFS.com and sinenomine.com
- A couple hundred kCHF seem a realistic
 - shareable with among HEPiX sites?
- Code should be in mainline OpenAFS
 - supported by core developers
 - published to share debugging
- But: delivery and timelines unclear
 - we never bought a feature so far
 - impact of AFS standardization unknown



DSS Option 3: Implement it ourselves

- The HEPiX community probably has the skills to implement IPv6 for OpenAFS
- The effort stays the same
 - read: same price
 - resources? (needs an AFS expert)
- Means some kind of OpenAFS fork
 - large changes are hard to get into main code tree
 - support from core OpenAFS team?
 - maintenance burden afterwards
- Hire a current core developer to work with us?



DSS Option 4: Look for alternatives

- To be combined with "Do nothing" option
- Re-evaluate the use cases for AFS
 - dropbox-like for home/docs, CVMFS for s/w, NFS for dev/batch ... ?
- Unclear if serious alternatives exist
 - 1.5 billion files, 50kHz access rate
 - risk to end up with a zoo of services
- What will users/experiments say to this?
 - AFS is pretty integrated in many workflows
 - AFS used space more than doubled during past year
 - Major performance improvements in the past months

- There is no IPv6 in OpenAFS as of today
- IPv6 in OpenAFS will not come (by itself or for free)
- Private IPv4 addresses may be easy way out
- If we (decide we) want IPv6, we should buy it and start the process now!

