

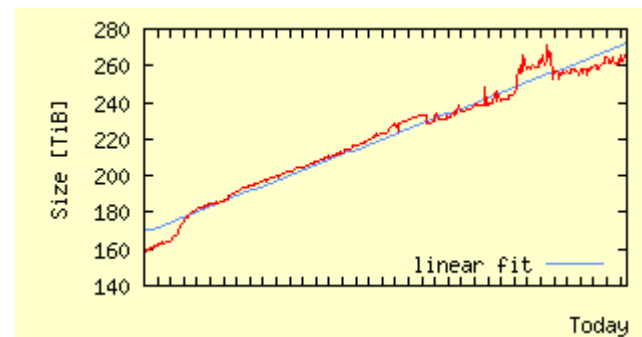
CERN AFS Site Report

Arne Wiebalck
(for the CERN AFS team)

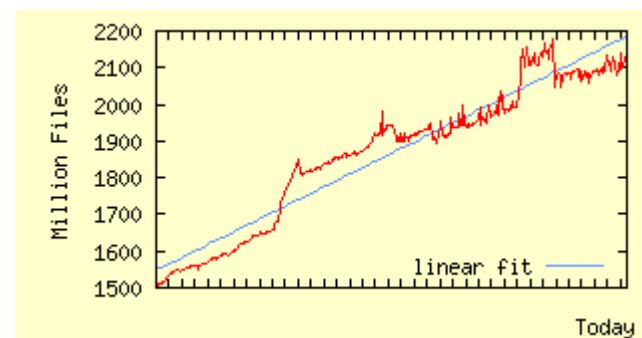
EAKC 2014
CERN, 26 March 2014

- 260TB of data
 - massive growth in the past 2 years
 - crossed 1PB capacity in 2013
 - user get 110GB by default
- 2.1 billion files
- 30k accounts / 15k clients
 - User IDs > 64k
 - Remember: 32766 is anonymous!
- ~50k accesses/sec
- ~82k volumes
- ~50 servers w/ 450 partitions
 - going down slightly


AFS used space in past 12 months



AFS #files in past 12 months



- AFS@CERN traditionally used “special” hardware
 - Fibre Channel fabrics, SAS drives, Sun/Oracle servers/arrays, ...
- Then moved to dual-head SAS enclosures
 - two servers, two disk arrays, fail-over to VMs, ...
 - SSDs with Facebook’s flashcache
- Now fully migrated to CC standard hardware
 - Simplification (procurement, setup, ops procedures, ...)
 - No more expert knowledge needed
 - Quad-servers: 64GB, 10GigE, 24x 3TB SATA



- **Clients: all on 1.6**
 - Tested during summer 2013 on a couple of hundred batch nodes
 - Driven by an incompatibility of the 1.4 tree with recent RH kernels
 - 1.6.5 our first version, moved to 1.6.6 after kernel crashes
- **Servers: 1.4.15-cern**
 - Some 1.6.6 (file) servers as well
 - Completion of migration to 1.6 expected in the coming months
- **CERN patches**
 - Mostly removed for 1.6.x, rest sent for review to Gerrit
 - Some CERN specifics (startup, packaging) kept
 - “Big loop retry”, “Volume throttling”
- **Configuration: Puppet!**
 - Moving away from a 2k line ksh script, and a frozen config
 - Puppet managed servers in production since Dec 2013
 - Only minor issues ... so far, so good ...

- AFS@CERN has been re-key'ed in Nov 2013
- Why did that take almost 6 months after the announcement was made?
 - Server updates to 1.4.15 in the cell
→ including porting ~140 patches ... from “everywhere”
 - AFS service kept Heimdal despite IT dep consolidation under AD
→ admin accounts
→ server patch
→ host keytabs
→ ops keytabs & scripts
 - Batch authentication system relied on old AFS DES key (and libdes)
→ needed patching & testing (~500k batch authentications/day!)
 - With AD there is “no way back” (to the old key)
→ needed careful testing & understanding of steps

- Replaced DB server hardware
 - Kept IP addresses, swapped hardware
 - Only every ~4 years
 - Worked w/o problems :)
- Replaced our backup system
 - Overcome various issues with the old system (scalability, maintainability)
 - See Kuba's talk on Friday for details
- Conducted an “OpenAFS and IPv6” survey in the High Energy Physics Community (HEPiX)
 - “dual stack” seems a valid option for now
 - funding IPv6 developments for OpenAFS would be difficult
 - mid- to long-term need for AFS (or any distributed FS) unclear

- CERN's new data center in Hungary?
 - 23ms seconds away
 - login and batch scheduling?
 - shadow volumes for disaster recovery?

- AFS and the Cloud?
 - CERN IT is moving towards an Agile infrastructure (virtualized services on OpenStack VMs configured with Puppet)
 - We have a patch that can move data into an S3 backend (tested w/ OpenStack Swift and Huawei)
 - We have started testing an OpenStack AFS server VM w/ a Ceph volume (make servers of arbitrary size, don't worry about h/w failures)

DSS

That's it, thanks!

