

dapnia



saclay

# From IBM AFS to OpenAFS

# Contents

---

dapnia



saclay

- Presentation
- First contacts with AFS
- DAPNIA cell
- Migration
- Future
- Issues

# CEA

dapnia



saclay

• Commissariat à l'Énergie Atomique : French Atomic Energy  
Commission created in 1945 :

- Nuclear power generation program
- Nuclear deterrence
- Micro- and nanotechnologies
- Medical imaging
- Toxicology
- Biotechnologies
- Astrophysics
- Nuclear physics

• 15 000 people, 1600 patents

• 9 main sites including Saclay

<http://www.cea.fr>

# DAPNIA

dapnia



saclay

•Département d'Astrophysique, de physique des Particules, de physique Nucléaire et de l'Instrumentation Associée :

- Astrophysics
- Particle physics
- Nuclear physics
- Detectors conception
- Computing
- ...

•Collaborations with other labs:

- BNL
- CERN
- DESY
- FNAL
- SLAC
- TJNL

<http://www-dapnia.cea.fr>

# First contact with AFS

dapnia



saclay

- The CERN decided to use AFS in 1992.
- AFS was quickly used to distribute CERN software (like the CERN library). CERN got the right from TRANSARC to « re-distribute » the client software to small labs to avoid them using ftp
- So we « bought » clients from CERN to install in Saclay as part of the CERN cell in 1994 :
  - rs\_aix32,
  - hp700\_ux90
- We never got the third one for DEC alpha running OSF 1 (just a beta version).

# Towards a « Saclay » cell

dapnia



saclay

- In 1995 we bought an IBM SP 2 with 6 nodes ( + 2 )
- I convinced my boss to also acquire an AFS license to use the distribution mechanism (rs\_aix32 sysname) through the AXLAN french society
- I decide to choose the cell name dapnia.saclay.cea.fr and, to minimize the cost, to elect my workstation « dphrsg » unique server of this new cell
- We did not get any other version of the software at that time; so we did continue to use the CERN AFS clients on our HP UX 9 machines (and later HP UX 10.2 and SunOS 4 and 5).
- As we already had NFS home directories, it was decided that the cell should not be widely opened to our users ...

dapnia



saclay

- The first server :
  - IBM RS 6000 220
  - AIX 3.2.5
  - 2 gigabytes external disk
  - AFS 3.4  $\beta$
- Major upgrade in 1996 :
  - IBM 43P (my new workstation)
  - AIX 4.1.4
  - 2 gigabytes disk
  - AFS 3.4a build 5.00
- Second upgrade
  - AIX 4.2.1
  - AFS 3.5 build 3.45
  - 9 gigabytes disk
  - Linux Red Hat 6 and Windows NT 4

# From AFS 3.6 to OpenAFS 1.0

---

dapnia



saclay

- IBM decision to make a copy of the source of AFS 3.6 for community development and maintenance

- Creation of OpenAFS

- But lack of a mechanism to manage the CellServDB file between the two products ...

- And nothing about the fsck for the AIX JFS v3 file system



# First contact with OpenAFS

dapnia



saclay

- Compilation on an IBM 43 P of the release 1.0 of OpenAFS : no problems
- Creation of a test cell : hep.cea.fr
  - IBM 43 P
  - AIX 4.3.2
- But the compilation failed for the 1.0.1, 1.0.2 and 1.0.3 ...
- So I decided to stay with IBM AFS :
  - Upgrade to AIX 4.3.2
  - AFS 3.6
- But It was decided to no more maintain the SP2 ( 9 nodes ) ...

# First crisis

---

- IBM Software withdrawal 902-085, April 9, 2002
  - « Withdrawal Effective July 8, 2002, with Service Extension to June 30, 2003 for AFS® 3.6 »
- Solutions ?
  - DSF ?
  - OpenAFS ?
- First of all get the latest version of AFS 3.6 :
  - Upgrade to AIX 4.3.3.10
  - And wait to see if we should go further with our cell or stop it with the IBM SP2 ...

dapnia



saclay

## Second crisis

---

dapnia

---

- After a lot of discussions, a group of particle physics searchers convince the “management” they needed AFS space in Saclay !

cea

---

- So we decide to migrate to OpenAFS using PC running Linux.

saclay

# Compatibility

dapnia



saclay

- I have first try to use the DAPNIA cell with an OpenAFS client :
  - AIX 4.3.2 client
  - Linux RH 6 and 7
- I have then added a file server running OpenAFS software :
  - AIX 4.3.2 server
- I have then added a file server running Linux Red HAT 7.3 with OpenAFS : DELL GX 240 with its own IDE disk.
- And finally I installed the database functionality on the Linux file server.

# Migration

dapnia



saclay

- I first got a new DELL GX 260 with Windows 2000 installed :
  - Installation of a SCSI card
  - Installation of a 80 Gbytes SCSI external disk
  - Installation of Red Hat 7.3.1 (kernel 2.4.18-3)
- I installed OpenAFS 1.2.7 on the DELL
- I then stopped the first server (IBM 43P) and restarted it with OpenAFS installed
  - Two database/file servers :
    - IBM 43P : AIX 4.3.3
    - DELL GX 260 : Linux 7.3.1 (kernel 2.4.18-3)
  - One additional file server
    - IBM 43P : AIX 4.3.3
  - OpenAFS 1.2.7
- I finally replaced the 80 Gbytes disk by a « box » containing 5 160 Gbytes disk ...

# Migration phase 2

dapnia

- As AIX was not obviously the main target of the developers of the OpenAFS, I asked for a second PC (without additional disk) to replace the IBM 43P. I got quickly a new DELL GX 260.

ced

- As I got some driver problems with Red Hat 7.3, I decided to install another distribution on this new machine.

saclay

- Migrate the Linux database/file server to OpenAFS 1.2.8
- Move the AFS volumes from the AIX partition to the EXT2 Linux partitions,
- Stop the first database server,
- Install Red Hat 8.0 on the new machine and OpenAFS 1.2.8 using the same IP address and name as the IBM machine and copying the files (/usr/afs/etc) from the remaining server ...
- And it worked !

# Red Hat new policy

---

dapnia



saclay

- End of the maintenance of the release 6, 7 and 8 at the end of 2003
  - Announce a new distribution Red Hat 9.0 with a dead line of end of April 2004 for its maintenance
  - Announce a new policy based on Red Hat Enterprise Linux
  - Announce of the Fedora project
- As a lot of people in France take their Holidays at the end of July and beginning of august, I decided to upgrade the two servers to Red Hat 9.0 and OpenAFS :
- Stop the servers
  - Uninstall the previous version of the OpenAFS RPMs
  - Upgrade the two machines to Red Hat 9.0
  - Install the new RPMs (1.2.9)
  - Restart the servers

# New tests

---

dapnia



saclay

- As I was not sure of the decision of the management about RHEL, I decided to do some tests with Fedora Core 1 as soon as possible and to restart my test cell `hep.cea.fr`

- IBM 43 P with AIX 4.3.3
- OpenAFS 1.2.10 and then 1.2.11
- DELL GX 240 with Fedora Core 1

- It works fine but I am the only user ...



# New users

---

dapnia



saclay

- A group of users of the Astrophysics service needing local space was pushed by their system administrator to use AFS !
- A new PC was installed with disk space dedicated to them
  - PC managed by another people :
    - Red HAT 7.3.1
    - SCSI drives
  - Full installation (so Ubik problem because of the IP address lower than our first database/file server one)
  - OpenAFS 1.2.8

# Next operating system ?

---

dapnia



saclay

- Wait for Fedora Core 2 ?
- The HEP UNIX community is discussing with Red Hat management to get better prices ... so RHEL 3 ?
- If this discussion should finish in a dead end, some labs are willing to create their own distributions using the RHEL 3 sources:
  - CERN project CEL
  - FNAL proposition HEPL
- Debian ?

# Issues

dapnia



saclay

- Difficult to use RPM distribution schema of OpenAFS and the (old) AFS upgrade mechanism (third PC needed ?)

- Bigger dependencies of OpenAFS of the different upgrades of the kernel and operating system than AFS on AIX

- Difficult to compile the sources

- I first decided to use EXT2 file systems and I followed the IN2P3 schema (partitions up to 8 Gbytes) so I have a lot of lost space on « my » big disks ...

- After a discussion during meeting with Alf, he told me I could use JFS or XFS file system, so I am using JFS on tests.