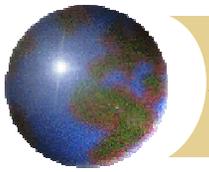


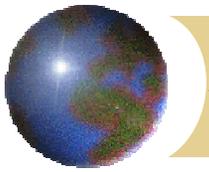
# *AFS Briefing*

CSS Department  
March 20<sup>th</sup>, 2003



# *Agenda*

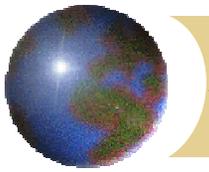
- ✚ Why We Are Here
- ✚ What is the problem
- ✚ Options
- ✚ Recommendation
- ✚ Usage
- ✚ Statistics
- ✚ Review



## *Why we are here?*

We need a strategic decision on AFS support.

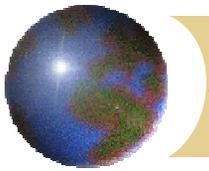
Our support contract ends on 5/1/2003.



# *What is the problem?*

## ❖ IBM

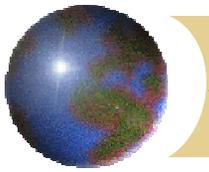
- ❖ IBM AFS support for 2003-2004 will cost \$399K, an increase of > 2000%.
- ❖ Current support contract ends 5/1/2003.



# *What is the problem?*

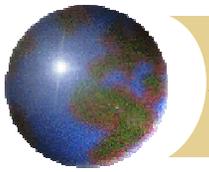
## ⊕ AFS deficiencies

- ⊞ Performance
- ⊞ Volumes > 8GB
- ⊞ Kerberos5 authentication
- ⊞ GRID awareness
- ⊞ Management Tools



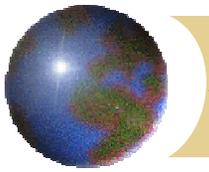
## *Options Considered*

- ⊕ Global Filesystem Unnecessary
- ⊕ Alternative GFS
- ⊕ Freeze AFS
- ⊕ IBM AFS support
- ⊕ 3<sup>rd</sup> Party AFS support
- ⊕ Community AFS support

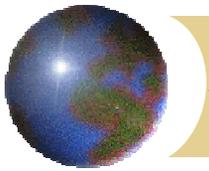


## *Recommendation: two-pronged approach*

- ✿ Secure the present
  - ▣ Move to OpenAFS
  - ▣ 3<sup>rd</sup> party support or local support
- ✿ Address needs of the future
  - ▣ Address AFS deficiencies
  - ▣ Look at alternatives



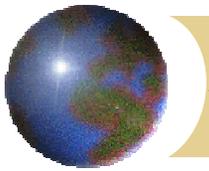
*Current Status: Is anyone using AFS?*



## *Current Usage*

- ⊕ Control Room Logbook
- ⊕ Log archive for k5 and w2k logs
- ⊕ UNIX/Windows/Mac file-sharing
- ⊕ UPS product and patch distribution

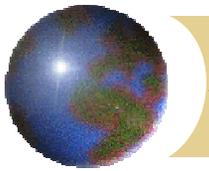
	6/2002	3/2003
Products	245	283 (+16%)



## *Current Usage*

- Central web servers keep pages in AFS
- Allows web authors to edit content directly w/o accounts on the central web servers

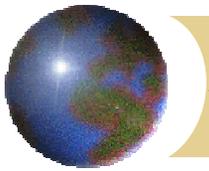
	6/2002	3/2003
Main WebServer	1	1
Departmental/ Experiment	54	58 (+7%)
Pubs	1	2 (+100%)



# *Off-Site Usage*

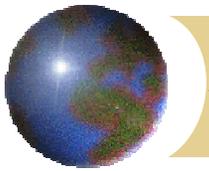
## ☉ Top users include:

- ☒ Obvious HEP labs: LANL, CERN, SLAC, Rutherford, DESY, INFN, BNL, IN2P3, ANL, LLNL
- ☒ Large US universities: Iowa, Caltech, William & Mary, UCSD, UW, Michigan, CMU, Notre Dame, Stanford, Indiana, NC, MIT
- ☒ Home (?) users: AT&T, WideOpenWest, Ameritech, EarthLink
- ☒ Many European universities: UK, SE, DE, CH, SL, IT, CZ
- ☒ Some non-European universities: MX, CN, BR



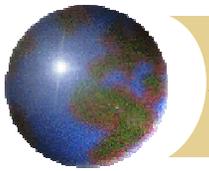
# Statistics

	6/2002	3/2003
AFS Users	3427	3773 (+10%)
AFS Groups	640	696 (+9%)
Experiments/ Projects	33	33
Clusters	9	12* (+33%)



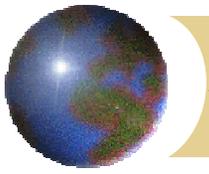
# Statistics

	6/2002	3/2003
Servers	11	11* (Many infrastructure changes)
Storage	2.6 TB RAID5 (1.6TB on SAN)	4.2TB RAID 5 (all on SAN)
Network	100Mb/s Switched	100Mb/s Switched

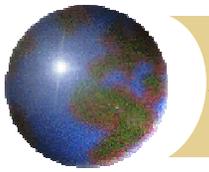


# *Statistics*

	6/2002	3/2003
Disk in Use	1.2TB	1.5TB (+25%)
Volumes Served	11766	12564 (+7%)

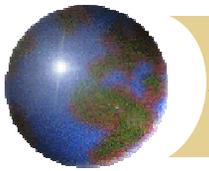


# *Review*



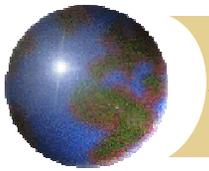
## *Options Considered*

- ⊕ Global Filesystem Unnecessary
- ⊕ Alternative GFS
- ⊕ Freeze AFS
- ⊕ IBM AFS support
- ⊕ 3<sup>rd</sup> Party AFS support
- ⊕ Community AFS support



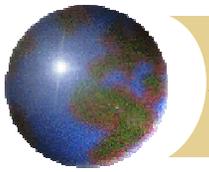
# *Global Filesystem Unnecessary*

- Increase of effort lab wide
  - Purchasing of infrastructure on a per cluster basis (duplication)
  - May require retraining of staff and users
  - Increased management
- No cross-platform file sharing
- AFS client access still needed by experiments



## *AFS Alternatives*

- Migration will be long & painful
  - Reorganizing code/web areas/products to fit new namespace
  - Security migration of users, groups and ACLs
  - Retraining of staff and users
  - Best if coordinated with other HEP sites



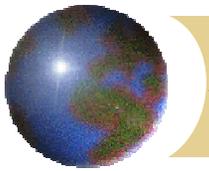
## *AFS Alternatives (cont)*

### ⊕ NFS

- ⊞ The only other well known FS in the lab.
- ⊞ No security yet.
- ⊞ Does not scale well as clients increase.
- ⊞ Unreliable WAN performance.
- ⊞ Can be as expensive as AFS is now.

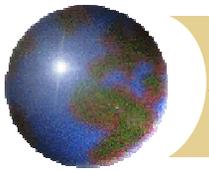
### ⊕ WebDAV, WebFS, GFS, Ibrix, NFSv4

- ⊞ Requires research to determine feasibility



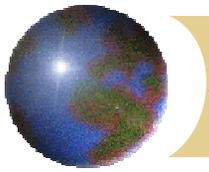
## *Freeze IBM AFS*

- ❖ Servers stuck at frozen OS
- ❖ Potential interoperability problems as other sites move away from IBM AFS
- ❖ No updates for security problems
- ❖ Does not address shortcomings of AFS



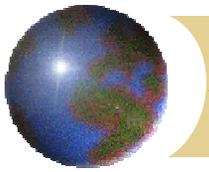
## *IBM Support of IBM AFS*

- ⊕ Cost is \$399K, an increase of almost 2000%.
- ⊕ Bug-fixes only. No enhancements.
- ⊕ Does not address shortcomings of AFS



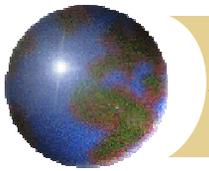
## *3<sup>rd</sup> Party Support of OpenAFS*

- Commercial quote of ~\$33.6K/year.
  - Solaris 8 support only. +9.6K/year for each additional OS and version.
  - Bug fixes only. No enhancements.
  - 10 hours/month/ticket.
  - 24x7 phone and e-mail support.
  - 5 points of contact.
- Could address shortcomings of AFS.



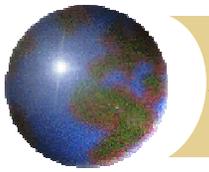
## *Community Support of OpenAFS*

- ✚ Allows for both bug-fixes and enhancements.
- ✚ Active participation in code development helps ensure our concerns are addressed and changes rolled back into main source tree.
- ✚ Can address shortcomings of AFS.



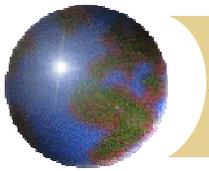
# *Community Support of OpenAFS*

- Two local developers
  - Based on recommendation from CERN and DESY
- Collaborative effort with other HEP sites
  - Ensure compatibility
  - CERN, DESY, SLAC have moved to OpenAFS
- Leverage OpenAFS community effort
  - K5 authentication (now available)
  - >8GB volume, >2GB file size



## *Points we want to stress...*

- Global access is important
  - Both for on-site and off-site
- AFS alternatives – requires further research
- AFS is needed, but requires improvements to ensure longevity



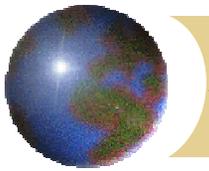
## *Points we want to stress...*

### ✚ Move to OpenAFS

- ▣ Consistent strategy with other HEP labs

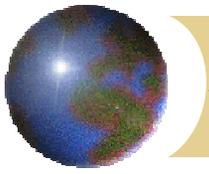
### ✚ Support of OpenAFS

- ▣ Two-person local support or
- ▣ 3<sup>rd</sup> party support
- ▣ DESY has a desire to pool AFS support

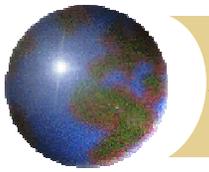


## *Recommendation: two-pronged approach*

- ✿ Secure the present
  - ▣ Move to OpenAFS
  - ▣ 3<sup>rd</sup> party support or local support
- ✿ Address needs of the future
  - ▣ Address AFS deficiencies
  - ▣ Look at alternatives

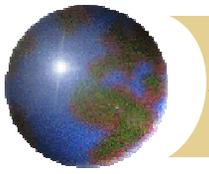


*End*

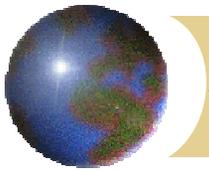


# *Reference Information*

- ⊕ History
- ⊕ Comments from Experiments (2002)



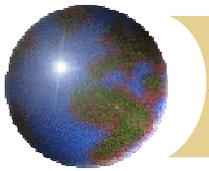
# *History*



# *AFS History at FERMI*

## Original Mission

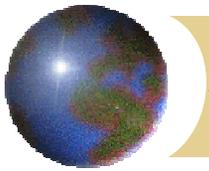
- ❑ Pilot program started on FNALU (the new central Unix cluster)
- ❑ Used for login areas, data areas, products and experiment project areas
- ❑ Expected to eventually deploy servers at local work group sites
- ❑ Many multi-flavored “smaller” systems
- ❑ FNALU to be restricted to physics applications
- ❑ AFS seen as a method to share physics data site-wide and to other HEP labs



# *The Last Few Years*

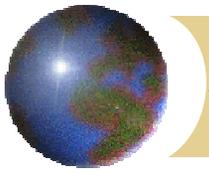
## ✚ Usage increase

- ✚ Introduction of Linux a large factor
  - Linux kernel/AFS version compatibility increasingly time consuming
- ✚ Introduction of a stable Windows client
  - provides a method for file sharing between Windows and UNIX
- ✚ Many more central web servers hosted from AFS
- ✚ Scientific users finally appreciate a single login environment and password

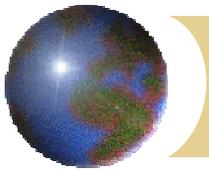


## *The Last Few Years*

- ❖ Management manpower decreases
  - ❑ Still able to keep up with hardware improvements
  - ❑ Much needed management software still not available
- ❖ Transarc purchased by IBM
  - ❑ New feature additions stopped
  - ❑ Writing on the wall – AFS will be dropped
  - ❑ Move to “Open” the software



# *Comments from Experimenters*



# *Comments from Experiments*

## ⊕ CMS (*Hans Wenzel*)

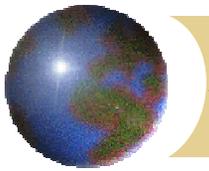
### ⊞ Home areas

- Currently CMS has ~ 100 users with accounts on CMSUN1.

## ⊕ CMS

### ⊞ Code distribution

- CMS software resides in AFS. All local CMS desktops have AFS and access CMS software this way.
- AFS ACL's are e.g. used to regulate access to licensed software (e.g. Objectivity).

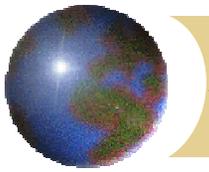


# *Comments from Experiments*

## ⊕ CMS

### ⊞ Code distribution

- Present usage status:
  - Total usage: 32GB
  - Total quota: 80GB
- 25% growth since 25 Apr 2002

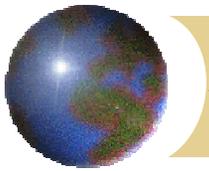


# *Comments from Experiments*

## ● CMS

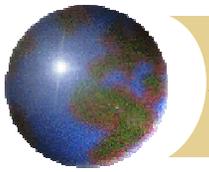
### ■ Code distribution

- About 70 requests for access to CMS software on AFS.
- University of Iowa, for e.g., successfully use AFS to access CMS software.
- Different solutions for code distribution possible, so a global namespace via WAN is not a necessity.



# *Comments from Experiments*

- **Minos (*Liz Buckley-Geer*)**
  - Build Offline code
  - Web space on central web server
  - Data Distribution



# *Comments from Experiments*

- *Auger (Lukas Nellen, Aaron Chou)*
  - Code Distribution
  - Data Distribution for Fermi users only