


# Status of the CERN Analysis Facility

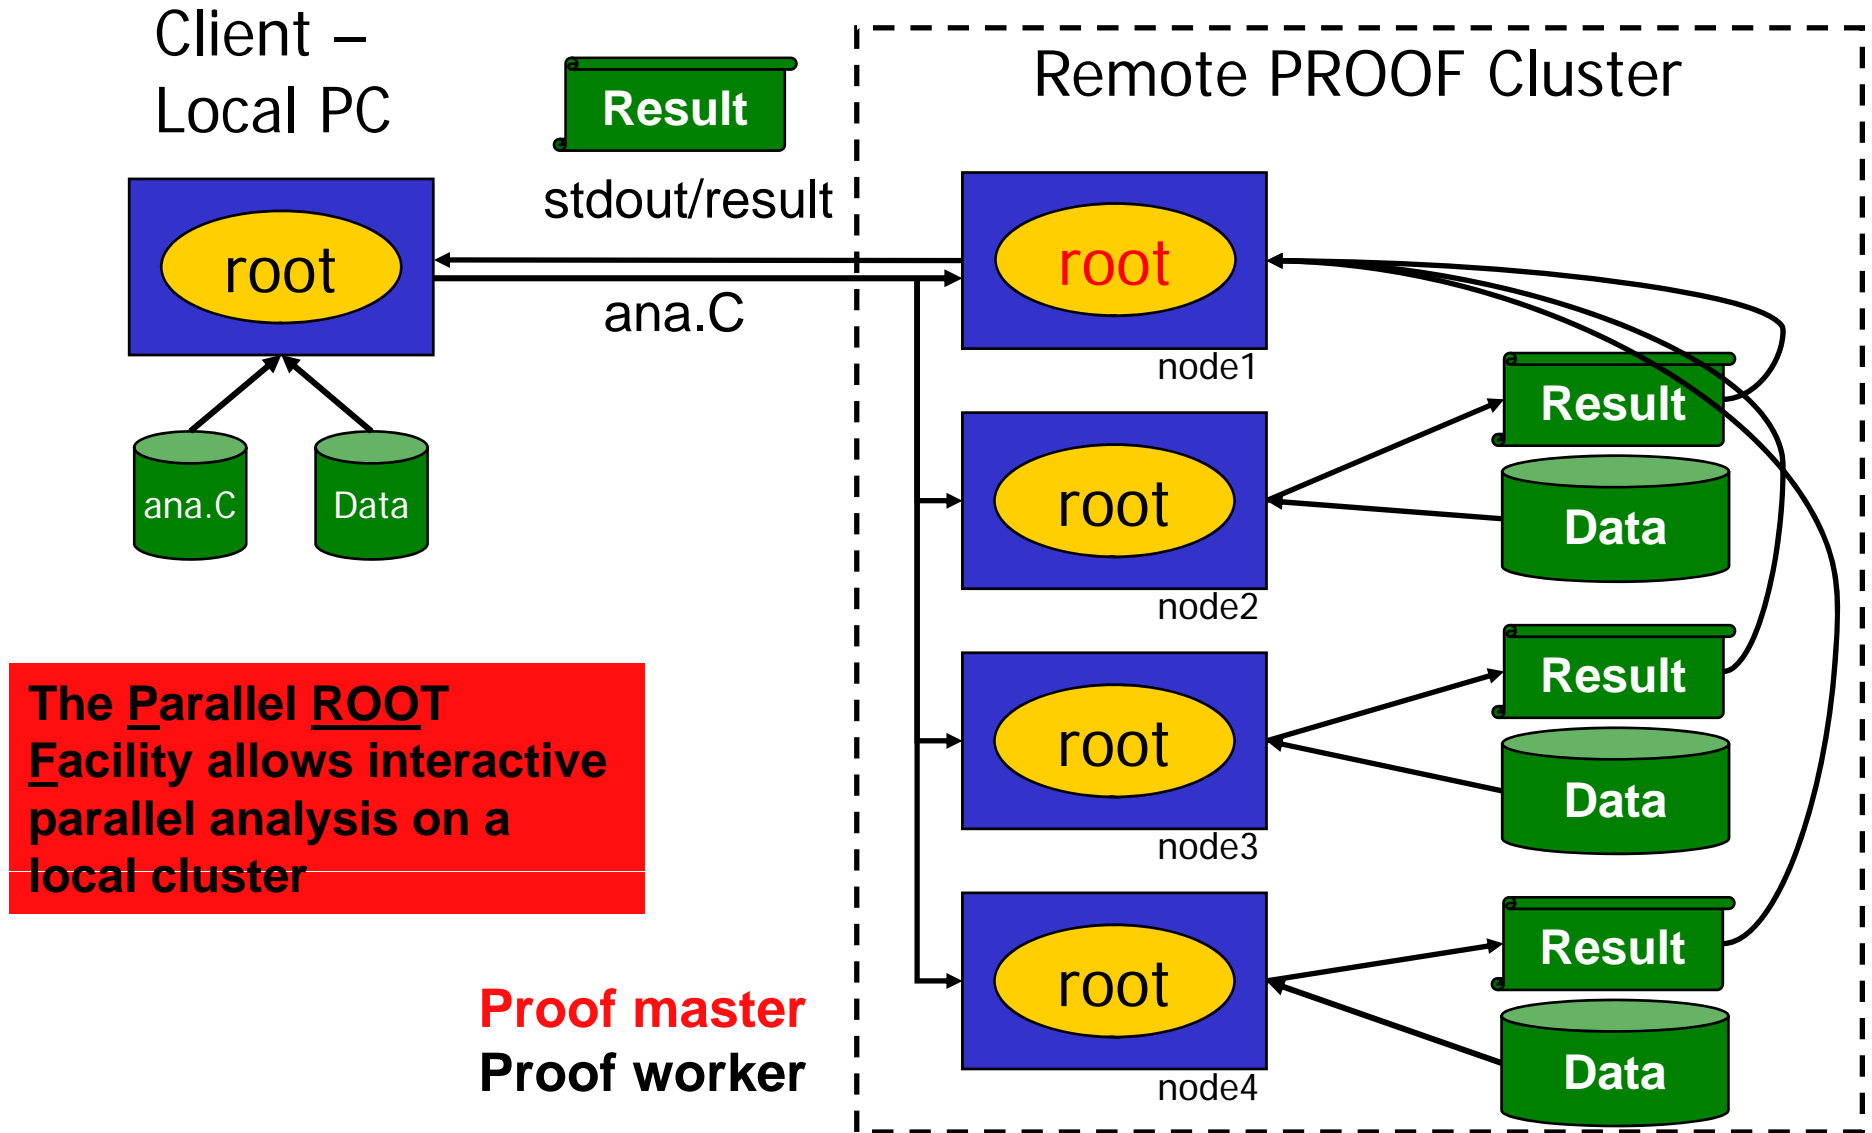
J. F. Grosse-Oetringhaus,  
ALICE Offline Week, October 2007



# Content



- CAF / PROOF
- Status
  - Current version
  - Available data
  - Problems
- Current Development
  - Datasets and automatic staging from AliEn
  - Groups, disk quotas, CPU fairshare
- CAF workshop announcement

# PROOF Schema







A cluster called CERN Analysis Facility (CAF) runs PROOF for

- Prompt analysis of pp data
  - Pilot analysis of PbPb data
  - Fast simulation/reconstruction
- Calibration & Alignment


### CAF computing cluster




# Current ROOT version



- ROOT version from 30th July
- Since this week under test: new version with SVN tag 20268, probably will get default next week




# Available data




- PDC06
  - 14 TeV pp minbias: 3 million events
  - 0.9 TeV pp minbias: 200k events
  - AliRoot v4-04-Release available
- PDC07
  - 0.9 TeV pp minbias: 3 million events staged
  - AliRoot v4-05-Release not there yet, but ESD.par can be used
- See  
<http://aliceinfo.cern.ch/Offline/Analysis/CAF/index.html#data>

# Recent Problems



- Full cluster failure around 20th Sep (single incident)
  - Reason still not fully understood
  - "Hard" reset required by IT
- User problems
  - Memory exceeded (especially when filling trees)
    - Could affect other users and system processes (failing new, malloc calls)
    - New: Hard limit of 2 GB per process
    - Under development
      - Memory tracking directly from PROOF progress window
      - Creation of files to deal with large output on workers (in new CAF version)
  - Running over long chains failed (~ 20.000 files)
    - Fixed by Gerri Ganis

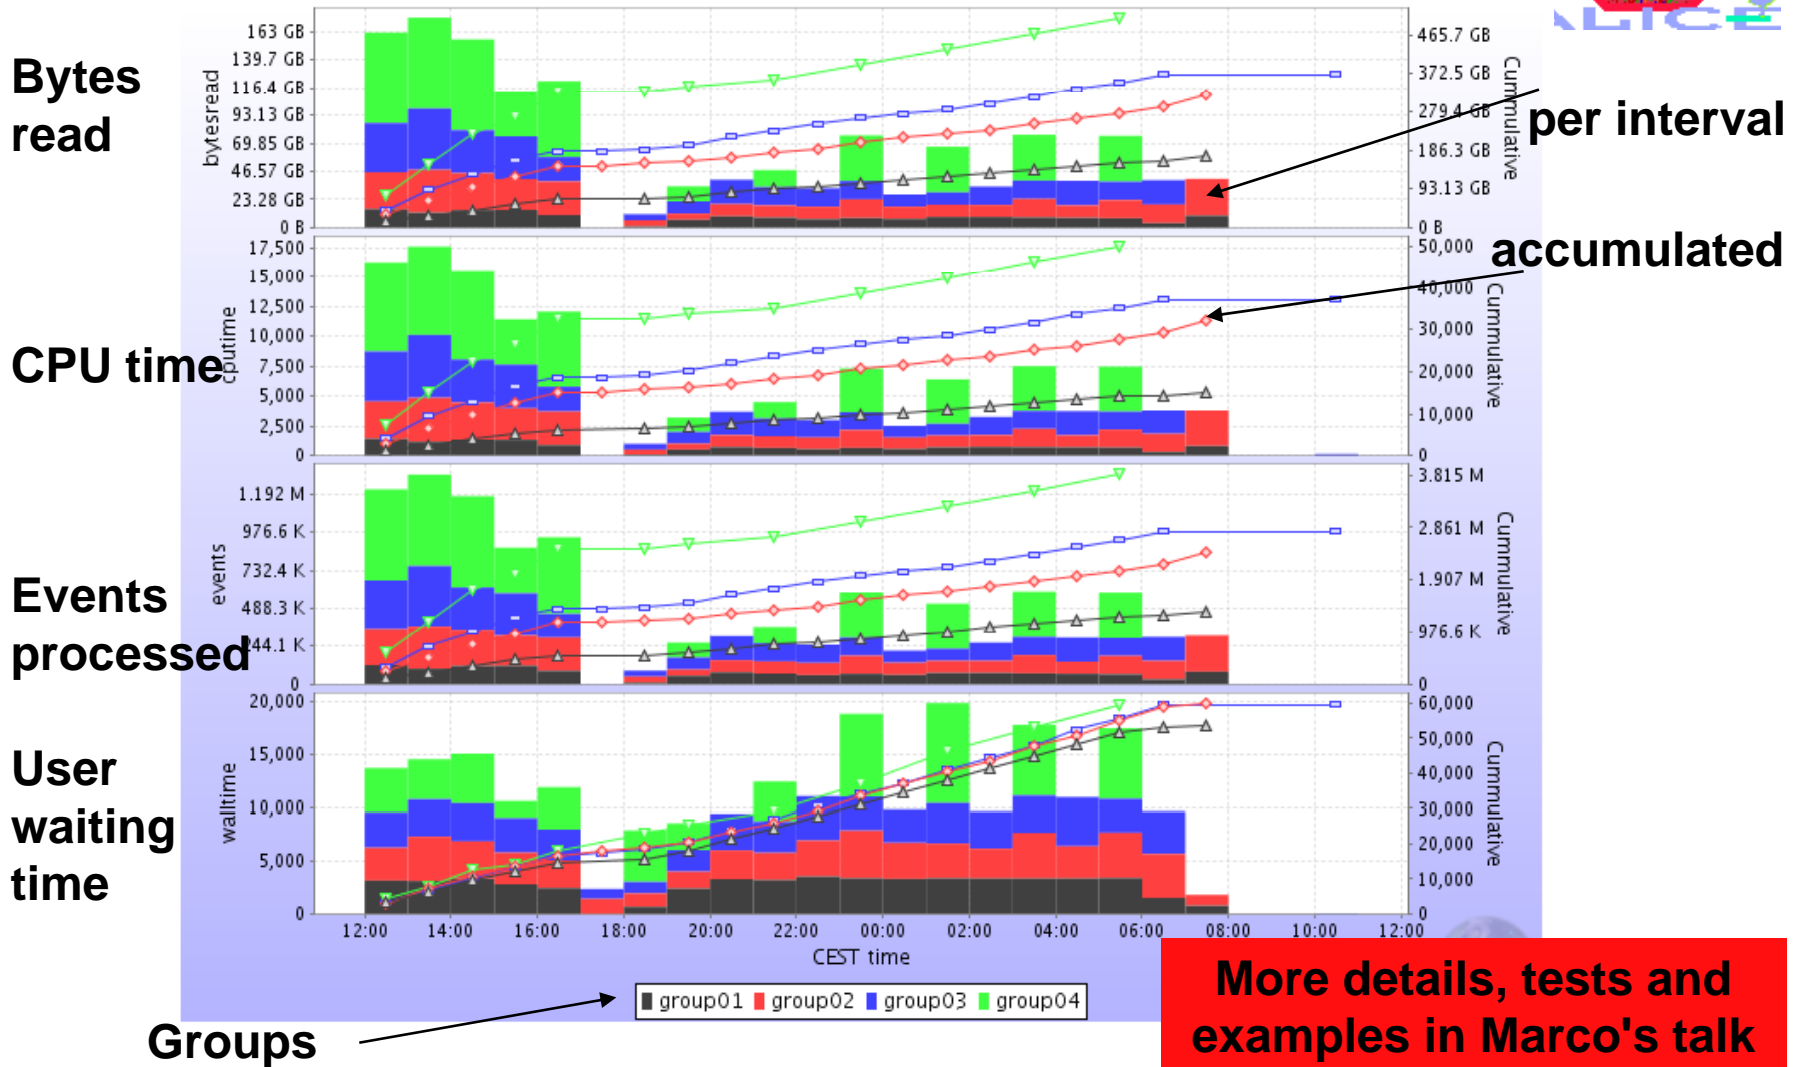

# Groups, Quotas, Fairshare




- ALICE wants to assure a fair use of the CAF resources
- Users are grouped
  - E.g. sub detectors or physics working groups
  - Users can be in several groups
- Each group has a quota on the disk
  - This space can be used for datasets staged from AliEn (see next slides)
- Each group has a CPU fairshare target
  - This governs the priority of concurrent queries
  - More details in Gerri's and Marco's talks



# Monitoring of cluster usage



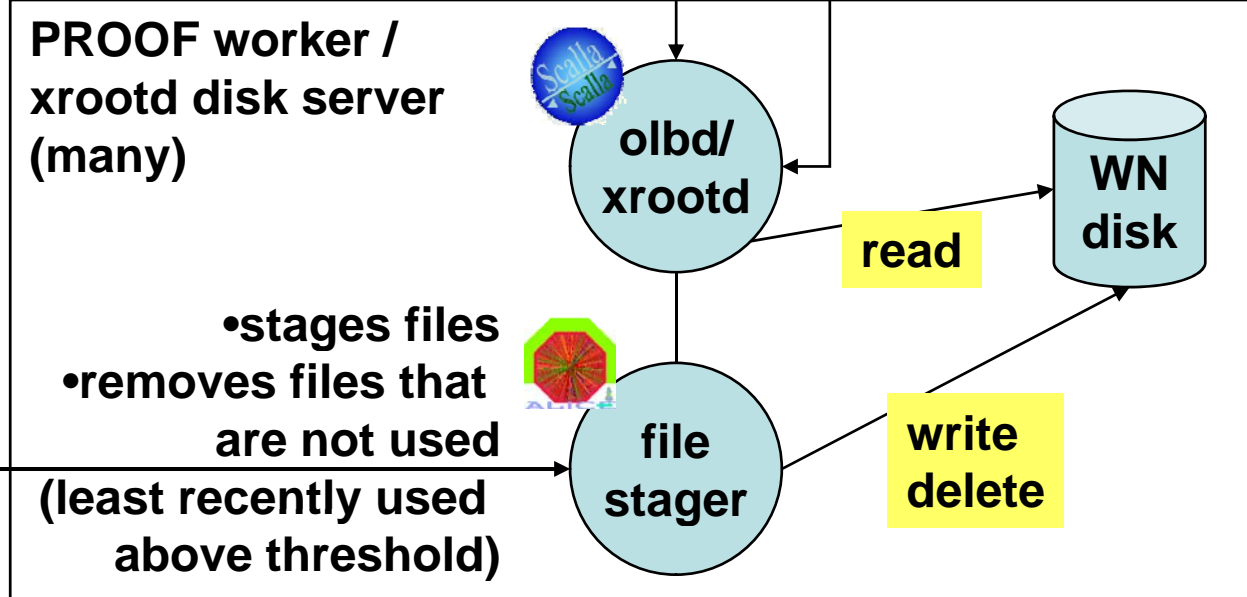
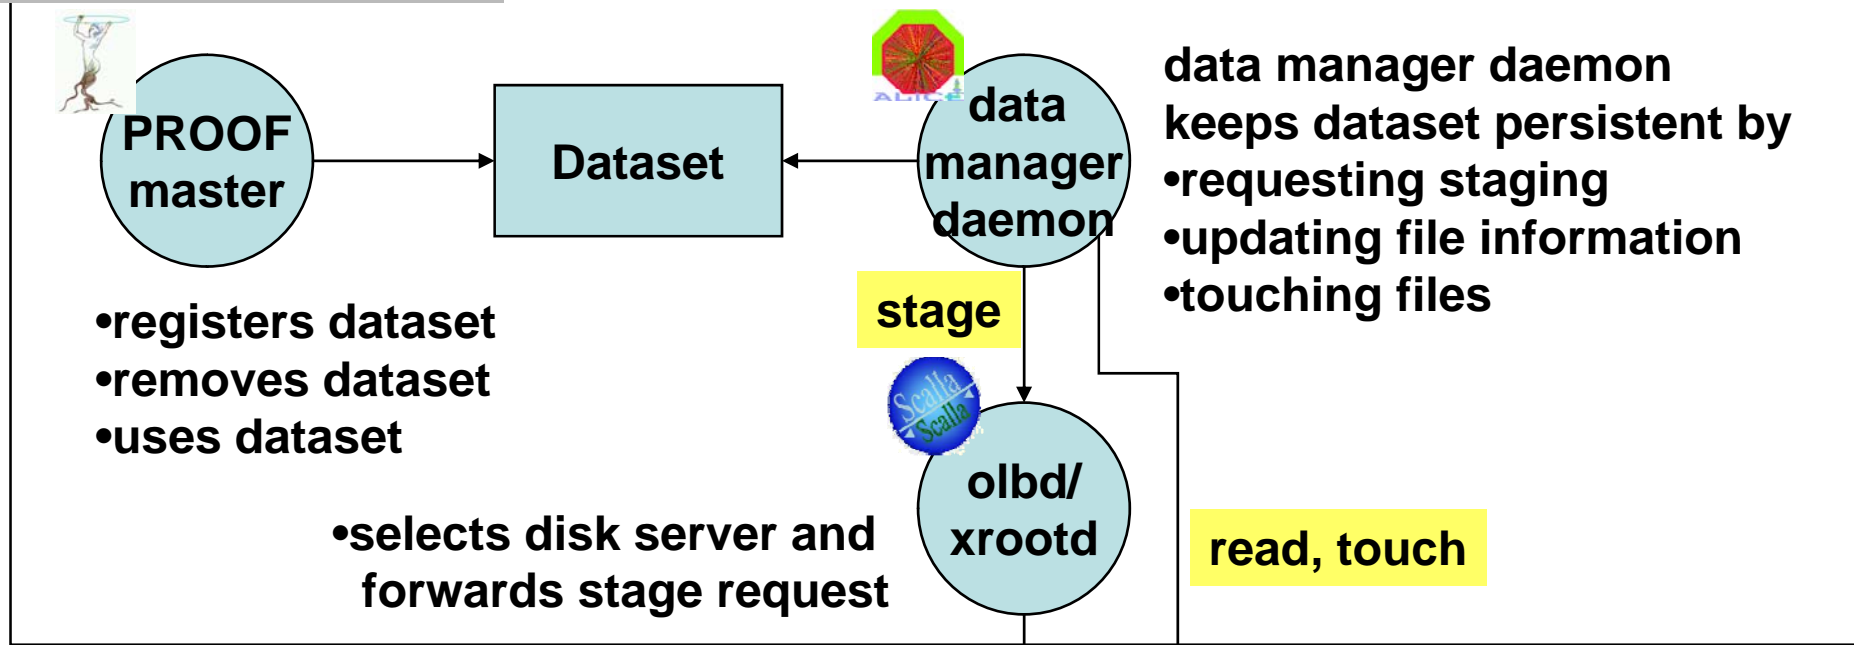
# Easing the Staging of Data – PROOF Datasets



- A dataset represents a list of files (e.g. physics run X)
  - Correspondence between AliEn collection and PROOF dataset
- Users register datasets
  - The files contained in a dataset are automatically staged from AliEn (and kept available)
  - Datasets are used for processing with PROOF
    - Contain all relevant information to start processing (location of files, abstract description of content of files)
- File-level storing by underlying xrootd infrastructure
- Datasets are public for reading (you can use datasets from anybody!)
- There are common datasets (for data of common interest)

# Dataset concept

## PROOF master / xrootd redirector



...

# Datasets in Practice

- Create DS from AliEn collection
  - TFileCollection\* ds = TGridResult::GetFileCollection()
- Upload to PROOF cluster
  - gProof->RegisterDataSet("myDS", ds)
- Check status: gProof->ShowDataSet("myDS")

```
root [9] gProof->ShowDataSet("ESD5000")
Info in <TXProofServ::GetDataSet> on master-0: uri=ESD5000

TFileCollection ESD5000 - title contains: 21899 files with a size
of 1177346326985 bytes, 100.0 % staged
The files contain the following trees:
Tree /esdTree: 2189800 events
Tree /HLTesdTree: 2189800 events
```

- Use it: gProof->Process("myDS", "mySelector.cxx+")  
(not yet implemented)

# Dataset in Practice (2)

- List available datasets: `gProof->ShowDataSets()`

```
root [3] gProof->ShowDataSets()
Dataset URI                                     |# Files|Default tree|# Events|  Disk  | Staged
/default/jgrosseo/ESD100                       |  6764|/esdTree   |  676400|  343 GB| 100 %
/default/jgrosseo/run82XX_part1                 | 10000|/esdTree   |  998900|  288 GB|  99 %
/default/jgrosseo/run82XX_part2                 | 10000|/esdTree   |  944700|  272 GB|  94 %
/default/jgrosseo/run82XX_part3                 | 10000|/esdTree   |  987900|  285 GB|  98 %
/default/jgrosseo/ESD600                       |  1844|/esdTree   |  184400|   51 GB| 100 %
/default/jgrosseo/ESD_FullMisalignment          |   944|/esdTree   |   92100|   47 GB|  97 %
/default/jgrosseo/run12000                      |    62|/esdTree   |    49  |    4 GB|  79 %
/default/jgrosseo/ESD5000                      | 21899|/esdTree   | 2189800| 1096 GB|  99 %
/default/jgrosseo/ProofSessionFiles             | 25438|/esdTree   | 4460900| 1640 GB| 100 %
```


- You always see common datasets and datasets of your group
- This method was used to stage 30 M events of PDC07
- Demo today in Gerhard's talk

# Monitoring of datasets

Number of files per host

15. lxb6055.cern.ch	2797	133.2 GB
16. lxb6056.cern.ch	2795	133.2 GB
17. lxb6057.cern.ch	2785	131.9 GB
18. lxb6058.cern.ch	2772	129.1 GB
19. lxb6059.cern.ch	1704	64.2 GB
20. lxb6060.cern.ch	1917	82.55 GB
21. lxb6061.cern.ch	2766	130.3 GB
22. lxb6062.cern.ch	1228	40.67 GB
23. lxb6063.cern.ch	2789	131.9 GB
24. lxb6064.cern.ch	2777	131.3 GB
25. lxb6065.cern.ch	2755	128.2 GB
26. lxb6066.cern.ch	2364	118.2 GB
27. lxb6067.cern.ch	2745	127.9 GB
28. lxb6068.cern.ch	2740	128.2 GB
29. lxb6069.cern.ch	2778	129.6 GB
30. lxb6070.cern.ch	2749	128.9 GB
31. lxb6071.cern.ch	2756	130.4 GB
32. lxb6072.cern.ch	2699	124.8 GB
33. lxb6073.cern.ch	2741	128 GB
34. lxb6074.cern.ch	2737	127.2 GB
35. lxb6075.cern.ch	2755	128.9 GB
36. lxb6076.cern.ch	2741	130.6 GB
37. lxb6077.cern.ch	2731	125.8 GB
38. lxb6078.cern.ch	2741	128.5 GB
39. lxb6079.cern.ch	2739	126.8 GB
40. lxb6080.cern.ch	601	16.8 GB
<b>Total</b>	<b>96964</b>	<b>4.431 TB</b>
<b>Average</b>	<b>2620</b>	<b>122.6 GB</b>


## Data set usage per group



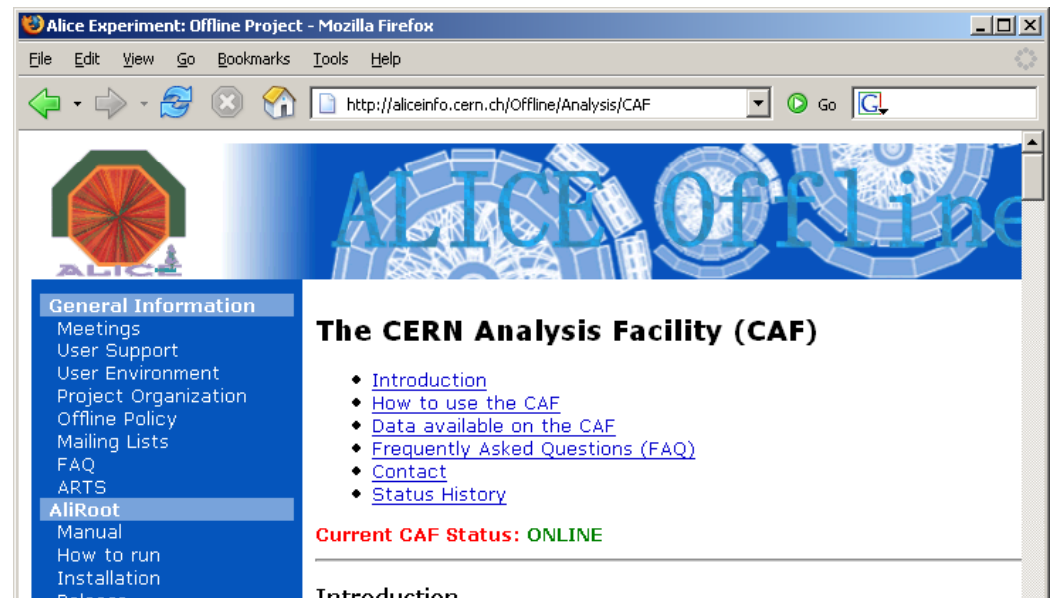
# Using CAF @ Ixplus

- CAF can be accessed by all ROOT production releases
- ROOT/AlRoot version compatible with CAF for Ixplus users is on AFS
  - source /afs/cern.ch/alice/caf/caf-ixplus.sh (bash shell)
  - NB: Built for Ixplus architecture, gnu/linux distribution
  - Soon a parameter is needed, to indicate the required AliRoot version
  - Is kept up to date, if CAF version is changed

# Information




- <http://aliceinfo.cern.ch/Offline/Analysis/CAF>
- Monthly tutorials at CERN, contact Yves Schutz
- Slides of the last tutorial (AliRoot, PROOF, AliEn)  
<http://aliceinfo.cern.ch/Offline/Analysis/Tutorial/>
- Server installation at  
<http://root.cern.ch/twiki/bin/view/ROOT/ProofInstallation>






# CAF workshop



- We will have a CAF workshop between November 28-30 (probably 1 day)
  - Please mark in your calendars
    - Feedback from users
    - Feature requests for PROOF and CAF
    - Discussion about staging to CAF
    - The analysis framework on CAF
- Not a tutorial

# Summary & Plans



- CAF up and running with minor problems
- Several new developments coming up
  - Users need to be grouped. We will soon send a call to assign yourself to a group
  - Datasets functionality to be made public in the next weeks
    - You will be able to decide yourself what is staged to the system
- There will be a CAF workshop at the end of November (around 28-30). Please mark in your calendars

# Backup

