"Introduction to the FNAL UAF "

Hans Wenzel Fermilab

- Introduction. Why would you want to use the UAF?
- What are the different components of the computing system at FNAL?
- How do you get an account?
- General setup and how do you set up the CMS environment?

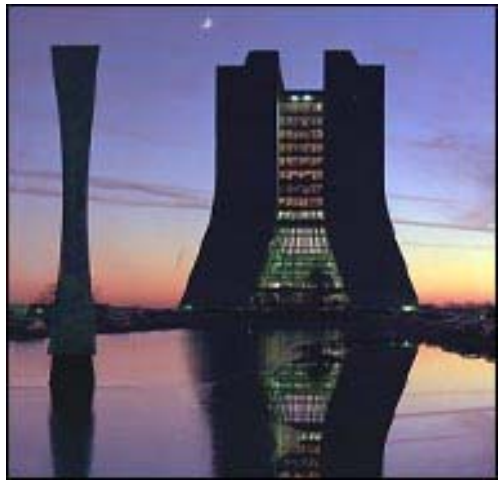

Why do you want to use the UAF?

- Reference Platform.
- interactive system.
- CMS software is available and maintained.
- System administrators keep the system up. (be aware Thursday morning is the slot for system maintenance and upgrades.)
- Fast access to mass storage, plenty of disk space.
- Fast networking (GigaBit Ethernet).
- Plenty of CPU to do some production projects.
- Condor Batch system to automate this projects
- Data samples.
- Basically provide all the things you “don’t” want to do or worry about yourself!

What is the UAF?

- The User Analysis Farm (UAF) is part of the Tier 1 facility at Fermilab. The UAF is a farm running the LINUX operating system providing interactive computing for users.
- Separate interactive and batch computing:
  - Pool of interactive nodes (CMSUAF), load balancing is done on the networking switch, tested solution support Kerberos, afs ... Comes at no extra cost since the switches we are using are already equipped with the option.
  - Big batch farm (shared with production,grid) running condor.

Configuration of User Computing as of November 2005:

Enstore: Powderhorn STK Silo:Drives: 9 x 940B, speed 30 MB/s,

Capacity: 200 GB/cartridge

ENSTORE

dCache

dCache:

- **>25 TB read Pools (backed up to tape)**
- **>9 TB write Pools (backed up to tape, can also read from them)**
- **>13 TB resilient Pools (no tape backup, multiple copies, not deleted)**
- **/pnfs/cms/WAX**
- **/pnfs/cms/WAX/resilient**

CMSUAF

cmswn

cmswn

cmswn

cmswn

cmswn

cmswn

Condor High Throughput Computing

Gigabit Ethernet

IBRIX

/uscms/home/<username>
/uscms_data/d1/<username>
/uscms_scratch/<physicsgrp.>

> 240 registered users

kinit –n <username>@FNAL.GOV
ssh <username>@ cmsuaf.fnal.gov

Hardware and people behind the UAF

Besides doing many other things keeping Hardware and OS alive are:

- **Jon Bakken (not shown)**
- **Lisa Giacchetti**
- **Joe Kaiser**
- **Gary Stiehr**
- **Tim Messer (not shown)**
- **Dave Fagan**

racks with computers

- **CMS Software**: Natalia Ratnikova is our code librarian. She is doing a terrific job installing the software and making sure that it works.
- LPC Helpdesk: Patrick Gartung knows it all

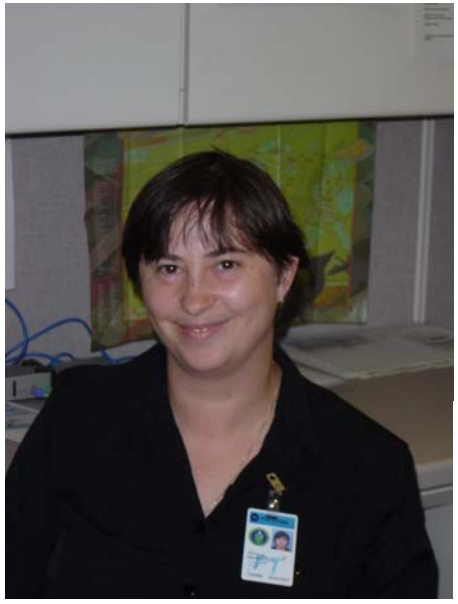

US CMS

The Compact Muon Solenoid

The UAF Web sites

http://www.uscms.org/

Then on the left select: Software & Computing

Then: User Computing

Getting help: or contact LPC helpdesk (aka Patrick) Ext. 3832

http://csdserver1.fnal.gov/HelpDesk/cd/

or send mail to: helpdesk@fnal.gov

User Computing - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://www.uscms.org/SoftwareComputing/UserComputing/UserComputing.html

Getting Started | Latest Headlines | CVS | ROOT | Personal | CMS | Fermilab | CERN | Red Hat | Search | C++ | My Projects

USCMS
THE COMPACT MUON SOLENOID

Home | Overview | News | Links & Resources | Contact | Search

Software & Computing
User Computing

News

Scheduled Downtimes

UAF Systems Status:
FBSNG Status
Ganglia Status
Enstore Status
dCache Status

Get an Account
Introduction
Environment Setup
Connecting to the UAF

Tutorials:
Jetmet
Muon
New

Disk Space:
IBRIX File System
Request Disk Space
AFS Home Areas

Mass Storage
Batch System
Data Sets
Data Transfer
HowTos
FAQ
Help

CMS Software
Grid
Project Management
Organization
Calendar

Introduction

The User Analysis Farm (UAF) is part of the Tier 1 facility at Fermilab. The UAF is a small farm running the LINUX operating system providing interactive and batch computing for users. Load balancing, resource and process management are realized with FBSNG , the batch system developed at Fermilab. The system consists of scripts on top of FBSNG which control the interactive login and provide a mechanism to obtain kerberos credentials and to securely forward x-window socket connections. We chose this solution for the following among other reasons:

- it requires no special hardware (e.g. switches).
- it requires no special kernels (like e.g. LVS or Mosix).
- no special software (besides ssh) has to be installed on the client side.
- the system can be extended easily by just adding new nodes.
- if resources are available in other CMS clusters we can easily temporarily add these resources to the cluster (dynamic partitioning) and release them as soon as they are needed again.
- kerberos is supported.
- easy to apply policies.

Low level monitoring of the hardware like CPU usage, memory usage etc. is provided by GANGLIA and for the UAF have a look here. To check the status of the batch system please have a look here. Below is a snapshot of what you will find on this web page in the Queues section.

Below you find a schematic view of the UAF and mass storage as of February 2005:

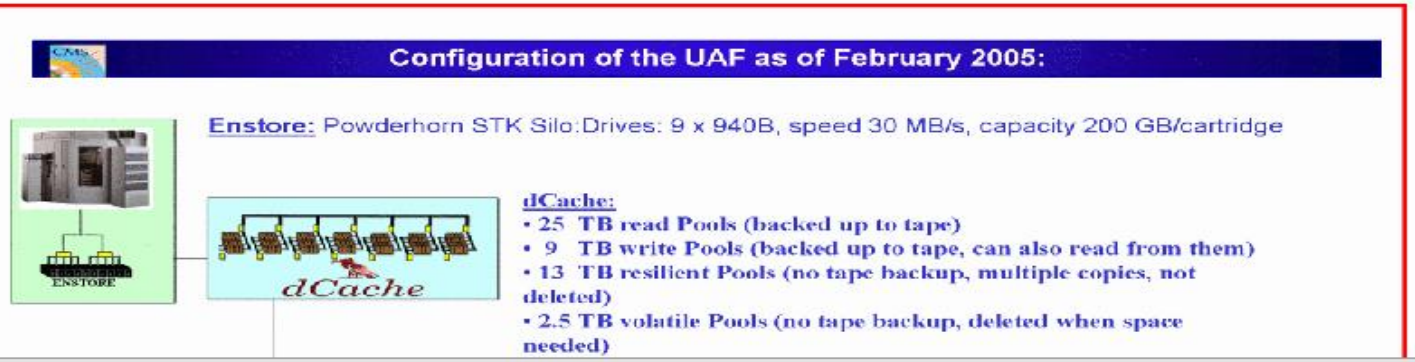

http://cmsuaf.fnal.gov:8080/fbswww

What do you find on the UAF web pages:

News
Scheduled Downtimes
(every thursday morning)
UAF Systems Status:
- *CONDOR Status*
- *Ganglia Status*
- *Enstore Status*
- *dCache Status*

Get an Account
Introduction
Environment Setup
Connecting to the UAF

Tutorials:
- *Jetmet*
- *Muon*
- *New*

Disk Space:
- *IBRIX File System*
- *Request Disk Space*
- *AFS Home Areas*

Mass Storage
Batch System
Data Sets
Data Transfer
HowTos
FAQ
Help

Obtaining a CMS account at FNAL

- Step 1: Get a Valid Fermilab ID

  (easy send fax or use the new web procedure for off/on site users we are working with the lab to make this as painless as possible.)

- Step 2: Get CMS account using the easy Webform.

To get in:

ssh cmsuaf.fnal.gov

(if you don't use kerberized ssh you will be presented with a kerberos challenge)

For more information see:

http://www.uscms.org/SoftwareComputing/UserComputing/ConnectUAF.html

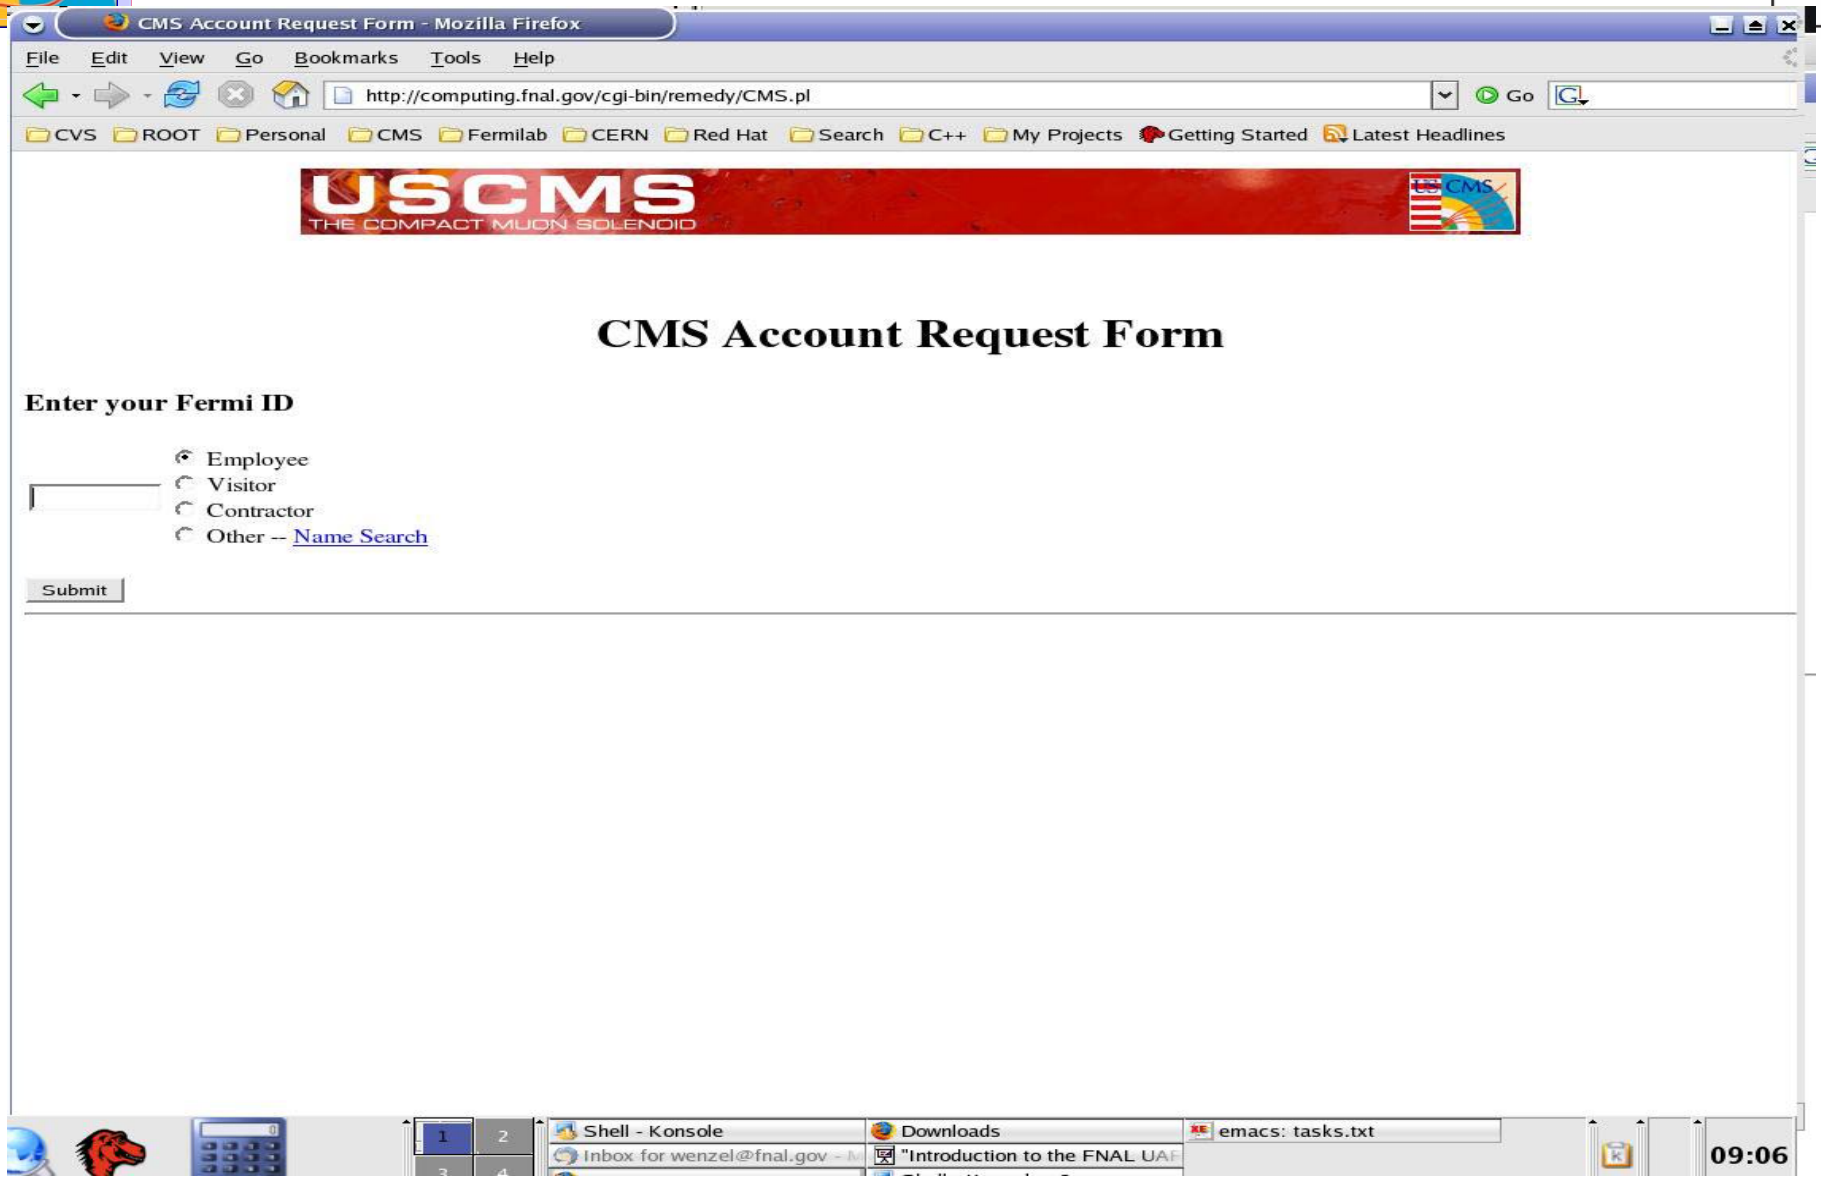
CMS Account Request Form - Mozilla Firefox
File Edit View Go Bookmarks Tools Help
http://computing.fnal.gov/cgi-bin/remedy/CMS.pl
Go
CVS ROOT Personal CMS Fermilab CERN Red Hat Search C++ My Projects Getting Started Latest Headlines
USCMS
THE COMPACT MUON SOLENOID
CMS Account Request Form
Enter your Fermi ID
Employee
Visitor
Contractor
Other -- Name Search
Submit
Shell - Konsole
Downloads
emacs: tasks.txt
Inbox for wenzel@fnal.gov
"Introduction to the FNAL UA
09:06

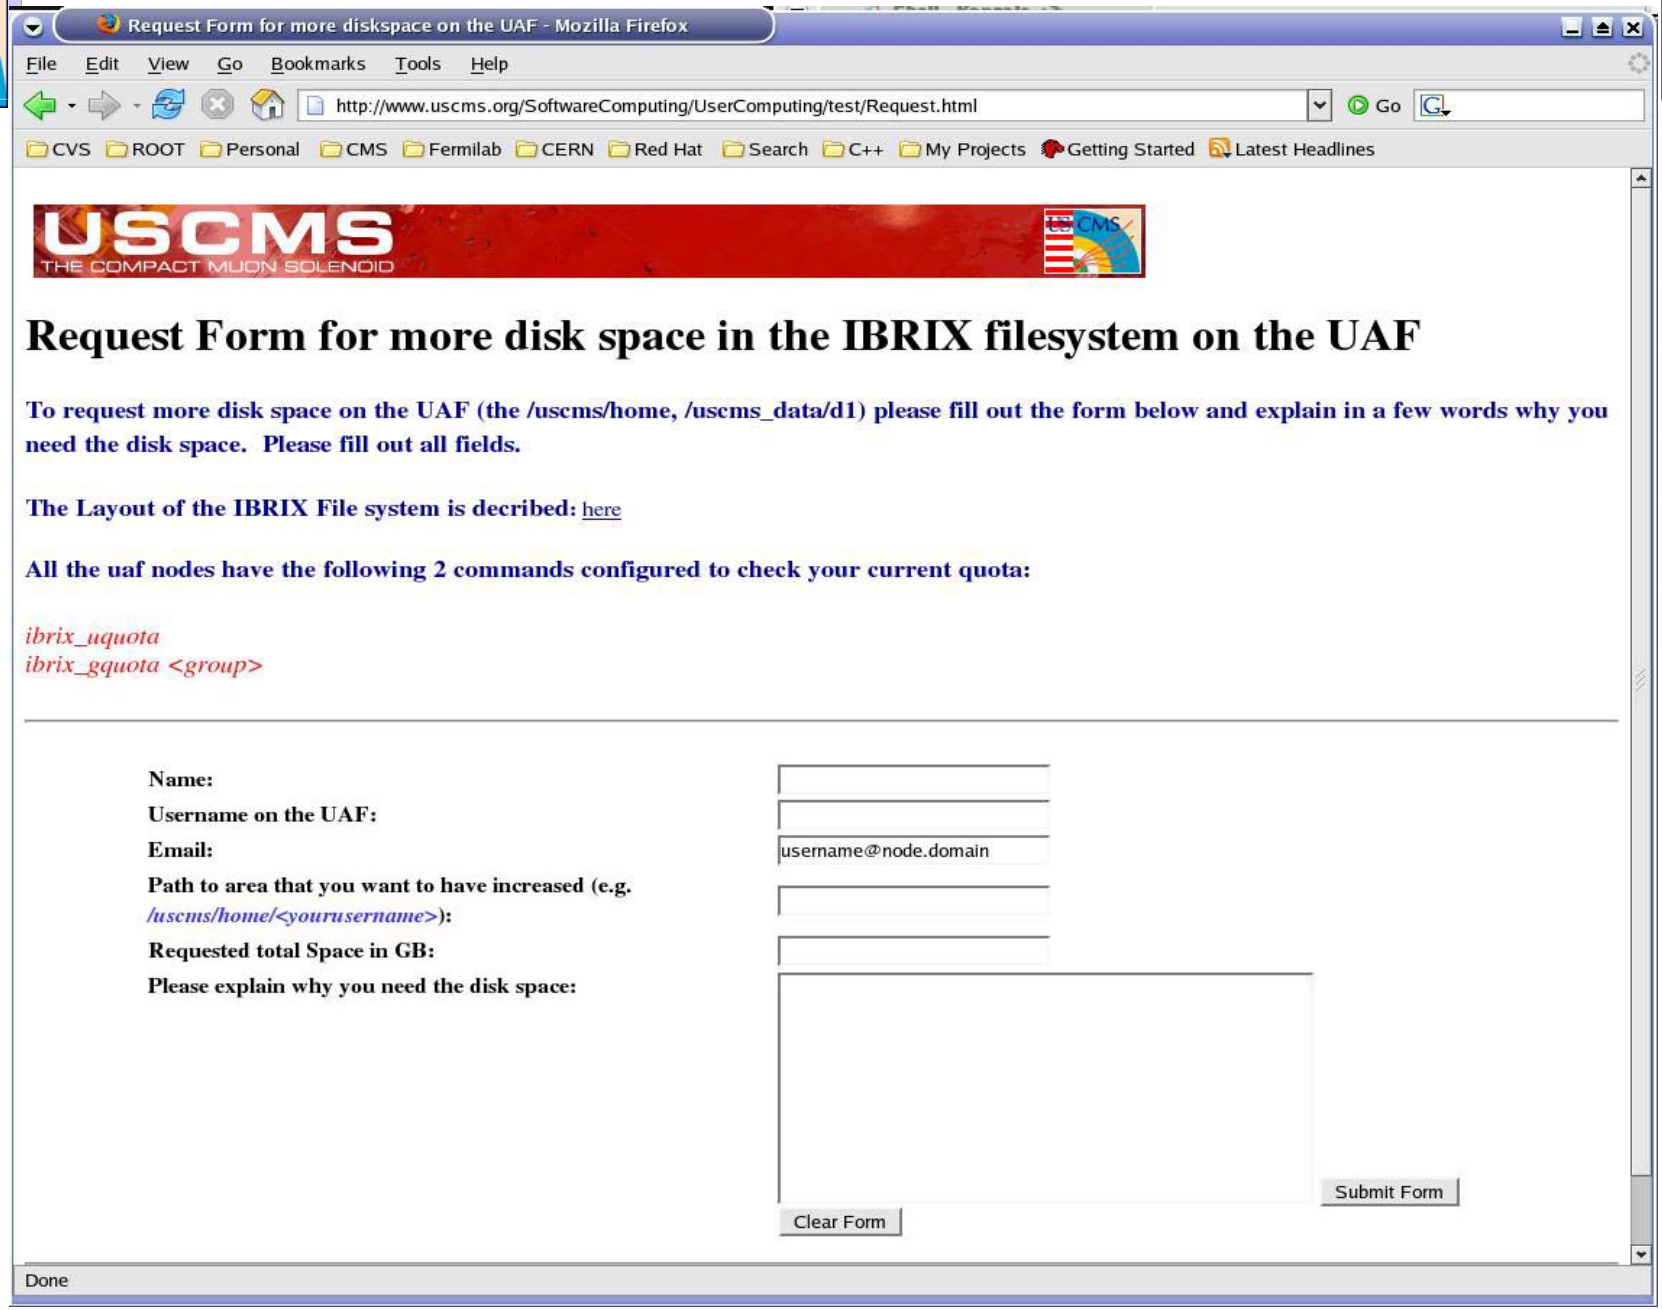
Request Form for more diskspace on the UAF - Mozilla Firefox
File Edit View Go Bookmarks Tools Help
http://www.uscms.org/SoftwareComputing/UserComputing/test/Request.html
Go
CVS ROOT Personal CMS Fermilab CERN Red Hat Search C++ My Projects Getting Started Latest Headlines
USCMS
THE COMPACT MUON SOLENOID
Request Form for more disk space in the IBRIX filesystem on the UAF
To request more disk space on the UAF (the /uscms/home, /uscms_data/d1) please fill out the form below and explain in a few words why you need the disk space. Please fill out all fields.
The Layout of the IBRIX File system is decribed: here
All the uaf nodes have the following 2 commands configured to check your current quota:
ibrix_uquota
ibrix_gquota <group>
Name:
Username on the UAF:
Email:
username@node.domain
Path to area that you want to have increased (e.g. /uscms/home/<yourusername>):
Requested total Space in GB:
Please explain why you need the disk space:
Submit Form
Clear Form
Done

General Setup of runtime environment:

In SL3 (Linux) tools like: dccp, kerberos, condor … are now part of the OS and there is no set up needed (no UPS/UPD). The .cshrc file should look like this:

```
#!/bin/sh
#
# - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
# Place any items that you want executed even for non-interactive use here
#skip if not interactive shell
if ( $?prompt ) then
   set noclobber          #prevent overwrite when redirecting output
   set ignoreeof          #prevent accidental logouts
endif
source /afs/fnal.gov/files/code/cms/setup/cshrc uaf
setenv SCRAM_ARCH `scramv1 arch`
```

setenv SCRAM_ARCH `scramv1 arch`
scram list

Tutorial:

- http://agenda.cern.ch/displayLevel.php?fid=584
- setenv CVSROOT lpccvs@cdcvs.fnal.gov:/cvs/lpc
  cvs co -r SL3_1_0 newtutorial
  cd newtutorial
- ./setup.pl -h

MCPS (Dave Evans, Marc Pollack)

- http://lynx.fnal.gov/runjob/MCPS_20ExRootAnalysis

What happens to the signal:

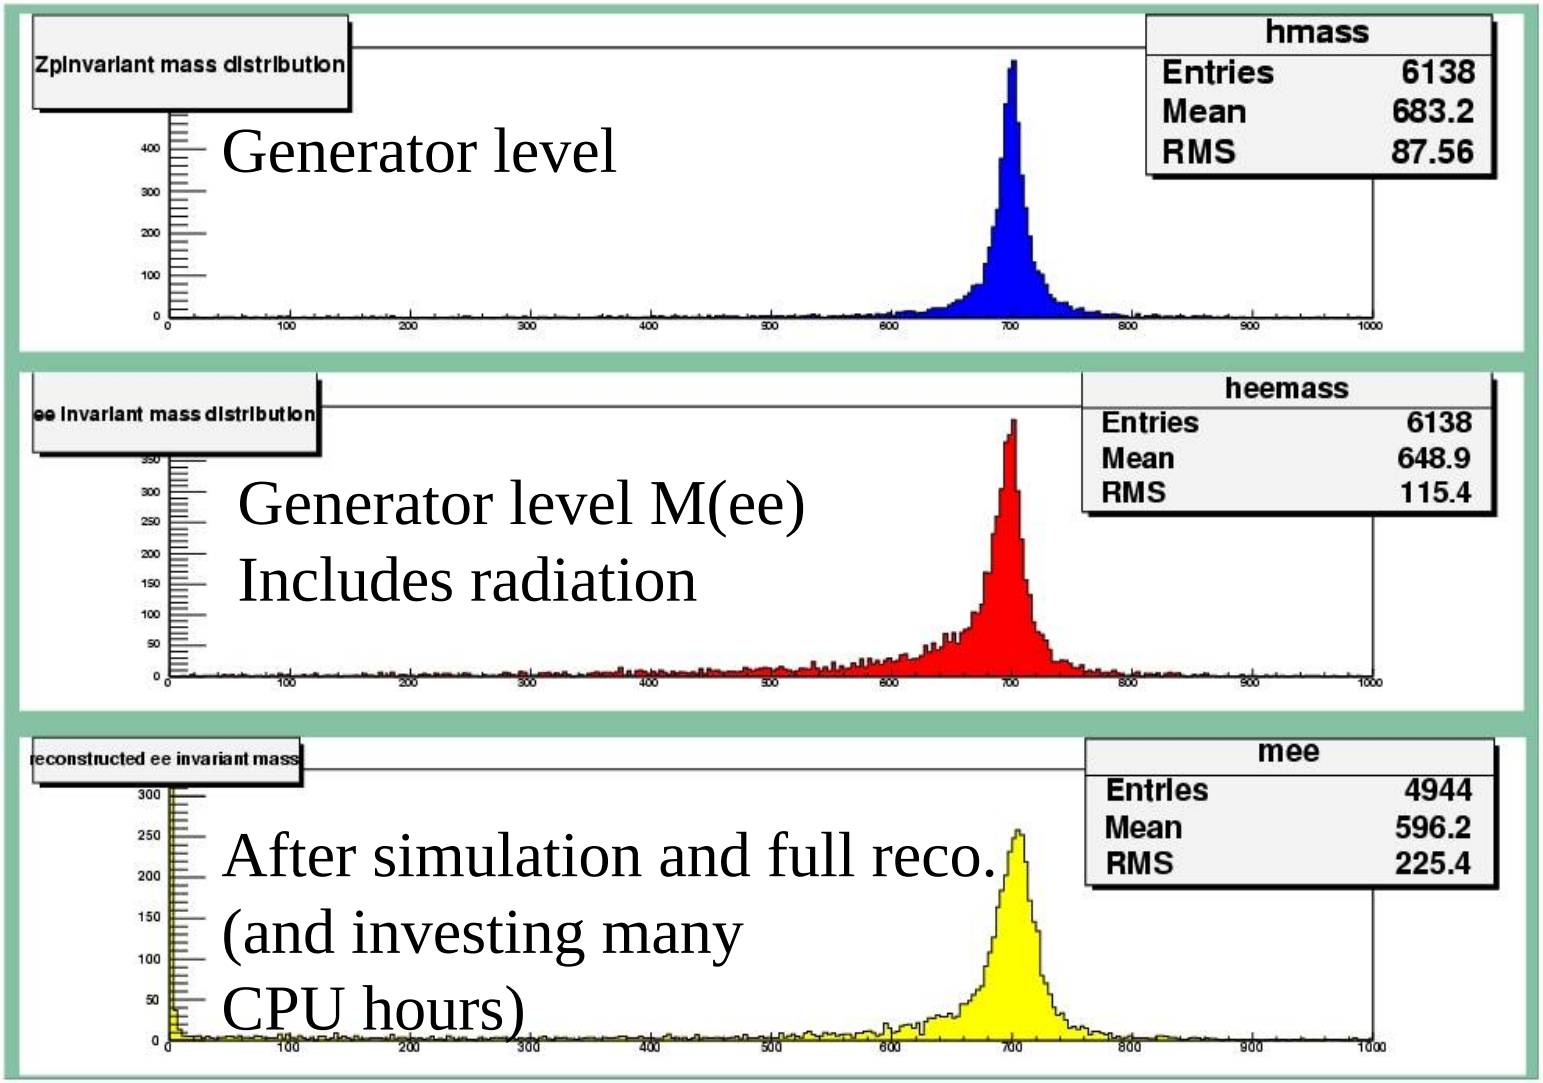

Condor
High Throughput Computing

How to submit a job into condor:

Typical job description file as used in the new tutorial:

zprime700_mumu_ROOTDIGIS_condor

```
universe = vanilla
Executable =  /uscms_data/d1/wenzel/Tutorial/scripts/run_rootdigis_zprime700_mumu.csh
Requirements   = Memory >= 599 && OpSys == "LINUX" && (Arch =="INTEL" || Arch =="x86_64")
Should_Transfer_Files = NO
Output  = /uscms_data/d1/wenzel/Tutorial/logs/zprime700_mumu_ROOTDIGIS_$(Cluster)_$(Process).stdout
Error = /uscms_data/d1/wenzel/Tutorial/logs/zprime700_mumu_ROOTDIGIS_$(Cluster)_$(Process).stderr
Log = /uscms_data/d1/wenzel/Tutorial/logs/zprime700_mumu_ROOTDIGIS_$(Cluster)_$(Process).log
notify_user = wenzel@FNAL.GOV
Arguments = $(Process) 1000 400

+LENGTH="SHORT"
Queue 5
```

Condor
High Throughput Computing

Each node is a submitter (queue into condor). Job (cluster) number not sufficient to uniquely identify a job, need also submitter. Good documentation on our web site.

- get a list of all the jobs and their status that you submitted into condor

condor_q -submitter wenzel@fnal.gov

condor_q

- Display information about request submitters

condor_status -submitters

- delete a job:

Submitter node

Job(cluster)

condor_rm 1

condor_rm -name cmswn088.fnal.gov 1

Mass Storage:

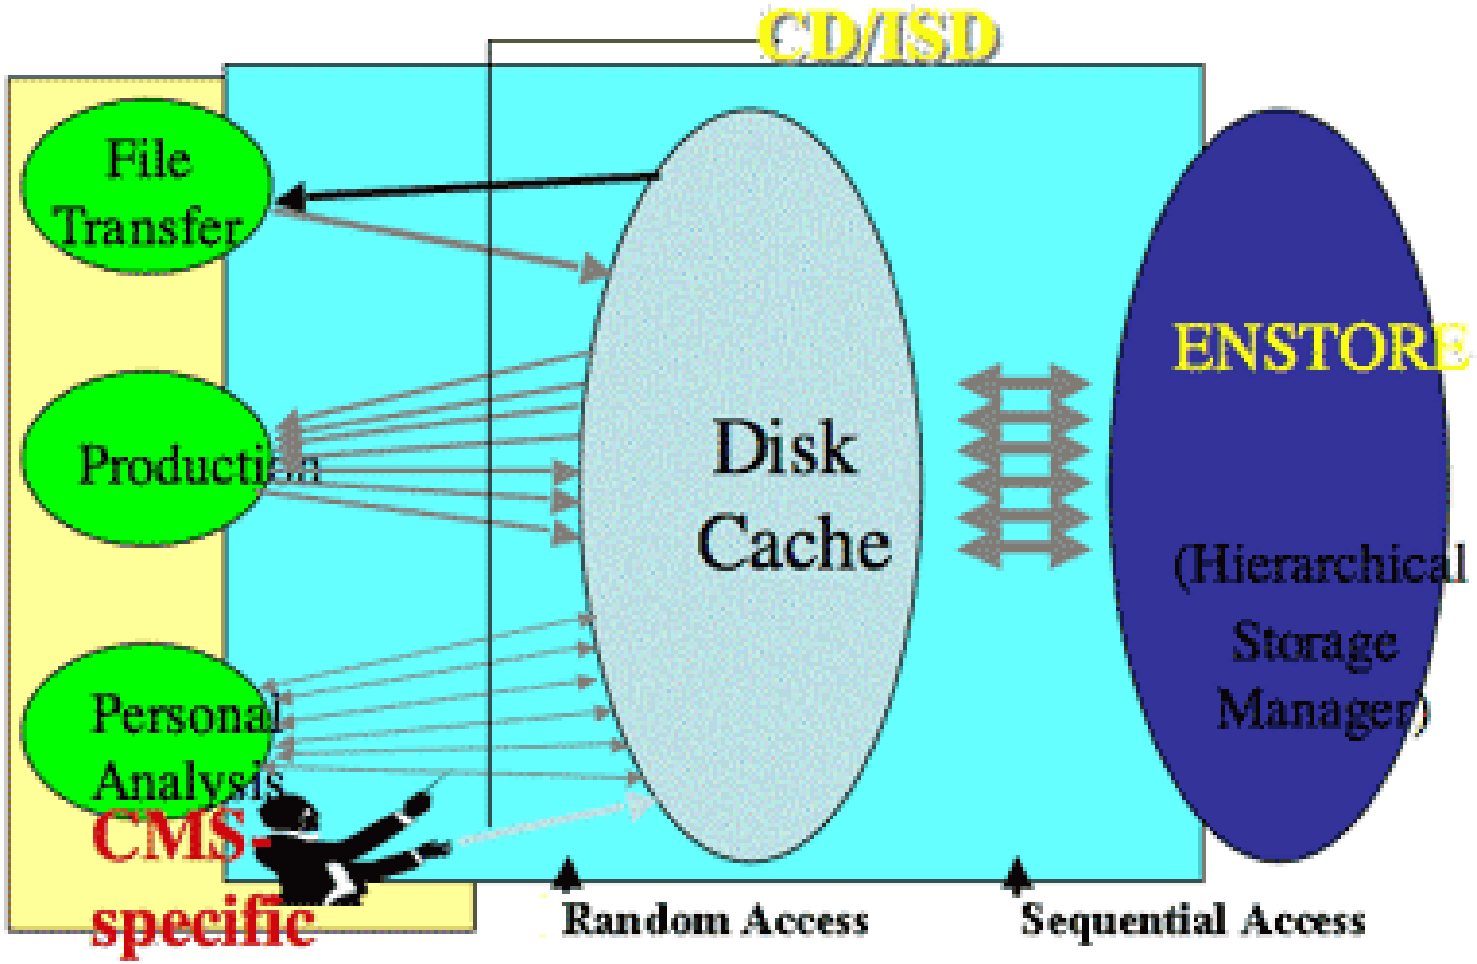

http://www.uscms.org/SoftwareComputing/UserComputing/MassStorage.html

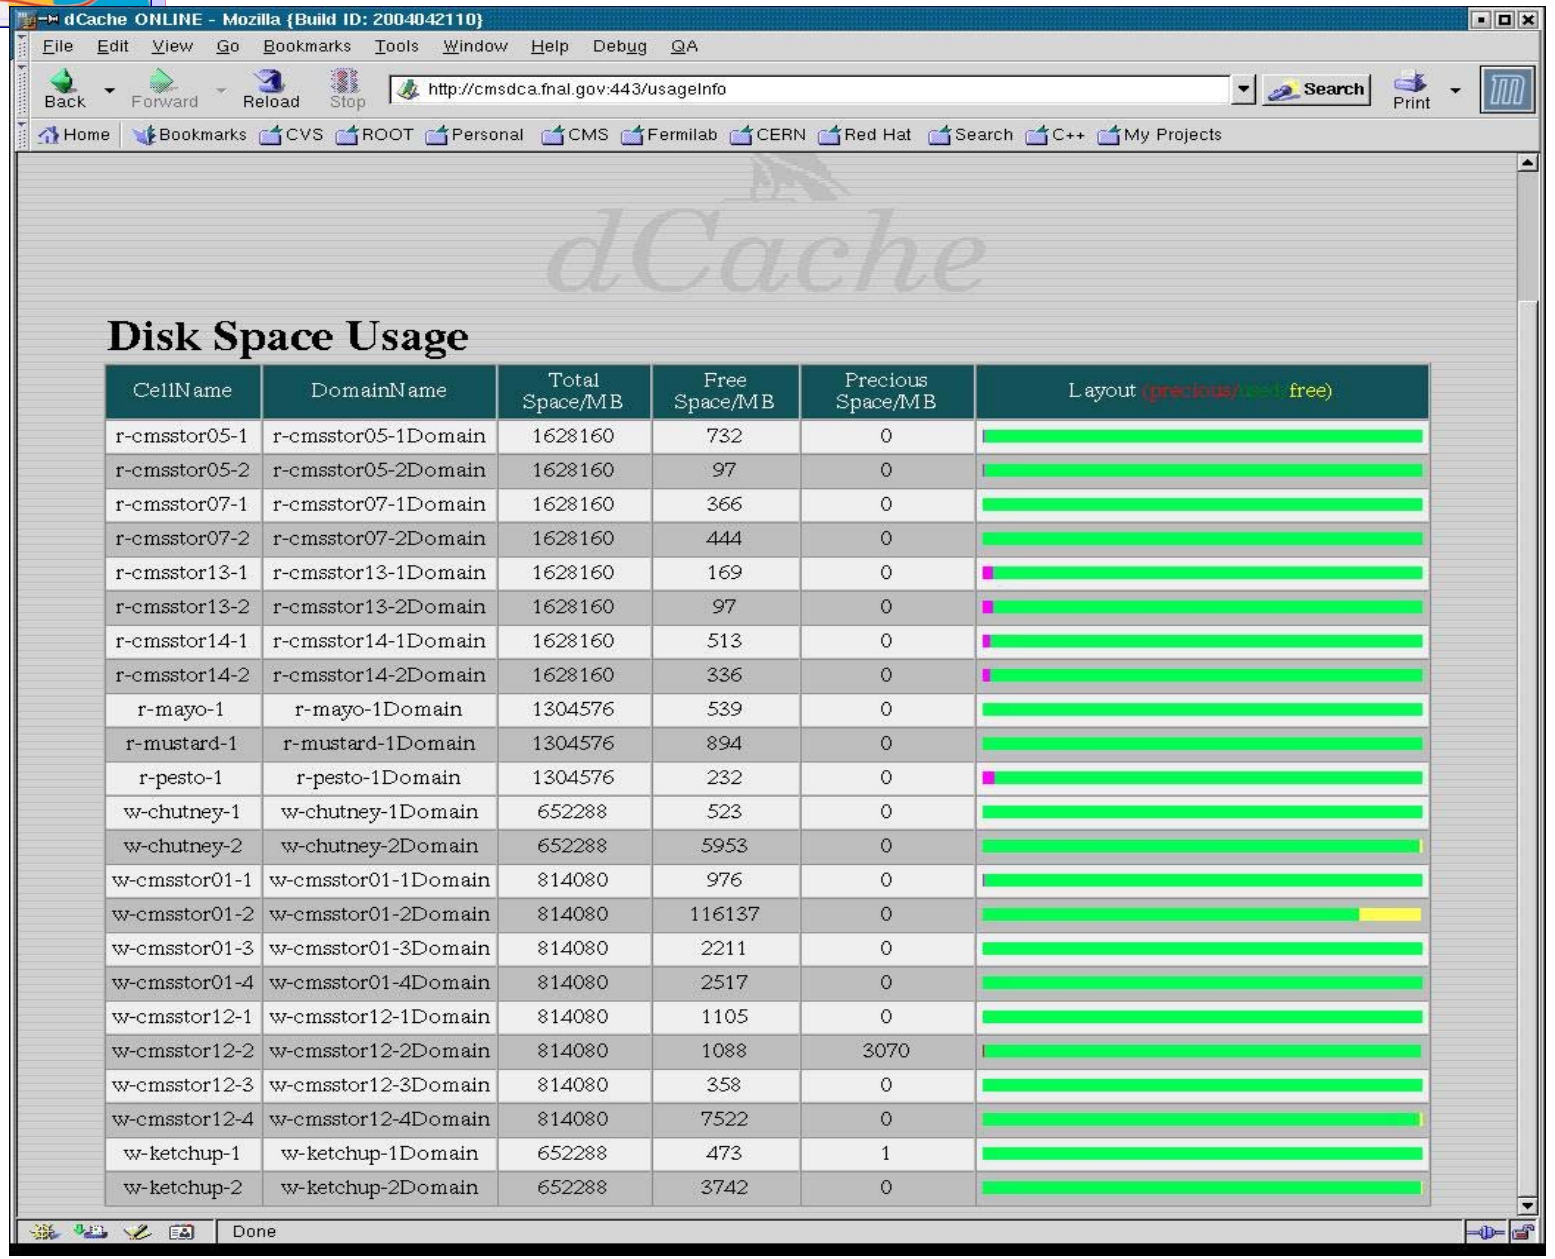
dCache ONLINE - Mozilla {Build ID: 2004042110}
File Edit View Go Bookmarks Tools Window Help Debug QA
Back Forward Reload Stop
http://cmsdca.fnal.gov:443/usageInfo
Search Print
Home Bookmarks CVS ROOT Personal CMS Fermilab CERN Red Hat Search C++ My Projects
dCache
Disk Space Usage
CellName | DomainName | Total Space/MB | Free Space/MB | Precious Space/MB | Layout (precious/used/free)
r-cmsstor05-1 | r-cmsstor05-1Domain | 1628160 | 732 | 0
r-cmsstor05-2 | r-cmsstor05-2Domain | 1628160 | 97 | 0
r-cmsstor07-1 | r-cmsstor07-1Domain | 1628160 | 366 | 0
r-cmsstor07-2 | r-cmsstor07-2Domain | 1628160 | 444 | 0
r-cmsstor13-1 | r-cmsstor13-1Domain | 1628160 | 169 | 0
r-cmsstor13-2 | r-cmsstor13-2Domain | 1628160 | 97 | 0
r-cmsstor14-1 | r-cmsstor14-1Domain | 1628160 | 513 | 0
r-cmsstor14-2 | r-cmsstor14-2Domain | 1628160 | 336 | 0
r-mayo-1 | r-mayo-1Domain | 1304576 | 539 | 0
r-mustard-1 | r-mustard-1Domain | 1304576 | 894 | 0
r-pesto-1 | r-pesto-1Domain | 1304576 | 232 | 0
w-chutney-1 | w-chutney-1Domain | 652288 | 523 | 0
w-chutney-2 | w-chutney-2Domain | 652288 | 5953 | 0
w-cmsstor01-1 | w-cmsstor01-1Domain | 814080 | 976 | 0
w-cmsstor01-2 | w-cmsstor01-2Domain | 814080 | 116137 | 0
w-cmsstor01-3 | w-cmsstor01-3Domain | 814080 | 2211 | 0
w-cmsstor01-4 | w-cmsstor01-4Domain | 814080 | 2517 | 0
w-cmsstor12-1 | w-cmsstor12-1Domain | 814080 | 1105 | 0
w-cmsstor12-2 | w-cmsstor12-2Domain | 814080 | 1088 | 3070
w-cmsstor12-3 | w-cmsstor12-3Domain | 814080 | 358 | 0
w-cmsstor12-4 | w-cmsstor12-4Domain | 814080 | 7522 | 0
w-ketchup-1 | w-ketchup-1Domain | 652288 | 473 | 1
w-ketchup-2 | w-ketchup-2Domain | 652288 | 3742 | 0
Done

How to use the mass storage

to list files in mass storage do:

ls /pnfs/cms

Create your directory in mass storage:

mkdir /pnfs/cms/WAX/2/username

Then, to copy a file into dCache, do:

dccp myfile.dat /pnfs/cms/WAX/2/username/myfile.dat

or, to copy the file over to your local disk, do

dccp /pnfs/cms/WAX/2/username/myfile.dat .

To pre stage a file into dCache (from Enstore):

dccp -P /pnfs/cms/WAX/2/username/myfile.dat

To check if a file is in dCache:

dc_check /pnfs/cms/WAX/2/username/myfile.dat

Note pnfs might look like a regular File system but it's not! restrict commands to:

mkdir, rmdir, ls, rm (don't cd into /pnfs, don’t run find commands)

Pool (CMS persistency) know how to deal with files in dCache!

Integrating the Desktop

- Besides central computing we want to make use of the powerful PC's running linux (plenty of disk, cpu, freedom, it's all yours ..)
- Created CMS desktop workgroup containing everything you need to run CMS software on your PC (afs ....). CMS software is kept up to date in AFS and IBRIX.
- Uses the same setup as the UAF for the CMS software environment. This will allow you to run (edit, compile debug) CMS software on your Desktop and store data locally.
- SL3: (e.g common LPC desktop cluster on the 11th floor) condor, dcache … part of the system, IBRIX available via NFS.

The LPC desktop cluster

- http://www.uscms.org/SoftwareComputing/UserComputing/lpcdt/LPC.htm

- Synchronized with UAF. Meaning if you have an account on the UAF you can log into the common desktops with your kerberos password. CMS Software,IBRIX, dCache are availabe, Local condor batch system to run jobs. It's work in progress we are still trying to make it more useful.