

# *Computing Developments at HERMES*

*A. Kisselev  
PNPI, St. Petersburg, Russia*

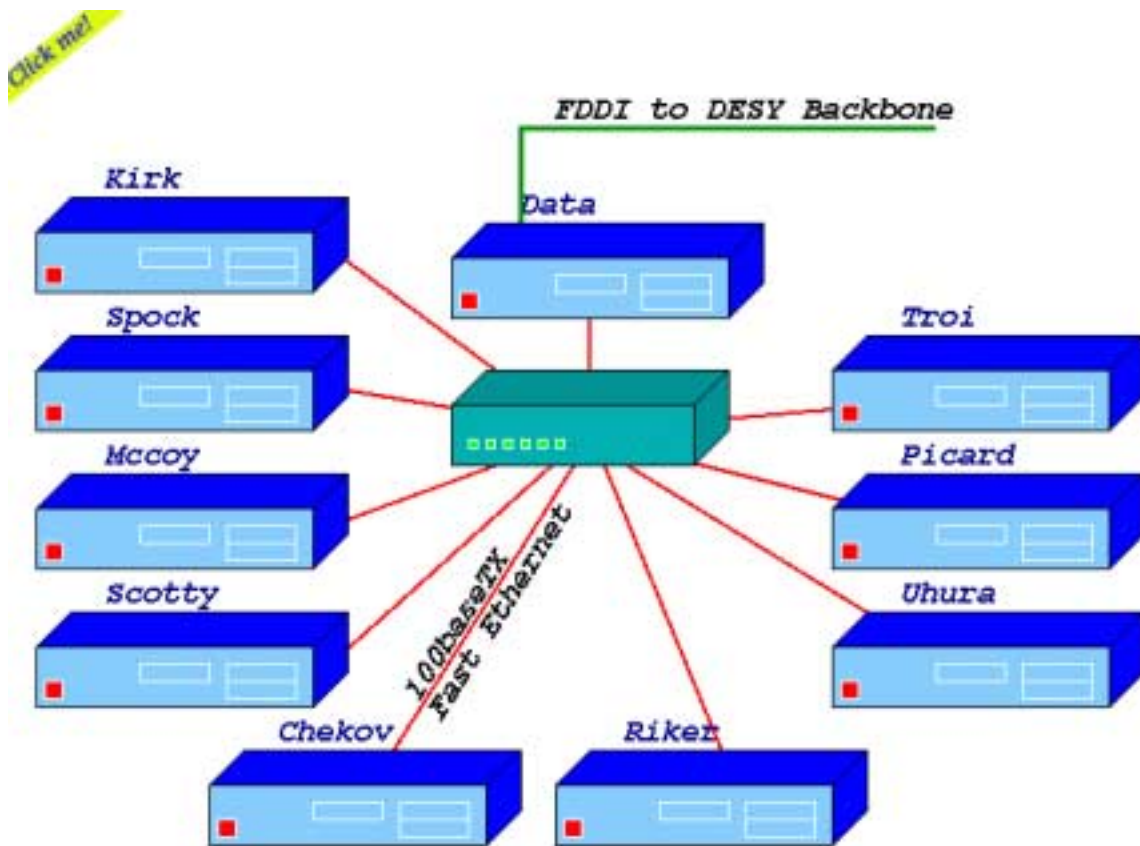
*April 29th, 2002, DESY  
DESY Computing Seminar*

- *HERMES PC farm evolution*
- *Present setup overview*
- *Conclusions and outlook*

## Summer'1996

- *HERMES offline computing is a 28CPU SGI Challenge XL*
- *Lack of CPU power for reconstruction*
- *Solution: move data production to a cluster of Linux-based PCs*
  - *Production software ported to Linux*
  - *Promising performance measurements*

## First HERMES PC farm - Fall'1996



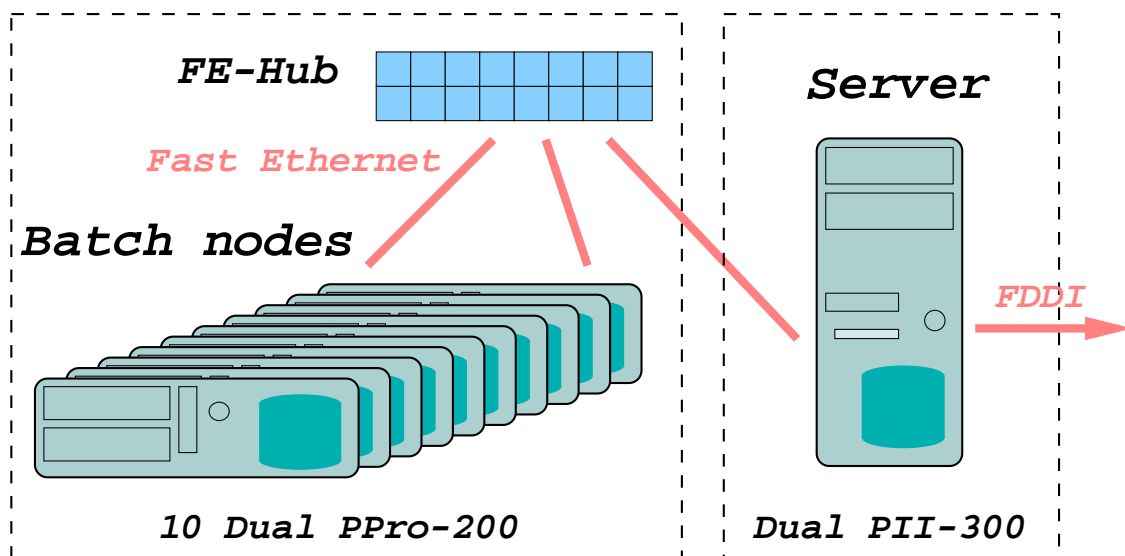
- *First Linux PC farm in HEP*
- *10 single-CPU Intel PPro-200*
- *At least CPU-equivalent to the HERMES SGI*
- *Custom batch system*
- *Home and HERMES software shared via NFS*

## *Troubles of 1997*

- *Disastrous upgrade to Dual CPU boards in spring*
  - *on average 1 week of uptime*
  - *production chain screwed up*
- *First SMP-safe Linux kernel released half a year later*
- *In general it was a good lesson*
  
- *Wolfgang Wander left HERMES*

## Reconfiguration - Fall'1998

- *First time beaten by IBM disks*
- *Looking how to improve reliability*



- *Clients mount \*everything\* via NFS*
- *Custom cloning software*
- *RARP for node identification*
  
- *Linux software RAID-1 on the server*
- *Regular ADSM backup*

## *Beginning of '2000*

- *Tough discussions on upgrade plans start*
- *Ageing SGI Challenge should be phased out*
- *Available PC farm is useless for users*
- *General requirements to the successor*
  - *Order of magnitude CPU power increase*
  - *Up to 10TB disk space*
  - *Some batch system*
  - *Reliable hardware and software*
  - *Good price/performance ratio*
  - *Scalability (incremental upgrade)*
  - *Use central DESY IT services as much as possible*
- *For the time being upgrade existing PC farm with 7 Dual PIII-550 PCs and fast ethernet switch*

## *Suggested configuration*

- *Single mainframe is not feasible*
  - *Compute server*
  - *Login server*
  - *File server*
  - *Anything else?*
- *Distribute functionality between differently configured PCs*
  - *Batch nodes - x86 hardware under Linux*
  - *Interactive nodes should be the same platform*
  - *“Master” server can also be x86 Linux*
- *Aggange shared disk space*
- *Glue everything together with a batch system*
- *... So what to do with the File server?*
  - *SGI Origin 2000*
  - *Some Linux-based solution*

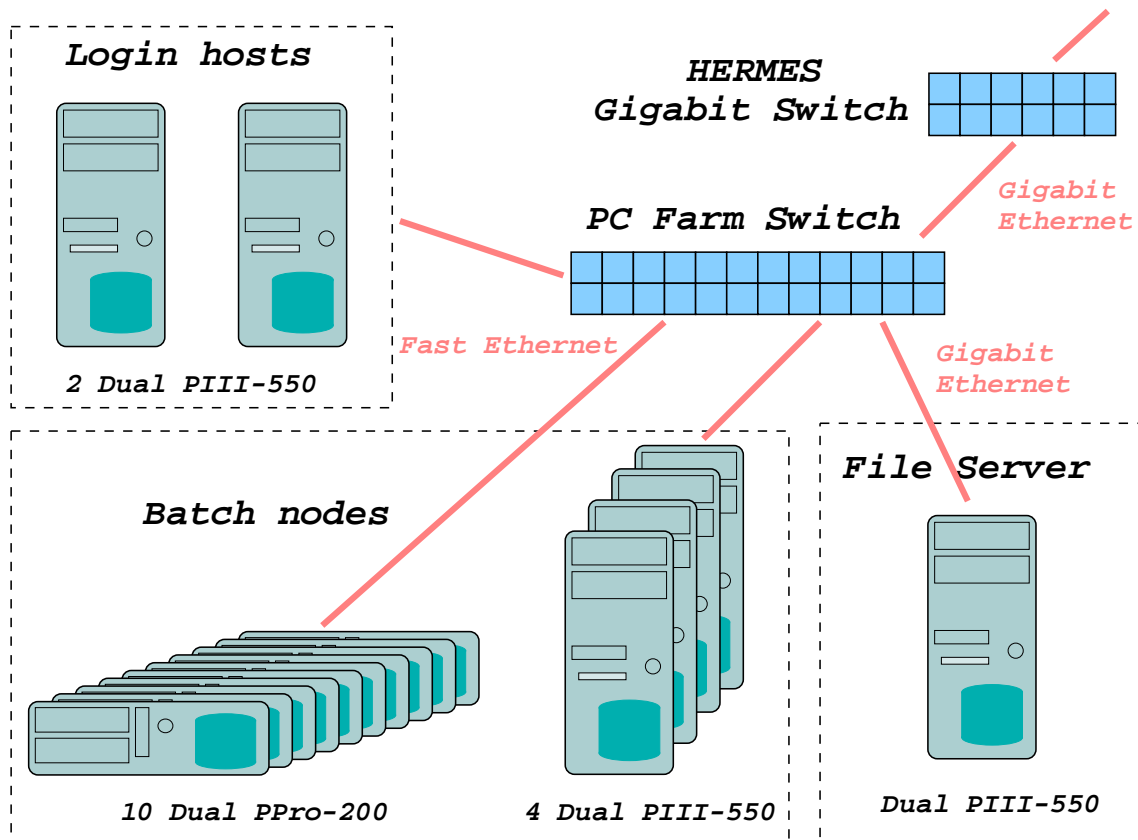
## *Linux IDE File servers?*

- *Intel x86 platform reliability in general*
- *Specially designed commodity hardware*
  - *RAID controllers*
  - *Hot swappable disk bays*
  - *Performant, reliable and cheap disks*
  - *Ability to pack a lot of disk space in one box*
  - *Appropriate boxes*
  - *Redundant power supplies*
  - *Remote access*
  - *Gigabit network*
- *Software components*
  - *Journalling filesystem*
  - *File access protocol*
  - *AFS client availability*
- *Positive experience at CERN*
- *Initiative coming from H1*

## *Software alternatives*

- *Journalling filesystem type*
- *Linux NFS vs all the rest*
- *Software vs hardware RAID*
- *Linux kernel version*
- *Batch system choice*

# Test Setup - Fall'2000

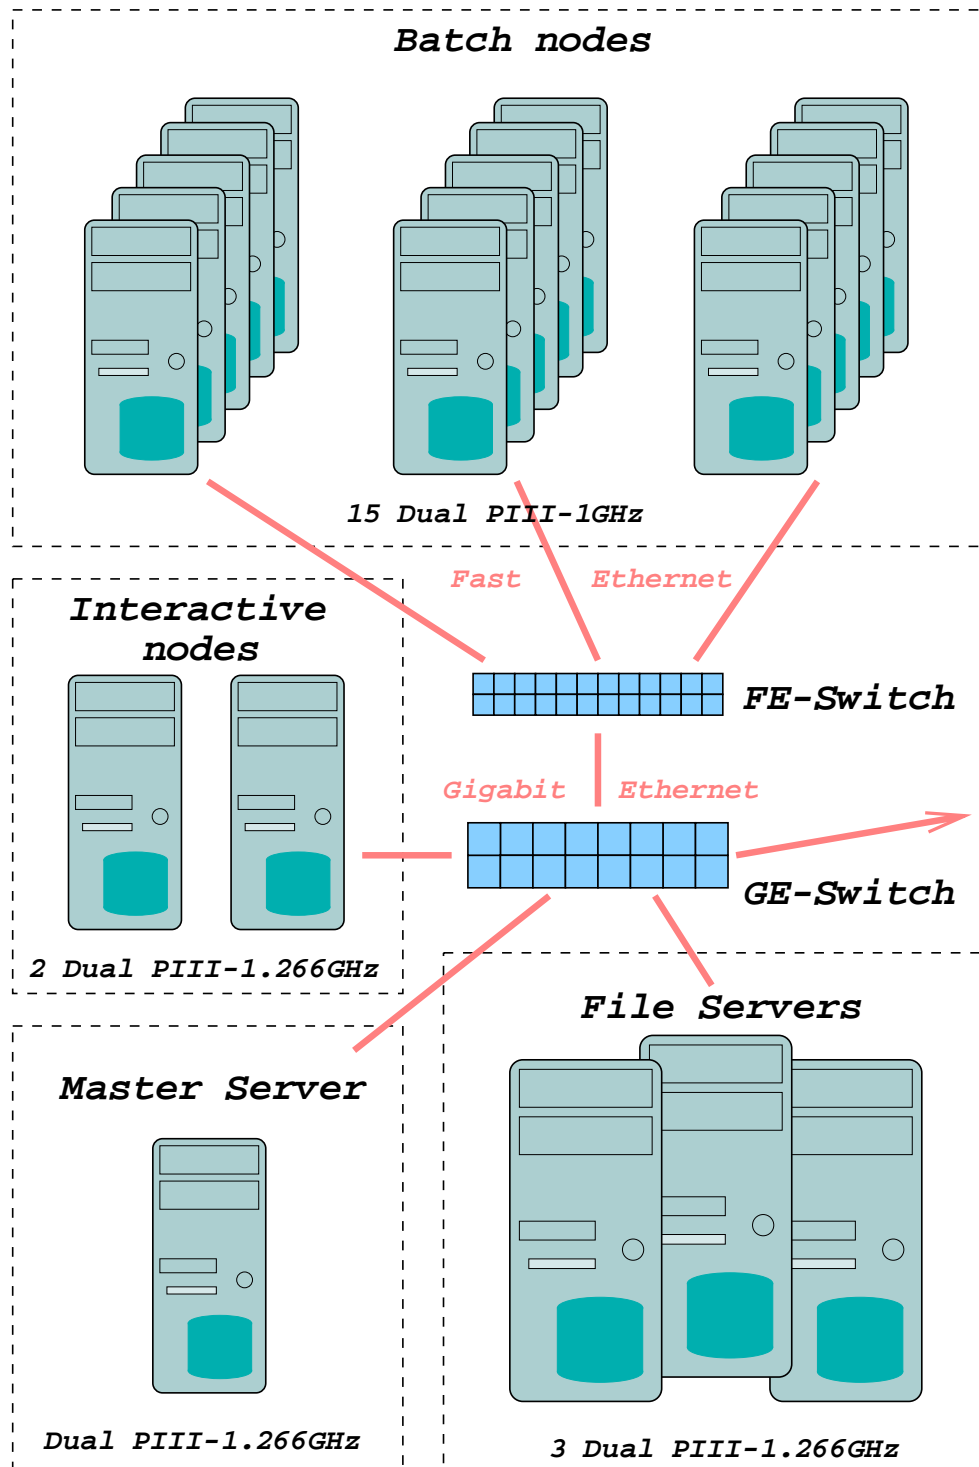


- *Times more powerful than HERMES SGI*
- *Served as a testbed for more than a year*
- *Was reconfigured few times in order to try out possible options*
- *January'2001 - 500GB NFS server added*

## *HERMES software on the PC farm*

- *Major part was ported in 1996*
- *Few g77 specific fixes done*
- *Reconstruction chain and micro-DST production configured for PBS*
- *Assisted users with software porting*
- *Floating-point exceptions*
- *File copy/allocation scripts mostly rewritten*

# HERMES PC farm - January'2002



# *Real estate - April'2002*



## *Year 2000 data sample*

- *About 15TB of raw EPIO files*
- *About 600GB of reconstructed ADAMO output*
- *Reconstruction takes 1 month on 10 Dual PIII-1GHz batch nodes*
- *Less than 200GB micro-DST files*
- *Light-weighted user code may require up to 3-5MB/s network bandwidth*

## General overview

- Overall CPU power of the batch nodes order of magnitude higher than on the HERMES SGI
- 3.6TB of disk space on the File Servers
- Identical file system layout on all nodes
- File sharing through Linux kNFS
- Kernel 2.2.20 (+ a dozen of patches)
- ext3 journalling file system (v0.0.7a)
- DESY Linux 3 everywhere but on the file servers

## *Common hardware*

- *Supermicro P3TDLE mainboards*
- *Intel PIII 1GHz or 1.266GHz CPUs*
- *3Ware 3W-6200 and 3W-7850 disk controllers*
- *Western Digital 40GB and 100GB disks*
- *Hot swap disk bays*
- *Chieftec chassis*
- *Netgear GA620 gigabit network cards*
- *Intel EEpro fast ethernet cards*




# *Fail Safe Operation and Manageability*

- *Redundant interactive nodes*
- *Hardware features*
  - *RAID-1 and RAID-5 disk space only*
  - *Hot swap/repair available for all the disks*
  - *Redundant PSU on all critical nodes*
  - *Health monitoring (to be done)*
  - *Custom alarm system with SMS enabled*
- *Remote reset option*
- *Remote BIOS level access for experts*
  
- *Backups*
- *Custom cloning software*

## Security model

- *Inherited from DESY Linux 3 with few exceptions*
  - */etc/hosts.equiv: “@hermesfarm” instead of “@afs”*
  - *No /root/.rhosts*
- *NFS volumes exported in “root squash” mode*
- *Monitor security updates*

# *“hermes.desy.de” stays with HERMES*



3) *mail addr@hermes* → *kirk...*

3a) *list@hermes* → *list members*


3b) *user@hermes* → *user@mail.desy.de*

## Master server

- *Core component of the whole setup*
- *Configured to be lightly loaded*
- *Expected uptime - several months*
- *Nightly incremental backup (TSM)*
- *Nightly local backup of system disk*
- *Services*
  - *Geometry and calibration database*
  - *Apache Web server*
  - *Virtual Server client*
  - *PBS server*
  - *AFS token server for PBS system*
  - *MySQL server*
  - *NFS server (common software, links database)*
  - *Majordomo mailing lists*
  - *Image database for cloning*
  - *MRTG-based monitoring*
  - *DHCP server*


## *Interactive nodes*

- *The only 2 nodes available for user login (via SSH, XDMCP)*
- *2GB of RAM, to be increased to 4GB*
- *Few kernel limits tuned for multi-user operation*




- *Example weekly CPU usage pattern*
- *Not overloaded, but uptimes of about 4 weeks only*

## *Interactive nodes, cont'd*



- *worf.desy.de*



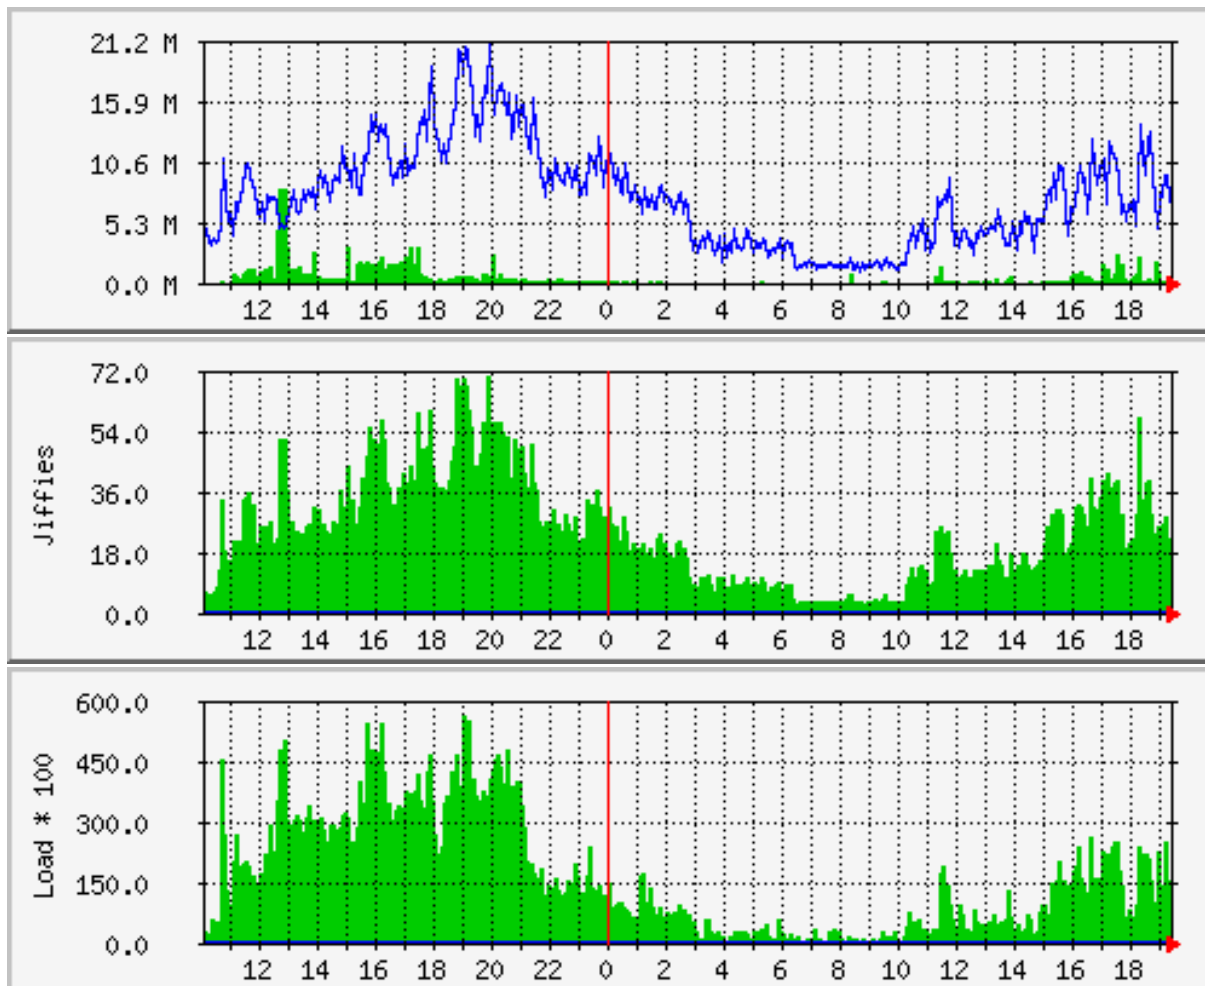
- *geordi.desy.de*
- *Example monthly “population” plot*
- *Reasonable distribution of user load*

## *Jumbo NFS servers*




- *4x4 300GB RAID5 volumes*
- *All volumes bootable*
- *Persistent mount points*
- *Unified root file system images*

## *NFS servers, cont'd*



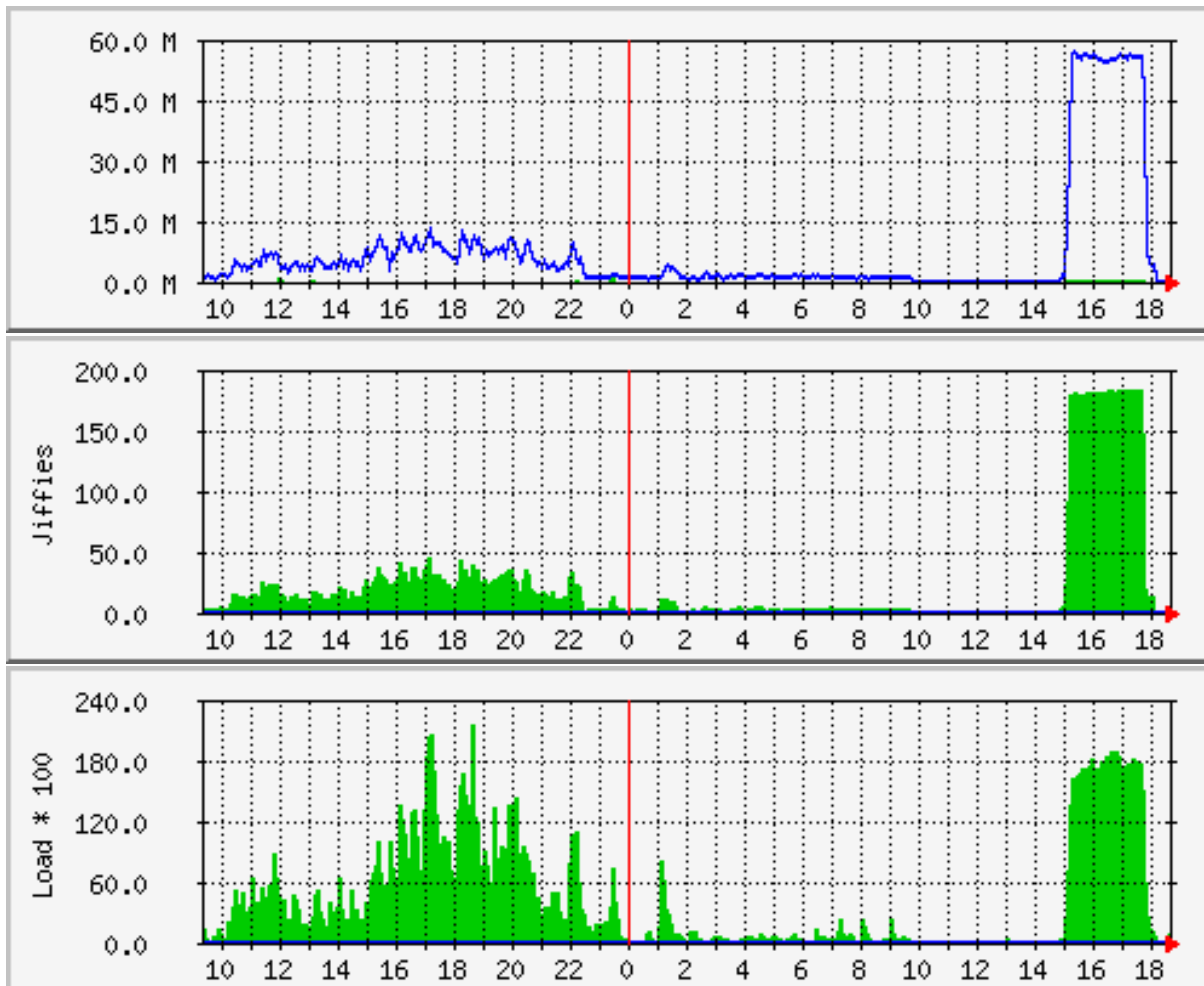
- *Example daily plot of traffic and CPU load*
- *Typical “write once/read many” access pattern*

## *NFS servers, cont'd*



- *Typical weekly traffic graph*
- *Mostly read access, below 10MB/s on average*
- *Small number of heavy network I/O streams*

## *NFS servers, HERMES extreme :-)*



- *NFS memory-to-memory copy over network peaks at 55MB/s*
- *Still limited by available CPU power (Dual Intel PIII 1.266GHz)*

## *Disk space allocation*

- *In total 12 huge (300GB) RAID5 volumes*
- *3x “user” volumes (/user01..03)*
  - *Single entry point for a given user (like /user/kisselev) created manually*
  - *Soft quota and permanent disk usage monitoring*
  - *3GB of a “guaranteed” disk space for every user*
  - *Quota is increased per request*
  - *At present volumes are occupied by 30%*
- *(Almost) no group-based disk space allocation*
  - *A single shared volume for groups*
  - *Small shared volume for group Web pages*
  - *UNIX group structure inherited from the hermes SGI*

## *CPU power sharing*

- *Dedicated batch nodes for data productions*
- *Same for Monte-Carlo productions*
- *PBS tuning (or patching :-)*
- *Queue limits for users*
- *Killer script on interactive nodes*

## Design Limitations

- *Linux kernel limits in 2.2 series*
  - *Switch to kernel 2.4*
- *Disk access slows down during concurrent read/write operations*
  - *Replace 8-port cards by 4-port ones*
  - *Switch to kernel 2.4 and tune disk I/O elevators*
  - *Put user volumes in a separate file server and increase RAM (cacheing on the server)*
- *PAW-based interactive analyses suffer from slow network I/O*
  - *Increase RAM on interactive nodes up to 4-8GB (cacheing on the client)*

## *Troubles during first 4 months*

- *Master Server*
  - *“Drive timeout” on one of the disks*
- *Interactive nodes*
  - *Hanged up 4 and 5 times respectively*
- *Batch nodes*
  - *One solid freeze*
  - *Strange problems with one of the nodes*
  - *CRC error on one of the old IBM DTLA disks*
- *File servers*
  - *One solid freeze*
  - *One 100GB WD disk died*
  - *4 “drive timeouts”, all related to cabling*
  - *3Ware “design features”*

# *Extention Options and Scalability*

- *Interactive nodes*
  - *Install more than 2 boxes*
  - *Upgrade existing ones*
  - *Replace by more modern hardware*
- *Batch nodes*
  - *Can be added infinitely as long as they don't kill file servers*
- *File servers*
  - *Add more when disk space becomes tight*

## *Conclusions and Outlook*

- *Setup in general works as projected*
- *Look forward to fix remaining issues*
- *Look for time to write decent description*
- *We are already in upgrade phase :-)*