Computing news

Szymon Gadomski, May 13th, 2011

- monitoring the Network File System
- PROOF on interactive machines

Monitoring NFS response times

- Motivation
  - occasional complaints of “everything slow”
  - sometimes it is clear that the NFS is slow
  - we suspected batch jobs doing I/O
  - suggestion of Will (early March)
    - monitor the response time systematically
    - look for correlation to number of jobs by user
- monitoring setup Fri Mar 4th
  - monitor response time of ls
  - a few cycles of looking at results and adding information

Raw data

the logging script:

```
dList=['/atlas/users','/atlas/software',
       '/atlas/data1','/atlas/data2',
       '/atlas/data3','/atlas/data4',
       '/neutrino/data1','/neutrino/data4',
       '/afs/cern.ch/atlas/software','/home']
```

file systems to monitor

```
for d in dList:
    cmd='time ls '+d
```

measure response time

+ # grid data copy processes (dq2-get, lcg-cp, globus-url-copy)

+ load, swap used, time to run 'date'

the log file:

```
[gadomski@atlas013]~% tail /var/log/nfs-response.log
Thu May 12 11:15:01 CEST 2011 0.003 0.003 0.002 0.002 0.002 0.003 0.003 0.003 0.009 0.002 2 0 0  0.00 104 0.001
Thu May 12 11:20:01 CEST 2011 0.003 0.002 0.002 0.002 0.003 0.002 0.002 0.002 0.009 0.002 2 0 0  0.00 104 0.001
Thu May 12 11:25:01 CEST 2011 0.003 0.003 0.002 0.002 0.003 0.003 0.003 0.003 0.008 0.002 2 0 0  0.29 104 0.001
Thu May 12 11:30:01 CEST 2011 0.003 0.003 0.002 0.003 0.004 0.002 0.002 0.002 18.281 0.002 2 0 0  0.71 104 0.001
Thu May 12 11:35:01 CEST 2011 0.003 0.003 0.002 0.002 0.003 0.002 0.003 0.003 0.009 0.002 2 0 0  0.83 104 0.002
```

Distribution of times to do ls

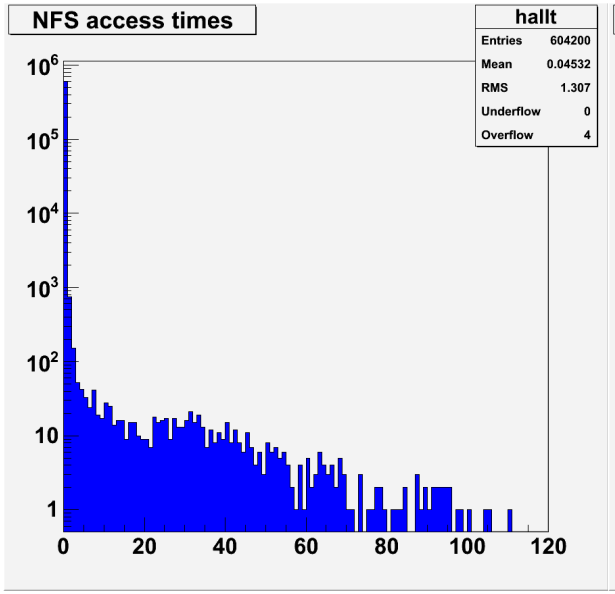

time to do ls [s]

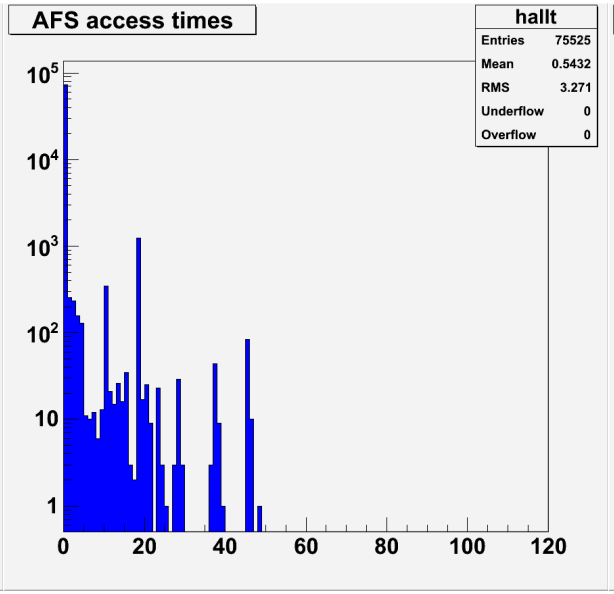

time to do ls [s]

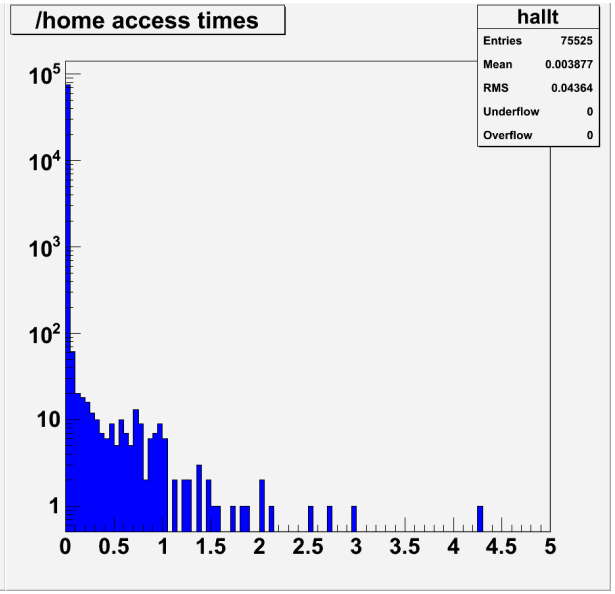

time to do ls [s]

Average times: **NFS (file server in the setup): 45 ms**
AFS (file server at CERN): 540 ms
/home (disk in the machine): 4 ms

The problem we expected to see

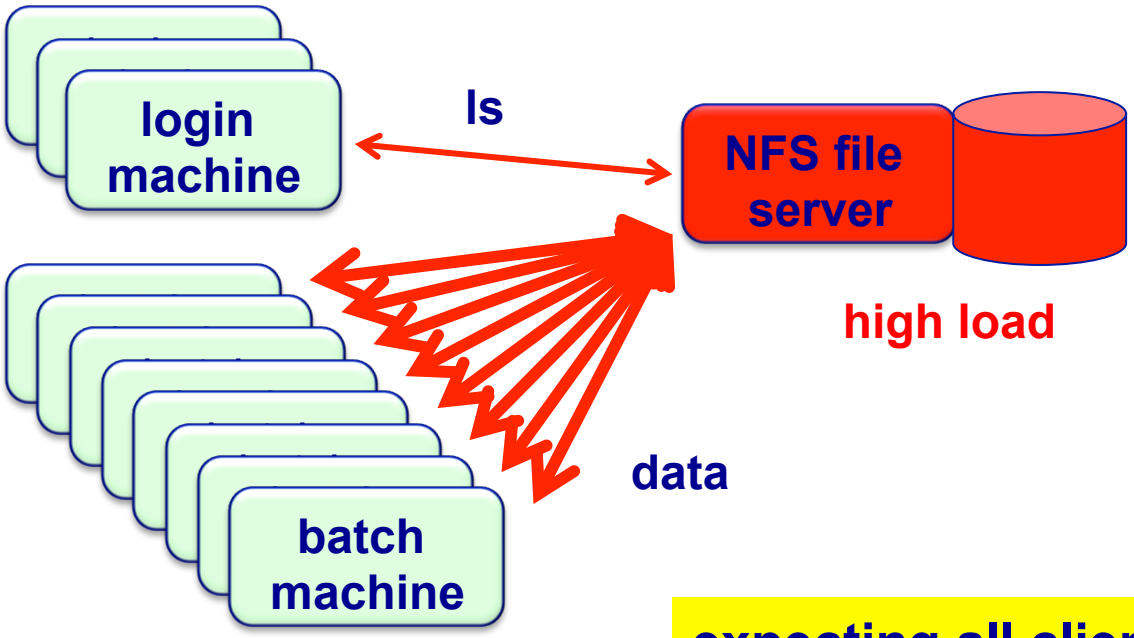

expecting all clients logging large access times

We do see overloaded file servers

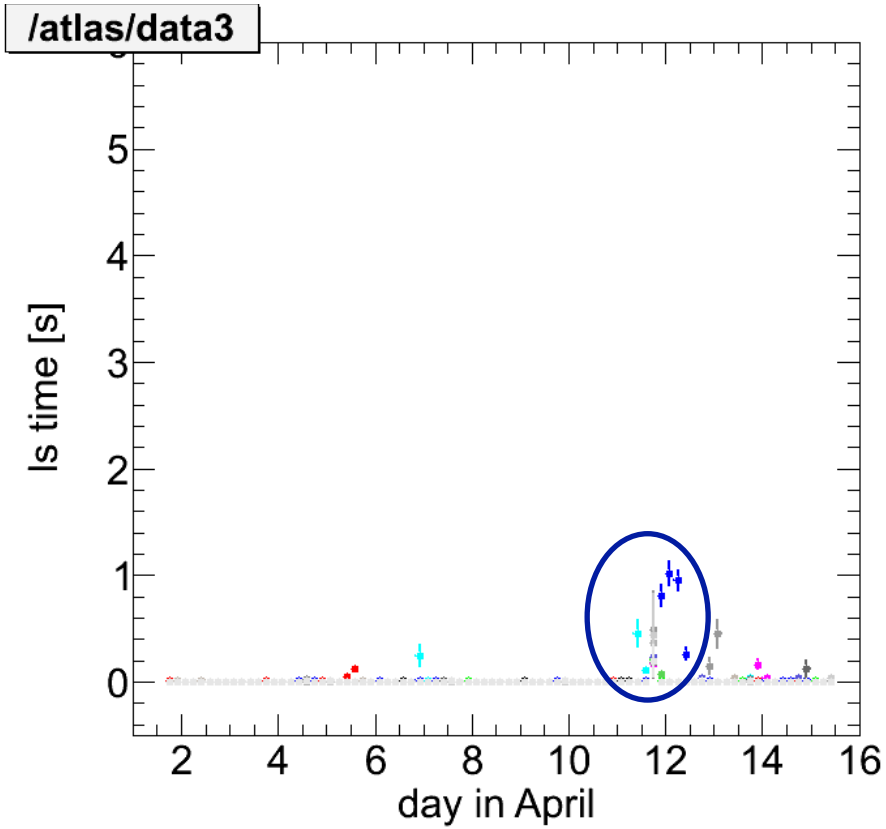

- but rarely
- in the period of early March to mid April this was not the most frequent reason of slow response

Usually no correlation between clients

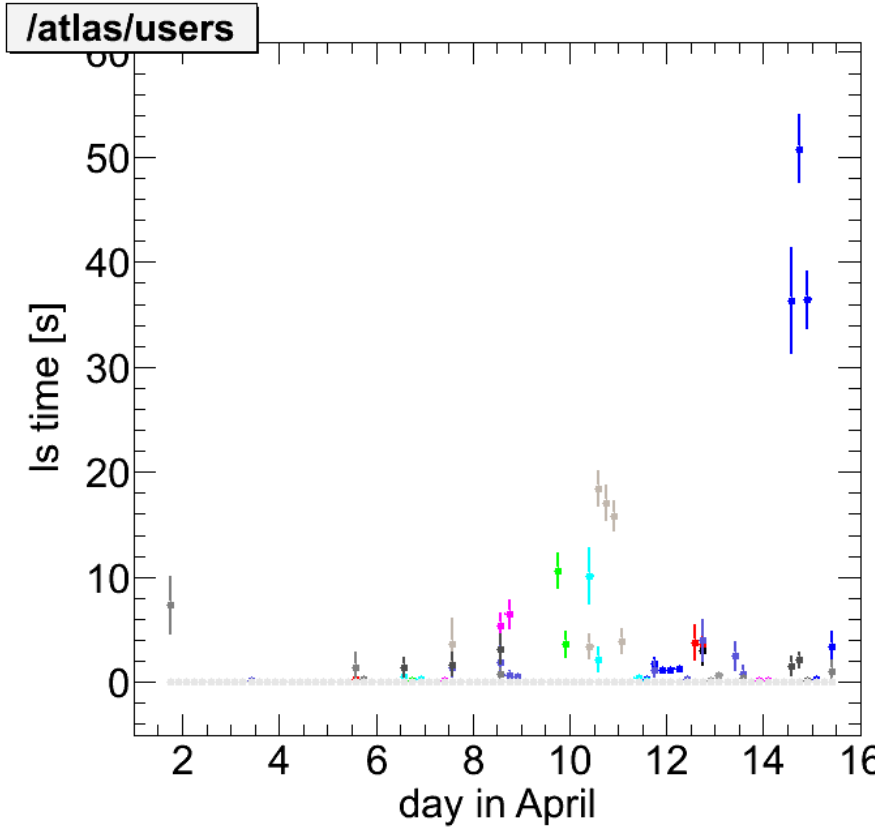

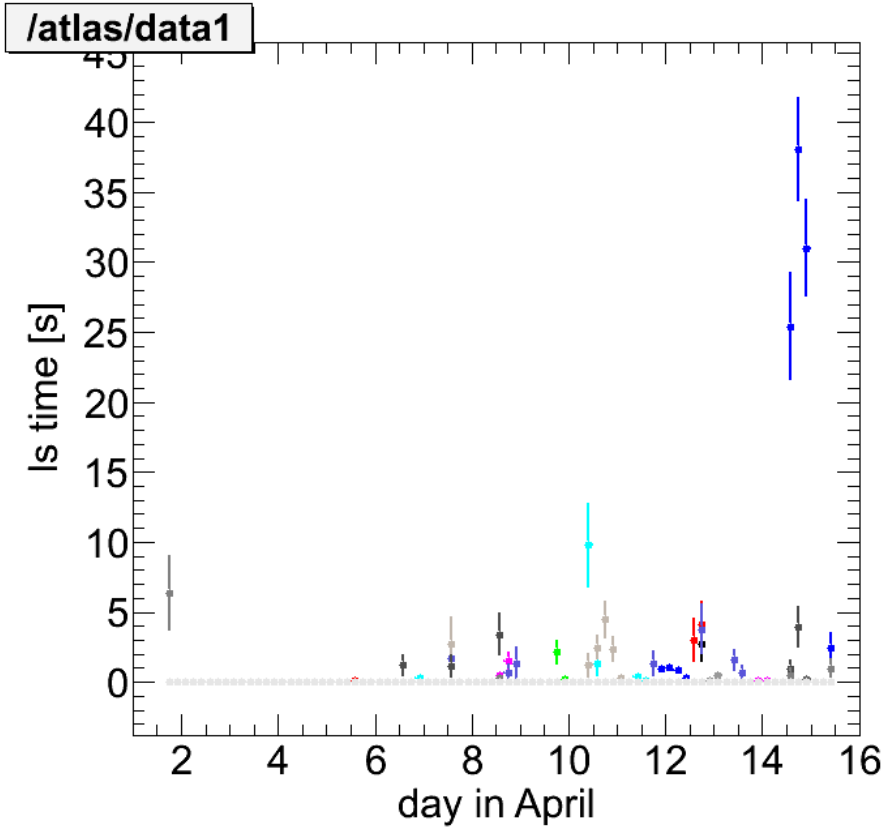

When the NFS response is slow:

- **other clients are OK**
- **another file system on the same server, seen from the same client, is slow too**

Another example

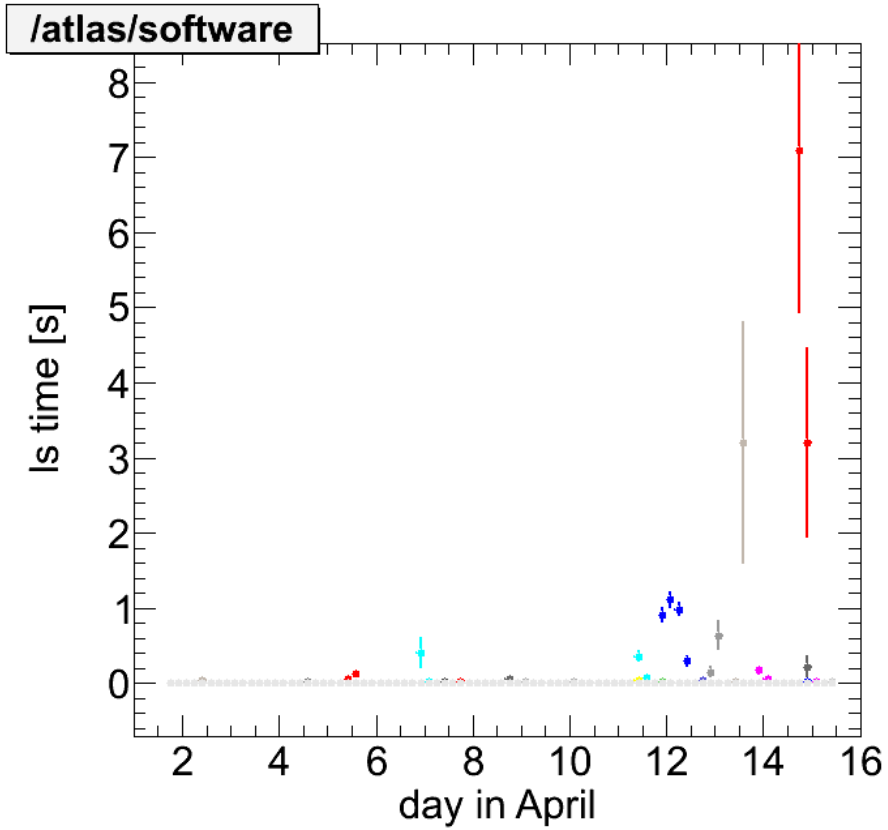

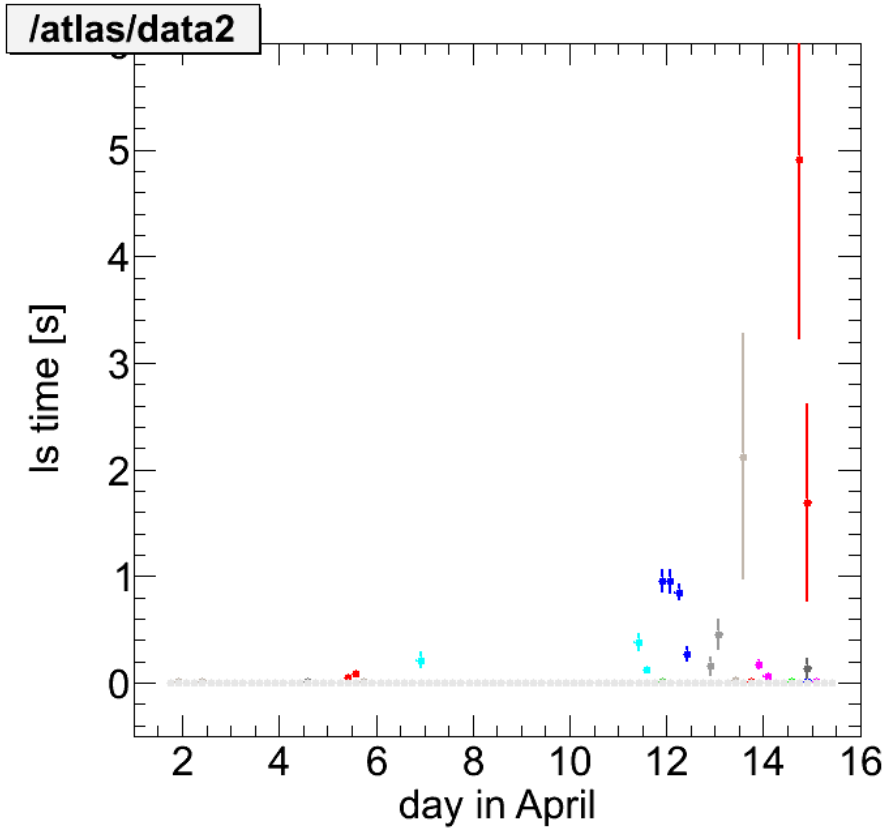

History of one machine

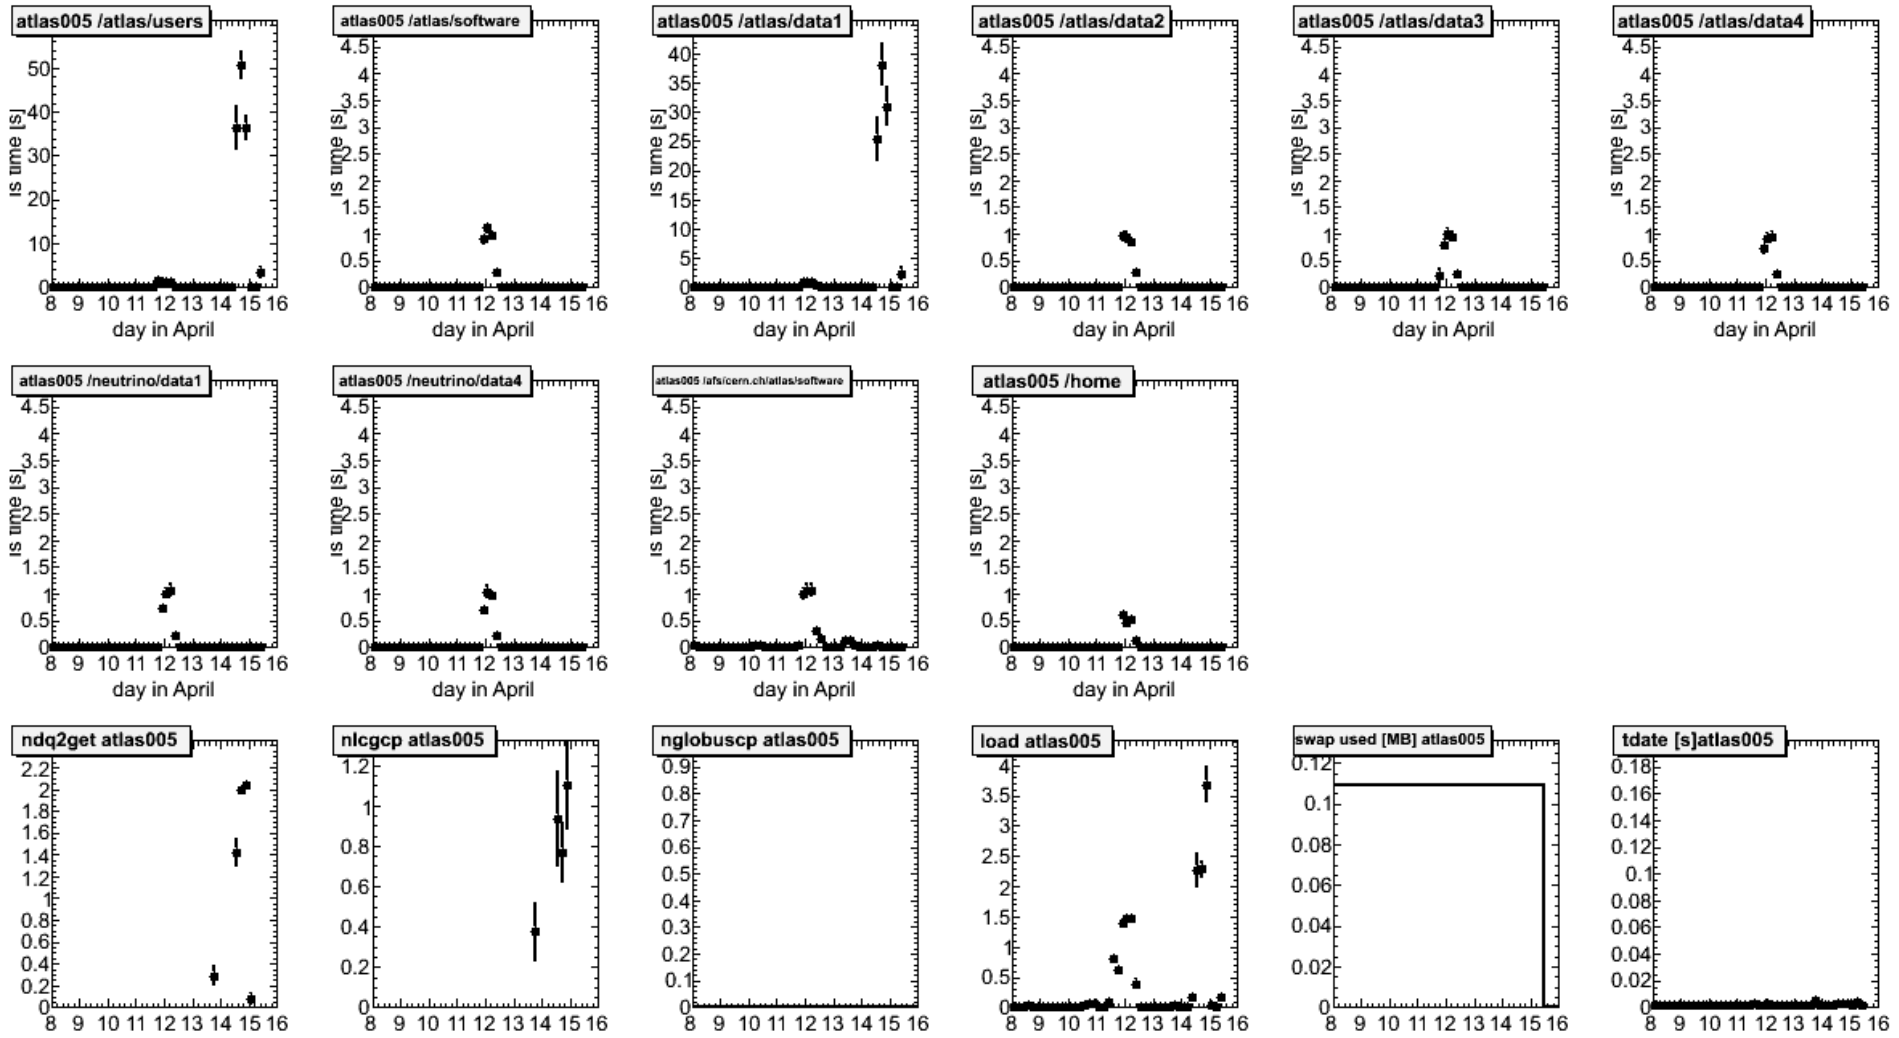

The problem we see

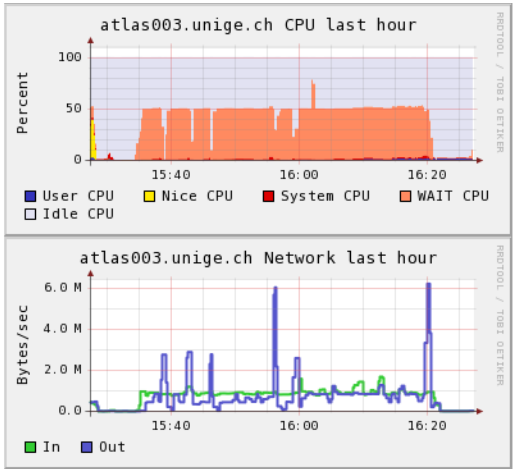

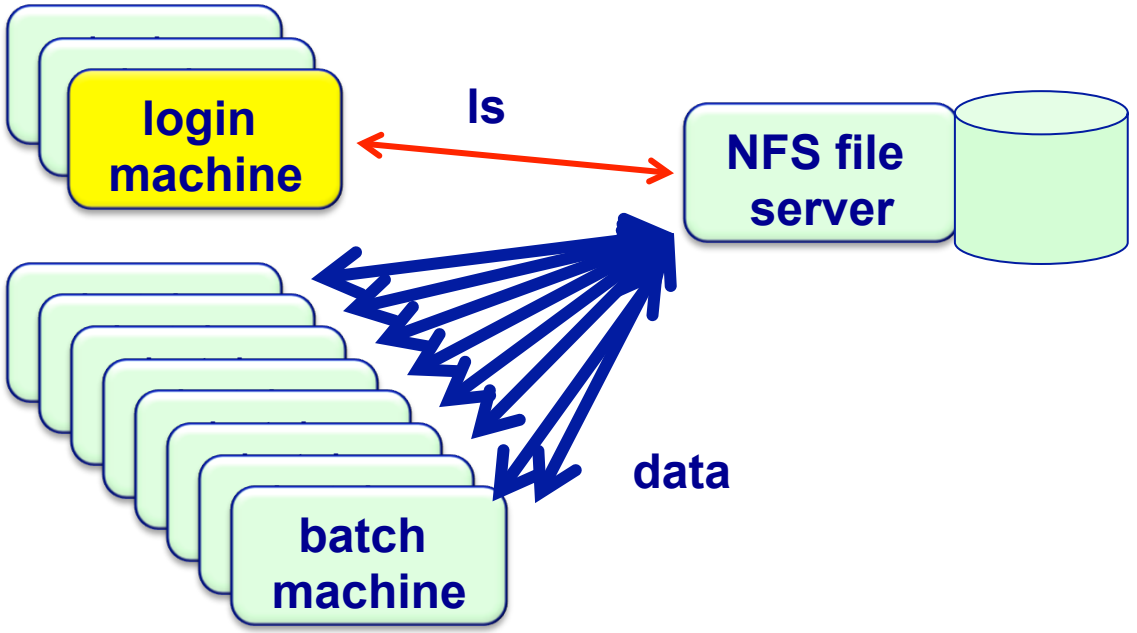

342 CPU cores

Time to run the ‘date’ command

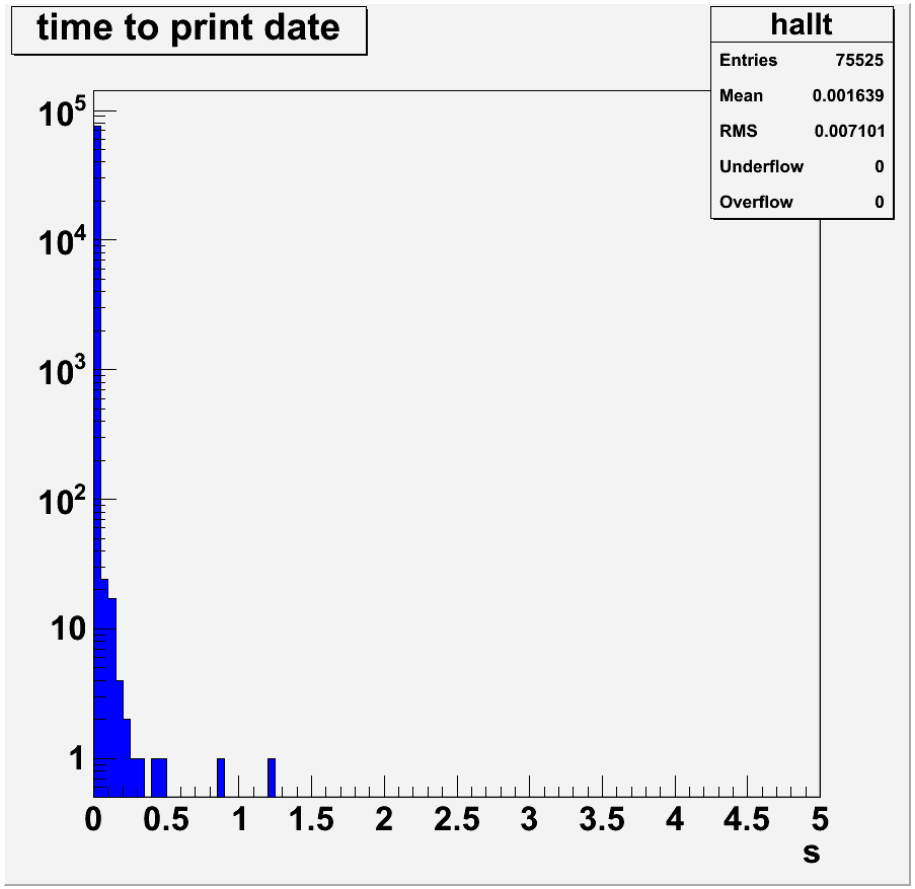

- The tail above 1s is at the 10^{-5} level.
- User experience would not be sensitive to it.
- We do experience waiting, in particular when the session was idle for some time.
- There must be something else going on.

Another possibility…

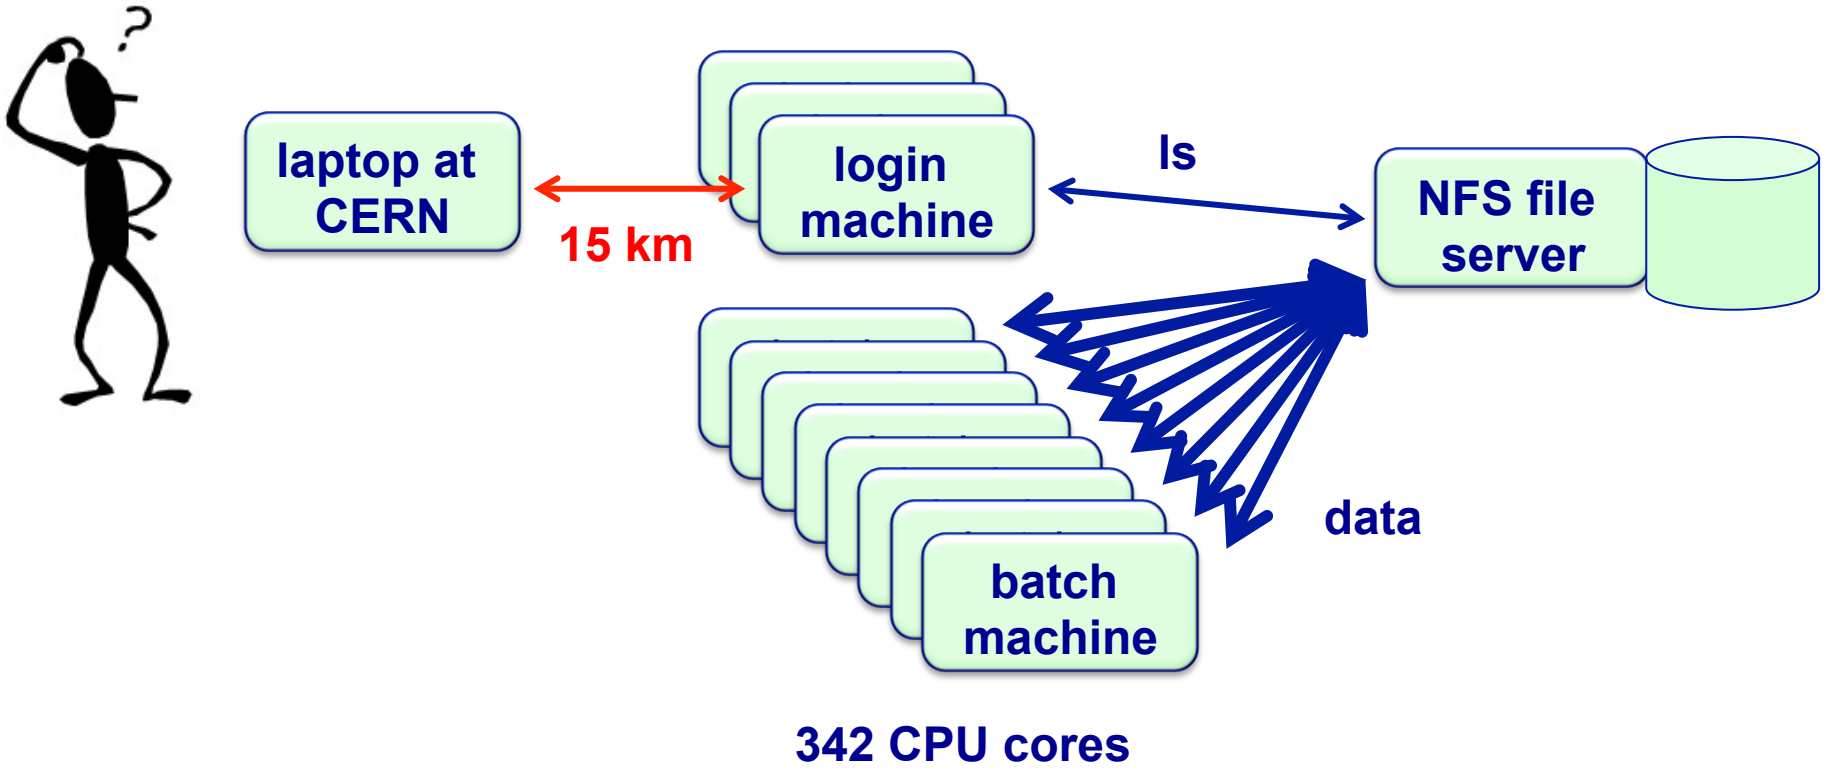

Network monitoring

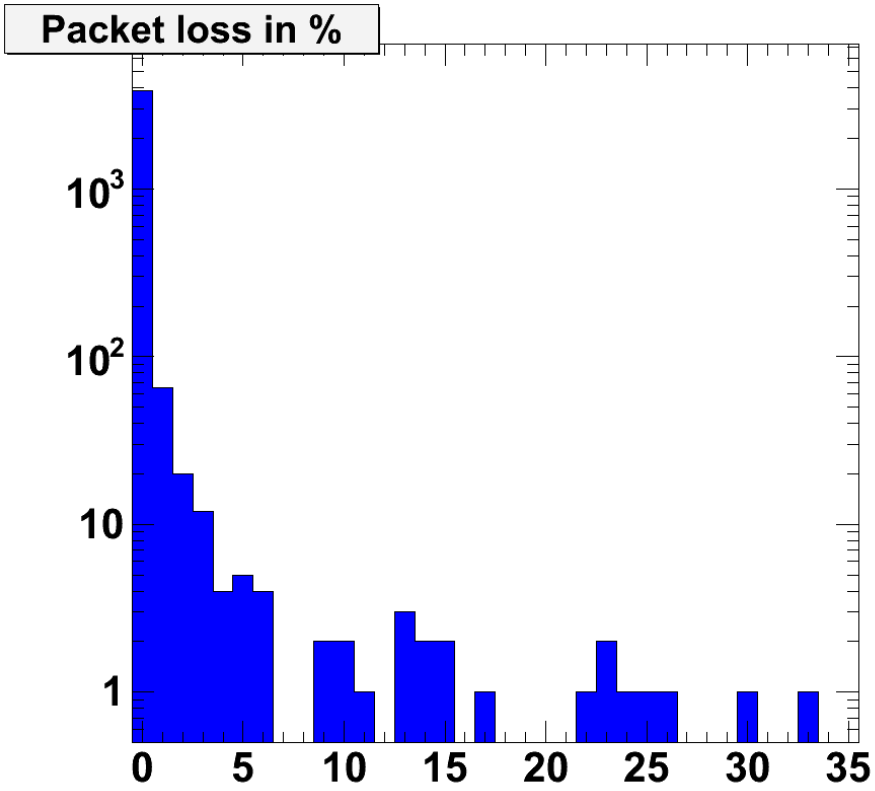

ping between my desktop at CERN and the cluster every 5’

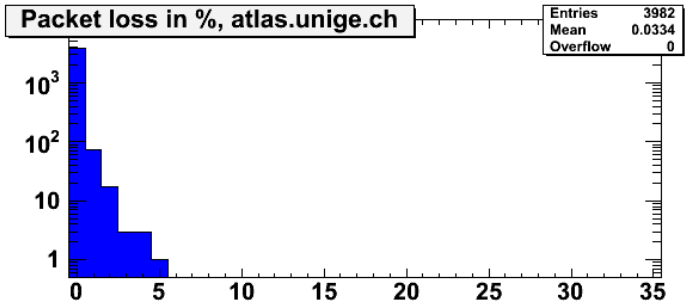

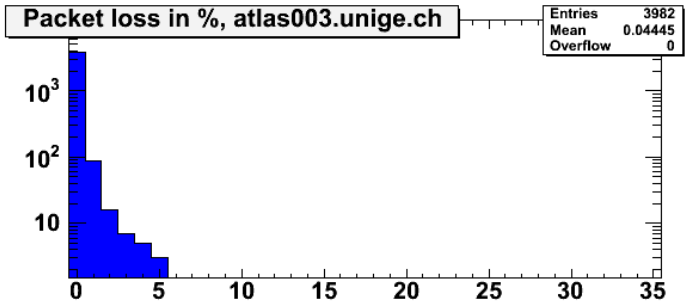

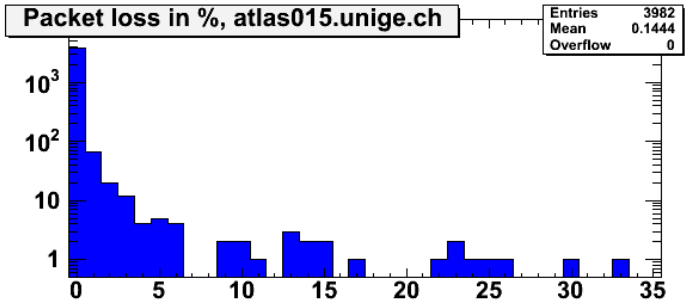

Network monitoring (2)

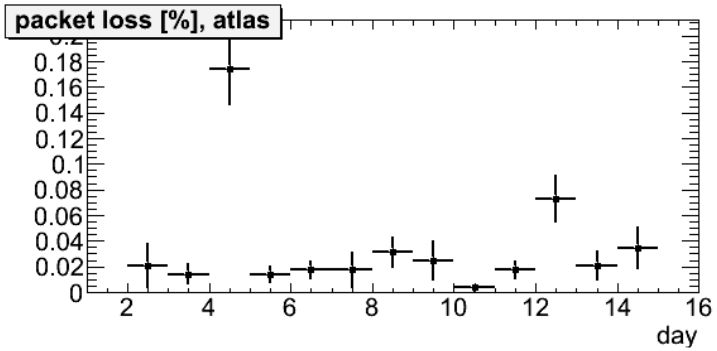

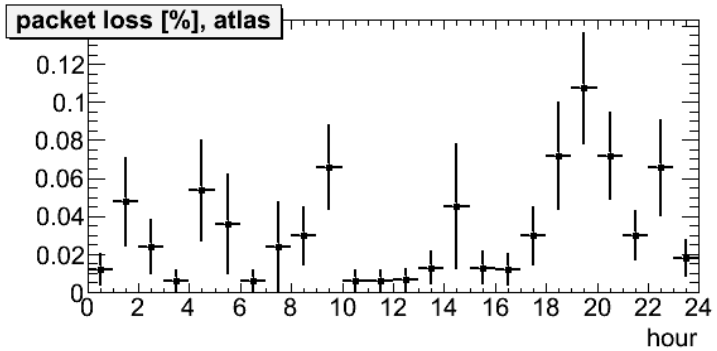

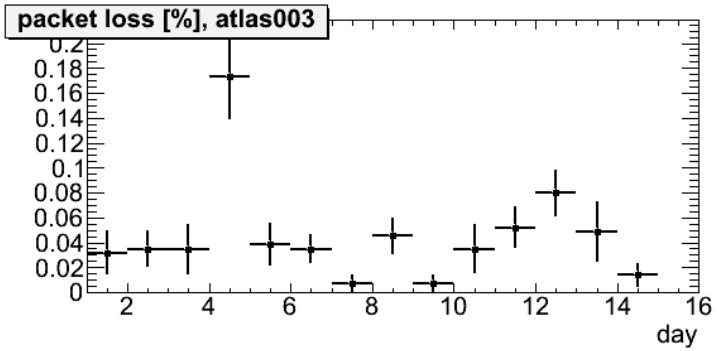

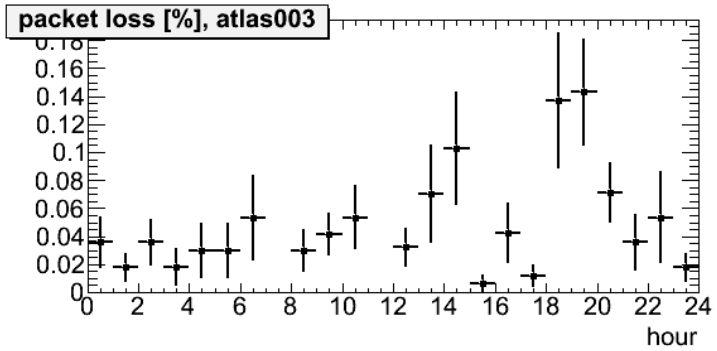

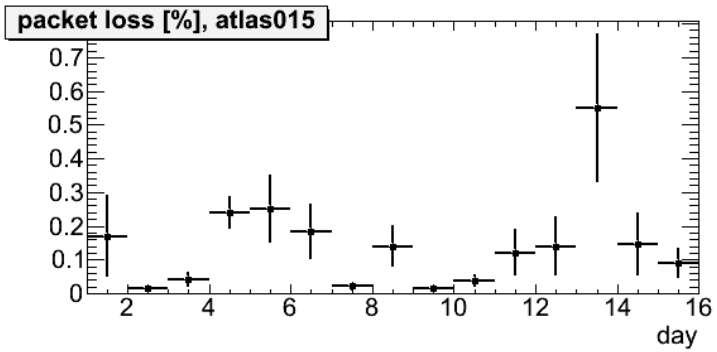

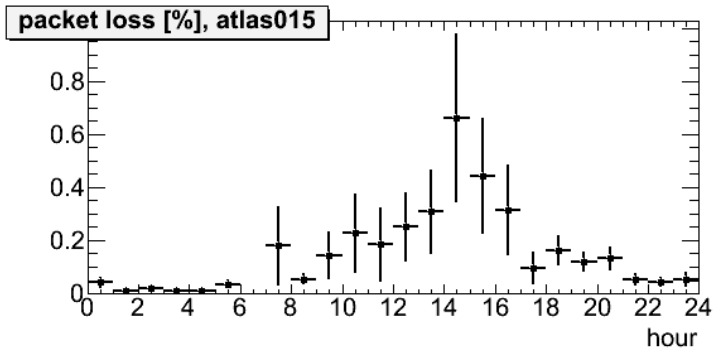

Summary

- From previous experience, we expected all clients having problems with the same server the same time.
  - Little evidence of this problem since early March.
- Occasional slow response limited to one client-server pair.
  - Other clients are OK, even talking to the same server.
  - Next step: catch it in real time and try a fix (unmount and mount).
  - It is rare. There was no chance last 2-3 weeks.
  - Keep the logging running.
- I suspect that is probably not all. Keep monitoring the network as welll.

Other news: PROOF

- Parallel ROOT facility.
- Interesting tests done by Allison on a 16-core batch machine.
- If more people are interested, more batch machines can be made login machines.
- Applicable only to CPU-bound processes!