ATLAS Report

NAF User Committee Face-to-Face Meeting — 12th January 2011

Marcello Barisonzi - University of Wuppertal
Wolfgang Ehrenfeld - DESY

Content

1. Introduction
2. Infrastructure
3. ATLAS
4. Summary

ATLAS NAF CPU Usage

- Current snapshot of ATLAS NAF batch usage
  - accounting interval: 1.8. - 29.11.2010
  - total CPU time: 21362 days
    --> ~12% of NAF CPUs
  - total wall clock time: 58228 days
    --> ~33% of NAF CPUs
  - nominal ATLAS share: ~25%
  - ATLAS is working on the NAF

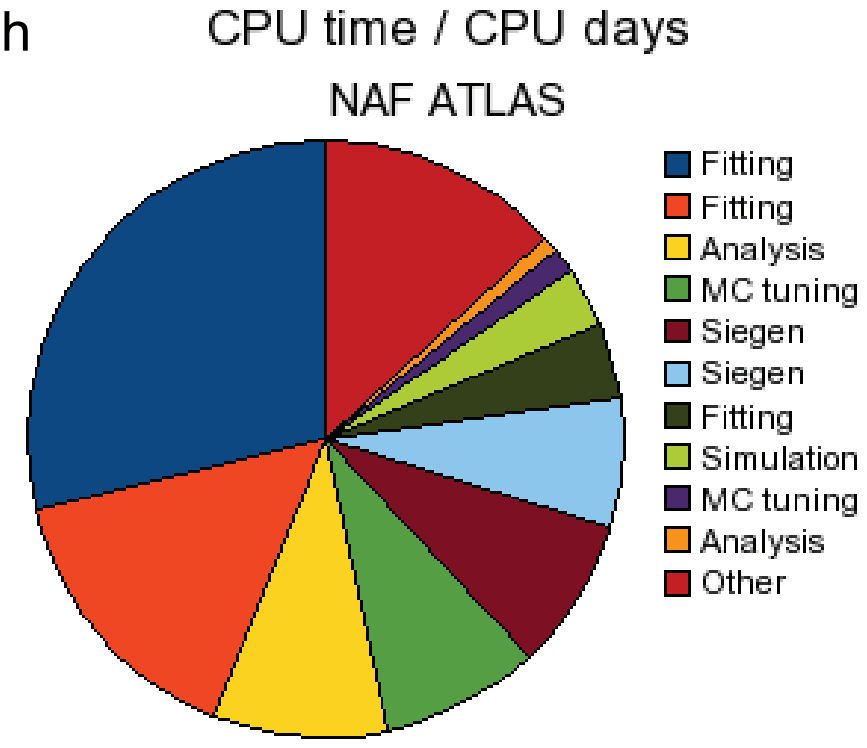

- Analysis type jobs are becoming more prominent within the group of power users
  - Other category includes ~100 users
  - As an example: September 2010 - 71 users total, 28 from DESY/HUB

Login and Load Balancing

- We had some problems there, are they all fixed?

Work Group Server

- ATLAS server have extreme high load sometimes (most kswapd problems?)
- NAF admin say that server have usually high load
  - we want to understand and solve this! How?
  - we provide dq2-get-naf (dq2-get on batch)
- problems with servers are usually triggered by ATLAS support or users $\rightarrow$ should be done by admins
- Request: Access to monitoring of uptime, CPU, memory, network (for ATLAS admins)
- Request: One more work group server in HH
- Request: ATLAS support should be informed about reboots (with brief diagnostics)

AFS

- No further comments on AFS instabilities (see NUM)
- resources are fine at the moment
- `naf_token` for Ubuntu and MacOSX support community effort?

Lustre

- usually full:
  HH: 90% ZN1: 95% ZN2: 10%
- current space: 135 TB
  HH: 79 TB, ZN1: 20 TB (until summer 2011), ZN2: 36 TB
- user quota not working (HH, ZN2)
- better administration needed:
  - user quota to quickly access usage
  - still interested into group quota for working groups
  - clean up (mail to all users are not working efficiently)
- access speed:
  - When will we use Infiniband in Hamburg for payload again?
  - What is the max expected Lustre throughput?
- stability

dCache

- For user analysis total turn around time is important for a given data set/ run period. This is dominated by
- DESY-HH_LOCALGROUPDISK or DESY-HH_GROUPDISK are often congested.
  - queued movers (due to ATLAS user jobs)
  - non-optimal file distribution over pools
- already in contact with dCache ops team to improve, but progress is slow
- provided bench mark tool to access sustained

dCache Test

1.5 TB from DESY-HH_LOCALGROUPDISK (dccp) in 76 jobs

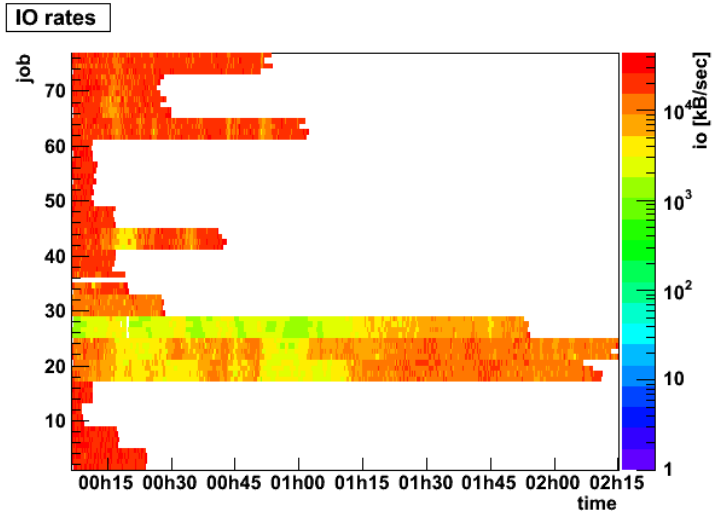

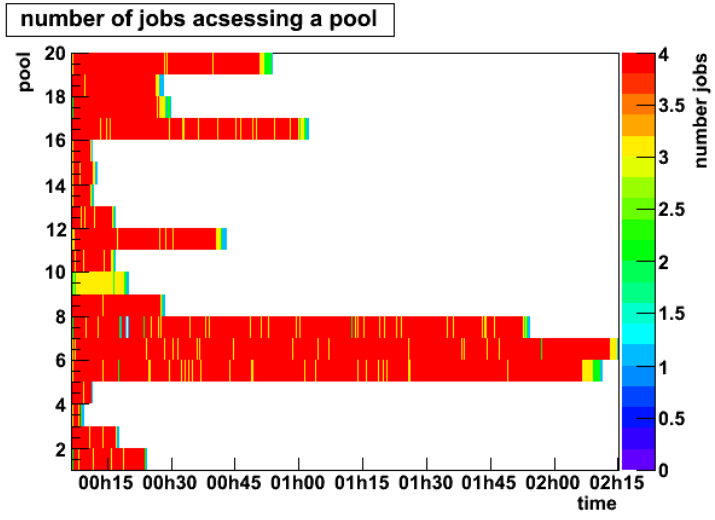

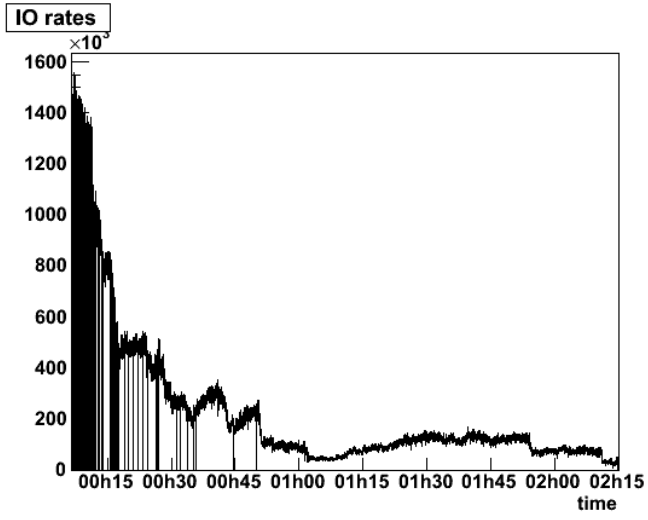

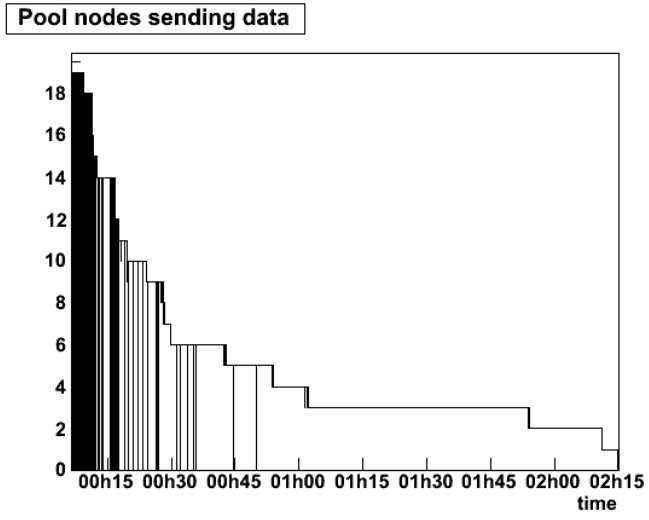

Support, Announcements etc.

- announcements are not as good as they should be (ATLAS SE problems from recent downtime)
- news section of naf.desy.de not updated timely
- What happened to NAF Twitter?
- Request tracker for support lists?
- Open list for support requests/questions (NUM)?
- Show sub system status on web (dashboard)
- Request: motd for ATLAS users (for new information and problems)

CMT

There are two options now:

- RPM installation of ATLAS software on work group servers
  - need some final tests
- ATLAS software via CVMFS
  - only needed files are caches on local client
  - uses local squid proxy for caching
  - will be used on Grid sites
  - setup need to be tested
  - performance (compile time) needs to be tested
  - no local tweaks possible
- monitoring of AFS volume reads would help to understand release usage

NX/VNC

Wikipedia:

> *NX technology is a computer program that handles remote X Window System connections, and attempts to greatly improve on the performance of the native X display protocol to the point that it can be usable over a slow link such as a dial-up modem.*

- NX will speed up remote X connections, this is needed for interactive ROOT/PROOF
- NX allows to export the full desktop but also single apps
- `http://www.nomachine.com/`
- VNC is a different approach to the same problem, it transfers part of the desktop as bitmap
- VNC sessions can easily be shared between users to allow interactive, collaborative work

This is a request to provide both. Details can and should be discussed.

Summary

- if the system works, the throughput is fine
- at some points usability need to be improved
- administration of the ATLAS resources needs to be improved