NAF ATLAS Report

NAF User Committee Meeting

11. November 2009

Wolfgang Ehrenfeld (DESY)
Jan Erik Sundermann (Universität Freiburg)

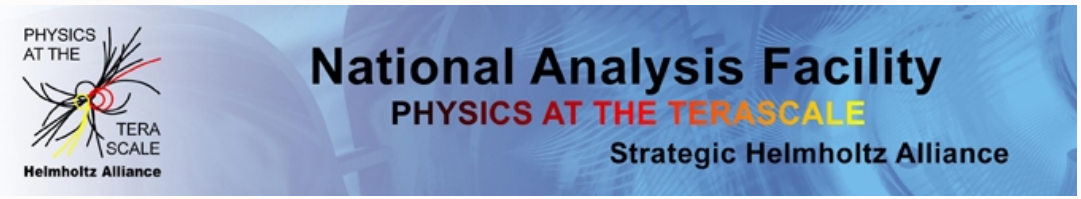

Outline

- Exemplary ATLAS use cases
- Open issues

Overview

- NAF usage during the last half year:

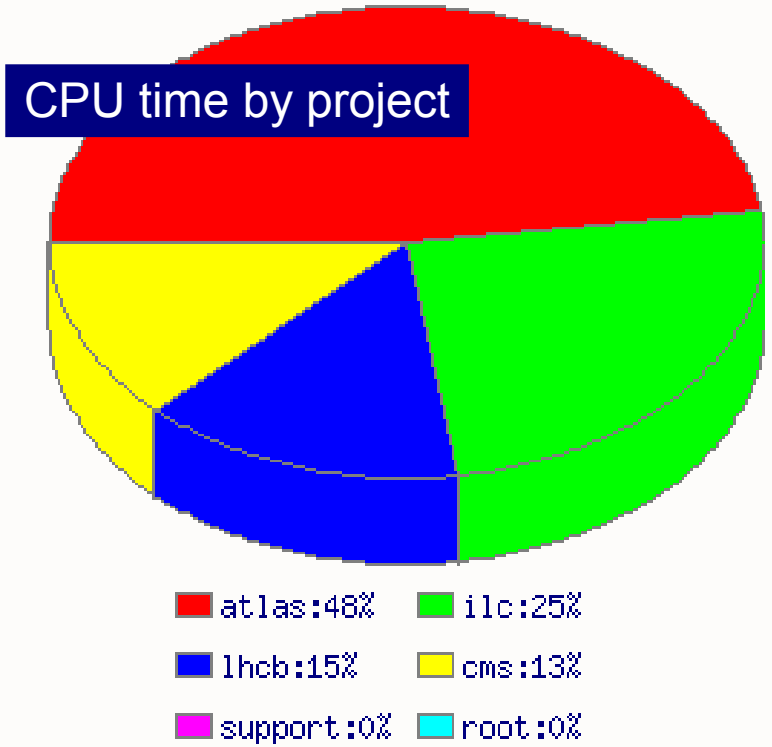

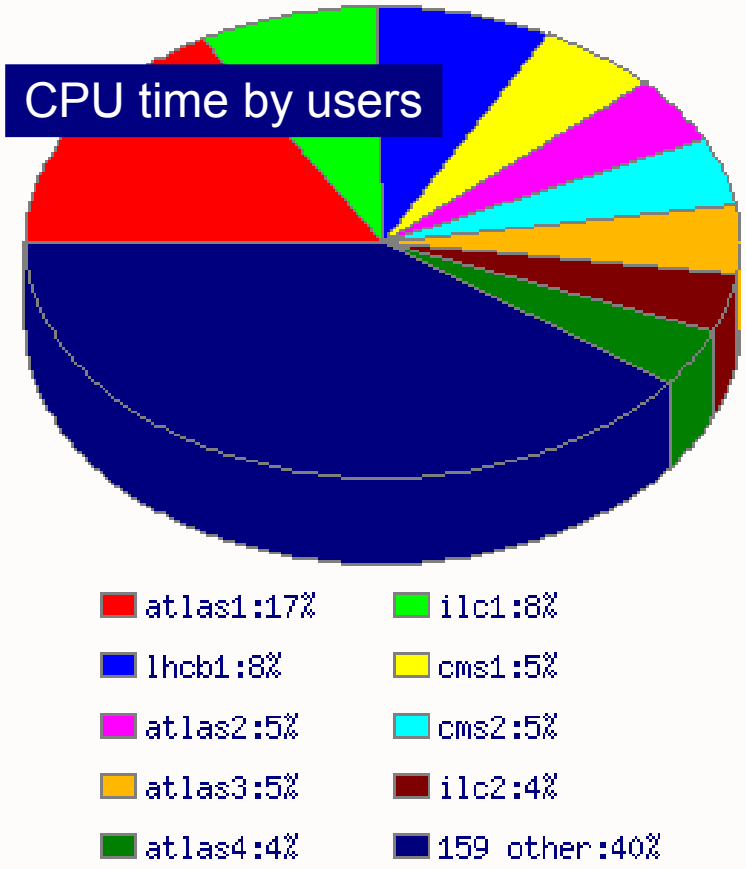

→ *Some ATLAS power users but in general larger number of users*

Exemplary Use Cases

- Examples for ATLAS usage of NAF resources:
  - User analysis using NAF grid and batch resources
  - Private Monte Carlo production
  - Monte-Carlo tuning using Tevatron data
  - Calculation of Higgs-production cross sections
  - Model fitting using Fittino

Use Case: User Analysis

- Analysis of officially produced ATLAS MC
  - Input data stored on NAF grid storage element
- Jobs submitted to the grid or local batch system
  - Using ganga or custom scripts for job submission
  - Input data copied to local disk
  - Jobs produce (flat) ntuples for subsequent analysis
- Output stored on the NAF grid storage element, NAF Lustre space or remote AFS space
- Output copied afterwards to the home institute
- Typical user problems with grid submission:
  - “Software too large for input sandbox”
  - “Failure rate higher than for local jobs”

Use Case: Monte-Carlo Tuning

- Tuning of ATLAS Monte-Carlo using Tevatron data:
  - Using Rivet toolkit installed by user
    - Some problems running SL4,32bit software on SL5,64bit WNs
    - Software currently installed on Lustre (too large for homes)
    - User would prefer central installation
  - Produces large number of MC events
  - Repeat Tevatron analysis on truth level
- Use case characteristics:
  - Almost no input data, only configuration
  - Very small output data (histograms) stored on Lustre
  - Jobs 100% CPU-limited

Use Case: Cross Section Calculations

- Calculation of Higgs-production cross sections in vector boson fusion
  - Calculations done for many different Higgs masses and beam-energies taking PDF uncertainties into account
- Use case characteristics:
  - Calculation quite slow due to electro weak corrections
  - 100% CPU-limited
  - Output very small, basically one number

Open Issues: Lustre

- Most of the time full (user do not clean up)
- Automatic clean up needed now (or sudo access)

- Some files should stay longer in /scratch
  - A: large space for ntuples and log files from production (should be removed after prod → automatic cleanup)
  - B: small space for code, histograms, ganga repository … (should stay longer)
- Separate instances? Quota? AFS? But now is the time!

Open Issues: Lustre/AFS

- Some of the use cases for smaller files with longer life time can be solved by AFS
  - What is the limit for the home directory?
  - What is the limit for user AFS scratch directory?
  - Can user AFS scratch directory be created automatically? Mounted into home?

Open Issues: Storage

- Concerning storage the NAF is two sites!
  - dCache: DESY-HH, DESY-ZN
  - Lustre: /scratch/hh, /scratch/zn
  - Need to optimize data to cpu matching in the future (batch system)

Other Open Issues

- cmt problem:
  - cmt is slow on multi user machines and AFS
  - User see this at atlas setup and compilation time
  - This is a major show stopper!
  - How can we solve this problem?
    - Copy of ATLAS software on Lustre or local?
- User notification:
  - News web-page does not contain information about planned down-times, software and operations are mixed
  - motd (or similar) would be useful

Summary and Conclusions

- Wide range of ATLAS applications ranging from typical user analysis on ATLAS data to CPU-intensive calculations
- NAF well accepted by users
- Feedback by users very positive
- The remaining open issues should be approach before LHC experiments start taking data