

# ATLAS NAF user report.

Some use cases, feedback, and requests from the DESY HH ATLAS group

**Karl-Johan Grahn, DESY**

**6th Annual Workshop of the Helmholtz Alliance "Physics at the Terascale"**

**4.12.2012**

## > Some use cases from the HH ATLAS group:

### > Analysis

- Top, SUSY, Standard model, Jet calibration
- ROOT-based analysis: SFrame, TopRootCore. Some Athena.
- Typically run on D3PDs stored on local group disk (HH or ZN).
- I/O driven: Large input, typically tens of terabytes for one analysis iteration.
- Thousands of jobs not uncommon.
- Need file-based scratch space for output. dCache/dq2 data sets not flexible enough.
- Output stored to Lustre/Sonatas. Order ~ 1 TB.
- Smaller subsequent jobs run with both input and output in Sonas. Either interactive or batch.

### > Event generation, theory calculations

- E.g., Alpgen
- No/small input
- Low(er) I/O, CPU intensive



- > With current NAF setup, input split between the Hamburg and Zeuthen sites.
- > Some analyses experienced problems running cross-site jobs (ZN dCache input on HH nodes or vice versa).
- > Input usually a lot larger output. Need to run where input is.
- > With the harder split of NAF 2.0, highly desirable to have all ATLAS resources in one place.
- > Awaiting NAF 2.0, Sonas mount in ZN would be desirable.

- > Asked for feedback within DESY HH ATLAS group. In general, people reasonably happy with current NAF performance.
- > Clear improvement seen in NAF Helpdesk response times and problem resolution.
- > [atlas-germany-naf-users](#) mailing list proven very useful for sharing info ("Am I the only one seeing this problem?").
- > AFS: No problems reported.
- > Sonas
  - In general, no complaints.
  - Migration from Lustre seems to have worked well.
  - General feeling: Faster and more stable than Lustre.
  - Contrary to Lustre, quotas in place. People have so far gotten the quotas they need.



- > How to handle runaway interactive jobs in WGSs?
- > Observation: Job congestion before major conferences. Probably hard to avoid...
- > Request: Schedule and notify of downtime well in advance (weeks).
- > Request: Not schedule downtime during run-up to main conferences (Moriond + summer).
- > Observation: Highly differing network throughput CERN->NAF between different WGSs.
- > Request: Direct access to CERN EOS possible?
- > Observation/request: DaTRI transfers from grid to local group disk sometimes slow (~days for a few TB). More monitoring?

# BACKUP