

# Göttingen Report

27.02.23

Sebastian Wozniowski

# Overview

- New site contact for ATLAS (me) taking over from Gen Kawamura
- Solved remaining network issues after GoeGrid-Cluster had been moved to new building
- Slurm had been replaced by HTCondor in 2021. Further optimized configuration of HTCondor.
- Supporting analysis of ATLAS data on tape at GridKa in the course of tape migration
- Disk firmware update on a 800TB disk array
- Set up monitoring dashboards for ATLAS Panda Jobs in ElasticSearch-Kibana instance hosted at CERN
- Other exceptional events: Two power cuts causing downtimes but luckily no loss of hardware.

# Resources

GoeGrid is a joint cluster: ATLAS Tier 2 and local Institute of Theoretical Physics

- Disk storage (ATLAS only): 2.65 PB pledged + 150 TB LocalGroupDisk + ~200 TB spare from old systems
- CPU: Total 17.000 slots (40% ATLAS, 60% theory)

Opportunistic usage of theorist's fraction allowed for a significant excess in CPU power provided to ATLAS over the past year, as soon as proper fairshare mechanisms were implemented.

=> Pledged in 2022: 317,640 kHS06\*h; delivered: 1,019,417 kHS06\*h

Drawback: Theorist's requirements for scratch disk on worker nodes << ATLAS requirements.

=> ATLAS jobs do not reserve disk and disk may fill up in case of strong opportunistic usage.

Issue came up occasionally and required tuning of submit policy (requiring minimum free disk at submit time), trying to optimize throughput of successful jobs (e.g. currently again).

# Analysis of ATLAS datasets on GridKa tape

Created statistical overview of datasets and files by data type and file size  
in order to help Haykuhi with optimally placing data on new tape system at GridKa.

# Disk firmware update

One disk array had noticeably many disk failures in the past. According to vendor related to bad disk firmware version  
=> Recommended update and provided tool for that.

Due to problems with the tool, the operation extended over three months. Thanks to ATLAS for patience and reduction of cached data at GoeGrid!

Update is investment in post-warranty period of the system when disks are not replaced anymore.

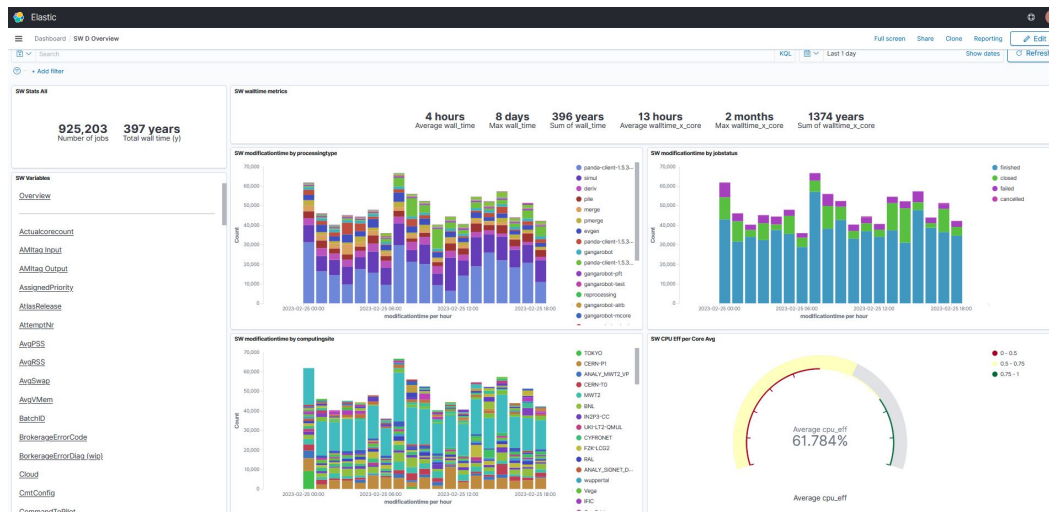
# ATLAS Panda Job Monitoring

Set up a collection of 128 dashboards + corresponding data flow from central database to ES/Kibana instance.

- Monitoring of Panda Job properties and runtime statistics
- Work still ongoing and expanding to JEDI tasks

Supports central ATLAS DDM but also our site contacts and can help identifying issues and correlations.

Got recently contacted by Ben (DESY) and Rod (Munich) about adding dashboards on CO<sub>2</sub> consumption of sites and am happy to join efforts.



# Outlook towards HPC usage

In the FIDIUM project, our group is working together with representatives from the local HLRN + NHR HPC cluster Emmy and GWDG in order to establish a connection between Emmy and GoeGrid that allows us to forward ATLAS jobs to Emmy reading data from GoeGrid dCache.

Recently handed in NHR application at percent level of current pledge for development and test purposes.

Hoping to profit from opportunistic usage soon and being ready for a regular usage in the next funding period.