

OPEN ACCESS

The German National Analysis Facility as a tool for ATLAS analyses

To cite this article: W Ehrenfeld *et al* 2011 *J. Phys.: Conf. Ser.* **331** 072053

View the [article online](#) for updates and enhancements.

You may also like

- [The DESY Grid Centre](#)
A Haupt, A Gellrich, Y Kemp et al.
- [Logical OR operation and magnetic field sensing based on layered topology](#)
Junyang Sui, Dan Zhang and Haifeng Zhang
- [Improvement of High Voltage Cycling Performances of \$\text{Li}\[\text{Ni}_{1/3}\text{Co}_{1/3}\text{Mn}_{1/3}\]\text{O}_2\$ at 55 °C by a \$\(\text{NH}_4\)_2\text{AlF}_6\$ Coating](#)
Yang-Kook Sun, Seung-Taek Myung, Chong Seung Yoon et al.



ECS
The
Electrochemical
Society
Advancing solid state &
electrochemical science & technology

DISCOVER
how sustainability
intersects with
electrochemistry & solid
state science research

The German National Analysis Facility as a tool for ATLAS analyses

W. Ehrenfeld^a, K. Leffhalm^a, S. Mehlhase^{1a}

^a*Deutsches Elektronen-Synchrotron DESY, Hamburg and Zeuthen, Germany*

Abstract

In 2008 the German National Analysis Facility (NAF) at DESY was established. It is attached to and builds on top of DESY Grid infrastructure. The facility is designed to provide the best possible analysis infrastructure for high energy particle physics of the ATLAS, CMS, LHCb and ILC experiments.

The Grid and local infrastructure of the NAF is reviewed with a focus on the ATLAS part. Both parts include large scale storage and a batch system. Emphasis is put on ATLAS specific customisation and utilisation of the NAF. This refers not only to the NAF components but also to the different components of the ATLAS analysis framework.

Experience from operating and supporting ATLAS users on the NAF is presented in this paper. The ATLAS usage of the different components are shown including some typical use cases of user analysis. Finally, the question is addressed, if the design of the NAF meets the ATLAS expectations for efficient data analysis in the era of LHC data taking.

¹contact: sascha.mehlhase@desy.de

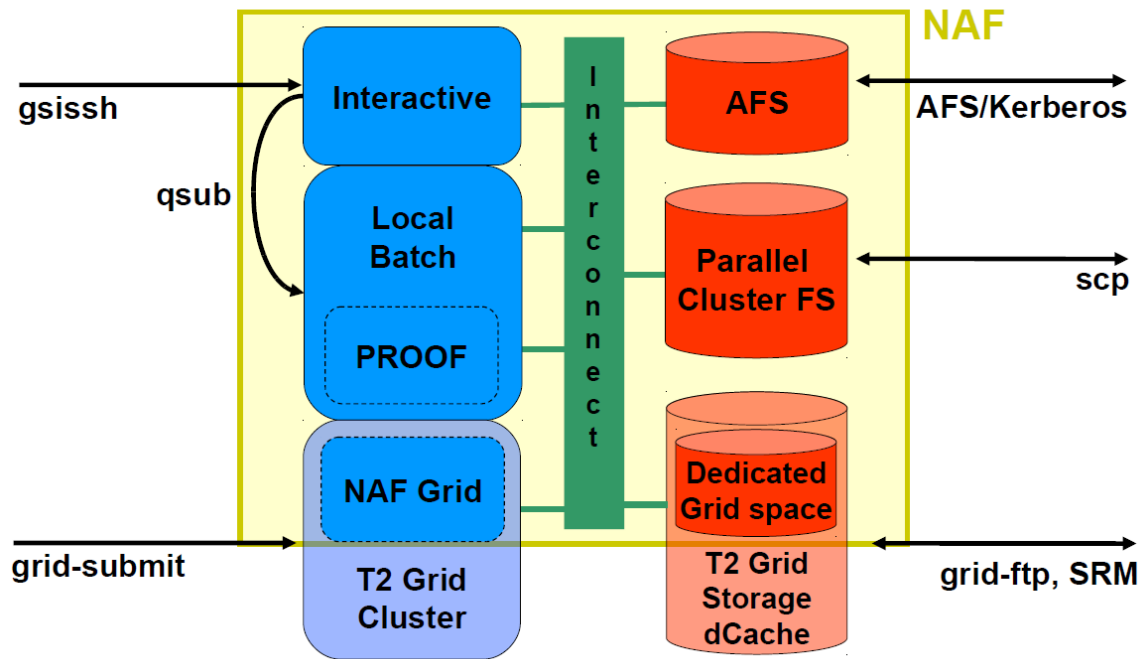


Figure 1: Illustration of the general NAF infrastructure.

1 Introduction

This paper gives an overview on the use of the German National Analysis Facility (NAF) as a tool for ATLAS analyses. In Section 2 the general structure and building blocks of the NAF are introduced followed by a rough evaluation of the usage statistics in Section 3. The availability of data and Monte Carlo (MC) samples at the NAF is discussed in Sections 4 and 5, respectively. The paper concludes with a selection of exemplary use cases of ATLAS analyses on the NAF in Section 6 and an operations review in Section 7.

2 The National Analysis Facility

The NAF is set up within the framework of the Helmholtz Alliance: “Physics at the Terascale”. It is designed as an analysis facility for experiments both at the Large Hadron Collider (LHC) and a possible International Linear Collider (ILC). Its services are offered exclusively to physicists of German particle physics institutes within the Terascale Alliance.

The NAF infrastructure is build and operated jointly by both sites of the Deutsches Elektronen-Synchrotron (DESY) in Hamburg (HH) and Zeuthen (ZN). Besides interactive sessions on dedicated work group servers for each experiment, the NAF offers a SGE batch system and dedicated Grid Computing Elements for its users, along with AFS, Lustre and dedicated Grid storage both in HH and ZN. More details about the general infrastructure of the NAF, illustrated in Figure 1, can be found e.g. in [1, 2].

The German ATLAS community has in total four work group servers (for interactive sessions) and a share of 25% of the 1400 cores in the SGE batch system. In addition a high priority batch queue allows for short turn-around times in exceptional cases.




Figure 2: ATLAS NAF CPU time for power users exceeding 300 CPU days in 2010. The different German institutes are encoded by a number in the user name. Different users from the same institute are encoded by a letter.

3 Usage Statistics

The NAF currently (Oct 2010) supports about 500 users in total, with about 300 from the ATLAS community. About 100 ATLAS users can be considered active, having logins within the last 30 days. The 1170 TB of available data storage are composed of 740 TB of Tier 2 Grid storage, 340 TB NAF Grid storage and 90 TB of local storage. Each user is provided with an AFS home directory (the size is between one and five GB) and some AFS scratch space (the size is between one to ten GB). On average the users utilise about 195 MB of the former and about 620 MB of the latter. Figure 2 shows the CPU usage (wall-time) for various ATLAS users along with their job type.

4 Data Availability

ATLAS physics data as well as important Monte Carlo samples are provided centrally for all ATLAS NAF users, while further storage is provided for individual and group data. The full ATLAS AOD data set is available at the NAF, using official ATLAS space tokens at the DESY Tier 2 extended by dedicated NAF storage. LOCALGROUPDISK space tokens are provided as local Grid storage for analysis data (e.g. ntuples) and archive (e.g. older AOD versions) purposes, while all Grid storage is included into the common ATLAS distributed data management.

Via dCache the data is directly accessible not only for Grid jobs but also from the work group servers and the batch system. In addition fast local storage is provided using the Lustre file system, offering short access times for large files and short analysis turn-around.

5 ATLAS Analysis Support

To allow for a broad range of ATLAS analyses with up-to-date tools, both common and ATLAS specific software are installed centrally for all users at the NAF. Common software; such as ROOT, Python, a Latex environment and AutoProxy; and Grid tools; such as Panda and Ganga clients, DQ2 tools; are provided by DESY-IT in a (semi)automatic fashion. These tools can be accessed directly or using simple and automated setup scripts. The deployment of ATLAS software and tools, ranging from a wide selection of ATLAS releases and caches to flat ATLAS condition data taken from the DESY Tier 2, is also performed in a (semi)automatic way and no individual user setup is necessary.

Besides the plain supply of software and tools, users are supported in all ATLAS software NAF related questions and are forwarded to the appropriate ATLAS help forums for non NAF specific issues.

6 ATLAS Analysis Use Cases

In this section four exemplary analysis use cases, depicted in Figure 3, are described in more detail.

6.1 N-tuple production on the Grid using Ganga job submission (I)

With direct access to general and dedicated Grid resources, production jobs are submitted using local installations of Ganga and Panda. Skimmed/slimmed/thinned data sets (D2PDs) or flat n-tuples (D3PDs) are created using AODs or ESDs as input and are stored at the local Grid Storage Elements. These derived format allow for shorted turnaround times in final analyses and can be accessed directly via dCache on all NAF systems.

6.2 SGE batch based analysis using data from Grid storage element via dCache (II)

Medium sized jobs, e.g. production of ntuples from derived formats stored locally on the Grid Storage Elements, can be performed with short turn-around times using the batch system. The output can be stored e.g. on scratch space available on AFS or Lustre. Job submission and monitoring can done with standard SGE tools or Ganga. If very fast turn-around is needed, jobs can be submitted to the ATLAS high priority batch queue. This is useful before conferences.

6.3 PROOF cluster on demand on a batch system (II*)

Support of an easy to use and simple script to setup a fully working PROOF cluster on demand for a user. The script starts a PROOF master on the local work group server and submits multiple PROOF clients to the batch system. PROOF can be used for data analysis based directly on TSelector or TSelector based frameworks as SFrame. More details can be found in [3].

6.4 Interactive analysis on a work group server using fast Lustre storage (III)

Using the fast Lustre storage element, end user analyses and the production of final plots can be carried out using interactive sessions on the dedicated ATLAS work group servers. These provide all tools necessary, such as ROOT, editors, LaTeX and more.

7 Operations Review

The German ATLAS NAF team has collected extensive experiences in supporting the German ATLAS users during the first three years of NAF operations.




Figure 3: Illustration of possible analysis use cases on the NAF. Depicted is the work flow and data access for four different scenarios: I) using a work group server to submit job to the grid using data on and writing results to the grid storage element; II) using a work group server to send jobs to the local batch farm using data from Lustre and writing results to AFS; II*) same as before but using a Proof setup; and III) running jobs directly on a work group server using data from Lustre and writing to AFS.

- It has been shown that software installation can easily be done via scripts and automated by cron jobs, while support of experiment software for forward compatibility (operating system) is time consuming and does not always work out of the box. A close collaboration with the ATLAS software librarians is needed to solve problems in a timely manner.
- AFS storage management is easy, using tools such as `afs.admin`. Local storage is always tight and users are typically not cleaning up on their own. As a result of this, automatic cleanup, preferentially based on access time, is considered a possible improvement.
- User Grid storage is almost unlimited, while efficient space administration and management is only possible within the ATLAS distributed data management system.
- User support is time consuming.

8 Summary and Outlook

As one of the first communities, the German ATLAS user group has been using the National Analysis Facility for almost three years now. The strong support from both DESY sites, the availability of data sets ranging from official data and MC AODs to privately produced samples, the centrally provided ATLAS software and finally the possibility to choose between various analysis work flows have made the NAF a useful tool for a large and still growing number of German ATLAS users. It has been shown that both ATLAS specific software installations and storage management can be fairly well automatised and covered by a reasonable amount of people, while software and especially user support take up a decisive amount of time. Even though there is room for improvements, the NAF has proven itself to meet the expectations for efficient data analysis in the era of LHC data taking.

References

- [1] Haupt A and Kemp Y 2010 *J. Phys. Conf. Ser.* **219** 052007
- [2] Aplin S, Ehrenfeld W, Haupt A, Kemp Y, Langenbruch C, Leffhalm K, Lucaci-Timoce A and Stadie H 2010 The National Analysis Facility at DESY: Status and Use Cases by the Participating Experiments [these proceedings]
- [3] Behrenhoff W, Ehrenfeld W, Samson J and Stadie H 2010 PROOF on a Batch System [these proceedings]