

Integration of the Goettingen HPC cluster Emmy into the WLCG Tier-2 centre GoeGrid and performance tests

Inga Łakomic¹, Saidev Polisetty^{1,*}, Arnulf Quadt¹, Daniel Schindler¹, and Sebastian Wozniowski¹

¹II. Physikalisches Institut, Georg-August-Universitaet Goettingen, Friedrich-Hund-Platz 1, 37077, Goettingen, Germany

Abstract. The German university based WLCG Tier-2 centres successfully contributed a significant fraction of computing power for the Runs 1-3 of the LHC. But for the upcoming Run 4, with its increased need for computing resources for the various HEP computing tasks in terms of storage and computing power, there is a necessity for a transition to a new model. In this context, the German community under the FIDIUM project is making interdisciplinary resources of the National High Performance Computing (NHR) usable within the WLCG and centralising mass storage at the Helmholtz centres. At the Goettingen campus, there is both a WLCG Tier-2 site (GoeGrid) and a large HPC cluster Emmy by NHR and the North German Supercomputing Alliance (HLRN). The integration is done by virtually extending the GoeGrid batch system with containers, turning the HPC nodes into virtual worker nodes with their own partitionable job scheduling in order to run the HEP jobs which serve the ATLAS collaboration in the case of GoeGrid. Submission and management of these containers is automated using COBaD (the Opportunistic Balancing Daemon) and TARDIS (The Transparent Adaptive Resource Dynamic Integration System). Data is provided via the GoeGrid mass storage for which a dedicated network connection has been established. Continuous production of the ATLAS jobs is currently tested in a one-year pilot phase. The setup, experiences, performance tests and an outlook are presented.

1 Introduction

The usage of compute resources for data analysis and storage increased exponentially in the field of scientific research with advancement of the experiments. The Large Hadron Collider (LHC) experiments are examples of such mega-scaled experiments in the field of high energy physics (HEP) where an enormous amount of data is generated and distributed around the world for research via the Worldwide LHC Computing Grid (WLCG) sites [1].

Currently, the WLCG Tier-2 centres in Germany are run at DESY (Deutsches Elektronen-Synchrotron), MPP (Max Planck for Physics) and universities. In the context of the upgrade to the HL-LHC (High Luminosity - Large Hadron Collider) and its large computing requirements, the fraction of Tier-2 compute power provided currently by universities is

*e-mail: u.polisetty@uni-goettingen.de

being moved to the National High Performance (NHR) computing sites in Germany. This aids in the provisioning of sufficient and sustainable compute power with efficient energy consumption, proficient usage of the compute resources and co-operative collaboration. This is achieved through the R&D project, "Federated digital infrastructure for research in universe and matter" (FIDIUM), funded by the Federal Ministry of Education and Research (BMBF) of Germany.

1.1 The FIDIUM project

In the FIDIUM project, research and development work is classified in three areas:

1. Development of tools for the integration of heterogeneous resources.
2. Data lakes, distributed data, caching.
3. Customisation, testing and optimisation on production and analysis environments.

All the three work areas focus on different topics which are interlinked and aim to provide a setup to run HEP jobs. Work area 1 focuses on utilising external computing sites that are not part of the LHC computing grid as compute resources for HEP research. Work area 2 deals with mass storage. Storage is moved from the university Tier-2 sites to DESY and KIT. Currently implemented methods for data access are tested. Alternate methods are also developed and validated. Work area 3 is dependent on the other two work areas. The primary goal of work area 3 is to test the methods developed in the other two work areas in realistic environments, and customisation for optimal performance.

2 Integration of the compute resources at Goettingen

2.1 WLCG Tier-2 and NHR in Goettingen

The WLCG Tier-2 and NHR sites are spread across the country. Goettingen is one of the rare cases where there is a WLCG Tier-2 centre, GoeGrid, and an NHR site, Emmy (HPC - High Performance Computing) in the same city. GoeGrid is a dedicated WLCG Tier-2 site providing 3.6 PB of dCache [2] storage and a large share of its 22,000 logical cores for the ATLAS [3] collaboration. Emmy is an HPC cluster that is a federal resource with over 140,000 CPU and 8400 GPU cores.

2.2 Integrated cluster layout

The layout for the job flow at GoeGrid is described as follows. Jobs submitted to the Production and Distributed Analysis (PanDA) [4] workload management system for ATLAS are sent to GoeGrid, which runs them on an HTCondor batch system [5]. HTCondor manages and schedules the jobs received from the PanDA queue on the worker nodes according to the job requirements. The jobs have access to the GoeGrid dCache storage pool. The input data are transferred to the jobs, and the job results written out at the end of job execution.

This setup in GoeGrid is virtually extended into Emmy by using containers. A container (Singularity/Apptainer) [6] is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another [7]. These container environments where the jobs run on Emmy are called

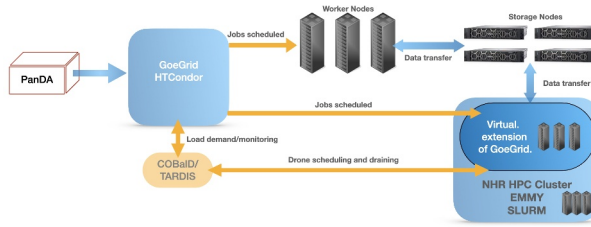


Figure 1. The complete layout of the setup at Goettingen, integrating NHR Emmy with Tier-2 GoeGrid. The layout depicts the job flow and the data access in both clusters.

"drones". There is a dedicated omni-path network connection between GoeGrid and Emmy for running these drones for HEP jobs (4x100 Gbit/sec). This connection is only accessible for the drones sent from GoeGrid to Emmy. The launch of the drones is automated by using the COBalD/TARDIS (CT) resource manager [11]. Drones are scheduled, launched and killed/drained with the CT resource manager. Emmy uses the Slurm batch system [10] rather than HTCondor, and Slurm is what launches the drones on Emmy. Inside the drones, HTCondor runs the actual jobs, and this is important for Emmy because only whole-node reservation is possible, and HTCondor can take care of partitioning and allocating the slots required by the jobs.

CVMFS (CernVM File System) [8], the repository source containing all the required software for the running jobs, requires FUSE (Filesystem in Userspace) permissions to install on Emmy. This is avoided by installing the unprivileged user CVMFS package "cvmfsexec" on the nodes. This also helps operators of drones in retaining full control of cvmfs to troubleshoot in case of errors. The completed layout of the integrated setup is shown in Fig. 1.

2.3 CPU Benchmarks

The HEPscore23 (HS23) [9] benchmark is used to obtain a standard reference score for the performance of the processors in running HEP jobs. It runs various typical HEP jobs that are used for analysis in various experiments to determine the performance of the computing clusters in units of HS23. The higher the number, the better the performance. GoeGrid employs various processors. The majority are AMD Epyc 7 series. The Benchmark numbers are as follows:

- GoeGrid: 2.4k HS23 on 256 logical cores.

Emmy upgrades processors in phases using Intel Cascade Lakes for Phase 2 and Intel Sapphire Rapids for Phase 3 (the most recent phase). The Phase 2 nodes are installed with shared SSD and HDD disks.

- Emmy Phase 2: 1.9k HS23 on 192 logical cores.
- Emmy Phase 3: 2.9k HS23 on 192 logical cores.

All the benchmarks are taken on the processors with hyper threading enabled.

3 Emmy performance tests

The final integrated setup is tested to analyse the job performance in the cluster, monitored for any limitations of the containerised setup and to optimise for better performance. These tests provide a deeper insight into the performance of the containerised setup. Emmy uses different types of storage and filesystems for its range of nodes. They ultimately have an effect on the CPU performance in running the HEP jobs. The testing of these various types helps us to determine the optimal resources to use for the experiments. The cluster is also tested with heavier job loads to monitor scalability of nodes. The tests comprise various ATLAS job types, including standard event generation, data processing and user analysis, which is a mix of jobs of low and high IO intensity. The jobs themselves run for a maximum period of 4 days, whereas the environment they run in (drones) have a lifetime of 7 days.

3.1 CPU Efficiency: Shared SSD vs Shared HDD

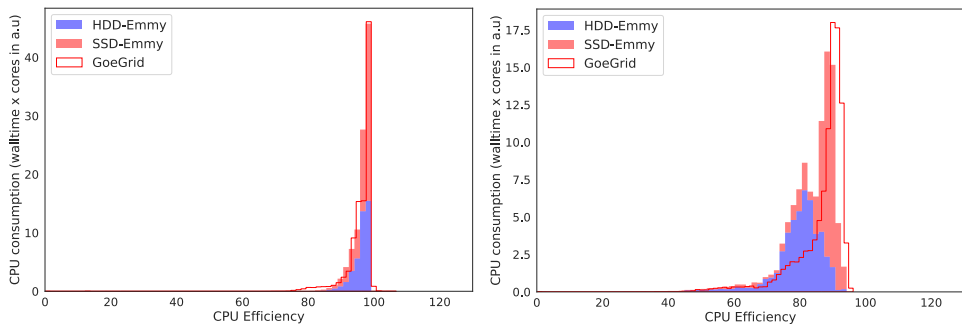


Figure 2. CPU consumption of ATLAS jobs; simulation (left) and pile up (right) shown against the efficiency in the two computing sites. Emmy Phase-2 shared disks are shown separately for SSD and HDD, with the two plots stacked.

The Phase 2 nodes of Emmy employ both SSDs and HDDs with a lustre filesystem. Multiple factors affect the CPU performance. In the lustre filesystem, storage is shared between LHC jobs and non-LHC jobs. When a large number of non-LHC jobs are running, this can increase the waiting time for the LHC jobs to read/write data, affecting the CPU efficiency. The type of storage disk used also affects the reading/writing speeds. SSDs are faster than HDDs, which makes a difference when running high IO jobs. This is observed as expected in Phase 2 nodes. Figure 2 shows CPU consumption with respect to CPU efficiency of pileup jobs which are IO intensive jobs. A clear difference between the performances of SSDs and HDDs in Emmy is observed. GoeGrid and Emmy SSDs have comparable performance. The effect of the storage is seen clearly in pileup jobs. While SSDs of Emmy are comparable to GoeGrid with average efficiencies 86% for both, HDDs show worse performance with 76% in a realistic case. In contrast the simulation job type, which is not IO intensive, performed similarly in both SSDs and HDDs and quite comparable to GoeGrid as well. They have similar average efficiencies (95% for Emmy and 96% for GoeGrid).

Jobs run at CPU efficiencies around 90% for all storage types and SSD performs better than HDD as expected for the overall mix of jobs, especially for IO intensive jobs. A small

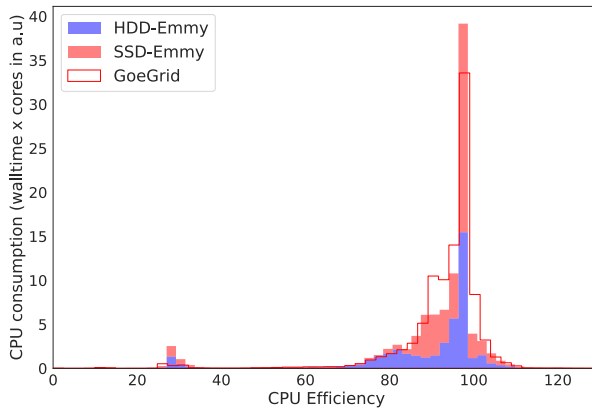


Figure 3. CPU consumption of GoeGrid and Emmy are shown against the CPU efficiency of all ATLAS job types. Emmy Phase-2 shared disks has been segregated to SSD and HDD stacked plots.

peak observed at 30% is due to smaller stuck test jobs, while efficiencies above 100% seen in Fig. 3 are due to event generation (evgen) jobs which used more cores than booked.

3.2 Scaling tests

Scaling tests are done using Phase 2 and 3 nodes to study the effects on the performance and their sustainability on different heavy loads. Gradually increasing loads are tested, each for a period of two days, incrementing from 50 drones running 9600 logical cores up to 100 drones running 19,200 logical cores. The performance is studied by observing the consumed CPU resources of the ATLAS jobs against the CPU efficiency and compared to GoeGrid.

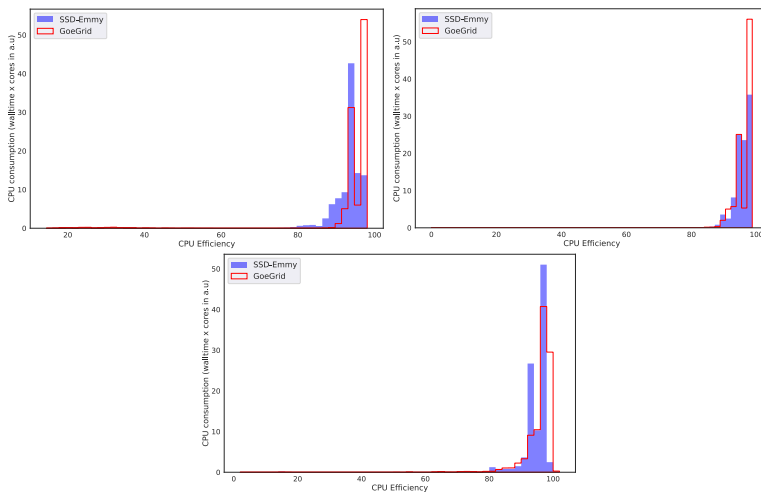


Figure 4. CPU consumption against the CPU efficiency is shown for ATLAS simulation jobs on a load scaling from 9600 cores (top left) to 15,360 cores (top right) and 19,200 cores (bottom).

Jobs from the ATLAS job types that are least IO intensive jobs such as simulation have a sustainable constant performance with minor fluctuations over the increased load. They run with $\sim 95\%$ average efficiencies and use most of the CPU resources as seen in Fig. 4. With all the ATLAS job types combined, including high IO jobs, deviation in efficiency is very minimal. It is still comparable to GoeGrid and the cluster withstands heavy loads well. The average CPU efficiency is almost constant as observed in Fig. 5.

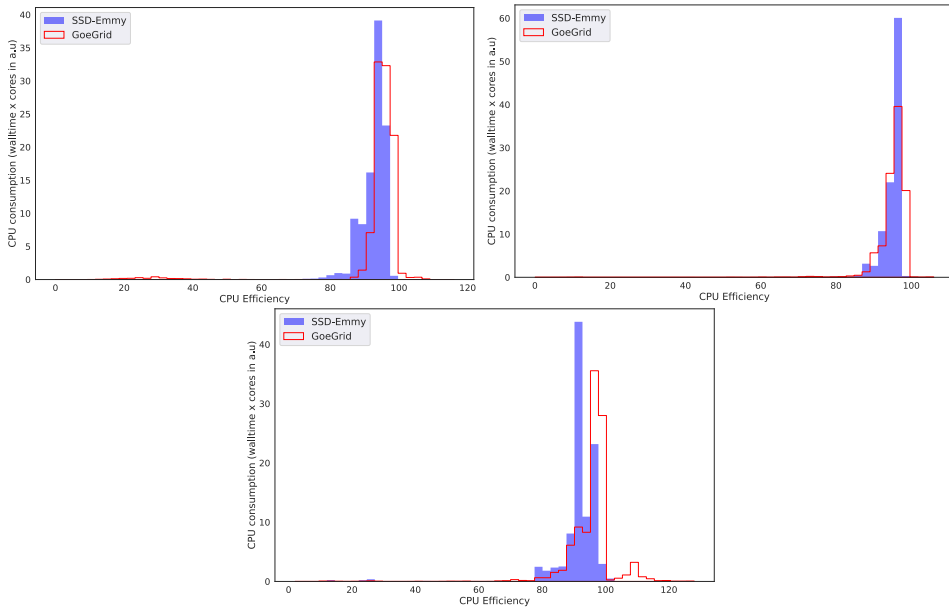


Figure 5. CPU consumption against the CPU efficiency is shown for all types of ATLAS jobs combined on a load scaling from 9600 cores (top left) to 15,360 cores (top right) and 19,200 cores (bottom).

3.3 Data caching

It must be determined whether the data access methods currently used can still be employed when the data are moved to the new remote mass storage sites. In the usual case, ATLAS runs the job near the data. In this paper, pre-caching, where small cache instances are used near sites, and the data are transferred there for the running jobs, is tested and results of these first tests are discussed. This method of data transfer for the jobs corresponds to the current ATLAS standard. Once the disk is filled and jobs are finished, the data are deleted and refilled with new data for a new set of jobs that are waiting in the queue. PanDA sets the rules for copying/moving the data upon which Rucio [12] acts and takes care of the data transfer and deletion. This method is tested at Emmy and compared to GoeGrid. The results are seen in Fig. 6, where the data are shown filling 50 TB of test disk space at Emmy as the jobs progress. Short-term spikes of Dark data are an artifact, resulting from the small size of the test disk, and can be ignored. There is a periodic trend of regular transfers from remote sites and deletion of data. No drawbacks are observed in the initial tests and Rucio kept pace well with rapid transfers for the initial disk size of 50 TB.

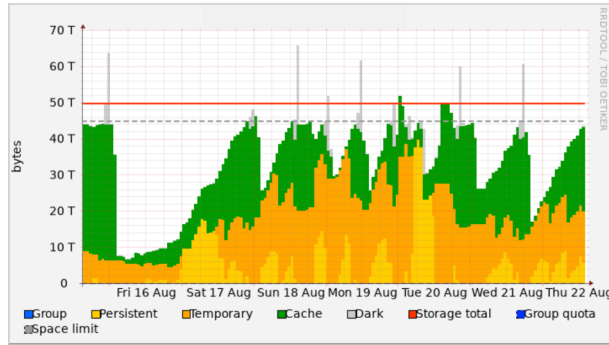


Figure 6. Data transfer and deletion on a 50 TB test disk over a period of a week.

4 Conclusions

This paper is focused on the methods used for the integration of two computing clusters (a Tier-2 WLCG centre, GoeGrid, and an NHR HPC site, Emmy), analysis of the integrated setup performance and R&D to adapt and optimise it for realistic conditions in running HEP jobs for the upcoming HL-LHC. Performance of the two clusters in the integrated setup is quite comparable. The latest nodes on the HPC outperform GoeGrid for typical HEP workflows, measured using the HEPsScore23 benchmark. The performance under realistic conditions is yet to be tested and analysed. The integrated setup sustained heavy loads larger than a Tier-2 WLCG site with only a minor decline in the average CPU efficiency. No drawbacks are observed with utilising the containerised setup to run HEP jobs. Using a local pre-cache instance is quite promising, with no disadvantages found in the initial tests. Further tests will be conducted with various disk sizes and numbers of drones. In addition, a Wide Area Network method will be tested and analysed, where the data will be downloaded directly to the worker nodes.

References

- [1] Shiers, J, *The Worldwide LHC Computing Grid (worldwide LCG)*, CPC **177**, 219-223 (2007)
- [2] Mkrtchyan, T, et al. *dCache: Inter-disciplinary storage system*, EPJ Web Conf., **251**, 02010 (2021)
- [3] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3**, S08003 (2008)
- [4] Maeno, T, Alekseev, A, Megino, B, et al. *PanDA: Production and Distributed Analysis System*, Comput Softw Big Sci **8**, 4 (2024).
- [5] Basney, J and Livny, M, *High Performance Cluster Computing: Architectures and Systems*, (Prentice Hall PTR, 1999)
- [6] Gregory, K, Sochat, V, Bauer, M, *Singularity: Scientific containers for mobility of compute*, PLOS ONE, **12**, 1-20 (2017)
- [7] *Docker*, <https://www.docker.com/>, accessed: 2024-12-05
- [8] *CVMFS*, <https://github.com/cvmfs/cvmfs>, accessed: 2025-01-13
- [9] Giordano, D, et al. *HEPsScore: A new CPU benchmark for the WLCG*, EPJ Web Conf **295**, 07024 (2024)

- [10] Yoo, A, Jette, M and Grondona, M, *Job Scheduling Strategies for Parallel Processing*, (Springer Berlin Heidelberg, 2003) 44–60
- [11] Von Cube, R, Caspart, R, Fischer, M, Giéls, M, Kuehn, E, Quast, G and Schnepf, M, *DLC 2020: Proceedings of the 4th International Workshop on Data Life Cycle in Physics*. Ed.: A. Kryukov (RWTH Aachen, 2020), 51-58
- [12] Barisits, M, Beermann, T, Berghaus, F, et al, *Comput Softw Big Sci*, **3**, 11 (2019)