T 91: GRID computing

Time: Thursday 16:00–17:30 Location: Tp

T 91.1 Thu 16:00 Tp

Opportunistic cloud computing for German HEP — •R. Florian von Cube, René Caspart, Max Fischer, Manuel Giffels, Eileen Kühn, Günter Quast, and Matthias J. Schnepf — Karlsruhe Institute of Technology

For on-demand usage of resources not dedicated to HEP, KIT develops the resource management system COBalD/TARDIS. It allows for dynamic integration of resources temporarily made available from different providers. Those might be e.g. university HPC clusters, shared science computing centers or commercial cloud providers.

By using modern virtualization and containerization technologies, such as Docker or Singularity, heterogeneous resources can be used transparently by experiments e.g. for Monte-Carlo-production and end-user analysis.

In this contribution we introduce COBalD/TARDIS and present recent successes with exemplary setups.

T 91.2 Thu 16:15 Tp

Job Shaping with HammerCloud ATLAS — •Michael Böhler, David Hohn, and Markus Schumacher — Albert-Ludwigs-Universität, Freiburg, Deutschland

The functionality of the compute sites of the Worldwide LHC Computing Grid for the ATLAS and CMS experiments is verified by a large number of experiment specific test jobs. These jobs are steered, controlled and monitored by the HammerCloud testing infrastructure. HammerCloud ATLAS runs different functional tests, continuously checking the site status by representative MC simulation and analysis jobs. If these test jobs fail, the site is automatically excluded from central ATLAS job brokerage system: only test jobs will be send to the site until the test results succeed again. The auto-exclusion mechanism increases the success rate of the user jobs by only allowing job brokerage to healthy sites.

The aim of Job Shaping, which is discussed in this talk, is to speed up auto-exclude and re-include decisions made by HammerCloud. This is to be achieved by dynamically adjusting the frequency of test jobs based on latest test job results. Dedicated visualizations are developed to provide intelligible information. Additionally, specialized debug test jobs can be sent to problematic sites to identify root causes of problems like failing or missing test job results. The additional information of the debug jobs will provide more detailed data in order to help problem solving and identifying failure patterns. Therefore new test templates are developed which focus on testing specific components of the site functionality.

T 91.3 Thu 16:30 Tp

Benchmarking of GRID resources using the HammerCloud service — Michael Böhler, •David Hohn, Benoît Roland, Benjamin Rottler, and Markus Schumacher — Albert-Ludwigs-Universität Freiburg

The high luminosity LHC is a major upgrade to the current accelerator and will provide an order of magnitude more high energy physics data in the coming decades than exists now. To enable the analysis of this wealth of data the computing resources need to follow suit and be upgraded as well.

Reliable performance information is the foundation for planning of upgrade efforts as well as efficient operation of the complex computing systems. To this end, benchmarks provide essential insight. This talk will present fast benchmarks of CPUs which can enable prompt matching between available and required resources, as is mandatory to incorporate opportunistic and heterogeneous resources.

The benchmarks are performed by the HammerCloud service within the ATLAS GRID computing network. This constitutes an extension of the HammerCloud use cases from functional to additional performance testing.

T 91.4 Thu 16:45 Tp

Performance monitoring of the opportunistic resource NEMO at ATLAS-BFG — Michael Böhler, Anton J. Gamel, •Stefan Kroboth, Benoit Roland, Benjamin Rottler, and Markus Schumacher — Albert-Ludwigs-Universität Freiburg

The workload of computing clusters is typically unpredictable and tends to alternate between over- and under-utilization of the available resources. The software COBalD/TARDIS provides an easy way to opportunistically make under-utilized resources of one cluster available to another cluster. Fine-tuning of the involved software infrastructure to optimize efficiency and user experience needs to be performed in a production environment and is therefore difficult without continuous monitoring of logs and meaningful metrics. In this work we present the current situation at Freiburg University where resources of the NEMO cluster are used to extend the WLCG-Tier-2/3 cluster ATLAS-BFG in an opportunistic fashion using COBalD/TARDIS. The talk covers the tools involved in the collection and analysis of logs and metrics acquired from different sources within the ATLAS-BFG and opportunistic NEMO. Examples of how the aggregation of logs and the monitoring of metrics aids decision-making are shown. Besides fine-tuning of the involved tools, this setup can also be used to detect problems and anomalies early on. It furthermore serves as a basis for the future development of an accounting system for compute infrastructure which involves opportunistically integrated resources.

T 91.5 Thu 17:00 Tp

Performance gain in HEP workflows via coordinated caches in heterogenous distributed systems — Rene Caspart, Tabea Fessenbecker, Max Fischer, Manuel Giffels, Christoph Heidecker, •Maximilian Horzela, Eileen Kühn, Günter Quast, and Paul Skopnik — Karlsruhe Institute of Technology

The steadily increasing demand on compute resources due to an explosively growing amount of data and demand for simulations for analysis purposes particularly in the context of the High Luminosity LHC can only be covered by harnessing novel computing concepts.

Besides the integration of additional resources, a promising approach to meet the need for resources is to improve the efficiencies of the workflows. In the HEP community these are often limited by the I/O bandwidth. Nevertheless, to utilize the full available computational capacities of distributed computational systems, aiming for optimal efficiency, coordinated caching solutions in combination with data-locality aware batch systems can provide a significant contribution.

This talk outlines basic concepts and visions for coordinated caching solutions in heterogeneous distributed batch systems.

T 91.6 Thu 17:15 Tp

Implementation and benchmarking of a caching solution in the ATLAS Freiburg environment — Michael Böhler, Anton Gamel, Stefan Kroboth, •Dirk Sammel, and Markus Schumacher — Albert-Ludwigs-Universität Freiburg

The near future of particle physics will be an era of high luminosity: the High Luminosity Large Hadron Collider (HL-LHC), which is expected to be operational in 2027, will enable the collection of data corresponding to a total luminosity of around 3000 fb^{-1} by the ATLAS experiment. In addition to this data, a large number of simulated events will be needed by a multitude of analyses. This results in the requirement for huge amounts of disk space. In order to cope with these enhanced storage requirements, scenarios with fewer, but larger storage sites and several compute sites without long-term storage are discussed.

Therefore, caching solutions are being developed. In this approach, the required data is transferred from external storage sites and then stored on local cache spaces. This enables fast access if the data is subsequently needed. The data is automatically deleted if it is not used for some predefined amount of time.

In this talk, the implementation of a caching setup, Disk-Caching-On-The-Fly (DCOTF), in the ATLAS Freiburg environment and the development of a benchmark that mimics a typical user-analysis are presented. The results of the benchmark testing the performance in Freiburg, especially in comparison to non-cached data access, are discussed.