



Contribution ID: 106

Type: Poster

Paving the Way for HPC: An XRootD-Based Approach for Efficiency and Workflow Optimizations for HEP Jobs on HPC Centers

Wednesday 13 March 2024 16:15 (30 minutes)

Today, the Worldwide LHC Computing Grid (WLCG) provides the majority of compute resources for the High Energy Physics (HEP) community. With its homogeneous Grid centers all around the world trimmed to a high throughput of data, it is tailored to support typical HEP workflows, offering an optimal environment for efficient job execution.

With the future German HEP computing strategy, however, there will be a shift away from dedicated resources to official shares on national HPC centers.

This bears many challenges, since these more heterogeneous resources are designed for performance and security rather than transferring and processing large amounts of data. The different focus and certain limitations can lead to higher failure rates and worse efficiency of HEP jobs running at such centers, i.e. because of tendentially slower WAN connections.

Monitoring data collected at the HoreKa HPC Center at KIT further confirmed that assumption. Lower CPU efficiency and an increased failure rate compared to the KIT Tier-1 center indicated a bandwidth limitation, in particular for data intensive workflows.

An efficient resource utilization, however, is the main objective for the success of the German HEP strategy in the future, not only in terms of sustainability, but also to cope with the anticipated data rates of the HL-LHC era. To tackle these challenges, we developed an XRootD/XCache based solution - in close contact with the XRootD/XCache developers - that aims to maximize computational throughput and mitigates the limitations of contemporary HPC centers. The currently operational setup at HoreKa leverages the parallel filesystem and a transfer node of the cluster as some sort of XRootD caching 'buffer', leading to a more stable performance and better utilization of the cluster.

In this contribution, the prerequisites and challenges associated with HEP workflows on HPC centers are pointed out and our solution for a more efficient utilization, backed by a fully operational proof of concept on HoreKa, is presented.

Significance

Since HPC centers gain importance in the HEP computing environment in the future, especially in Germany in the next year, a well prepared and fluent transition to such resources is important for the efficient operation of the WLCG. We are ensuring this with our work and other sites can benefit from our experience of incorporating HPC efficiently.

References

Experiment context, if any

The current P.o.C is running with CMS jobs, but the presented concepts are not specifically aimed at CMS and can in principle be generalized to other experiments.

Primary author: HOFSAESS, Robin (KIT - Karlsruhe Institute of Technology (DE))

Co-authors: STREIT, Achim (KIT - Karlsruhe Institute of Technology (DE)); PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE)); GOTTMANN, Artur Il Darovic (KIT - Karlsruhe Institute of Technology (DE)); QUAST, Gunter (KIT - Karlsruhe Institute of Technology (DE)); GIFFELS, Manuel (KIT - Karlsruhe Institute of Technology (DE)); SCHNEPF, Matthias Jochen

Presenter: HOFSAESS, Robin (KIT - Karlsruhe Institute of Technology (DE))

Session Classification: Poster session with coffee break

Track Classification: Track 1: Computing Technology for Physics Research