



Contribution ID: 329

Type: **Poster**

## XRootD caching for Belle II

*Tuesday 25 October 2022 16:10 (30 minutes)*

The Belle II experiment at the second generation e<sup>+</sup>/e<sup>-</sup> B-factory SuperKEKB has been collecting data since 2019 and aims to accumulate 50 times more data than the first generation experiment, Belle.

To efficiently process these steadily growing datasets of recorded and simulated data that end up on the order of 100 PB and to support Grid-based analysis workflows using the DIRAC Workload Management System, an XRootD-based caching architecture is presented.

The presented mechanism decreases job waiting time for often-used datasets by transparently adding copies of these files at smaller sites without managed storage.

The described architecture seamlessly integrates local storage services and supports the use of dynamic computing resources with minimal deployment effort.

This is especially useful in environments with many institutions providing comparatively small numbers of cores and limited personpower.

This talk will describe the implemented cache at GridKa, a main computing centre for Belle II, as well as its performance and upcoming opportunities for caching for Belle II.

### Experiment context, if any

Belle II

### References

### Significance

This is the first application of the XRootD caching technology for the DIRAC Workload Management System and a novel application of caching for the Belle II experiment. Owing to the unique challenges of Belle II, solutions like these may prove to be essential to analyze the large dataset we aim to collect and increase the efficiency of the resources available to the experiment.

**Authors:** QUAST, Gunter (KIT - Karlsruhe Institute of Technology (DE)); GIFFELS, Manuel (KIT - Karlsruhe Institute of Technology (DE)); SCHNEPF, Matthias; FISCHER, Max (Karlsruhe Institute of Technology); BAUER, Moritz David

**Presenter:** BAUER, Moritz David

**Session Classification:** Poster session with coffee break