



Contribution ID: 527

Type: Oral

Nordugrid ARC cache: Efficiency gains on HPC and cloud resources

Tuesday 5 November 2019 12:00 (15 minutes)

The WLCG is today comprised of a range of different types of resources such as cloud centers, large and small HPC centers, volunteer computing as well as the traditional grid resources. The Nordic Tier 1 (NT1) is a WLCG computing infrastructure distributed over the Nordic countries. The NT1 deploys the Nordugrid ARC CE, which is non-intrusive and lightweight, originally developed to cater for HPC centers where no middleware could be installed on the compute nodes. The NT1 runs ARC in the Nordugrid mode which contrary to the Pilot mode leaves jobs data transfers up to ARC. ARCs data transfer capabilities together with the ARC cache are the most important features of ARC.

HPCs are getting increased interest within the WLCG, but so are cloud resources. With the ARC CE as an edge service to the cloud or HPC resource, all data transfers required by a job are downloaded by data transfer nodes on the edge of the cluster before the job starts running on the compute node. This ensures a highly efficient use of the compute nodes CPUs, as the job starts immediately after reaching the compute node compared to the traditional pilot model where the pilot job on the compute node is responsible for fetching the data. In addition, the ARC cache gives a possible several-fold gain if more jobs need the same data. ARCs data handling capabilities ensures very efficient data access to the jobs, and even better for HPC centers with its fast interconnects.

In this presentation we will describe the Nordugrid model with the ARC-CE as an edge service to an HPC or cloud resource and show the gain in efficiency this model provides compared to the pilot model.

Consider for promotion

No

Authors: PEDERSEN, Maiken (University of Oslo (NO)); KONYA, Balazs (Lund University (SE))

Presenter: PEDERSEN, Maiken (University of Oslo (NO))

Session Classification: Track 3 –Middleware and Distributed Computing

Track Classification: Track 3 –Middleware and Distributed Computing