



Contribution ID: 151

Type: **not specified**

Advanced Data Staging in the ARC Computing Element

Tuesday, 27 March 2012 16:00 (30 minutes)

Description of the Work

One of the main problems with the old staging framework was a rigid design structure which made it difficult to introduce large architectural changes required to address the issues caused by heavy real-life loads. Therefore the system was completely redesigned with several significant changes. Instead of a simple first-in first-out queueing system for tasks, there is a single transfer queue of all files, and with the introduction of priorities and fair-share, an intelligent scheduler decides which transfers to process from the queue.

The scheduler sits in the middle of a three-layer system: the top layer accepts incoming requests for data transfer and directs them to the middle layer, which schedules individual transfers and negotiates with various intermediate catalog and storage systems until the physical file is ready to be transferred. The lower layer performs all operations which use large amounts of bandwidth, i.e. the physical data transfer. This layered structure provides flexibility to substitute other components for any given layer, allows more efficient use of the available bandwidth as well as enabling late-binding of jobs to data transfer slots. In addition, the system can easily be extended to improve throughput by adding extra nodes to which the physical transfers can be delegated.

Conclusions

The redesigned ARC CE data staging framework addresses all the concerns seen with the previous framework and brings improved throughput performance, while retaining the flexibility needed for future changes in the Grid environment.

Impact

The new framework is available in the most recent EMI release of the ARC CE and can be enabled through configuration options. Experience with early adopter sites has shown promising results, with increased throughput due to the more efficient queueing system ensuring that as one transfer finishes, there is another waiting to start immediately after. The addition of priority and fair-share combined with intelligent scheduling provides both users and site administrators with the flexibility they need to prioritise their workload.

Overview (For the conference guide)

The Advanced Resource Connector's Computing Element (ARC CE) is responsible for staging input and output data for tasks running on the computational resources it manages. This paper presents the newly redesigned staging framework within the ARC CE, which addresses several issues encountered as the data demands of

tasks have
increased.

Primary author: Dr CAMERON, David (University of Oslo)

Co-authors: KONSTANTINOV, Aleksandr (Vilnius University); KARPENKO, Dmytro (University of Oslo)

Presenter: Dr CAMERON, David (University of Oslo)

Session Classification: Data Technologies