



Contribution ID: 273

Type: Oral

Testing the performance of data transfer protocols on long-haul high bandwidth paths.

Thursday 7 November 2019 15:00 (15 minutes)

This paper describes the work done to test the performance of several current data transfer protocols. The work was carried out as part of the AENEAS Horizon 2020 project in collaboration with the DOMA project and investigated the interactions between the application, the transfer protocol, TCP/IP and the network elements. When operational, the two telescopes in Australia and South Africa that will form the SKA will each need to be able to transfer 1 Petabyte of data per day to various regional centres situated around the world that will extract the science from these data products. These requirements to transfer and access data are very similar to those of the current and future LHC experiments, and this has helped to form the collaboration between the particle physics and radio astronomy communities.

Measurements were made using high performance server nodes with NVME and RAM disks that were directly connected with 100 Gigabit links to the NRENs and the GÉANT in Europe, and to AARNet in Canberra and the Murchison Royal Observatory in Western Australia.

Some understanding of the disk to disk transfer rates was gained by measuring the read and write speeds between disk and memory for the file systems in use, and also the observed memory to memory network performance between the hosts used. In addition, some of the open source data transfer and network performance tools were instrumented to provide more insight on the observed rates and performance.

The measurements were made over various 100 Gigabit links, including dedicated and shared links, with round trip times (RTT) of 7.5, 15, 30 and 300 ms. The disk and file system configurations used included xfs over NVME disks in RAID0, a zfs pool of NVME disks, xfs on RAM disks, and systems at CERN and DESY. The following application and protocols were compared: GridFTP with globus-url-copy, Xrootd with xrdcp using xrd protocol, Xrootd with curl using http protocol, Xrootd with davix using http protocol, and dCache with davix using http protocol.

The results and analysis of the data from the sub-systems' and application performance are presented and discussed in relation to the end-to-end disk transfer rates observed. The paper concluded with some recommendations for improving performance and welcomes open discussion. The authors acknowledge the help from members of the DOMA Project.

Consider for promotion

Yes

Authors: Dr AMY, Shaun (CSIRO); Dr HUGHES-JONES, Richard (GEANT Association); Dr RAYNER, Tim (AARNet)

Presenter: Dr HUGHES-JONES, Richard (GEANT Association)

Session Classification: Track 4 –Data Organisation, Management and Access

