Pakistan-UK LHC Data Movement Issues Summary

Contact Point: Hans Addleman (addlema@iu.edu)

Audience: General

Last updated: August 17, 2022

The National Center for Physics (NCP) at the Quaid-i-Azam University Campus in Islamabad, Pakistan, is a Tier 2 Large Hadron Collector (LHC) Site. NCP contacted EPOC in September, 2018, for assistance with an ongoing data transfer performance issue to and from the GridPP (www.gridpp.ac.uk) Tier 1 site at the Queen Mary University in London. As a Tier 2 LHC site, NCP is responsible for downloading data sets from the Tier 1 site to share with LHC researchers in the region in a timely manner. NCP has a 1Gbps connection to their national R&E network, Pakistan Education and Research Network (PERN), however, transfer rates were as low as 40Mbps to some Tier 1 LHC sites. These performance problems were intermittent, and were ongoing for at least a year prior to EPOC's involvement . Our engagement with NCP started in October 2018 and ended in January 2019. We identified five main issues contributing to the low transfer speeds, which resulted in performance improvement to ~480Mbps. However, this rate is not consistent due to an existing congested link in the path, which is detailed in issue five below.

The performance goal requested by NCP was to achieve file transfer rates to LHC sites in line with transfer rates to other locations. At the time, NCP was able to successfully achieve transfers up to approximately 500Mbps to AARNet (Australia) and up to 280Mbps to ESnet (es.net) Data Transfer Nodes in the United States. Engineers concluded that an acceptable performance goal would be to achieve at least 160Mbps to the GridPP LHC site at Queen Mary University.

EPOC engineers and engagement specialists worked with engineers and systems administrators from five other institutions during the investigation including TEIN*CC (www.tein.asia), PERN (pern.edu.pk), GEANT (www.geant.org), and INFN (home.infn.it).

During the investigation, a perfSONAR (www.perfsonar.net) mesh and dashboard that included servers spanning the NCP, PERN, TEIN, and GEANT networks was created. This mesh allowed for easier ongoing identification of packet loss and network congestion. perfSONAR experts from the GlobalNOC systems team were able to guide the configuration of these hosts and help alleviate systems issues that arose. Adhoc latency and bandwidth testing were also used on an ongoing basis to see the effect of troubleshooting in real time. The corresponding perfSONAR MaDDash was left in place for future testing as well.

Five main issues were identified during the engagement:

1. A traffic shaping misconfiguration on the NCP connection to PERN caused research and education traffic to be limited to 50Mbps. PERN identified the error and fixed the issue.

This material is based upon work supported by the National Science Foundation under Grant No. 1826994

PERN engineers reported that their research and education traffic should not have had any traffic shaping applied.

2. A top-of-rack switch was identified as a bottleneck between NCP's file transfer node and edge router. When the file transfer node was directly attached to the edge router, performance increased from 40Mbps to an average close to 500Mbps during testing.
3. We also observed in the perfSONAR dashboard small amounts of ongoing packet loss within the campus network supported by NCP. NCP moved their data node closer to the edge of their network to alleviate the issue.
4. Packet loss was also identified by the perfSONAR mesh inside the PERN regional network. PERN network engineers worked to resolve the issue, however, they did not report the specific cause of the loss.
5. Additional bottlenecks were observed at the connection between PERN and TEIN networks. This connection currently only supports 1Gbps, making congestion inevitable, and causes packet loss which inhibits the rate of NCP's file transfer speeds. This cause was verified by temporarily changing the routing to a commercial path, at which time NCP achieved expected transfer rates. TEIN and PERN are currently discussing an interconnect upgrade to 10Gbps, which would likely alleviate the congestion and allow for greater transfer speeds. It was suggested that in the meanwhile, NCP could work with PERN to have their LHC traffic follow the commercial path while they wait for the interconnection upgrade.

The final result of this engagement was an overall increase in data transfers between NCP and Queen Mary University to approximately 480Mbps. As identified above, this performance is based on a test during an uncongested time. Additional areas for performance improvements have been identified, but require longer term changes in infrastructure to be achieved. EPOC will stay engaged with all parties and help NCP test when needed.

This material is based upon work supported by the National Science Foundation under Grant No. 1826994
info@epoc.global