Kyungpook National University Data Transfer Performance to CERN

Contact Point: Ed Moynihan, edmoyn@iu.edu
Audience: General
Last updated: November 12, 2021

ABOUT EPOC

Over the last decade, the scientific community has experienced an unprecedented shift in the way research is performed and how discoveries are made. Highly sophisticated experimental instruments are creating massive datasets for diverse scientific communities and hold the potential for new insights that will have long-lasting impacts on society. However, scientists cannot make effective use of this data if they are unable to move, store, and analyze it. The Engagement and Performance Operations Center was established in 2018 as a collaborative focal point for operational expertise and analysis and is jointly led by Indiana University (IU) and the Energy Sciences Network (ESnet). EPOC provides researchers with a holistic set of tools and services needed to debug performance issues and enable reliable and robust data transfers. By considering the full end-to-end data movement pipeline, EPOC is uniquely able to support collaborative science, allowing researchers to make the most effective use of shared data, computing, and storage resources to accelerate the discovery process.

EPOC supports six main activities:

- *Roadside Assistance and Consultations* via a coordinated Operations Center to resolve network performance problems with end-to-end data transfers;
- *Application Deep Dives* to work more closely with application communities and understand full workflows for diverse research teams in order to evaluate bottlenecks and potential capacity issues;
- *Network Analysis enabled by the NetSage* monitoring suite to proactively discover and resolve performance issues;
- *Data Transfer Testing/ Data Mobility Exhibition* to check transfer times against known good end points;
- *Provision of managed services* via support through the IU GlobalNOC and our Network Partners;
- *Coordinated Training* to ensure effective use of network tools and science support.

Kyungpook National University Data Transfer Performance to CERN

The Center for High Energy Physics at Kyungpook National University (KNU) in Daegu, South Korea is the home of both a Tier 2 and Tier 3 data center for the Compact Muon Solenoid (CMS), a general purpose detector that is part of the Large Hadron Collider at CERN in Geneva, Switzerland (http://t2-cms.knu.ac.kr/home.php). As such, it was designed primarily to receive data from the Tier 1 site for Asia, located in Taiwan, and to serve that data to local researchers.

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

However, KNU also occasionally sends data back to CERN for additional processing and analysis or to stress test the current network infrastructure.

During routine traffic analysis of the TransPAC-PacWave 100G circuit between Seattle and Tokyo using NetSage (https://transpac.netsage.global), EPOC engineers noticed a significant increase in traffic from KNU to CERN over the first 3 months of 2021. NetSage analysis showed that nearly 10TB was transferred from KNU to CERN, compared to only 117GB in the final 3 months of 2020. Transfer rates on these flows were low, averaging only 30Mbps. Additionally, while there was an increase in traffic from KNU to CERN, there was only one small flow in the opposite direction, CERN to KNU, indicating that traffic in the reverse direction might be traveling a different path, or that there might be asymmetric routing.

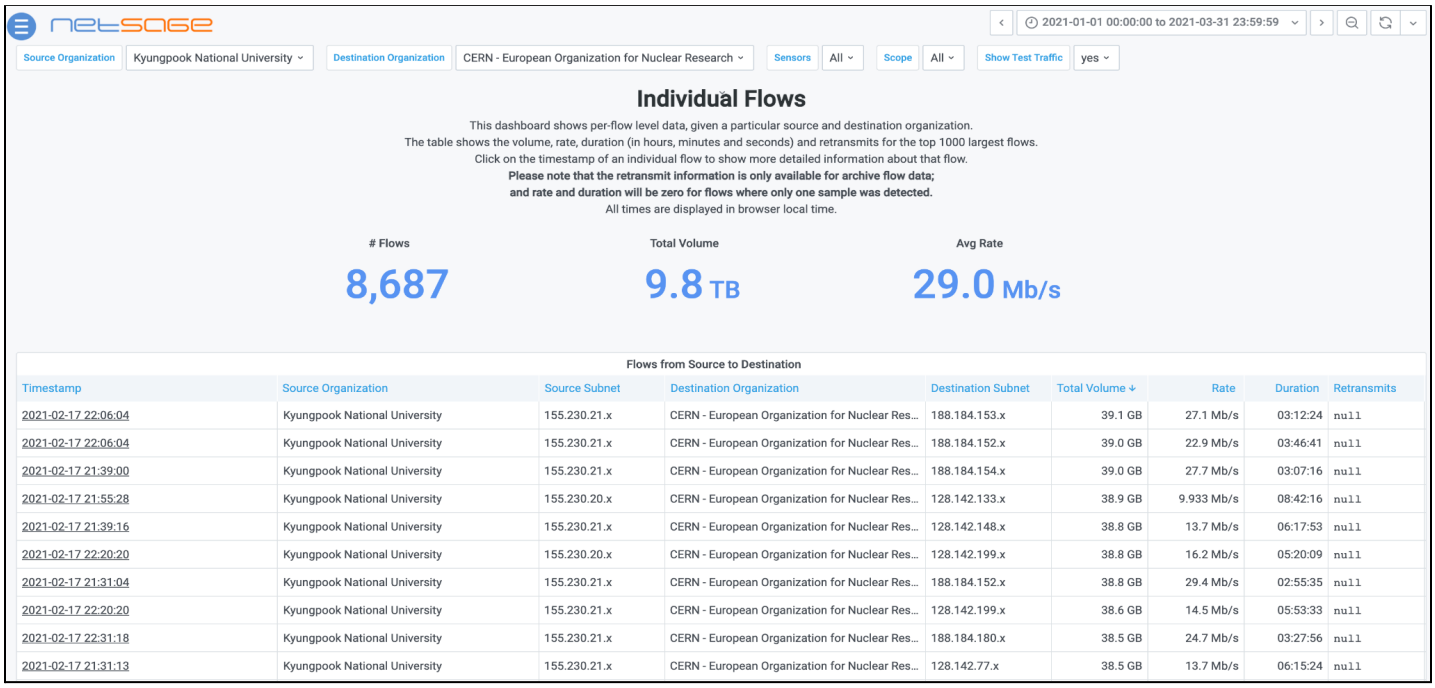

Timestamp	Source Organization	Source Subnet	Destination Organization	Destination Subnet	Total Volume	Rate	Duration	Retransmits
2021-02-17 22:06:04	Kyungpook National University	155.230.21.x	CERN - European Organization for Nuclear Res...	188.184.153.x	39.1 GB	27.1 Mb/s	03:12:24	null
2021-02-17 22:06:04	Kyungpook National University	155.230.21.x	CERN - European Organization for Nuclear Res...	188.184.152.x	39.0 GB	22.9 Mb/s	03:46:41	null
2021-02-17 21:39:00	Kyungpook National University	155.230.21.x	CERN - European Organization for Nuclear Res...	188.184.154.x	39.0 GB	27.7 Mb/s	03:07:16	null
2021-02-17 21:55:28	Kyungpook National University	155.230.20.x	CERN - European Organization for Nuclear Res...	128.142.133.x	38.9 GB	9.933 Mb/s	08:42:16	null
2021-02-17 21:39:16	Kyungpook National University	155.230.21.x	CERN - European Organization for Nuclear Res...	128.142.148.x	38.8 GB	13.7 Mb/s	06:17:53	null
2021-02-17 22:20:20	Kyungpook National University	155.230.20.x	CERN - European Organization for Nuclear Res...	128.142.199.x	38.8 GB	16.2 Mb/s	05:20:09	null
2021-02-17 21:31:04	Kyungpook National University	155.230.21.x	CERN - European Organization for Nuclear Res...	188.184.152.x	38.8 GB	29.4 Mb/s	02:55:35	null
2021-02-17 22:20:20	Kyungpook National University	155.230.21.x	CERN - European Organization for Nuclear Res...	128.142.199.x	38.6 GB	14.5 Mb/s	05:53:33	null
2021-02-17 22:31:18	Kyungpook National University	155.230.21.x	CERN - European Organization for Nuclear Res...	188.184.180.x	38.5 GB	24.7 Mb/s	03:27:56	null
2021-02-17 21:31:13	Kyungpook National University	155.230.21.x	CERN - European Organization for Nuclear Res...	128.142.77.x	38.5 GB	13.7 Mb/s	06:15:24	null

Figure 1: NetSage Dashboard showing data transfers from Kyungpook National University (KNU) to CERN between January 1, 2021, and March 31, 2021. Available online at:https://aponet.netsage.global/grafana/d/-l3_u8nWk/individual-flows?orgId=1&var-src=Kyungpook%20National%20University&var-dest=CERN%20-%20European%20Organization%20for%20Nuclear%20Research&var-sensors=TransPAC%20Seattle%20sFlow&var-country_scope=All&var-is_net_test=yes&from=1609477200000&to=1617249599000.

KNU connects to two advanced networks that serve the South Korean research and education community, the Korea Advanced Research Network (KOREN) and the Korea Research Environment Open Network (KREONET). KOREN generally handles domestic traffic between higher education institutions within South Korea, while KREONET generally carries larger,

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

international data transfers. In turn, both KOREN and KREONET connect to the Asia Pacific Advanced Network (APAN) and to the Trans-Eurasia Information Network (TEIN*CC). KREONET also operates a direct 100G connection from South Korea to Chicago and a 10G connection between Chicago and Amsterdam. With access to this advanced infrastructure, EPOC staff determined that KNU should have been getting much higher transfer rates on these flows.

Initial trace routes indicated that traffic being seen by the TransPAC NetSage sensor in Seattle from KNU was using the KOREN network and TEIN*CC to get to APAN in Hong Kong and then APAN-JP in Tokyo, then transiting the Pacific via the TransPAC-PacWave 100G Circuit, using ESnet circuits to transit the US and the Atlantic, before using the pan-European network GÉANT to get to CERN, as seen in Figure 2. There were no public perfSONAR nodes or router proxies at KNU for EPOC staff to be able to perform additional testing, so EPOC staff reached out to engineering contacts at KREONET to see if more information could be gathered on the cause of the traffic increase, poor performance, and asymmetrical routing, and to offer assistance with improving speeds and finding a more efficient route.

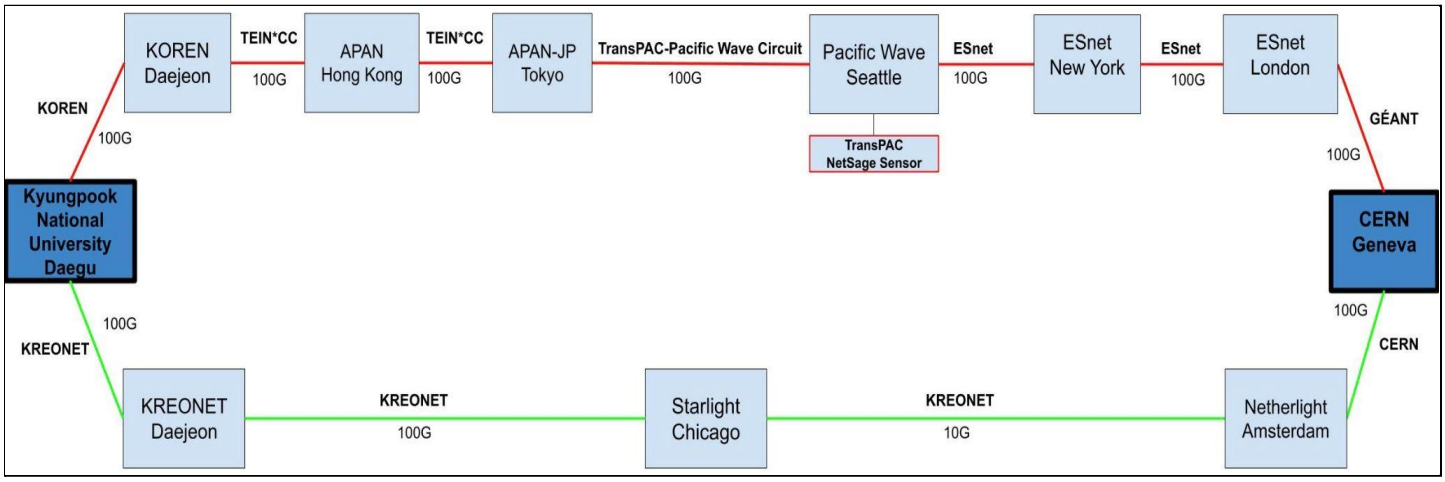

Figure 2: A diagram showing how the traffic flows from Kyungpook National University (KNU) in Daegu, South Korea, to CERN in Geneva, Switzerland, over the TransPAC Pac Wave circuit, in red, and the adjusted route using KREONET through Chicago in Green. Fewer hops in general means a more effective transfer.

KREONET engineers investigated and confirmed that some traffic from KNU to CERN was using KNU's KOREN connection while other flows used KNU's KREONET connection. The flows from KNU to CERN using KREONET's connection went directly from Daejeon to Chicago, and then used KREONET's 10G link from Chicago to Amsterdam, where KREONET peers directly with CERN. KREONET engineers ran further tests on traffic using the KREONET connection and saw much better overall performance as expected with a route that has significantly fewer hops.

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

After discussing possible routing options, engineers at both KREONET and KOREN decided that it would be more effective for all of the data from KNU to use the KREONET route rather than the KOREN route, and routing adjustments were made. With this change, the data transfers from KNU to CERN no longer used TransPAC and so could not be seen on the TransPAC NetSage Dashboards. KREONET engineers confirmed that all of the traffic had been shifted to the KREONET route.

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