National Center for Atmospheric Research Performance
EPOC Contact Point: Brenna Meade (meadeb@iu.edu)
Audience: General
Last Edit: May 2, 2022

EPOC
Engagement and Performance Operations Center

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 resolve 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.

National Center for Atmospheric Research Performance
The National Center for Atmospheric Research (NCAR) in Colorado supports the atmospheric and related Earth system science community with High Performance Computing resources , research aircraft, sophisticated computer models, and extensive data sets. NCAR uses tools like the Local Data Manager (LDM) to share weather data between 250 member institutions across the US that uses a multicast protocol to transfer data. NCAR engineers began testing a new version of LDM and discovered that they could not pass data faster than 50Mbps without significant packet loss. NCAR engineers reached out to staff at EPOC who helped determine the cause of this packet loss.

Troubleshooting

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

NCAR was running tests of the LDM transfer tool using a test bed that consisted of sites at University of Virginia (UVA), University of California San Diego (UCSD), University of Wisconsin Madison (UW-M), University of Washington (UW), and NCAR. The endpoints were connected via the Internet2 AL2S network, as shown in Figure 1. In the test setup, UCSD and NCAR had issues with packet loss when sending and receiving from the other three sites. UW-M, UVA, and UW are able to transmit and receive data as expected from each other, as seen in Figure 2.

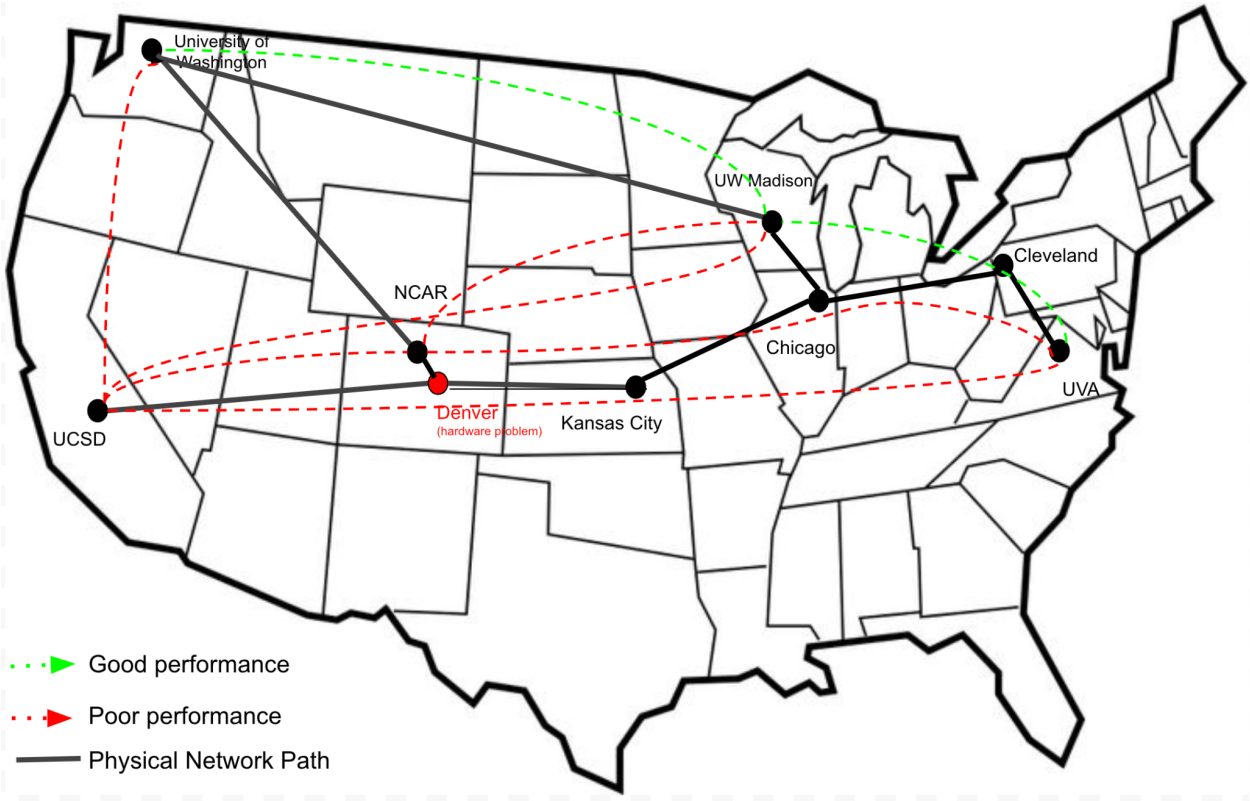

Figure 1: A map showing relevant portions of the Internet2 Network, and the end points that were part of the LDM testing. The green lines represent the transfers between endpoints with no issues, while the red represents the paths with poor performance.

		Receiver				
		NCAR	UWisc	UVA	UWash	UCSD
Sender	NCAR					
	UWisc					
	UVA					
	UWash					
	UCSD					

Figure 2: A table illustrating the performance between each of the sites that were part of LDM testing.

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

The initial investigation focused on whether this was related to the multicast protocol. Multicast is used to send a set of data across a network to many users at the same time from a single source. Initial testing involved using different multicast group addresses, or what a user connects to when receiving the multicast stream. Depending on the platform, some network devices hard-code the behavior of some well-known group addresses, regardless of other configuration settings, so it was thought perhaps a faulty group address was the issue. However, changing the group address did not alleviate the issue.

The next step was to run tests using iPerf multicast tools between the various sites. It was discovered that tests that achieved transfer rates over 50Mbps experienced a large amount of packet loss when the transfer path included Denver. This information led engineers to determine that there might be a piece of equipment in the path between non-working endpoint pairs. A likely candidate was the Front Range GigaPop (FRGP) network near Denver, part of which could prevent multicast packets from being forwarded.

FRGP engineers monitored interface counters on switches along this path during continued testing, and an exchange point switch in Denver was found to be the problem. The output counters on this device showed that it would stop transmitting multicast packets after approximately 250 packets had been forwarded.

Conclusion

Once the issue was identified, work began to replace the switch. Testing by the researchers after replacement verified that the path was now clean and transfer rates were as expected.

The result of this engagement was an elimination of the packet loss for multicast transfers traversing the Denver FRGP switch. Though engineers initially thought it was a multicast issue, the process of troubleshooting led them to a hardware issue instead. As a result, multicast traffic is no longer limited between these sites to 50Mbps allowing the upgrade to the new LDM program to proceed.

Acknowledgement: The Engagement and Performance Operations Center (EPOC) is a joint collaboration of Indiana University and the Energy Sciences Network (ESNet), and is funded in part by NSF Award #1826994.