EPOC

Engagement and Performance Operations Center

Award #1826994

Newsletter | September 24, 2020

EPOC Data mobility Exhibition (DME)

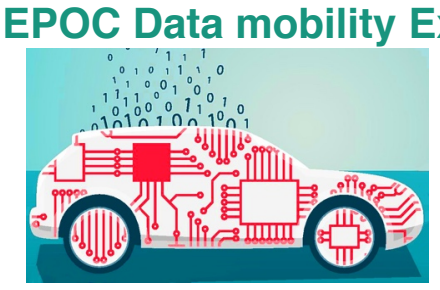

How long does a One Terabyte (1TB) research dataset download take to complete on your campus? A lunch break? An entire workday? What if, when we still went to offices, you clicked download, walked down to the water cooler, had a chat with a coworker, then upon return, noticed the download finished in around 20 minutes. This should be your reality with a 10G Data Transfer Node (DTN), Science DMZ, and parallel transfer application.

Unfortunately, this is not the case on many campuses where researchers download to their desktops, which can take most of the work day. Or if this is attempted connected to the campus WiFi on a laptop, it may take 24-48 hours. The EPOC team suggests that a goal for a campus is to download the 1TB dataset in about an hour, but how can a campus work toward measuring that type of performance, let alone meeting it?

The Data Mobility Exhibition (DME) was established to measure and understand data transfer baselines for performance expectations, to increase knowledge of data transfer hardware and software, and to share advanced design patterns of portal software as a benefit for scientific communities. The DME uses reference data sets, and existing campus CI components, to download, measure, and potentially improve scientific data movement capabilities. The DME is a multi-year experiment where we hope to understand and improve the research community's collective ability to address a critical problem in scientific technology use: transferring data sets using networks at consistent high performance speeds.

During the past 12 months, there have been over 600 transfer tests completed. Is your campus next? This is an excellent opportunity to test campus DTNs, Science DMZ, and CC* infrastructure with real data transfers to known, tuned endpoints. For more information, past webinars, and contact information, please see: https://fasterdata.es.net/performance-testing/2019-2020-data-mobility-workshop-and-exhibition/2019-2020-data-mobility-exhibition/

Roadside Assistance and Consultations: Sometimes It's Not Just One Thing

Data transfer performance can often be impacted by multiple factors, and at times are not straightforward to resolve. An investigation to improve performance for data moving between two sites may involve multiple phases, as the first fix may uncover other, more significant problems.

The Ohio Supercomputer Center (OSC) recently expanded their perfSONAR deployment and testing to better understand the performance between

Upcoming Talks and Events (EPOC and others)

25 September, 2020, 2pm EDT: Jennifer Schopf, Indiana University
- International Networks, "NetSage Now", CI Engineering Brownbag Talks.

28 September, 2020, 11am EDT: Ron Hutchins, University of
- Virginia, "ACCORD: Integrating CI policy and mechanism to support research on sensitive data",Trusted CI Webinars.

2 October, 2020, 2pm EDT: Edward Colone, University of Michigan &
- perfSONAR, "Mobile perfSONAR nodes best practices", CI Engineering Brownbag Talks.

15 October 2020, TBD: Jens David Ohlin, Cornell University, For more
- information: CACR Security Speaker Series.

16 October, 2020, 2pm EDT: Vas Vasiliadis, University of
- Chicago & Globus, "GCS v5" , CI Engineering Brownbag Talks

23 October, 2020, 2pm EDT: Michael Smitasin & Jay Krous,
- Lawrence Berkeley National Laboratory, "Cloudflare DNS", CI Engineering Brownbag Talks.

26 October, 2020, TBD: "Enforcing Security and Privacy Policies to
- Protect Research Data", Trusted CI Webinars.

5 November, 2020, TBD: Duncan Hollis, Temple University, For more
- information: CACR Security Speaker Series.

6 November, 2020, 2pm EDT: Edward Colone, University of
- Michigan & perfSONAR, "Provisioning perfSONAR with

their resources and several commonly used sources and destinations, including several Department of Energy sites supported by ESnet. The initial data showed very low transfer rates, and EPOC was contacted to help troubleshoot the problems. Transfer performance from OSC to the National Energy Research Scientific Computing Center (NERSC), a common collaborator for OSC was below 60 Mbps, when the performance in the other direction was over 9Gbps. In addition, there were hundreds of packet retransmissions during the brief test, which often indicates firewall issues or a rate-limiting device.

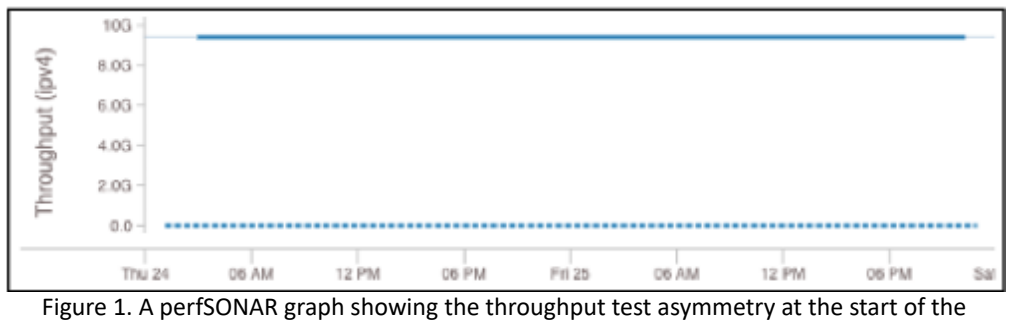

Figure 1. A perfSONAR graph showing the throughput test asymmetry at the start of the engagement.

In order to debug this issue, EPOC consulted with engineers from ESnet, NERSC, OSC, and OARnet to learn more about the end-to-end path the data was taking. This included understanding the routed paths of the networks involved using traceroute, examining the device capabilities and configurations of the network components, and evaluating if congestion patterns or time of day affected the observed throughout.

Three problems were discovered, although the last obfuscated the first two. Initially, the poor performance was believed to be caused by the fact that instead of using regional research and education (R&E) circuits that have been tuned for scientific data flows, the data transfers were being routed over commercial networks. ESnet engineers found out-of-date BGP (Border Gateway Protocol) policies on the ESnet routers, which meant traffic from OSC to NERSC wasn't taking advantage of the more expedient connections. In addition, it was discovered that traffic in the reverse direction was taking a different, but still commercial, path. This was resolved when OSC Engineers found a memory configuration issue affecting the routing table on their core router.

However, the performance was still not as expected, and there were still a large number of packet retransmissions taking place. Additional investigation found that several routers and switches in the path were set to have a 1500 byte maximum transmission unit (MTU). Reconfiguring these to have jumbo MTU, set to 9000 byte, enabled the final performance gain.

Once these issues were addressed, the network measurements were found to have increased to an average of 6Gbps in both directions, which far exceeded the expectations of the team when measuring just a short test transfer. Conversations with several end users showed that actual transfers for larger data sets and longer-lived flows between OSC and NERSC improved significantly.

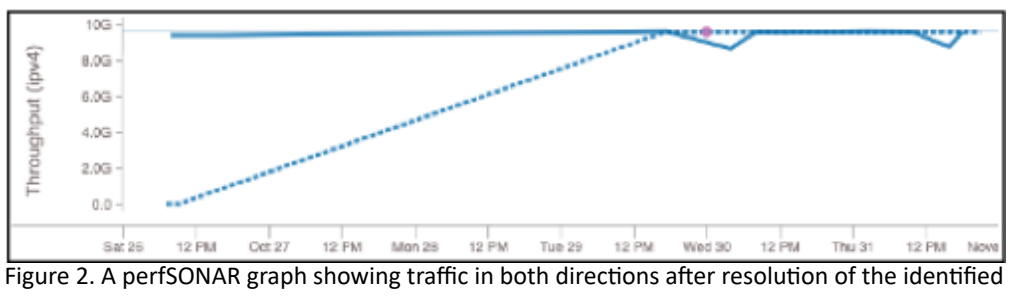

Figure 2. A perfSONAR graph showing traffic in both directions after resolution of the identified issues.

EPOC appreciated the contributions from staff at OSC, OARNET, ESnet, and NERSC for their assistance in resolving these issues. For more information about problems like this one, please see:

- Tuning your MTU: https://fasterdata.es.net/network-tuning/mtu-issues/
- PerfSONAR measurement tool: https://www.perfsonar.net/
- Normalization of Research and Education Routing: https://epoc.global/wp-content/uploads/2020/09/20200505-UofSC-Virtual-Workshop-BGP-Architectures-and-Best-Practices.pdf

More information about the Roadside Assistance and Consulting process is available at: EPOC Roadside Assistance.

Ansible" , CI Engineering Brownbag Talks.

https://epoc.global