100Gb Network Analysis: for High Energy Physics

Pamela Perkins
University of Notre Dame, IN 46556

The Purpose

There is an increasing need to share big data for science and research purposes on college campuses. Despite the availability of high-capacity networks, Notre Dame researchers are experiencing bottlenecks in their data transfers across their 100Gb network. Our goal is to tune the cyberinfrastructure components to optimize performance of the 100 Gb network, while also keeping information highly secure within the Science DMZ.

When tested, the highest single stream point to point was sometimes low and sometimes high, but with several locations were tested, within the United States and overseas, we found a significant correlation between growing latency and significantly dropping bandwidth. There was no significant correlation found in the number of hops.

PerfSONAR

With the help of PerfSONAR, we have tested the network's performance by attaining the following metrics: The amount of achievable bandwidth (throughput), available to move data, quantified in bits per second (Mbps); round trip latency (RTT), which is the amount of time it takes to receive a packet to and from across the network measured by milliseconds (ms); and path data transit (traceroute tool) which measures the number of hops between points in the network data takes to get to its destination.

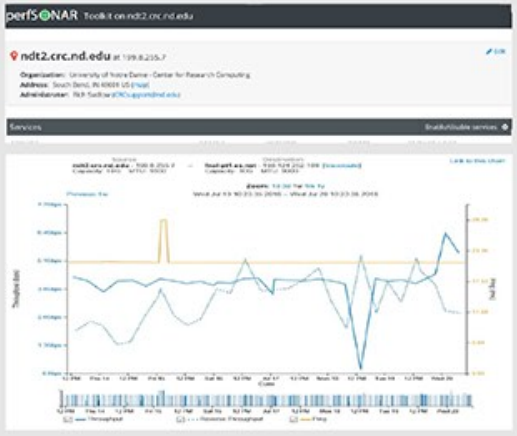

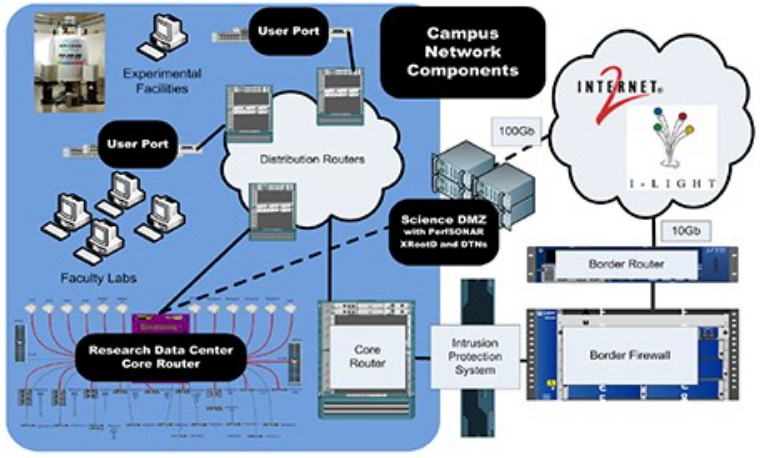

Findings

I identify factors that cause poor network performance: the presence of a bottleneck, high bandwidth-delay products on "long fat" networks, and network throttling (rate limiting) that could occur especially from outside the United States. To illustrate the latter issue, I contacted one of the test points in China to learn about our network connection-- The result was China was able to move more gig between us, than we could to them.

Correlation graph:

- When several locations were tested, within the United States and overseas, I found a significant correlation between growing latency and significantly dropping bandwidth.
- There was no significant correlation found in the number of hops.
- Correlation does not mean causation, but It's helpful to find out there was a connection between the two.

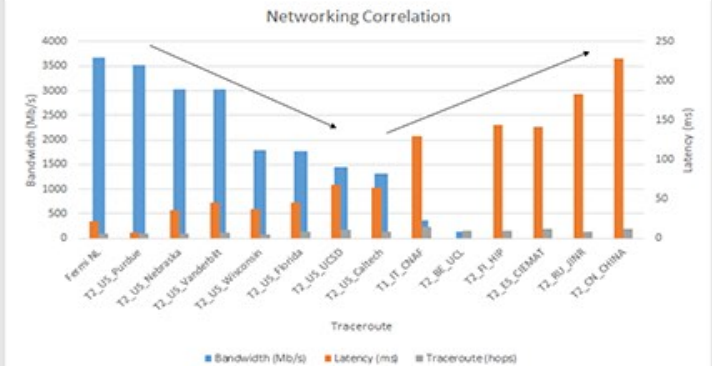

Future Work

Our future work entails clearing up the network's bottleneck issue, finding ways to reduce the throttling, reduce the cluttering by adding more servers, and keep our data safer by increasing security with our transfers. We hope to be a highly secure, but also a very fast 100Gb network.

Based on the 100Gb "Top 10" physics sites from which we pull data, and given the aggregate is 17Gbsp; I've come to the the conclusion that adding a couple more cache servers will enable the network to run more effectively.

Therefore, I recommend the Notre Dame High Energy Physics team bring in four additional xRootd servers.

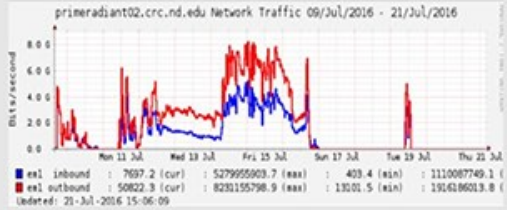

References

Science DMZ security – firewalls vs. router ACLs. (2016). https://fasterdata.es.net/

Hanemann, A., Boote, J. W., Boyd, E. L., Durand, J., Kudarimoti, L., Łapacz, R., ... & Zurawski, J. (2005, December). Perfsonar: A service oriented architecture for multi-domain network monitoring. In *International Conference on Service-Oriented Computing* (pp. 241-254). Springer Berlin Heidelberg.

Bauerdick, L. A. T., Bloom, K., Bockelman, B., Bradley, D. C., Dasu, S., Sfiligoi, I., ... & Yagil, A. (2012). XRootd monitoring for the CMS experiment. In *Journal of Physics: Conference Series* (Vol. 396, No. 4, p. 042058). IOP Publishing.

Acknowledgment

I would like to thank the National Science Foundations and my advisor, Dr. Paul Brenner, as well as Dr. Thain, Dr.Lannon, and Kenyi Paolo HurtadoAnampa , for all the help and guidance that they provided me. Also I would like to think Kallie O'Connell for her support and advice during these 10 weeks.

CRC
CENTER FOR RESEARCH COMPUTING