

Understanding Routing Behaviors on R&E Networks

Submitted by Jennifer M. Schopf, Indiana University; Brian Tierney, ESnet; Hans Addleman, Indiana University; Doug Southworth, Indiana University

As operators two of NSF's International Research and education Network Connections (IRNC) networks, we see two pragmatic and important research questions that we are unable to answer given the current state of network measurement. The first is that we have no real understanding of how changes in capacity to research and education (R&E) networks will affect network routing and traffic. The second is the need for better understanding of when routing policy is misconfigured.

Every time capacity is added to the R&E network space, the path the traffic takes between a given source and destination may change. For example, when an R&E path between Guam Open Research and Education eXchange (GOREX) and Hong Kong was put in place in 2018, traffic that had flowed between several institutions in China with receivers in South Korea was re-directed to prefer this path, even though it meant that the route included crossing the Pacific Ocean, twice [1]. Similarly, we saw unexpected changes when an R&E route became available between Fortaleza, Brazil, and Sangano, Angola [2]. When new capacity is added to the R&E networking fabric, unanticipated changes are likely to result. An open question is how to anticipate these consequences.

There is also a more general question of which R&E routes are inefficient in the current networks. Knowing that capacity changes can impact routing, one can examine the routes after the fact for egregious problems. But what about the routing issues that currently exist, impacting day-to-day throughput, but not to the extent that they are noticed? What is needed are datasets and tools to understand, in a more straightforward way, what routes currently exist and to make predictions of what other possible routes are more effective, both in practice and in light of possible infrastructure additions or changes.

Much of the needed data can be collected, but it's not available consistently, in formats that are easily shared, or with the breadth or coverage needed to be able to address our research questions. Current data includes:

- Hop counts – available via Traceroute, which can only be run by a resource owner, or the perfSONAR [3] Traceroute collection tool, which is only available to someone with whom a perfSONAR archive is shared, and generally has limited scope and coverage.
- Latency information, link congestion information – also available via perfSONAR, with the same constraints.
- Link Capacity – no general data set available
- Routing tables – Route Views [4] collects this information, but deployment is limited and data access for this use case would be quite complicated with the current implementation.
- BGP Configuration data – LookingGlass [5] can collect this data, but again, deployment is limited and data access for this use case would be quite complicated with the current implementation.

If this data were more easily available, there are still no common analytical techniques in practice today. And of course, the ultimate goal isn't to simply identify ineffective routes, but to fix them. Research into tools and techniques to address these issues are needed.

References

[1] NetSage page for TransPAC5 20-G LAG between Guam and Hong Kong showing traffic between KISTI (S. Korea) and the Computer Network Information Center at the Chinese National Academy of Science,

https://portal.netsage.global/grafana/d/xk26iFhmk/flow-data-for-circuits?orgId=2&var-Sensors=TransPAC%20Hong%20Kong%20sFlow&var-country_scope=All&var-is_net_test=yes&from=1525147200000&to=1546232399000

[2] Roderick Fanou, Bradley Huffaker, Ricky Mok, kc claffy, "Unintended consequences: Effects of submarine cable deployment on Internet routing," Proceedings of PAM 2020, April 2020.

[3] perfSONAR, <http://perfsonar.net>

[4] Route Views, <http://www.routeviews.org/routeviews/>

[5] Looking Glass, https://en.wikipedia.org/wiki/Looking_Glass_server