

August 17, 2010

NTAC peering/routing
8/17/10

In attendance: Jeff Bartig, Randy Brogle, Grover Browning, Cort Buffington, PJ Clayton, David Crowe, Cas D'Angelo, Dale Finkelson, Tom Hamilton, John Hernandez, John Hess, Akbar Kara, Michael Lambert, Caren Litvanyi, Darrell Newcomb, Mike Phillips, Dave Pokorney, Chris Robb, Linda Roos, Von Royal, Paul Schopis, Shannon Spurling, Steven Wallace, Kenny Welshons, Matt Zekauskas

Agenda

1. Update from Internet2 on peering
2. Update on CPS/TR consolidation implementation
3. Dallas peering node
4. Other topics
5. BTOP network design review discussion

1. Update from Internet2 on peering

Chris reported that it has been a quiet month. Internet2 has ordered the fourth 10G from McLean to Ashburn in anticipation of the start of the fall semester. The NY router is scheduled for delivery on September 9 and should be ready to be installed by September 15.

2. Update on CPS/TR consolidation implementation

There are some additional routes being exchanged. Have been working with FRGP, LONI and a few others to push traffic from legacy Transitrail to exchange point routers. A note will be sent to the community with an explanation of the additional routes that are being exchanged and the work is to be done early on the 19th. There was a hardware failure this morning that lasted several hours but is now resolved.

A question was posed regarding where congestion points are expected and Darrell indicated that congestion is most likely to occur in public peering points. There is substantial headroom on interconnects and although the public exchange port fabric could become a problem for some connector migrations they have a solution to mitigate this. There isn't a great deal of overhead during the connector migrations but it should be manageable with the additional waves.

3. Dallas peering node

Steve Wallace sent note to Peering/Routing with some high-level bullets on the Dallas peering node. When CPS and TR consolidation plans were being made, several maps to describe a reasonable end state for the service were drawn. The notion was to expand core nodes/peering points for service and it appeared that Dallas was the next logical one for balancing of topology and number of peers. At about the time of the SMM, LEARN requested that a peering point in Dallas be considered and they indicated that they were building out their metro capacity there. LEARN offered to provide LEARN member pricing at the facility. It was already envisioned that an Internet2 core node associated with BTOP funding would be located in Dallas. It has been noted that Houston has vulnerability due to extreme weather events so Dallas is attractive as a back-up. Further, generally, most CPS traffic drains locally with the exception of the congestion on the KC to Chicago route. Adding a peering facility in Dallas allows modification to the traffic engineering and a shorter path for those in that region. The Dallas Equinix facility is growing and they also own Telx facilities so it makes sense to establish a point in Dallas. LEARN is making it affordable to do the metro piece. There is a significant budget impact of \$230-325K depending on final configuration and this will be offset by planned BTOP router hardware. The timeframe for implementation is 6-10 months. Internet2 and CENIC staff support this and we are asking the peering /routing group to concur with this recommendation.

Michael Lambert asked when peering connections of greater than 10G are expected. Darrell indicated that economies of scale for 10G port prices decreasing will push this out.

Abar asked if the proposed hardware would be able to accommodate additional 10G ports. Steve indicated that these will be PNIs rather than through public switch.

Steve indicated that end state is that connectors will have layer 2 connectivity to more than one TR-CPS router so this will allow more connectors to use this as a secondary.

Akbar estimated that they will have better latency with Dallas as more than half of LEARN's traffic funnels into Dallas (about half of the state) so this a big plus for LEARN.

Steve mentioned that there will be connectivity from LA to Atlanta. There is a plan to insert this in middle of southern TR-CPS backbone. There may be connectivity from Dallas to Chicago depending on what happens in BTOP upgrade. It is possible that Houston core node will remain and that there will also be a core node in Dallas---that is yet to be seen.

Jeff indicated that the record should show that NTAC peering/routing group had no unresolved concerns and the group recommends that the plans for a peering node in Dallas be carried to the AOAC. Jeff indicated that he supported the plan and encouraged review of a Dallas-Chicago route as part of this.

4. BTOP US UCAN network design review discussion

This agenda item was added so that the group reviewing the US UCAN network design could gather and others from the Peering/Routing group were invited to participate. Jeff indicated that he sent an email to the network design review team earlier today. Here are the points to be considered.

- Use existing NLR fiber, augment with new fiber. Try to be as diverse as possible.
- Upgrade optical layer to 100 Gbps, migrate existing national R&E Layer 1 services to new optical infrastructure. There will still be at least two parallel optical networks and multiple optical vendors.
- Migrate existing national R&E IP networks to the new Layer 1 infrastructure, and upgrade routers to support a 100 Gbps backbone and a 10 Gbps low-latency layer. (Note: Is new traffic from US UCAN being added to the existing mix of R&E IP traffic: needs clarification)?
- As proposed, US UCAN IP network consists of red and blue networks.

Given that we know, we propose focusing on the following:

- What are the pros, cons and cautions related to building a nationwide IP backbone with a separate 100 Gbps backbone and 10 Gbps low-latency (or aggregation) layer? Are there traffic engineering implications?
- How are RONS going to connect? How is resiliency improved? Are there reliability or operational issues associated with the 10G routers? Do the proposed router locations make sense? Other issues?
- What implications does the proposed design have on how the backbone connects to commodity and commercial peers? Are changes required to the way the TR/CPS has evolved? Are there any implications to how the network connects to existing exchange points? Is there a need to separate the new CAI traffic from existing R&E traffic? Other issues?
- Given the above context, and if the group can come to a consensus, make a recommendation as to whether or not the red and blue IP networks should be combined.

Dave Pokorney suggested that the group poll all of the RONS to answer some of these questions. Jeff asked what specific information should be collected- ~~separate VRFs for commodity, CPS, etc and how does hand-off occur at the RON? It was suggested that as much flexibility as possible be built into the model. Dave also suggested that phone interviews be conducted with the RONS. Chris mentioned that a general needs assessment from the RONS would be critical--is speed important, etc?~~

Paul asked about annual operating costs of US UCAN and further asked if no other members were added, would there be additional fees to the connectors? Steve Wallace mentioned that the budgeting was done such that Internet2 base fees could support US UCAN.