Proposed Architecture for a Combined CPS and TR

by the ad hoc design committee
comprised of
Dave Reese, Darrell Newcomb (CENIC)
Jeff Bartig (Internet2 NTAC Peering Working Group Chair)
David Sinn (PNWGP)
Matt Davy, Caren Litvanyi (Internet2 NOC - GRNOC at Indiana University)
Chris Robb, Steven Wallace, Rob Vietzke (Internet2)

On September 3rd, the above committee met in Indianapolis, Indiana for the purpose of designing a draft architecture for a combined Internet2 Commercial Peering Service (CPS) and CENIC Transit Rail (TR). The scope of this meeting and the draft architecture described herein are limited to technical implementation and architecture, and did not include the important details for how this infrastructure would be collaboratively managed or the the underlying business model. However, since those issues were not worked out at the time of the meeting, preference was given to technical designs that provided flexibility regarding management and operations of the service. **Throughout this paper, "CPS&TR" will refer to the combined service.**

The design constraints the committee was given for combining these services included:

1. The ability to manage and operate CPS&TR as a separate infrastructure and a separate autonomous system from the Internet2 R&E network.
2. Maintain the current access mechanism, or very similar, for existing TR and CPS users.
3. Provide the union of the traffic offload value of CPS and TR to users of CPS&TR.
4. Leverage existing infrastructures where possible.
5. Provide for a rapidly growing service.

The consensus of the committee was that a design containing the following features best met the given constraints:

1. A layer3 router would be located at each exchange point facility. These may be new devices, the existing TR routers, or some combination thereof.
2. All exchange point routers will operate in a CPS&TR ASN. CPS&TR will operate under the same ASN as the current TR (i.e., 11164). AS11164's ownership will be transferred from CENIC to Internet2.
3. Circuits dedicated to CPS&TR will interconnect the exchange point routers. This will comprise the CPS&TR backbone. These circuits will be provided by Internet2, NLR, and CENIC. <*see diagram*>
4. Customers will connect to the CPS&TR backbone routers primarily via layer 2 services over the Internet2 Network or NLR. Direct customer circuits to the CPS&TR routers are also possible. <*see diagram*>
    1. Current TR users may continue to use their existing transport to the CPS&TR backbone.
    2. Internet2 users may continue to use their existing circuit(s) to the Internet2 Network for carrying CPS&TR traffic as well as R&E traffic. The transport to the CPS&TR backbone will be provisioned as layer2 virtual circuits over the existing Internet2 Network, interconnecting the Internet2 user's router(s) with one or more exchange-point-located CPS&TR routers. Internet2 users will receive this transport as additional vLANs on their existing connection (Internet2 will investigate how to best apply this architecture to non-Ethernet connected users).

5. Within the confidentiality constraints of a few commercial peering agreements, the CPS&TR infrastructure will be managed by the similar visible network tools as the Internet2 Network, providing similar levels of transparency.

For current TR users, this proposed architecture would require little to no change in configuration or operation.

For current CPS users, this proposed architecture would require the following changes:

1. CPS&TR would be delivered over two vLANs, unless the customer explicitly chooses only one (more than two would be negotiable). Similar to CPS today, these vLANs would be provisioned over the customer's existing Internet2 Network connection. Each vLAN would terminate on a diverse CPS&TR exchange point router. Transport for each vLAN would be provided via an MPLS layer2 point-to-point VPN. The purpose of multiple connections to CPS&TR is to provide resiliency for the customer.
2. For most, two new BGP sessions will need to be configured on the customer's router(s), and associated policy defined for their use - load balancing, failover, etc.
3. CPS&TR would operate as AS11164, not 11537.
4. To insure the Internet2 Network and CPS&TR are able to reasonably manage their traffic engineering and capacity planning, greater coordination between Internet2, CPS&TR, and the Connector will be required in deciding how the multiple CPS&TR layer2 circuits are routed over the Internet2 backbone.

This proposed architecture requires that Internet2 and the GRNOC design, implement, and operate as a production service MPLS LDP-signalled layer 2 circuits for each Connector, likely over an RSVP-signalled MPLS infrastructure to provide traffic engineering. Internet2 has some experience with this type of service (e.g, layer 2 VPNs were used in the HOPI network), however implementing this as a production service over the entire footprint will be a substantial effort. It's likely this effort will offer the Internet2 community additional service options in the future.

Combining CPS and TR into CPS&TR represents substantial engineering, implementation, and operation efforts and costs. During the September 3rd meeting we were able to discuss transition strategies for both combining CPS and TR as well as migrating users from their current service to the combined service. Assuming combining CPS and TR moves forward, we can take a first step by some carefully managed inter-connection of CPS and TR for the purpose of expanding each service's routing table. Additional capacity must be provisioned at some cost in order to provide this. We also will be committed to providing CPS and TR users with the technical support they may require for a successful transition to CPS&TR, should this project move forward.

CPS/TR Proposed Architecture

22-Sept-2009

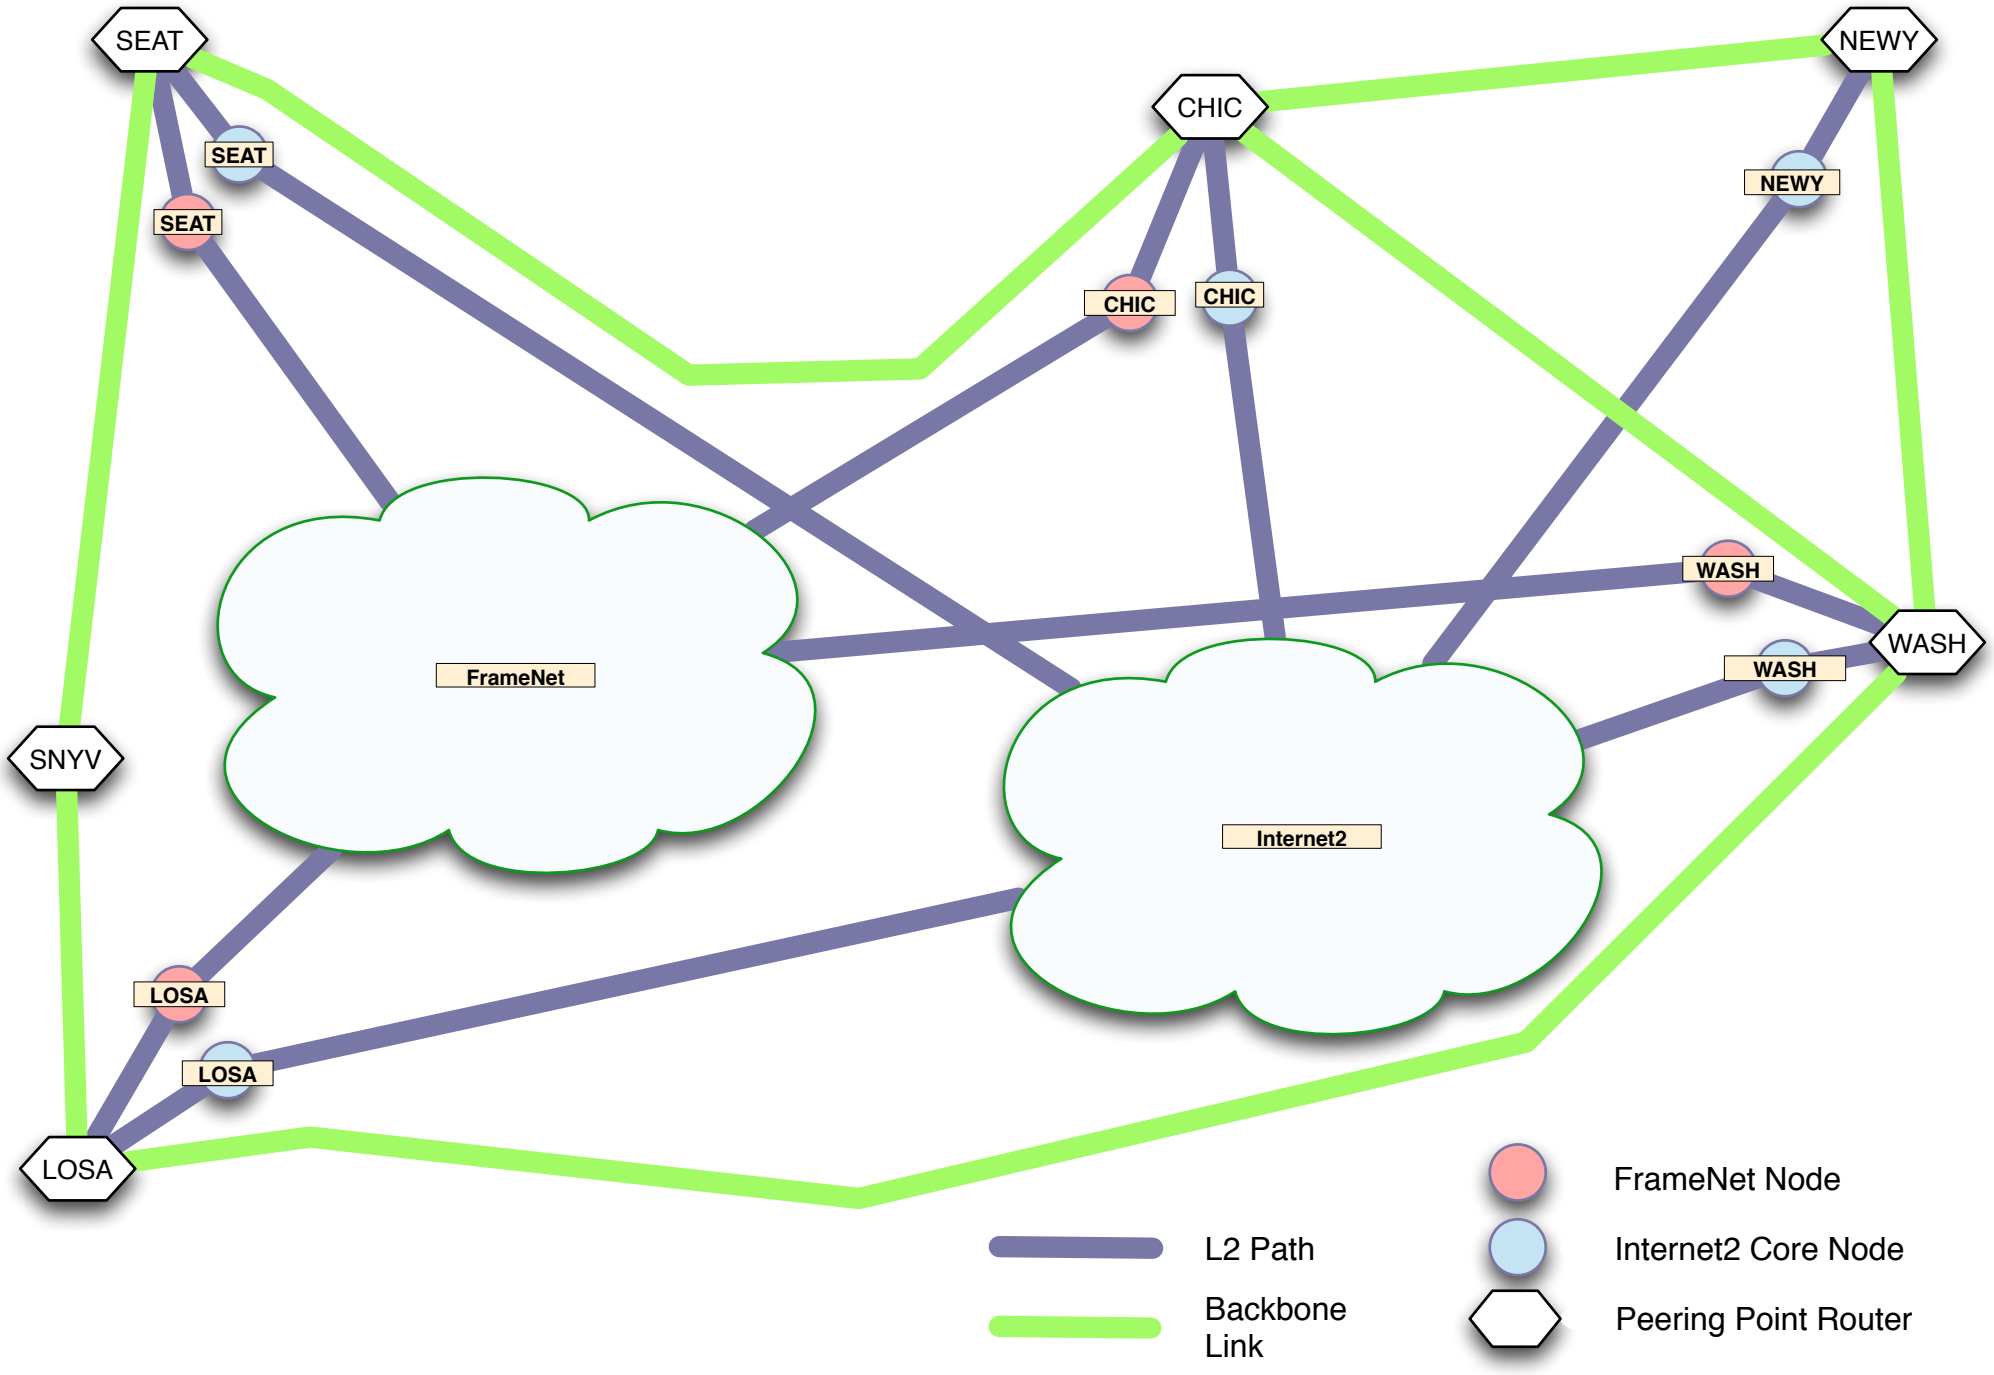

CPS/TR Connector Examples

22-Sept-2009