January 2010 Internet2 Traffic Snapshot

Prepared for the Joint Techs Connector BOF
On February 1st, 2010

Executive Summary

Internet2 consistently monitors and projects backbone IP network utilization to maintain a high performance uncongested national network for use by the research and education community. Due to growing demands on the Internet2 IP network from both researchers and the continued success of the Commodity Peering Service, Internet2 determined that the following actions would likely need to be implemented in the coming months:

- More efficient implementation of load balancing methodologies that recognize the reality of >5G R&E traffic flows
- Upgrade of the entire backbone to 20G between nodes
- Provide an updated backbone headroom document to the community that reflects the differential treatment of CPS-focused and R&E-focused backbone interconnects

Document Purpose

Internet2 has prepared this document in early February 2010 to provide a review of the available capacity of the Internet2 Network over the prior 4 months.

Background Information

There are two large classes of bandwidth usage on the backbone that are treated with different traffic engineering principals: Research and Education (R&E) and Commodity Peering Service (CPS).

In general, these two classes of traffic are blended together across shared interconnects between backbone routers. CPS traffic is carried in a full mesh of MPLS tunnels that provide the capability to route CPS Connectors and Peers in a separate routing table from the R&E-only Internet2 Connectors. In cases where Internet2 augments capacity between adjacent routers, the two classes are influenced toward separate parallel backbone interconnects. Internet2 has a published backbone upgrade practice of upgrading blended backbone interconnects when the weekly 95th-percentile metric is reliably above 30% of link capacity (3Gbps).[1] When a second 10G interconnect between routers is put in place, traffic is engineered such that CPS traffic is strictly routed over the new link and R&E traffic is loosely routed over the existing 10G link. In a failure scenario, R&E traffic will migrate to the CPS-focused backbone link, but CPS

---

1 https://wiki.internet2.edu/confluence/download/attachments/17383/Internet2+Headroom+Practice+8-14-08.pdf?version=1

traffic would not migrate to the R&E-focused backbone link. This is done to maintain adequate headroom for R&E traffic.

The diagram below illustrates the distribution of Internet2 IP Network backbone links and their support capabilities.

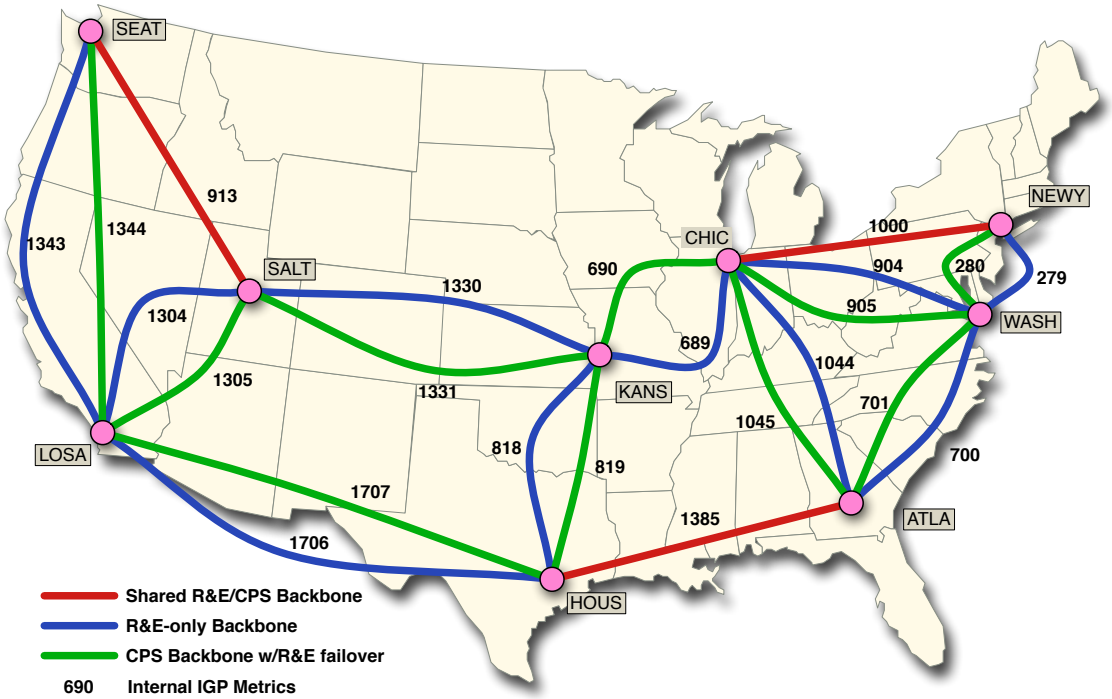

Internet2 Backbone Map with CPS/R&E Capabilities Highlighted

For parallel backbone interconnects that primarily carry CPS backbone traffic, Internet2 has, in practice, augmented the backbone when the weekly 95th percentile metric is reliably above 60% of link capacity (6Gbps).

Overall Traffic Load

The following section details overall Internet2 IP traffic loads. It takes into consideration *both* R&E and CPS-related traffic.

Connector Interface Utilization Breakdown

The table below details the aggregate peak IP load for the Internet2 Connector between January 25th, 2010 and February 1st, 2010.[2] Maximum values represent an average over 10-second intervals. The "Node Bandwidth" column represents the available bandwidth into that city from the rest of the Internet2 Network.

Connector	Speed	To Connector	From Connector	Router	Node Bandwidth
Indiana Gigapop	10Gbps	3.2G	1.7G	ATLA	50Gbps
KyRON	10Gbps	860M	200M	ATLA	50Gbps
SOX	1Gbps	560	160M	ATLA	50Gbps
MCNC	10Gbps	3.4G	833M	ATLA	50Gbps
USF/FLR	1Gbps	760M	360M	ATLA	50Gbps
MREN	2.5Gbps	1.5G	870M	CHIC	70Gbps
CIC	10Gbps	6.4G	2.2G	CHIC	70Gbps
Memphis	1Gbps	320M	60M	CHIC	70Gbps
LONI	10Gbps	1.8G	5.2G	HOUS	50Gbps
LEARN	10Gbps	810M	410M	HOUS	50Gbps
GPN	10Gbps	9.0G	6.8G	KANS	60Gbps
CENIC	10Gbps	2.9G	1.9G	LOSA	60Gbps
Oregon Gigapop	2.5Gbps	1.1G	840M	LOSA	60Gbps
Nysernet	1Gbps	960M	570M	NEWY	30Gbps
NOX	10Gbps	3.0G	8.4G	NEWY	30Gbps
MAGPI	2.5Gbps	830M	570M	NEWY	30Gbps
UEN	10Gbps	1.8G	960M	SALT	50Gbps
PNWGP	10Gbps	450M	290M	SEAT	30Gbps
Merit	10Gbps	8.6G	2.0G	WASH	60Gbps
MAX	2.5Gbps	920M	2.5G	WASH	60Gbps
Drexel	1Gbps	348M	70M	WASH	60Gbps
3ROX	1Gbps	382M	100M	WASH	60Gbps

R&E Exchange Point Utilization Breakdown

The table below details the aggregate peak IP load for the Internet2 interconnects to the Research and Education Exchange Points between January 25th, 2010 and February 1st, 2010.[3] Maximum values represent an average over 10-second intervals.

Exchange Point	Speed	To XP	From XP	Router	Node Bandwidth
Starlight	10Gbps	2.3G	1.9G	CHIC	70Gbps
Pacific Wave	10Gbps	6.5G	2.2G	LOSA	60Gbps
MANLAN	10Gbps	1.5G	1.2G	NEWY	30Gbps

[2] Graphs of connections to Internet2 Connectors can be accessed at the following URL: http://dc-snmp.wcc.grnoc.iu.edu/i2net/show-view-graphs.cgi?viewname=Internet2-Connectors-Aggregate&cmd=graphweek

[3] Graphs of connections to R&E Exchange Points can be accessed at the following URL: http://dc-snmp.grnoc.iu.edu/i2net/show-view-graphs.cgi?viewname=RE-Exchange-Points-Aggregate&cmd=graphweek

Pacific Wave	10Gbps	1.3G	1.8G	SEAT	30Gbps
NGIX-East	10Gbps	3.9G	4.2G	WASH	50Gbps
NGIX-West	1Gbps	1 Mbps	500Kbps	SALT	50Gbps

Internet2 Backbone Analysis

The table below details the aggregate peak IP load for the Internet2 backbone links between January 25th, 2010 and February 1st, 2010.[4] Maximum values represent an average over 10-second intervals.

City A	Peak Load	CityB
Atlanta	2.6G →	Chicago
	← 2.2G	
Atlanta	1.6G →	Washington
	← 3.7G	
Atlanta	3.5G →	Houston
	← 3.6G	
Atlanta	2.6G →	Washington
	← 5.2G	
Atlanta	652M →	Chicago
	← 600M	
Chicago	1.7G →	New York
	← 2.6G	
Chicago	8.0G →	Kansas City
	← 5.4G	
Chicago	1.9G →	Washington DC
	← 1.4G	
Chicago	4.8G →	Washington DC
	← 5.2G	
Houston	7.1G →	Los Angeles
	← 3.9G	
Houston	3.2G →	Kansas City
	← 802M	
Houston	556M →	Los Angeles
	← 1.8G	
Kansas City	1.9G →	Salt Lake City
	← 2.2G	
Kansas City	860M →	Salt Lake City
	← 1.9G	
Los Angeles	504M →	Seattle
	← 425M	
Los Angeles	1.3G →	Salt Lake City

[4] Graphs of backbone interconnects can be accessed at the following URL: http://dc-snmp.wcc.grnoc.iu.edu/i2net/show-view-graphs.cgi?viewname=backbone-interfaces&cmd=graphweek

	← 364M	
Los Angeles	306M →	Seattle
	← 311M	
Los Angeles	1.2G →	Salt Lake City
	← 235M	
New York	1.4G →	Washington DC
	← 1.6G	
New York	4.1G →	Washington DC
	← 3.35G	
Salt Lake City	1.8G →	Seattle
	← 2.9G	

Commodity Peering Service Traffic

In general, Internet2 is an “eyeball” network- that is, Internet2 Connectors are generally consumers of data from Commodity Peers rather than providers of data. Practically, this means that Internet2 will *generally* see a greater amount of traffic coming into its network from Commodity Peers than it sees going out of its network from CPS Connectors.

CPS Connector Interconnects

Internet2 maintains settlement free interconnects (SFI) with commodity peers at five major cities on its national footprint:

- Los Angeles (PAIX Palo Alto)
- Seattle (SIX)
- Chicago (Equinix)
- New York (PAIX)
- Washington DC (Equinix Ashburn)

In general, Internet2 Connectors that receive Commodity Peering Service in these cities will exchange most of their traffic with commodity networks within that same city. CPS Connectors on routers other than those listed above will generally exchange outbound traffic via the closest available peering with the commodity peer. Internet2 uses route miles as its internal routing metric to direct Connector-to-Peer traffic to the nearest commercial exchange point. While Internet2 attempts to balance peers across exchange points to make a local traffic drain possible, it’s sometimes not possible (a peer is not at an east-coast commercial exchange, for example). In the reverse direction, CPS Peer-to-Connector, Internet2 has less control over the path the traffic takes through the network. A peer may decide to offload traffic to an east coast CPS Connector on the west coast, for example. In general, Internet2 sees that peers tend to prefer the closest geographic exit point, but that’s highly

dependent on the traffic flow destination and the balancing of the Peer networks internal infrastructure.

The map below depicts the Internet2 CPS backbone network and the distribution of Internet2 Connectors and Layer2 Participants that take advantage of CPS.[5]

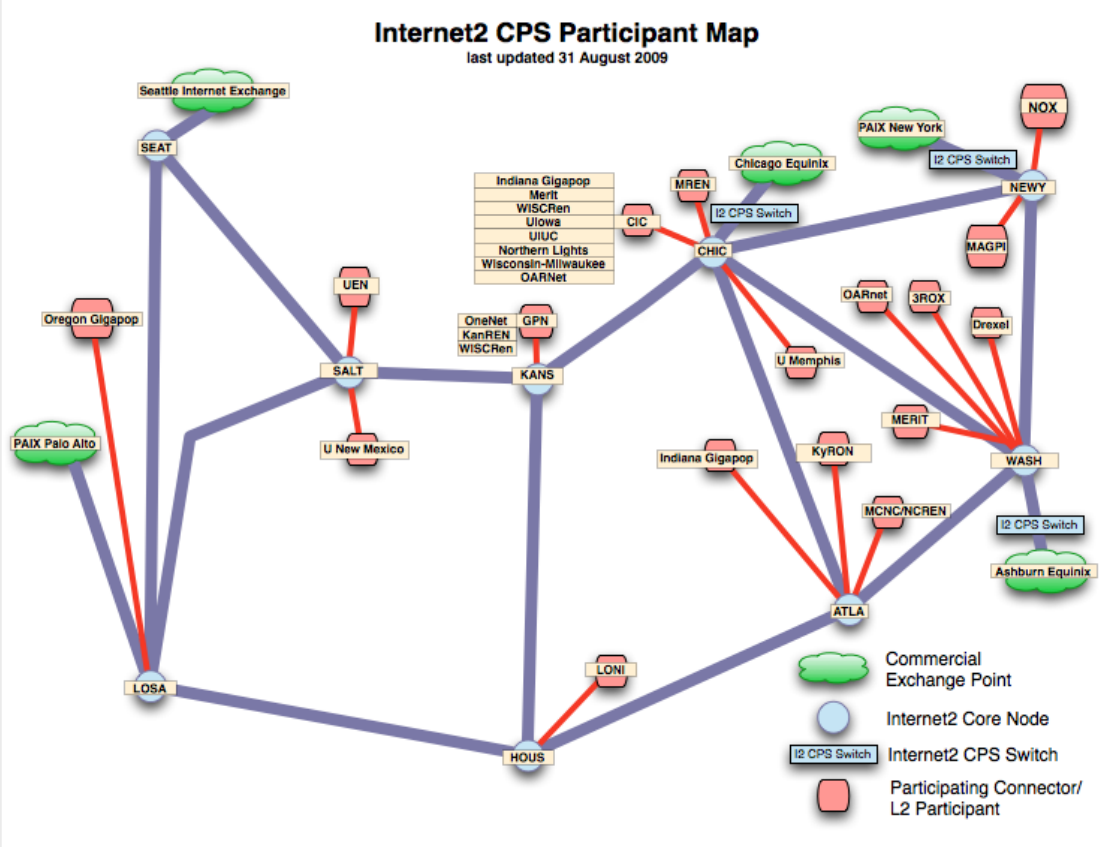

Internet2 CPS Participant Map

Internet2 has several large CPS Connectors connected to routers that don't have a local commercial exchange. *All* CPS traffic from these Connectors will traverse the Internet2 backbone to reach CPS Peers. The table below shows the peak values for CPS Participants that are homed to routers without local commercial peer connectivity. The peak values are representative of the highest measured 10 second average value measured between January 25th, 2010 and February 1st, 2010.

CPS Participant	To Connector	From Connector	Router	Node BW
Utah Education Network	923M	122M	SALT	50G
Front Range Gigapop	255M	98M	SALT	50G

[5] For a current version of this map, visit: http://noc.net.internet2.edu/i2network/commercial-peering-service/maps--documentation/internet2-cps-network-topology.html

ONEnet	617M	747M	KANS	60G
AREON	298M	110M	KANS	60G
KanREN	1.2G	184M	KANS	60G
UNL	2.2G	197M	KANS	60G
LONI	1.3G	300M (approx)	HOUS	50G
MCNC	2.6G	587M	ATLA	60G
South Carolina Gigapop	8.5M	3.0M	ATLA	60G
Indiana Gigapop[6]	3.0G	1.3G	ATLA	60G
KyRON	662M	128M	ATLA	60G

The data above suggests that Kansas City is the largest potential consumer of CPS *backbone* traffic, with most of the traffic likely coming from or going to the nearby Chicago Equinix exchange point. Indeed, the first Internet2 backbone augment was between Kansas City and Chicago. Atlanta provides additional CPS-related backbone potential, most likely on the backbone link to Washington DC. Internet2 augmented its backbone between Atlanta and Washington DC and between New York and Washington DC in early 2009.

The table below shows the peak values of all the CPS connectors that are homed to routers with local commercial peering exchange interconnects.[7] The peak values are representative of the highest measured 10 second average value measured between January 25th, 2010 and February 1st, 2010. CPS traffic to/from these Connectors are *more likely* to drain within the local router, though some traffic transiting backbone links to get to other commercial exchange points is expected and normal.

CPS Participant	To Connector	From Connector	Router	Node BW
Merit via MREN	0M	0M	CHIC	70G
MREN	1.2G	558M	CHIC	70G
WiscREN via CIC	0M	0M	CHIC	70G
U Minnesota via CIC	722M	334M	CHIC	70G
Wisc-Milwaukee via CIC	387M	625M	CHIC	70G
UIUC via CIC	722M	334M	CHIC	70G
U of Iowa via CIC	819M	320M	CHIC	70G
Indiana Gigapop via CIC	15M	5M	CHIC	70G
Merit via CIC	433M	97M	CHIC	70G
UIC via CIC	0M	0M	CHIC	70G
Oarnet via CIC	34M	31M	CHIC	70G
U of Memphis	275M	46M	CHIC	70G
Oregon Gigapop	1.1G	632M	LOSA	60G
NOX	7.3G	2.0G	NEWY	30G
MAGPI	0M	1M	NEWY	30G
Merit	2.7G	517M	WASH	60G
OARnet via Merit	5.3G	1.2G	WASH	60G

[6] The Indiana Gigapop sources most of its CPS traffic in Atlanta, but receives most of its CPS traffic via its CIC interconnect in Chicago.

[7] These graphs can be accessed here: http://dc-snmp.wcc.grnoc.iu.edu/i2net/show-view-graphs.cgi?viewname=CPS-Connectors-Local-Drainage&cmd=graphweek

Drexel University	0M	0M	WASH	60G
3ROX	325M	30M	WASH	60G

While most of the traffic from the above networks is expected to remain within the local core node, an outage on the local commercial exchange point or on one of the large private network interconnects could shift a significant portion of traffic to backbone links as the peer connectivity re-routes to alternate commercial exchanges.

CPS Backbone Traffic

Internet2 monitors router-to-router CPS backbone traffic in a full mesh to provide heuristics for backbone augments.[8] A scan of the afternoon hourly averages on Monday, February 1st show the following router-to-router paths with utilization levels above 1Gbps. Note that the transit shown between city pairs below may actually transit multiple backbone links:

From	To	Average
Chicago	Atlanta	1.1G
Washington DC	Atlanta	2.4G
Washington DC	Chicago	1.4G
Chicago	Kansas City	2.0G
Washington DC	New York	1.7G
Atlanta	Washington DC	1.2G
Chicago	Washington DC	1.8G
New York	Wasington DC	2.4G

The data above is largely expected. Most traffic is sourcing from Chicago and Washington DC, two of the most heavily trafficked commercial exchange points. When traffic is exchanged between one city with a commodity exchange point connection and another city with a commodity exchange point connection it is typically the result of one of the following scenarios:

- Internet2 may not peer with the same sets of peers at each of the locations (because the peer isn't available or doesn't have sufficient bandwidth)
- Peers may have opted to offload traffic in a location that's further away from the CPS connectors

The introduction of the Ashburn exchange point in late July 2009 has alleviated much of the pressure off the New York and Chicago commercial

[8] http://dc-snmp.grnoc.iu.edu/i2-mpls/rrd.cgi?time=hour

exchange interconnections by providing local CPS drainage to CPS Connectors on the Washington DC router. As a result, Internet2 is seeing less CPS-related backbone traffic between Washington DC and New York. Prior to the addition of Ashburn, Internet2 was experiencing CPS backbone loads that peaked at 8-9Gbps between Washington DC and New York.

Commodity Peer Interconnects

As mentioned above, Internet2 maintains settlement free interconnects (SFI) with commodity peers at five major cities on its national footprint:

- Los Angeles (PAIX Palo Alto)
- Seattle (SIX)
- Chicago (Equinix)
- New York (PAIX)
- Washington DC (Equinix Ashburn)

In addition to the commercial exchange point switch in each of the cities above, Internet2 maintains several direct Private Network Interconnects (PNIs) to large commercial peer networks when appropriate and technically feasible. These PNIs are often backhauled to the Internet2 core-node via a separate 10GigE interconnects to make more bandwidth available between Internet2 and the commercial exchange point switch. Internet2 analyzes netflow and MAC-based statistics to determine which peers on the commercial exchange will benefit from a PNI. Often, these statistics are used as justification to the commercial network to commit resources to a direct interconnect (interface costs, fiber cross-connects, etc.).

The graphs below depict the weekly utilization on the five commercial exchange point switch interconnects for the week of January 25th, 2010 and February 1st, 2010.[9]

[9] These graphs can be seen at http://noc.net.internet2.edu/i2network/commercial-peering-service/maps--documentation/internet2-cps-utilization-graphs.html

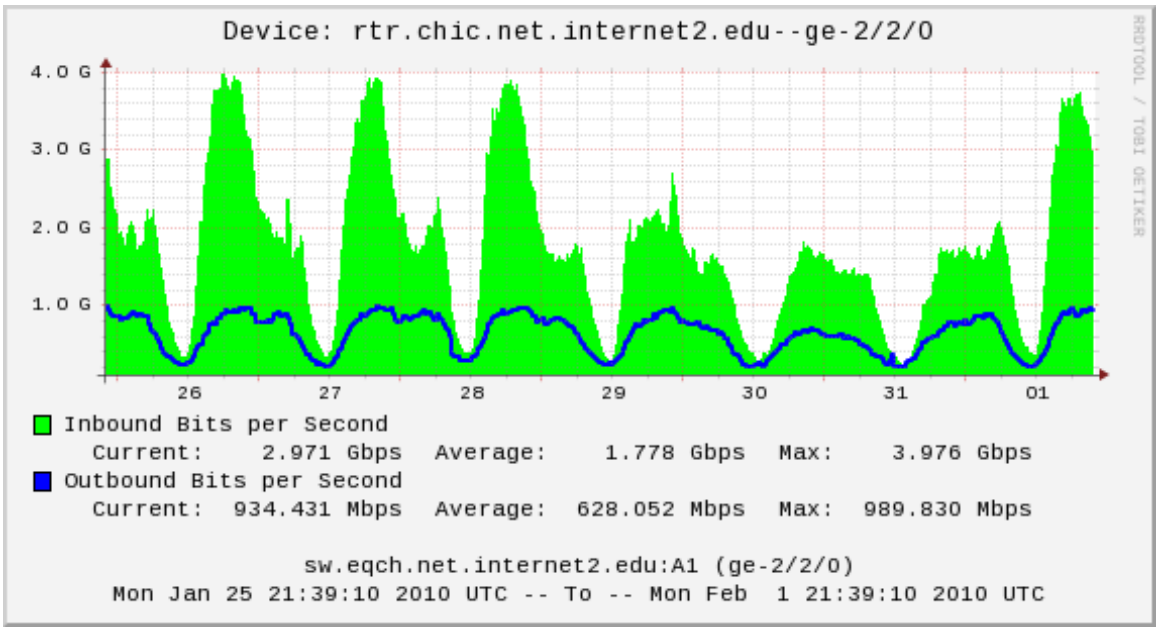

Chicago Equinix Exchange Point

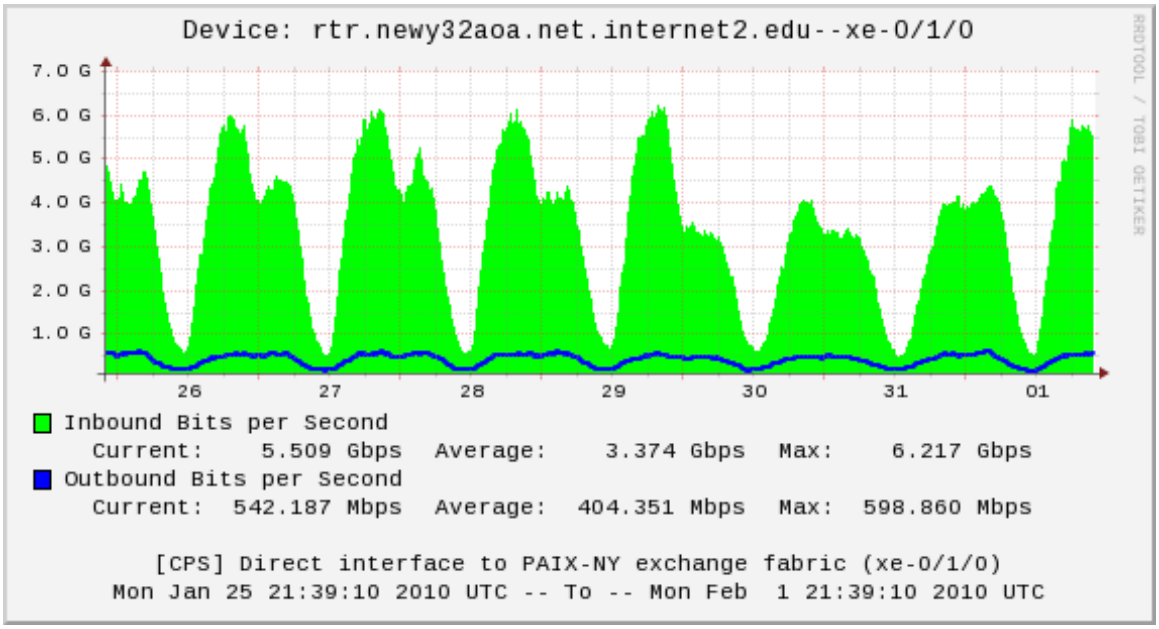

New York PAIX Exchange Point

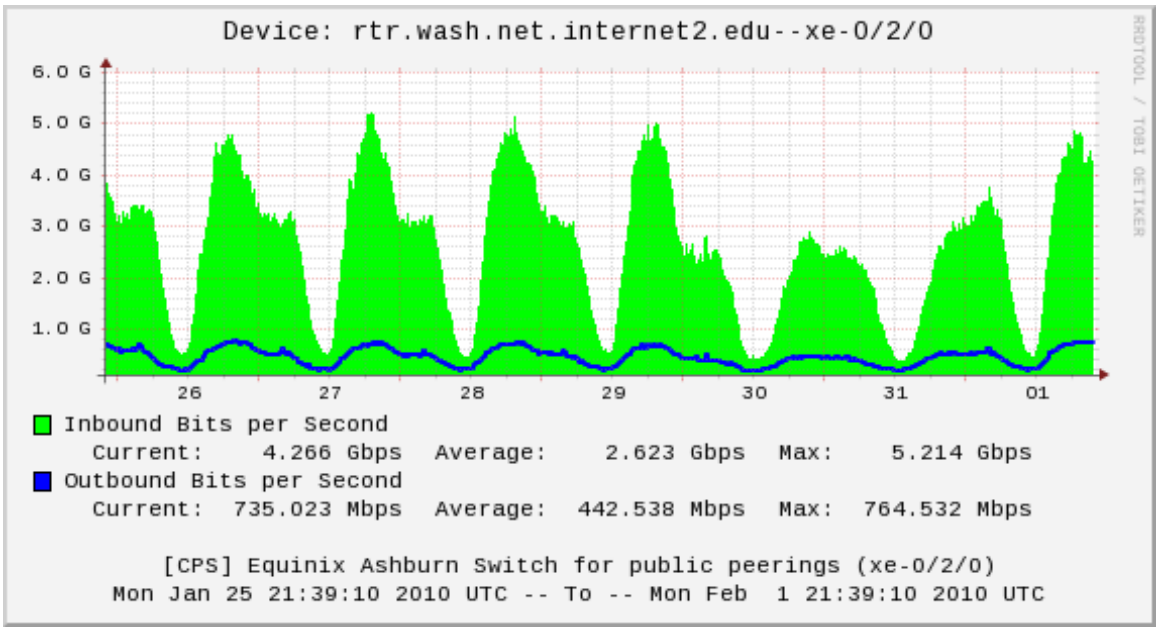

Washington DC Equinix Exchange Point

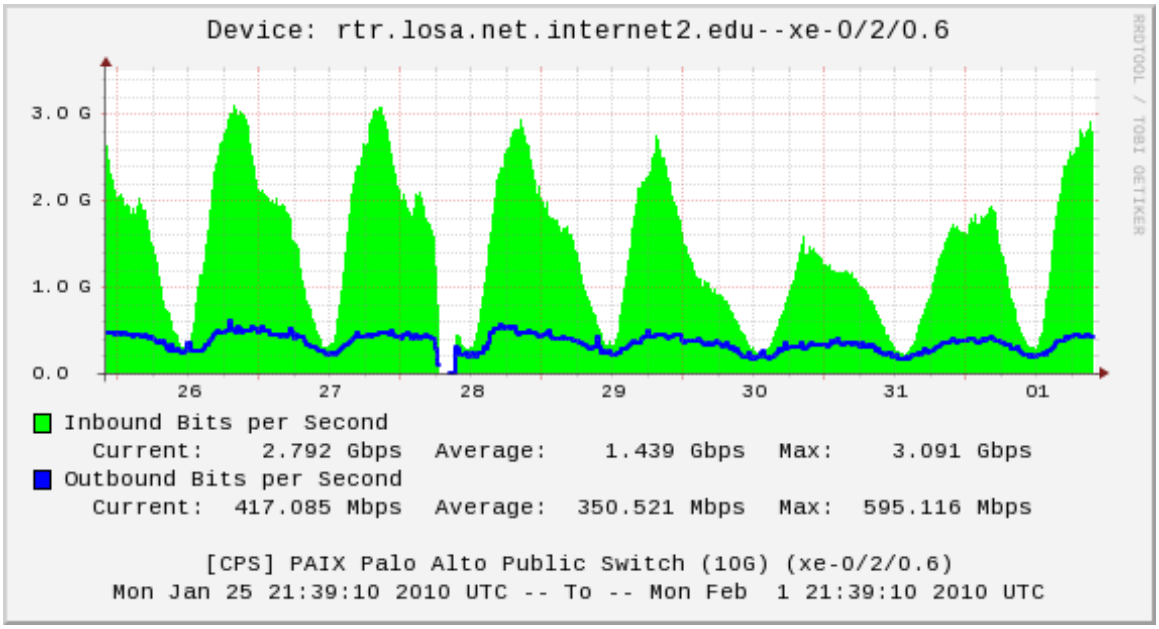

Los Angeles PAIX Palo Alto

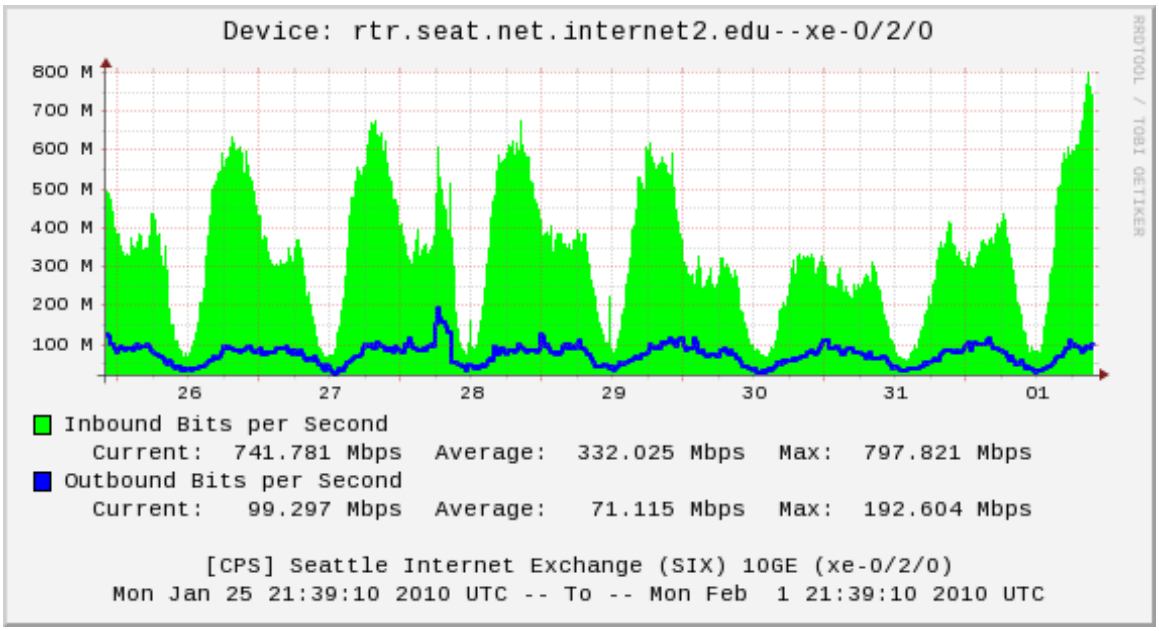

Seattle SIX

The graphs above indicate a concentration of traffic on the east coast exchange points. This is unsurprising given that the majority of large CPS Connectors are concentrated in the eastern half of the network and many commercial peers have concentrated facilities in those locations.

TransitRail Interconnects

As of January 30th, 2010, Internet2 and TransitRail have entered into an initial peering arrangement that will allow participants of both services to enjoy the benefits of the combined sets of commercial peer routes before the services themselves are combined. During the initial period, TransitRail and Internet2 will share their respective commercial peer base with each other. TransitRail and Internet2 will distribute these additional routes to their commercial participants, thus providing an increase in the overall commercial reach.

On January 30th, 2010, Internet2 and TransitRail began exchanging a subset of routes across 10GigE interconnects in the following cities:

- Los Angeles
- Chicago
- Seattle
- Washington DC

In order to avoid overloading the 10G interconnects, a subset of each networks routes were exchanged. The expectation is that the following list of networks will grow as capacity allows:

AS11164/TransitRail -> AS11537/CPS (~50k prefixes):

- Abovenet(AS6461)
- Charter(22212)
- ChinaTelecom(4134)
- Cogent(174)
- Comcast(7922)
- DaCom(3786)
- Earthlink(4355)
- Frontier(5650)
- IIJ (AS2497/Japan)
- Integra Telecom(7385)
- TWCable(7843)

AS11537/CPS -> AS11164/TransitRail (~2k prefixes initially):

- Rogers Communications(812)
- Partial Global Crossing (3549)

While it's early in the process to derive any long-term traffic trends, the initial data suggests that Internet2's CPS participants are producers of data for the list of TransitRail peer networks rather than consumers of data- as is generally the case with the Internet2 direct peers. The graph below shows a daily snapshot of the aggregation of all the data traversing the four TransitRail-CPS interconnects

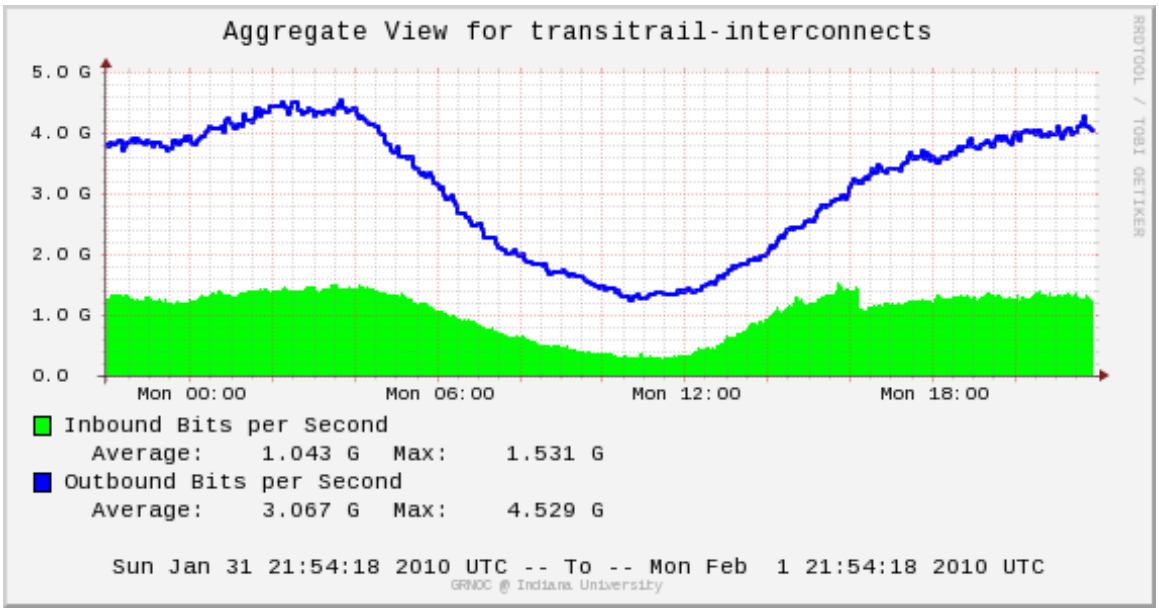

Aggregate graph of four TransitRail-CPS Interconnects

In general, there will be two types of traffic crossing the 10G links between TransitRail and Internet2 CPS:

- CPS Customer <-> Transitrail Peer

- TransitRail Customer <-> CPS Peer

Given the amount of broadband and international peers being shared by TransitRail, the greater emphasis on Connector traffic toward the TransitRail Peers makes sense and is in line with expectations.

LHC Traffic

The methodology for measuring LHC traffic is still evolving. Internet2 is working with its partners to determine the best way to identify LHC-related traffic. The graph below shows an extrapolation of sampled netflow data to and from FERMI National Lab and Brookhaven National Lab as identified by their AS numbers. The timescale is the month ending February 1st, 2010.

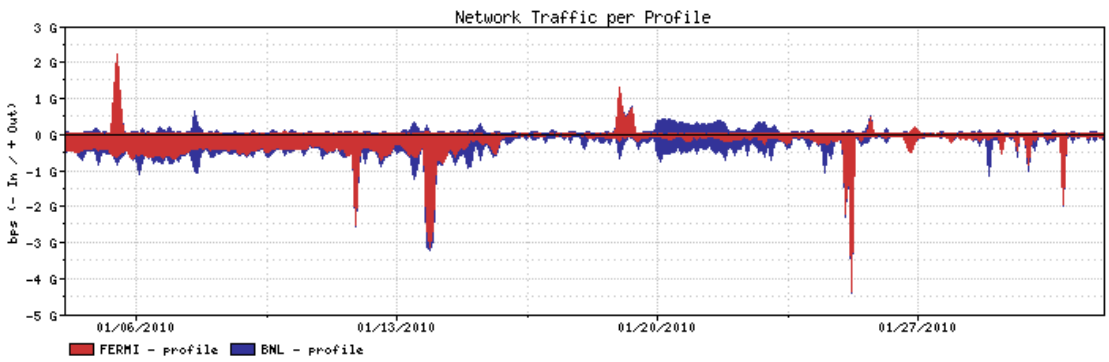

Over the month of January 2010, the traffic levels were as follows:

Network	Average To Internet2	Average From Internet2	Maximum From	Maximum To
Fermi	307M	45M	4.4G	2.2G
BNL	129M	63M	1.0G	642M

ESNet was able to correlate network traffic to FNAL during a 2.36 TeV run of the LHC on December 13th. Their graph is below:

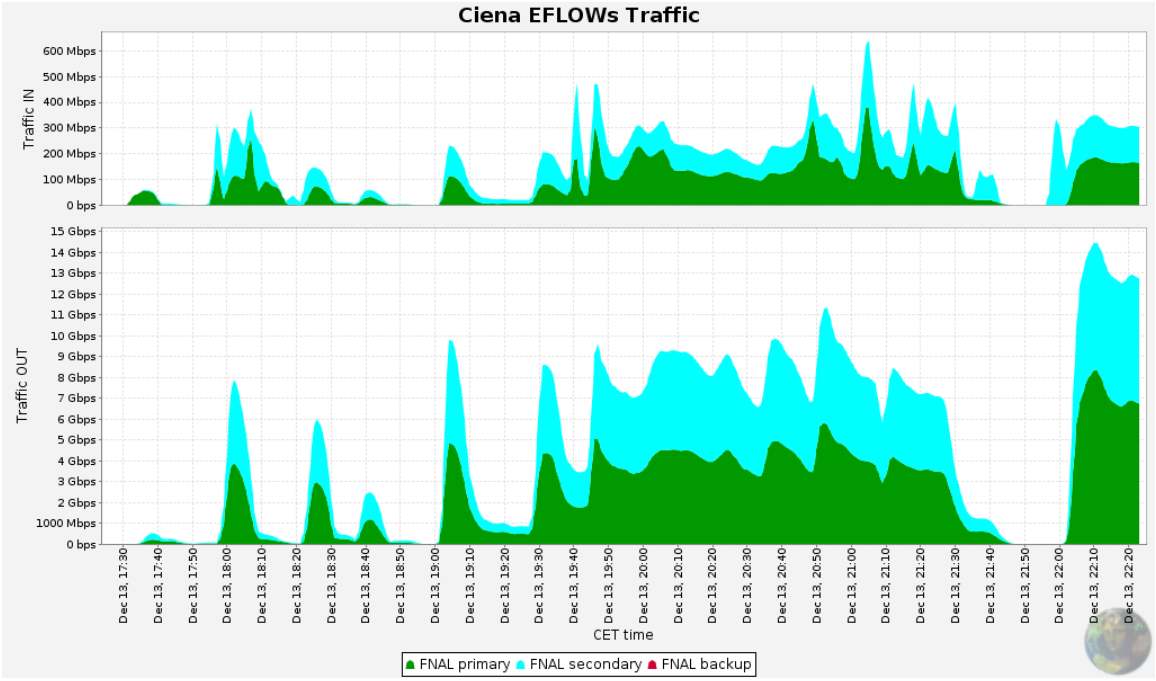

Internet2 didn't see nearly the amount of data to FERMI, but it did exist:

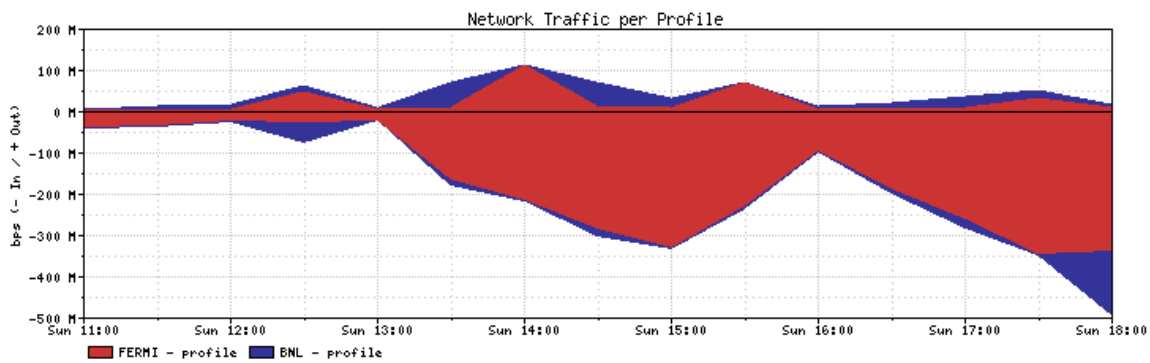

The peak traffic from FERMI over the same time period was 343Mbps. Some further work will need to be done to determine on which Connector interfaces the traffic drained.