

Back to School: A Snapshot of the Internet2 IP Network as Students Return to Class

**Prepared for the Internet2 AOAC
September 1st, 2009**

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:

- Augmenting the IP connections to 20Gbps for the Los Angeles to Houston and Houston to Kansas City links
- Aggressively monitor the Ashburn Exchange point and continue to offload traffic from large peers to PNIs as appropriate
- 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 late August 2009 to provide a review of the capacity state of the Internet2 IP Network as most Connectors prepare to support increased traffic loads associated with returning students and faculty.

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).

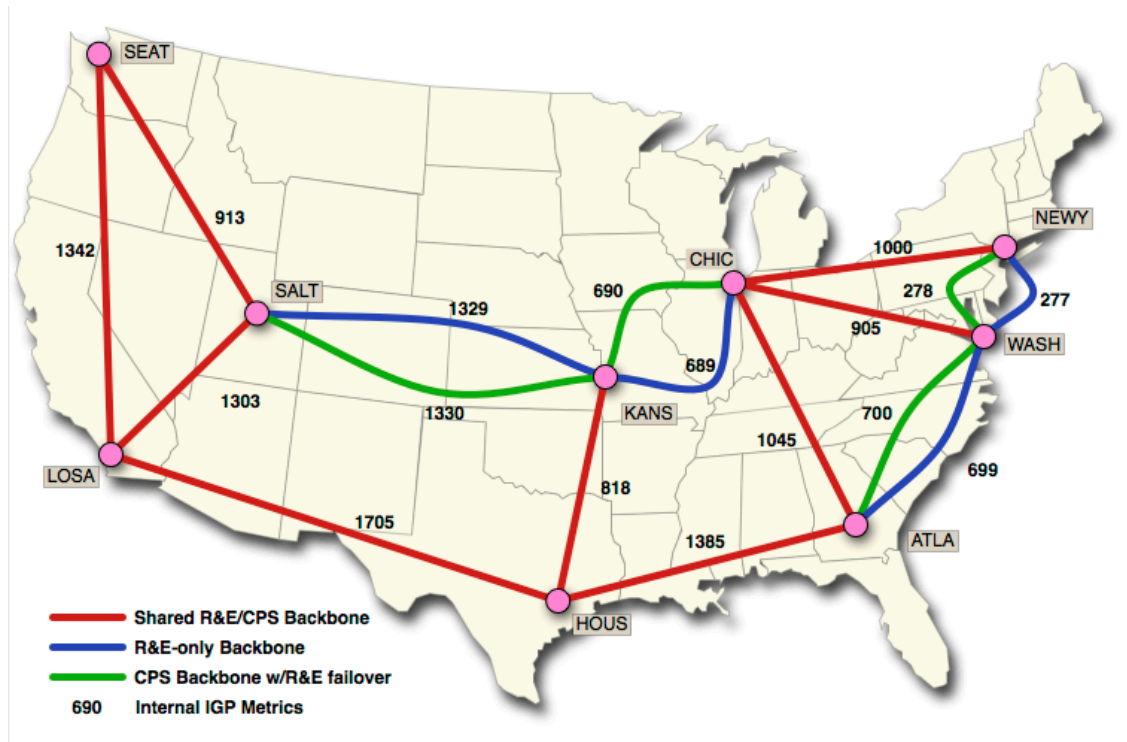
In general, these two classes of traffic are blended together across shared interconnects between backbone routers. CPS traffic is carried in a full mesh of MPLS tunnels that provide the capability to route CPS Connectors and Peers in a separate routing table from the R&E-only Internet2 Connectors. In 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

1

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

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 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 during the week of August 24th.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	480Mbps	1.39Gbps	ATLA	40Gbps
KyRON	10Gbps	635Mbps	144Mbps	ATLA	40Gbps
SOX	1Gbps	346Mbps	243Mbps	ATLA	40Gbps
MCNC	10Gbps	3.27Gbps	261Mbps	ATLA	40Gbps
USF/FLR	1Gbps	585Mbps	169Mbps	ATLA	40Gbps
MREN	2.5Gbps	823Mbps	737Mbps	CHIC	50Gbps
Nysernet	1Gbps	199Mbps	177Mbps	CHIC	50Gbps
CIC	10Gbps	5.75Gbps	1.94Gbps	CHIC	50Gbps
Memphis	1Gbps	158Mbps	104Mbps	CHIC	50Gbps
LONI	10Gbps	675Mbps	613Mbps	HOUS	30Gbps
LEARN	10Gbps	824Mbps	339Mbps	HOUS	30Gbps
GPN	10Gbps	6.56Gbps	4.5Gbps	KANS	50Gbps
CENIC	10Gbps	1.8Gbps	826Mbps	LOSA	30Gbps
Oregon Gigapop	2.5Gbps	590Mbps	597Mbps	LOSA	30Gbps
Nysernet	1Gbps	278Mbps	256Mbps	NEWY	30Gbps
NOX	10Gbps	4.19Gbps	2.31Gbps	NEWY	30Gbps
MAGPI	2.5Gbps	549Mbps	447Mbps	NEWY	30Gbps
U New Mexico	1Gbps	487Mbps	251Mbps	SALT	40Gbps
UEN	10Gbps	965Mbps	573Mbps	SALT	40Gbps
PNWGP	10Gbps	271Mbps	271Mbps	SEAT	20Gbps
Merit	10Gbps	5.37Gbps	1.13Gbps	WASH	50Gbps
MAX	2.5Gbps	913Mbps	2.47Gbps	WASH	50Gbps
Drexel	1Gbps	116Mbps	42Mbps	WASH	50Gbps
3ROX	1Gbps	907Mbps	219Mbps	WASH	50Gbps

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 during the week of August 24th.³ Maximum values represent an average over 10-second intervals.

Exchange Point	Speed	To XP	From XP	Router	Node Bandwidth
Starlight	10Gbps	4.75Gbps	4.01Gbps	CHIC	50Gbps
Pacific Wave	10Gbps	3.89Gbps	3.24Gbps	LOSA	30Gbps

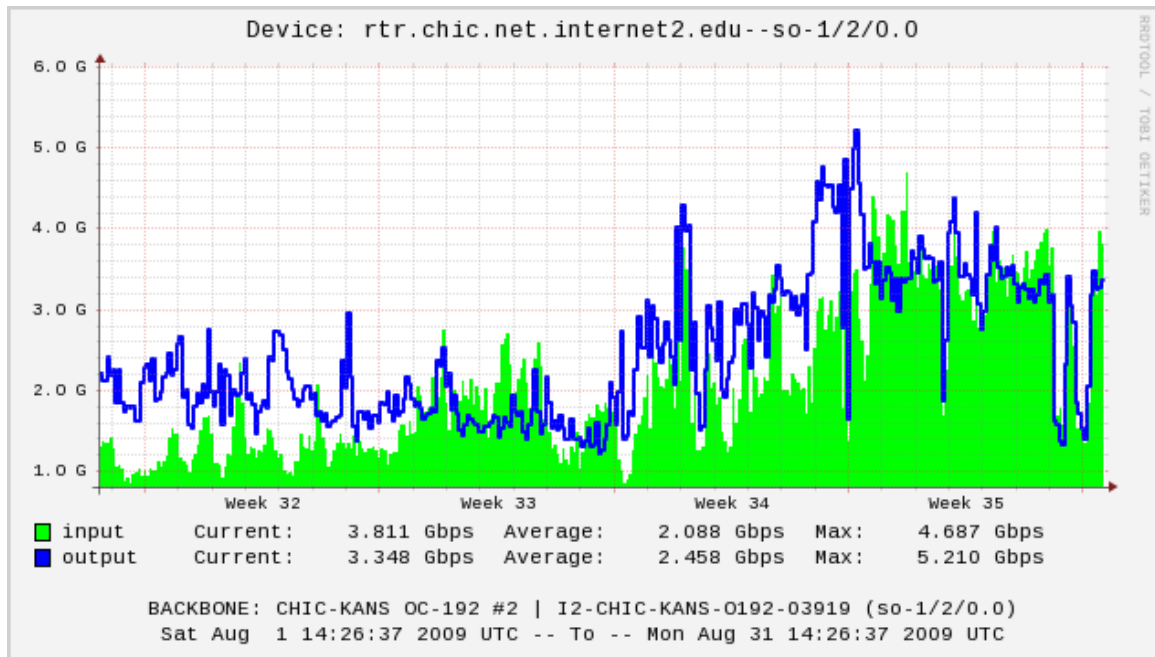
² Graphs of connections to Internet2 Connectors can be accessed at the following URL: <http://dc-snmpp.wcc.grnoc.iu.edu/i2net/show-view-graphs.cgi?viewname=Internet2-Connectors-Aggregate&cmd=graphweek>

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


MANLAN	10Gbps	615Mbps	1.6Gbps	NEWY	40Gbps
Pacific Wave	10Gbps	824Gbps	798Mbps	SEAT	20Gbps
NGIX-East	10Gbps	2.41Gbps	1.85Gbps	WASH	50Gbps
NGIX-West	1Gbps	1Mbps	500Kbps	SALT	40Gbps

Aggregate Internet2 backbone load


As the school term begins, portions of the Internet2 backbone are showing a higher utilization level than we've seen in prior years during this time frame. Of the R&E-only and blended backbone interconnects, Internet2 is seeing *daily* 95th-percentile metrics in the 5-6Gbps range on the Los Angeles-Houston-Kansas City-Chicago backbone path. The monthly graphs below show the increase in traffic during the month of August:



Chicago to Kansas City as measured in Chicago




Houston to Kansas City as measured in Houston

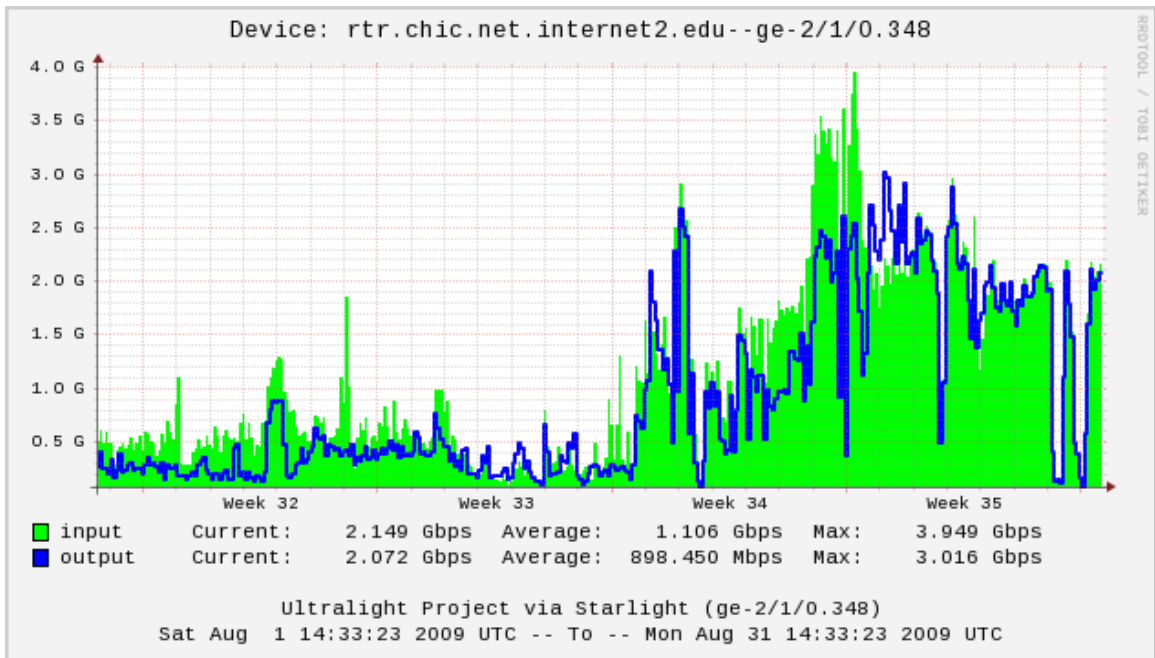


Los Angeles to Houston as measured in Los Angeles

While a portion of this traffic is attributable to CPS traffic (see discussion below), the vast majority is being driven by Internet2's connectivity to the Ultralight project via the Pacific Wave and Starlight R&E exchange points. The graph below shows the VLAN to Ultralight on the Los Angeles routers over the same timeperiod displayed in the graphs above. There is an inverse matching graph in Chicago that correlates to many of the backbone increases.




Ultralight VLAN to Pacific Wave in Los Angeles



Ultralight VLAN to Starlight in Chicago

It's important to note that the traffic samples in these monthly graphs are averaged across 6-hour timeframes. The graph below is zoomed into a 20 minute window in Chicago to provide a more granular view of the traffic bursts on a sub 1 minute average.



Zoomed-In Ultralight VLAN to Starlight in Chicago

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

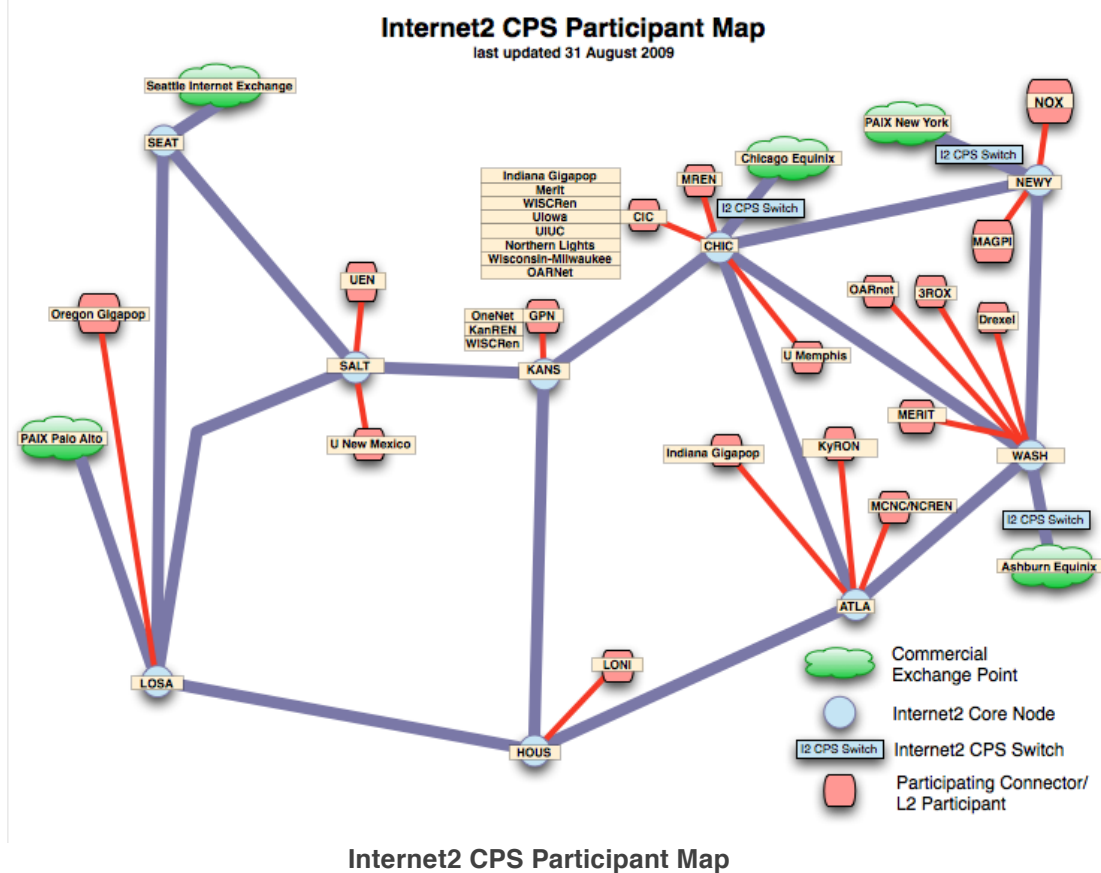
Internet2 maintains settlement free interconnects (SFI) with commodity peers at five major cities on its national footprint:

- 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.⁴



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

⁴ For a current version of this map, visit: <http://noc.net.internet2.edu/i2network/commercial-peering-service/maps--documentation/internet2-cps-network-topology.html>

highest measured 10 second average value measured between August 24th, 2009 and August 31st, 2009. This is likely representative of an increase with students returning, but many organizations had not yet begun classes during this period.

CPS Participant	To Connector	From Connector	Router	Node BW
Utah Education Network	450Mbps	45Mbps	SALT	30Gbps
U of New Mexico	420Mbps	125Mbps	SALT	30Gbps
ONEnet	900Mbps	550Mbps	KANS	30Gbps
KanREN	700Mbps	100Mbps	KANS	30Gbps
GPN	3Gbps	600Mbps	KANS	30Gbps
LONI	600Mbps	350Mbps	HOUS	30Gbps
MCNC	3.2Gbps	100Mbps	ATLA	30Gbps
Indiana Gigapop ⁵	150Mbps	1.1Gbps	ATLA	30Gbps
KyRON	450Mbps	70Mbps	ATLA	30Gbps

The data above suggests that Kansas City is the largest potential consumer of CPS backbone traffic, with most of the traffic likely coming from 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.⁶ The peak values are representative of the highest measured 10 second average value measured between August 24th, 2009 and August 31st, 2009. 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	0Kbps	0Kbps	CHIC	40Gbps
MREN	695Mbps	434Mbps	CHIC	40Gbps
WiscREN via CIC	0Kbps	0Kbps	CHIC	40Gbps
U Minnesota via CIC	1.22Gbps	173Mbps	CHIC	40Gbps
Wisc-Milwaukee via CIC	713Mbps	428Mbps	CHIC	40Gbps
UIUC via CIC	614Mbps	254Mbps	CHIC	40Gbps
U of Iowa via CIC	577Mbps	154Mbps	CHIC	40Gbps
Indiana Gigapop via CIC	1.82Gbps	599Mbps	CHIC	40Gbps
Merit via CIC	267Mbps	209Mbps	CHIC	40Gbps
UIC via CIC	0Kbps	0Kbps	CHIC	40Gbps

⁵ The Indiana Gigapop sources most of its CPS traffic in Atlanta, but receives most of its CPS traffic via its CIC interconnect in Chicago.

⁶ 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>

Oarnet via CIC	0Kbps	0Kbps	CHIC	40Gbps
U of Memphis	141Mbps	59Mbps	CHIC	40Gbps
Oregon Gigapop	394Mbps	511Mbps	LOSA	20Gbps
NOX	3.54Gbps	1.25Gbps	NEWY	20Gbps
MAGPI	375Kbps	140Kbps	NEWY	20Gbps
Merit	1.45Gbps	235Mbps	WASH	30Gbps
OARnet via Merit	3.19Gbps	744Mbps	WASH	30Gbps
Drexel University	12Kbps	547Kbps	WASH	30Gbps
3ROX	581Mbps	112Mbps	WASH	30Gbps

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.⁷ A scan of the afternoon hourly averages on Monday, August 31st 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	Kansas City	1.8 Gbps
Chicago	Washington DC	1.1 Gbps
Kansas City	Chicago	1.0 Gbps
Washington DC	Atlanta	1.8 Gbps
Washington DC	Chicago	1.8 Gbps
Washington DC	Kansas City	1.0 Gbps
Washington DC	New York	1.0 Gbps

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)

⁷ <http://dc-snmp.grnoc.iu.edu/i2-mpls/rrd.cgi?time=hour>

- Peers may have opted to offload traffic in a location that's further away from the CPS connectors

The traffic sourced from Kansas City to Chicago is not surprising given the 1.2Gbps peak traffic from the CPS Connectors on the Kansas City router.

The introduction of the Ashburn exchange point in late July 2009 has alleviated much of the pressure off the New York and Chicago commercial 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 August 24th, 2009 through August 31st, 2009- a period in which most universities were just beginning to welcome students back to campus.⁸


⁸ 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.

The interconnect to the Ashburn exchange point has proven to be phenomenally successful in reducing CPS backbone load, although it continues to be a focus as the exchange point interconnect nears saturation. Internet2 is working with some of the largest commercial peers in Washington DC to offload traffic to a second 10GigE to be used for PNIs. This will reduce traffic volume on the commercial exchange point interconnect while providing increased bandwidth to large commercial peers.