LHCONE Flow Analysis

Mike O’Connor, Network Engineer

ESnet Engineering Group

LHCOPN and LHCONE Joint Meeting

Paris (FR)

June 17, 2012

Netflow

NetFlow services provide network administrators with access to information concerning IP flows within their data networks

For example:

Input/output SNMP interface index, timestamps for flow start and finish time number of bytes and packets observed in the flow

Layer 3 headers:

- Source & destination IP addresses
- Source and destination port numbers for TCP,UDP, SCTP
- ICMP Type and Code
- IP protocol
- Type of Service (ToS) value

For TCP flows, the union of all TCP flags observed over the life of the flow

Layer 3 Routing information:

- IP address of the immediate next-hop along the route to the destination
- Source & destination IP masks (prefix lengths in the CIDR notation)

NetFlow version 9 can include all of these fields and can optionally include additional information such as Multiprotocol Label Switching (MPLS) labels and IPv6 addresses and ports,

Netflow V9

The basic output of NetFlow is a flow record. Several different formats for flow records have evolved as NetFlow has matured. The most recent evolution of the NetFlow flow-record format is known as Version 9. The distinguishing feature of the NetFlow Version 9 format is that it is template based. Templates provide an extensible design to the record format, a feature that should allow future enhancements to NetFlow services without requiring concurrent changes to the basic flow-record format.

IPFIX is an IETF standards track version based on Netflow V9, RFC 5101, RFC5102

Netflow at ESnet

ESnet routers export two versions of Netflow records both sampled only in the recommended inbound direction at a rate of 1/1000.

•Juniper MX exports V5, jflow or cflowd

•Alcatel SR-7750 exports V9, cflowd

V5 is most common version, available as of 2009 but restricted to IPv4 flows.

We use the OSU Flow-tools package flow-cat, flow-filter, flow print etc. and the V9 based, nfdump tools that support flows IPv6 and MPLS flows.

Commercial reporting tools typically import Routing tables as well as SNMP MIB information to add a rich context to the raw Netflow export stream.

ESnet uses:

Arbor - Peakflow

Packet Design - Traffic Explorer

ESnet LHCONE
Netflow Monitoring Locations

- BNL, FNAL, SLAC
- MANLAN
- PNWG
- STARLIGHT
- WIX

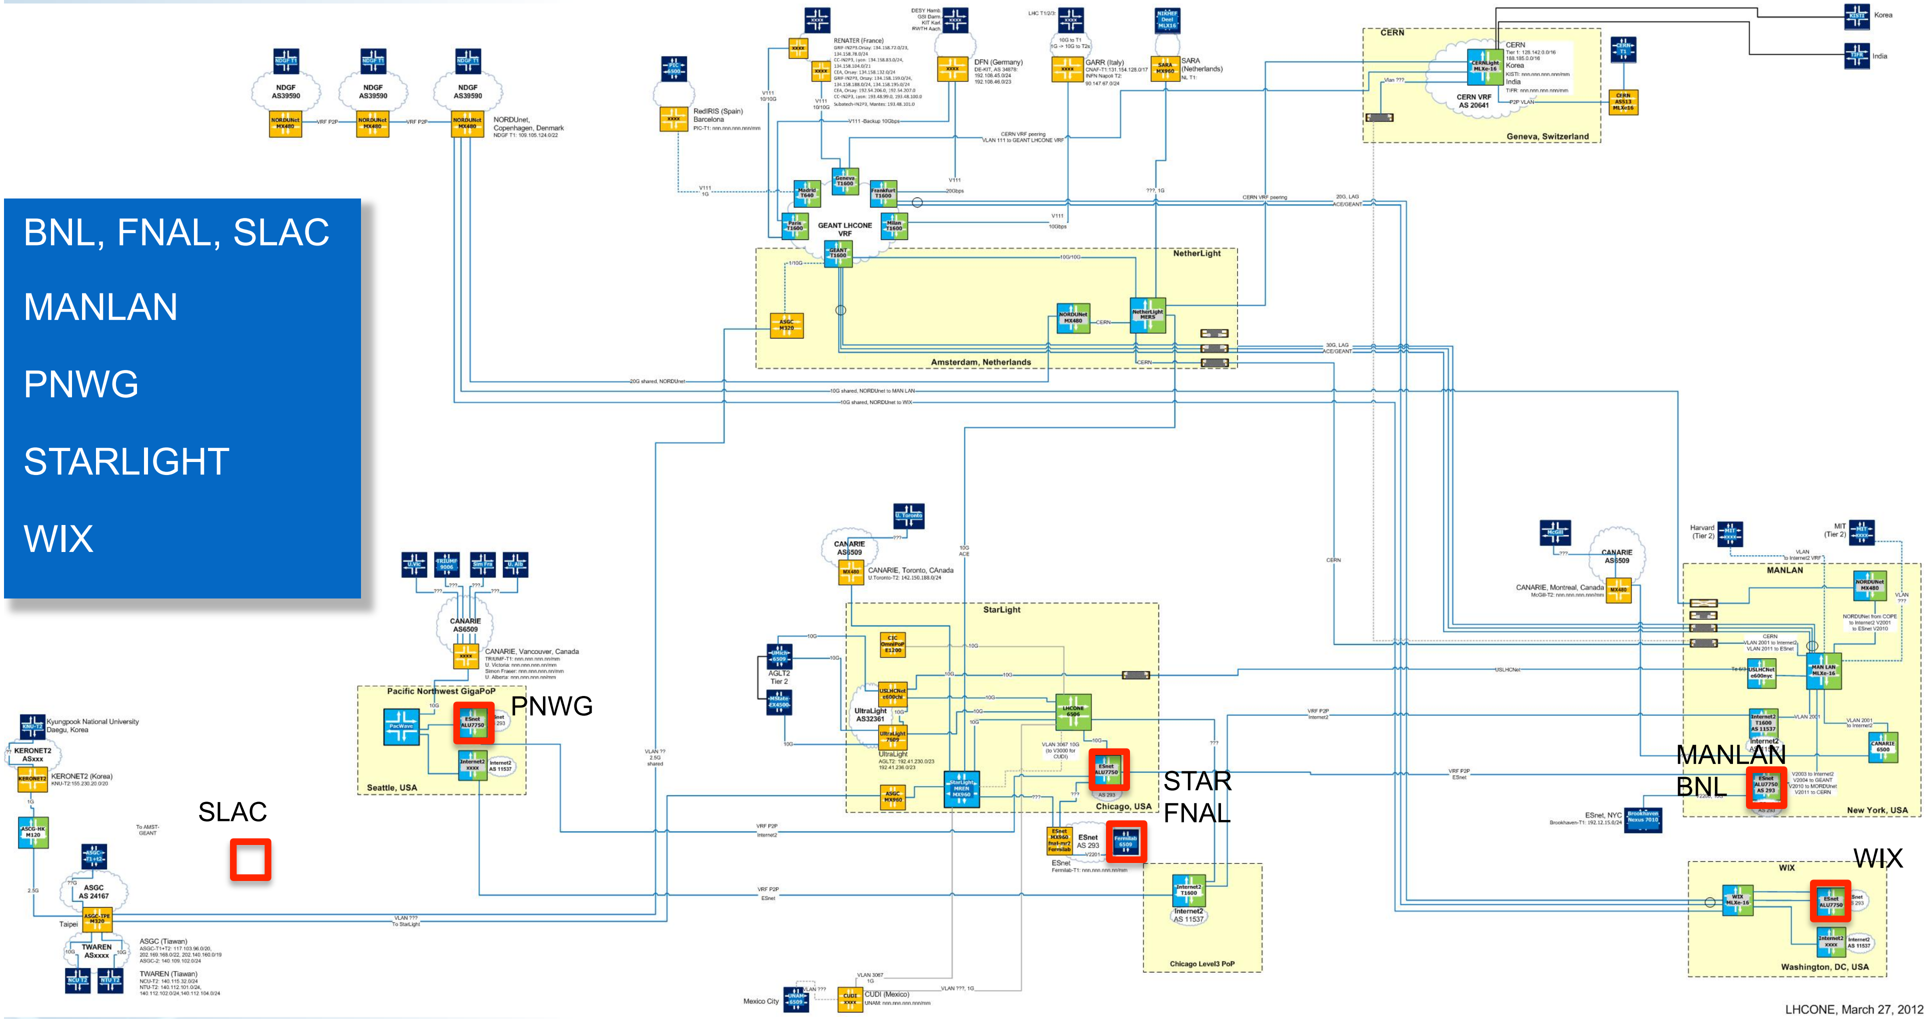

LHCONE, March 27, 2012

LHCONE Routing Table Analysis

- CANET(6509)
  - BCNET(271)
  - UTORONTO(239)
  - UVIC(16462)
  - MCGILL(15318)
  - TRIUMF(36391)
  - UALBERTA(3359)
- CSUNET(2153)
- ULTRALIGHT(32361)
- ESNET(293)
  - FNAL(3152)
  - BNL(43)
  - SLAC(3671)
- I2(11537)
  - UIUC(38)
  - UNL(7896)
  - MIT(3)
  - MERIT(229)
  - UOC(160)
- INDIAN(19782)
  - IUPUI(10680)

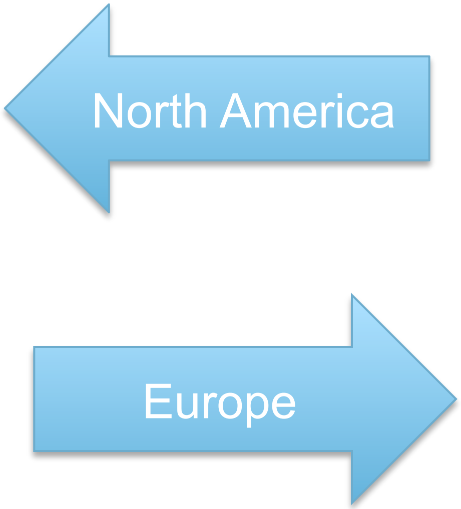

North America

Europe

- CERN-LHC1(20641)
  - CERN-WIGNER(61339)
  - CERN(513)
- DFN(680)
  - KIT(34878)
  - DESY(1754)
- GEANT(20965)
  - ROEDUNET(2614)
  - ASGARR(137)
  - ARNES-NET(2107)
- LHC1-RENATER(2091)
  - IN2P3(789)
  - CEA-SACLAY(777)
- NORDUNET(2603)
  - NDGF(39590)

Combining Netflow and Routing Information

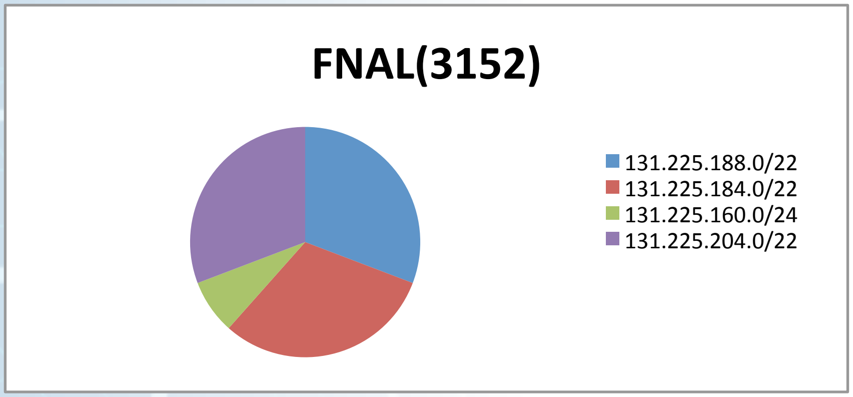

FNAL LHCONE CIDR prefixes

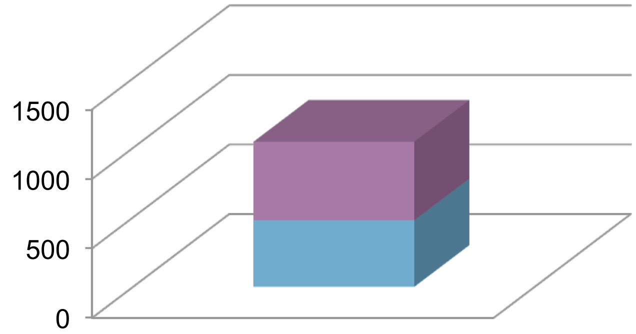

Active hosts in each prefix

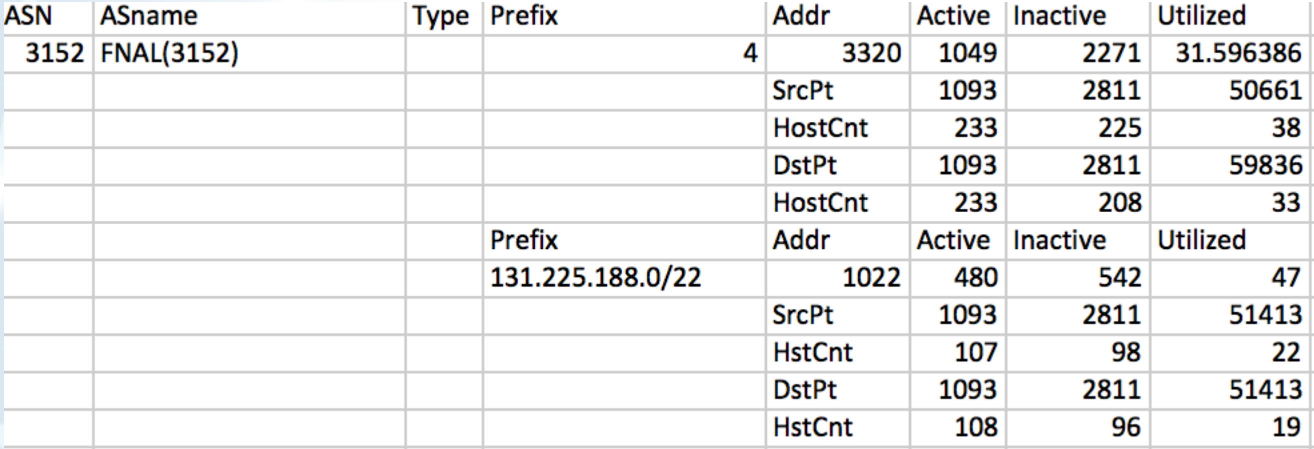

ASN	ASname	Type	Prefix	Addr	Active	Inactive	Utilized
3152	FNAL(3152)		4	3320	1049	2271	31.596386
				SrcPt	1093	2811	50661
				HostCnt	233	225	38
				DstPt	1093	2811	59836
				HostCnt	233	208	33
			Prefix	Addr	Active	Inactive	Utilized
			131.225.188.0/22	1022	480	542	47
				SrcPt	1093	2811	51413
				HstCnt	107	98	22
				DstPt	1093	2811	51413
				HstCnt	108	96	19

FNAL Active hosts and most popular ports

ESnet LHCONE Interfaces/Sub-interfaces

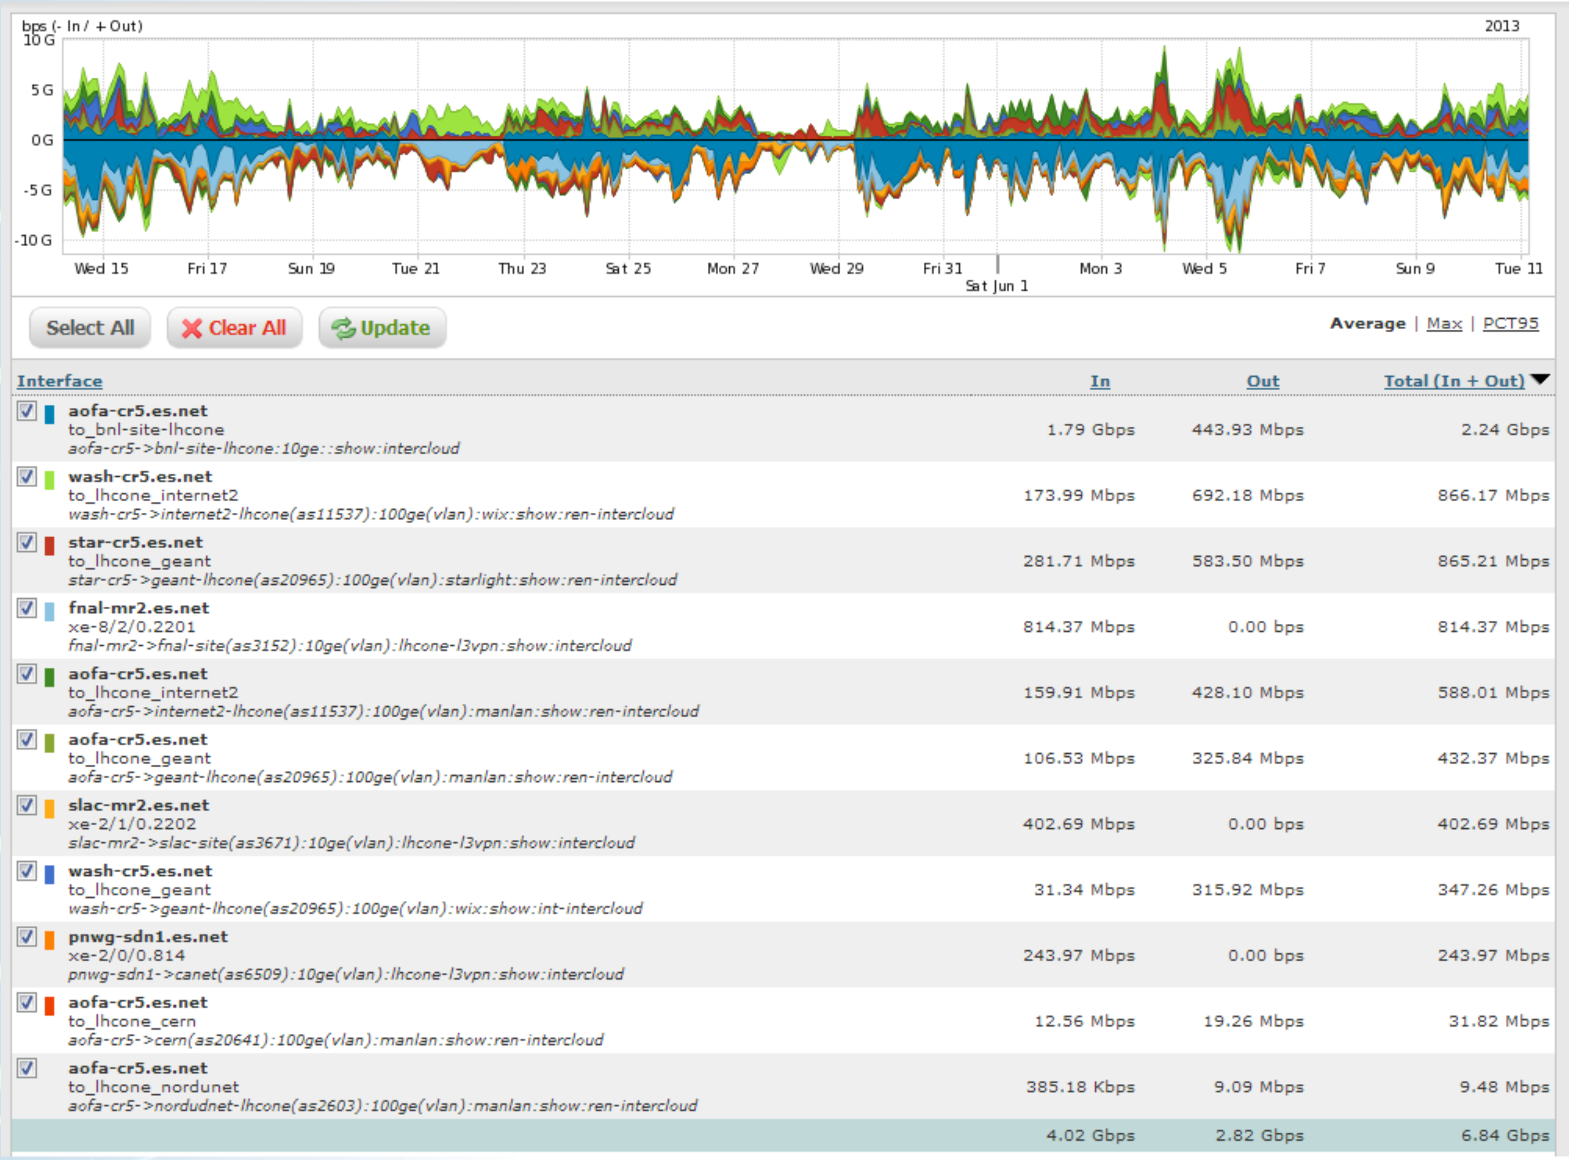

Interface	In	Out	Total (In + Out)
☑ 🟦 **aofa-cr5.es.net** to_bnl-site-lhcone *aofa-cr5->bnl-site-lhcone:10ge::show:intercloud*	1.79 Gbps	443.93 Mbps	2.24 Gbps
☑ 🟩 **wash-cr5.es.net** to_lhcone_internet2 *wash-cr5->internet2-lhcone(as11537):100ge(vlan):wix:show:ren-intercloud*	173.99 Mbps	692.18 Mbps	866.17 Mbps
☑ 🟥 **star-cr5.es.net** to_lhcone_geant *star-cr5->geant-lhcone(as20965):100ge(vlan):starlight:show:ren-intercloud*	281.71 Mbps	583.50 Mbps	865.21 Mbps
☑ 🟦 **fnal-mr2.es.net** xe-8/2/0.2201 *fnal-mr2->fnal-site(as3152):10ge(vlan):lhcone-l3vpn:show:intercloud*	814.37 Mbps	0.00 bps	814.37 Mbps
☑ 🟩 **aofa-cr5.es.net** to_lhcone_internet2 *aofa-cr5->internet2-lhcone(as11537):100ge(vlan):manlan:show:ren-intercloud*	159.91 Mbps	428.10 Mbps	588.01 Mbps
☑ 🟩 **aofa-cr5.es.net** to_lhcone_geant *aofa-cr5->geant-lhcone(as20965):100ge(vlan):manlan:show:ren-intercloud*	106.53 Mbps	325.84 Mbps	432.37 Mbps
☑ 🟧 **slac-mr2.es.net** xe-2/1/0.2202 *slac-mr2->slac-site(as3671):10ge(vlan):lhcone-l3vpn:show:intercloud*	402.69 Mbps	0.00 bps	402.69 Mbps
☑ 🟦 **wash-cr5.es.net** to_lhcone_geant *wash-cr5->geant-lhcone(as20965):100ge(vlan):wix:show:int-intercloud*	31.34 Mbps	315.92 Mbps	347.26 Mbps
☑ 🟧 **pnwg-sdn1.es.net** xe-2/0/0.814 *pnwg-sdn1->canet(as6509):10ge(vlan):lhcone-l3vpn:show:intercloud*	243.97 Mbps	0.00 bps	243.97 Mbps
☑ 🟥 **aofa-cr5.es.net** to_lhcone_cern *aofa-cr5->cern(as20641):100ge(vlan):manlan:show:ren-intercloud*	12.56 Mbps	19.26 Mbps	31.82 Mbps
☑ **aofa-cr5.es.net** to_lhcone_nordunet *aofa-cr5->nordudnet-lhcone(as2603):100ge(vlan):manlan:show:ren-intercloud*	385.18 Kbps	9.09 Mbps	9.48 Mbps
	4.02 Gbps	2.82 Gbps	6.84 Gbps

Filtering by ASN

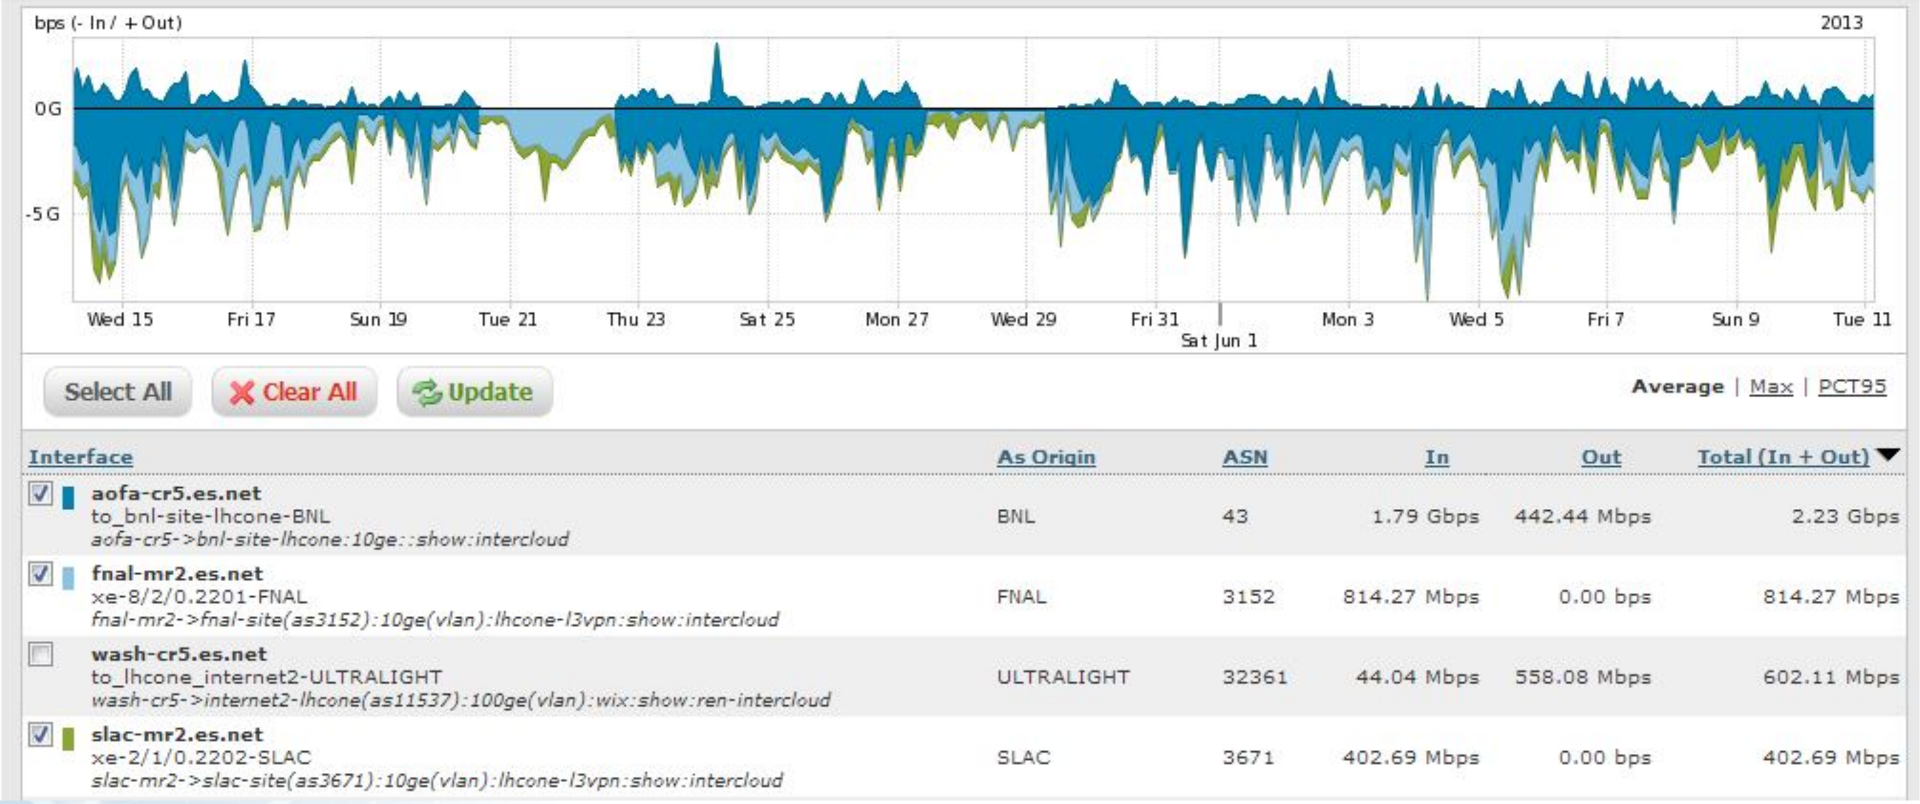

Interface	As Origin	ASN	In	Out	Total (In + Out)
☑ **aofa-cr5.es.net** to_bnl-site-lhcone-BNL *aofa-cr5->bnl-site-lhcone:10ge::show:intercloud*	BNL	43	1.79 Gbps	442.44 Mbps	2.23 Gbps
☑ **fnal-mr2.es.net** xe-8/2/0.2201-FNAL *fnal-mr2->fnal-site(as3152):10ge(vlan):lhcone-l3vpn:show:intercloud*	FNAL	3152	814.27 Mbps	0.00 bps	814.27 Mbps
□ **wash-cr5.es.net** to_lhcone_internet2-ULTRALIGHT *wash-cr5->internet2-lhcone(as11537):100ge(vlan):wix:show:ren-intercloud*	ULTRALIGHT	32361	44.04 Mbps	558.08 Mbps	602.11 Mbps
☑ **slac-mr2.es.net** xe-2/1/0.2202-SLAC *slac-mr2->slac-site(as3671):10ge(vlan):lhcone-l3vpn:show:intercloud*	SLAC	3671	402.69 Mbps	0.00 bps	402.69 Mbps

ESnet site data extracted using their ASN exported in netflow records.

Load Balancing Data, DESY via GEANT

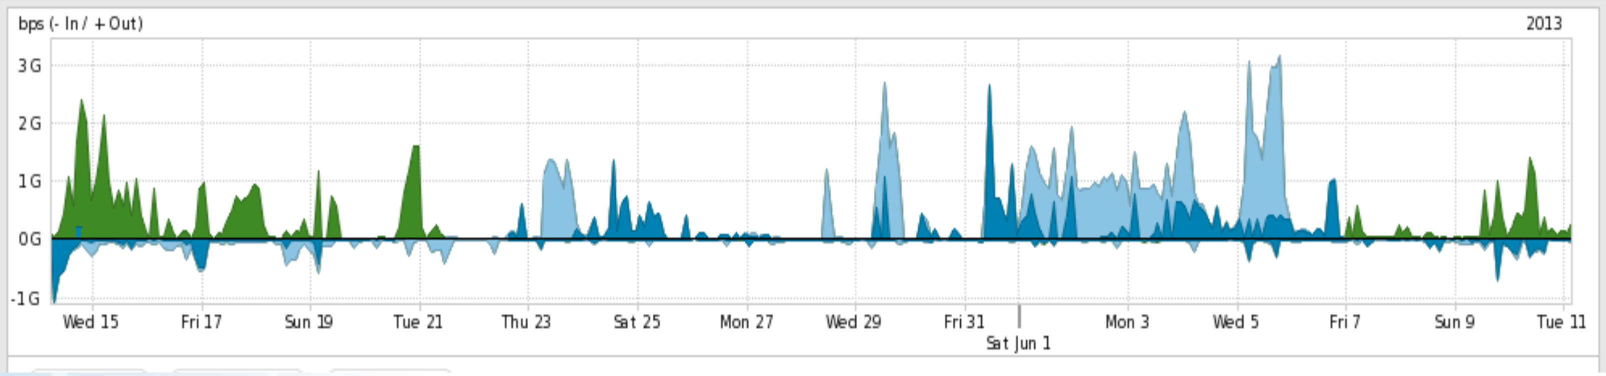

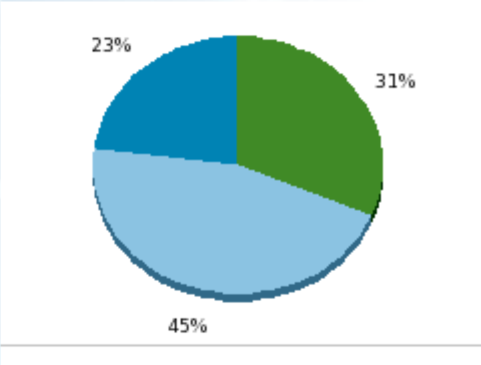

Toward GEANT

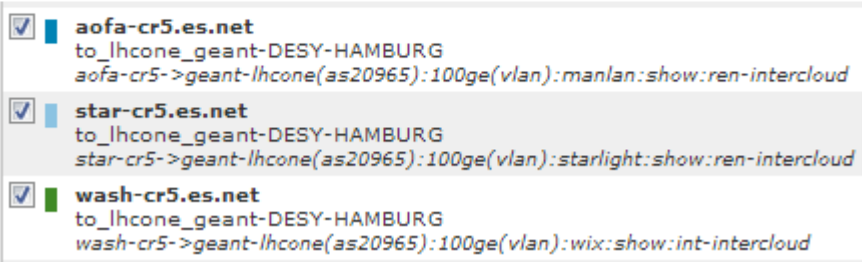

☑ 🟦 **aofa-cr5.es.net**
to_lhcone_geant-DESY-HAMBURG
aofa-cr5->geant-lhcone(as20965):100ge(vlan):manlan:show:ren-intercloud

☑ 🟦 **star-cr5.es.net**
to_lhcone_geant-DESY-HAMBURG
star-cr5->geant-lhcone(as20965):100ge(vlan):starlight:show:ren-intercloud

☑ 🟩 **wash-cr5.es.net**
to_lhcone_geant-DESY-HAMBURG
wash-cr5->geant-lhcone(as20965):100ge(vlan):wix:show:int-intercloud

Observation from June 11, GEANT is sending ESnet a MED of 1 at AOFA and STAR. Traffic is following the MED of 0 received at WASH. This graph implies that GEANT was shifting load between their three LHCONE peerings with Esnet over the last month.

GARR ASN 137

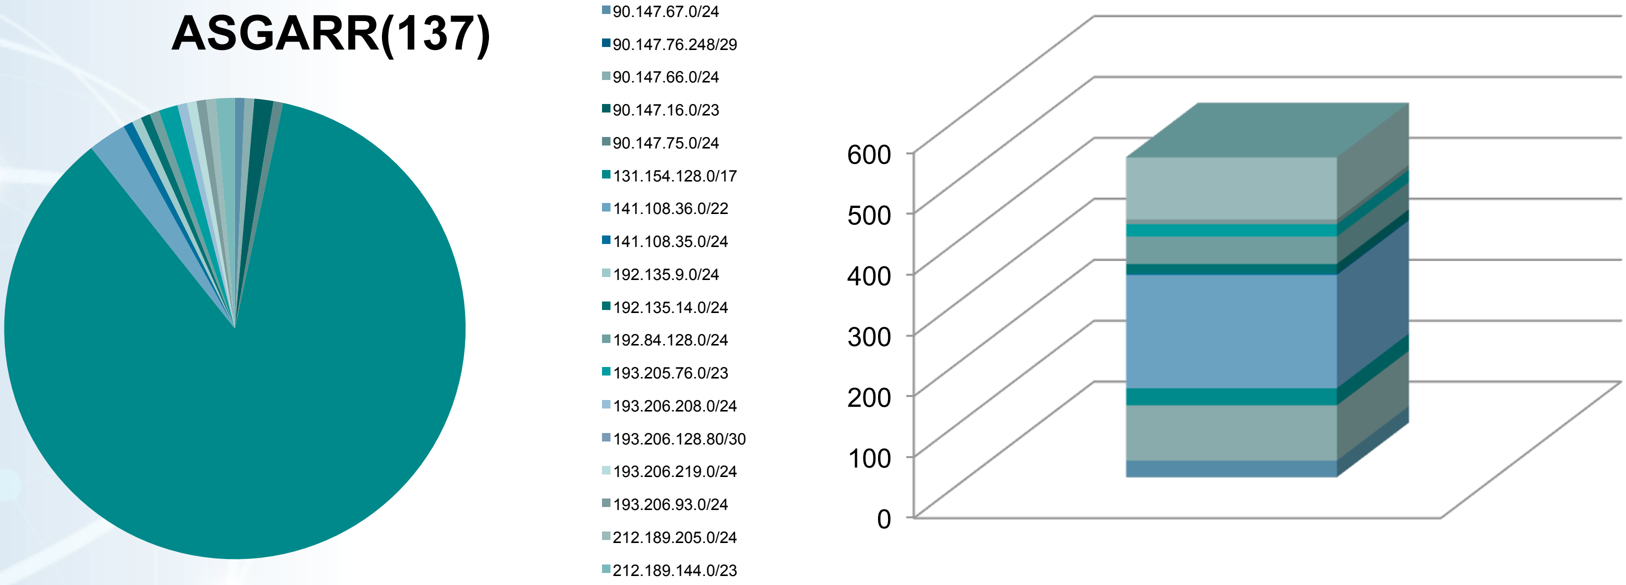

The LHCONE routes for the GARR NREN are sourced from AS 137.
ASN filtering would always group all of the GARR sites as one entity.

Questions?

Michael O’Connor

ESnet Network Engineer

moc@es.net

631 344-7410