ESnet

ENERGY SCIENCES NETWORK

LHCONE Operations Update

Michael O'Connor moc@es.net

ESnet

Network Engineering

LHCOPN-LHCONE meeting

Umeå university

Umeå (SE)

June. 4, 2019

U.S. DEPARTMENT OF ENERGY

Office of Science

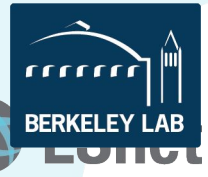

BERKELEY LAB

ESnet Traffic Volumes

ESnet LHCONE traffic has increased 81% in the past year.
The ESnet LHCONE service carries 54% of ESnet volume.

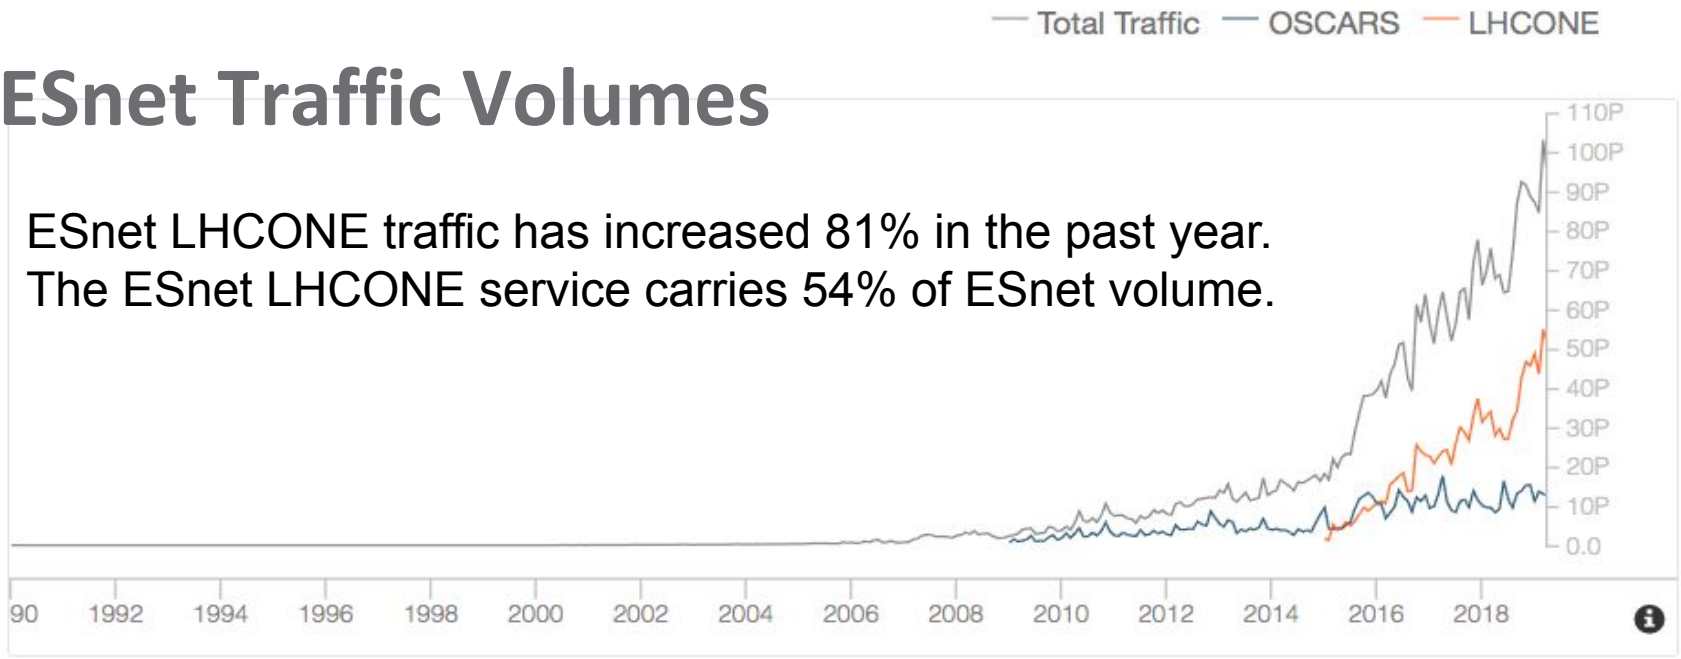

April 2019

	Bytes	Percent of Total	One Month Change	One Year Change
OSCARS	12.42PB	13.3%	-4.44%	+48.5%
LHCONE	50.53PB	54.3%	-7.94%	+81.3%
Normal traffic	30.19PB	32.4%	-14.2%	-4.34%
Total	93.14PB		-9.63%	+37.4%

LHCONE continues to drive annual traffic volume increases.

ESnet

LHCONE CERN/ESnet

April 1, 2019 through May 22, 2019

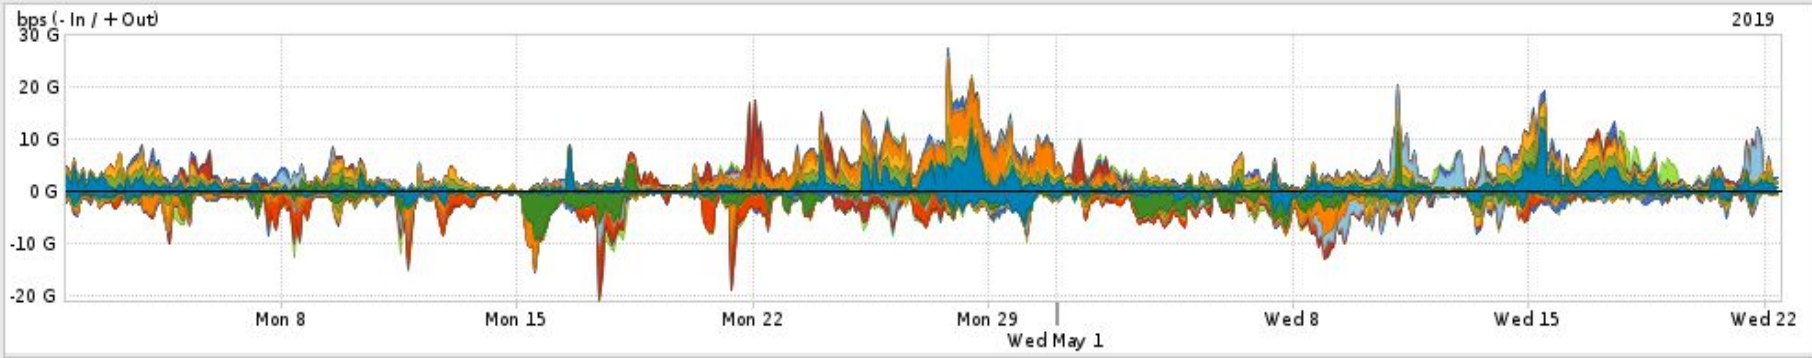

Select All | Clear All | Update

Average | Max | PCT95

Name	In	Out	Total (In + Out)	% Total
UOC(160)	537.58 Mbps	1.22 Gbps	1.76 Gbps	19.92%
MIT(3)	786.32 Mbps	172.20 Mbps	958.53 Mbps	10.86%
UTA(18515)	176.73 Mbps	741.49 Mbps	918.22 Mbps	10.40%
UTARLINGTON(18515)	176.73 Mbps	741.49 Mbps	918.22 Mbps	10.40%
CALTECH(32361)	326.30 Mbps	581.96 Mbps	908.25 Mbps	10.29%
WISC-MADISON(59)	261.08 Mbps	393.03 Mbps	654.11 Mbps	7.41%
VANDERBILT(7212)	554.25 Mbps	80.63 Mbps	634.87 Mbps	7.19%
PURDUE(17)	243.51 Mbps	327.80 Mbps	571.31 Mbps	6.47%
NET2(111)	161.34 Mbps	403.28 Mbps	564.63 Mbps	6.40%
UFL(6356)	165.40 Mbps	179.04 Mbps	344.44 Mbps	3.90%
unl	166.57 Mbps	72.45 Mbps	239.03 Mbps	2.71%

CERN LHCONE traffic crossing ESnet

ESnet

LHCONE US CMS Tier1 Fermilab

April 1, 2019 through May 22, 2019

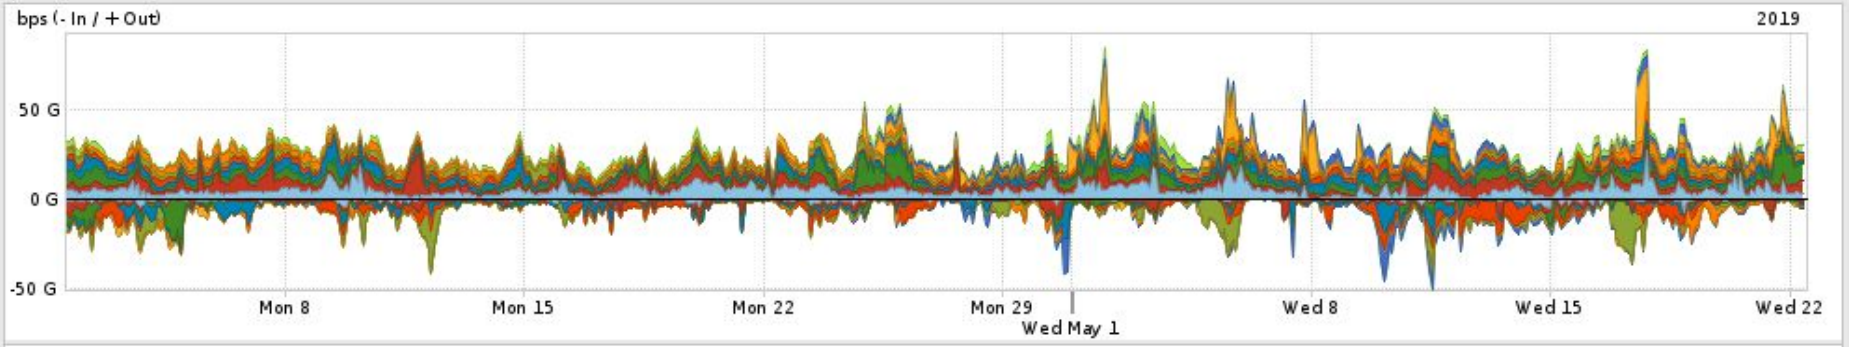

Select All | Clear All | Update

Average | Max | PCT95

Name	In	Out	Total (In + Out)	% Total
WISC-MADISON(59)	526.81 Mbps	4.56 Gbps	5.09 Gbps	11.14%
PURDUE(17)	1.09 Gbps	3.84 Gbps	4.93 Gbps	10.78%
MIT(3)	950.23 Mbps	3.56 Gbps	4.51 Gbps	9.87%
unl	1.51 Gbps	2.96 Gbps	4.46 Gbps	9.77%
VANDERBILT(7212)	2.04 Gbps	2.34 Gbps	4.37 Gbps	9.57%
KIT-GridKa(58069)	1.68 Gbps	1.79 Gbps	3.47 Gbps	7.60%
CALTECH(32361)	597.74 Mbps	2.32 Gbps	2.92 Gbps	6.39%
JANET(786)	422.15 Mbps	1.93 Gbps	2.35 Gbps	5.15%
UNL(7896)	723.70 Mbps	1.51 Gbps	2.23 Gbps	4.88%
UFL(6356)	1.77 Mbps	1.89 Gbps	1.89 Gbps	4.13%
ASGARR(137)	378.14 Mbps	1.07 Gbps	1.45 Gbps	3.18%
DESY(1754)	1.02 Gbps	406.85 Mbps	1.43 Gbps	3.13%

FNAL LHCONE traffic crossing ESnet

ESnet

LHCONE US CMS Tier2 Wisconsin Madison

April 1, 2019 through May 22, 2019

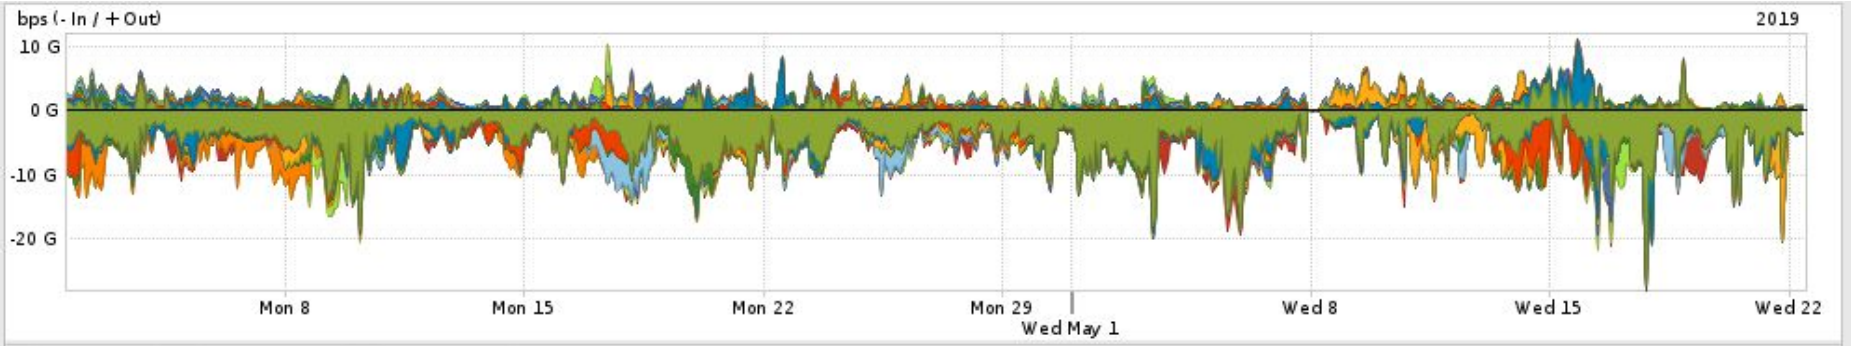

Select All | Clear All | Update

Average | Max | PCT95

Name	In	Out	Total (In + Out)	% Total
FNAL(3152)	4.56 Gbps	526.92 Mbps	5.09 Gbps	42.03%
ASGARR(137)	412.79 Mbps	371.81 Mbps	784.60 Mbps	6.48%
VANDERBILT(7212)	469.30 Mbps	249.41 Mbps	718.71 Mbps	5.94%
CERN(513)	393.13 Mbps	261.31 Mbps	654.44 Mbps	5.41%
MIT(3)	449.02 Mbps	205.11 Mbps	654.13 Mbps	5.40%
unl	355.51 Mbps	157.03 Mbps	512.54 Mbps	4.23%
UCSD(26397)	148.13 Mbps	241.81 Mbps	389.94 Mbps	3.22%
IN2P3(789)	363.26 Mbps	15.86 Mbps	379.11 Mbps	3.13%
UFL(6356)	208.01 Mbps	163.52 Mbps	371.53 Mbps	3.07%
PIC(43115)	327.83 Mbps	859.06 Kbps	328.69 Mbps	2.71%
JINR(2875)	318.64 Mbps	8.49 Mbps	327.13 Mbps	2.70%

ESnet

LHCONE US ATLAS Tier1 and Belle-II at BNL

April 1, 2019 through May 22, 2019

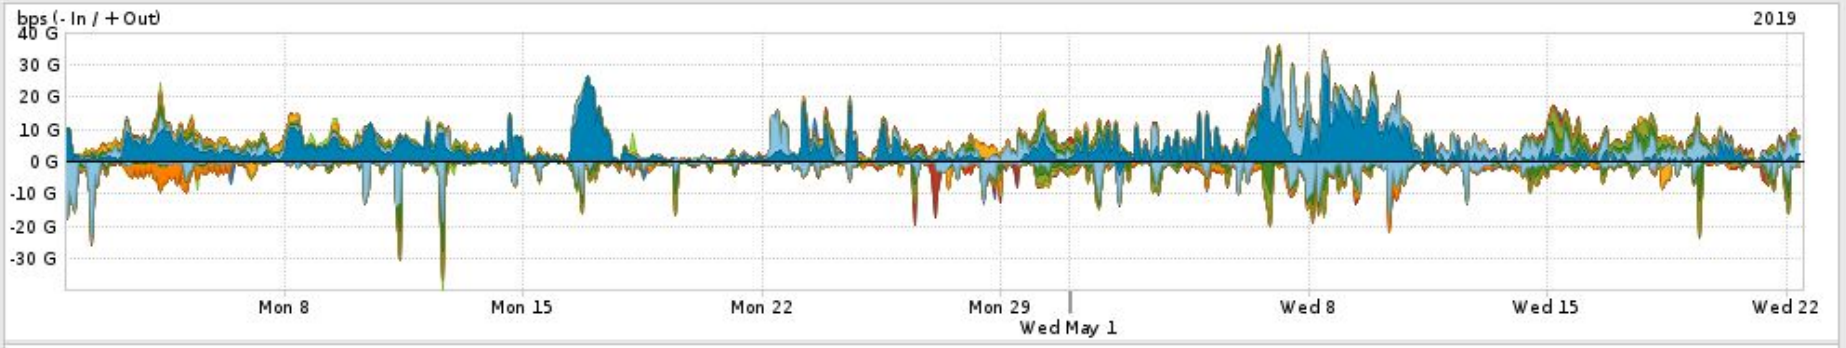

Select All | Clear All | Update

Average | Max | PCT95

Name	In	Out	Total (In + Out)	% Total
AGLT2(229)	23.70 Mbps	3.82 Gbps	3.84 Gbps	36.24%
UOC(160)	1.33 Gbps	1.74 Gbps	3.07 Gbps	28.92%
UTA(18515)	459.05 Mbps	531.73 Mbps	990.78 Mbps	9.34%
UTARLINGTON(18515)	459.05 Mbps	531.73 Mbps	990.78 Mbps	9.34%
NET2(111)	151.30 Mbps	436.43 Mbps	587.73 Mbps	5.54%
HEPNET-J(2505)	514.67 Mbps	16.89 Mbps	531.56 Mbps	5.01%
SLAC(3671)	53.47 Mbps	146.31 Mbps	199.78 Mbps	1.88%
SFU(11105)	131.33 Mbps	63.79 Mbps	195.12 Mbps	1.84%
DESY(1754)	78.89 Mbps	105.68 Mbps	184.56 Mbps	1.74%
UTNET(2501)	67.57 Mbps	93.72 Mbps	161.28 Mbps	1.52%
RoEduNet(2614)	26.95 Mbps	123.42 Mbps	150.37 Mbps	1.42%

BNL LHCONE traffic crossing ESnet

ESnet

LHCONE US ATLAS Tier2 U of Chicago

April 1, 2019 through May 22, 2019

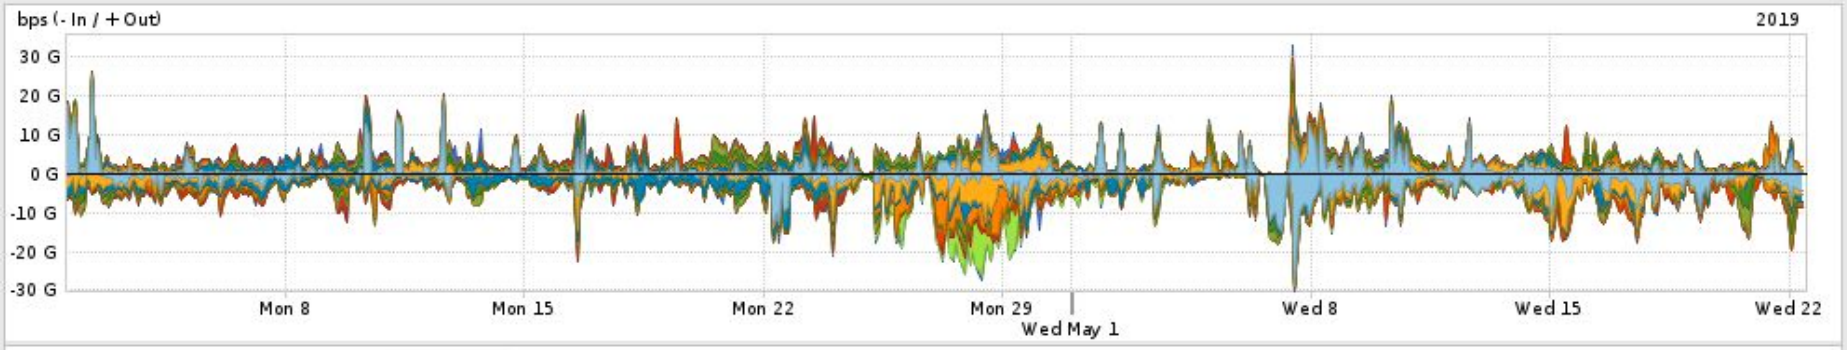

Select All | Clear All | Update

Average | Max | PCT95

Name	In	Out	Total (In + Out)	% Total
BNL(43)	1.74 Gbps	1.33 Gbps	3.07 Gbps	20.37%
CERN(513)	1.22 Gbps	537.61 Mbps	1.76 Gbps	11.67%
AGLT2(229)	1.12 Gbps	623.33 Mbps	1.74 Gbps	11.55%
ASGARR(137)	616.14 Mbps	389.57 Mbps	1.01 Gbps	6.67%
UTA(18515)	397.74 Mbps	460.90 Mbps	858.64 Mbps	5.70%
UTARLINGTON(18515)	397.74 Mbps	460.90 Mbps	858.64 Mbps	5.70%
KIT-GridKa(58069)	412.79 Mbps	224.46 Mbps	637.25 Mbps	4.23%
IN2P3(789)	264.11 Mbps	315.93 Mbps	580.04 Mbps	3.85%
TRIUMF(36391)	411.42 Mbps	49.64 Mbps	461.06 Mbps	3.06%
DESY(1754)	134.97 Mbps	223.61 Mbps	358.58 Mbps	2.38%
KIAE(59624)	255.71 Mbps	91.63 Mbps	347.35 Mbps	2.31%

ESnet

ESnet Core Saturation Event

ATLAS Cache Testing

- April 2, 2019 the 100G core circuit between the two ESnet POPs in New York City saturated in one direction for approximately 2 hours.
- Testing of BNL's caching facility between US ATLAS T1 and UChicago (Midwest T2) saturated the link
- UOC pushed data into the ATLAS cache server at BNL and then drew the data into the Google Cloud from BNL cache using several thousand cloud instances

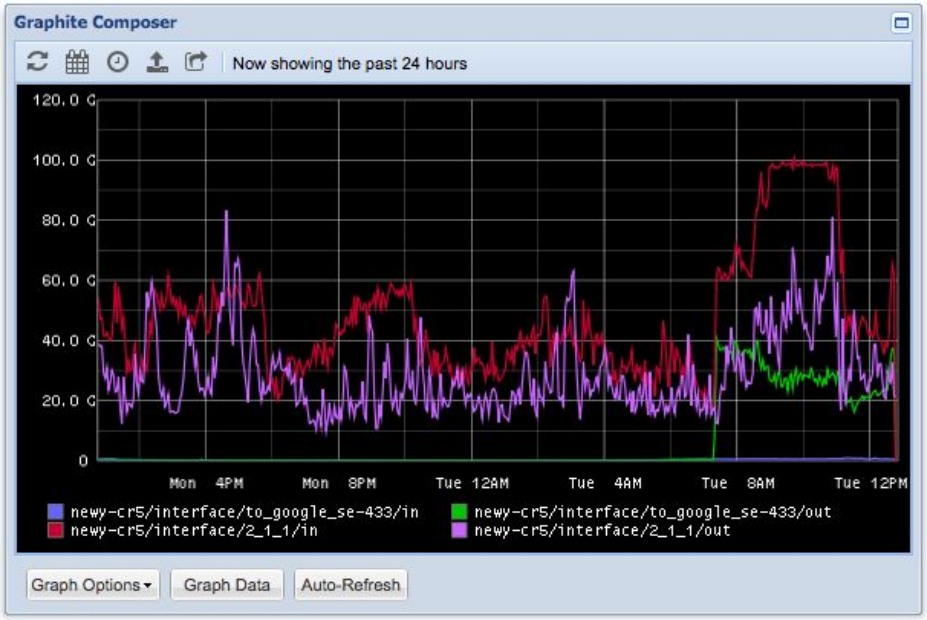

An LHC workflow connects ATLAS Tier centers across both LHCONE and General IP to Google Cloud with unintended consequences.

ESnet

Events Leading to Core Link Saturation

Best Laid Plans Go Awry

Positive steps to engineer the network from a local perspective created an environment that enabled this unforeseen issue.

Traffic engineering steps leading up to this event:

- LHC Tier centers BNL, UOC and others connect to LHCONE
- BNL Network Engineering distributes services across it's two 100G links to ESnet, in order to avoid 100G local loop saturation.
- ESnet adds a new 100G direct Google peering on the east coast at 111 8th Ave (ESnet NEWY POP)

These steps were all individually successful in achieving their goal

- LHCONE did not saturate
- BNL local loops did not saturate
- ESnet/Google peering did not saturate

The combined, General IP and LHCONE load saturated the core circuit

ESnet

BNL Local Loop

- Previous saturation events on their primary 100G uplink caused BNL to distribute load by shifting the primary LHCONE path from AOFA to NEWY
- This step successfully prevents congestion on the BNL links to ESnet

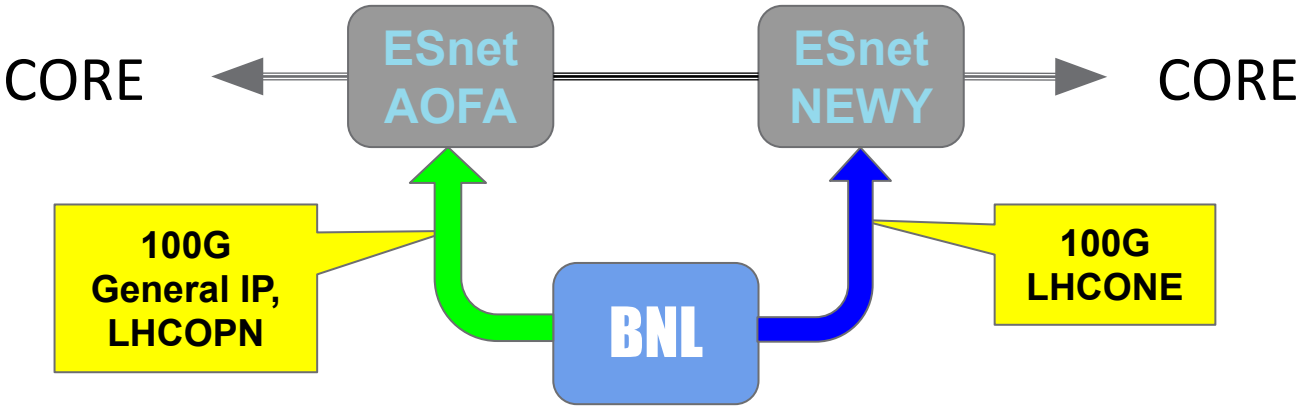

BNL shifts LHCONE primary path to NEWY

ESnet

LHCONE Path UOC to BNL

Loading data into the ATLAS Cache server at BNL on LHCONE

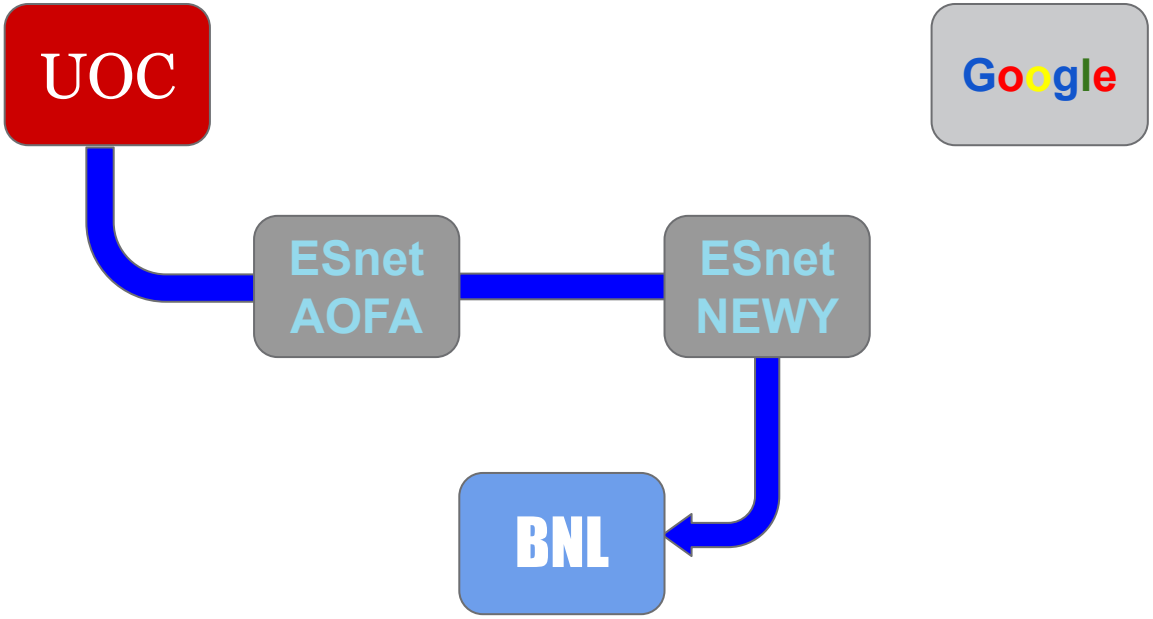

ESnet

General IP Path BNL to Google

Reading data from the ATLAS Cache server at BNL into Google Cloud

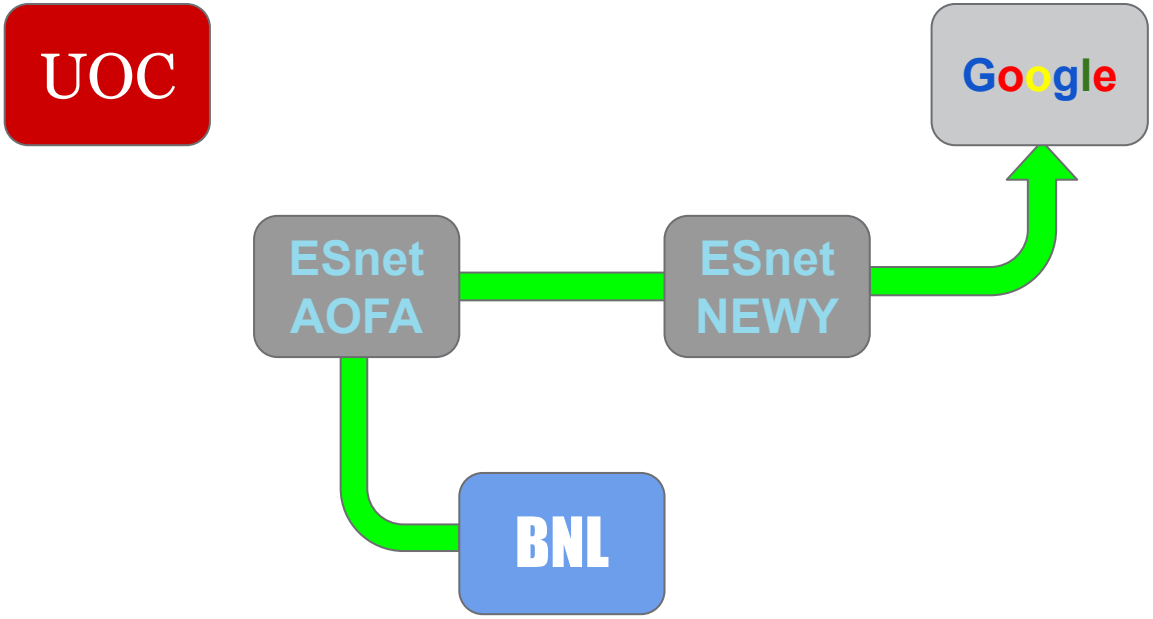

ESnet

General IP Path BNL to Google

- UOC encountered a limit just under 50G
- The link between the NYC POPs saw the traffic twice, once on LHCONE and then on General IP

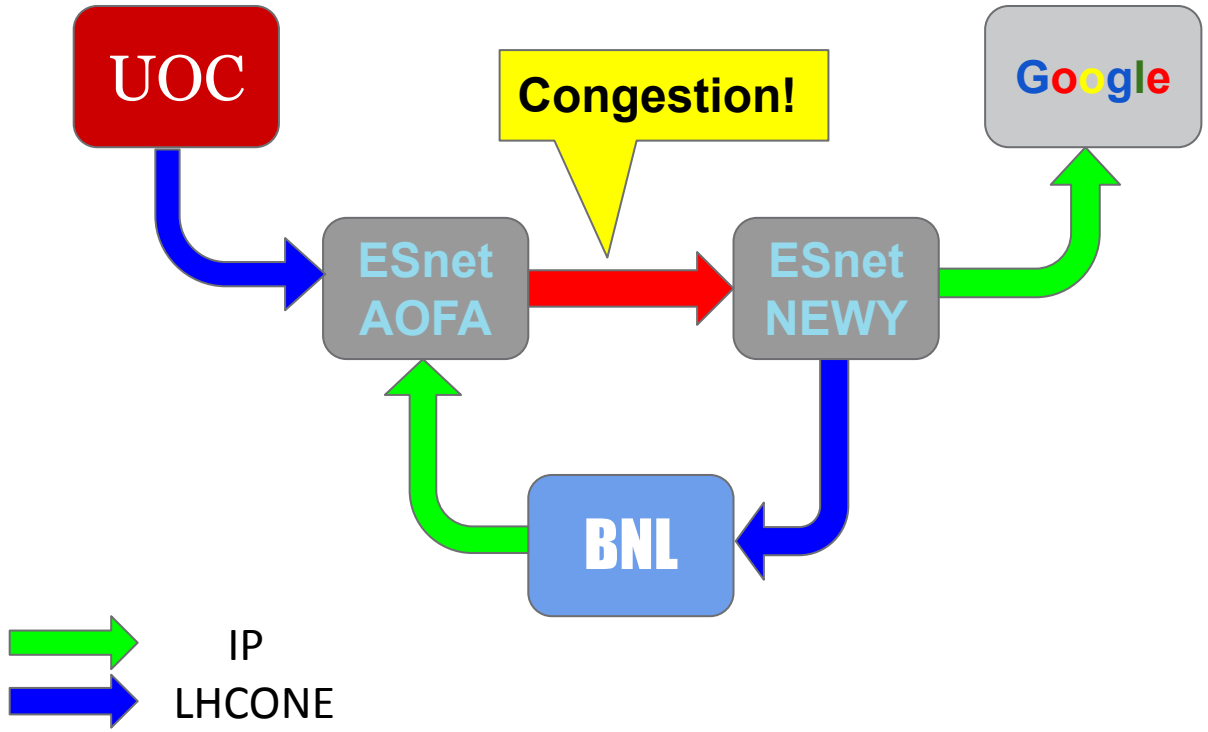

The link between the NYC POPs has not crossed the ESnet monitoring threshold used by the capacity planning process to recommend an upgrade.

ESnet

ESnet

ENERGY SCIENCES NETWORK

Questions?

Michael O'Connor moc@es.net

ESnet

Network Engineering

LHCOPN-LHCONE meeting

Fermi National Laboratory

Batavia, IL

Oct. 30, 2018

U.S. DEPARTMENT OF ENERGY

Office of Science

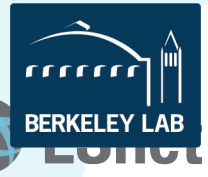

BERKELEY LAB