



ESnet

ENERGY SCIENCES NETWORK

A closer look at ESnet's LHCONE traffic

Aligning numbers with CMS and ATLAS

Richard Cziva, Ph.D.
Software Engineer
ESnet, Berkeley Lab

LHCONE - LHCOPN Meeting
Umea, Sweden
4 June 2018



U.S. DEPARTMENT OF
ENERGY
Office of Science



Motivation

- At ESnet, we analyse network traffic using
 - SNMP interface counters and
 - NetFlow records
- ESnet needs deep understanding of traffic patterns for:
 - capacity planning
 - determine free and available capacity
 - anticipate link upgrades well in advance due to long lead time
 - resiliency analysis
 - simulate the network under failure conditions
 - network optimization
 - topology: identify bottlenecks
 - routing: find inefficiencies, unroutable packets, asymmetric routing

Goal of this analysis: **common understanding of traffic patterns and future needs between ESnet and HEP (CMS/ATLAS) coordinators**

Data at ESnet

- Timeframe analyzed: 2018/01/01 - 2019/04/20 (15 months)
- Network interfaces analyzed:
 - In the US:
 - all edge interfaces with LHCONE interface tags
 - Transatlantic:
 - Boston - Amsterdam 100GE
 - New York (newy) - London 100GE
 - New York (aofa) - London 100GE
 - New York (aofa) - London NEEAR 100GE
 - Washington - CERN 100GE
- Data sources:
 - Netflow data (collected at ESnet for all interfaces)
 - sampling: 1 to 1000 (~ up to 5% error compared to SNMP)
 - SNMP counters, publicly available on my.es.net
 - LHCONE US T2 prefixes (manually edited list)
- Processing software: “ESflow” - ESnet flow analyzer

LHCONE US T2-T2 Traffic Matrix - Heatmap

Month: Sept 2018

From \ To	MIT	SWT2/	MWT2/	MWT	UFL	AGL	MWT2/	AGLT	NET2	Caltec	UCSD	Vander	U Wisc.	Purdue	AGLT2	UNL	SLA	SW
MIT (T2)	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	1.1%	2.0%	1.6%	0.0%	0.0%	1.5%	0.0%	0.0%
SWT2/OU (T2)	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
MWT2/UChicago	0.0%	0.0%	0.0%	0.4%	0.1%	0.0%	0.0%	0.9%	1.2%	0.4%	0.1%	0.0%	0.3%	0.0%	1.0%	1.2%	0.6%	1.3%
MWT2/IU (T2)	0.0%	0.1%	0.0%	0.0%	0.0%	0.0%	0.7%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.0%	0.0%	0.1%
UFL Univ. Florida (T2)	1.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.5%	1.2%	0.6%	0.6%	0.0%	1.1%	0.0%	0.0%
AGLT2	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
MWT2/UIUC (T2)	0.0%	0.1%	0.0%	0.7%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.2%	0.0%	0.1%	0.0%	0.0%	0.1%
AGLT2/UM (T2)	0.0%	0.0%	0.4%	0.0%	0.0%	0.0%	0.0%	0.0%	0.5%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.5%	0.5%
NET2 (T2)	0.0%	0.0%	0.6%	0.0%	0.0%	0.0%	0.0%	0.4%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.4%	0.0%	0.2%	0.6%
Caltech (T2)	0.5%	0.0%	0.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.8%	0.8%	0.5%	0.0%	0.9%	0.0%	0.0%
UCSD	0.4%	0.0%	0.0%	0.0%	0.4%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.6%	0.5%	0.3%	0.0%	0.6%	0.0%	0.0%
Vanderbilt University (T2)	0.3%	0.0%	0.0%	0.0%	0.4%	0.0%	0.0%	0.0%	0.0%	0.3%	0.3%	0.0%	0.4%	0.3%	0.0%	0.4%	0.0%	0.0%
U Wisc. Madison (T2)	1.2%	0.0%	0.0%	0.0%	1.5%	0.0%	0.0%	0.0%	0.0%	1.6%	1.4%	2.6%	0.0%	1.5%	0.0%	1.8%	0.0%	0.0%
Purdue CMS T2	1.8%	0.0%	0.0%	0.0%	2.3%	0.0%	0.0%	0.0%	0.0%	2.0%	1.6%	3.3%	3.2%	0.0%	0.0%	3.4%	0.0%	0.0%
AGLT2/MSU (T2)	0.0%	0.0%	0.4%	0.1%	0.0%	0.0%	0.1%	0.0%	0.5%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.5%	0.5%
UNL	4.0%	0.0%	0.3%	0.1%	6.7%	0.0%	0.1%	0.1%	0.0%	3.4%	3.2%	6.5%	2.4%	3.4%	0.0%	0.0%	0.0%	0.0%
SLAC (T2)	0.0%	0.0%	0.9%	0.4%	0.0%	0.0%	0.4%	0.0%	0.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.2%
SWT2/UTA (T2)	0.0%	0.0%	0.6%	0.1%	0.0%	0.0%	0.1%	0.3%	0.7%	0.0%	0.0%	0.0%	0.0%	0.0%	0.3%	0.0%	0.4%	0.0%

UNL - UFL


UNL - Vanderbilt

Observation: There are particular T2 pairs that communicate more than others in the US.



LHCONE US T2-T2 Traffic Matrix - Dynamic Heatmap

Month: March 2019



LHCONE US T2-T2 Top Talkers

Month: May 2018

From site	Total TB	% of T2-total
MIT (T2)	355.50 TB	6.17%
SWT2/OU (T2)	0.00 TB	0.00%
MWT2/UChicago	431.39 TB	7.48%
MWT2/IU (T2)	53.58 TB	0.93%
UFL Univ. Florida (T2)	297.86 TB	5.17%
AGLT2	0.03 TB	0.00%
MWT2/UIUIC (T2)	66.11 TB	1.15%
AGLT2/UM (T2)	110.69 TB	1.92%
NET2 (T2)	126.28 TB	2.19%
Caltech (T2)	213.88 TB	3.71%
UCSD	163.59 TB	2.84%
Vanderbilt University (T2)	142.34 TB	2.47%
U Wisc. Madison (T2)	670.30 TB	11.63%
Purdue CMS T2	1,016.02 TB	17.63%
AGLT2/MSU (T2)	117.12 TB	2.03%
UNL	1,735.42 TB	30.11%
SLAC (T2)	124.65 TB	2.16%
SWT2/UTA (T2)	139.39 TB	2.42%
Total	5,764.15 TB	100.00%

T2 Top Talkers

LHCONE US XRootD traffic

Month: May 2018

From site	Total TB	% of T2-total
MIT (T2)	315.02 TB	7.71%
SWT2/OU (T2)	0.00 TB	0.00%
MWT2/UChicago	66.93 TB	1.64%
MWT2/IU (T2)	0.03 TB	0.00%
UFL Univ. Florida (T2)	185.38 TB	4.54%
AGLT2	0.00 TB	0.00%
MWT2/UIUIC (T2)	0.12 TB	0.00%
AGLT2/UM (T2)	0.07 TB	0.00%
NET2 (T2)	0.01 TB	0.00%
Caltech (T2)	178.01 TB	4.36%
UCSD	139.35 TB	3.41%
Vanderbilt University (T2)	93.99 TB	2.30%
U Wisc. Madison (T2)	588.38 TB	14.41%
Purdue CMS T2	935.40 TB	22.90%
AGLT2/MSU (T2)	0.03 TB	0.00%
UNL	1,581.12 TB	38.72%
SLAC (T2)	0.06 TB	0.00%
SWT2/UTA (T2)	0.11 TB	0.00%
Total	4,084.00 TB	100.00%

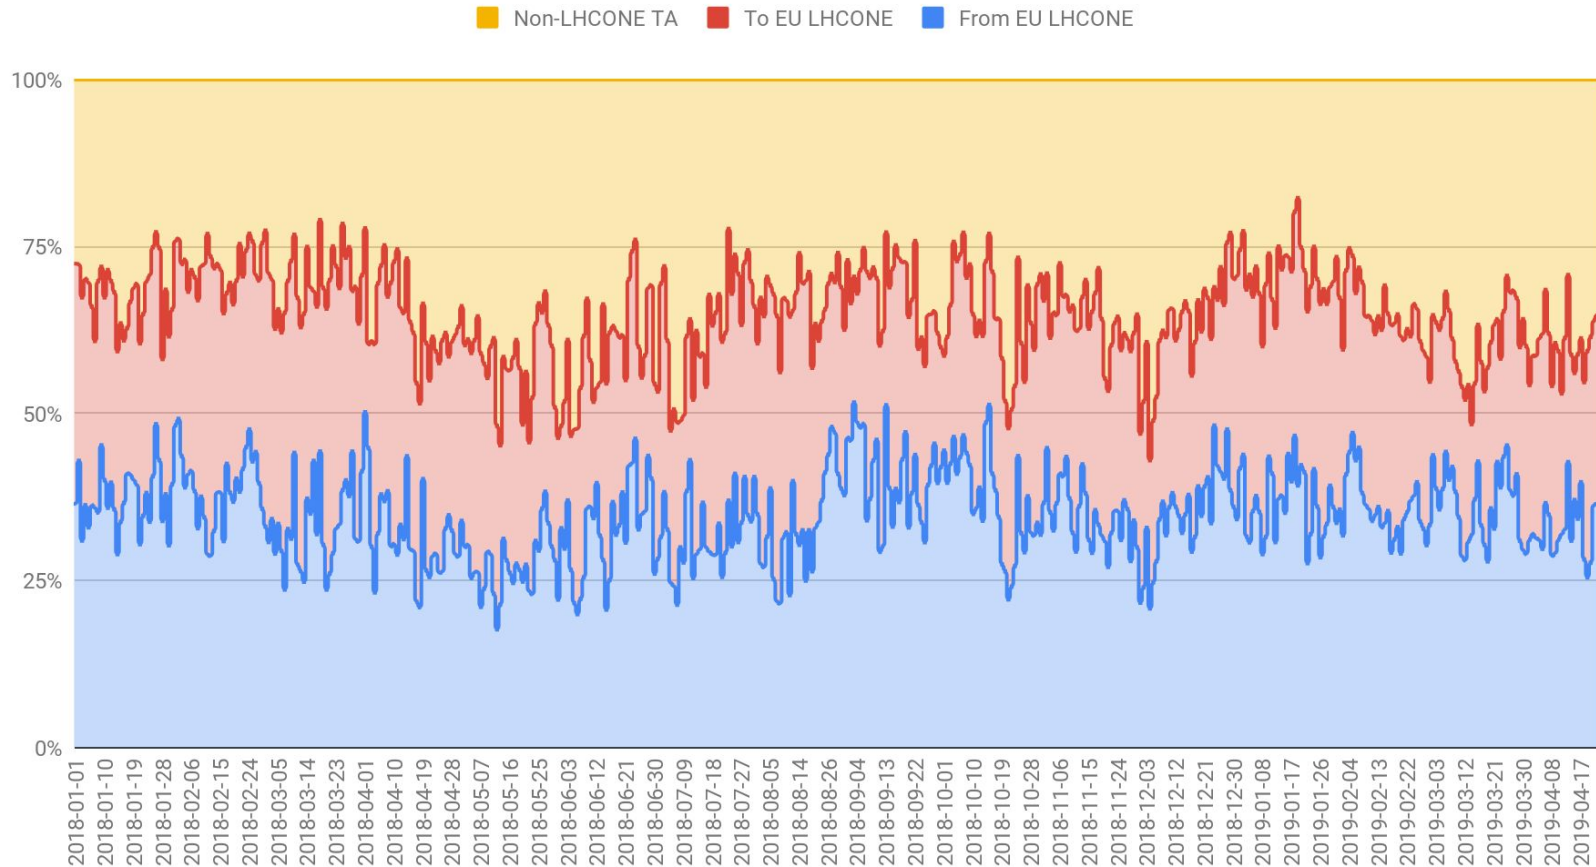
T2 Top Talkers

Note: XRootD is everything port 1094 in these numbers.



Transatlantic - who is using the links?

ESnet TA Traffic Distribution LHCONE vs Other Traffic




Observation: most of our TA traffic is LHCONE (58.8% on average).



Transatlantic - by month

LHCOPN traffic is in the non-LHCONE


Total Volume on ESnet TA Links by Month



Observation: LHC maintenance did only have a slight effect on the data volume.




Transatlantic - direction analysis



Observation: EU->US is only slightly higher than US->EU.

Transatlantic - day of week

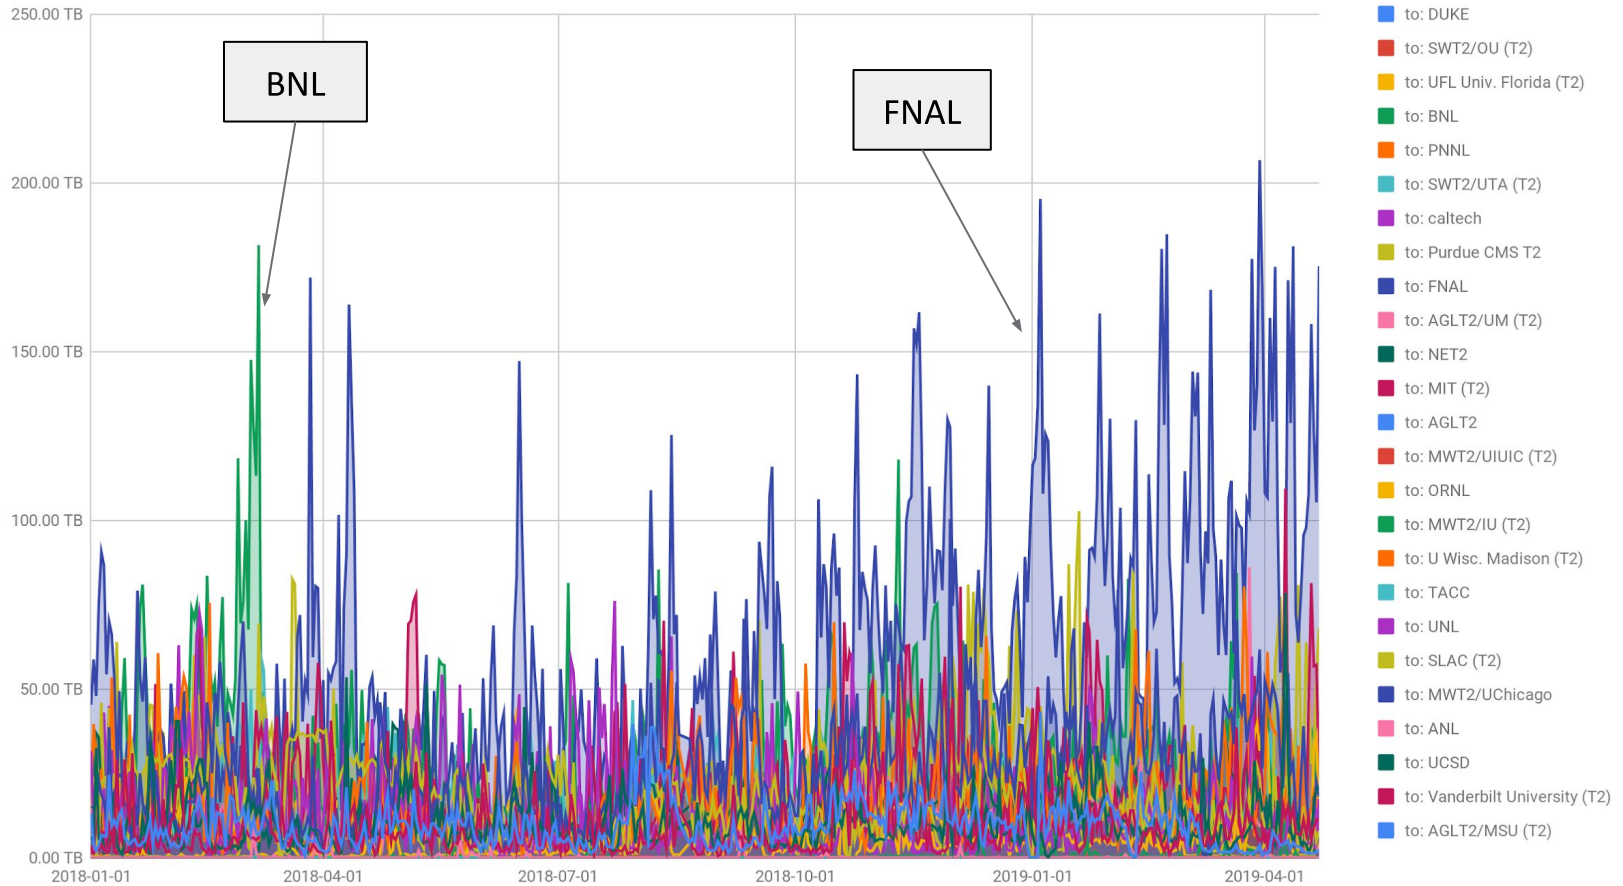
TA Data Volume - Day of Week



Observation: While Thursdays are little bit weaker than other days, there is no significant difference in data transfers between days.

Transatlantic - from EU to US by site

Daily Volume from EU to US




Fermilab is by far the top consumers of TA bandwidth from EU -> US.



Transatlantic - from US to EU by site

Daily Volume to EU from US sites

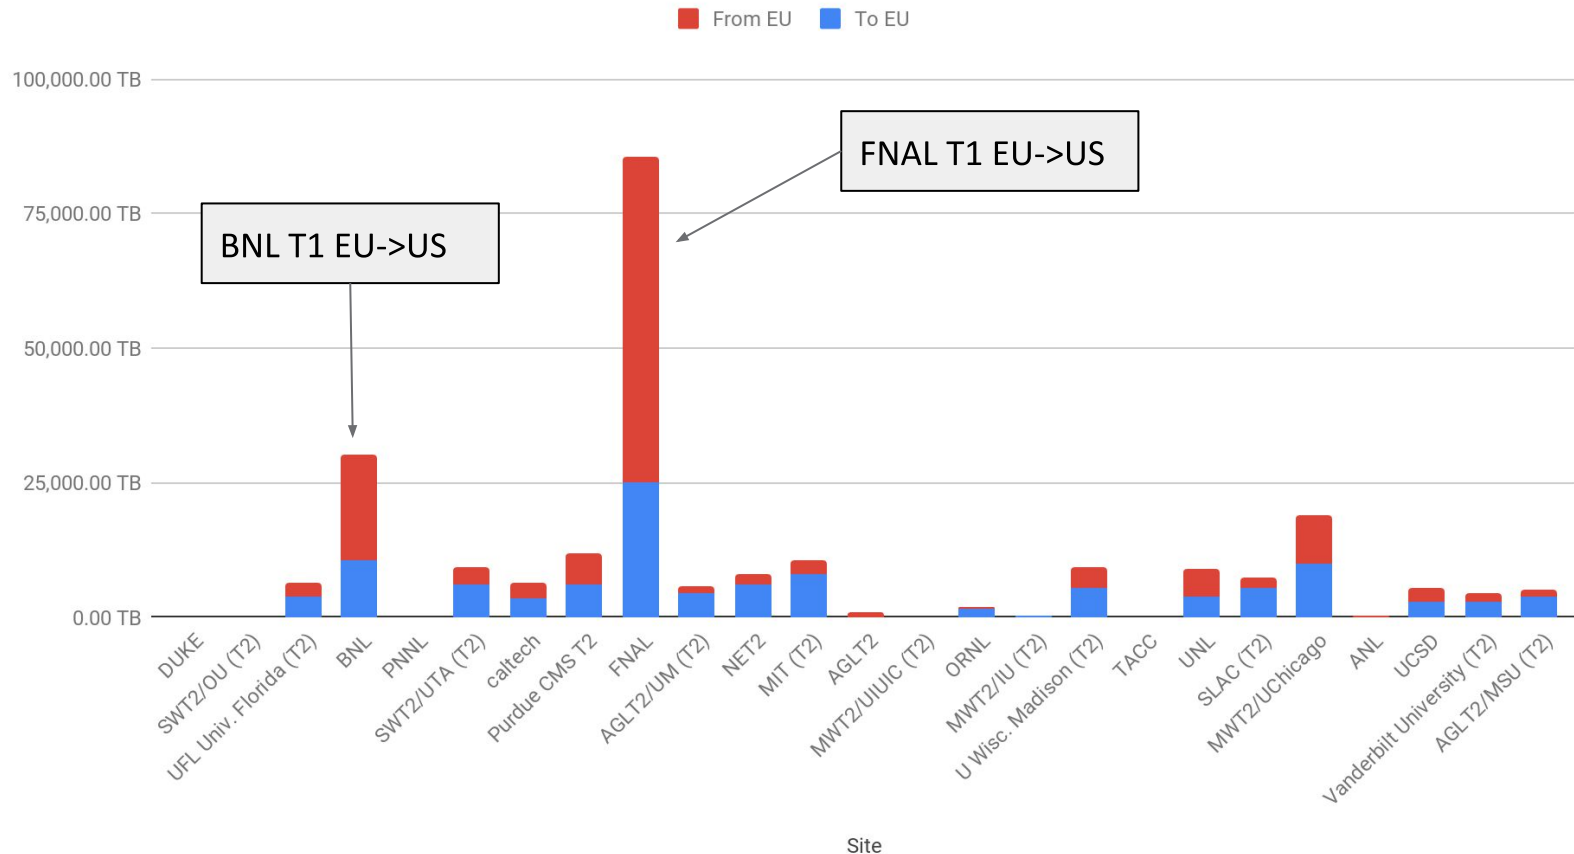


Fermilab and BNL are the top consumers of TA bandwidth from US -> EU.



Transatlantic - LHCONE TA

To EU and From EU




Tier model in operation: most traffic from EU comes to BNL and FNAL.



Transatlantic Forecasting

European Demand and Capacity Forecasts



Transatlantic Forecasting - Observations

- ESnet's TA usage continues to grow - LHCONE traffic is the main contributor. Approximating exponential growth on TA:
 - ~50% increase yearly - based on last year
 - ~40% increase yearly - based on last 3 years
- **Traffic growth changes continuously** - we need to constantly monitor and adjust predictions
 - we have seen a general slowdown in growth during the last 5 years
 - but it does not always apply to LHCONE
 - re-start of LHC will likely increase data rates
 - new LHCONE participants and possible expansion of the LHCONE network will also increase traffic

Conclusions

- ESnet performs extensive flow analysis on LHCONE traffic
 - ESnet's numbers *serve as a reference* for TA traffic
- ESnet will continue to generate **monthly reports** on:
 - LHCONE US utilization and traffic matrix
 - LHCONE TA utilization and forecasting
- HEP-ESnet data sharing provides benefit for both parties
 - multiple inefficiencies, monitoring, accounting and routing issues were identified by comparing CMS/ATLAS numbers with ESnet numbers already
 - Some common understanding of traffic patterns has been achieved, but we need to *keep this effort alive*

Thank you!

Richard Cziva, ESnet - richard@es.net

