1986–2016
ESnet
30 YEARS OF NETWORKING
AT THE SPEED OF SCIENCE

ESnet-LHC: Opportunities

Joe Metzger
Team Lead
ESnet Network Planning Team
Lawrence Berkeley National Laboratory

Pre-GDB
CERN
January 10 2016

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

Overview

- Supporting LHC science is a key part of ESnet's mission, and one of the motivations for ESnet6
- Changes in the LHC network flows in the WAN can have significant impact on all users of the network
- ESnet believes that a forward looking technical engagement with a few key LHC folks will be strategic to our partnership and contribute to optimizing LHC workflows

ESnet Support for LHC

- Over half of ESnet's total traffic is LHC related.
- ESnet's historic traffic growth rates have been 10X every 4 to 5 years.
- Some future LHC traffic growth rate estimates are above, others are below this rate, but nothing definitive, so ESnet is assuming 10X growth every 5 years at this point.
  - Is this reasonable?
- The ESnet6 project will face multiple challenges trying to meet this growth rate
  - Funding
  - Complexity
  - Vendor equipment delivery timelines & speed of evolution in the market
  - Space & power
  - The unknown
    - There are often surprises scaling by a factor of 10

ESnet6

- Address the next ten years of science requirements
- Three key design goals
  - **Capacity** - Handle exponential traffic growth at reasonable cost
  - **Reliability and Resiliency** - distributed science facilities, computing, data – scientists depend on the network for their science research to work
  - **Flexibility** - support changing compute models, near real-time analysis, 'superfacility', etc.
- Working/Proposed Schedule
  - 2016-2017 - Research architecture options, requirements, capabilities & services
  - 2017 - Architecture and proposal development
  - 2017-2019 - Reviews & Approvals on scope, budget & schedule, get funding & continue architecture refinement
  - 2019-2020 - Implementation & transition
-

1986–2016 ESnet 30 YEARS OF NETWORKING AT THE SPEED OF SCIENCE

Historic ESnet Growth Rates from my.es.net

ESnet
ENERGY SCIENCES NETWORK

Fasterdata | MyESnet | ESnet

HOME »

Traffic Volume

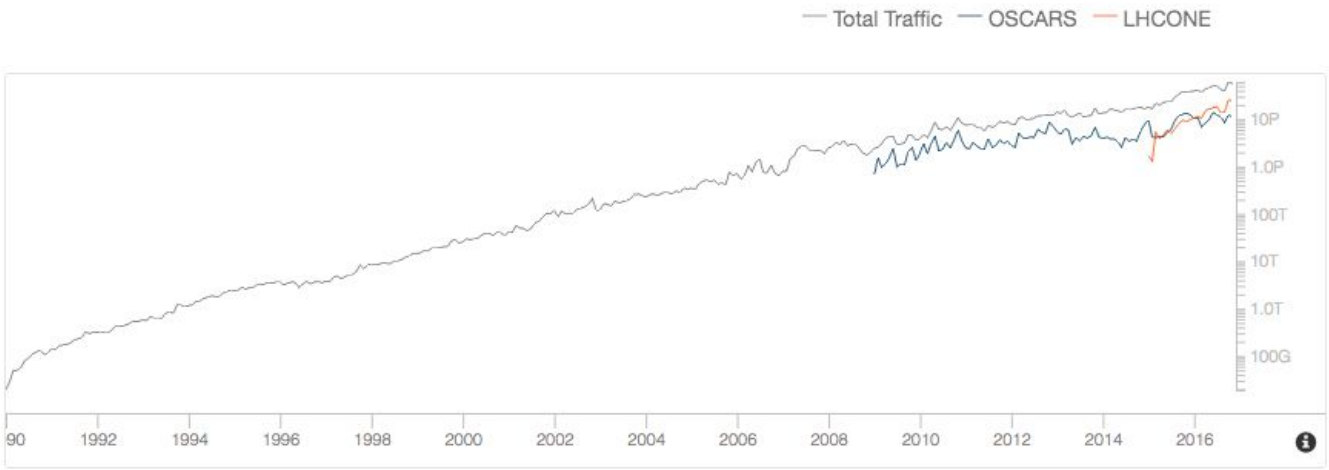

Options

MONTH

December 2016

TIME RANGE

All

VIEW

Overlay Summary

Interface Detail

SCALE TYPE

Log

Linear

December 2016

	Bytes	Percent of Total	One Month Change	One Year Change
OSCARS	12.69PB	19.9%	+13.3%	+2.87%
LHCONE	22.92PB	36.0%	-4.44%	+136%
Normal traffic	28.05PB	44.1%	+31.4%	+72.2%
Total	63.65PB		+12.6%	+66.0%

Forecast of ESnet 2020 link capacities required to meet demand based on historic traffic statistics

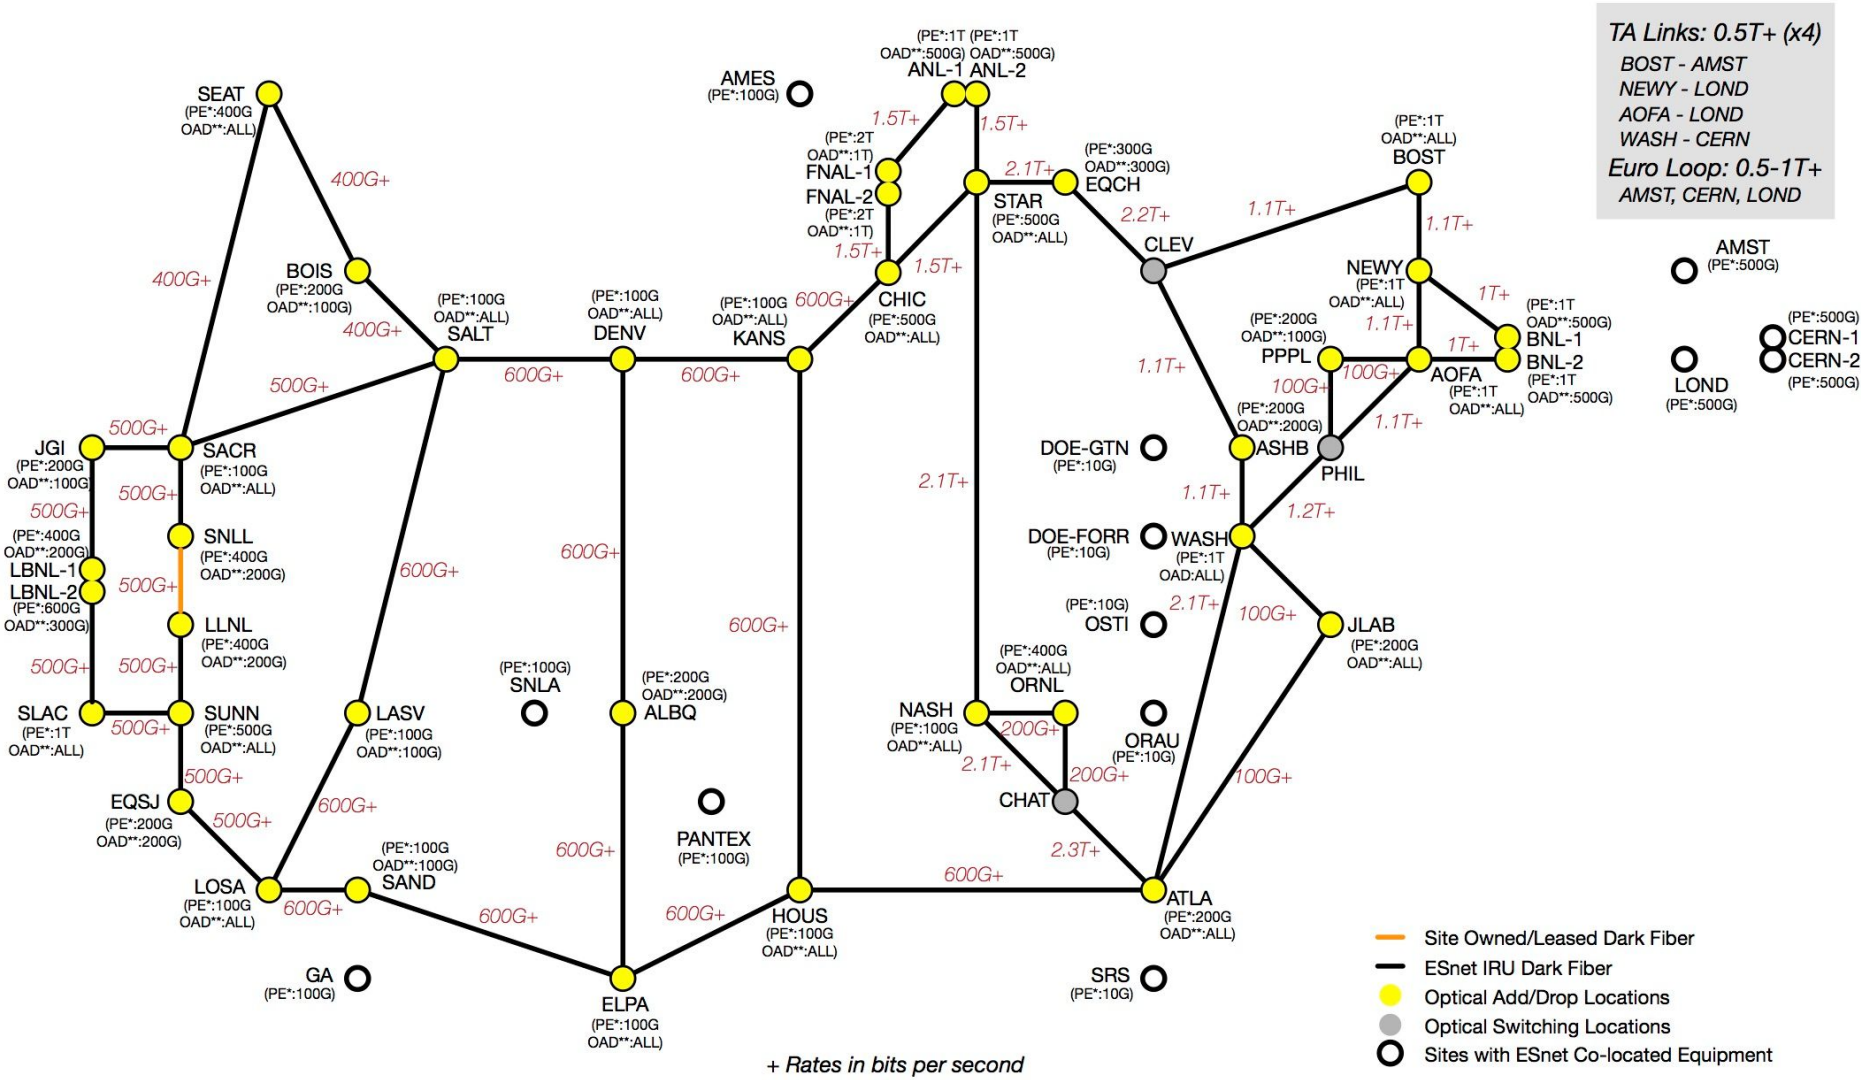

+ Rates in bits per second

** Provider Edge: N (where N denotes aggregate bandwidth to customers connected to node, includes redundant bandwidth to sites) [Jan 2020]*

*** Optical Add/Drop: N (where N denotes add/drop bandwidth on **EVERY** degree of node, including connection between nodes at 2-node sites) [Jan 2020]*

NB: Some sites without PE routers on local premises are not shown in this diagram

Forecast of ESnet 2025 link capacities required to meet demand based on historic traffic statistics

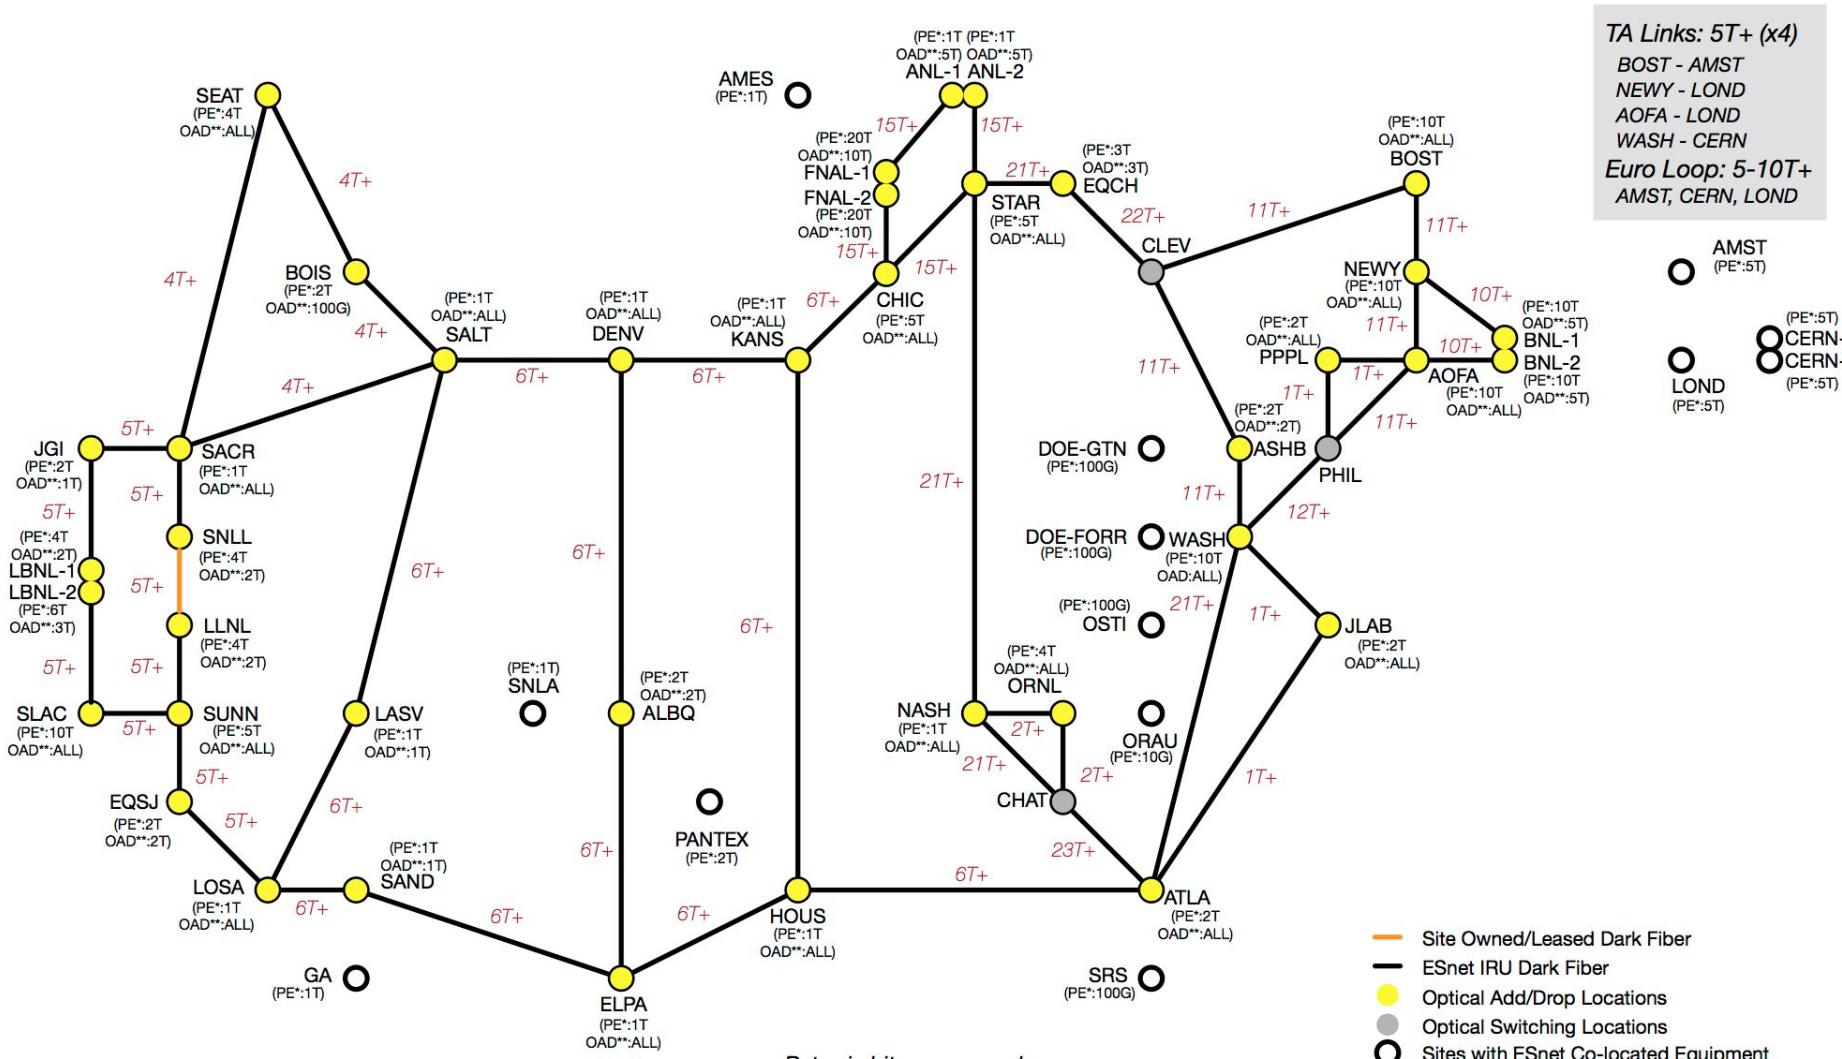

+ Rates in bits per second

** Provider Edge: N (where N denotes aggregate bandwidth to customers connected to node, includes redundant bandwidth to sites) [Jan 2025]*

*** Optical Add/Drop: N (where N denotes add/drop bandwidth on **EVERY** degree of node, including connection between nodes at 2-node sites) [Jan 2025]*

NB: Some sites without PE routers on local premises are not shown in this diagram

How LHC workflows use the network impacts Wide Area Network (WAN) operations & planning.

Some of the aspects of the LHC workflows that have significant impact on the networks supporting them are:

- Structure of the network flows
  - Large number of small streams vs small number of large ones
  - Congestion protocols - both aggressiveness & fairness
  - Sensitivity to interruptions, re-routing, loss, packet reordering & changes in latency
  - Burstiness - on all timescales
- The location of the endpoints and the networks that serve them

One example: Where will the 'extra' compute cycles to analyze the 'extra' CY16 data come from?

- Capacity or efficiency upgrades to existing Tier [123]
  - No significant impact on the networks. Continue business as usual
- HPC centers
  - Some are not currently integrated into LHCOPN or LHCONE, so will need some effort, but probably not too much new infrastructure investment
- Cloud Providers
  - May require significant new investment to get high throughput networking to Virtual Private Clouds depending on which cloud providers and which reliability zones/data centers.

ESnet Europe to US Traffic October 2016 Utilization sorted by Time

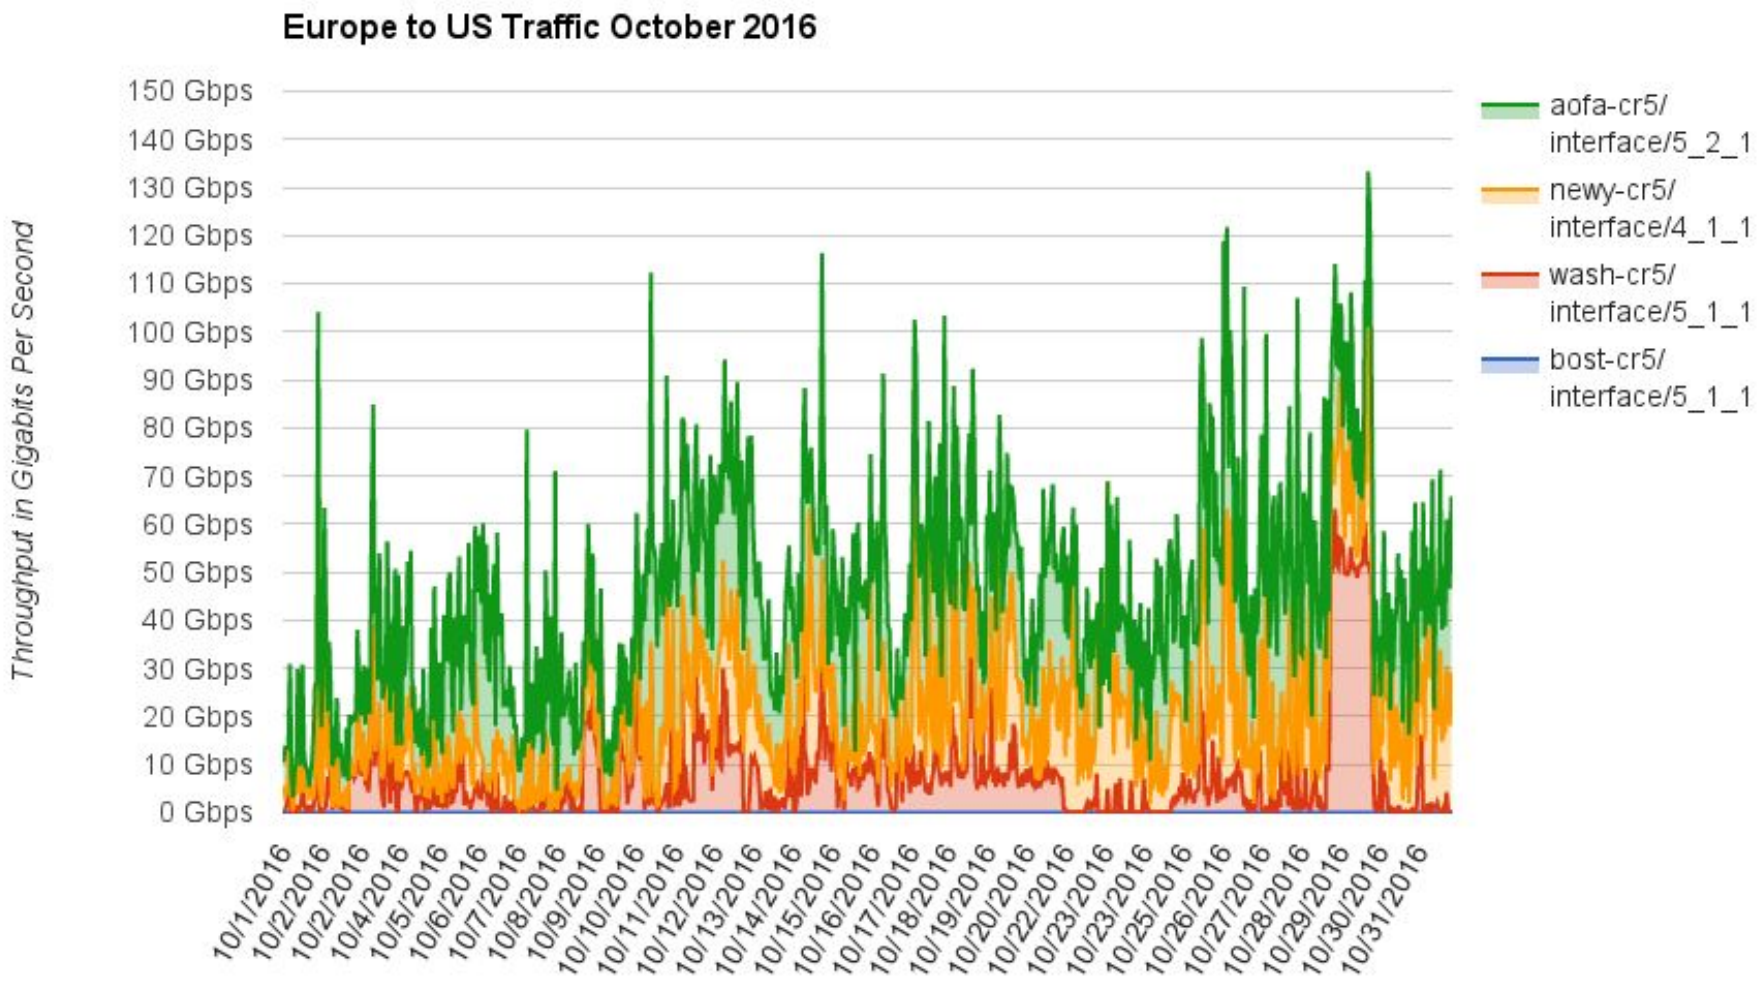

Same data points sorted by Utilization instead of Time

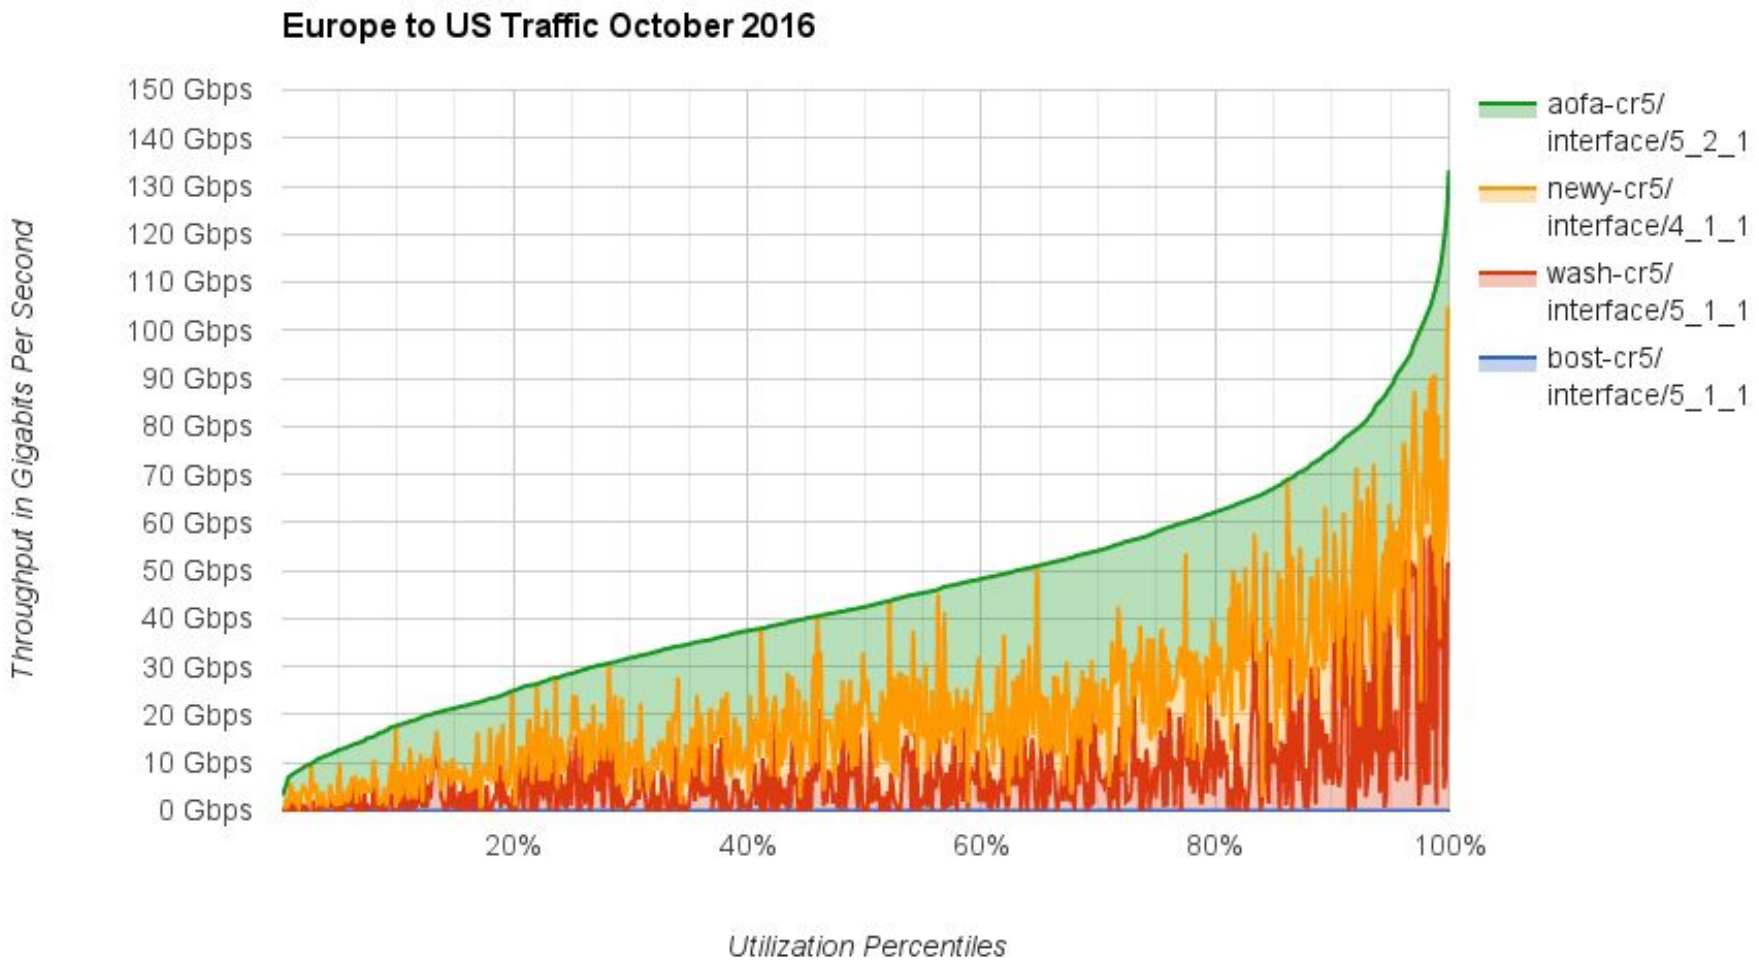

Peak Utilization Drives Network Operations & Planning

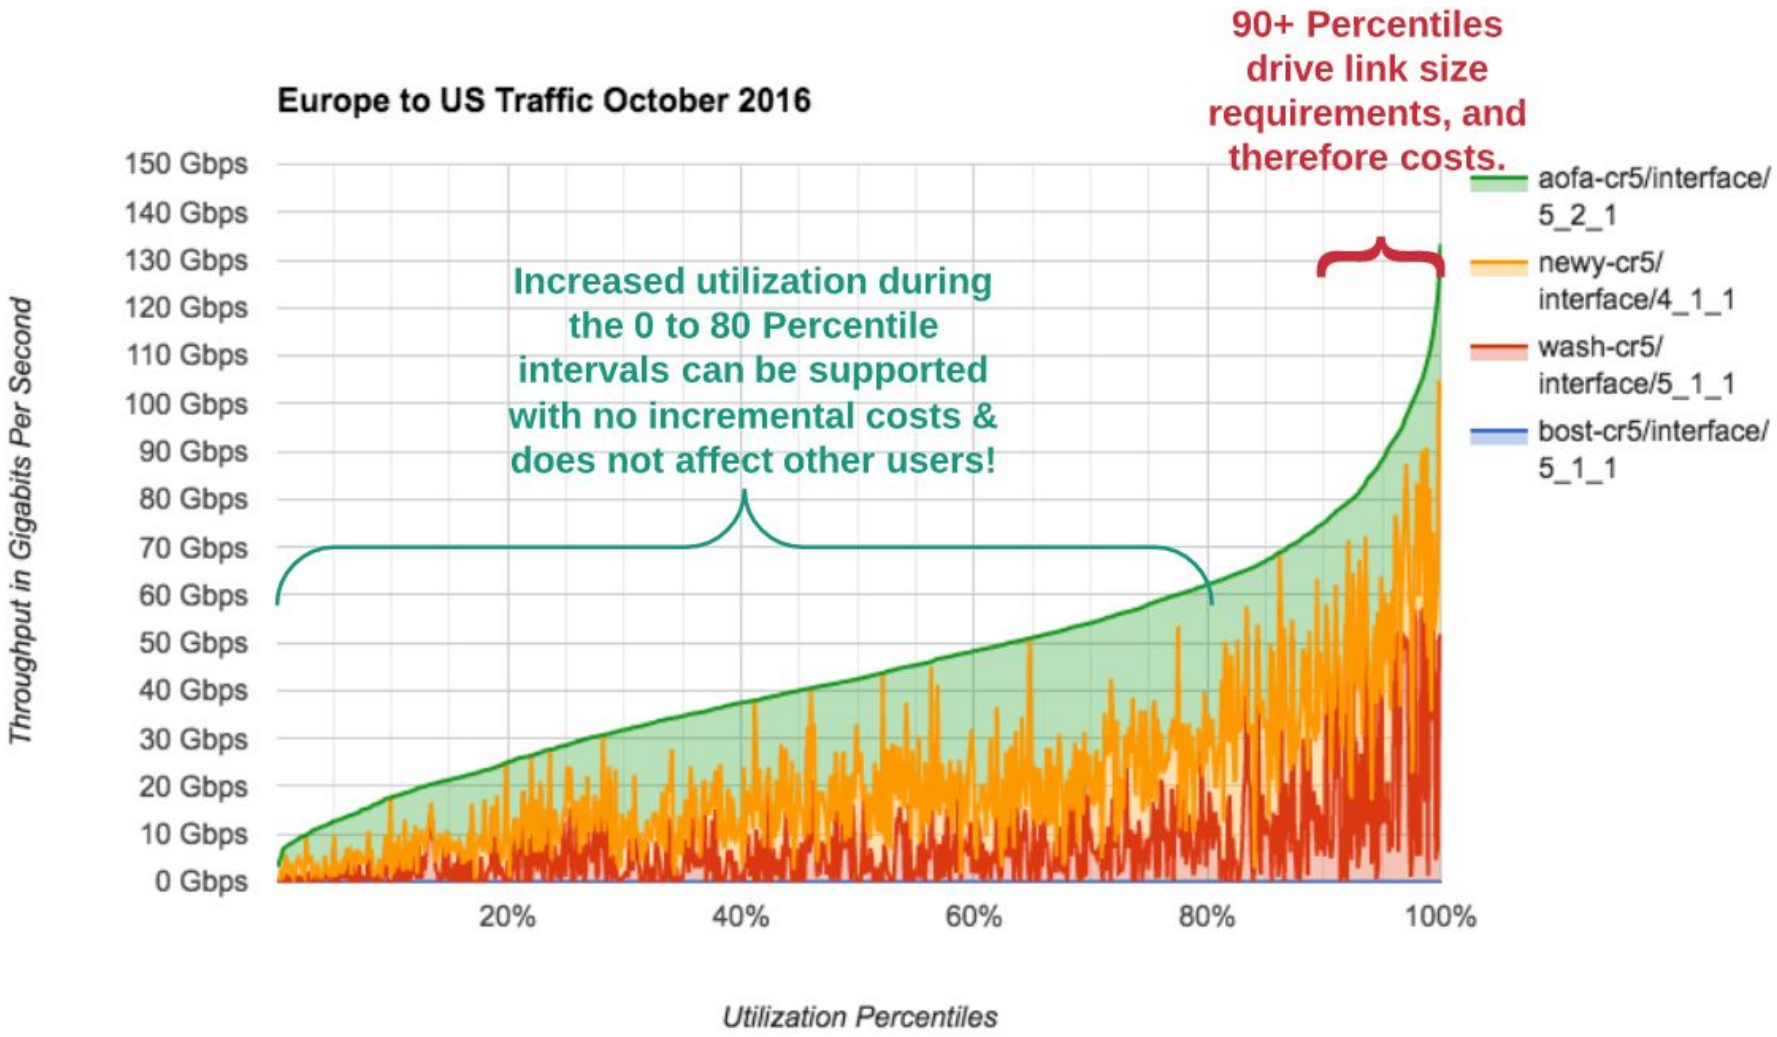

Running at higher levels of link utilization is possible, but communication is key to prevent unintended consequences

We need to make sure we don't get stuck in feedback loops like the following:

1. Scientists tune data transfer systems to run as fast as possible which generates more bursty behaviour with higher peak transfer rates.
2. Network Operators see the higher rate bursts as the new norm, and trigger upgrading the bottleneck links.
3. Repeat

Or,

1. Scientists discover idle network capacity, identify opportunistic uses of the capacity which will just go to waste if it isn't used anyway, and so they consume it.
2. Network Operators see the new higher sustained utilization, and trigger upgrading bottleneck links to get 'comfortable headroom' back.
3. Repeat

Summary

1. Supporting LHC science is a key part of ESnet's mission, and as such ESnet goes to considerable lengths to ensure that advanced network capabilities – both physical and logical – are available to the LHC community
   - Our understanding of the LHC networking requirements growth rate is around 10x every 4-5 years
   - This matches ESnet's 30 year history.
   - It will be challenging (especially funding), but should be possible.
2. Changes in the LHC network flows in the WAN can have significant impact on all users of the network
   - Potentially small changes in the way LHC community uses the network could have significant impacts on WAN providers. Especially as it relates to provisioning for peak capacity.
   - Care and planning are necessary so that the network can provide LHC traffic the capacity to meet the needs of the LHC community and at the same time allows the other science discipline uses of the network to coexist.
3. ESnet believes that a forward looking technical engagement with a few key LHC folks will be strategic to our partnership and contribute to optimizing LHC workflows
   - What is the best way to engage the LHC community so that ESnet, and other WAN network operators can provide the most valuable network services at the lowests costs optimized to support the LHC science mission?

The End

The following slides might be useful to support discussion, or answer questions.

ESnet6 R&D Architecture and Technologies Matrix

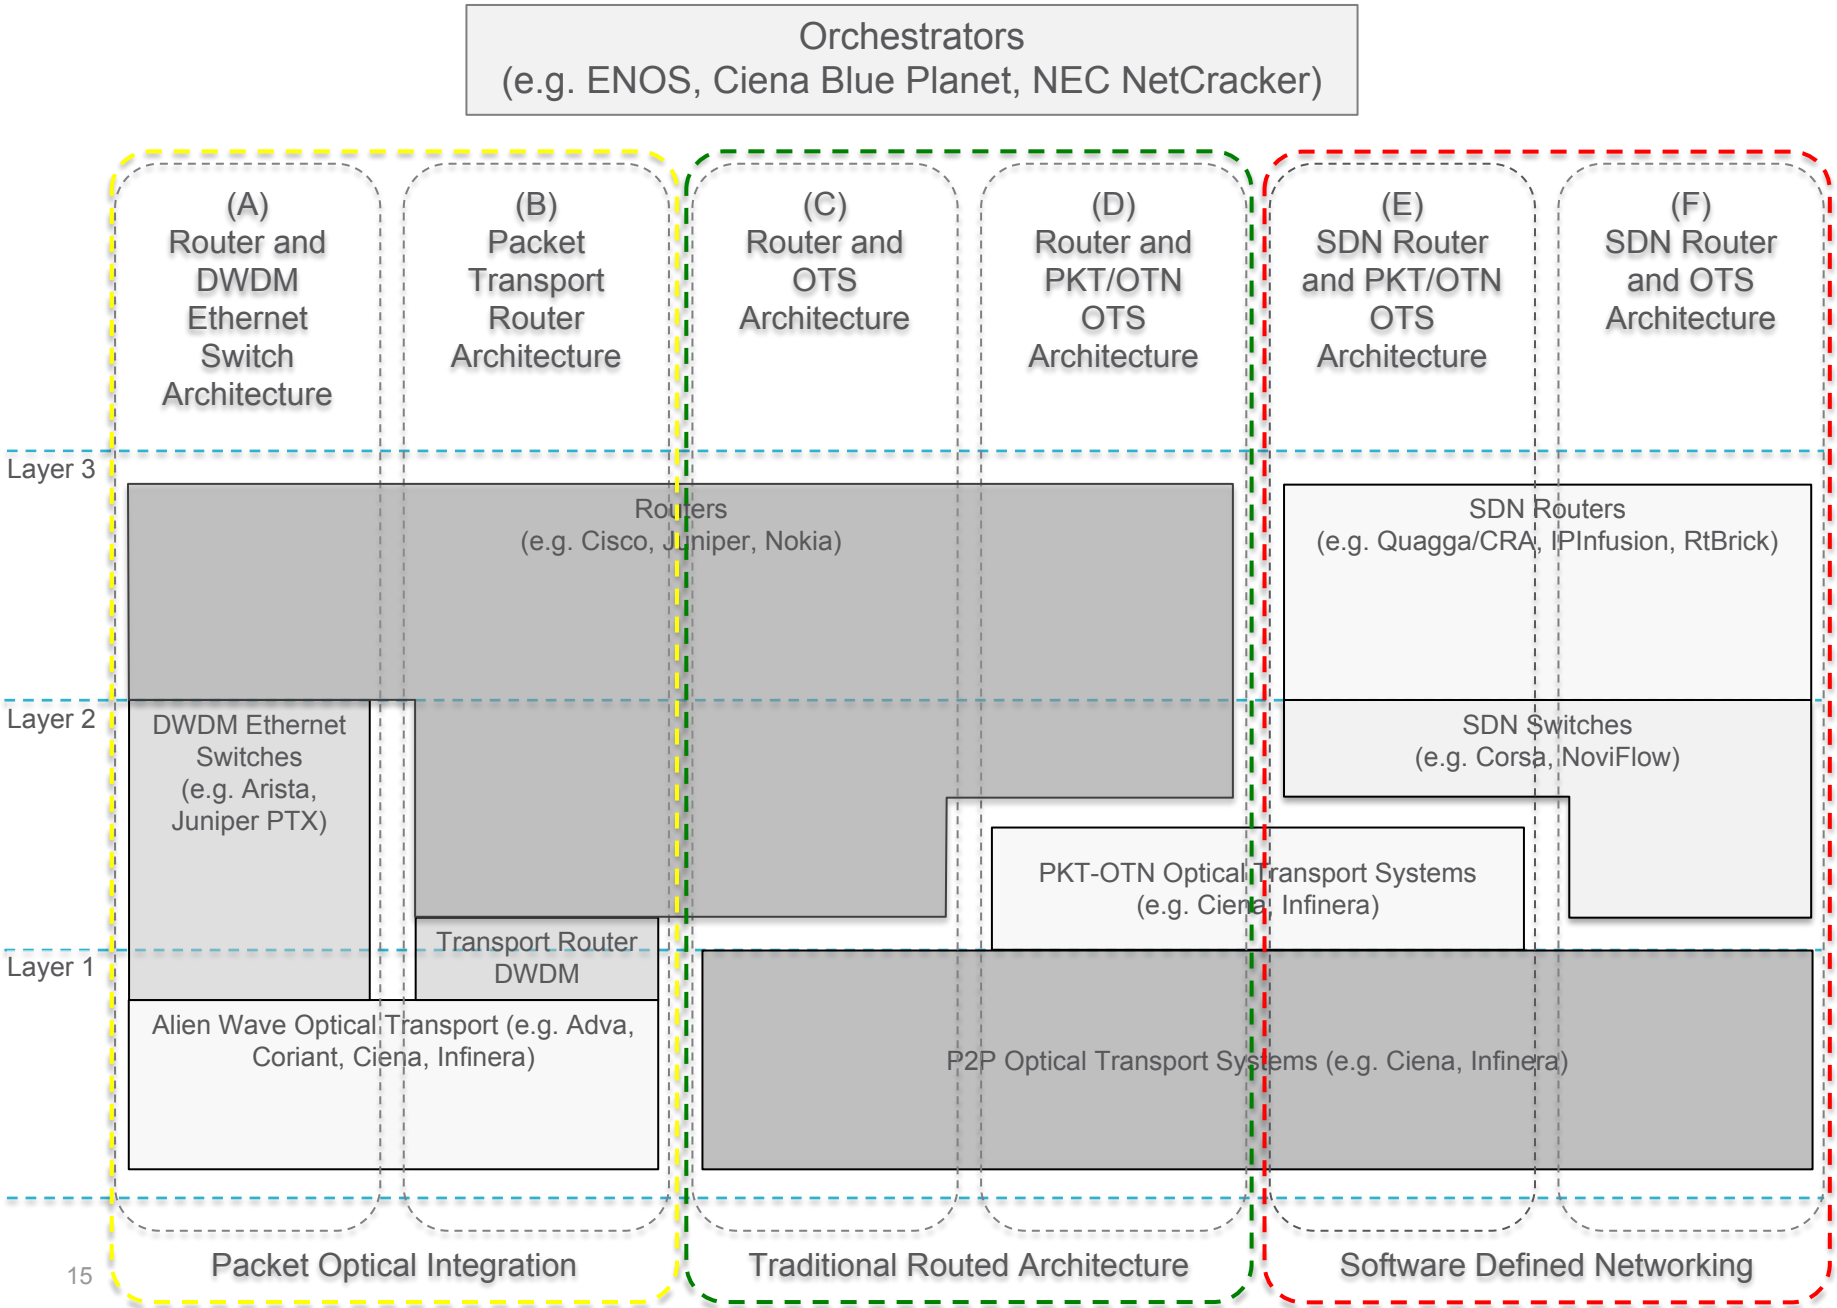

Some of the ESnet SDN Efforts

- SDN is just one of the many tools in the toolbox!
- ESnet's main SDN efforts are currently focused on how we might be able to use SDN to reduce costs and increase flexibility deploying network services in ESnet6
- Other ongoing ESnet SDN Activities
  - Collaboration with GEANT - defining SDN Operational Requirements
  - SENSE
    - CMS Workflows are one of the SENSE use cases
  - SDX

ESnet has a long history of SDN investigation

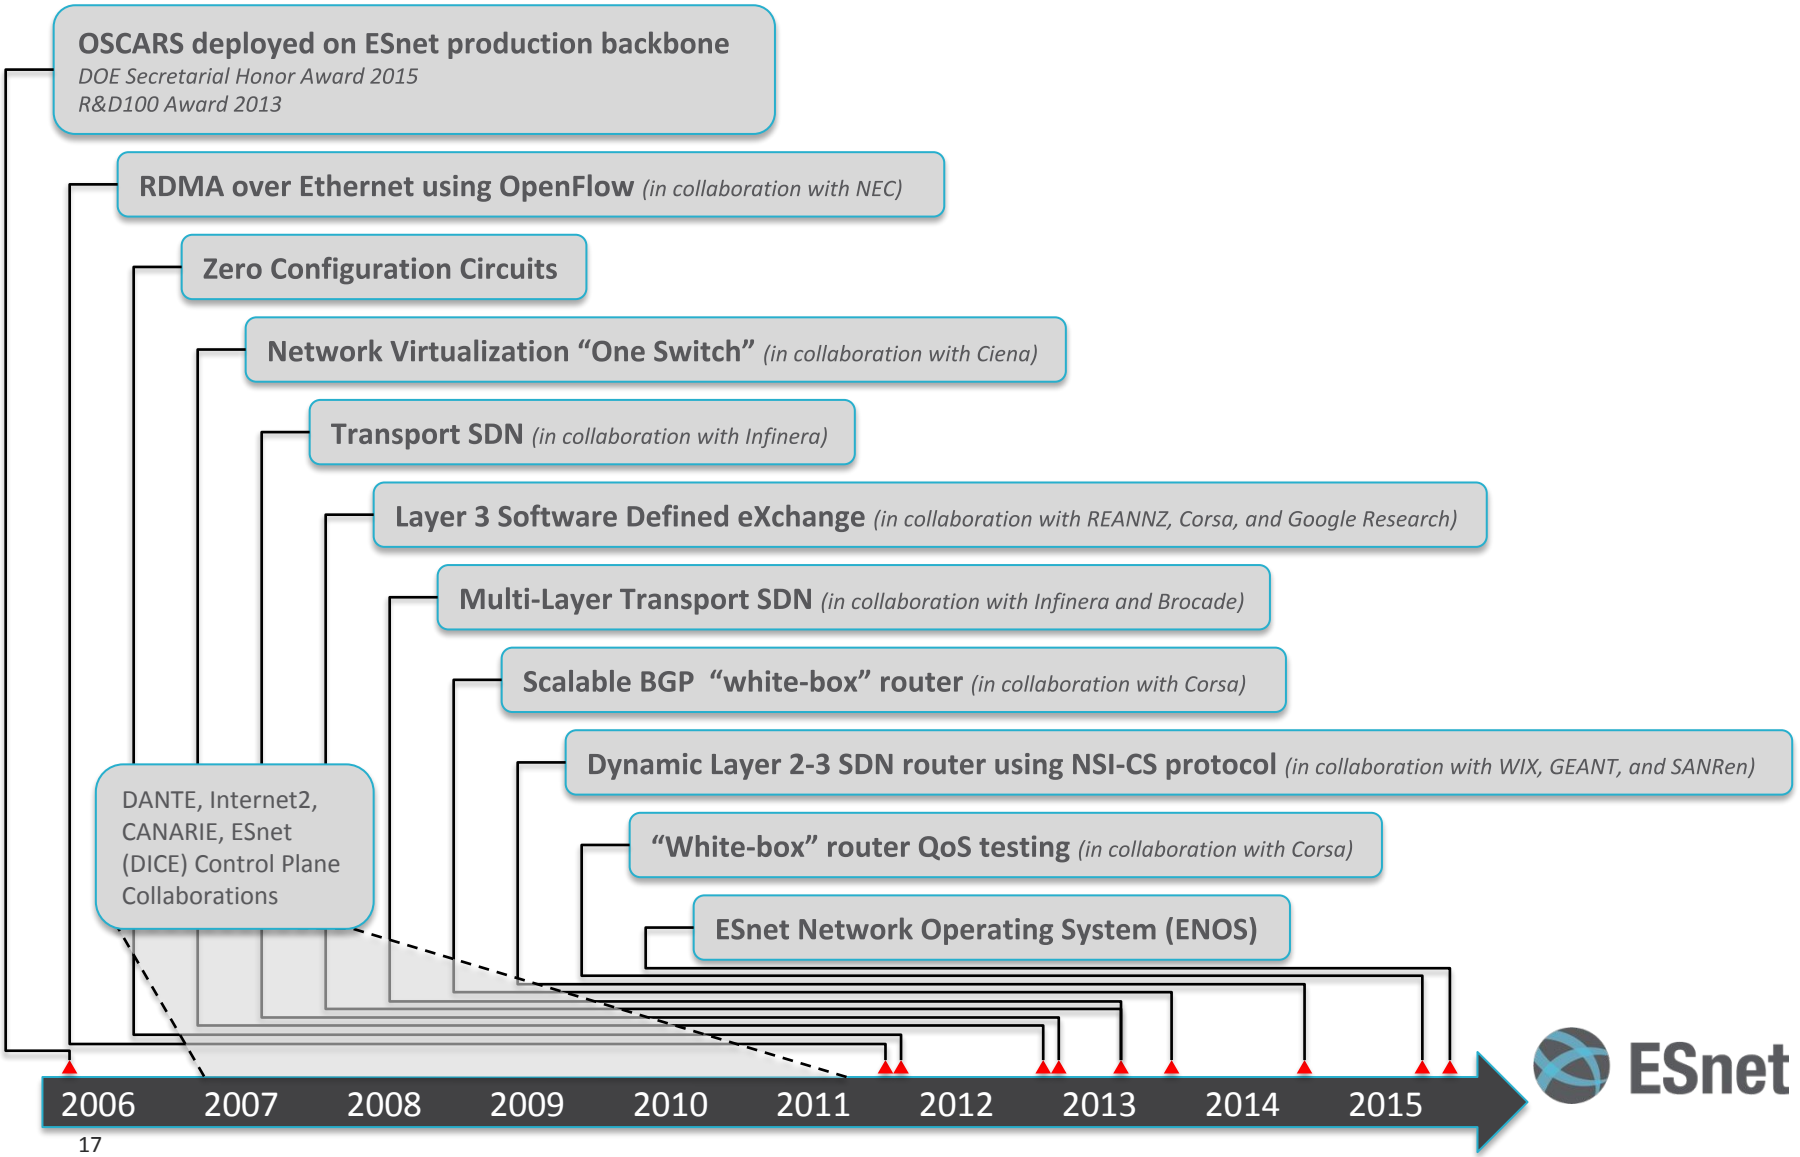

SENSE: SDN for End-to-end Networked Science for Extreme-Scale Science (FERMI, ANL, Caltech, UMD/Max, NERSC, ESnet)

Mar 2016 – Feb 2019

Goals

- **Leverage the emerging Software Defined Network (SDN) capabilities to develop intelligent, federated, end-to-end, science networking architecture friendly to data-intensive and network-aware distributed science applications.**
- **Allowing labs and universities to request and provision end-to-end intelligent network services from their application workflows.**

Impacts

- **Present geographically distributed resources (datacenters, instruments, etc.) as components of a local facility**
- **Simplifies complex massive datasets distribution with coordinated, multi-domain, smart and secure services**
- **Enable seamless application-network interaction for new near real-time distributed computing and data analytics**

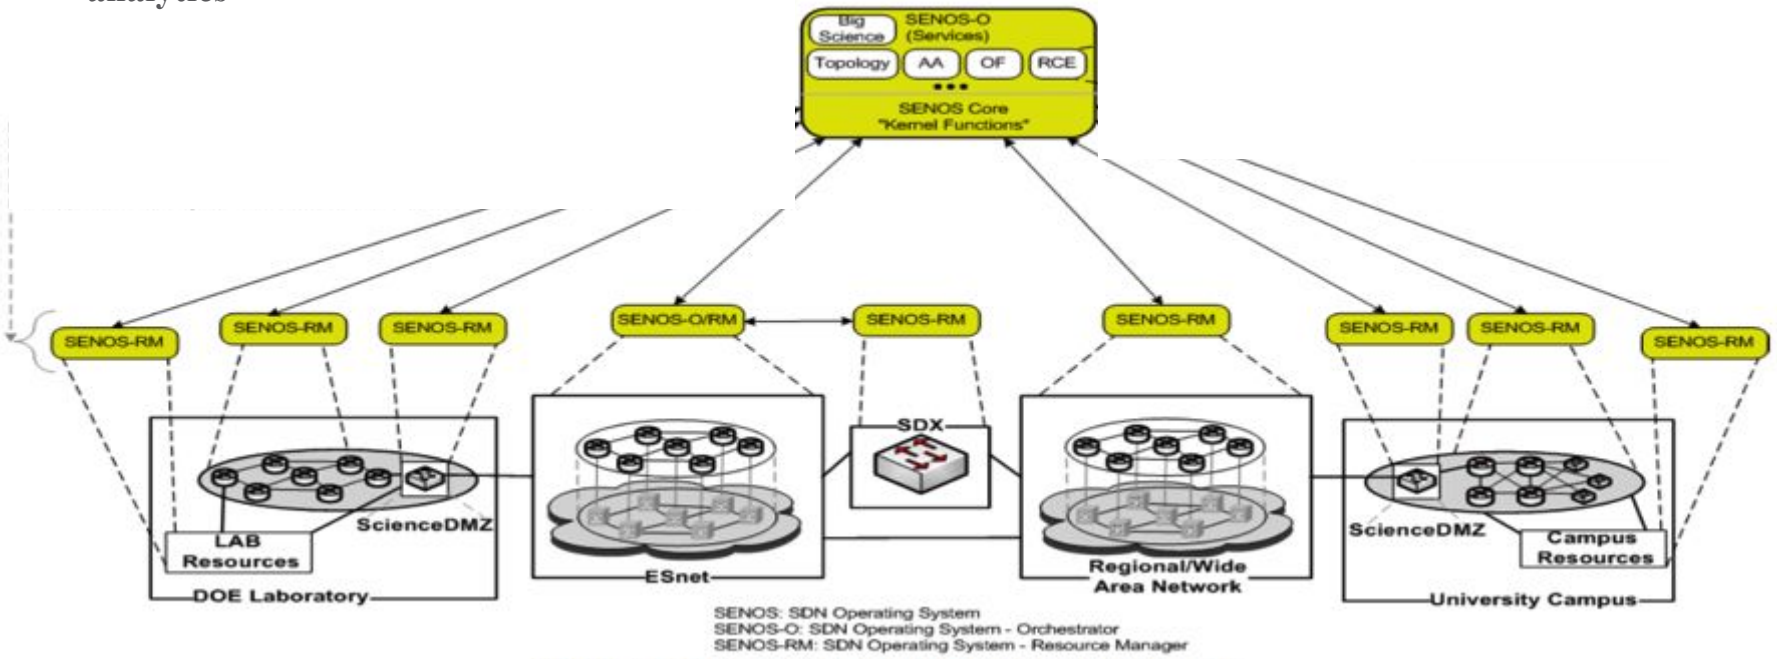

Figure 1. SENOS End-to-End Orchestration

LHCONE and ESnet TA Graphs

ESnet LHCONE Traffic - Aug 2016 (13.7 PBytes)

Breakdown of the Netflow data from ESnet's LHCONE implementation August 2016.

University Participants are US Universities only.

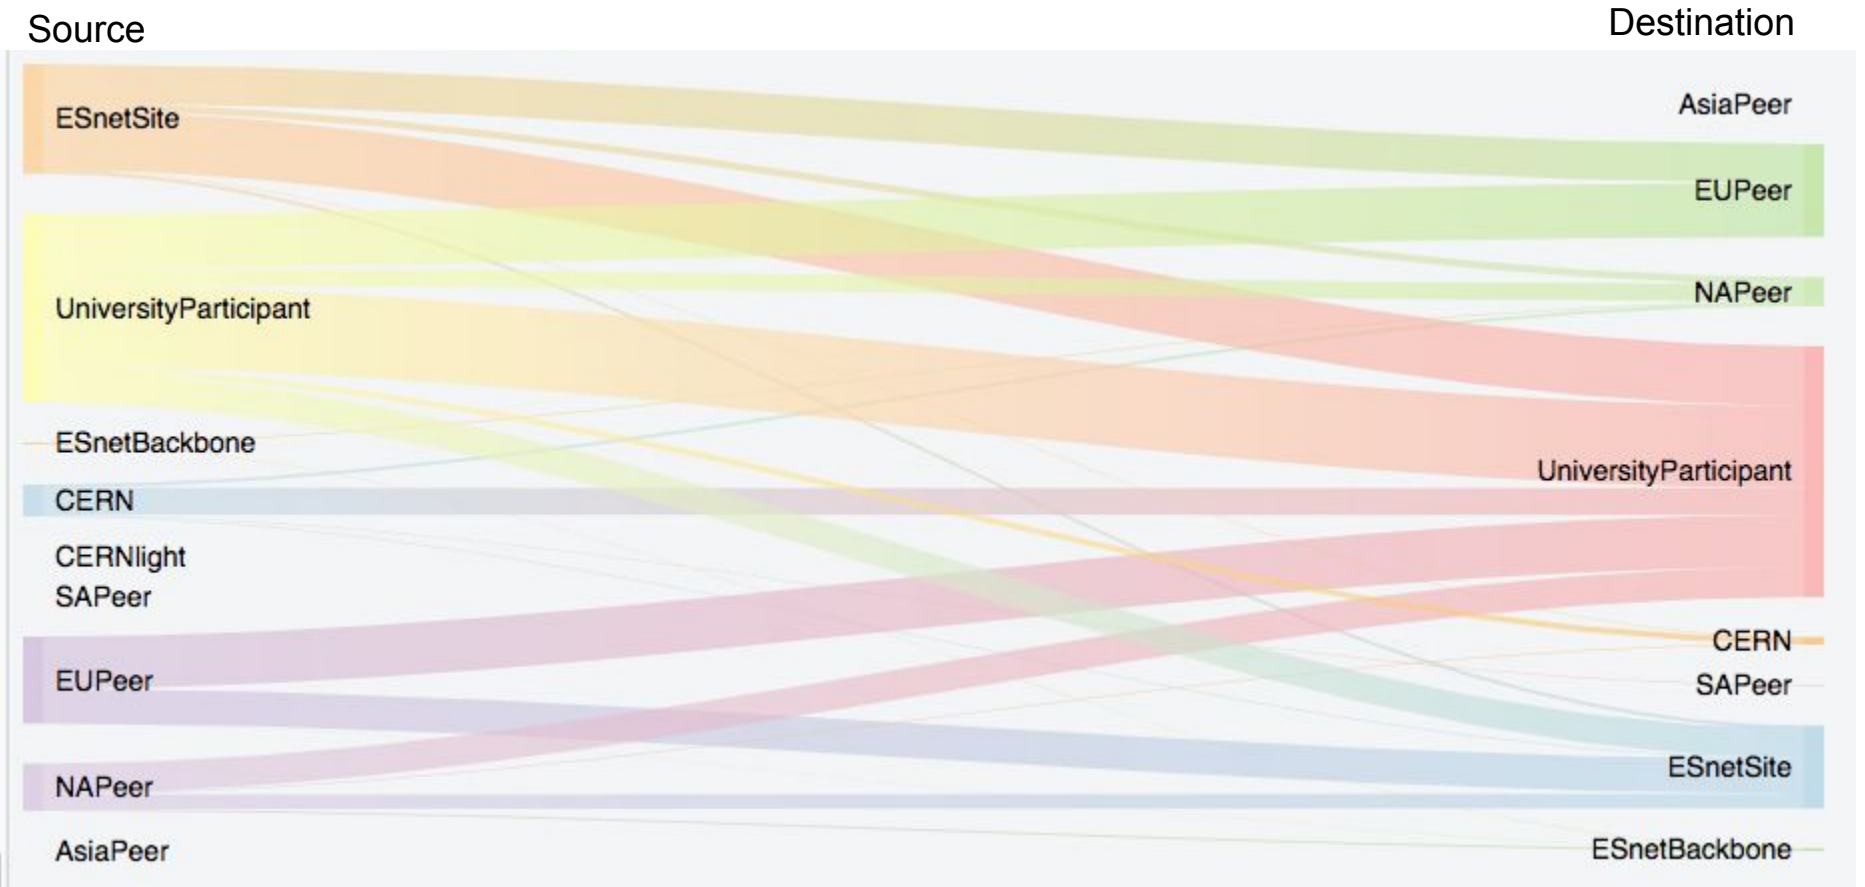

Breakdown of ESnet's Transatlantic Traffic

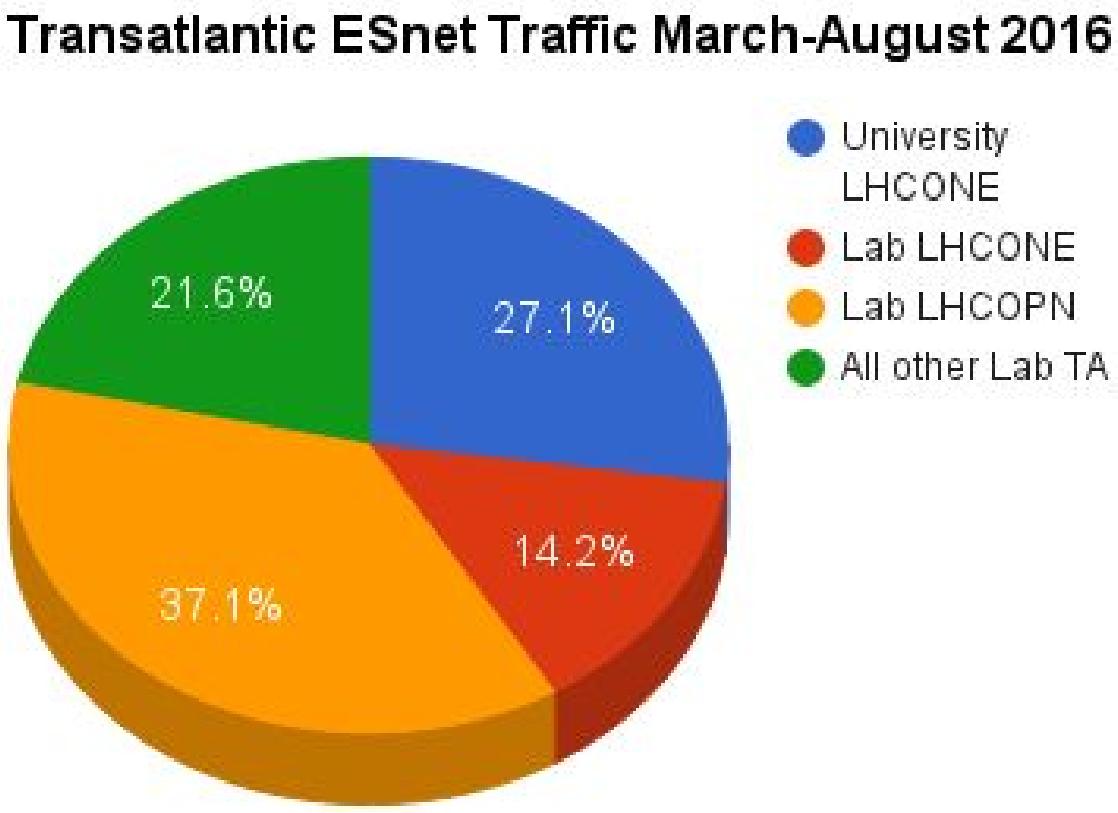

ESnet TA Circuit Reliability

More than 2100 Outages between January 2015 and November 2016.

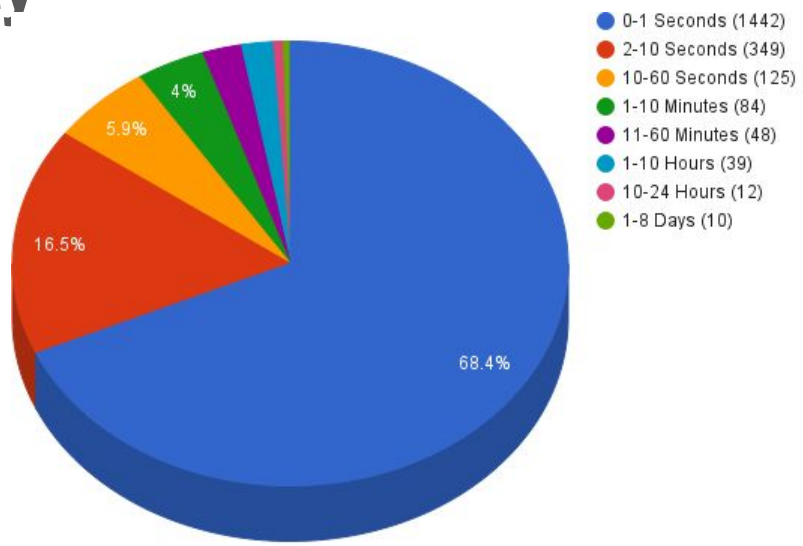

ESnet Available TA Capacity Jan 2015-Nov 2016

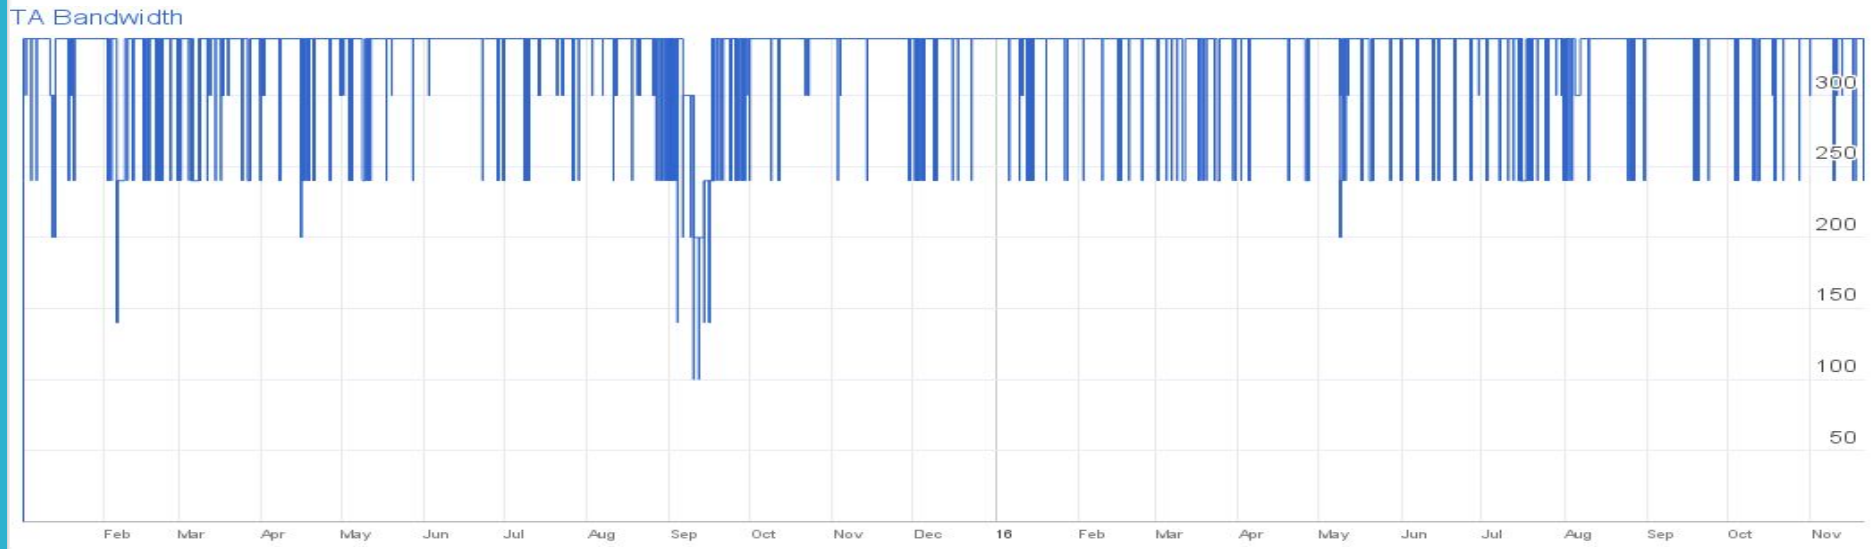

North America LHCONE Architecture (Oct 2016)

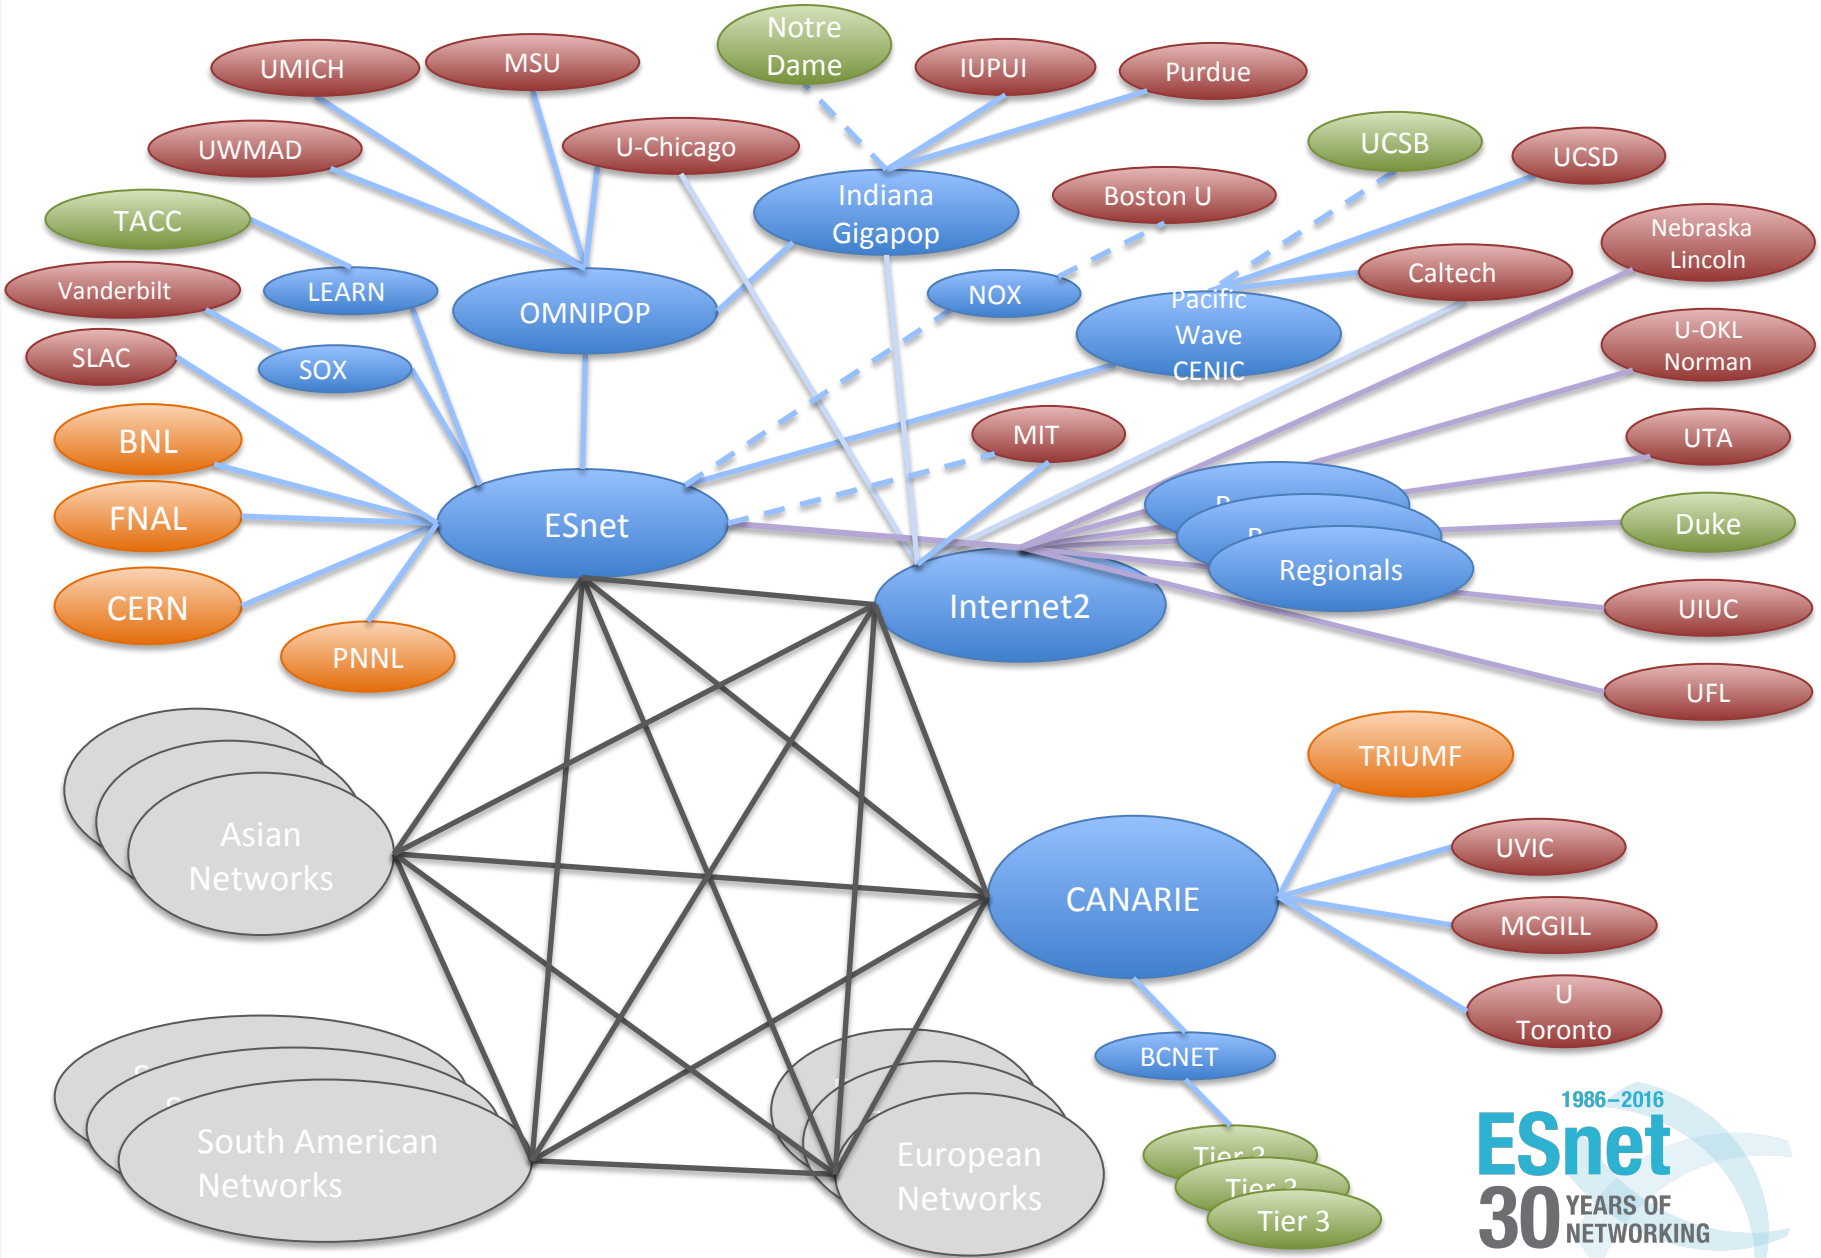

1986–2016 ESnet 30 YEARS OF NETWORKING AT THE SPEED OF SCIENCE