

Discovery, unconstrained by geography.



Inder Monga,
Interim Director, CTO, ESnet



ESnet is starting Run 3 in August



ESnet by Facts



The basic facts (new or notable):



High-speed **international** networking facility, optimized for DOE science missions:

- connecting 50 labs, plants and facilities with >150 networks, universities, research partners globally
- **340Gbps transatlantic extension in production (Dec 2014)**
- **university connections to better serve LHC science**
- \$35M in FY15, 42FTE
- older than commercial Internet, growing ~twice as fast
- the DOE user facility that serves all others

\$62M ARRA grant funded 100G upgrade in 2011:

- fiber assets + access to spectrum, shared with Internet2
- new era of optical networking, abundant capacity
- world's first 100G network at continental scale

Culture of urgency:

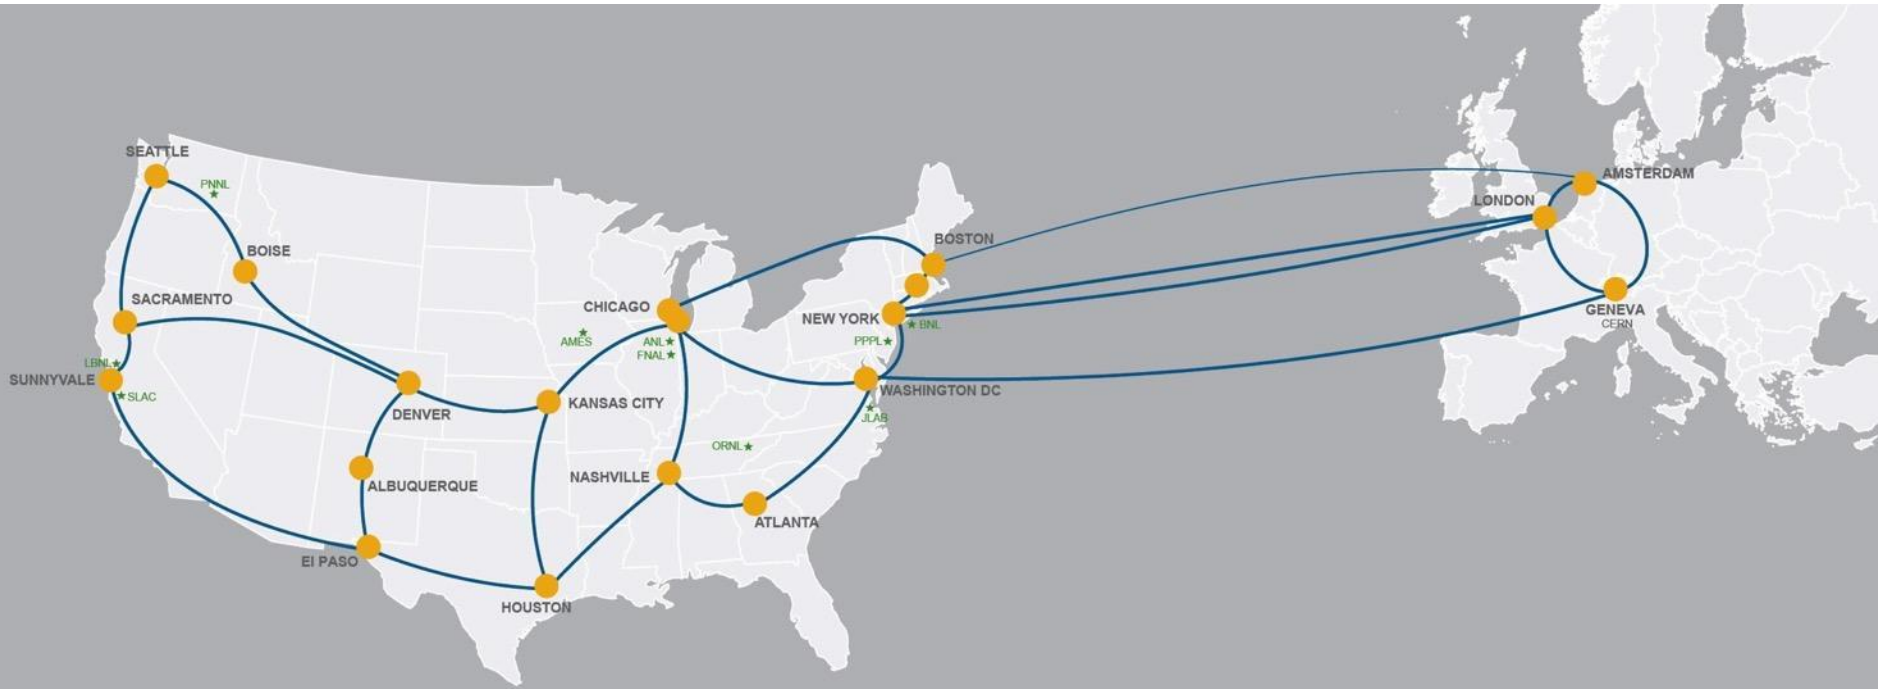
- **several recent awards**
- 80% engineers, highly motivated



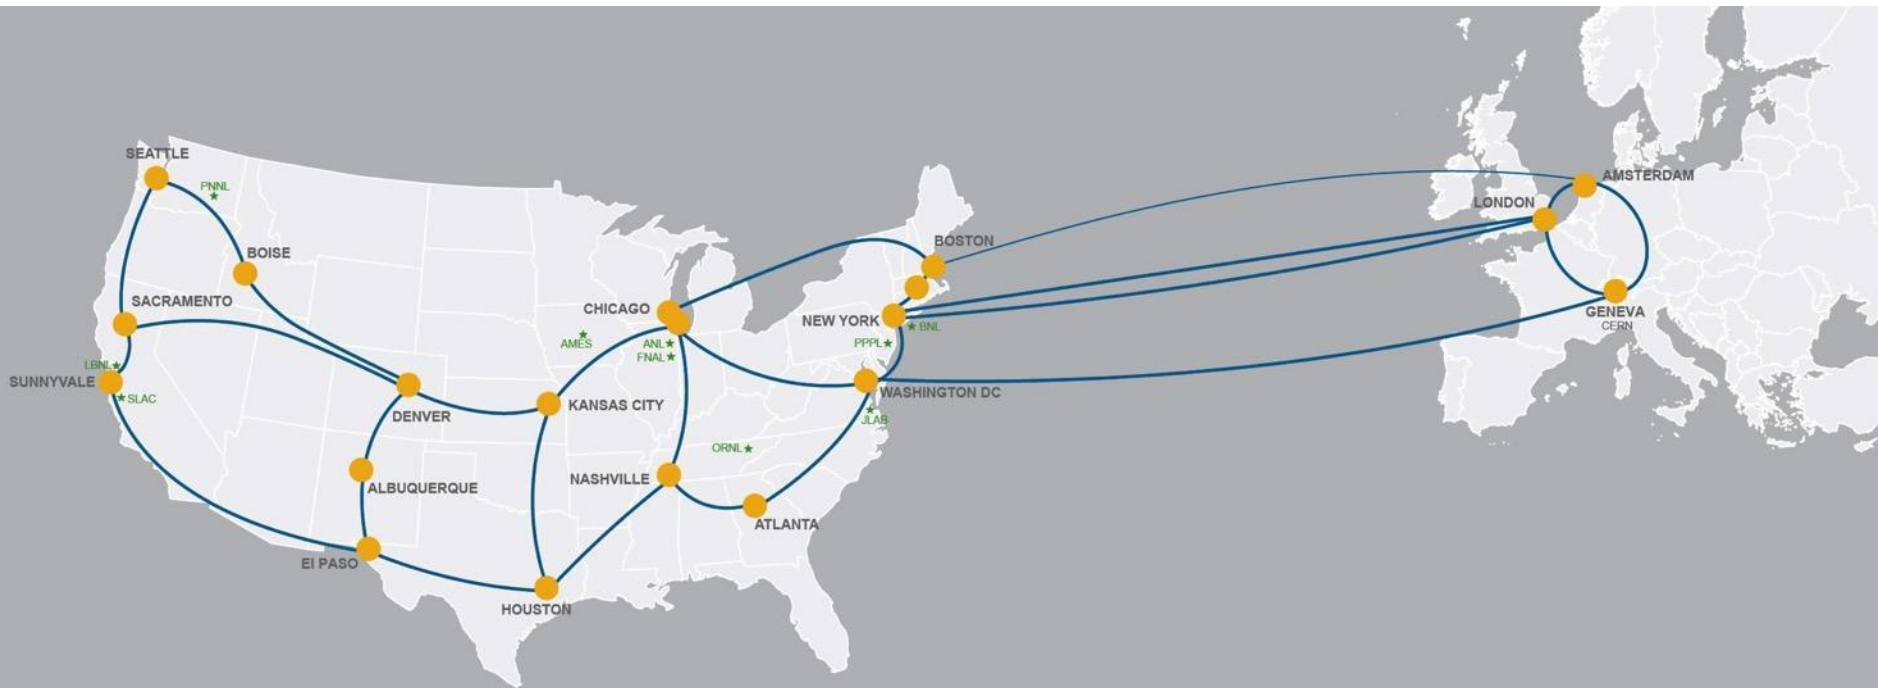
Our vision:

Scientific progress will be **completely unconstrained** by the physical location of instruments, people, computational resources, or data.

5th Generation ESnet: Current Network Footprint




ESnet is a **dedicated mission network** engineered to accelerate a broad range of science outcomes.



We do this by offering unique capabilities, and optimizing the network for data acquisition, data placement, data sharing, data mobility.

One of the oldest, fastest parts of the Internet.

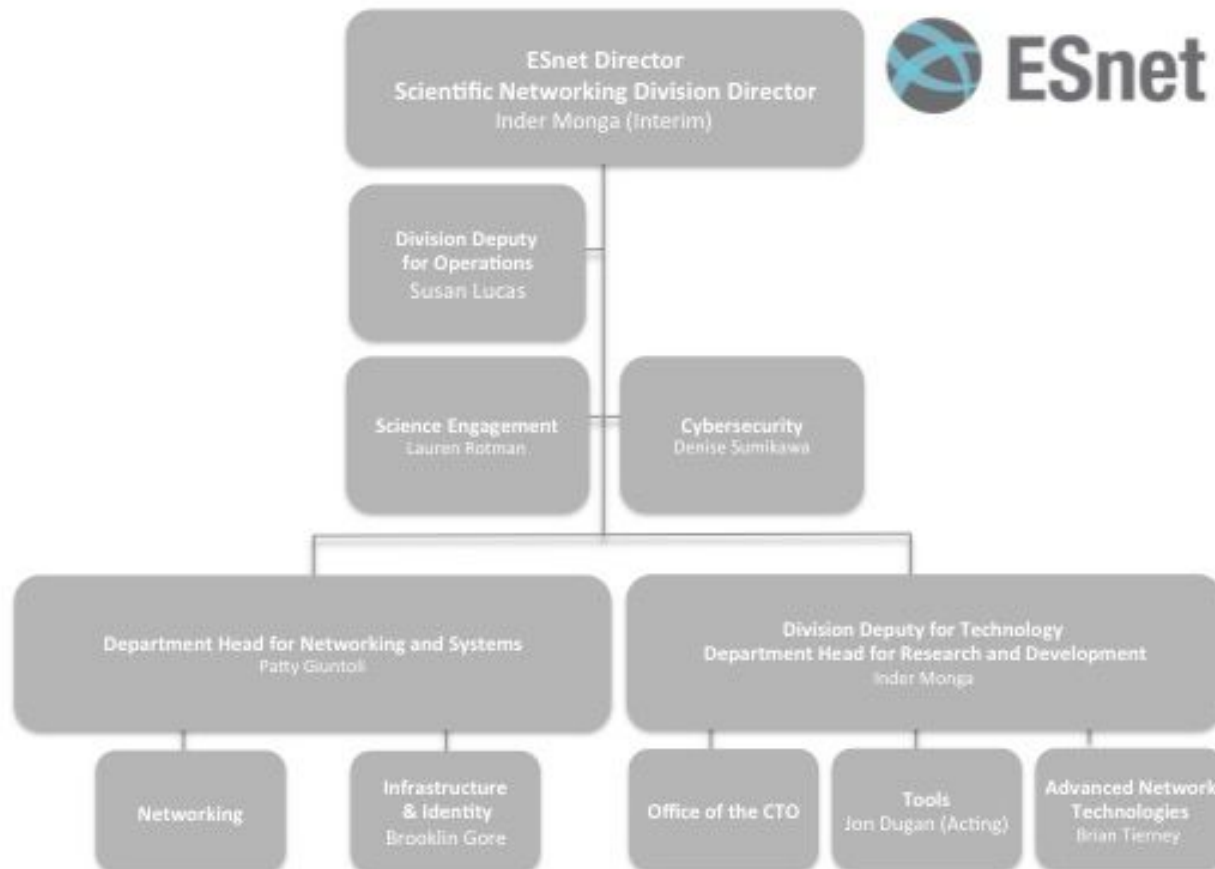


GLIF Map 2011: Global Lambda Integrated Facility Visualization by Robert Patterson, NCSA, University of Illinois at Urbana-Champaign Data Compilation by Maxine D. Brown, University of Illinois at Chicago Texture Retouch by Jeff Carpenter, NCSA Earth Texture, visibleearth.nasa.gov www.glif.is

Densely interconnected to institutions, R&E networks, commercial service providers, cloud providers etc.

150+ peers, 300+ peering points, 1.4 Tbps of peering capacity

Small but effective organization



Punching above our weight.

- Brazil (422)
- UK (180)
- Czech Republic (142)
- Netherlands (138)
- Croatia (112)
- Internet2 (>100)
- Hungary (94)
- Australia (80)
- Norway (78)
- Switzerland (76)
- Greece (68)
- France (66)
- Italy (61)
- Germany (54)
- Ireland (52)
- Slovenia (51)
- Belgium (50)
- Portugal (47)
- ESnet (42)

Caveats: varying service and business models make comparisons difficult, but the large-scale pattern is instructive. (Headcount numbers interpolated from bar graph on page 80 of the most recent *GÉANT Association Compendium*, or described elsewhere in that report.)

A reputation for innovation and excellence.




CENIC

InformationWeek
Government



OPEN NETWORKING
FOUNDATION




“The entire staff conscientiously and continually lead their field.”

[report from recent operational review]

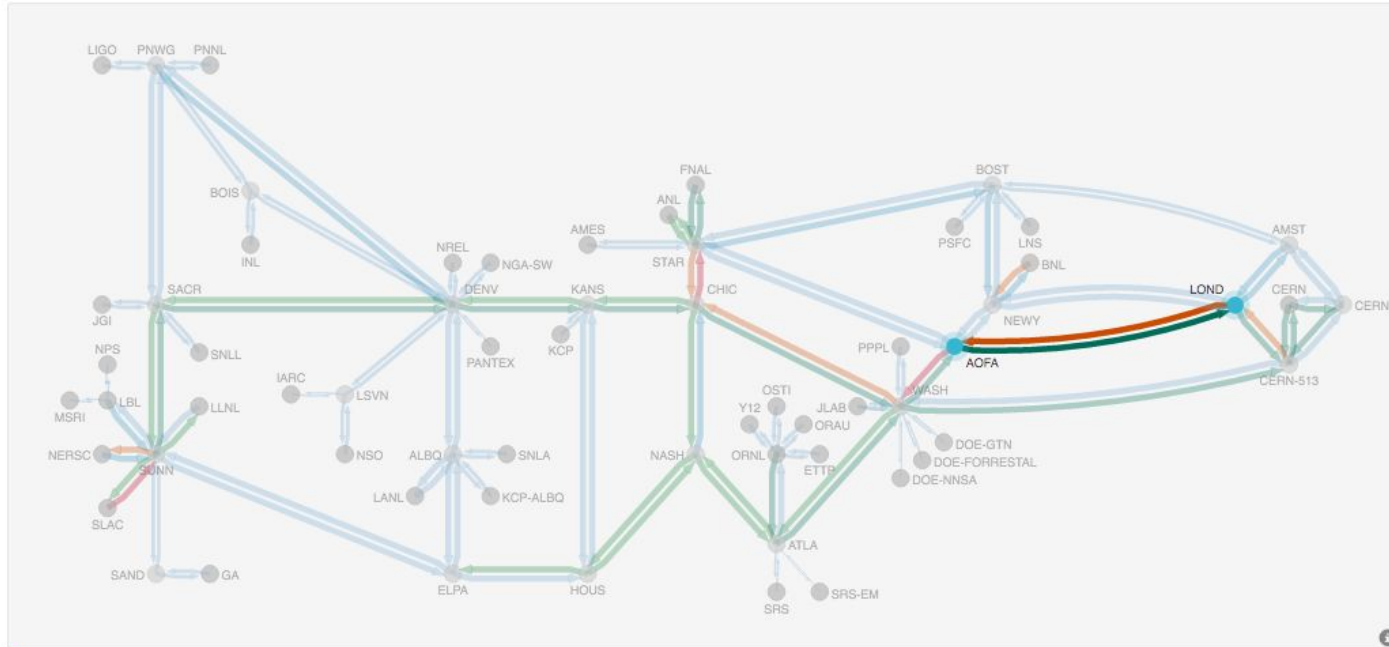
Big Data on ESnet



Overall Traffic doubled, LHCONE up 1500+% in Run 2 [January 2015 – May 2016]




Transatlantic traffic is healthy




AOFA ↔ LOND
100G

Interfaces
aofa-cr5 to_lond-cr5_ip-a



AOFA → LOND (orange) LOND → AOFA (blue)



We even backup partner NREN's trans-atlantic traffic...

MyESnet

Home

Network ▾

Sites ▾

Facilities ▾

Collaborations ▾

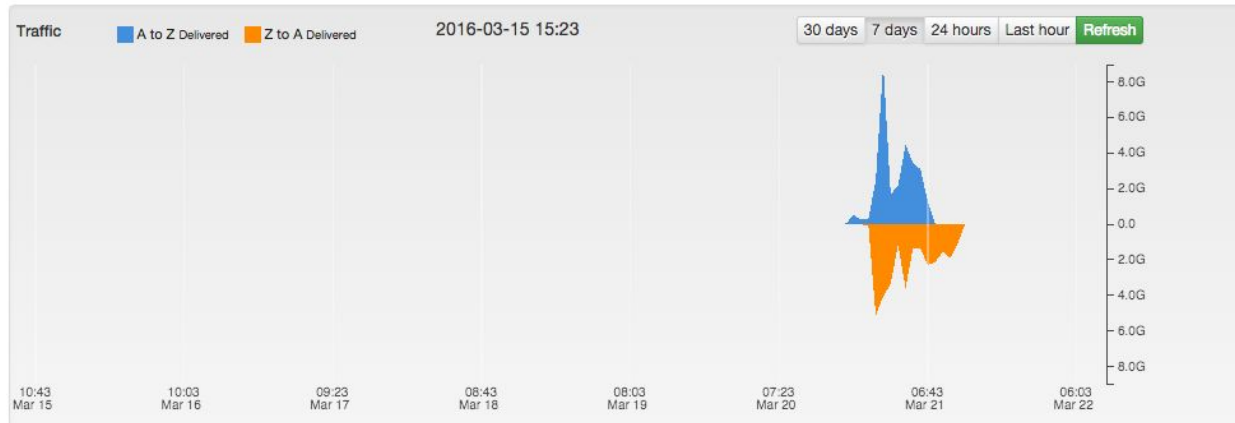

OSCARS / es.net-5891

es.net-5891

ANA-BACKUP: NORDUNET MANLAN to NETHERLIGHT 3020

09-03-2015 To 09-01-2020

OSCARS Circuit



..across multiple exchange points

MyESnet

Home

Network ▾

Sites ▾

Facilities ▾

Collaborations ▾

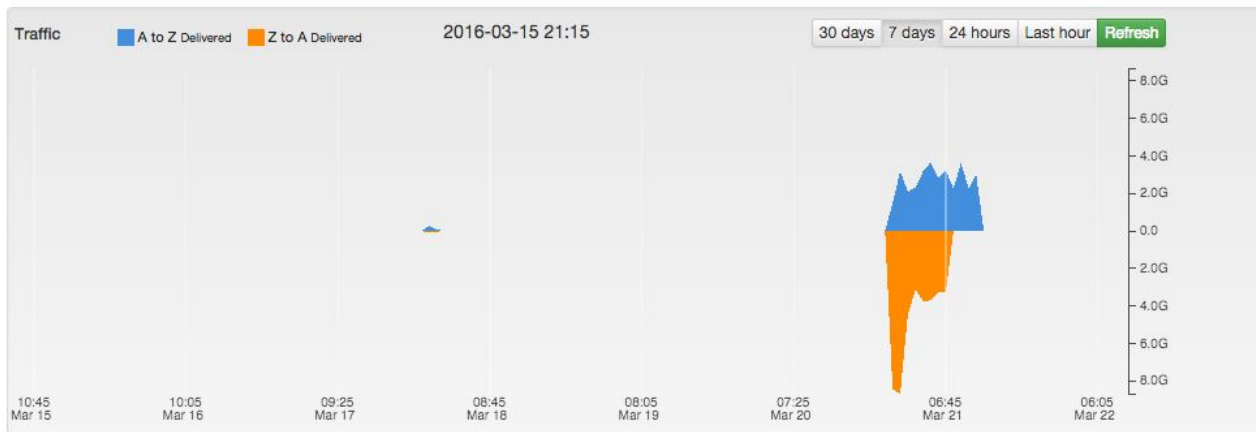

OSCARS / es.net-5893

es.net-5893

ANA-BACKUP: NORDUNET WIX to GEANTOPEN 2031

09-03-2015 To 09-01-2020

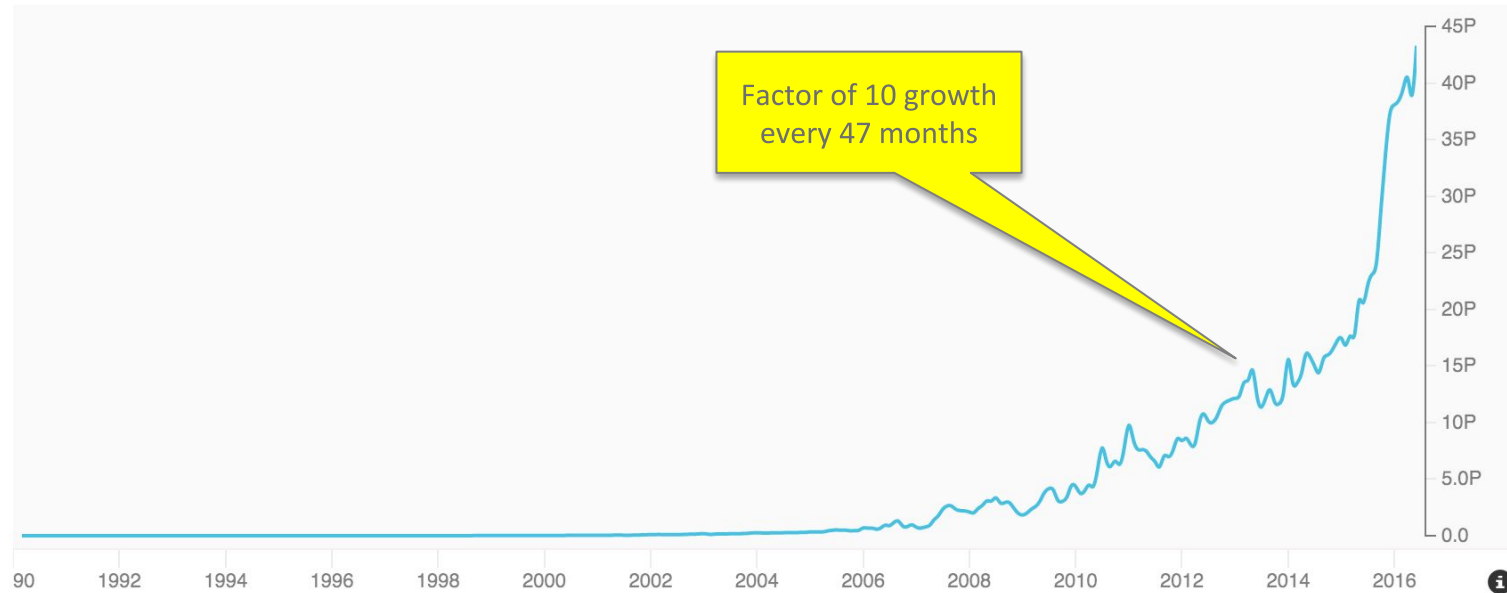
OSCARS Circuit



Continually planning for growth

Traffic Volume

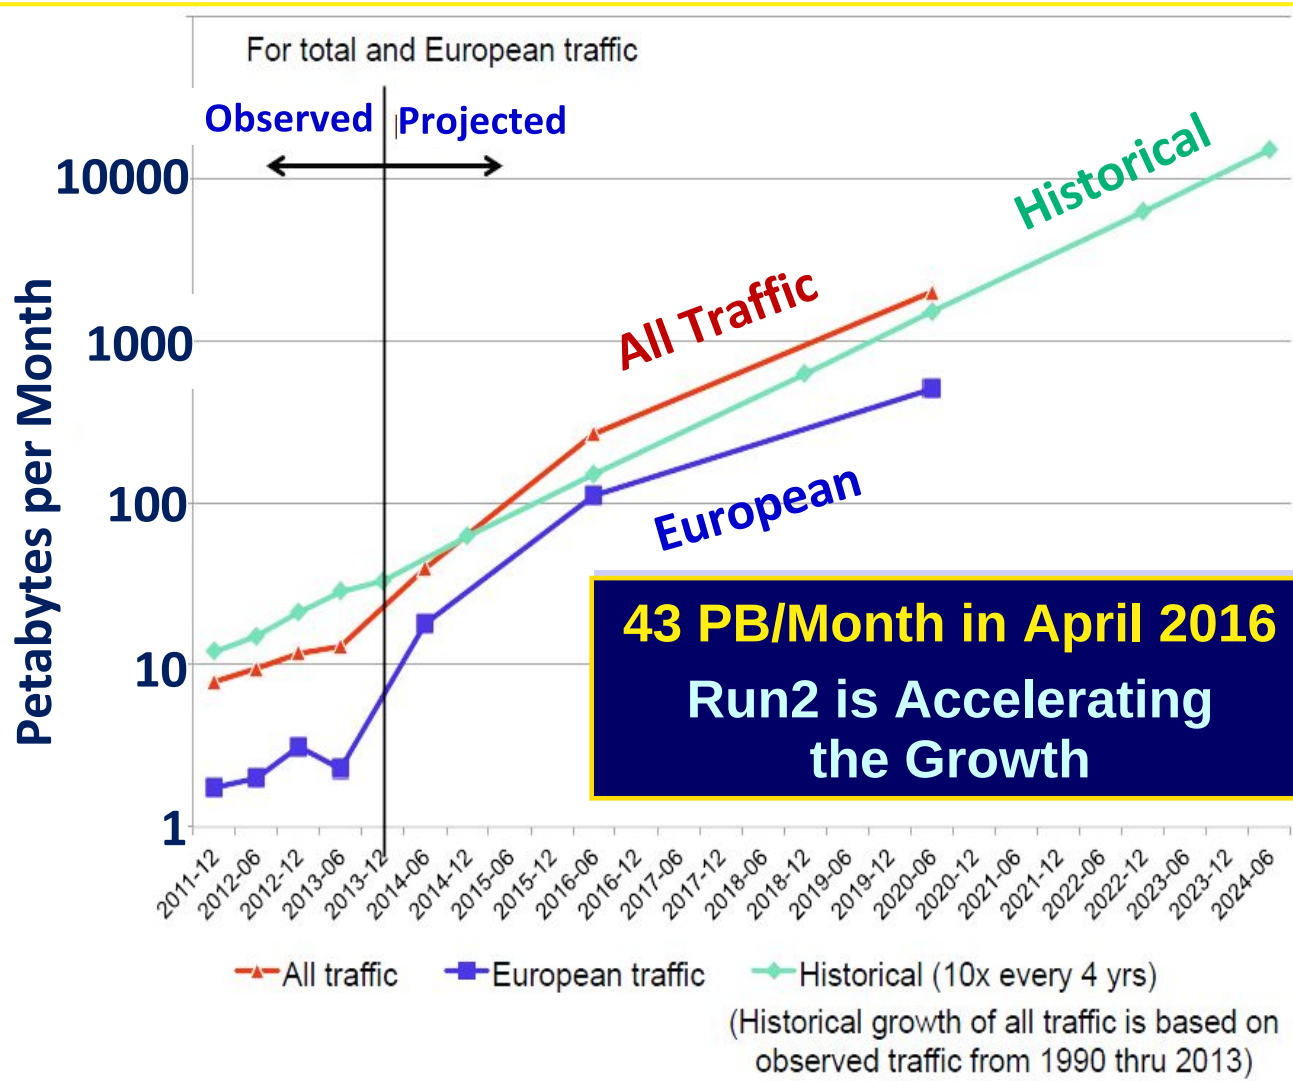
Available at <https://my.es.net/network/traffic-volume>



◀ February 2016 ▶

	Bytes	Percent of Total	One Month Change	One Year Change	
OSCARs	10.46 PB	25.2%	+0.0147%	+148%	Pt-to-pt circuits
LHCONE	11.22 PB	27.0%	+3.90%	+770%	LHCONE (T1-T1/2) traffic
Normal traffic	19.82 PB	47.8%	+9.49%	+77.7%	
Total	41.49 PB		+5.44%	+149%	

ESnet: an Exascale facility in 2020



Projected Traffic Reaches
1 Exabyte Per Month. by ~2020
10 EB/Mo. by ~2024


Slide from
 Harvey Newman



Strategy and Investments



ESnet Research and Software Development Portfolio



Emerging global consensus around Science DMZ architecture.



>120 universities in the US have deployed this ESnet architecture.

NSF has invested >>\$60M to accelerate adoption.

Australian, Canadian universities following suit.

<http://fasterdata.es.net/science-dmz/>

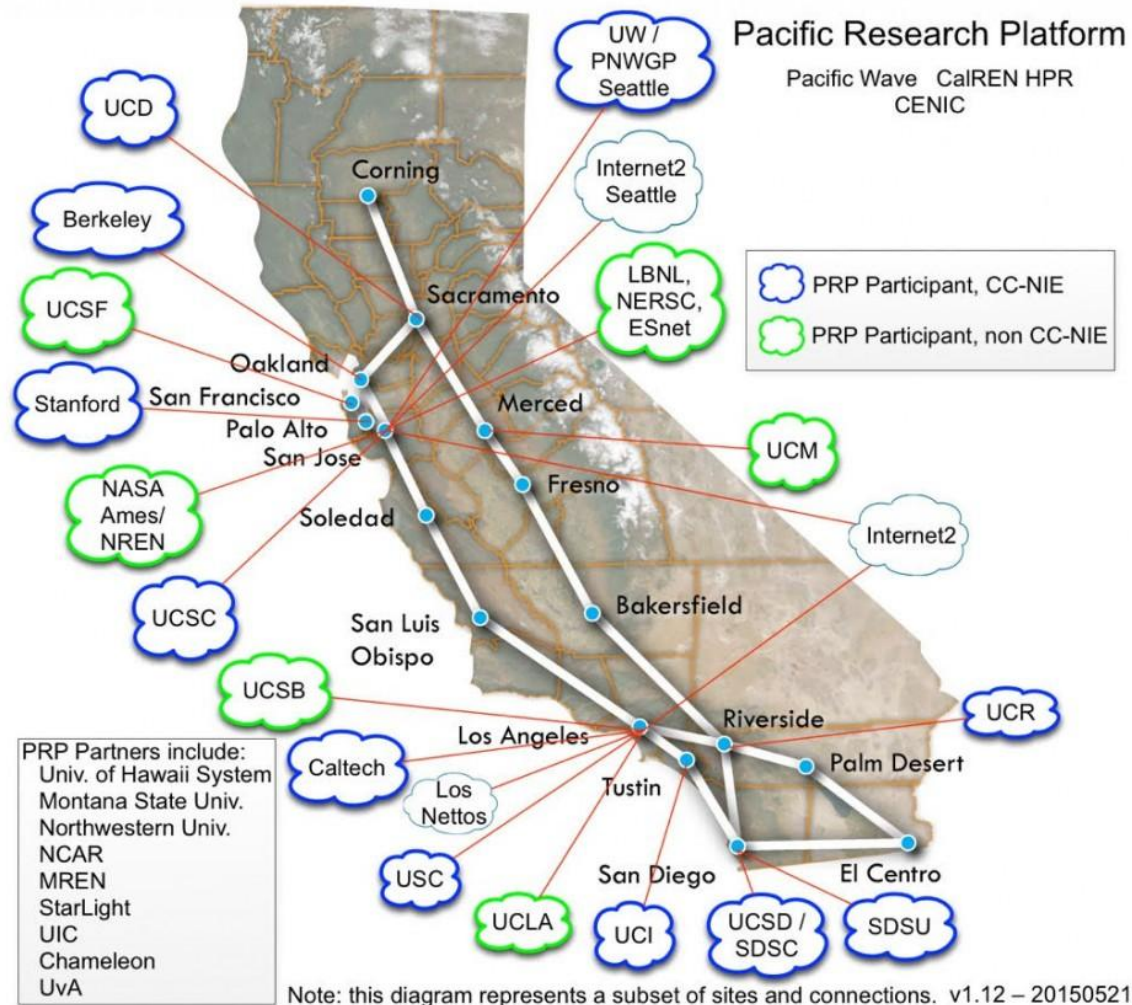
1. Friction-free network path
2. Dedicated data transfer nodes (DTNs)
3. Performance monitoring (perfSONAR)



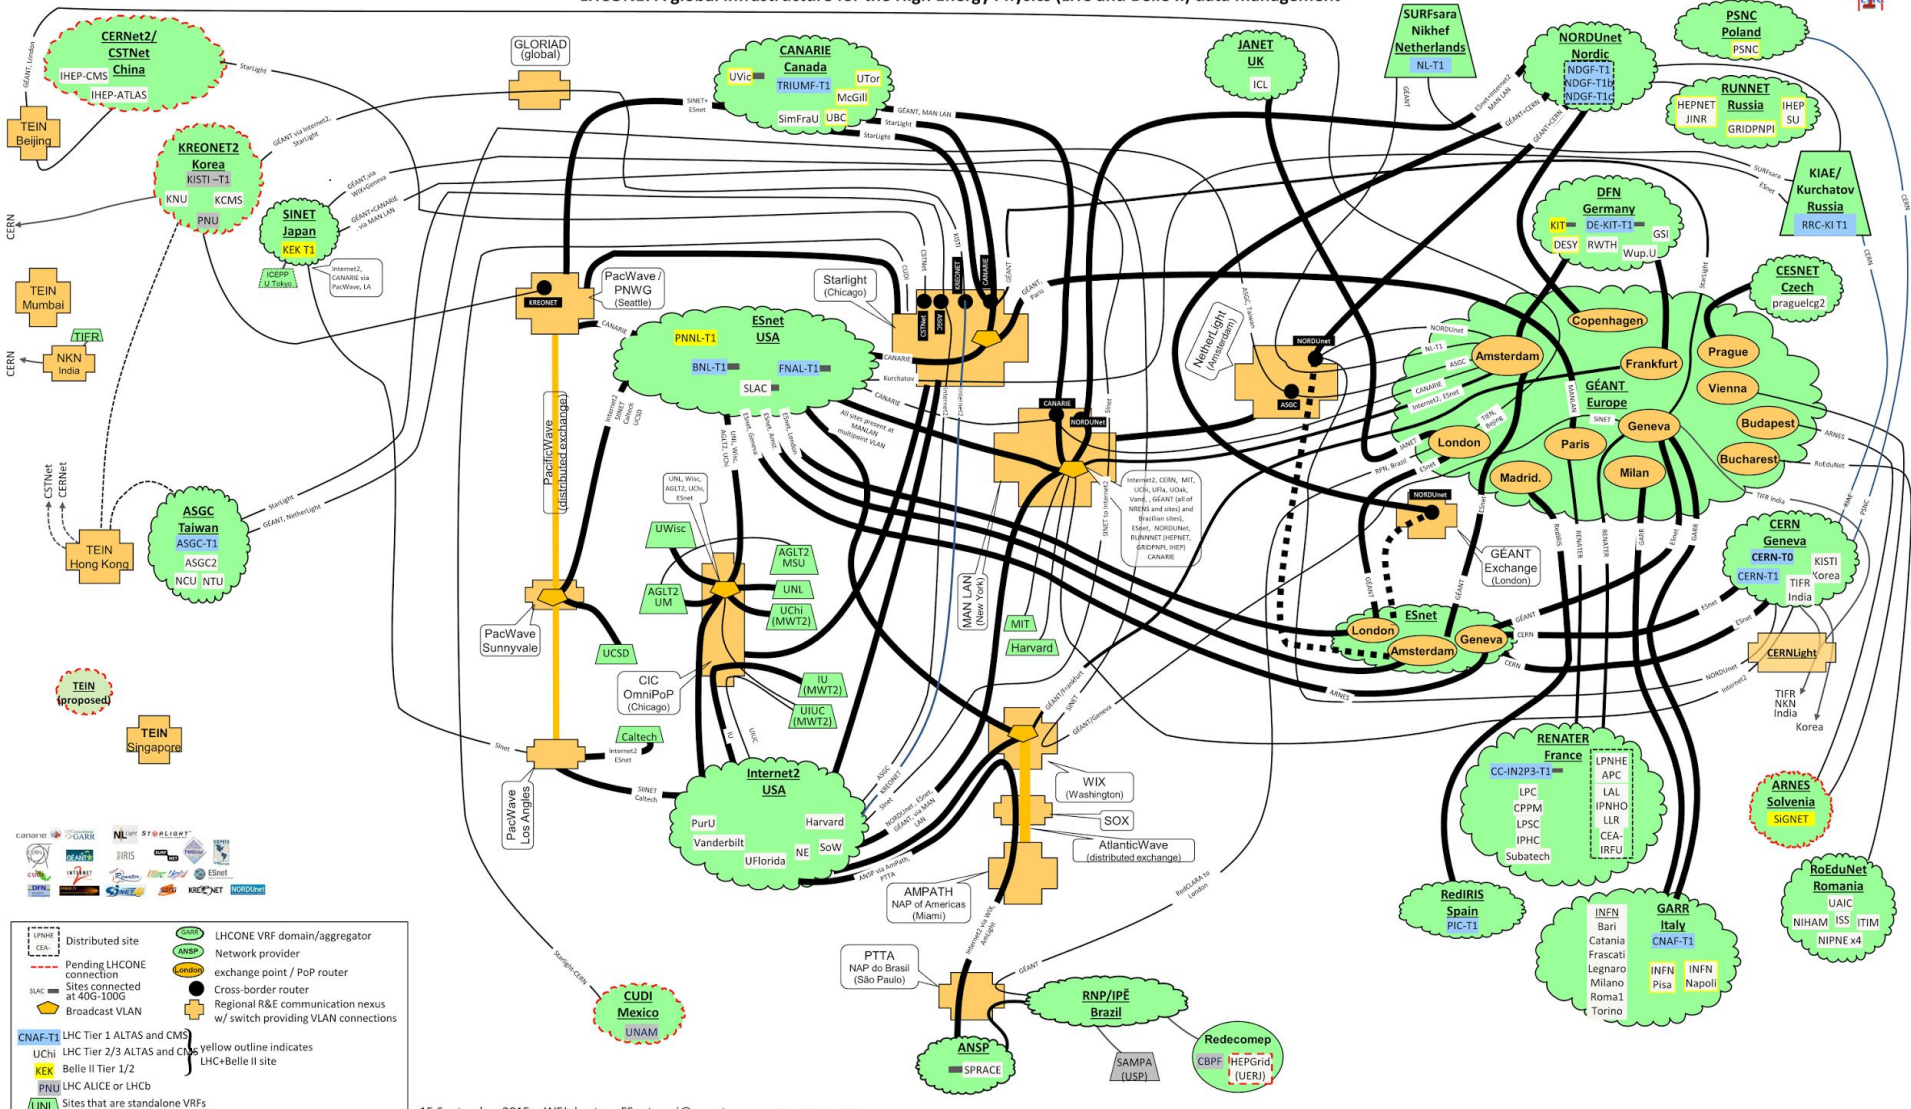
Coming next? Evolution of Science DMZ as a *regional and national* platform.

Pacific Research Platform
initiative, lead by Larry Smarr
(Calit2/UCSD)

- first large-scale effort to coordinate and integrate Science DMZs
- participation by all major California R&E institutions, CENIC, ESnet



LHCONE: A global infrastructure for the High Energy Physics (LHC and Belle II) data management



Isolated apps are less flexible than platforms.

Single-purpose apps:



Programmable platform:



Password authentication

Password:

Welcome to NetShell

```
admin@NetShell> cd /lib/layer2/demo  
changed to: /lib/layer2/demo
```

```
admin@NetShell> vpn create vpn1  
VPN vpn1 is created successfully.
```

```
admin@NetShell> vpn vpn1 addpop amst  
Pop amst is added into VPN vpn1 successfully.
```

```
admin@NetShell> vpn vpn1 addpop cern  
Pop cern is added into VPN vpn1 successfully.
```

```
admin@NetShell> vpn vpn1 addsite amst  
The site amst is added into VPN vpn1 successfully
```

```
admin@NetShell> 
```

An architecture we foresee: ESnet as *platform* for concurrent, domain-specific network apps.

Requires 'network operating system' for science.

- early-stage project (LBNL LDRD)
- multiple challenges in creating flexible, stable app execution environment

We envision apps that will express high-level *intentions*:


- create and manage virtual networks
- enable programmatic resource allocation
- optimize link utilization

Future apps could support:

- NDN for climate; data management for CMS, ATLAS, Belle-II; security overlay for KBase; replication for ESGF; detector / HPC coupling for light sources
- workflows we haven't imagined



ESnet's 100G SDN Testbed – significant footprint growth and dedicated bandwidth. **Unique within community**



SENSE: SDN for End-to-end Networked Science for Extreme-Scale Science

Inder Monga [Lead-PI] (ESnet), Phil Demar (FNAL), Harvey Newman (Caltech),
Linda Winkler (ANL), Tom Lehman (UMD/MAX), Damain Hazen (NERSC)
Mar 2016 – Feb 2019

Goal

- 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

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

SENSE SDN Control Plane Architecture for End-to-End Orchestration

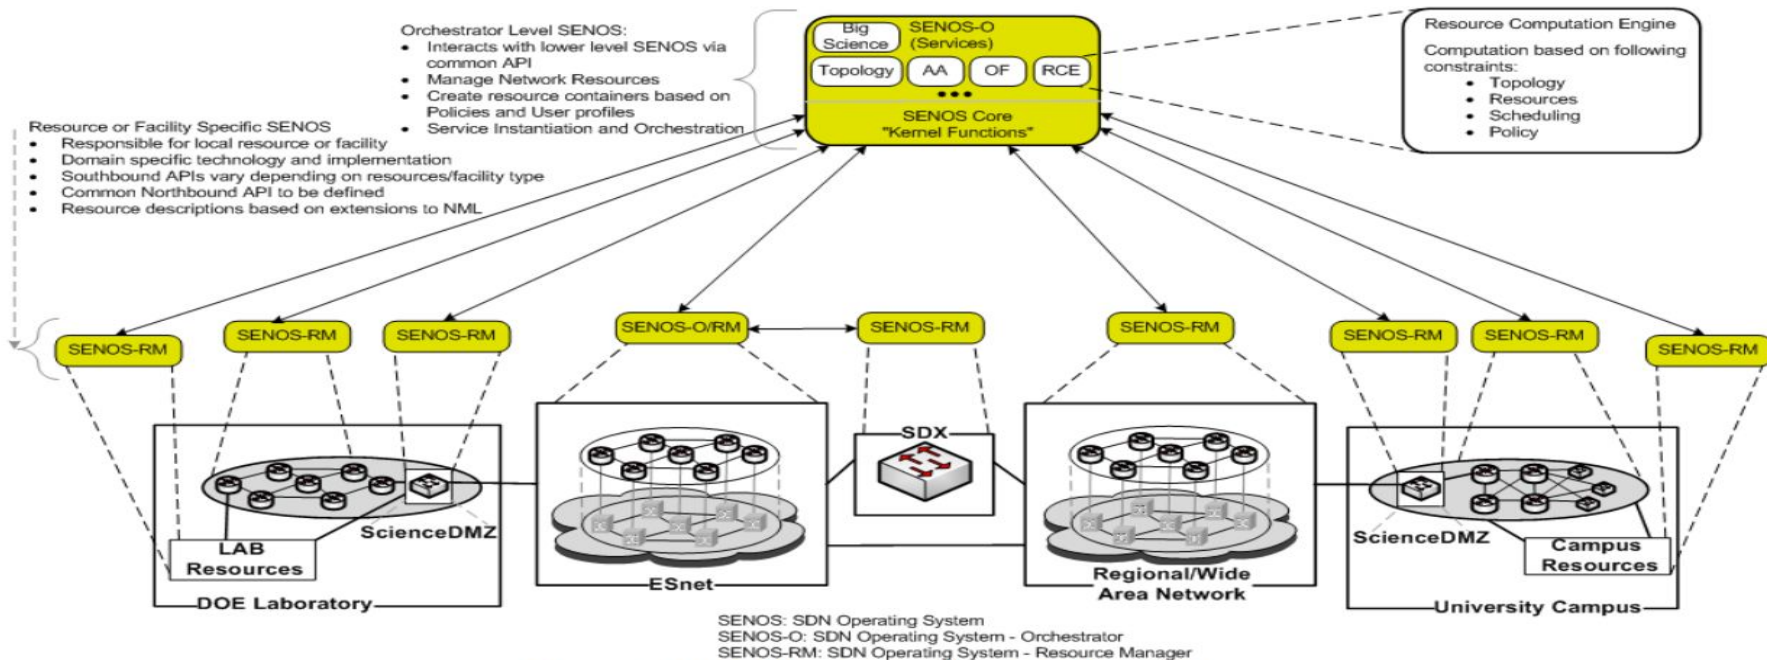
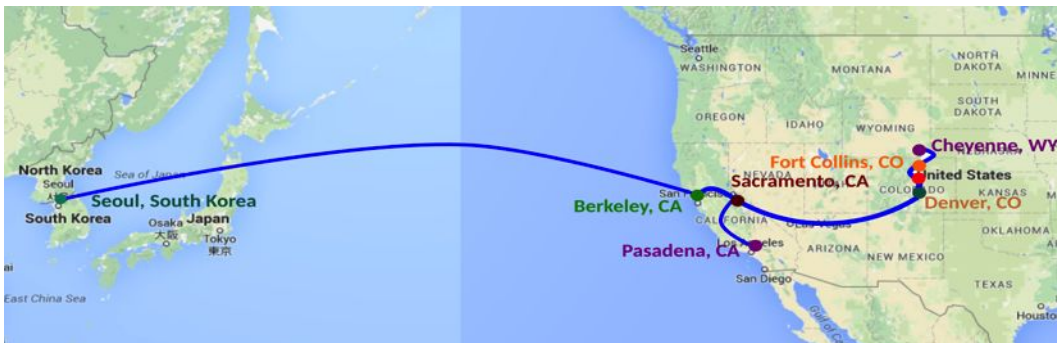


Figure 1. SENOS End-to-End Orchestration


Named-Data Networking collaboration with NSF funded Colorado State CC* project continues to engage DOE science initiatives

- An **NSF-funded Future Internet Architecture** in which clients express intent for data objects to the network
- Collaborating with Colorado State University, PI Christos Papadopoulos, to investigate applicability for climate data sets.
 - Mentoring the project and providing focus: enabling high-speed data transfer, data management and caching strategies for large scientific data sets.
- CSU collaborators and ESnet invited to presented NDN results at three conferences, CHEP 2015, NDN Comm 2015 and ACM Conference on Information Centric Networks
 - ESnet invited talks to give talks at ACM and NDNComm
- ESnet is hosting a large part of an NDN testbed and dedicated 10G to NDN research.



Superfacility: our global science complex can become more than the sum of its parts.

Real-time analysis of 'slot-die' technique for printing organic photovoltaics, using ALS + NERSC (SPOT Suite for reduction, remeshing, analysis) + OLCF (HipGISAXS running on Titan w/ 8000 GPUs).



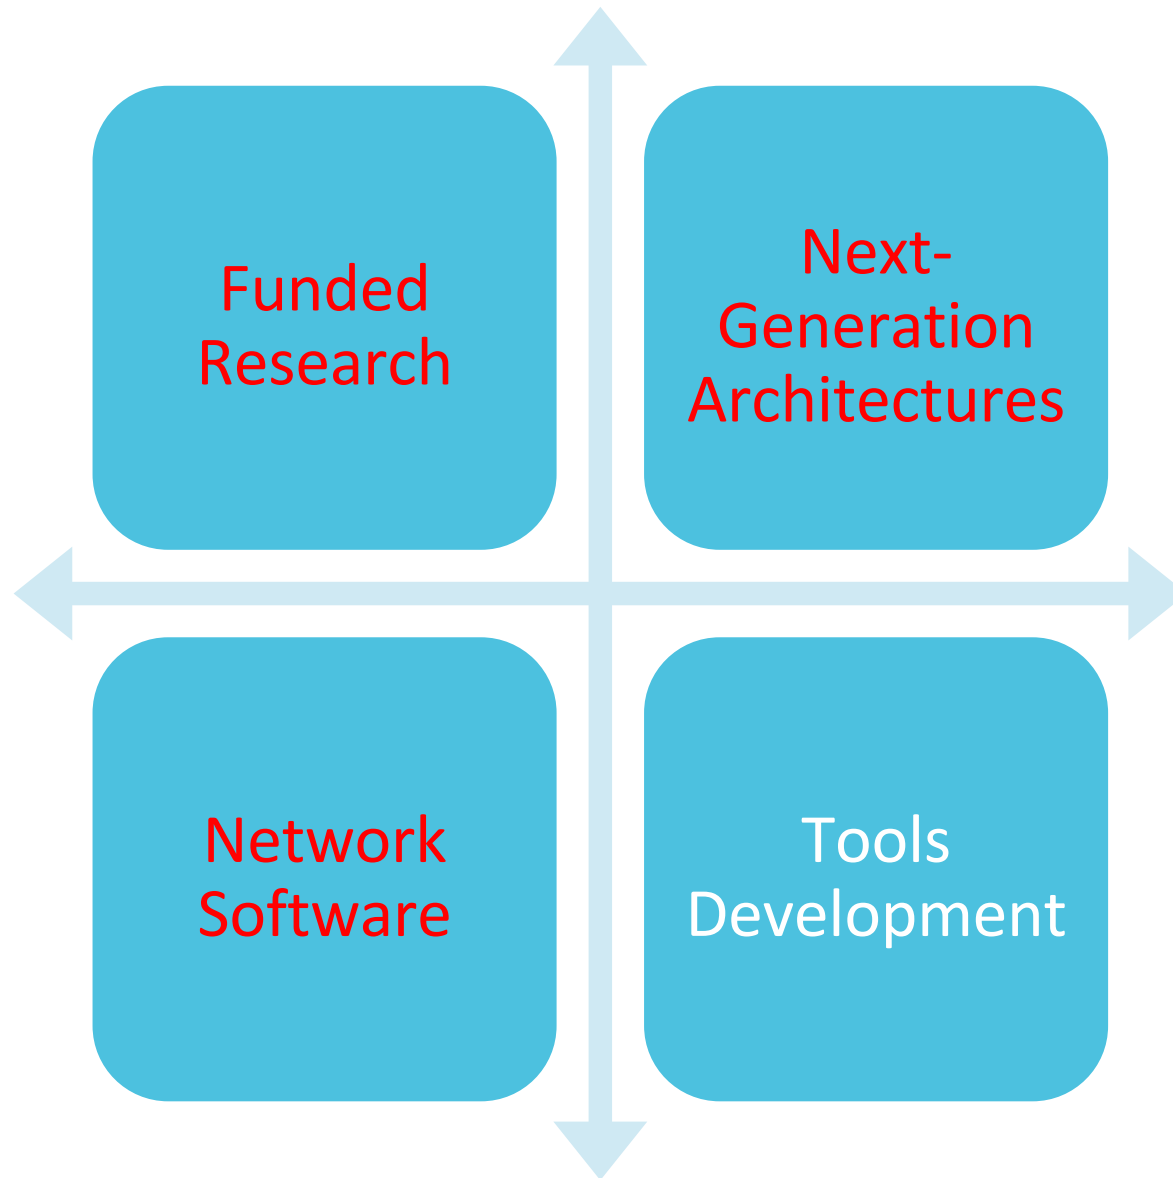
<http://www.es.net/news-and-publications/esnet-news/2015/esnet-paves-way-for-hpc-superfacility-real-time-beamline-experiments/>

Results presented at March 2015 meeting of American Physical Society by Alex Hexemer.

Additional DOE contributions: **GLOBUS** (ANL), **CAMERA** (Berkeley Lab)

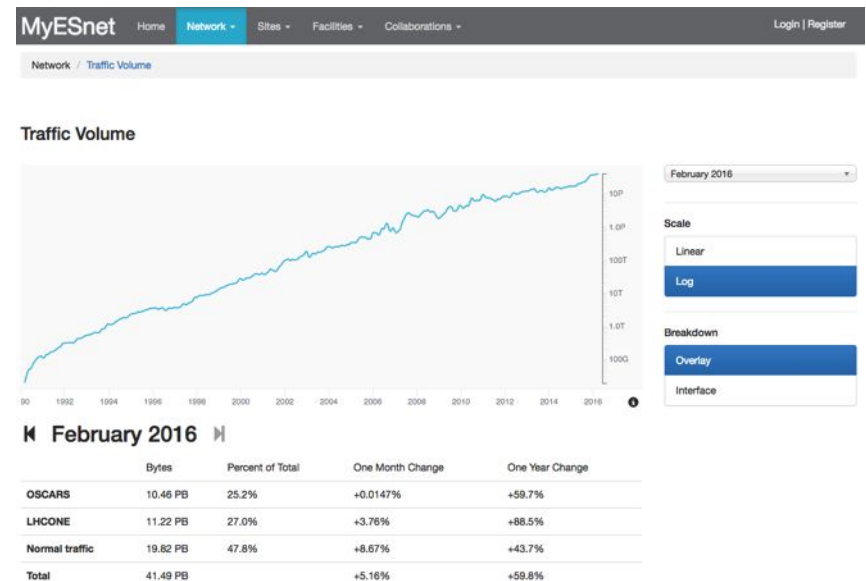


ESnet Research and Software Development Portfolio



New Portal Feature: Monthly Traffic Volume

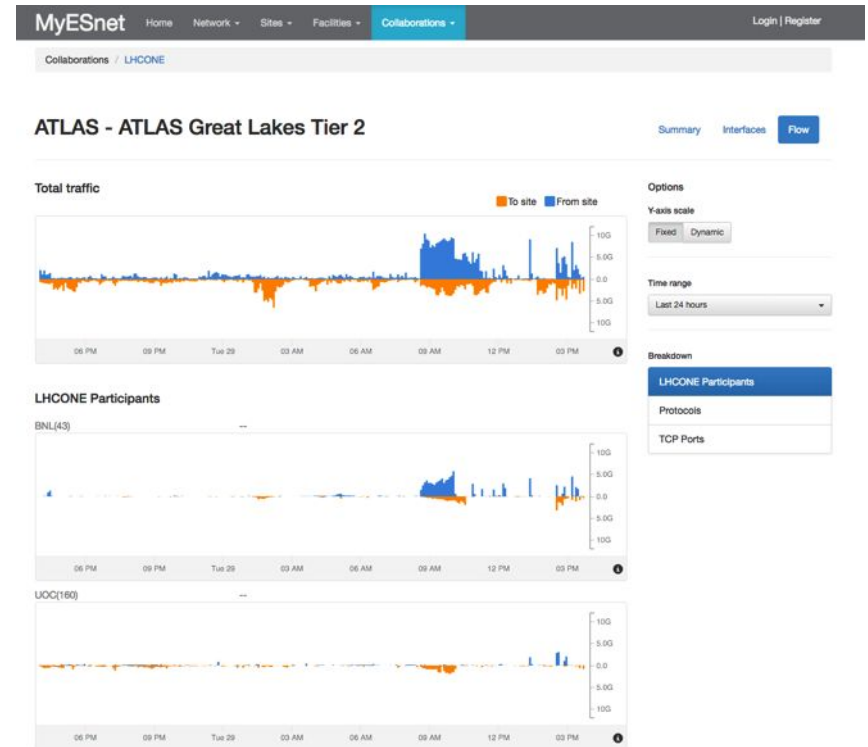
- Portal version of monthly stats
- Same data, new visualization
- Provides summary view
 - Percent change vs. last month
 - Percent change vs. same month last year
- Explore contribution of OSCARS or LHCONE
- Explore contribution of interfaces that match a pattern



<https://my.es.net/network/traffic-volume/>

New Portal Feature: LHCONE

- Provide a place for all ESnet connected LHCONE participants to see traffic
- Flow data provides insight on how traffic is flowing over the VPRN
- Showcases traffic from Universities onto ESnet



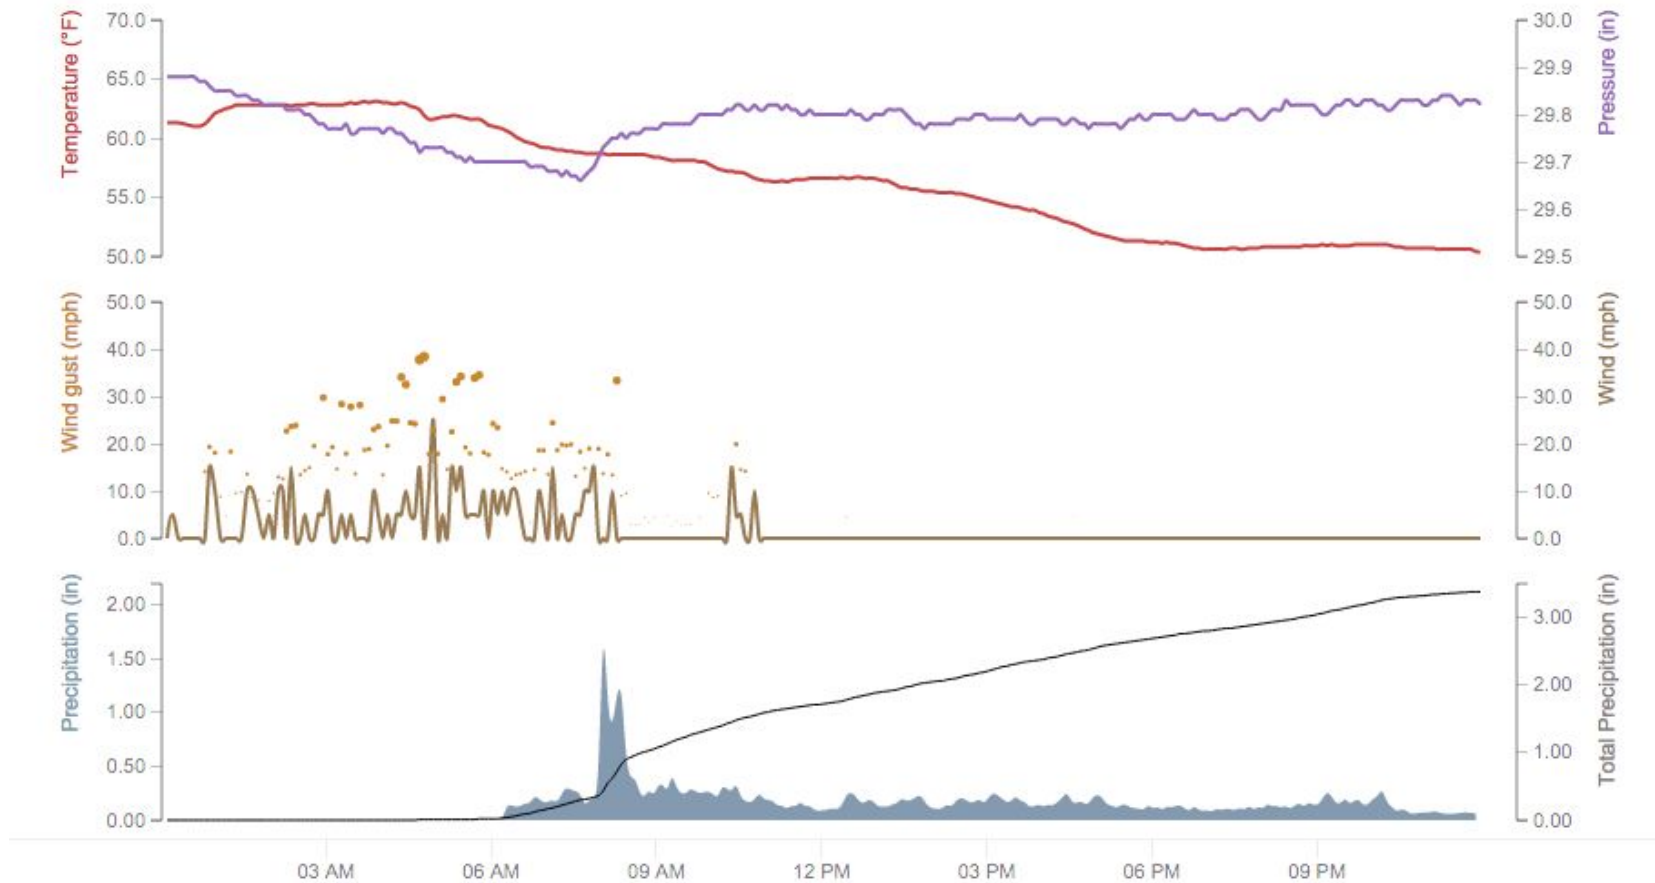
<https://my.es.net/collaborations/lhcone/>




Time Series Charts



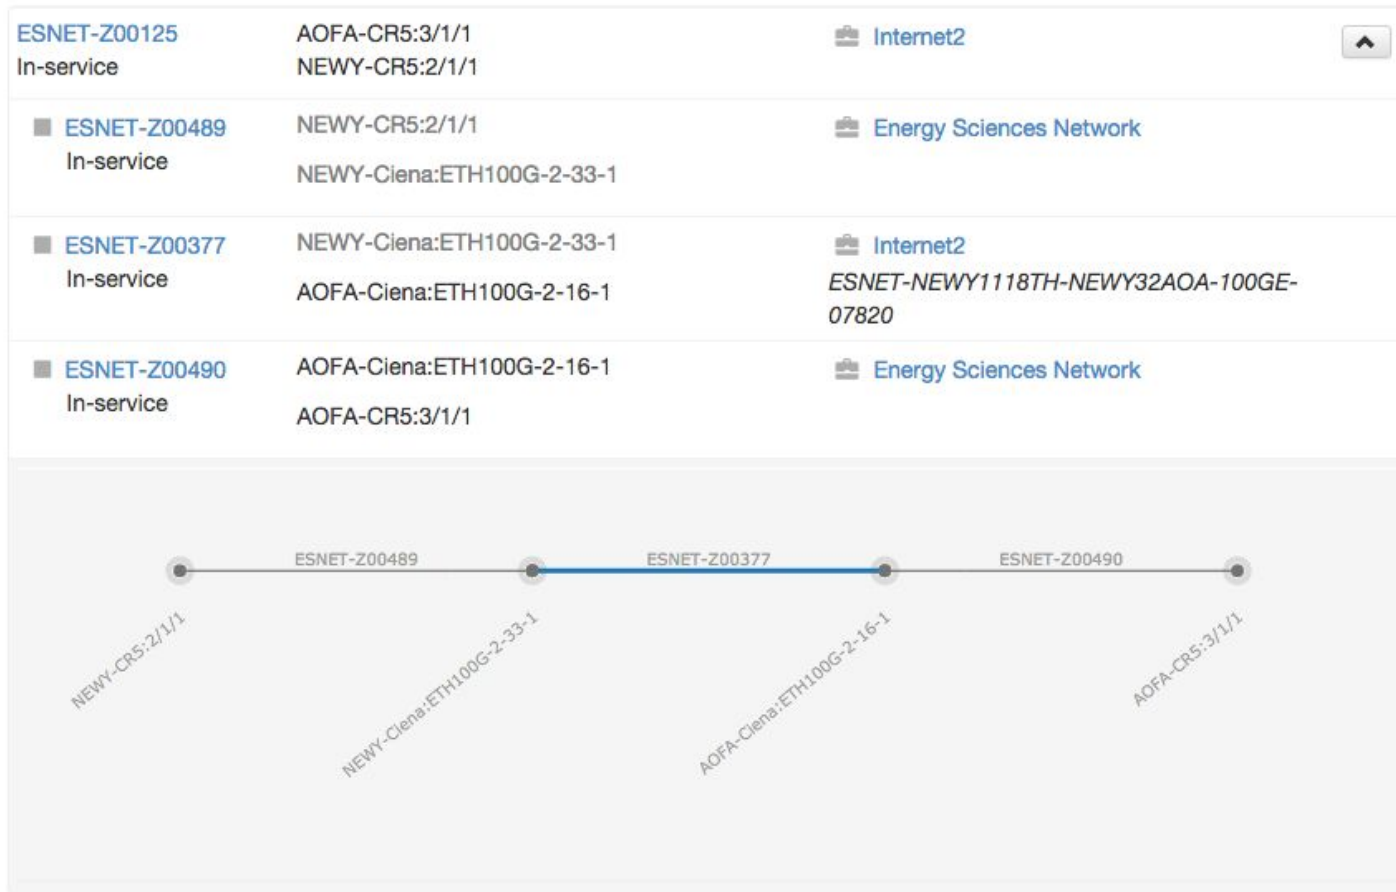
Time Series Charts



Network Diagrams




Network Diagrams



Pond: Time Series Operations

- We have lots of time series data: SNMP, netflow, log messages, etc
- Need a common abstraction and wire format for this data
- Provides common operations

Latest: Thu Jan 01 2015 05:29:00 GMT-0800 (PST)



Animation showing how Pond generates summary rollups

Open Source

Good news everyone, it's all open source!

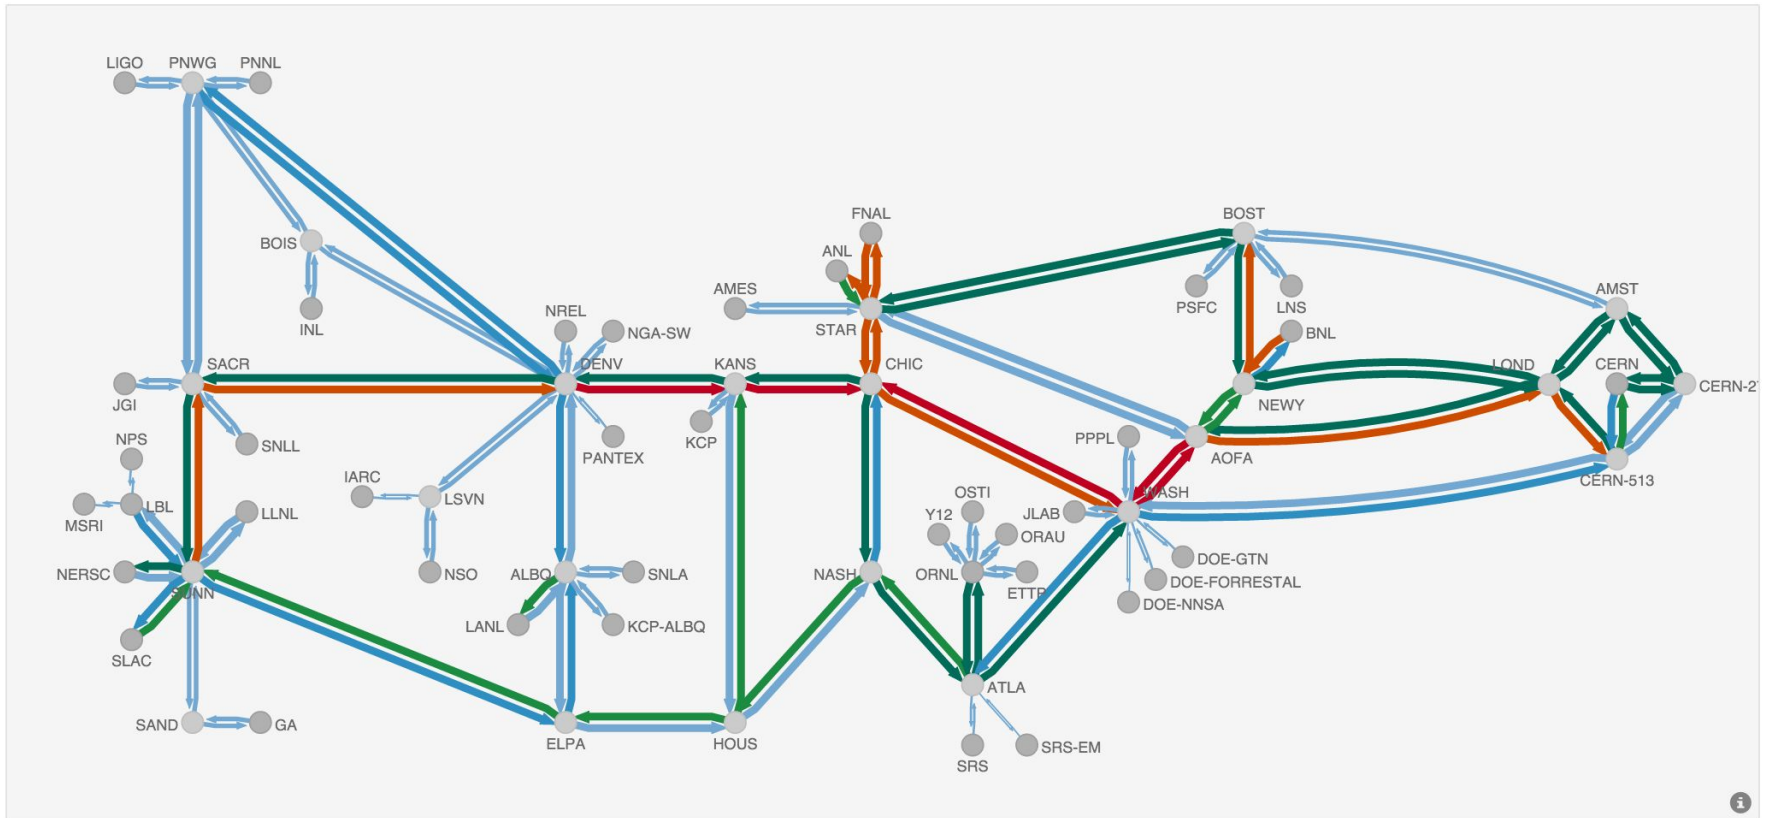
- <http://software.es.net/pond/>
- <http://software.es.net/react-timeseries-charts/>
- <http://software.es.net/react-network-diagrams/>

Motivation

- Develop a common toolkit for visualizing networks
- Lower barrier to entry
- Allow people to use what we've built so they don't have to build their own
- Encourage others to share their visualizations
- Give back to the community



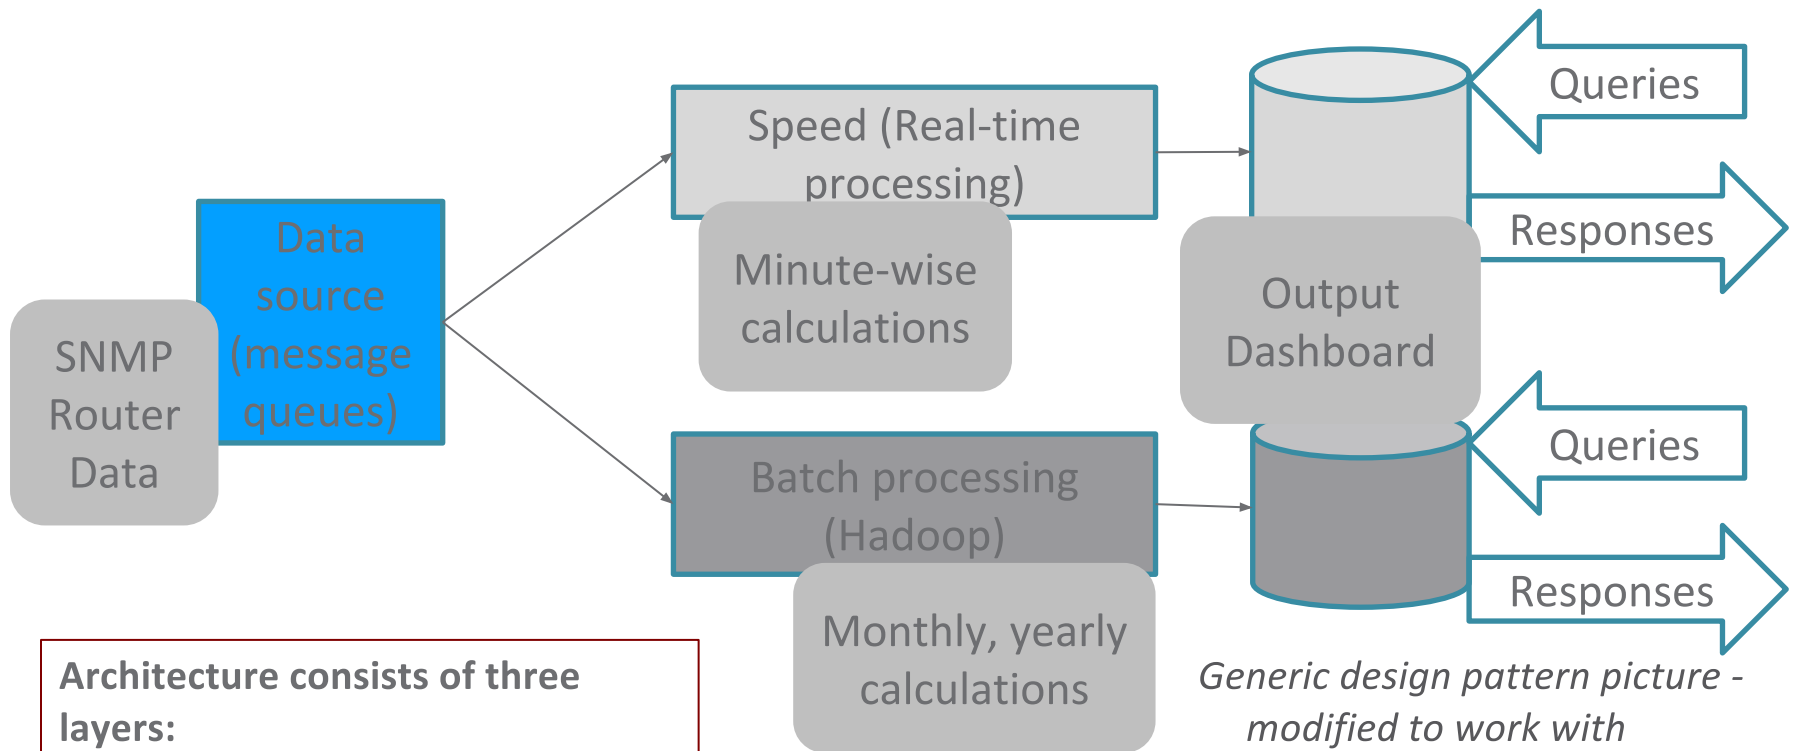
Visualization of data is good, but not enough



Network Analytics

- Data being generated by the network every few seconds but not analysed or available for real-time analysis
 - The ability to ask questions of historical network data, and get answers
 - The answers updated with new data in near real-time
 - SNMP data, Flow data, Topology data, etc..
- Smart Cities, IoT, Smart Grid – have common problems

Leverage cloud computing tools to put together a *network analytics* pipeline



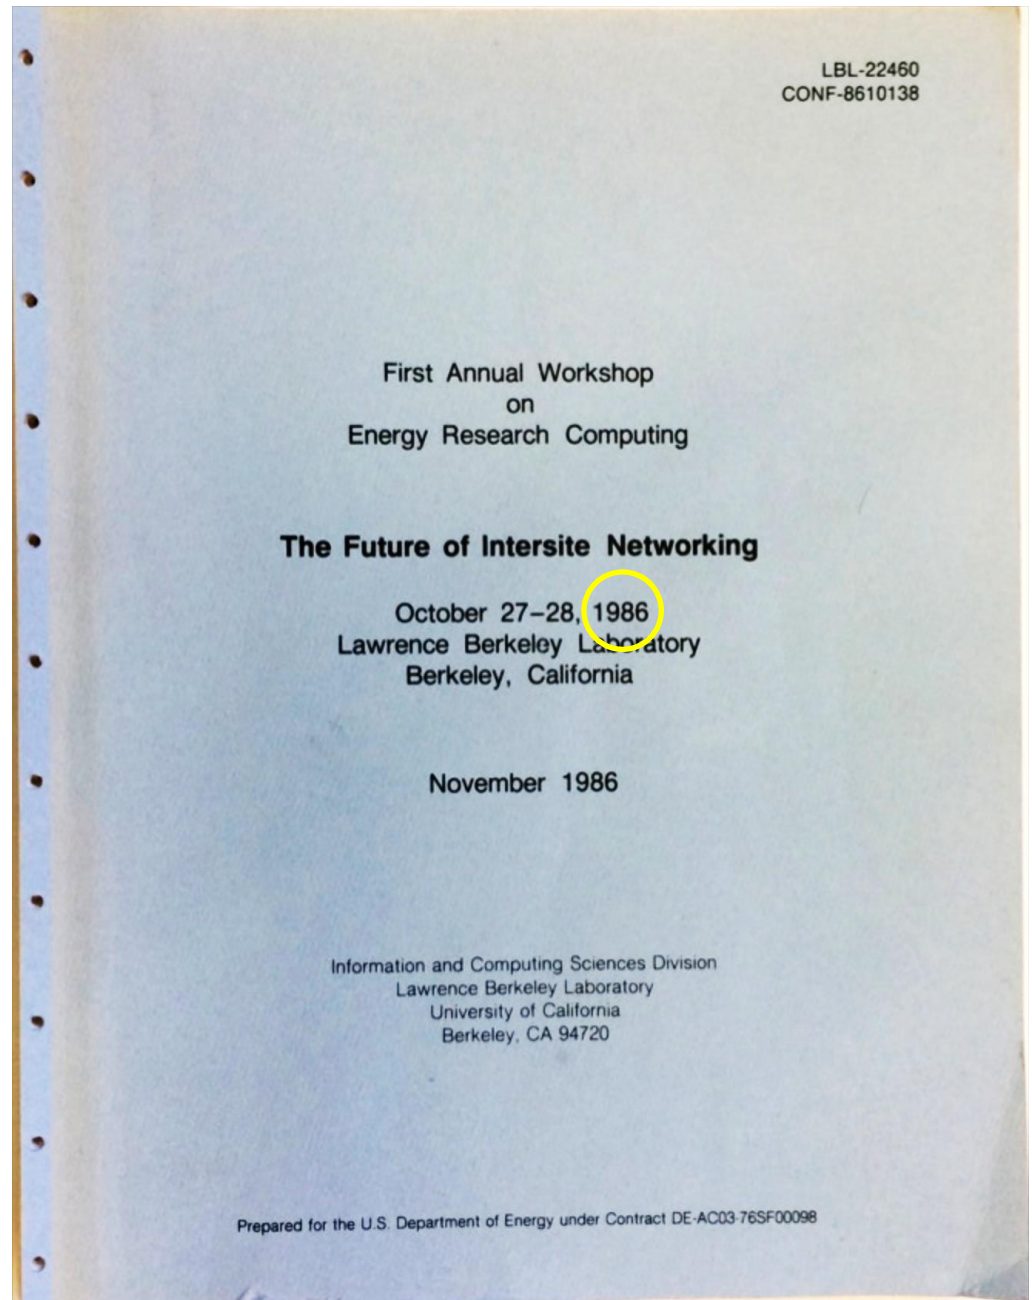
Architecture consists of three layers:

- Batch processing
- Speed or real time
- Layer to respond to queries

Generic design pattern picture - modified to work with specific cloud computing technologies

In 1986, physics community made two 'modest requests':

1. "What we can do on **LANs** today is indicative of what we wish to be able to do on **wide area networks**."
2. "Just as we expect a **computer** to perform as if we are the **only user**, we expect the **network** to give that same appearance."



In conclusion:

- Global science networks are not ISPs – rather, extensions of your discovery instruments.
- HEP's 'modest requests' from 1986 are in the process of being fulfilled.
 - 'discovery unconstrained by geography' is explicit vision of ESnet
 - and virtualization is a global initiative
- Design patterns, architectures, workflows and challenges from HEP science are now crossing over to domains.

In conclusion:

- Global science networks are not ISPs – rather, extensions of your discovery instruments.
- HEP's 'modest requests' from 1986 are in the process of being fulfilled.
 - 'discovery unconstrained by geography' is explicit vision of ESnet
 - and virtualization is a global initiative
- Design patterns, architectures, workflows and challenges from HEP science are now crossing over to domains.
- HEP has been pushing the boundaries of networking for three decades.
 - please keep it up!
- It's time to make the next modest requests
 - we are taking orders for 2018.

In conclusion, a reminder about our vision:



Scientific progress will be **completely unconstrained** by the physical location of instruments, people, computational resources, or data.

Big Science Data in Motion = Elephant Flow!

IoT watching LOL Cats = Mice flow!



Elephant Data vs. Mice Data Behavior

ESnet Traffic (last 24 hours)

