



ESnet

ENERGY SCIENCES NETWORK

Supporting Scientific Research with the Energy Sciences Network

David Mitchell - Network Engineer

mitchell@es.net

Inder Monga - Acting Director

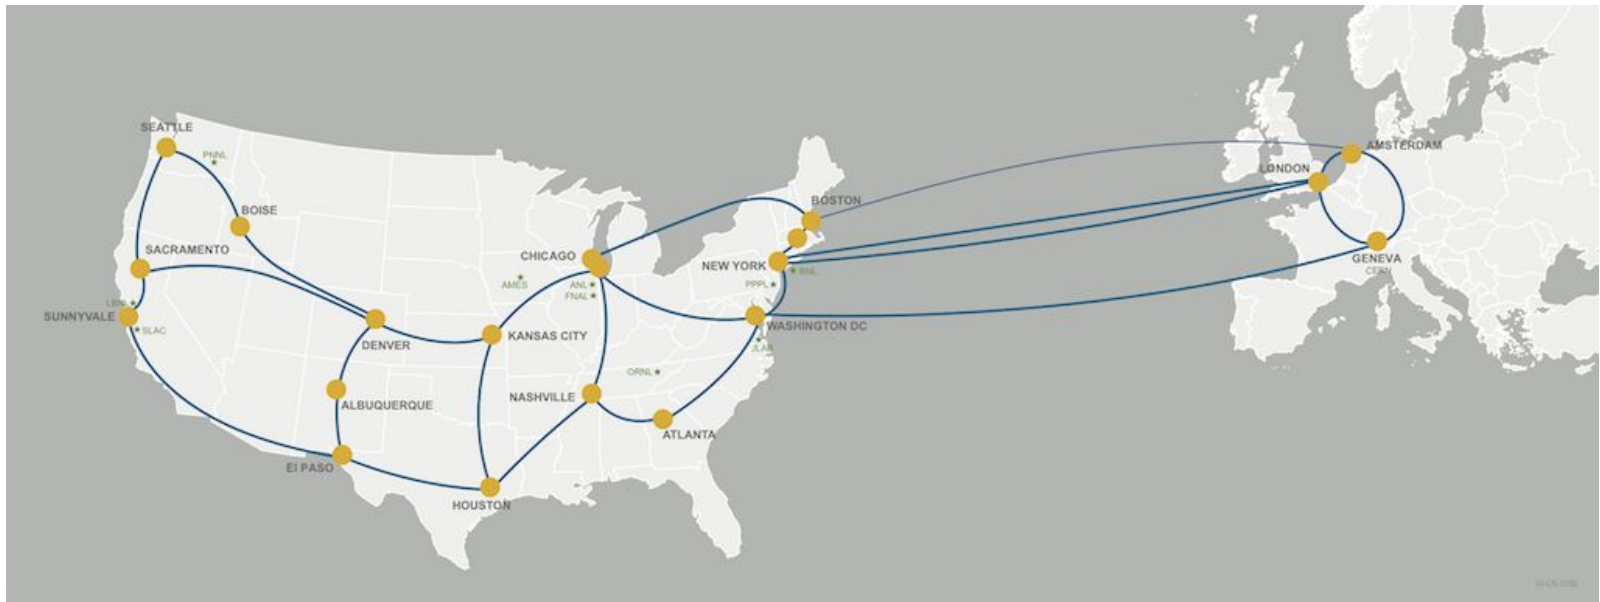
imonga@es.net



U.S. DEPARTMENT OF
ENERGY
Office of Science



ESnet Trans-Atlantic Services



ESnet
ENERGY SCIENCES NETWORK

★ Department of Energy Office of Science National Labs

- Ames Ames Laboratory (Ames, IA)
- ANL Argonne National Laboratory (Argonne, IL)
- BNL Brookhaven National Laboratory (Upton, NY)
- FNAL Fermi National Accelerator Laboratory (Batavia, IL)
- JLAB Thomas Jefferson National Accelerator Facility (Newport News, VA)

- LBLN Lawrence Berkeley National Laboratory (Berkeley, CA)
- ORNL Oak Ridge National Laboratory (Oak Ridge, TN)
- PNNL Pacific Northwest National Laboratory (Richland, WA)
- PPPL Princeton Plasma Physics Laboratory (Princeton, NJ)
- SLAC SLAC National Accelerator Laboratory (Menlo Park, CA)

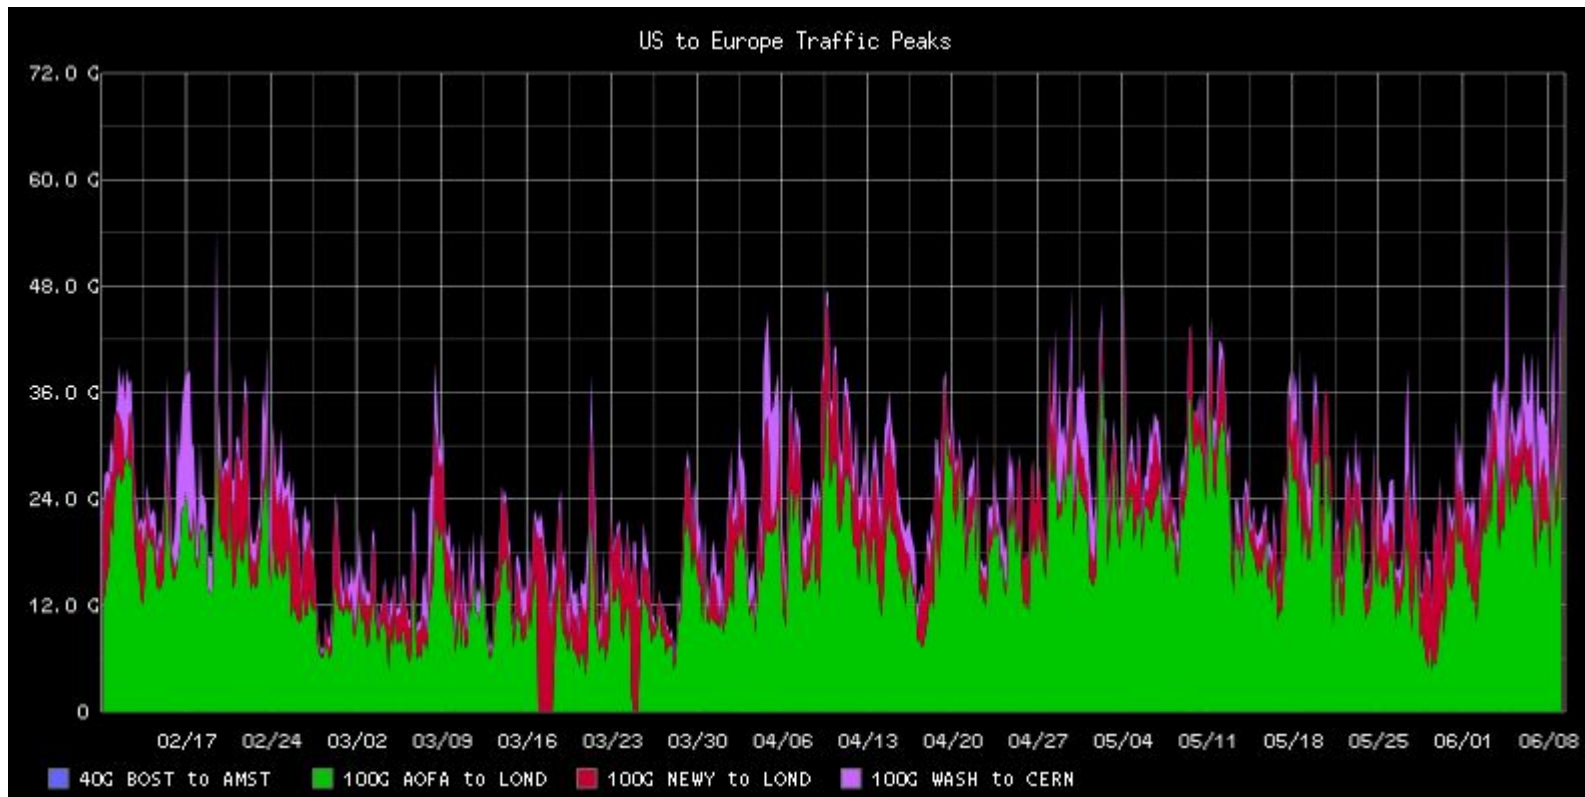


ESnet Trans-Atlantic Services

- Primary connectivity provided by three 100Gbps circuits and one 40Gbps trans-Atlantic circuits
- Three additional intra-Europe 100Gbps circuits and hubs at CERN, London and Amsterdam
- Supports the LHCOPN connections to US LHC Tier 1 sites
 - ATLAS Tier 1 at Brookhaven National Lab (BNL)
 - CMS Tier 1 at Fermi National Accelerator Laboratory (FNAL)
- Supports the LHCONE overlay network for connectivity to U.S. based Government and University LHC collaborators

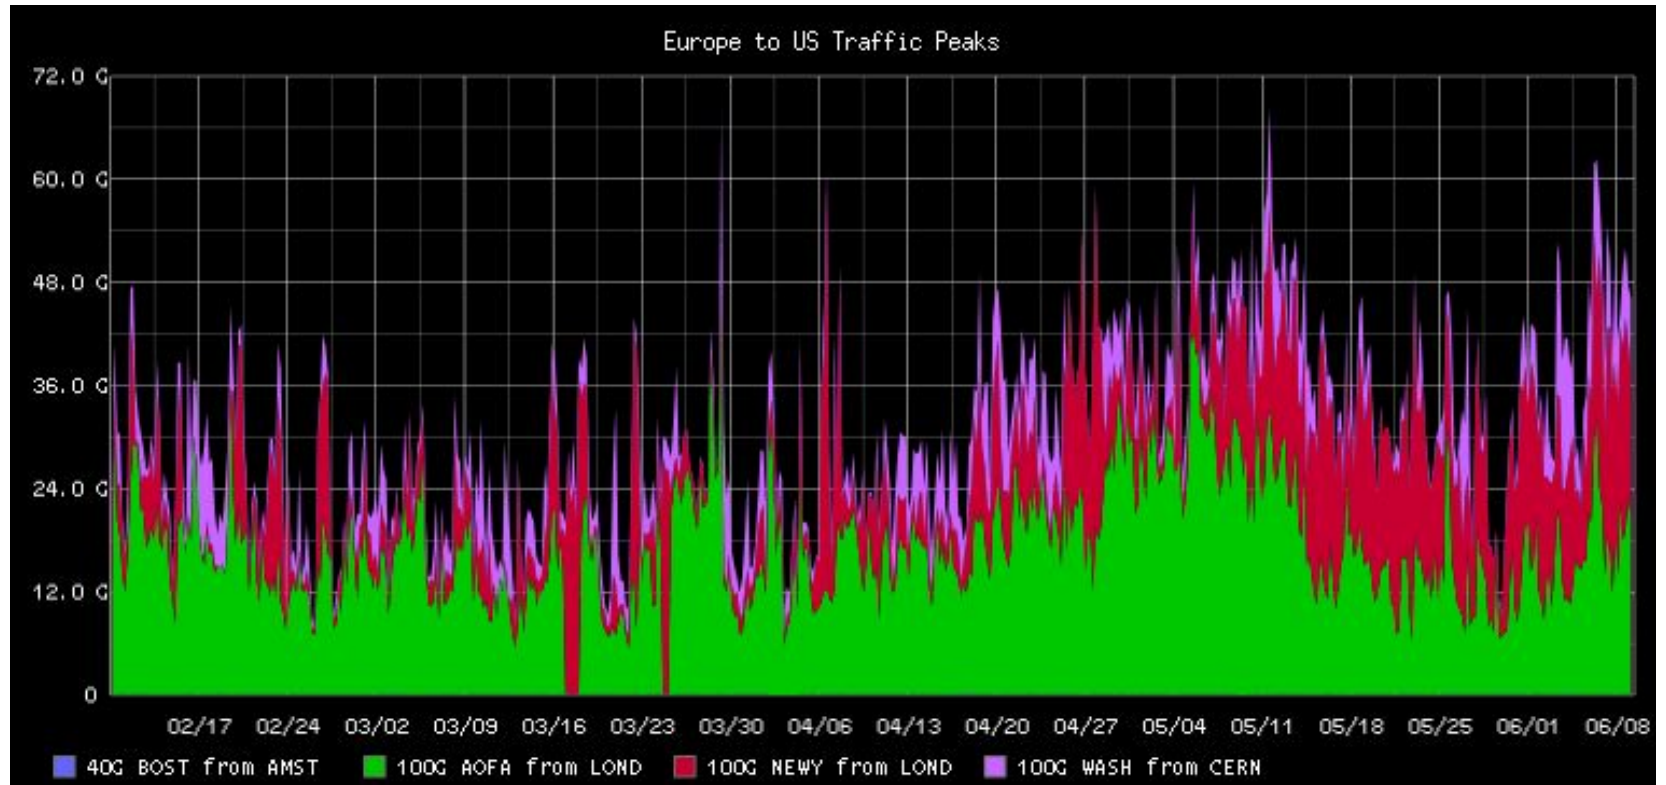
Trans-Atlantic Traffic

US to Europe traffic:



Trans-Atlantic Traffic

Europe to US traffic:



ESNet Cloud Services Pilot

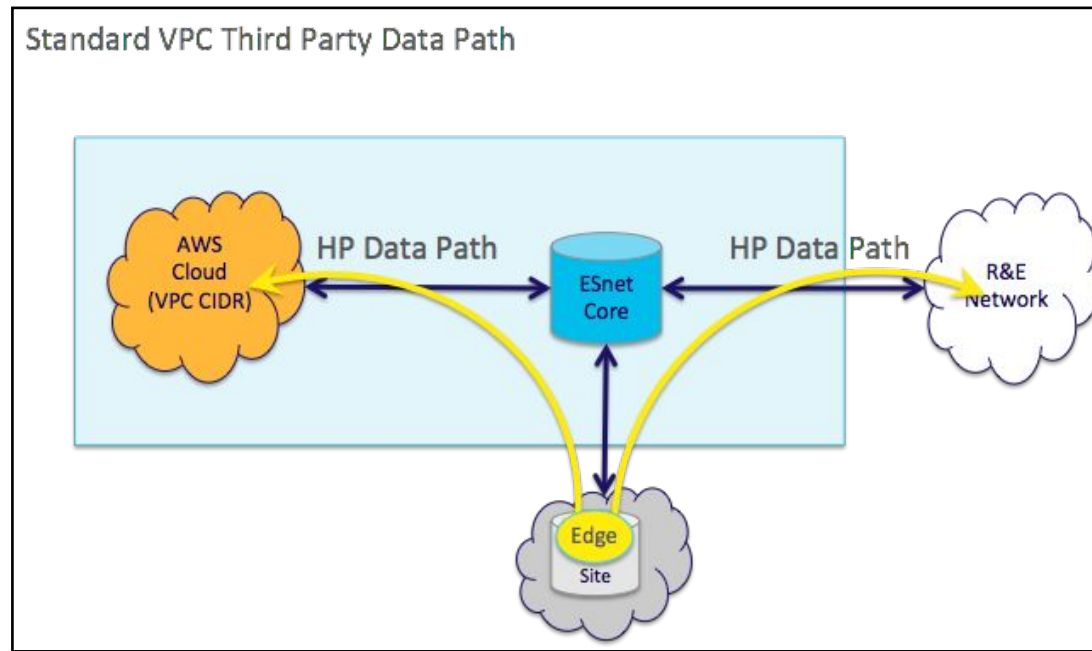
- Started in 2014 in cooperation with US ATLAS at BNL
- BNL was already a user of Amazon AWS and had encountered problems with network capacity.
- ESNet worked with BNL to set up a 10Gbps Direct Connection to AWS East. Later we switched to a 2x10G Direct Connect to the Amazon AWS West region.
- ESnet sites such as BNL utilized our OSCARS service to provision point to point circuits from their network edge to the Direct Connect links.
- BNL originally mainly used the AWS East instance.
- FNAL and PNNL (Pacific Northwest National Laboratory) also participated in the pilot.

Cloud Services Pilot Conclusions

- Data distribution is a challenge
- Connectivity between regions of AWS is not sufficient for R&E needs, never mind connectivity between different cloud providers
- Direct Connect worked well but is expensive. Over \$20K/year for a 10Gbps connection.
- Peering connections can provide reasonable bandwidth if they are sized appropriately.
- Virtual Private Cloud (VPC) services provide benefits similar to Direct Connect with respect to routing

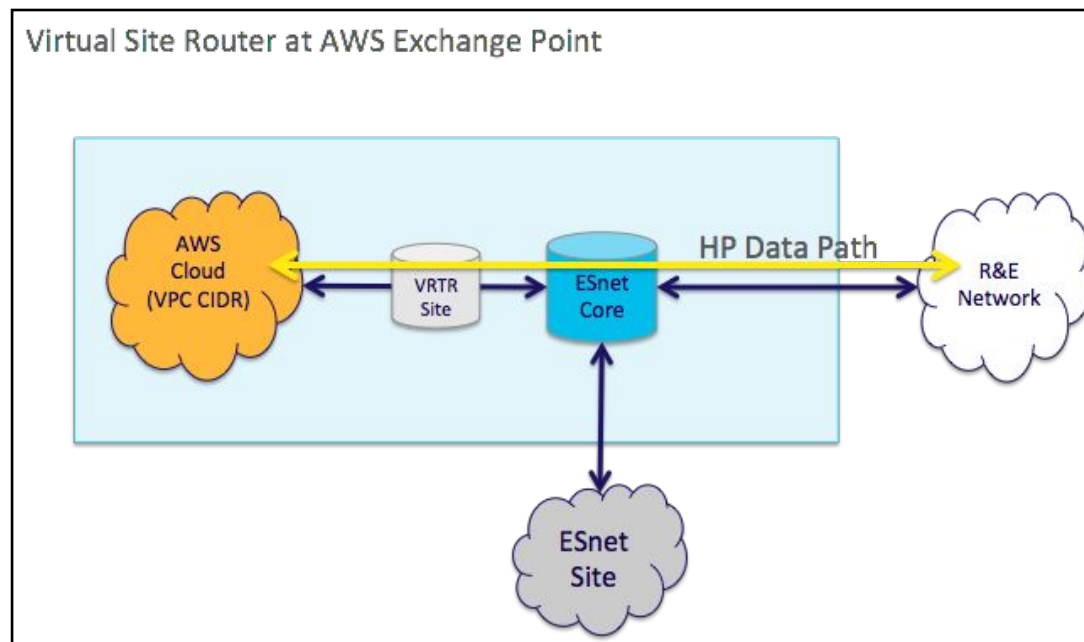
Cloud Services In Progress

- Utilize VPC functionality to provide routing control
- Upgrade public peerings to 100Gbps
- Tunnel traffic back to the site in IPsec or similar



Cloud Service Next Steps

- Implement methods of optimizing routing to R&E peers
- Create virtual 'site routers' at the cloud provider
 - iBGP back to the site and eBGP to ESnet



ESNet Performance Monitoring

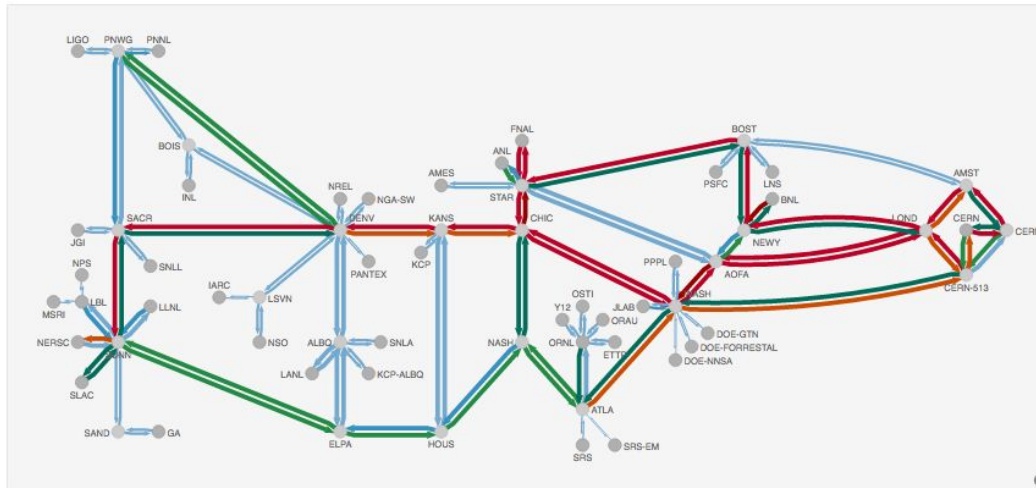
ESNet utilizes multiple systems to monitor and measure our network performance.

- Passive measurement
 - SNMP statistics
 - NetFlow statistics
- Active measurement
 - perfSONAR

ESNet SNMP Statistics



- ESNet collects SNMP from all managed devices
- Internally developed software (ESxSNMP) is used for collection
- 30 second interval for traffic counters
- In addition to collecting standard OIDs we collect some device-specific OIDs as well to monitor OSCARS circuits.
- Visualization is done via Graphite (<https://graphite.es.net>)
- SNMP traffic data is also presented on the portal (<https://my.es.net>)
- Active project underway to research processing SNMP data in the cloud using Google Compute Platform

SNMP Statistics Via The Portal



Network Wide Total Traffic

This map highlights the sites ESnet serves, the structure of the network and the current traffic load. Clicking on a node or edge will show details. Please note that this map does make some simplifications, click ⓘ on the map or chart for more details. The map of the transatlantic extension is a work in progress, more details are available in the info box.

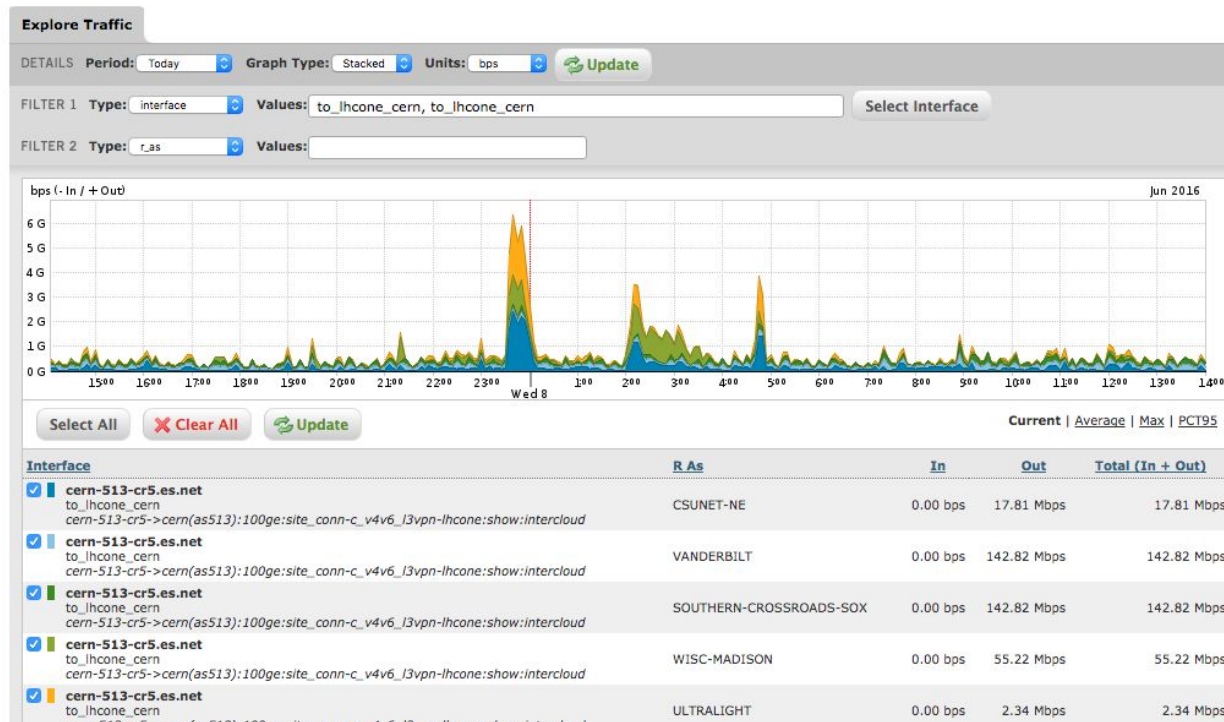


ESnet Monitoring via NetFlow

- NetFlow monitoring is utilized to help us understand now just how much traffic is on the network but what it is and where it is going.
- Primary collection and analysis tool is Arbor Networks Peakflow SP
- Site-facing information provide via the ESnet portal <https://my.es.net>
- Raw NetFlow data archived in case it is needed for ad-hoc analysis.

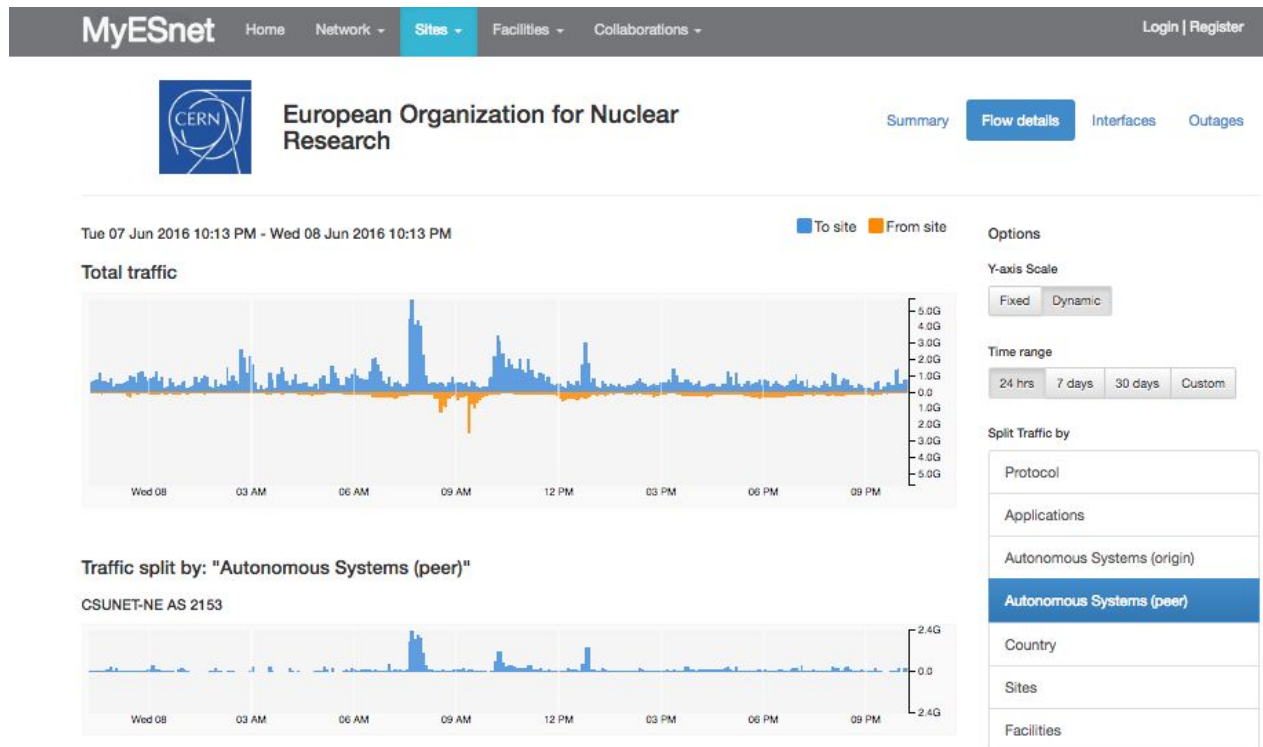
Arbor Peakflow

Allows us to build sophisticated queries and graphs in real-time but the interface isn't really suitable for external use.



Netflow in the ESNet portal

The portal provides sites with pre-configured views of NetFlow-based statistics. The data is obtained from Arbor on the back-end.



Active Measurement via perfSONAR

- ESNet is one of the main development partners on the perfSONAR project.
- ESNet has perfSONAR deployed at about 50 locations throughout the network.
- Some sites have separate OWAMP and BWCTL servers.
- The MaDDash package is used to maintain dashboards of tests.

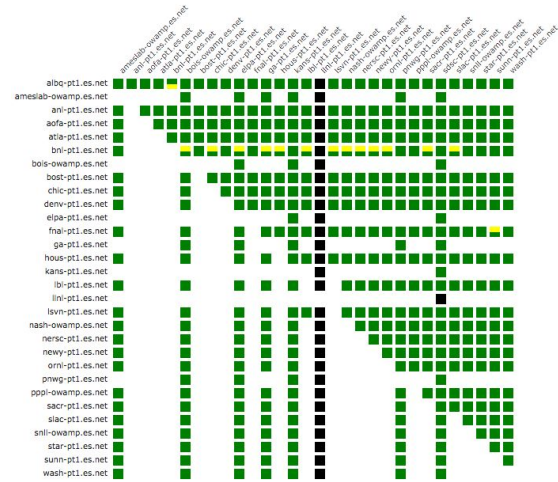
ESNet perfSONAR Dashboard

See <http://ps-dashboard.es.net>

1: ESnet to ESnet Packet Loss Dashboard


ESnet - ESnet to ESnet Packet Loss Testing

■ Loss rate is <= 0.001 ■ Loss rate is >= 0.001 ■ Loss rate is >= 0.1 ■ Unable to retrieve data ■ Check has not yet run




ESnet - ESnet Hub to Small DOE Site Border Packet Loss Testing

■ Loss rate is <= 0.001 ■ Loss rate is >= 0.001 ■ Loss rate is >= 0.1 ■ Unable to retrieve data ■ Check has not yet run




ESnet - ESnet to ESnet Intercontinental Packet Loss Testing

■ Loss rate is <= 0.001 ■ Loss rate is >= 0.001 ■ Loss rate is >= 0.1 ■ Unable to retrieve data ■ Check has not yet run



ESnet - ESnet to ESnet European Packet Loss Testing

■ Loss rate is <= 0.001 ■ Loss rate is >= 0.001 ■ Loss rate is >= 0.1 ■ Unable to retrieve data ■ Check has not yet run



Questions?