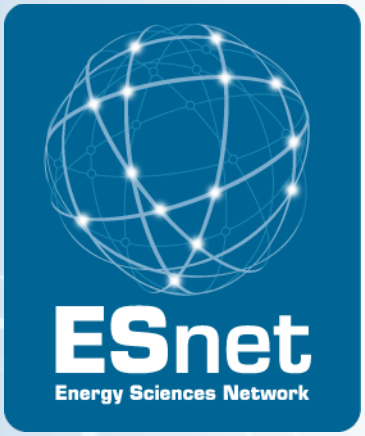

ESnet Update

Steve Cotter, Dept Head

Lawrence Berkeley National Lab

ESCC

Clemson, SC

Feb 2, 2011

DOE Network Requirements Workshops

Capacity
(Science Data > Planned ESnet capacity growth)

- Multiple 100G links by 2015
- Increased commodity Internet peering
- Increased capacity for private peering

Services

- Differentiated services on the IP routed network
- Expanded capability set for virtual circuits (L2, L3, etc.)
- Integrated Services Model

Operations

- End-site education, best practices and consulting services
- Performance management, personalized reporting
- Network Knowledgebase like fasterdata

Additional Topics

- Energy efficiency, reduced carbon footprint
- Security
- Research and Testbeds

Strategic Imperatives for ESnet

- More bandwidth to DOE facilities and Labs at lower costs

- Richer network services in support of distributed science
  - Develop ‘network aware’ integrated services that deliver ‘end-to-end’ high-performance data transfer, HPC/cloud computing, and science collaborative services
- Carrier-class network operations providing high network availability to all DOE facilities
  - Seamless network interoperability across multiple network domains
- Develop and deploy energy-aware and efficient networking infrastructure

- Provide a networking research testbed for DOE community
  - Conduct/enable groundbreaking research in new protocols/storage/ energy efficient networking

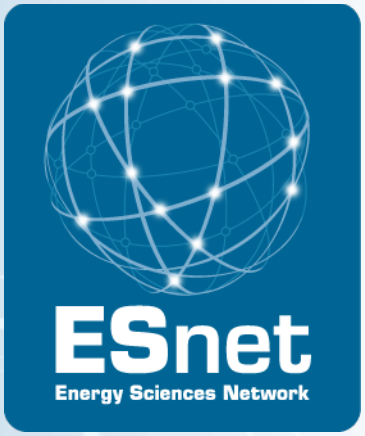

ESnet4 IP & SDN Networks

Current ESnet4 Topology

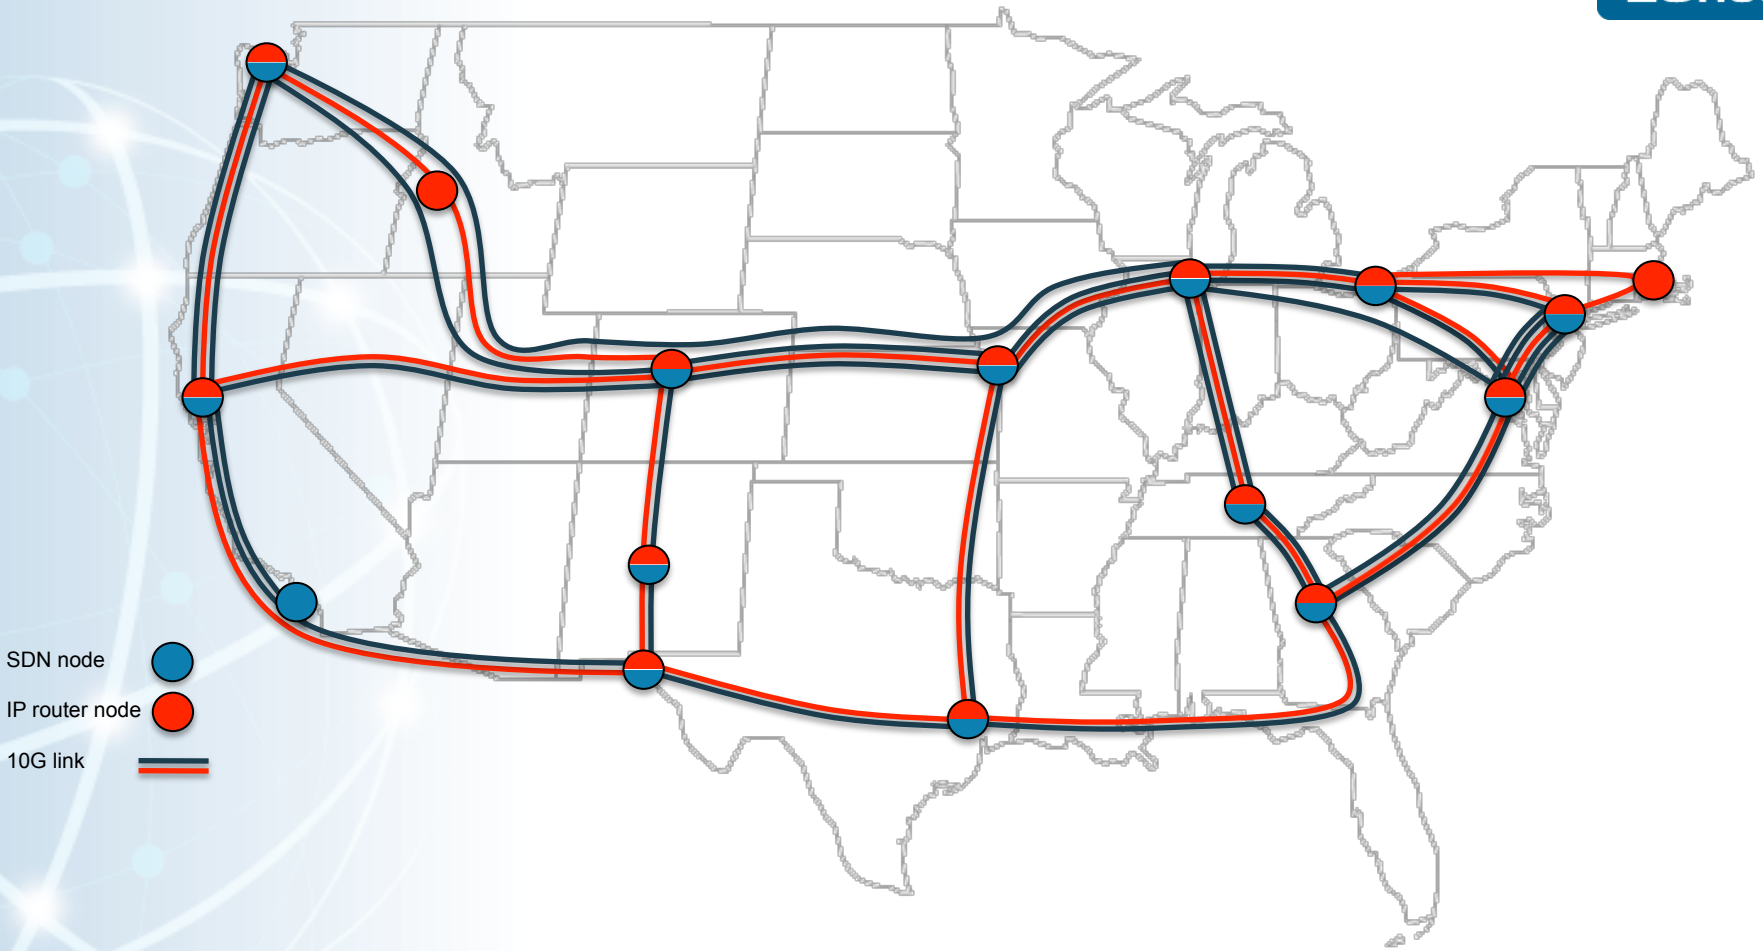

ESnet4 Topology End-2011

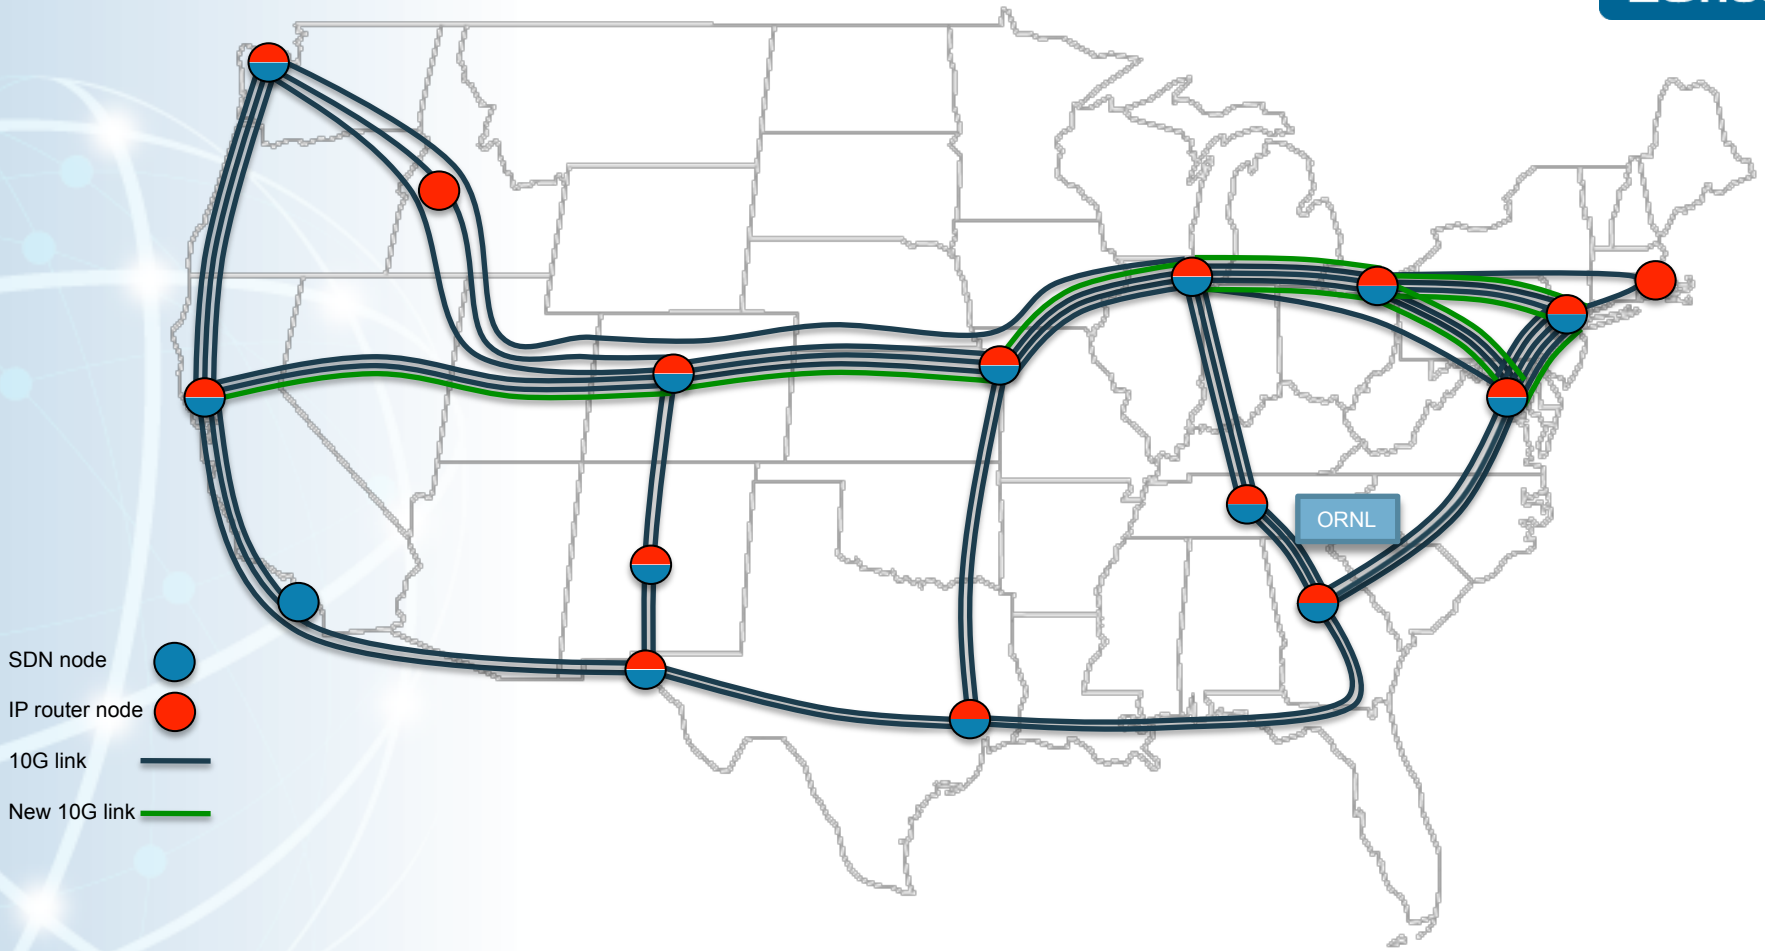

ESnet4 Network

Peaked at 10.6 PB in Nov

>1PB of genomics traffic between JGI – NERSC that month alone

More than 50% continue to go over OSCARS circuits

Instrumenting with perfSONAR all connections 1G and up

99.9983% network availability to Office of Science labs

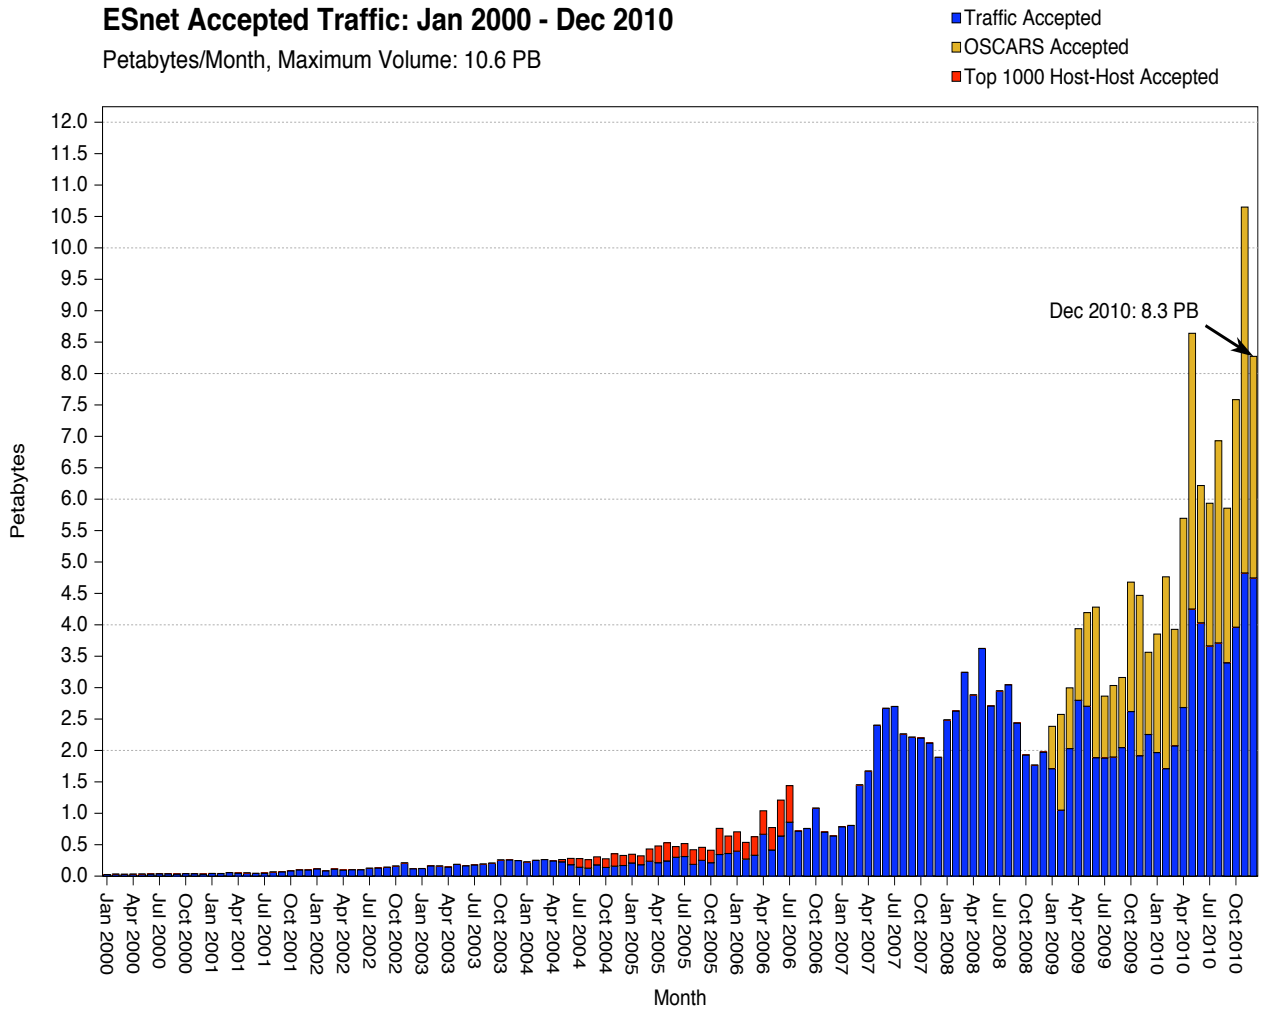

Site Upgrades

- AOFA-HUB 1 Infinera DTC chassis installed (Dec 2010)
- NEWY-HUB 3 Infinera DTC Chassis installed (Dec 2010)
  - Testbed servers at LBL to be installed (Mar/Apr 2011)
- BNL 4 Infinera DTC chassis installed (Dec 2010)
  - Testbed servers at LBL to be installed (Mar/Apr 2011)
- Adding perfSONAR servers to ESnet small sites
  - Completed PPPL, PANTEX< PAIX-PA, OSTI, ORAU, FORR, INL, GA,
  - Future installs: AMESLAB, DOE-GTN, EQX-CHI, EQX-SJ, LASV, LVK, NREL, NSO, SNLA, SRS
- PPPL M120: to be replaced with MX480 TBD

Site Disconnects

- YMP 1GE XC via LASV disconnected (Sept 2010)
- NOAA T1 via ORAU disconnected (Oct 2010)
- NSTec-Liv-OPS disconnected DS3 and removed router (Dec 2010)

Circuit Installs & Upgrades

- 1GE private peering at EQX-ASH with Level3 (July 2010)
- 10GE Level3 circuit replaced the 1GE circuit between Equinix SJ (SV1) (July 2010)
- 10GE MAXGigapop circuit replaced the Qwest 1GE link between Equinix Ashburn (DC2) (July 2010)
- 10GE peering with CICnet at CHIC (Aug 2010)
- 1GE peering with FNAL DOE-FSO (Aug 2010)
- 2 10GE Cross-connect for future peering with NOAA at DENV-HUB and CHIC-HUB (Aug 2010)
- 10GE private peering with MERIT at Starlight (Aug 2010)
- 10GE MERIT wave added to the existing 1GE MERIT link between Equinix Chicago (CH2) (Aug 2010)
- 3 1GE private peering with TATA Communications at all Equinix (EQX-ASH. EQX-SJ and EQX-CHI) (Sept/Oct 2010)

Circuit Installs & Upgrades

- 1GE private peering with SAVVIS at EQX-ASH (Nov 2010)
- 4 10GE SC2010 waves (2 from ATLA, 1 from HOUS, 1 from STAR) (Nov 2010)
- 1GE site ordered from KCP to KANS-HUB (replaced OC3 to ALBU) (Dec 2010)
- 10GE peering with LLNL LVOC (Jan 2011)
- 20G (2x10GE aggraded) connection to FNAL (Jan 2011)
- OC3c to LANL-SW peering in DENV-HUB (Jan 2011 Pending XC)
- NEWY to AOFA Metro OCG darkfiber installed and under test (Jan 2011)
- 100M → 1GE upgrade to Teliasonera peering at EQX-SJ (Jan 2011)
- OC3 to 1GE upgrade on the AOL peering at EQX-ASH (Oct 2010)
- 10M to 100M upgrades for the Sprint peering at EQX-ASH & EQX-SJ (Oct 2010)

In Progress

- 1GE backup PPPL-AofA MAGPI wave, waiting on hardware (Date TBD)
- 10GE SDN / ION peerings in CHIC and SUNN / SEAT (via PacificWave fabric) (date TBD)
- PNNL-OSTI 10M → 100M upgrade pending contract mod.
- SRS-EM 10M → 100M upgrade pending contract mod.
- Future 1GE for NNSA peering at FORR (date TBD)
- Future 1GE for OSTI-Y12 disaster recovery (TBD)
- 1GE wave in BOIS to INL via IRON (TBD)

Long Island MAN

Argh! No more snow! It's impossible to find manholes under 15' snow

Light between 111 8th & 32 AofA waiting for portions of the ring to be constructed

- Southern: (was) early Feb
- Northern: Mid-March

Will handle testbed & production traffic (eventually) using Layer 1 VPN capability

Dramatically reduces our costs to reach BNL

Learned some lessons that will be helpful when ANI gets going

OSCARS Development Status

PCE SDK (v01.11.11) is done

- SDK for development of PCEs for use in OSCASR (v0.6)
- ESnet has used the PCE SDK to develop bandwidth and VLAN selection PCEs
- Several DOE funded projects have identified interest in using the SDK
- PCE SDK is available for downloads to beta-user (contact chin@es.net)

OSCARS (v0.6 native) is close to code complete

- ~120,000 lines of code now (~130K when complete)
- Undergoing testing right now
- OSCARS v0.6 native (i.e. communicates with other OSCARS v0.6 servers) should be in beta in February
- Backward compatibility with v0.5.x is on target to be completed by April

OSCARS Deployment Status

OSCARS is now one of the core applications deployed by SCinet to support dynamic circuits

- First deployed in SCinet in SC09
- For SC10 over 150 dynamic circuits were configured in support of roughly 15 demonstrations

First deployments of OSCARS (v0.6) will be in 2Q2011

- Green field installations as part of I2's DYNES project

Proven use cases

- Extending site LAN over the WAN (Genomics: JGI/NERSC)
- Traffic engineering to avoid congestion points (Cloud testing: LBNL/Google)
- Guaranteeing network resources (LHC T0-T1, T1-T2)
- Deadline scheduling (Fusion DIII-D: GA/NERSC)

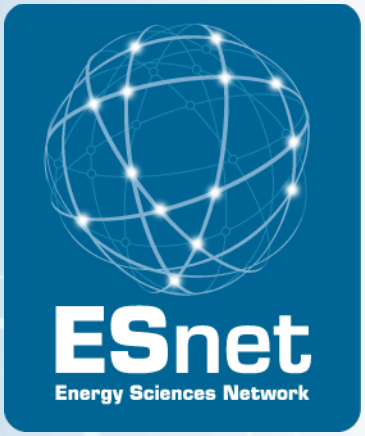

Fermi Challenge: Between 10G and 100G

Architectural Issues

The current ESnet backbone architecture has 3 types of links:

- IP links for carrying normally routed traffic
- SDN links for carrying traffic in engineered circuits (OSCARS)
- DMZ's for connections to sites & peers

The general philosophy has been:

- Move large and predictable flows to circuits over SDN
- 'Everything else' should fit in the 10G IP backbone

FERMI Challenge

They moved the bulk of their traffic to more than 15 OSCARS circuits distributed across multiple 10GE interfaces.

The remaining traffic was still approaching 10G from time to time.

They could not move it to Layer 2 OSCARS circuits because that requires a willing and able partner at the far end of the circuit!

They could not move it to Layer 3 OSCARS circuits because OSCARS Layer 3 circuits start and terminate on ESnet routers, and the bottleneck was on the 10GE link to ESnet on the FERMI egress router

Plan

1. Moved FERMI ‘IP’ DMZ to a 2x10GE aggregated connection
2. Deploy a minor enhancement to OSCARS
   - Allow Layer 3 circuits to dump packets into the standard forwarding table if the MPLS LSP supporting the circuit fails.
3. Re-route several large FERMI ‘IP’ flows onto SDN across the backbone
   - Note that re-routing a flow into a circuit is hit less!
4. Automate portions of the process of identify candidate large flows and re-routing them over OSCARS Layer3 circuits
5. Develop circuit life-cycle management processes
   - Documented processes for deciding when to turn up, change bandwidth, or turn down circuits

Before

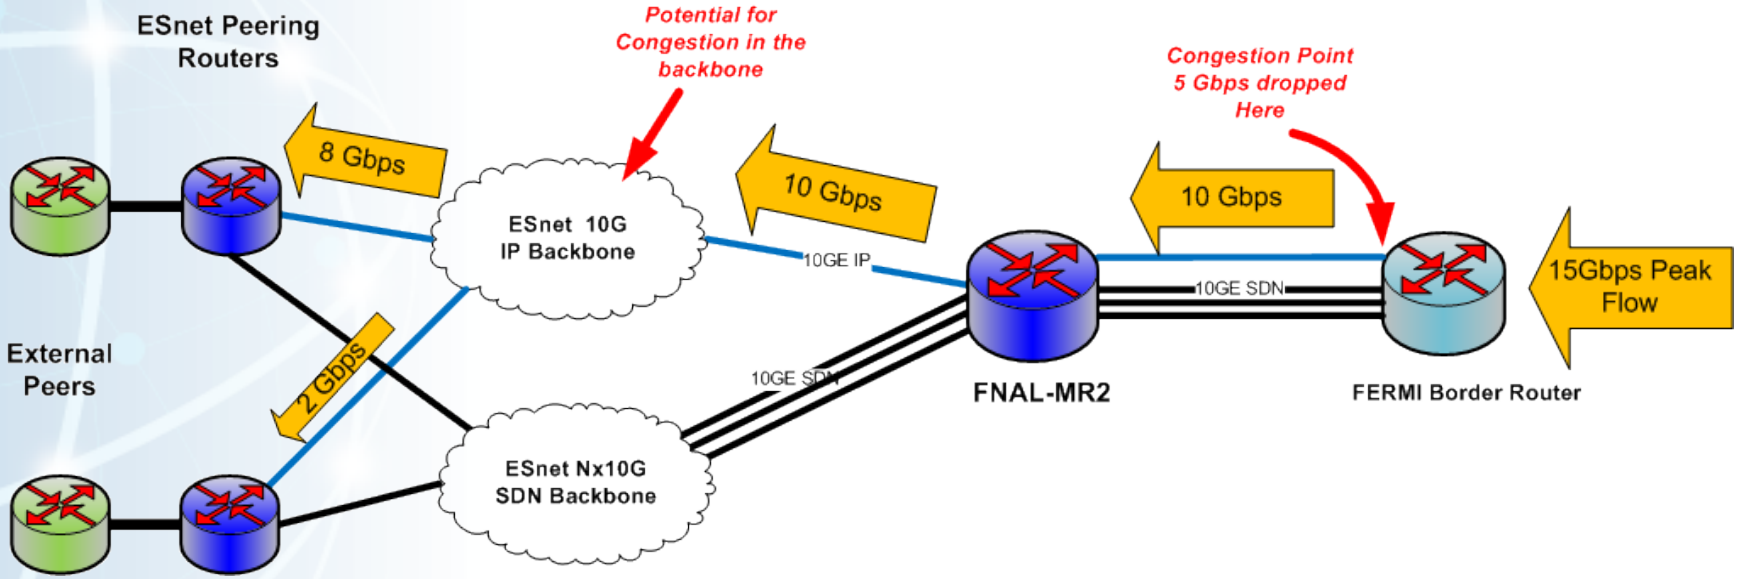
ESnet Peering Routers
Potential for Congestion in the backbone
Congestion Point 5 Gbps dropped Here
8 Gbps
ESnet 10G IP Backbone
10 Gbps
10 Gbps
10GE IP
10GE SDN
15Gbps Peak Flow
External Peers
2 Gbps
10GE SDN
FNAL-MR2
FERMI Border Router
ESnet Nx10G SDN Backbone

Transition

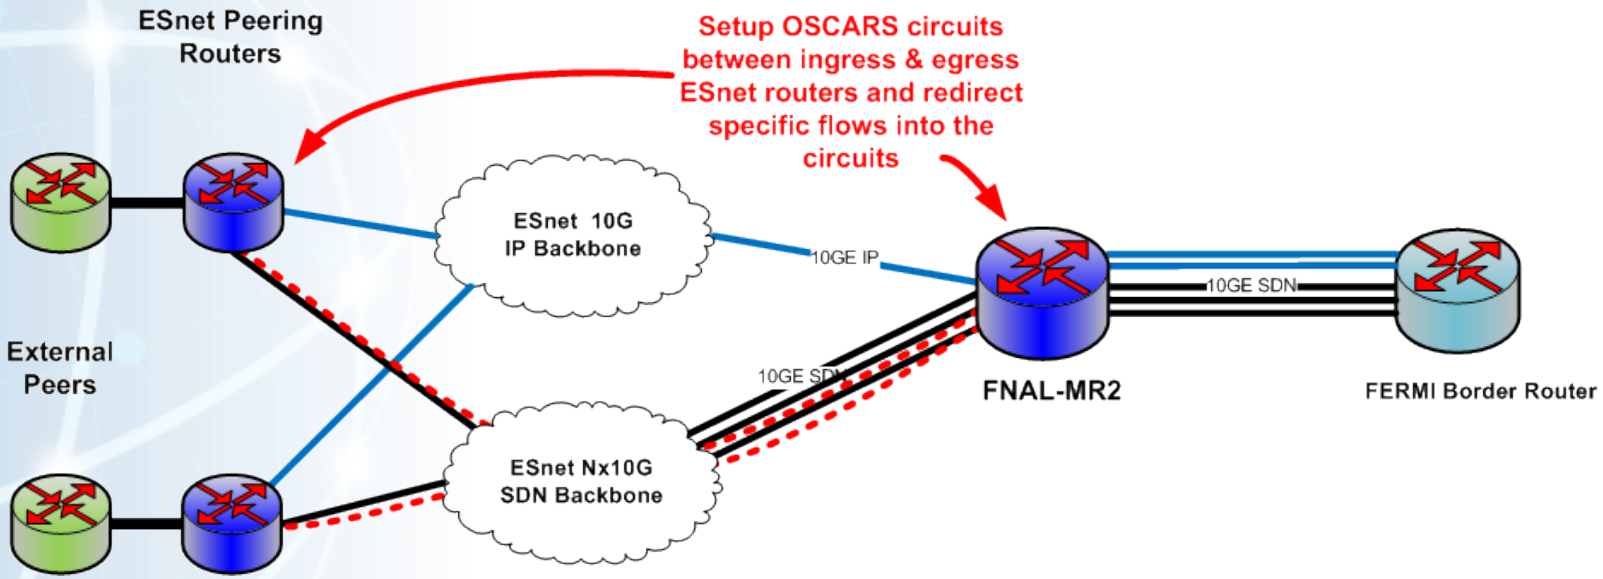
ESnet Peering Routers
Setup OSCARS circuits between ingress & egress ESnet routers and redirect specific flows into the circuits
ESnet 10G IP Backbone
10GE IP
10GE SDN
External Peers
10GE SDN
FNAL-MR2
FERMI Border Router
ESnet Nx10G SDN Backbone

Future

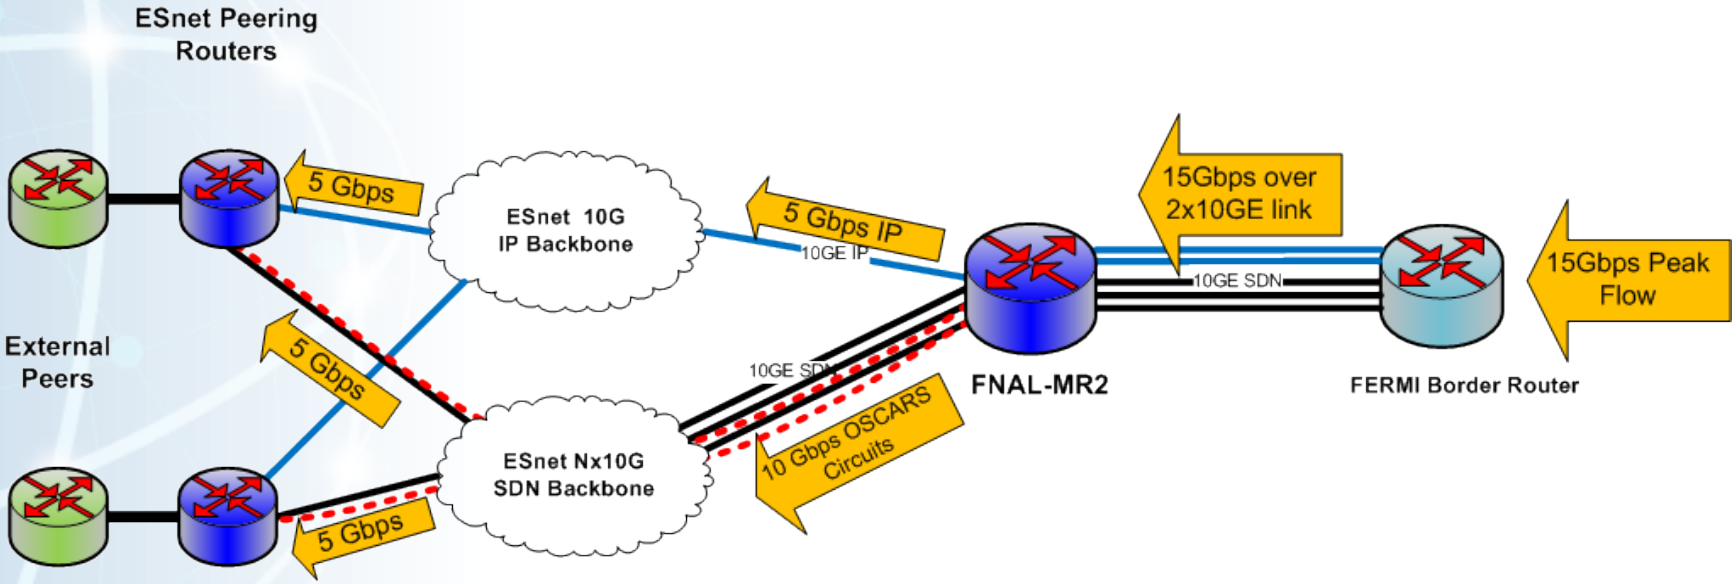
ESnet Peering Routers
External Peers
5 Gbps
ESnet 10G IP Backbone
5 Gbps IP
10GE IP
15Gbps over 2x10GE link
10GE SDN
15Gbps Peak Flow
FNAL-MR2
FERMI Border Router
5 Gbps
10GE SDN
10 Gbps OSCARS Circuits
ESnet Nx10G SDN Backbone
5 Gbps

Open Questions

The tools to re-route IP flows over SDN are available to OSCARS users including:

- ESnet Engineers
- Site Technical Contacts
- Others

Currently, ESnet is taking the lead creating the OSCARS circuits. Do others want to be involved?

What level of coordination should we have between ESnet engineers and site technical contacts when we re-route flows?

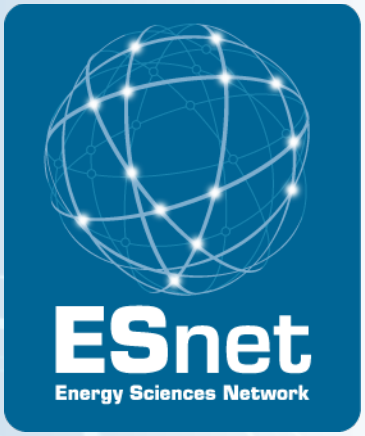

ANI & Testbeds

Advanced Networking Initiative

$62M ARRA funded 100G prototype network and testbed

- Leveraging this infusion of funds to lower long-term networking costs
  - Dark fiber in metro areas
  - Scalable, cost-effective 100G network
  - National dark fiber testbed
- Testbed already in use
  - Table-top being used by first round of research projects
  - Disassembled and shipped to Long Island as LI MAN fiber build completes
  - Second round of research projects selected by Advisory Committee comprising of R&E, lab, and commercial members

Minimizing the Impact of ANI

On staff:

- Hing Chow, PMP certified project manager
- Chris Tracy, Network Engineer
  - Optical experience

New hires:

- Michael Sinatra, Network Engineer (UC Berkeley) starts Feb 8th
  - DNSSec and IPv6 experience
- Patrick Dorn, Network Engineer (NCSA) starts Feb 22nd
  - Active w/Supercomputing
- Gopal Vaswani, Web Apps Dev (start-up) starts Feb 22nd
  - Web portal, working w/Dugan on network tool development

Open positions:

- Group Lead, Infrastructure Team
- Strategic Partnerships & Site Outreach Coordinator
- Software Developer

Upgrading infrastructure with an eye toward:

- Scalability
- Automation
- Reduced support requirements

Don't hesitate to bring it to our attention if you see issues!

New Ticketing System

- Moving to Service-Now.com SaaS (Software-as-a-Service) solution for ticketing
- Same vendor as UC Berkeley, NERSC, LBNL IT Division
- Benefits: easier to open, query, update tickets, track issues
- On our roadmap: integration with new website
  - ticketing portal + knowledge base for the entire community
- We're interested in your experience: have you adopted a new ticketing system recently?
  - please share your experience with greg@es.net

OpenDevNet

Platform for testing software/systems prior to deployment

Virtualization

- 40 cores, 128 GB RAM, 6TB storage

Performance hosts

- 2 Data Transfer hosts (10G disk to disk, RDMA capable)

Networking

- 1 OpenFlow switch

OpenDevNet

Users

- 13 users (9 ESnet, 4 external collaborators)
- 30 virtual machines

Projects

- OSCARS 0.5 testing
- OSCARS 0.6 development and testing
- perfSONAR development and testing
- ANI testbed control framework development and testing
- Iperf testing
- OpenFlow testing
- GENI/ORCA testing
- RDMA testing with WAN simulation

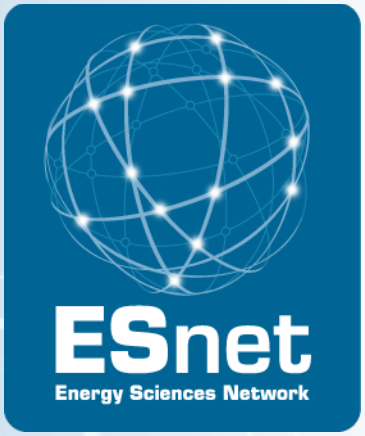

New ESnet Website

Our Vision and Goals

More relevant, engaging, up-to-date content

World-class tools for DOE networking community and beyond

- 'My ESnet portal' to offer custom view of information you care about: interface / circuit utilization, maintenance calendar, tickets, knowledge base…
- Talented tools developer (Gopal Vaswani) starting Feb 22

Easier to update, maintain, extend

- CMS (content management system) technology

No New Passwords

To customize your ‘My ESnet’ portal, a login is required

- But we won’t force you to get a new password…
- instead we’ll leverage federated identity, using InCommon or OpenID accounts

It’s likely you have one of these already

- Existing accounts with Google, Yahoo, AOL (and many other organizations) will work

We’ve been promoting federated identity throughout the DOE community

- Science Identity Federation

Status

Currently finalizing content and site architecture

Our web contractor is finishing styling, integrations

Hope to ‘go live’ in early March

We’ll be eager to get feedback on the new site

- CMS makes it very easy to update, refine, optimize

ESnet Website: now

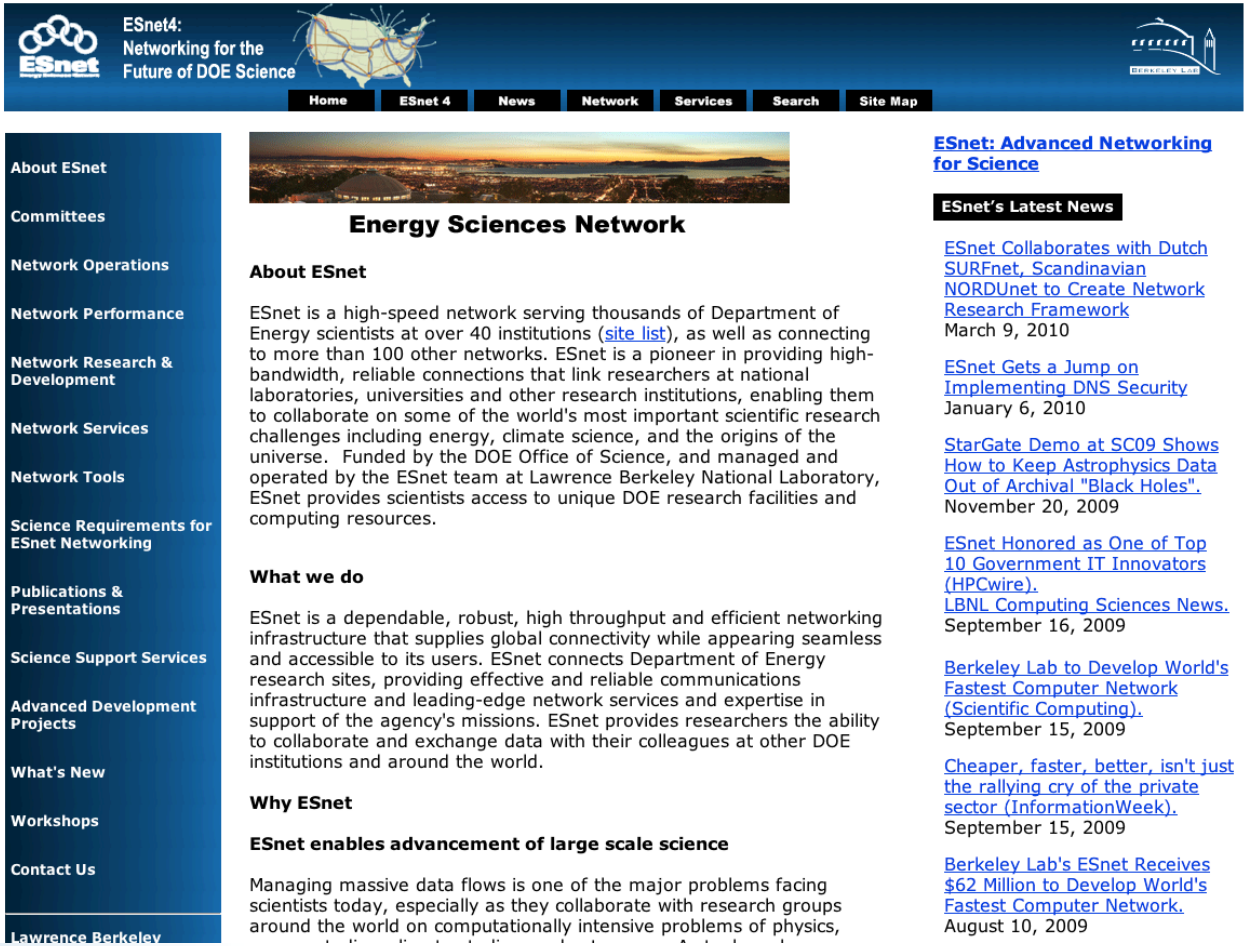

ESnet Website: soon

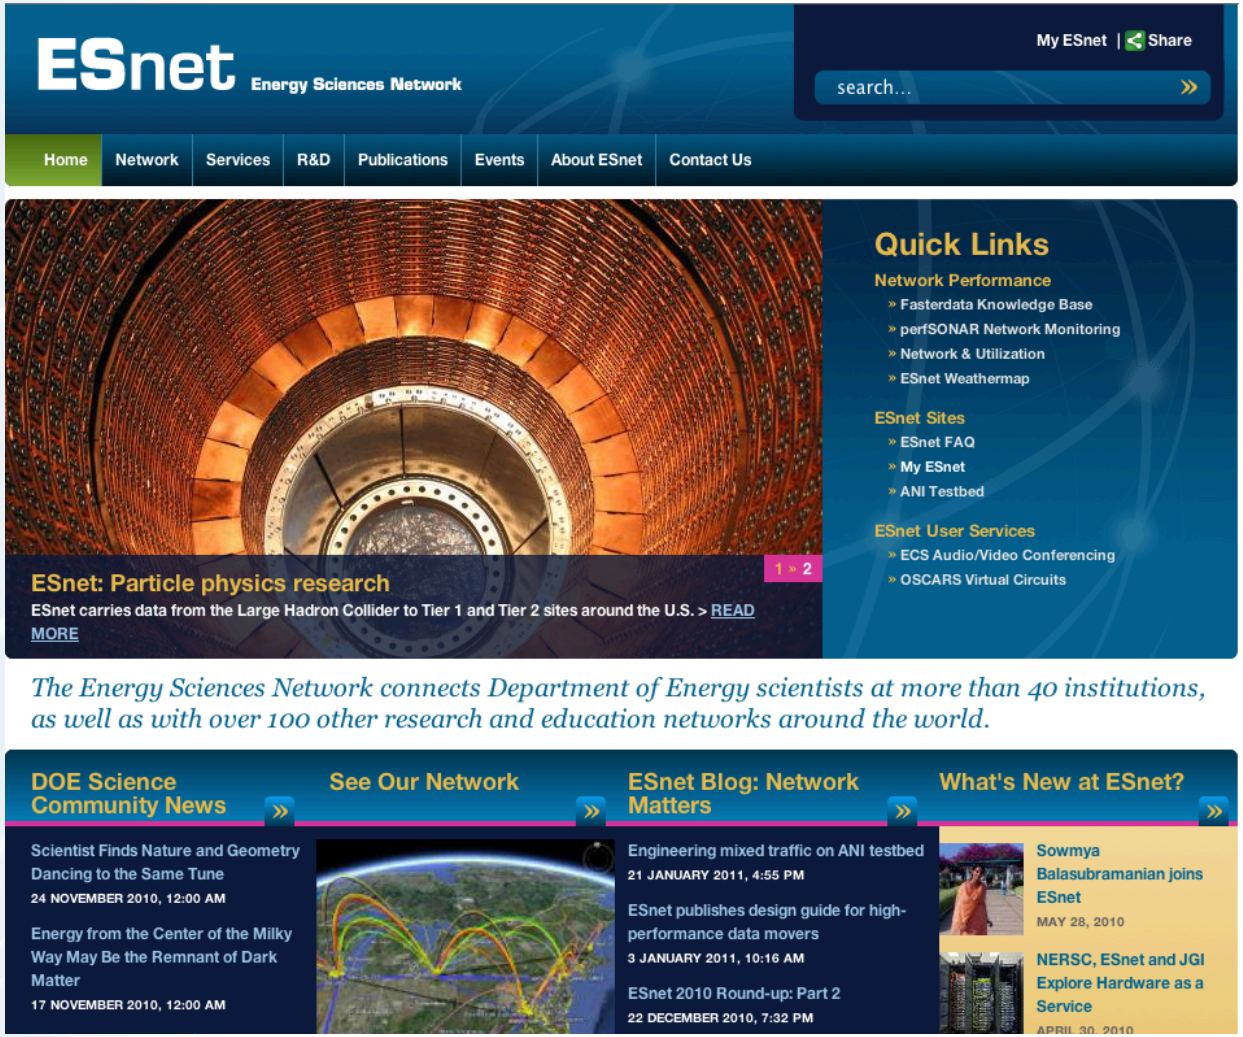

ESnet Website: Soon

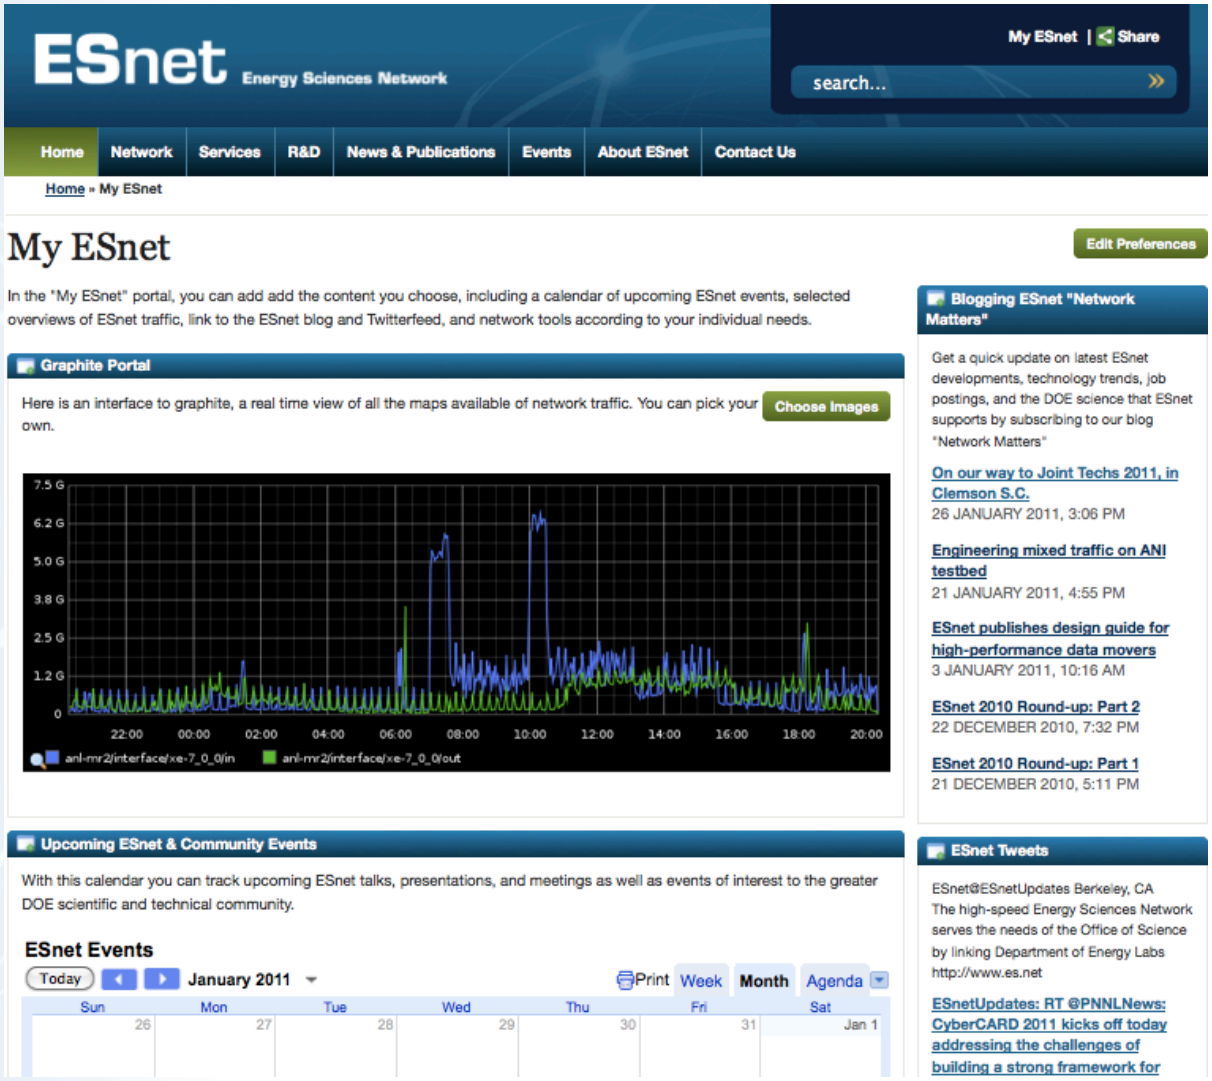

Web Content

Are existing sites like blog or fasterdata.es.net useful to you?

What sort of portal content, tools or visualizations would you like to see?

- Real-time, global perfSONAR app on Google Earth?
- OSCARS circuit visualizations?
- Network performance data/ interface utilization stats?
- Is there interest in social networking applications to share project specific information, publications, slides, maybe instruments etc.?

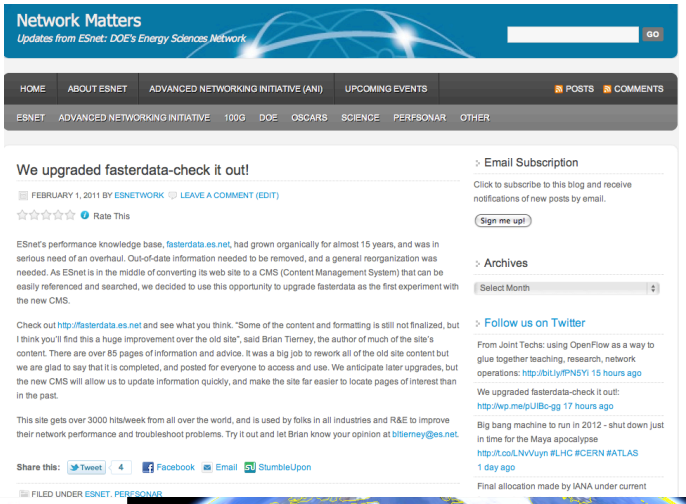

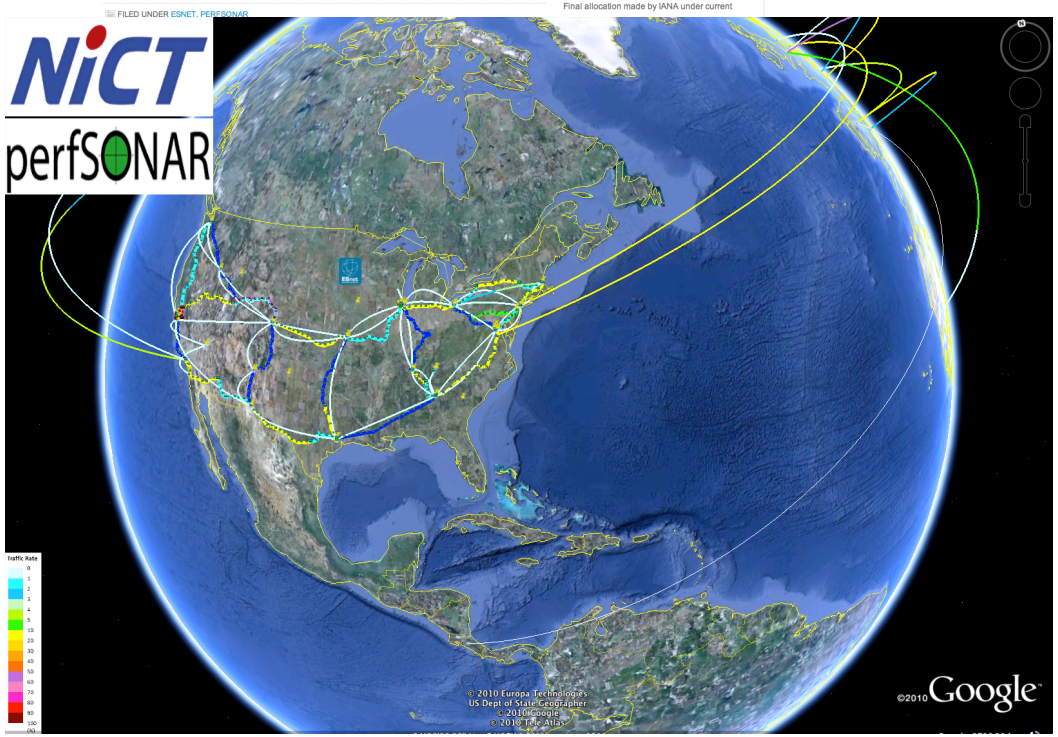

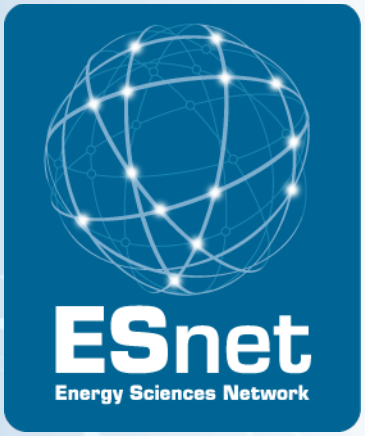

Authentication & Trust Fabrics Update

DOEGrids CA

DOEGrids v4 nears completion

- Distributed key management devices (netHSM) are online
  - Sunnyvale
  - Chicago
- Distributed key management in place
  - California: OSF (ESnet operations)
  - New York: ESnet remote engineer
- CA Cloning tested – current efforts
  - Smooth transition between old and new CA operation
  - Uniform certificate revocation management
  - Uniform agent and user interfaces

Science Identity Federation (SIF)

Focus: Get labs into InCommon

- Blanket agreement in place
  - 3 year memberships in InCommon for DOE labs, sites, user facilities
  - Primary: SC sites - but there should be sufficient funding for some non-SC
- Sites sign a participation agreement with InCommon
- Berkeley Lab will manage the procurement (billing)
- Agreement web sites will be available soon (maybe this week)

Future focus: applications; interoperability with DOE – ICAM

(Contact Mike Helm helm@fionn.es.net for more information)

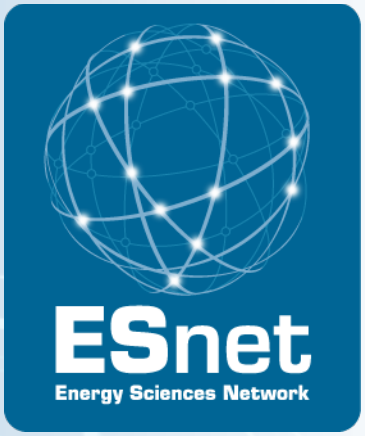

Network Analytics

Network Traffic and Routing Analysis

Key Benefits

- Provides a data-driven understanding of the conduct of science as seen by the network
- Observation of network traffic flow characteristics over time allows us to understand how the conduct of science is changing from the perspective of the network
  - Examples include:
    - Dramatic rise in the mid-2000s of the parallel data movers
    - Growing dominance of science traffic as the primary customer of ESnet
- Enables ESnet staff to trace large flows and understand where they are coming from and going to
- Improves productivity of staff, reduces network instability/anomalies

Network Traffic and Routing Analysis

Legacy NetFlow and BGP analysis system (2004-present)

- ESnet staff has developed and maintained an in-house, custom analysis system
  - Based on open-source tools
  - Limited number of reporting options
  - Significant overhead involved in maintenance, debugging and development of new features

Motivation for switching to a commercial system

- Provides more sophisticated reporting options, “what-if” analysis
- Minimize overhead in maintenance and development of legacy system (will be kept for purpose of raw data archiving)
- Commercial support, increased robustness

Network Traffic and Routing Analysis

Deployment of Commercial Appliances

- Arbor Networks: Peakflow SP
  - Models traffic across the entire network by monitoring NetFlow statistics and BGP attributes at the peering edge
  - Chosen for their extensive BGP analytics and automatic tracking mechanisms which provide many different ways of examining network traffic
  - Provides ESnet staff with the necessary data to make informed decisions about routing, transit, partners, customers and quality of service
  - Includes alerting capability to notify operations of significant changes to the network (DDoS attack, link saturation, etc.)
  - Includes support for a customer-facing web-based portal

Network Traffic and Routing Analysis

Arbor Networks: Peakflow SP

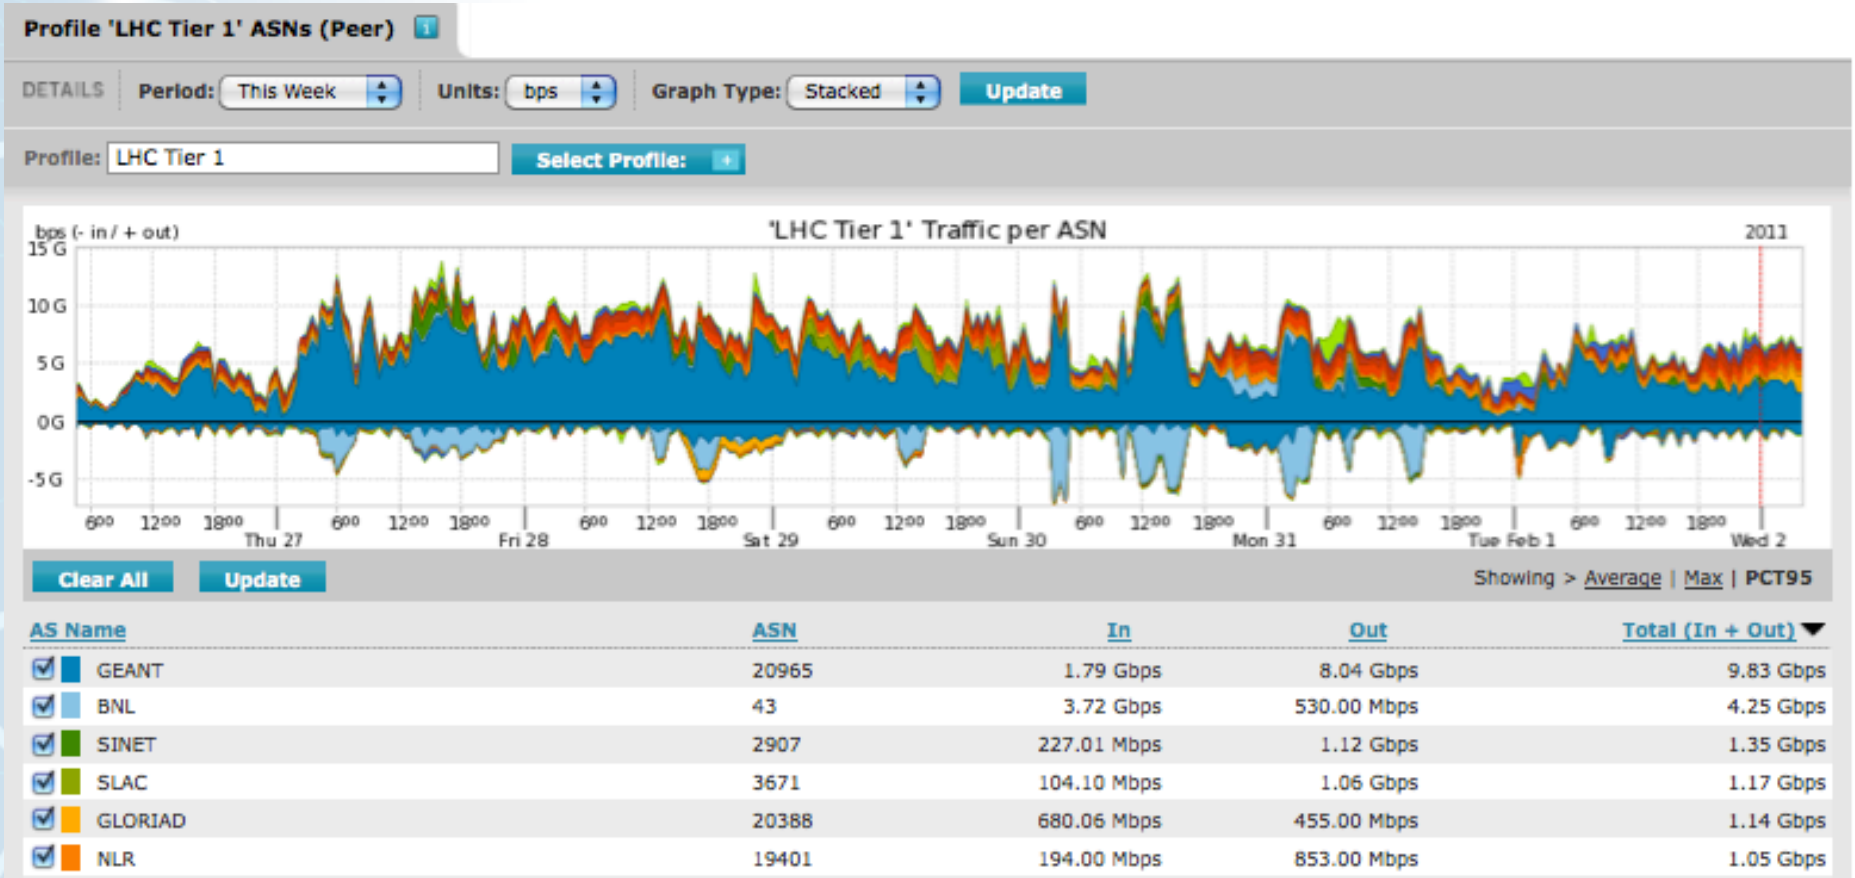

AS Name	ASN	In	Out	Total (In + Out)
GEANT	20965	1.79 Gbps	8.04 Gbps	9.83 Gbps
BNL	43	3.72 Gbps	530.00 Mbps	4.25 Gbps
SINET	2907	227.01 Mbps	1.12 Gbps	1.35 Gbps
SLAC	3671	104.10 Mbps	1.06 Gbps	1.17 Gbps
GLORIAD	20388	680.06 Mbps	455.00 Mbps	1.14 Gbps
NLR	19401	194.00 Mbps	853.00 Mbps	1.05 Gbps

Network Traffic and Routing Analysis

Deployment of Commercial Appliances

- Packet Design: Traffic Explorer & Route Explorer
  - Provides integrated, real-time view of network-wide routing and traffic behavior by monitoring NetFlow data and routing protocols
  - Chosen for their planning mode features and routing analytics
  - Enables ESnet staff to rapidly prototype network optimization schemes and perform “what-if” analyses with historical data
    - e.g., report on impact of link/node failures, understand impact of expected traffic growth for one particular science community
  - Traffic grouping features allow flows belonging to a particular customer or science community to be reported on separately
  - Improves productivity of network engineering staff by automating the analysis of complex problems previously done manually

Network Traffic and Routing Analysis

Packet Design: Traffic Explorer & Route Explorer

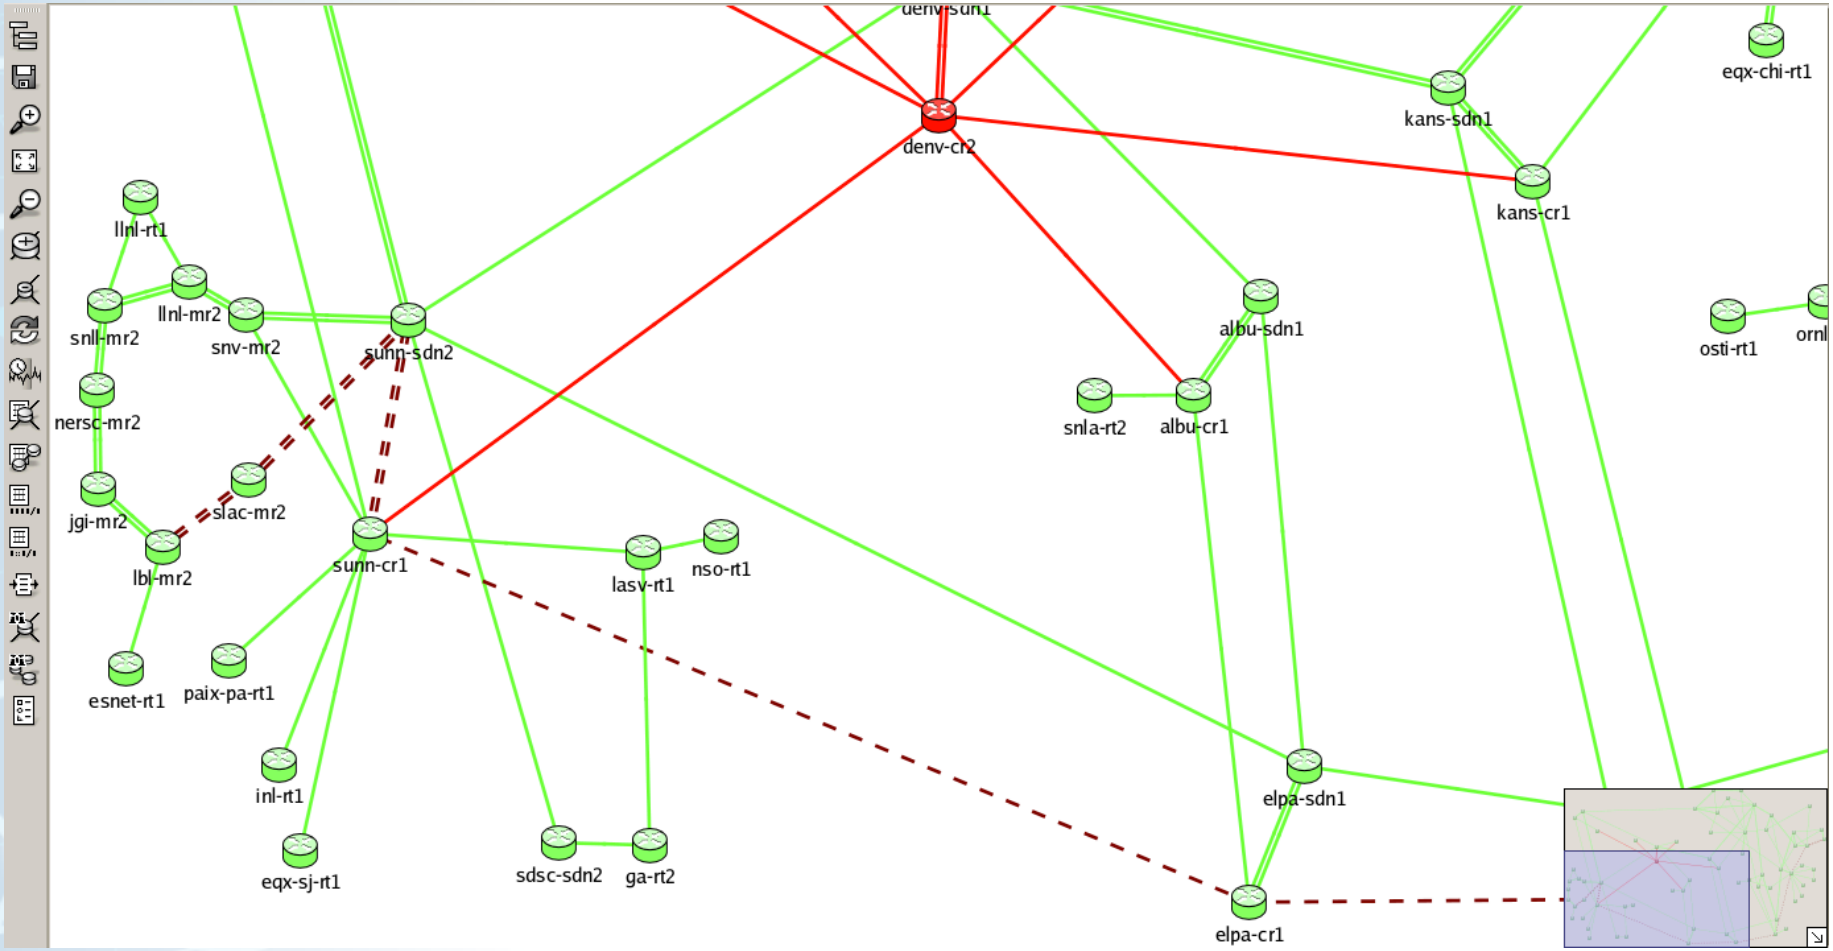

Network Traffic and Routing Analysis

Current Status of Deployment

- Both Arbor and Packet Design provided loaner systems which were evaluated by ESnet staff during a 60-day trial period
  - Evaluation period has given ESnet staff the opportunity to become familiar with how to deploy and manage each system
  - Insight provided by these systems during our evaluation helped to identify problems which had gone unnoticed
    - Asymmetric IP routing
    - Duplicate flow sampling
- Orders placed for Arbor and Packet Design appliances in Aug 2010
  - Deployed at LBNL in early Nov 2010
  - Now used by ESnet staff for traffic, routing and peering analysis

Arbor Managed Objects
Baseline traffic to/from GEANT (as20965)

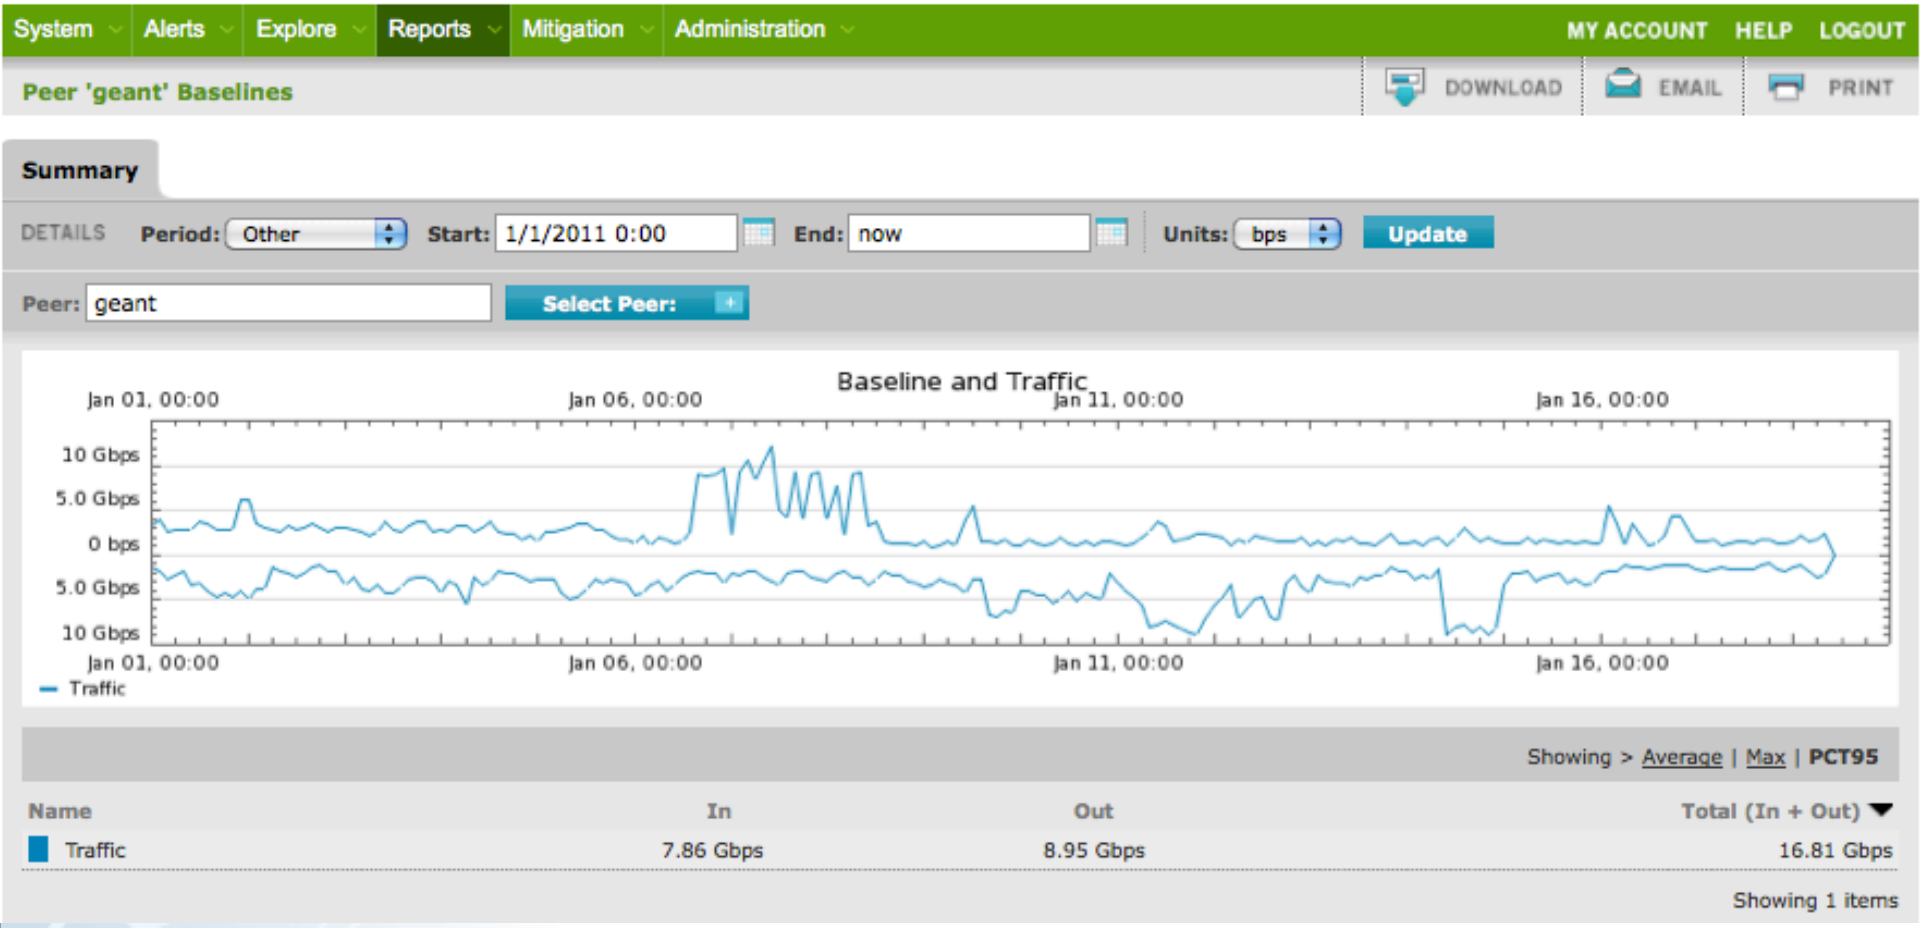

Arbor Managed Objects

Traffic to/from GEANT broken down by Origin AS

Peer 'geant' ASNs (Origin)

DETAILS | Period: This Week | Units: bps | Graph Type: Stacked | Update

Peer: geant | Select Peer:

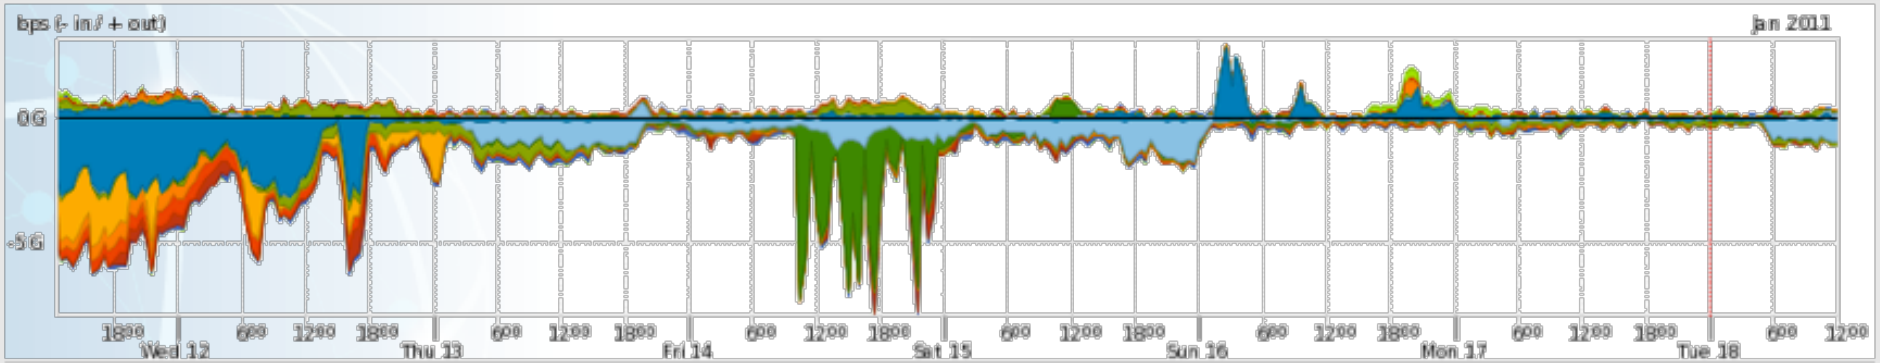

Clear All | Update

Showing > **Average** | Max | PCT95

AS Name	ASN	Into Peer	Out of Peer	Total
JANET	786	413.71 Mbps	145.78 Mbps	559.49 Mbps
DESY-HAMBURG	1754	336.31 Mbps	20.82 Mbps	357.13 Mbps
DFN-IP	680	269.27 Mbps	26.32 Mbps	295.59 Mbps
ASGARR	137	141.34 Mbps	103.37 Mbps	244.71 Mbps
ARNES-NET	2107	150.16 Mbps	4.73 Mbps	154.89 Mbps
IN2P3	789	115.51 Mbps	39.13 Mbps	154.64 Mbps
REDIRIS	766	107.85 Mbps	18.17 Mbps	126.02 Mbps
SURFNET-NL	1103	80.87 Mbps	25.87 Mbps	106.74 Mbps
EENAS	3221	76.27 Mbps	7.33 Mbps	83.60 Mbps
RCCN	1930	11.19 Mbps	28.38 Mbps	39.57 Mbps

Arbor Managed Objects

Traffic to/from Google broken down by ESnet customers (DOE sites)

Peer 'google' x Remote AS

DETAILS | Period: This Month | Units: bps | Update

Peer: google | Select Peer:

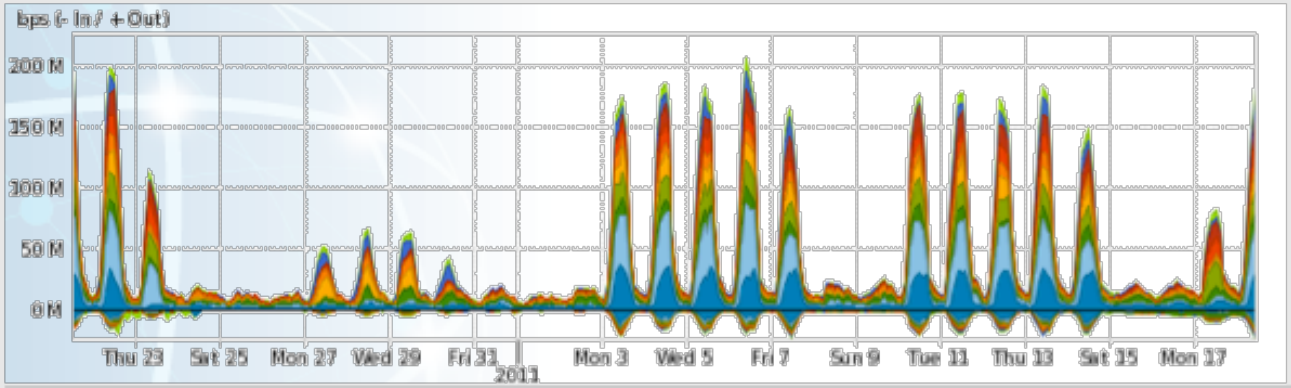

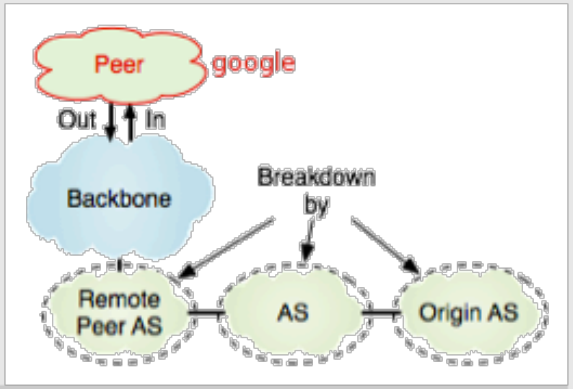

Clear All | Update

Showing > **Average** | Max | PCT95

AS Name	ASN	Into Peer	Out of Peer	Total (In + Out)	% Total
LBL	16	1.92 Mbps	6.90 Mbps	8.81 Mbps	14.01%
LANL-INET	68	527.20 Kbps	8.21 Mbps	8.74 Mbps	13.89%
BNL	43	672.56 Kbps	4.30 Mbps	4.97 Mbps	7.90%
SNLA-NET	377	419.94 Kbps	4.50 Mbps	4.92 Mbps	7.82%
LLL-TIS	45	300.92 Kbps	4.08 Mbps	4.38 Mbps	6.96%
FNAL	3152	452.89 Kbps	3.63 Mbps	4.08 Mbps	6.49%
INL	10702	306.80 Kbps	3.55 Mbps	3.86 Mbps	6.13%
ARGONNE	683	651.00 bps	3.39 Mbps	3.39 Mbps	5.39%
ORNL-MSRNET	50	310.47 Kbps	2.88 Mbps	3.19 Mbps	5.07%
ESNET-WEST	292	581.02 Kbps	2.26 Mbps	2.84 Mbps	4.51%

01/11/11

Arbor Managed Objects

Traffic to/from LHC Tier 1 sites broken down by ESnet peers

Profile 'LHC Tier 1' ASNs (Peer)

DETAILS | Period: This Month | Units: bps | Graph Type: Stacked | Update

Profile: LHC Tier 1 | Select Profile:

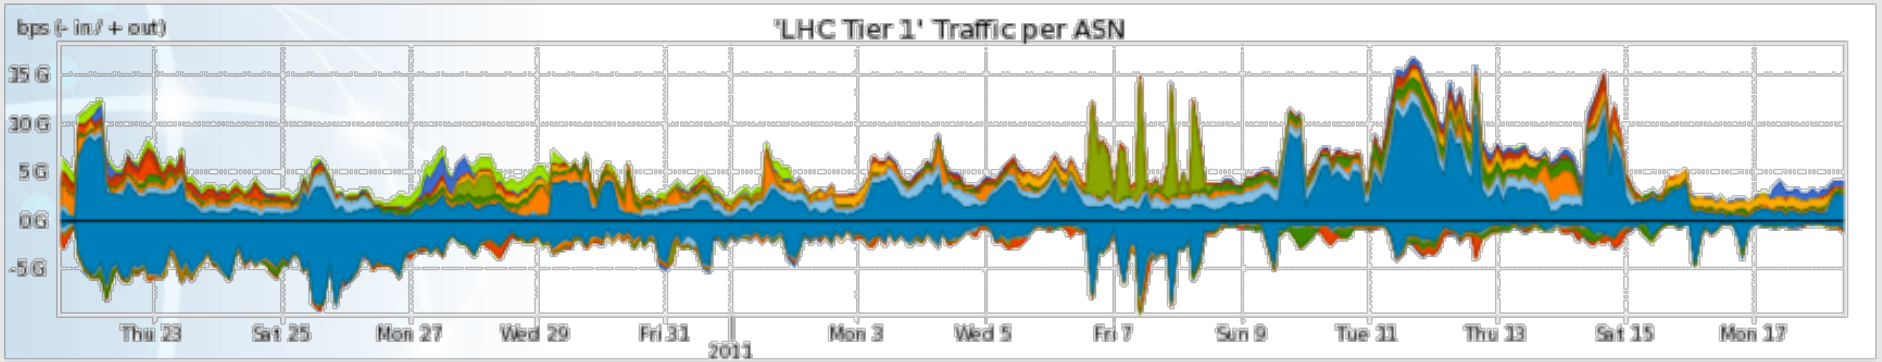

Clear All | Update

Showing > **Average** | Max | PCT95

AS Name	ASN	In	Out	Total (In + Out)
GEANT	20965	2.11 Gbps	2.44 Gbps	4.55 Gbps
NLR	19401	91.16 Mbps	609.48 Mbps	700.65 Mbps
GIGAPOP-NE	10578	112.78 Mbps	423.51 Mbps	536.29 Mbps
ABILENE	11537	144.35 Mbps	389.03 Mbps	533.38 Mbps
FNAL	3152	68.47 Mbps	321.25 Mbps	389.72 Mbps
CSUNET-NE	2153	3.15 Mbps	323.77 Mbps	326.92 Mbps
BNL	43	135.64 Mbps	146.99 Mbps	282.63 Mbps
CERN	513	17.52 Mbps	256.38 Mbps	273.90 Mbps
GLORIAD	20388	43.72 Mbps	213.91 Mbps	257.63 Mbps
AS14041	14041	5.52 Mbps	244.07 Mbps	249.58 Mbps

Arbor Peering Analytics

Verizon Business Peering in SJ is Warm - Who's using it most?

Arbor Peering Analytics

Is it worth it to peer directly with Telefonica (as12956) ?

Summary

DETAILS | **Period:** 4 Weeks | Update

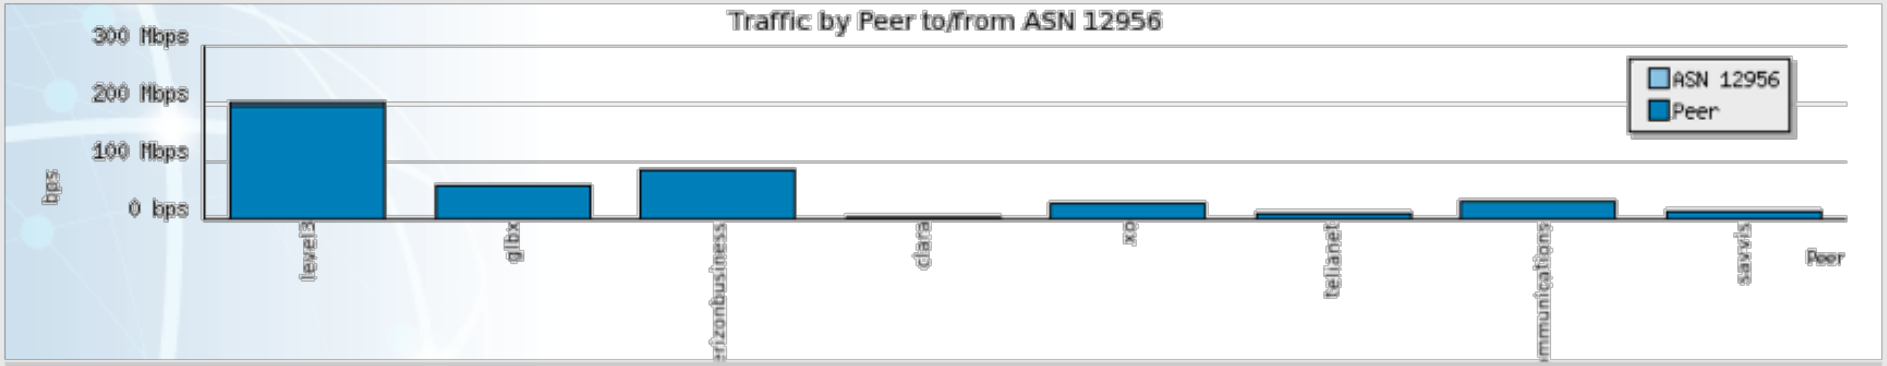

Clear All | Update

Showing > **Average** | Max | PCT95

Peer	Into Peer (ASN/Peer)	Out of Peer (ASN/Peer)	Total Savings
level3	7.54 Mbps / 117.91 Mbps	451.11 Kbps / 78.51 Mbps	7.99 Mbps
glbx	10.00 bps / 42.15 Mbps	29.50 Kbps / 16.60 Mbps	29.51 Kbps
verizonbusiness	23.00 bps / 44.58 Mbps	9.53 Kbps / 41.41 Mbps	9.55 Kbps
clara	0.00 bps / 1.16 Mbps	7.22 Kbps / 1.36 Mbps	7.22 Kbps
xo	0.00 bps / 21.79 Mbps	2.73 Kbps / 5.52 Mbps	2.73 Kbps
telianet	0.00 bps / 7.41 Mbps	1.09 Kbps / 3.41 Mbps	1.09 Kbps
tatacommunications	0.00 bps / 26.39 Mbps	437.00 bps / 5.33 Mbps	437.00 bps
savvis	0.00 bps / 9.22 Mbps	21.00 bps / 4.90 Mbps	21.00 bps

Arbor DDoS Alert

To single IP address at SLAC - really an attack, or slashdot effect?

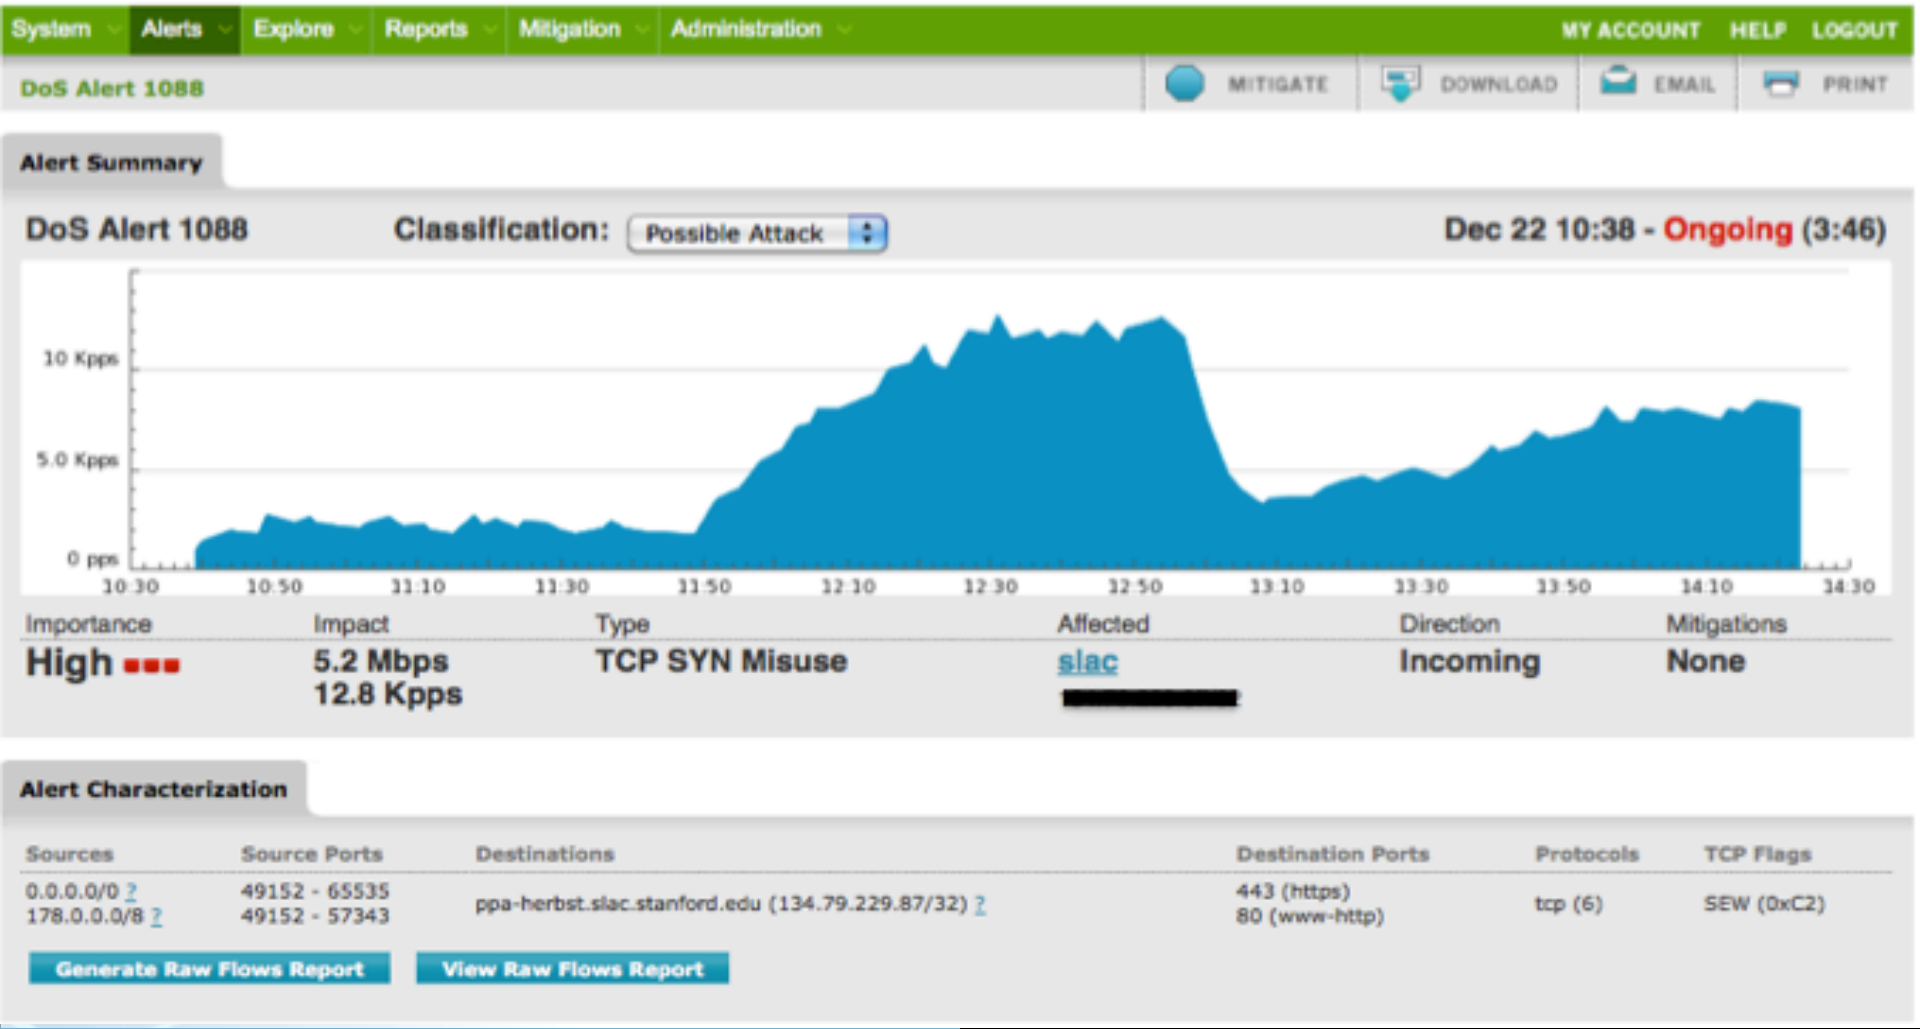

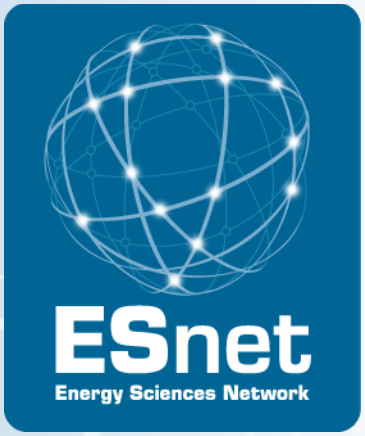

Experimental Capacity Planning Tool

Experimental Capacity Planning Tool

ESnet has grown to beyond the point where it is no longer feasible to use manual processes to determine when a circuit needs to be upgraded.

What kinds of traffic patterns can we detect that might give us an idea when a link is becoming saturated?

But what does saturated mean?

- There is no widely held idea of what percentage of utilization indicates that a circuit is becoming saturated. A thread on NANOG has ranges from 40% to 95% http://nanog.markmail.org/thread/vmkjrldzvsrbz3en
- The traffic profile of the network will effect this threshold. R&E networks are significantly different than the commercial Internet.

Approach

Taking some of the ‘art’ out of deciding when to add capacity

Process:

- Define metrics that can describe utilization patterns that may indicate circuits are busy
- Analyze all circuits each day and assign a score for each metric
- Initially the thresholds for the metrics and for which values of each metric indicate an “interesting” circuit will require human supervision
- Over time note which circuits are upgraded
- Use the decisions from step 4 along with the historical scores for each metric as an input for logistic regression

Peaks Metric

The number of sampling intervals where the utilization is over some threshold T.

In this example:

- T = 75%
- The value of the metric is 7 (there is a peak at time 0)

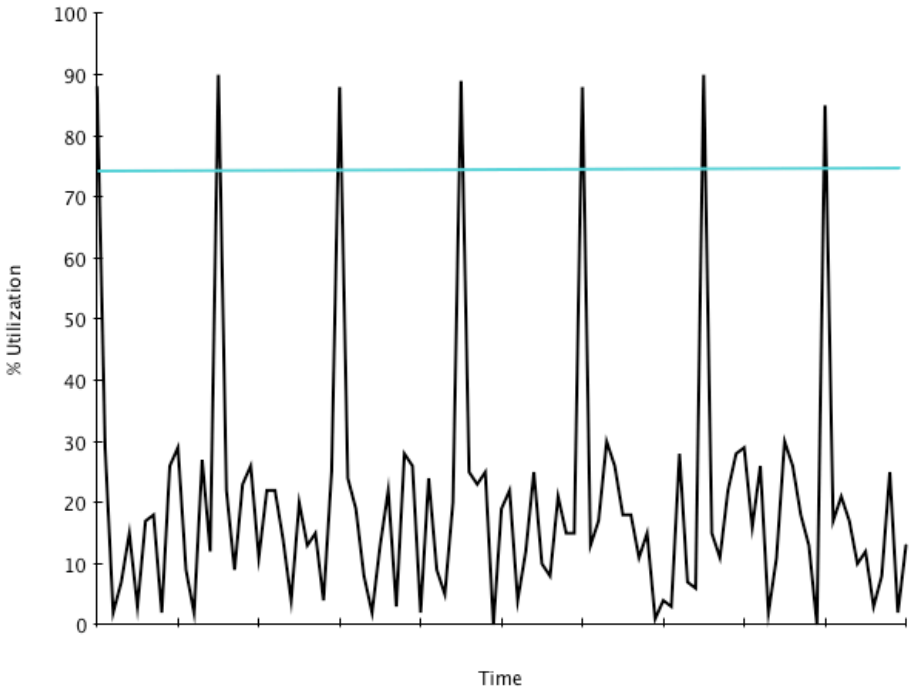

Plateaus Metric

The number of sampling intervals where the utilization is over some threshold T for at least t intervals.

In this example:

- T is 70%
- t is 20
- Metric value is 48 (left plateau is 27 time intervals wide, the right is 21 time intervals wide)
- Peaks which persist for less than t intervals are NOT counted

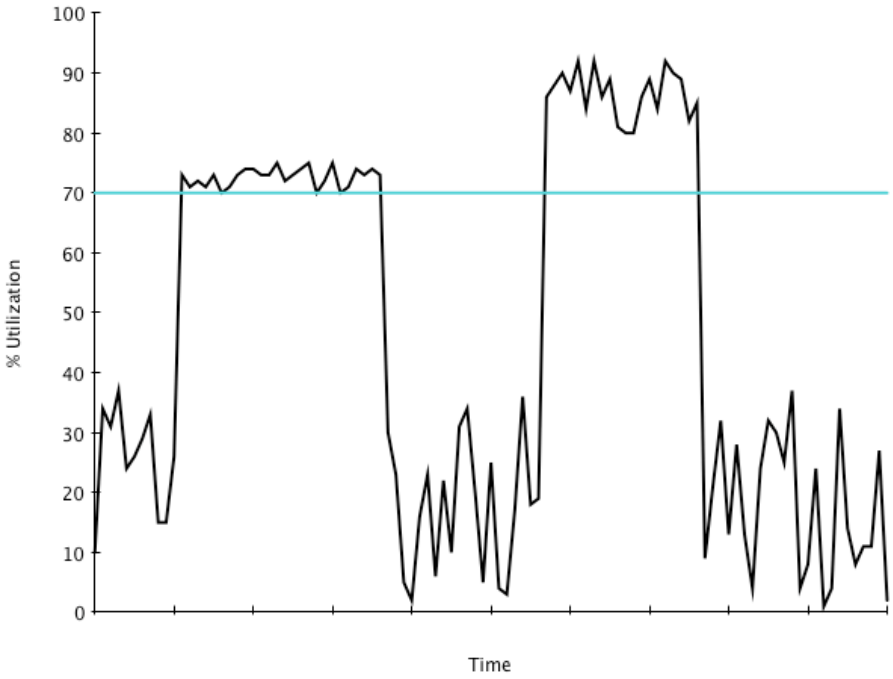

Current Status And Limitations

Status

Peaks and Plateaus metrics are both implemented

In the process of finding appropriate thresholds for each metric

Need more data regarding past upgrades to feed the logistic regression

The peaks and plateaus metrics when properly tuned may provide enough filtering to allow a human to evaluate the remaining circuits

Limitations of the regression

- Topology (both physical and routing) changes can complicate things
- The past may or may not be a good predictor of the future
- Additional metrics would be helpful

Contact Jon Dugan (jdugan@es.net) with questions, comments or suggestions!

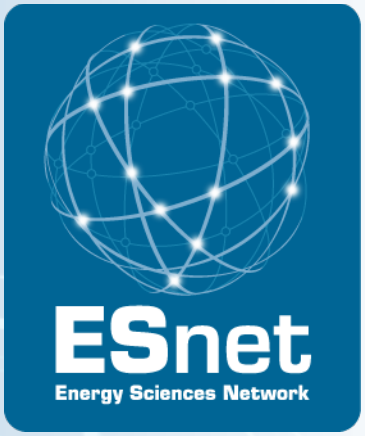

Focus Areas for the Next 12 Months

Finding Ways to Make the Network More Useful

If your scientists aren't effectively using the network, it's unlikely you are going to invest in the network because there are other areas with a better return on investment

- Science productivity and collaboration suffer

How do we get there?

- WAN: Improved efficiency (OSCARS), perfSONAR, security, etc.
- LAN: Science DMZs, perfSONAR, Openflow, data transfer nodes, security, etc.
- New tools for LANs that improve WAN effectiveness, visualization
- What else?

How do we become the Dept of (Less) Energy?

Ideas we're contemplating:

- Begin to understand how to measure the energy cost of moving data using existing infrastructure
  - Extend to 100G network
- Expand http://weathermap.es.net/ to include energy costs
- Test movement of VMs in terms of energy costs
- What is more energy efficient – computing in the cloud or locally? How can this decision be made?
- Use path computation engine SDK to create virtual networks on demand that satisfy not just BW allocation costs but also min energy cost, or max green energy costs
- Analyze the idleness of the paths – is there a chance for power management there? Does it make sense to purposefully create idle paths at a slight performance cost?
- Switch to move flows dynamically to create idleness along a path

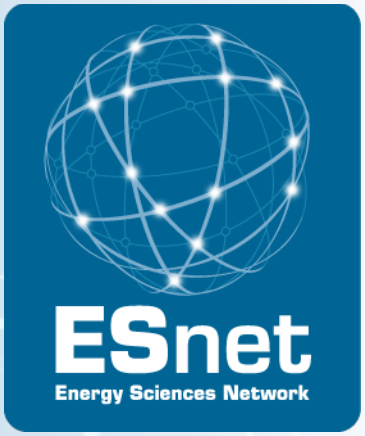

Twitter: ESnetUpdates

Blog: http://esnetupdates.wordpress.com/

Thank you