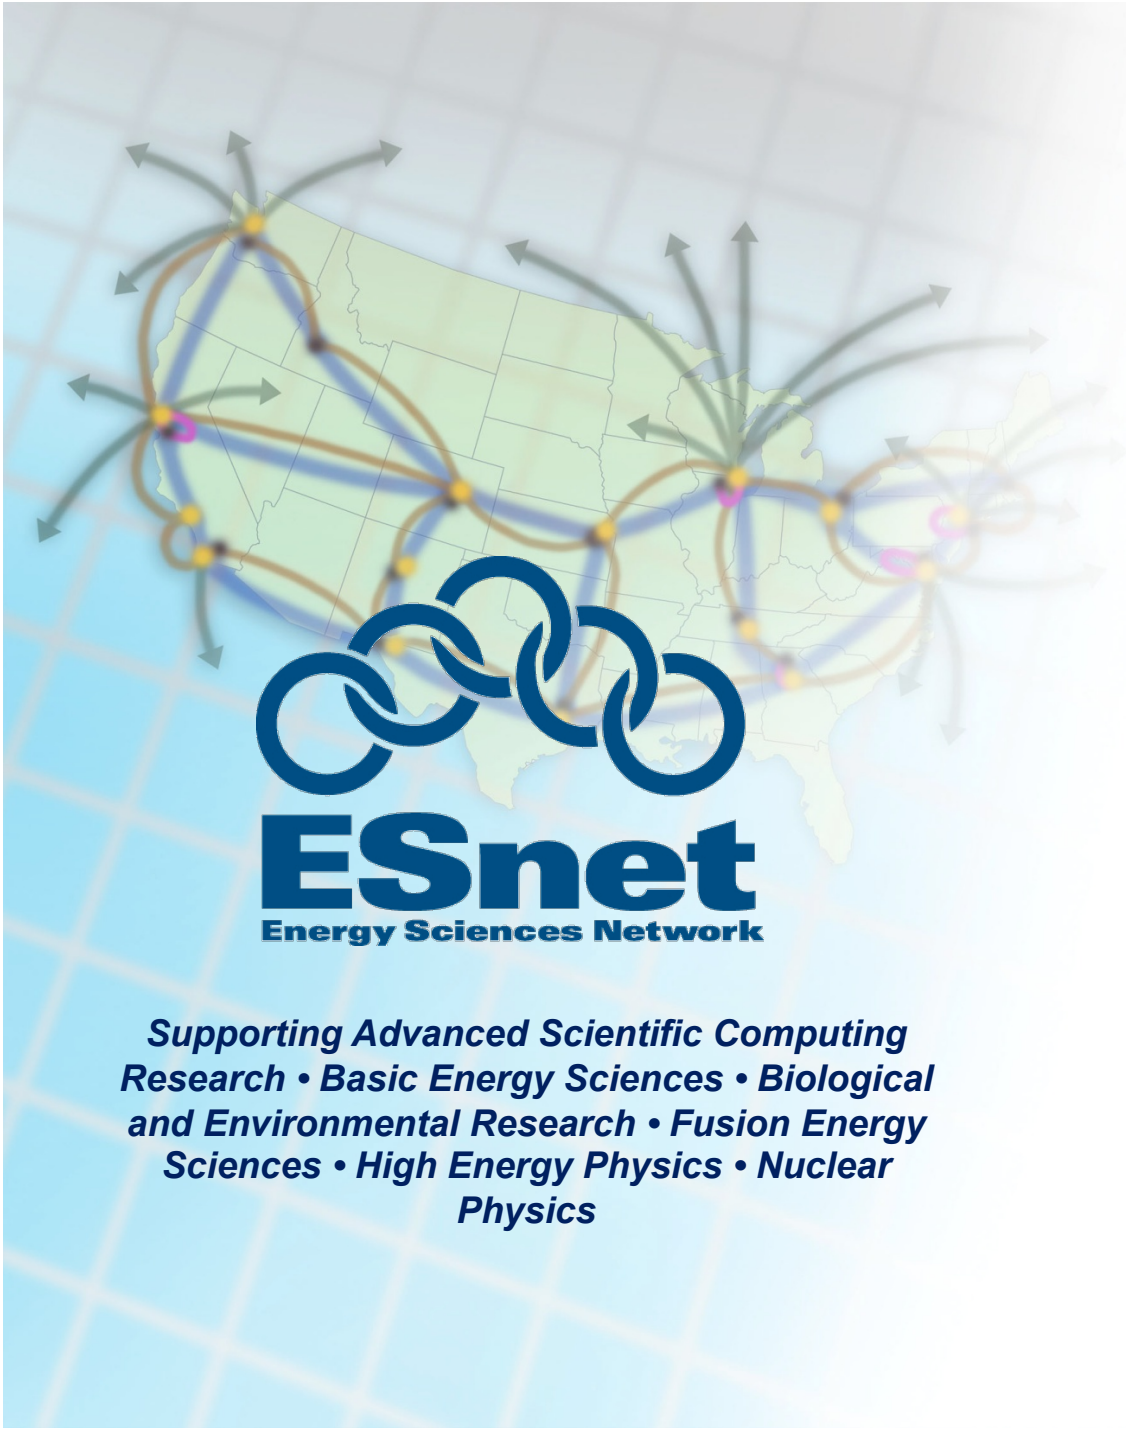

ESnet
Energy Sciences Network

Supporting Advanced Scientific Computing Research • Basic Energy Sciences • Biological and Environmental Research • Fusion Energy Sciences • High Energy Physics • Nuclear Physics

ESnet Update

Steve Cotter, Chin Guok, Joe Metzger, Bill Johnston

Berkeley Lab

U.S. Department of Energy | Office of Science

Agenda

- Network Update
- OSCARS
- perfSONAR
- Federated Trust

ESnet4 – Jan 2009

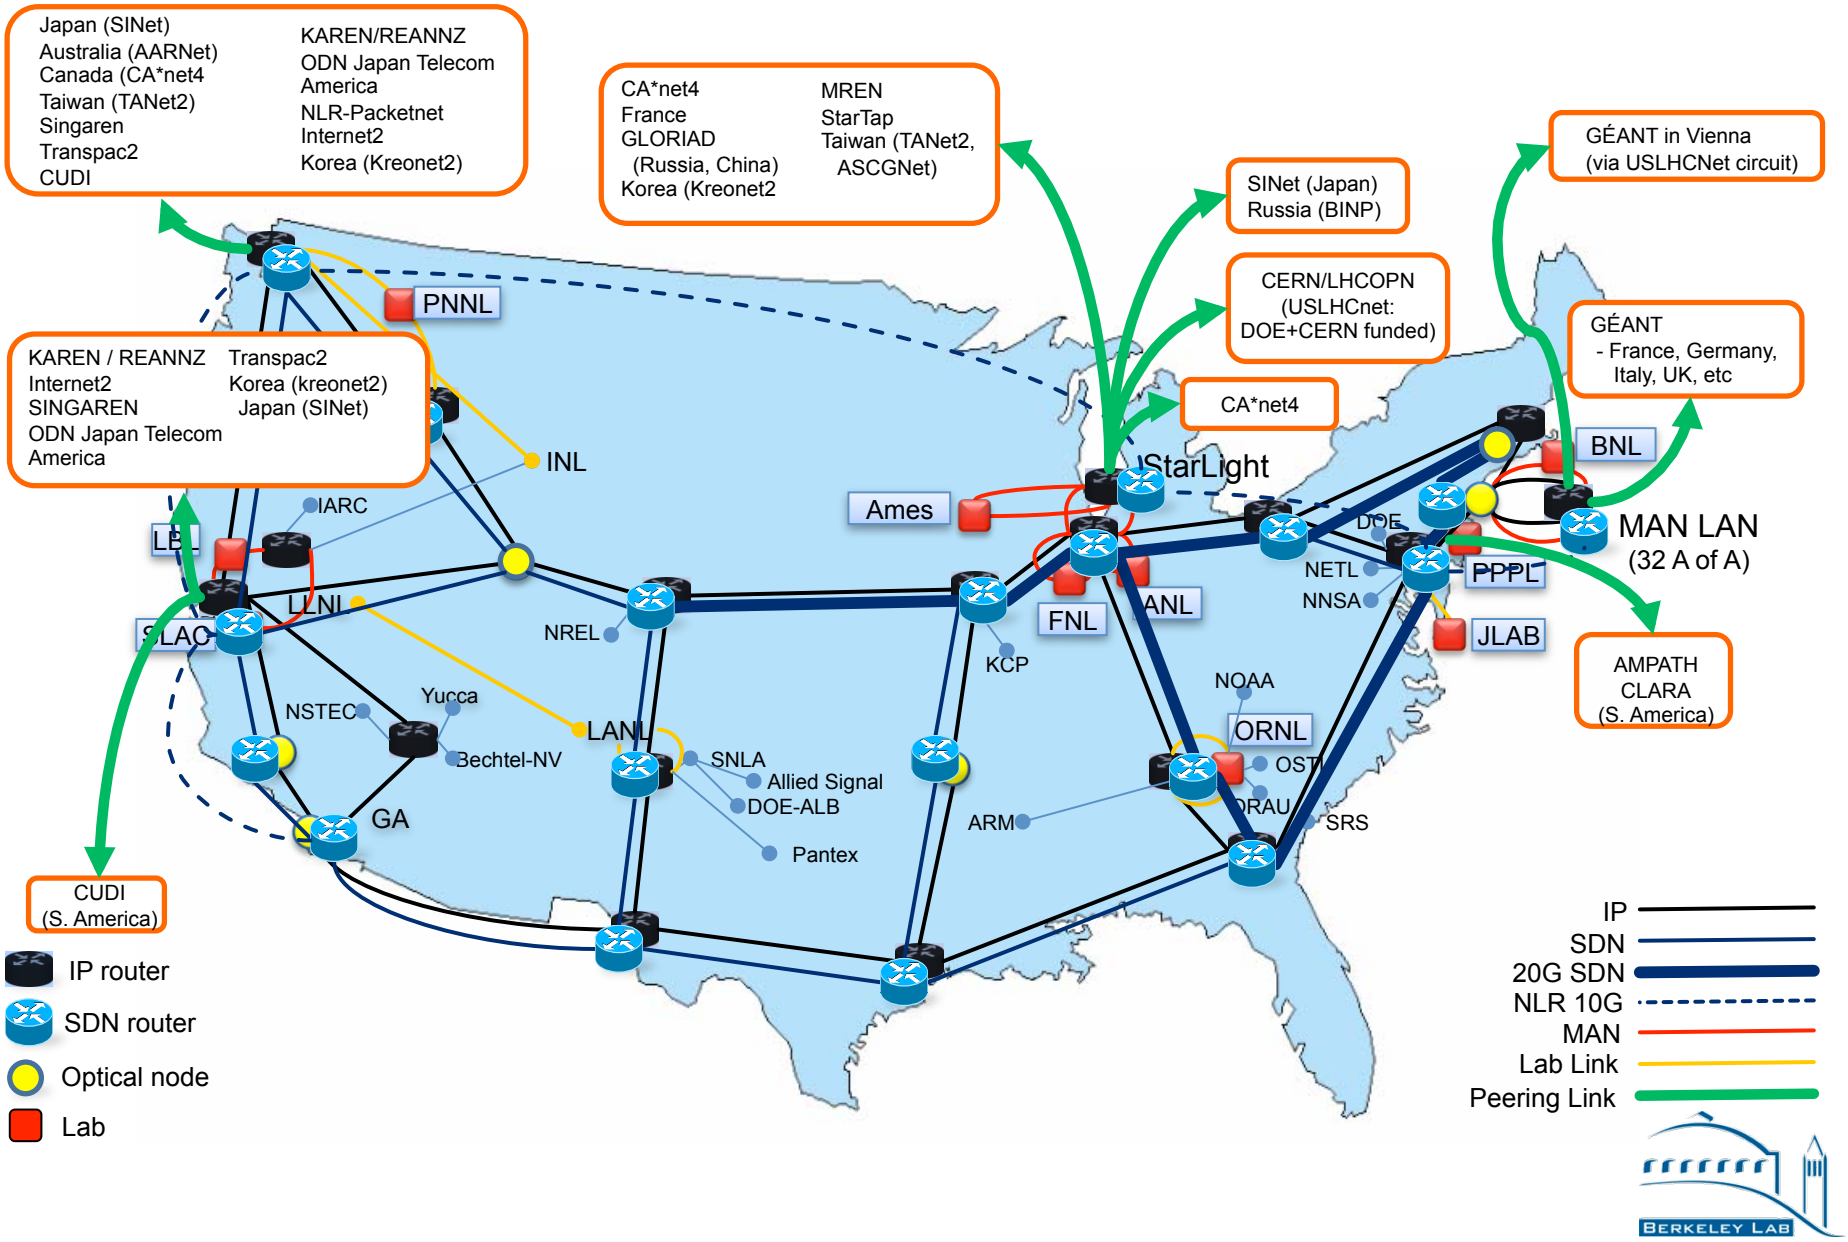

2008 Hub & Wave Install Timeline

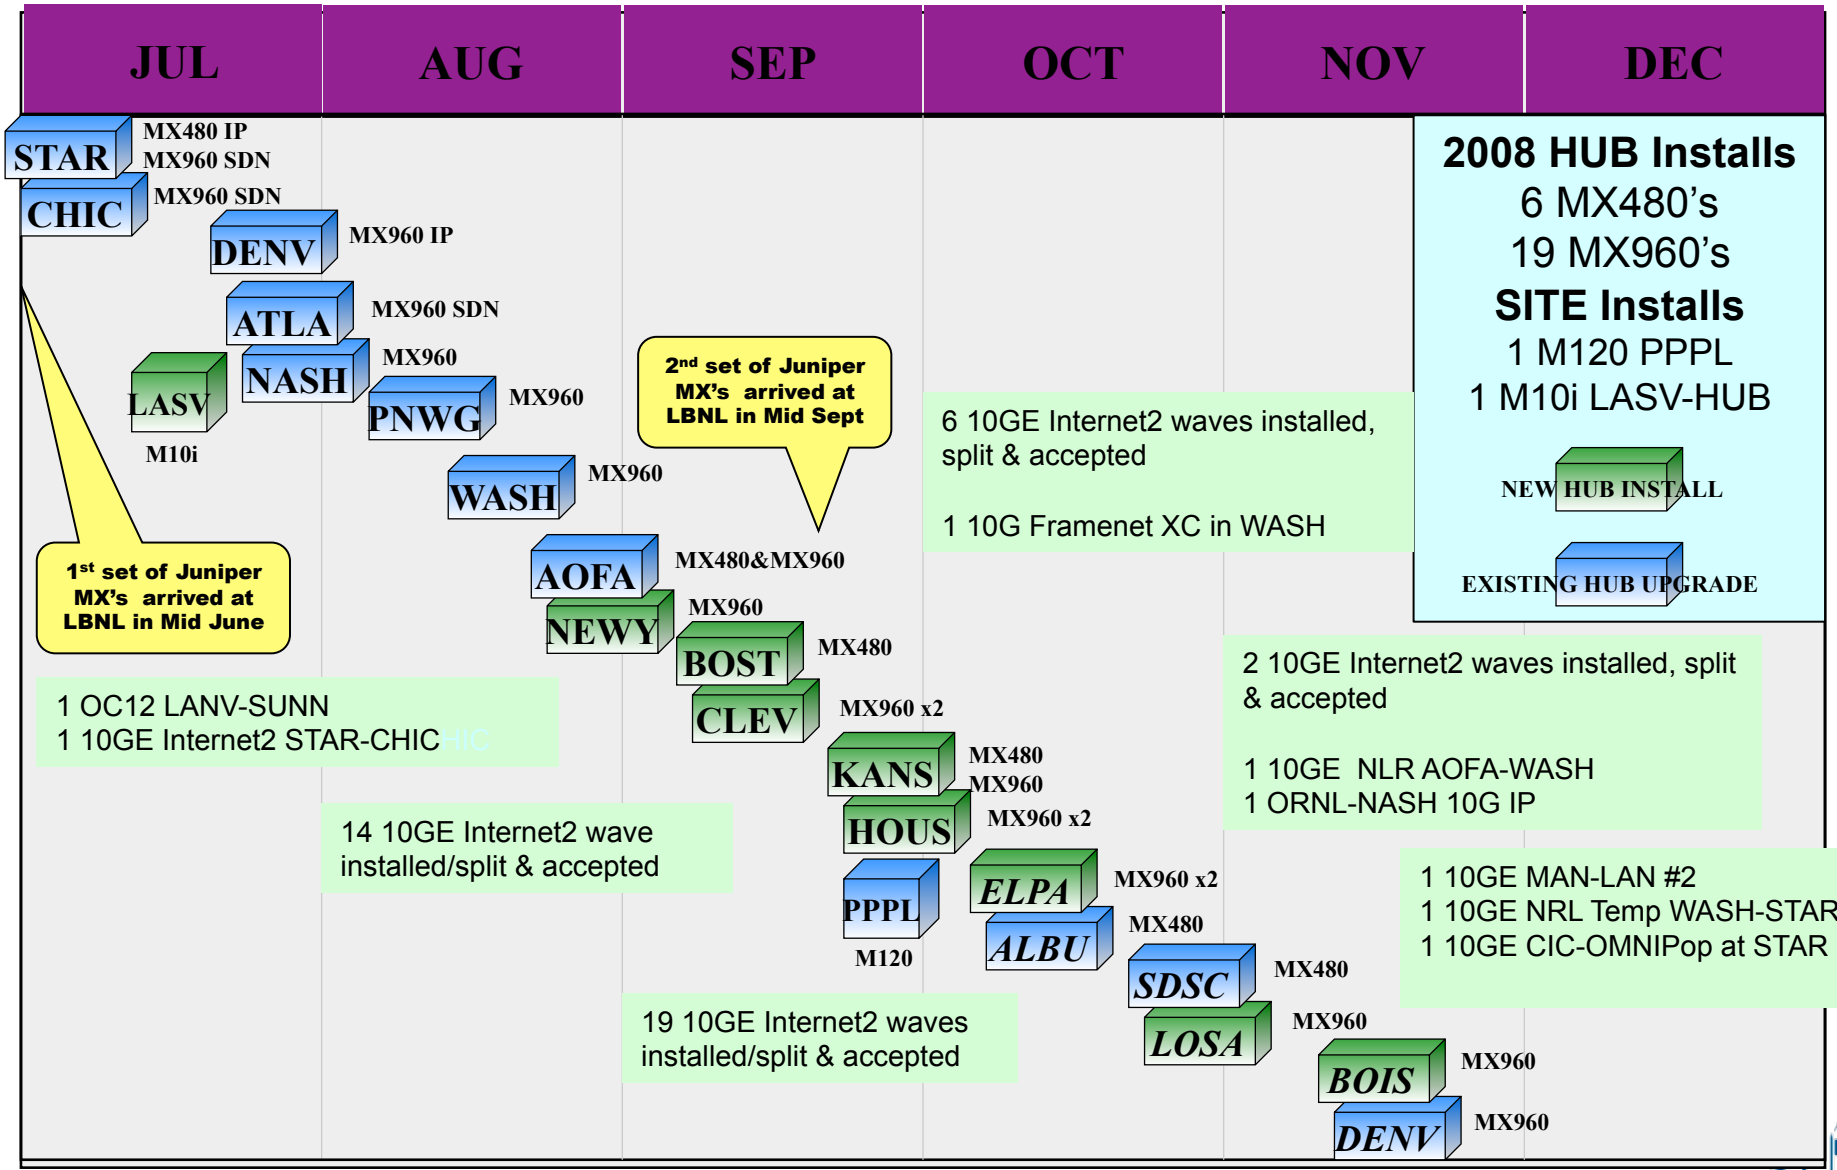

Created by Mike O'Connor Mod by JimG

Hub & Wave Count

Current Hub Count:

- 21 Completed: 32 AofA, NEWY*, WASH, ATLA, NASH, CLEV*, BOST*, CHIC, STAR, KANS*, HOUS*, ELPA*, DENV, ALBU, BOIS*, PNWG, SUNN, SNV(Qwest), LOSA*, SDSC, LASV(SwitchNap)*

  *9 New Hubs since July 2008

Current Backbone Wave Count:

- Internet2 / Level3 Waves:
  - IP Waves: 17 new/split for a total of 25
  - SDN Waves: 25 new/split for a total of 30
- NLR Waves:
  - 1 new wave for a total of 5
  - 1 temp wave (STAR-WASH) for used during NLR northern path upgrade

MAN Upgrades Timeline

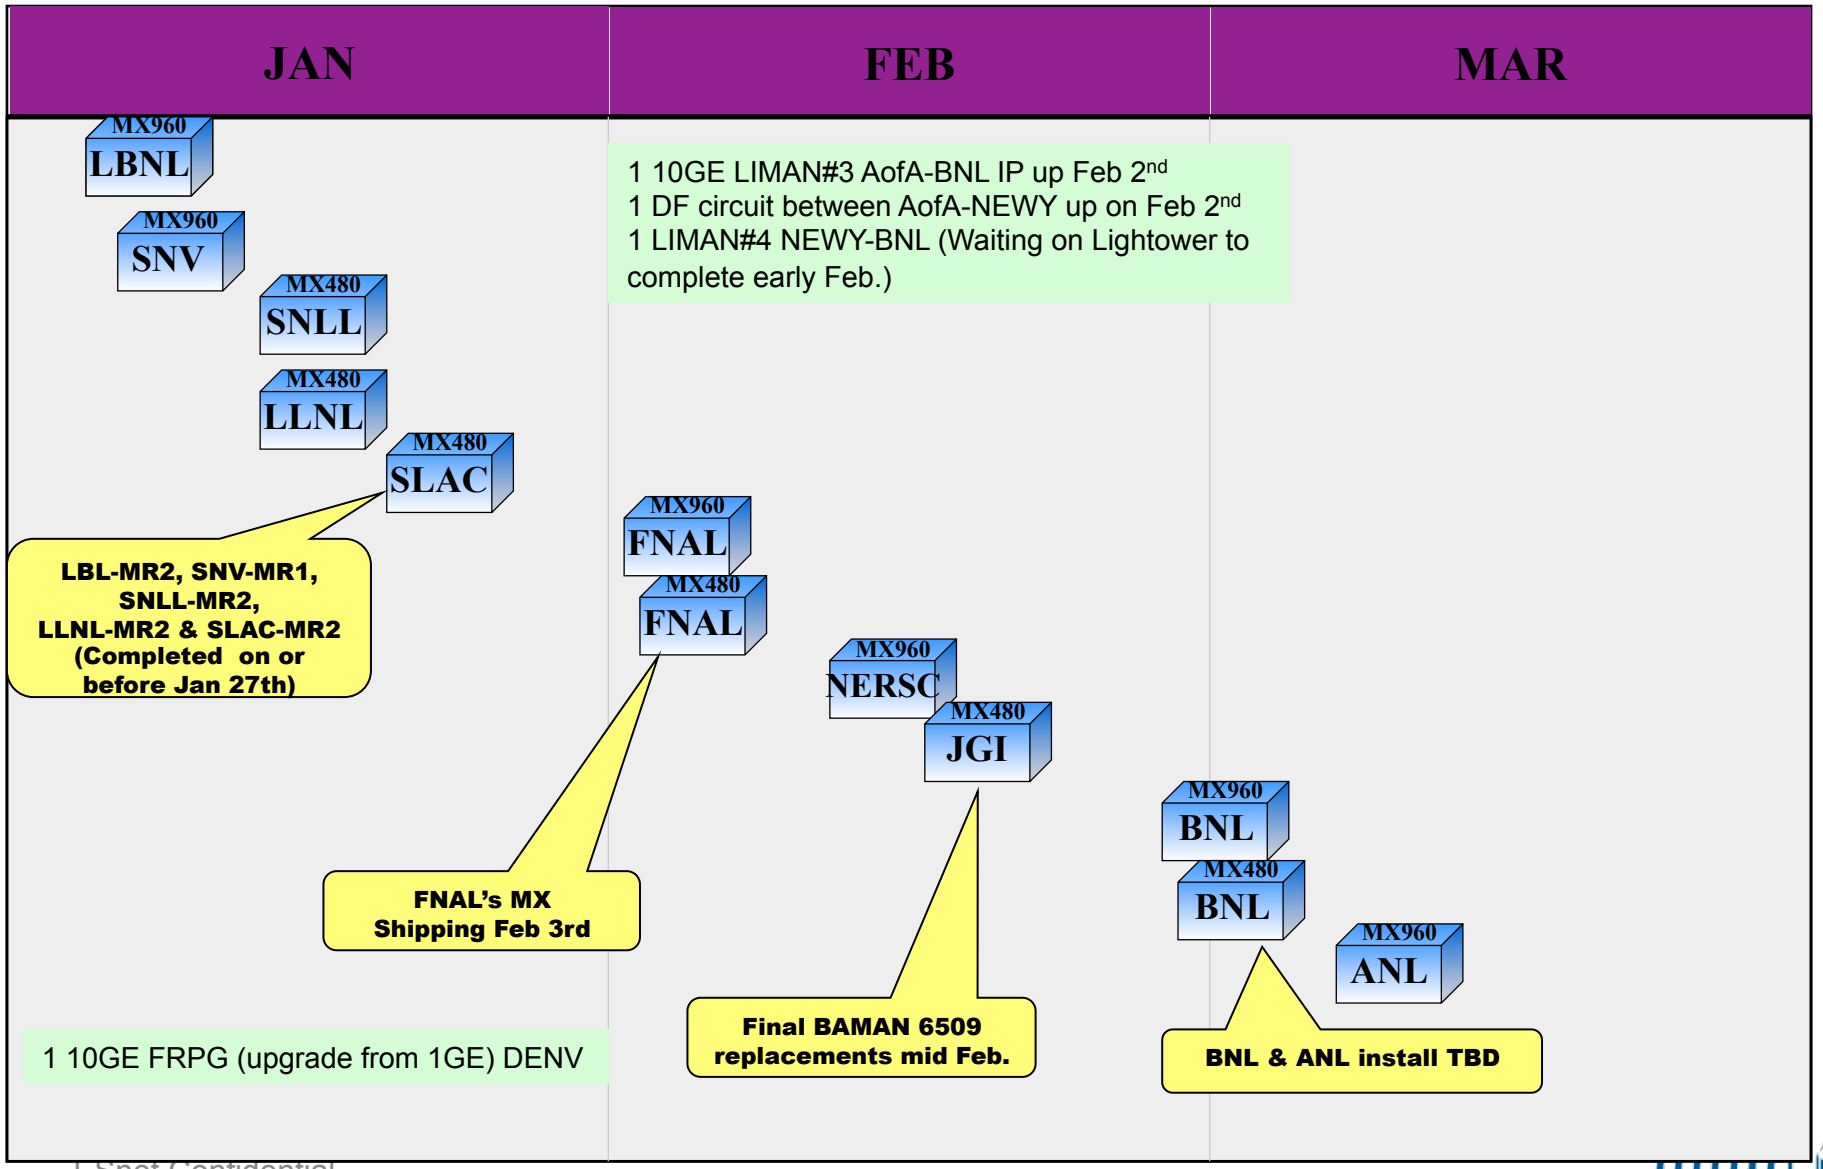

Created by Mike O'Connor Mod by JimG

Active ESnet Links as of 12/2008

Link speed	Description	Count
10 GE	National Core Waves (inter-hub)	61
10 GE	Metropolitan Area Network Circuits (SF Bay Area MAN, Chicago MAN, Long Island MAN)	33
10 GE	Circuits to ESnet sites	24
10 GE	Circuits to R&E peering points	24
	Total 10G WAN circuits	125
10 GE	Intra-hub connections (interconnecting ESnet equipment at the network hubs)	78
5 GE	GÉANT peering in Vienna, Austria (via USLHCNet and GÉANT circuits)	1
OC-192 SONET	special	1
OC-48 SONET	ORNL backup circuit	
1 GE	Mostly small site and commercial peering connections	83
Misc. slower	Mostly non-SC sites	64

Future Installs

- Replace site 6509s (FNAL, ANL & BNL) with MXs
  - FNAL (MX960 & MX480) shipped on Feb 3rd for site to install
  - BNL (MX960 & MX480) shipping & Install TBD
  - ANL (MX960) shipping & Install TBD
- Replace BAMAN 6509s with MXs
  - LBNL-MR3 (MX960), SNV-MR2 (MX960), LLNL-MR2 (MX480) & SNLL-MR2 (MX480) completed prior to Jan 22nd
  - SLAC-MR2 (MX480) Completed on Jan 27th
  - NERSC-MR2 & JGI-MR2 installs scheduled for Mid Feb.
- Future Circuits installs
  - New 10 G LIMAN wave & DF AOFA-NEWY End-2-end on Feb 2nd & #4 wave to BNL (Feb)
  - OC-12 between LASV hub and General Atomic (Feb)
  - 10 GE between BOST hub to MIT (Feb)
  - OC-12 between DENV hub and Pantex (TBD)
  - 1 GE wave in BOIS to INL via IRON (TBD)
  - 10 GE SDN wave between PNWG hub to PNNL (TBD)
  - 10 GE SDN wave between NASH hub to ORNL (TBD)

ESnet4 Metro Area Rings

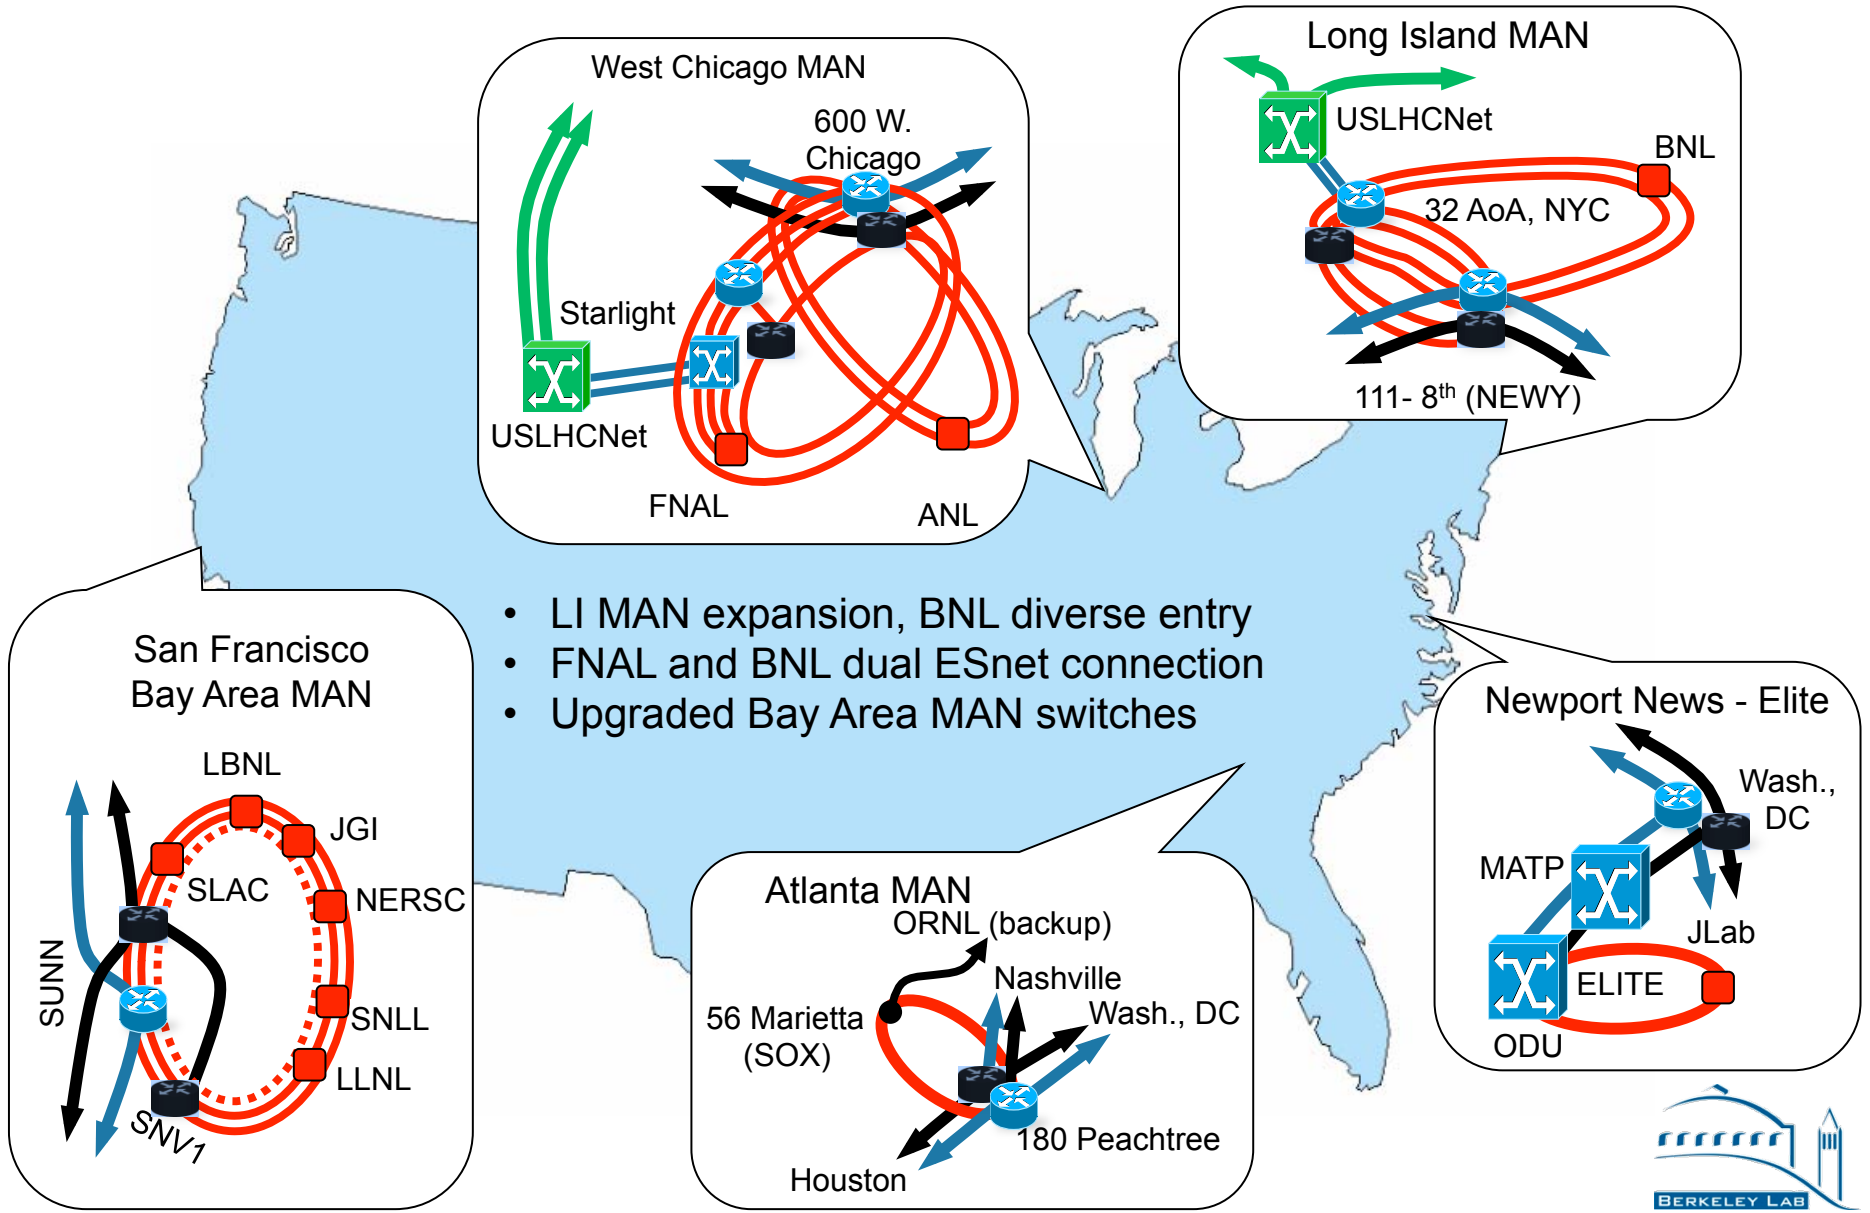

- LI MAN expansion, BNL diverse entry
- FNAL and BNL dual ESnet connection
- Upgraded Bay Area MAN switches

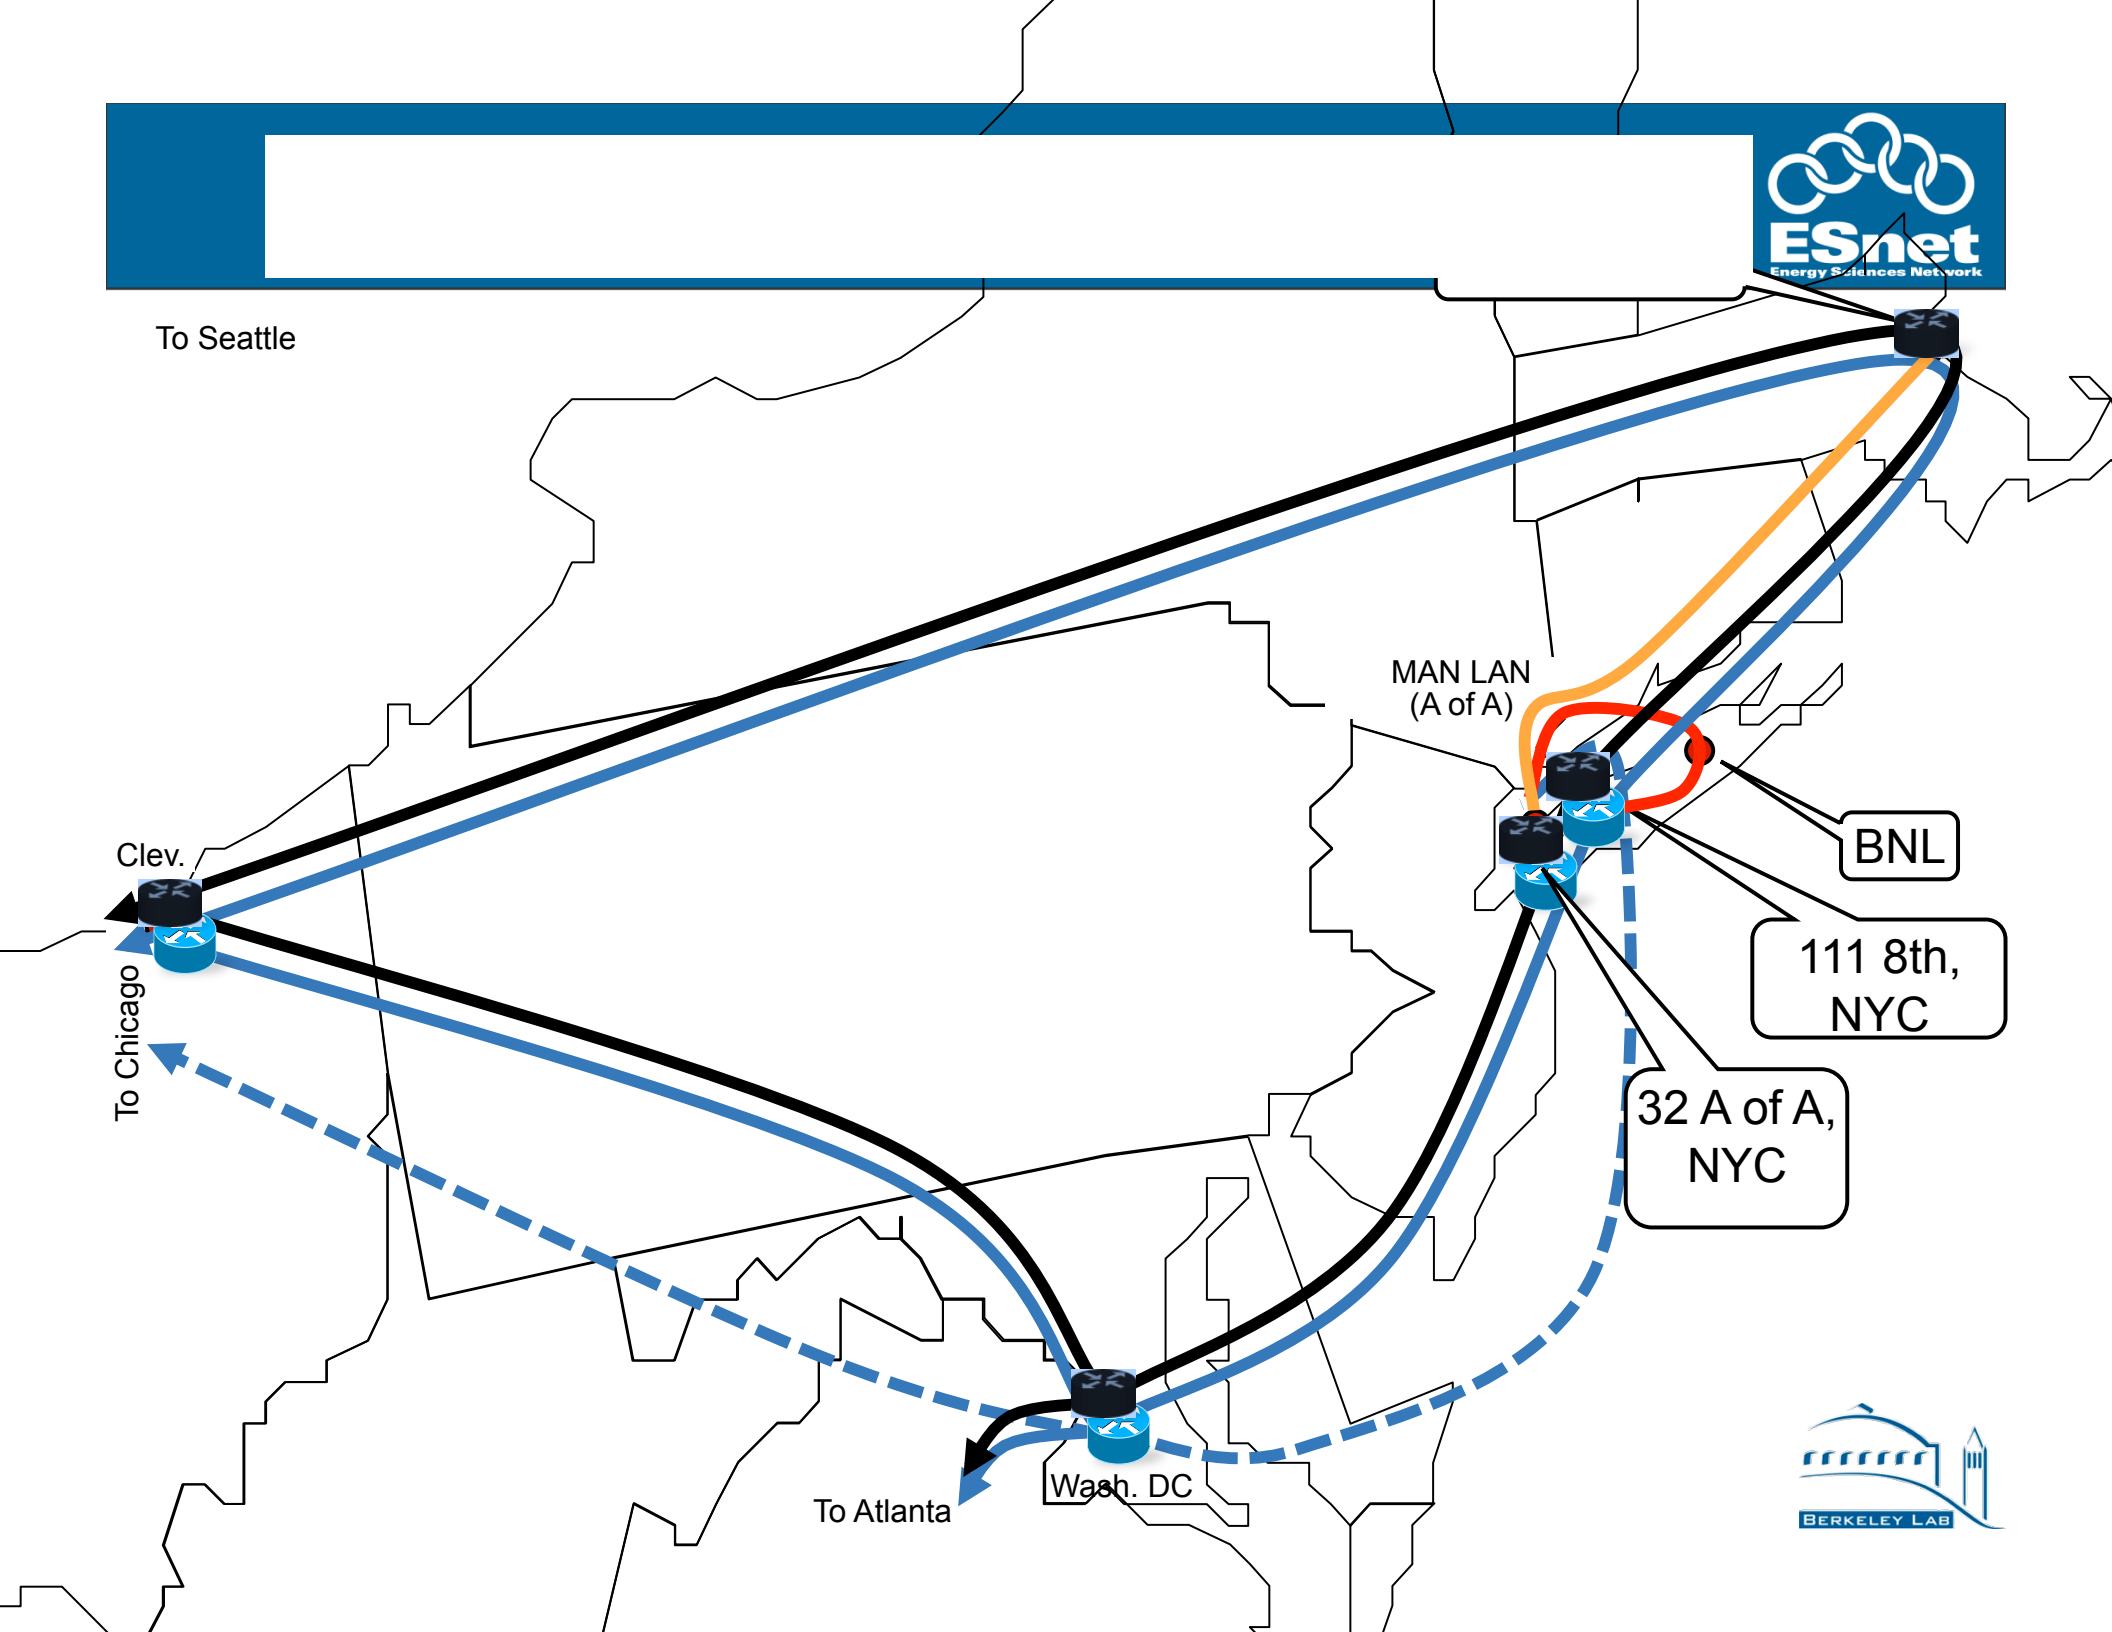
ESnet
Energy Sciences Network
To Seattle
MAN LAN
(A of A)
Clev.
BNL
111 8th,
NYC
32 A of A,
NYC
To Chicago
Wash. DC
To Atlanta
Berkeley Lab

Tier1 Redundancy: Long Island

Notes:

1) There are physically independent paths from R1 to Boston and from R2 to Washington
2) Only fiber paths are shown, wave counts are not indicated
3) The engineering and procurement for this configuration are complete, implementation is underway
4) An architecturally similarly situation is also being implemented for FNAL / Chicago

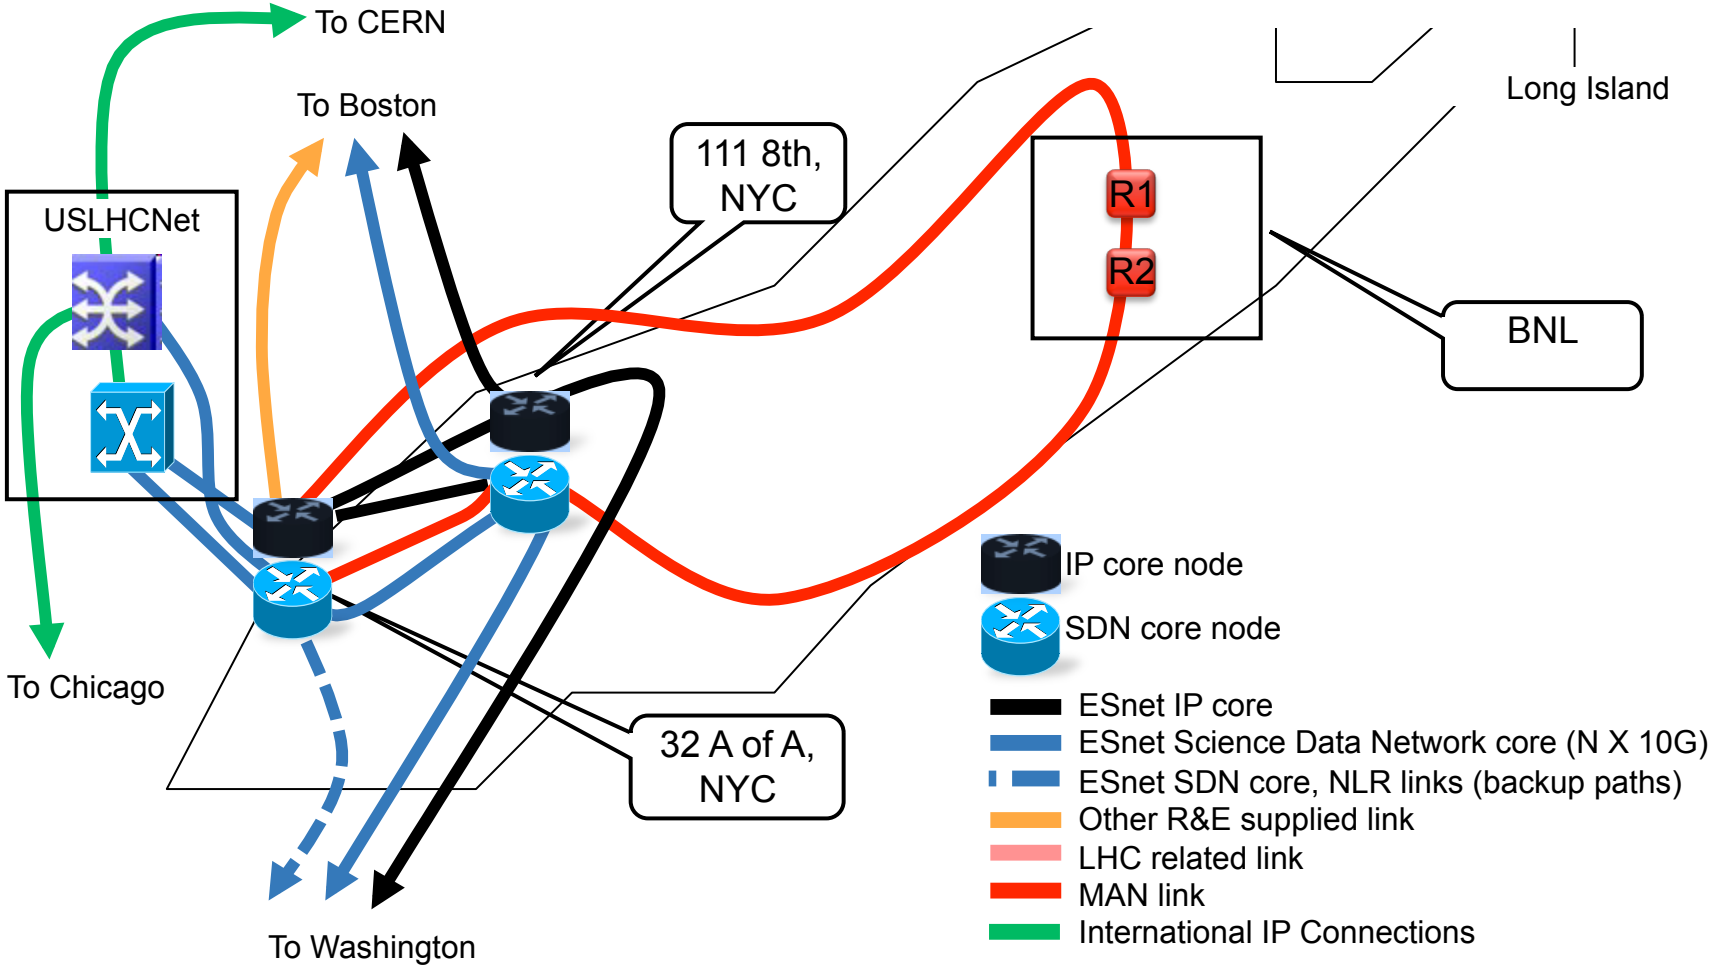

12 Month Circuit Availability 1/2009

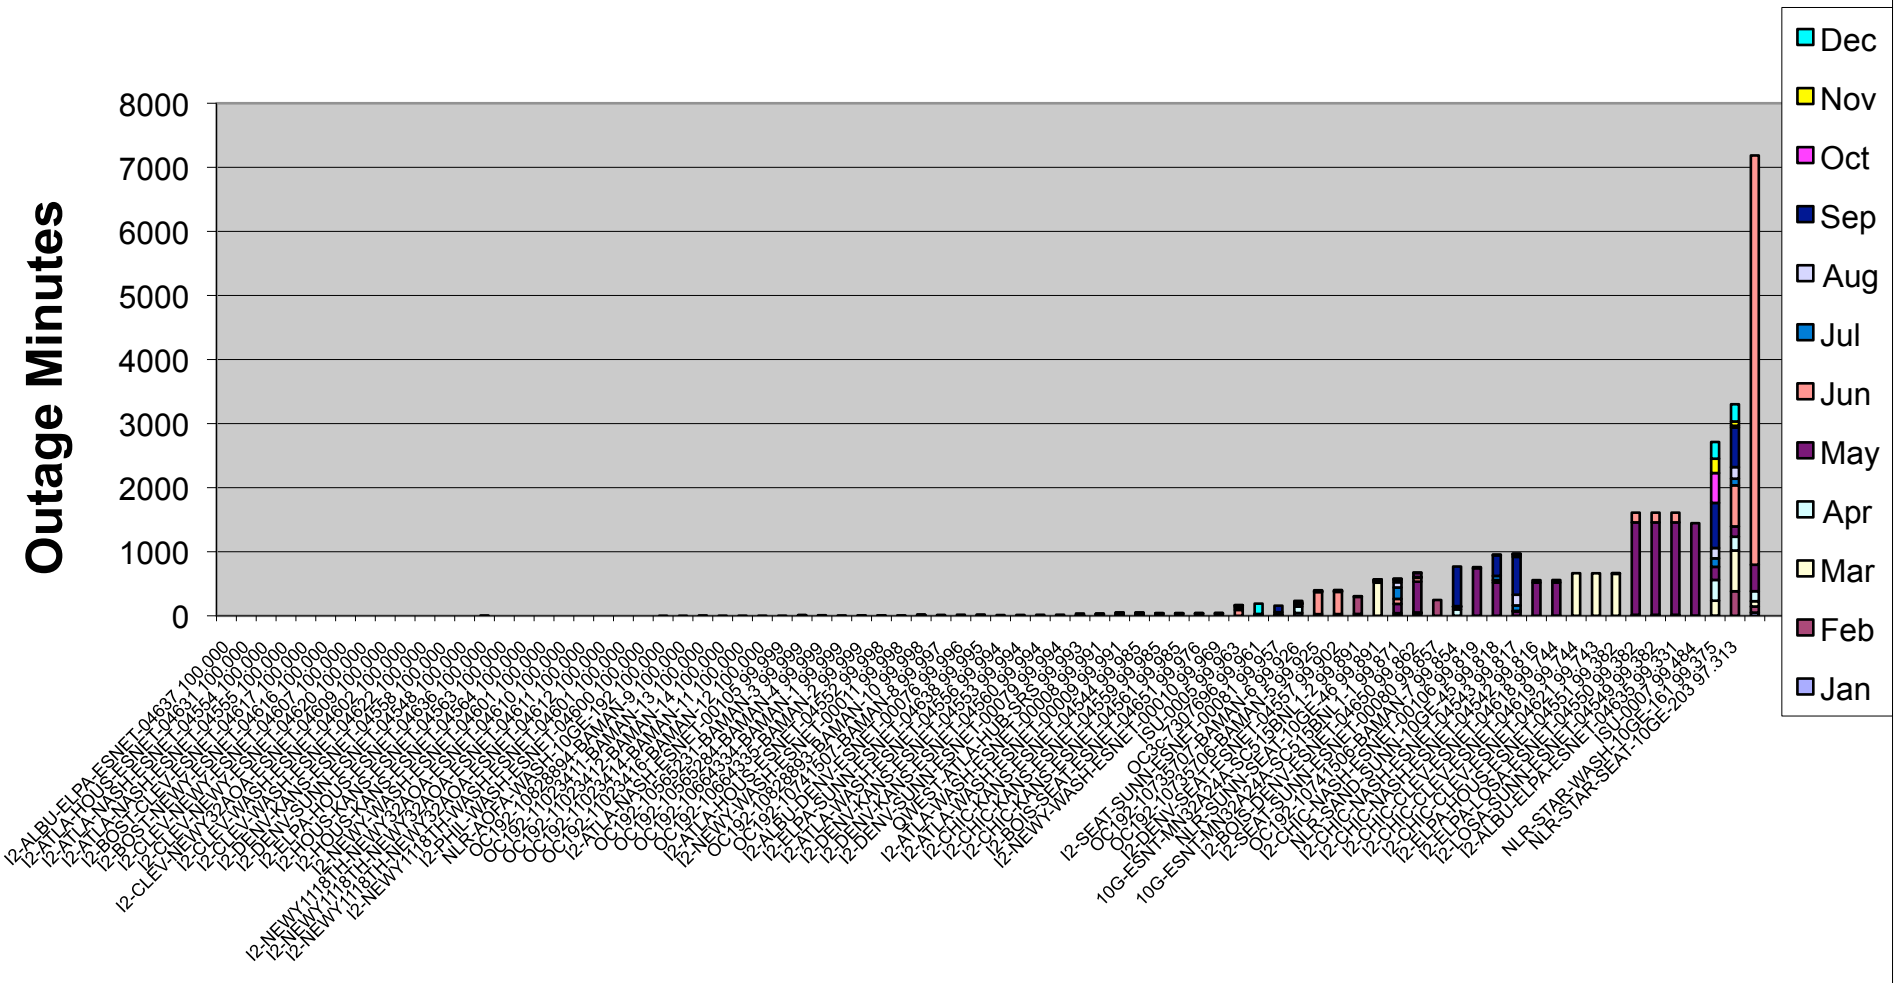

Improved Site Availability

Site Availability 2/2006 to 1/2007

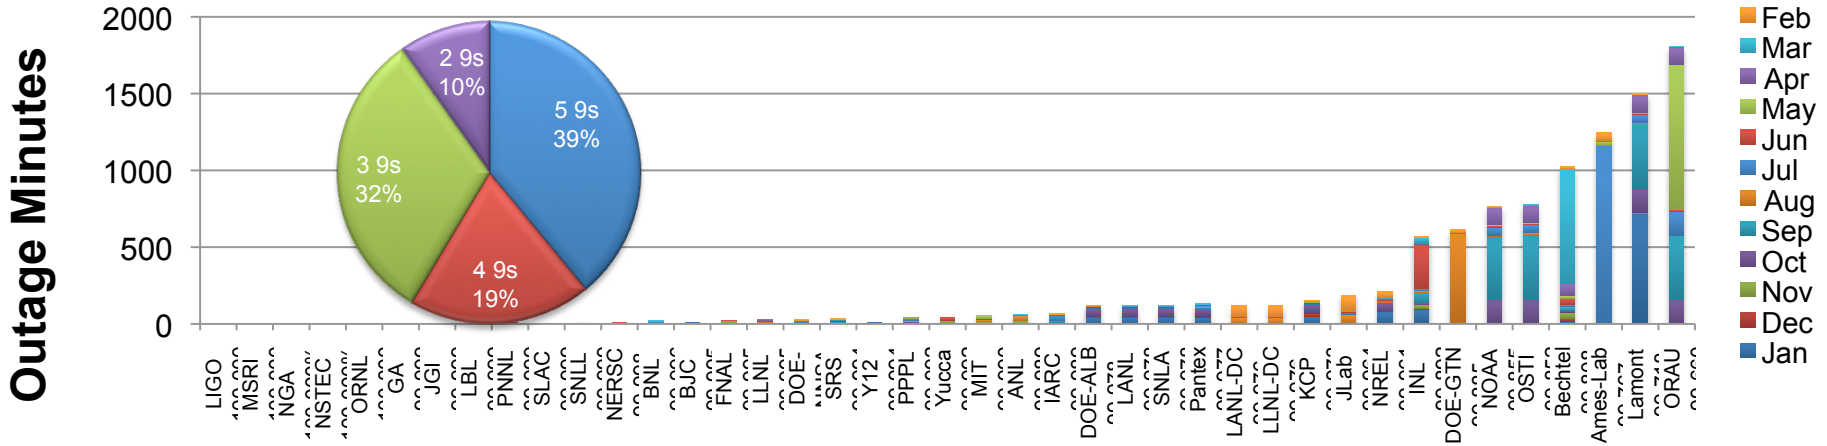

Site Availability 2/2008 to 1/2009

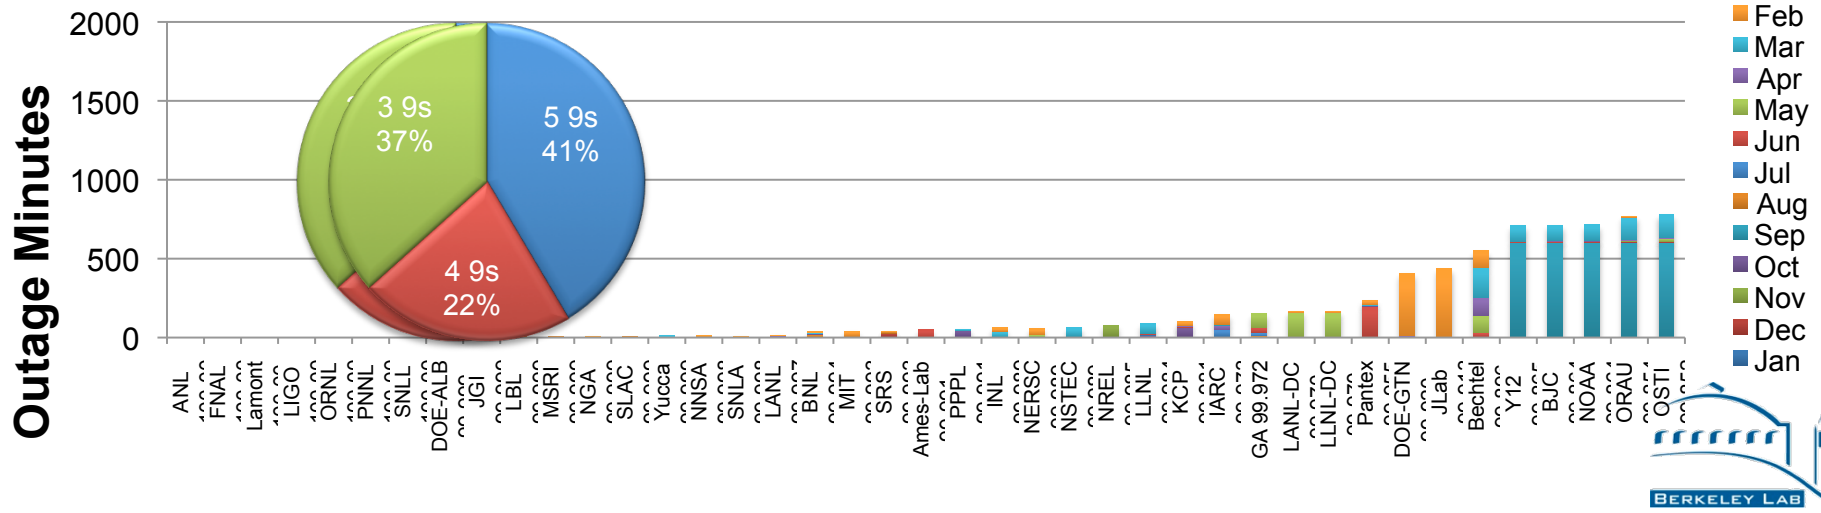

ESnet Accepted Traffic (Tby/mo)

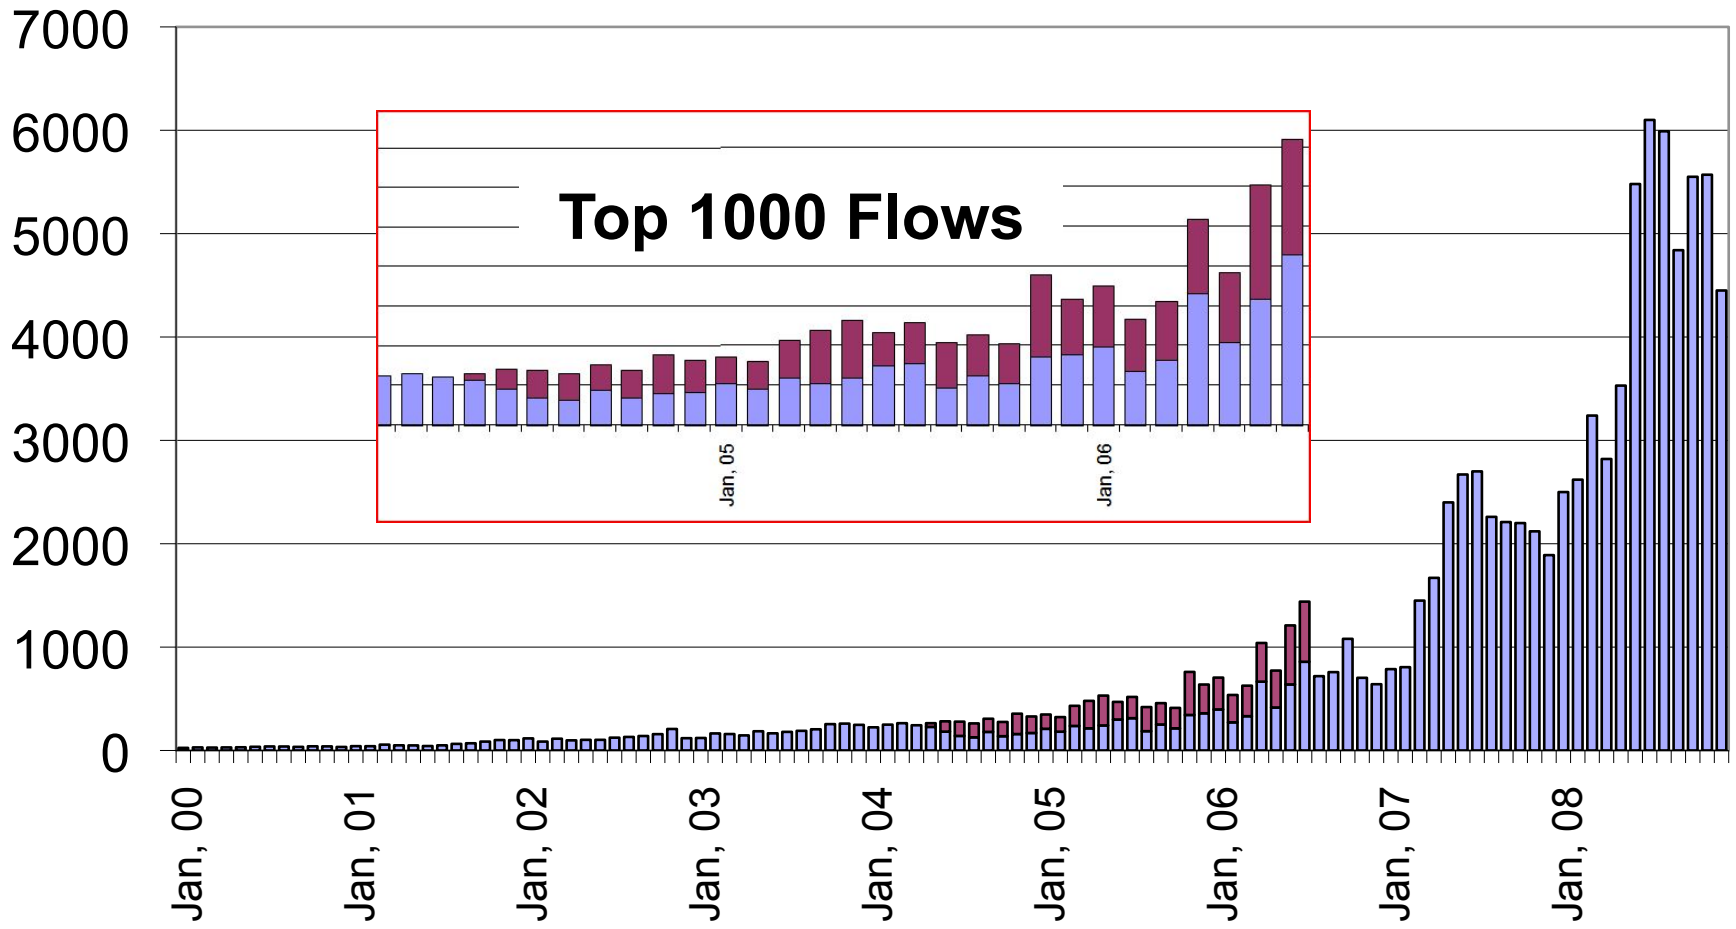

Historical ESnet Traffic Patterns

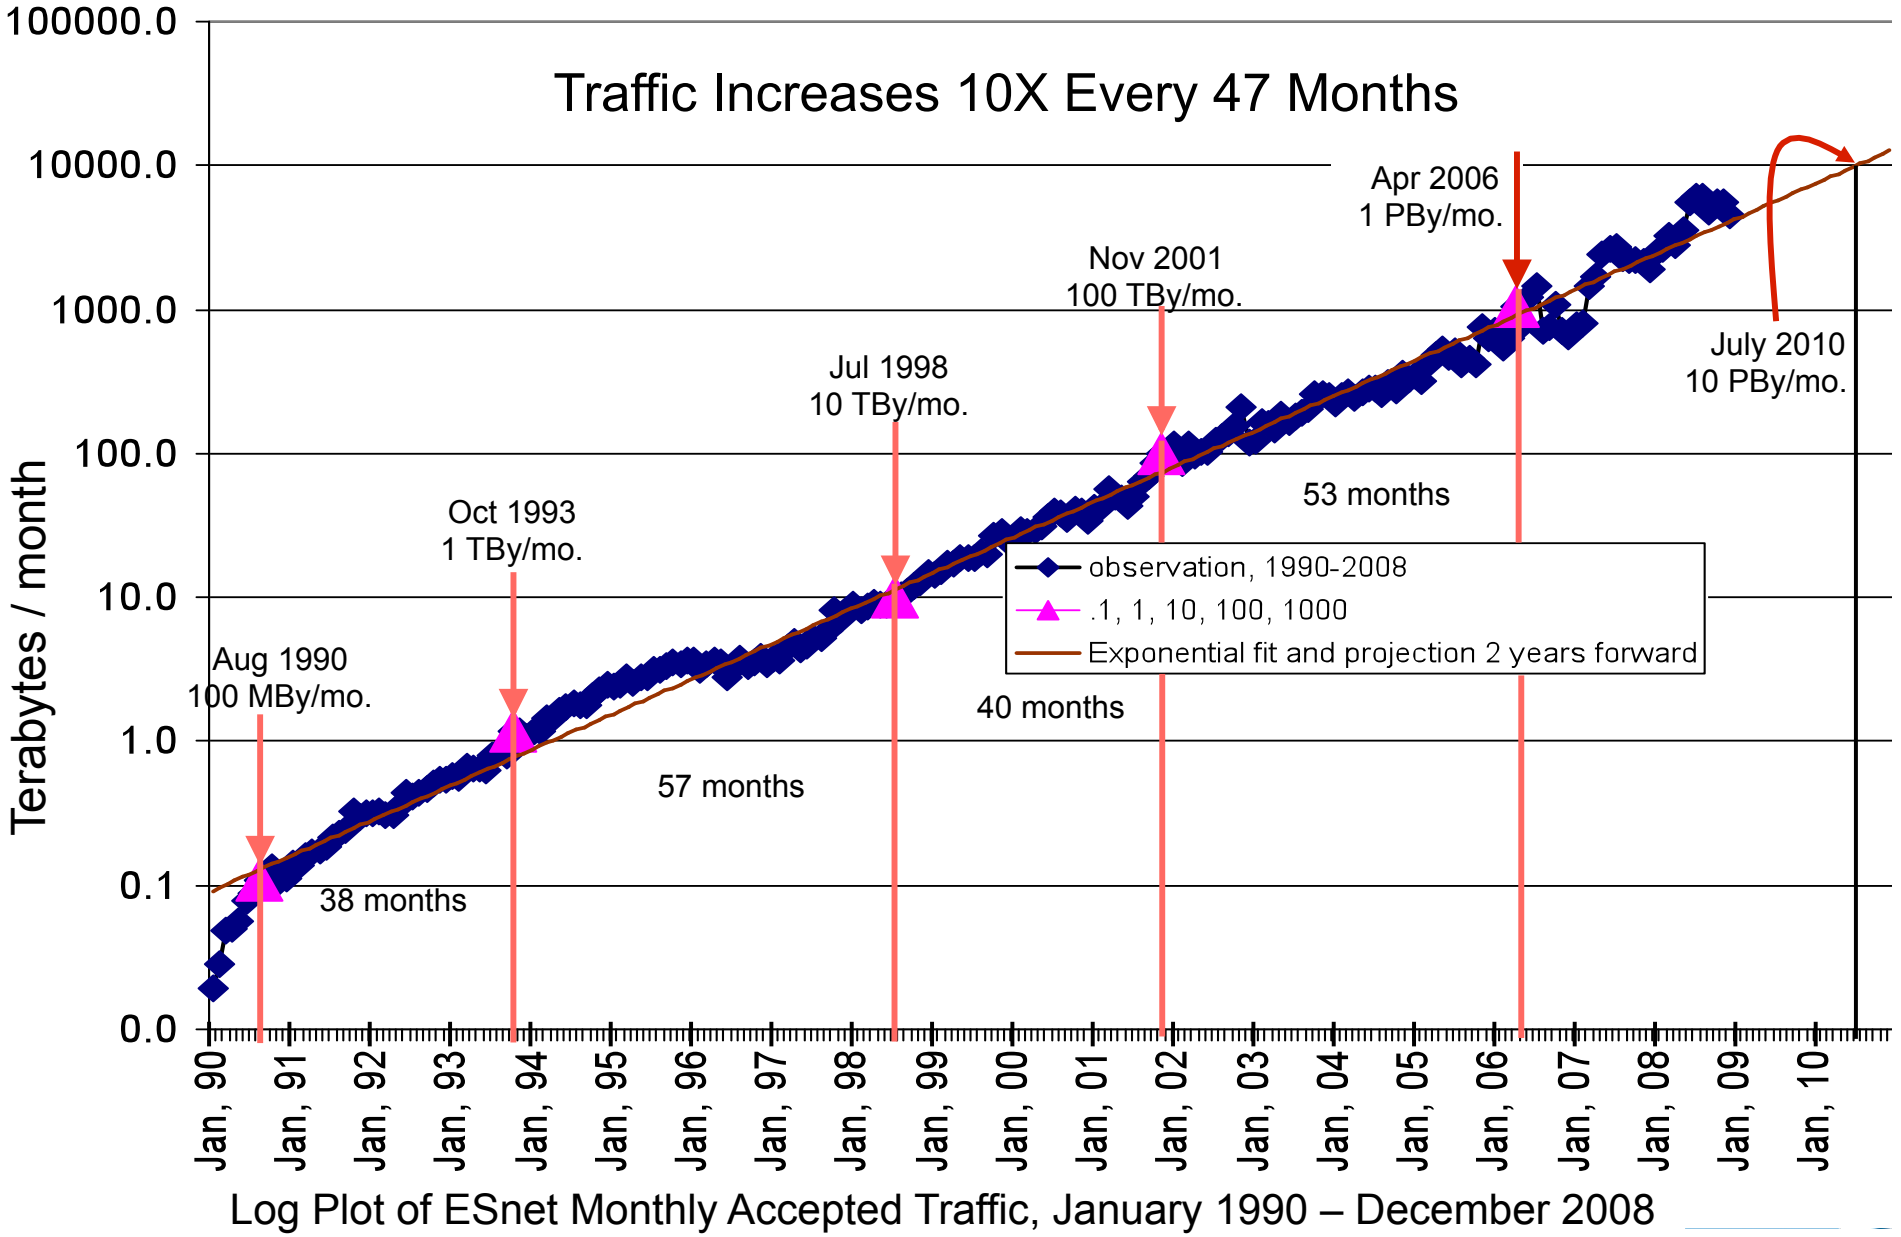

Network Traffic, Science Data, and Network Capacity

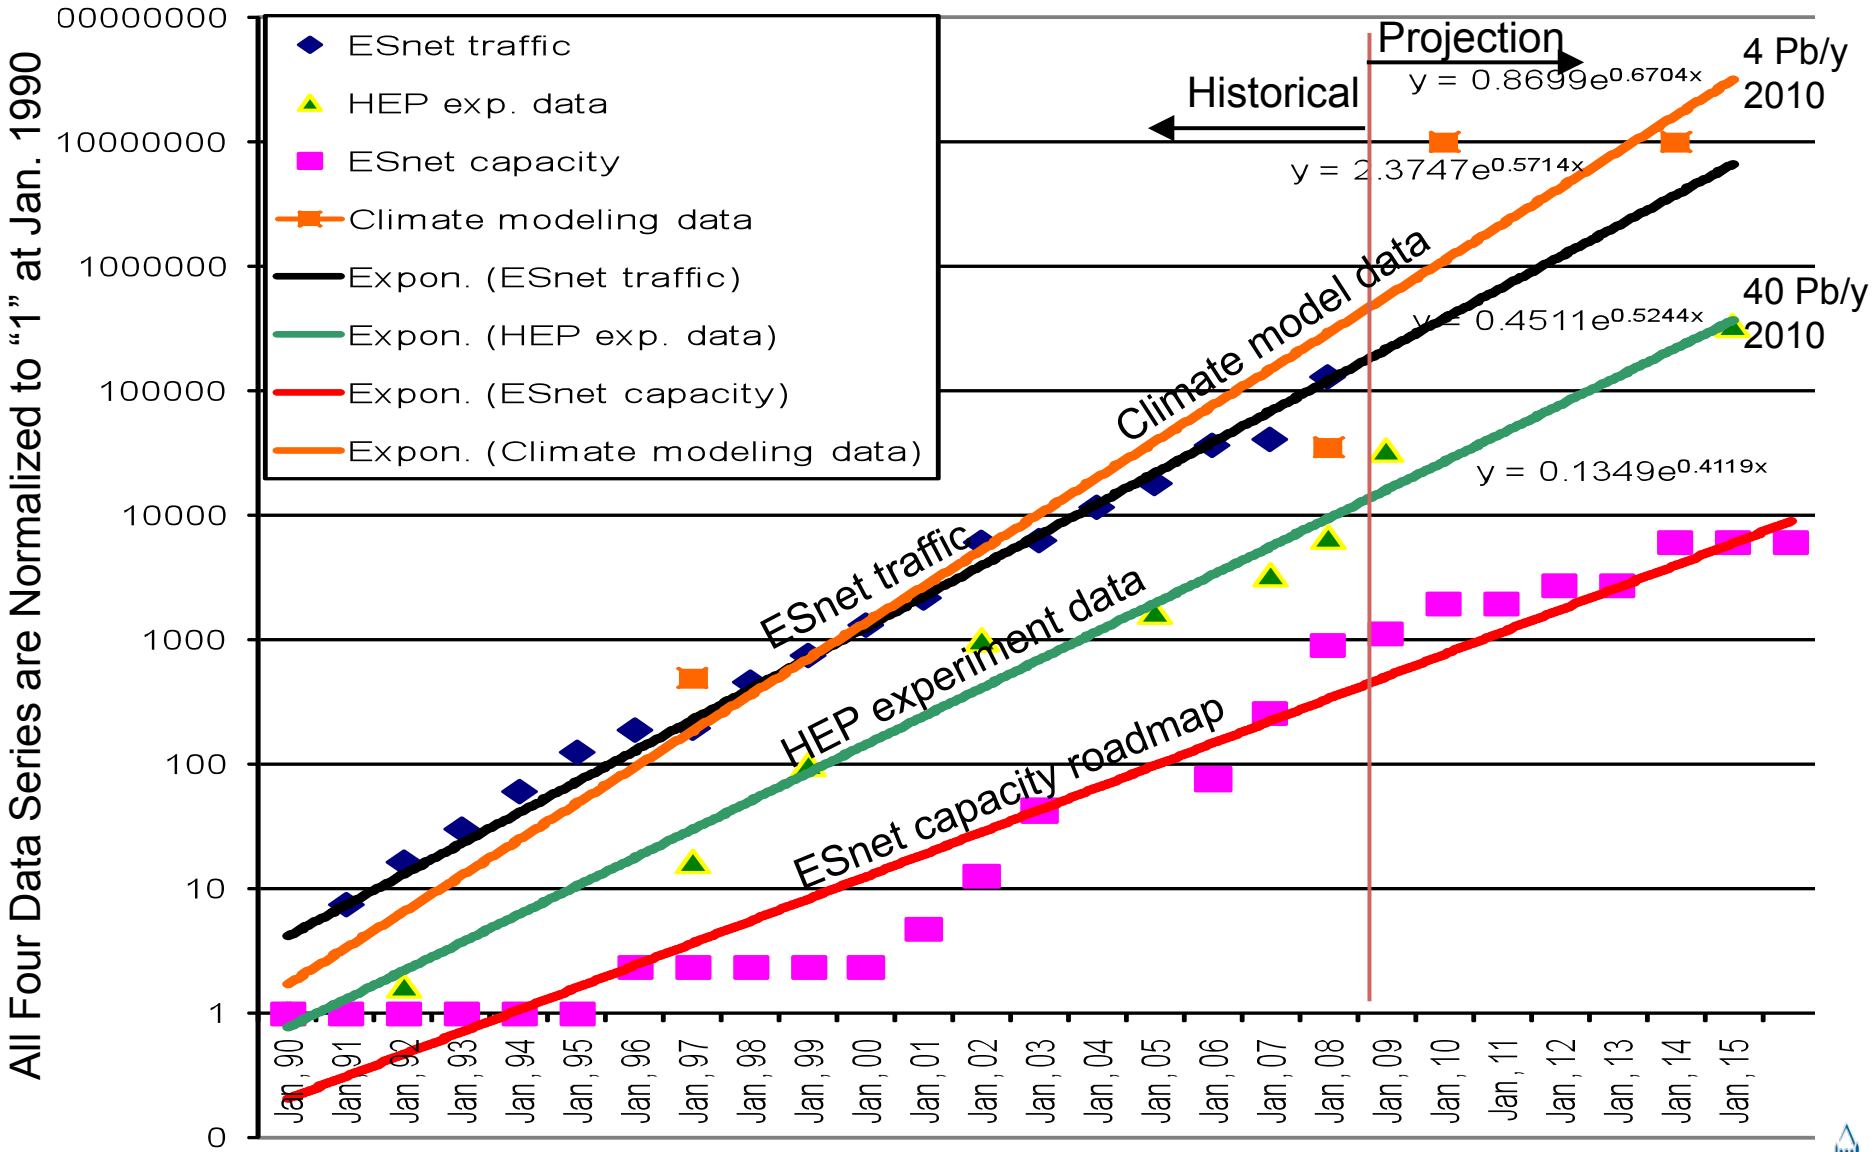

(HEP data courtesy of Harvey Newman, Caltech, and Richard Mount, SLAC. Climate data courtesy Dean Williams, LLNL, and the Earth Systems Grid Development Team.)

ESnet4 – 2010

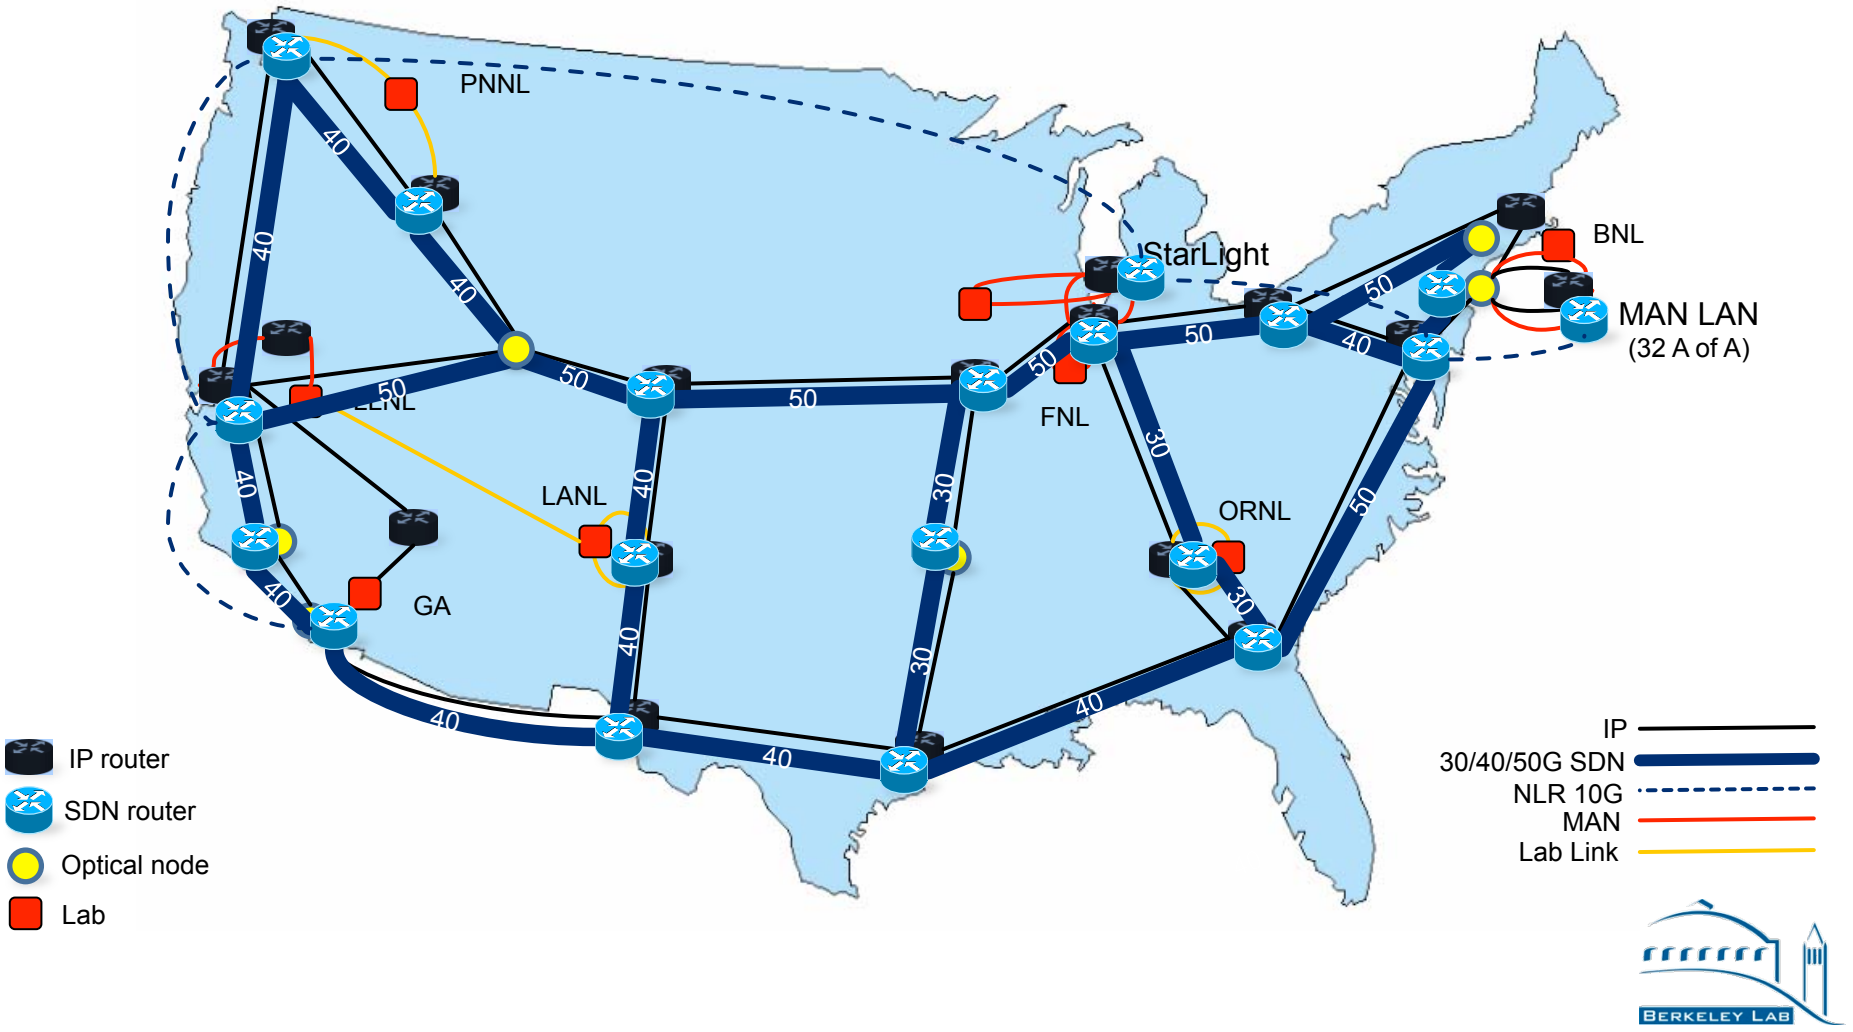
PNNL
StarLight
BNL
MAN LAN
(32 A of A)
FNL
LANL
ORNL
GA
40
50
30
IP router
SDN router
Optical node
Lab
IP
30/40/50G SDN
NLR 10G
MAN
Lab Link

Beyond 2010: 100 G

- ESnet4 planning assumes technology advances will provide 100 Gb/s optical waves (they are 10 Gb/s now)
- The ESnet4 SDN switching/routing platform (Juniper MX960) is designed to support new 100 Gb/s network interfaces
- With capacity planning based on the ESnet 2010 wave count, we can probably assume some fraction of the core network capacity by 2012 will require 100 Gb/s interfaces

- ESnet is involved in a collaboration with Internet2, Juniper Networks (core routers), Infinera (DWDM), and Level3 (network support) to accelerate its deployment and help drive down the cost of 100G components

ESnet Security & Disaster Recovery

- Advances in security at ESnet over the last 6 months:
  - Implemented Two-factor authentication for ESnet network engineers requesting privileged access to the network management plane. Reviewed and re-defined access to network management plane.
  - Upgraded Bro Intrusion Detection System
- ESnet Security Peer Review – Feb 11-12
  - Fed/R&E/Commercial experts reviewing ESnet security practices and procedures
- Disaster recovery improvements
  - Deployed Government Emergency Telecommunications Service (GETS) numbers to key personnel
  - Deploying full replication of the NOC databases and servers and Science Services databases in the NYC Qwest carrier hub

Website Redesign

- Goals
  - Better organization of information, easier navigation, searchable (not everything in pdfs) but don't want it to all be 'push'
  - Collaborative tool – upload best practices, video from conference, community calendar, staff pages
  - Integration of business processes into site
    - "My ESnet" portal for site coordinators / users
    - Exploring Google Earth or similar network visualization
      - IP / SDN / MAN representation
      - perfSONAR performance data
      - OSCARS virtual circuit status
  - Looking for ideas/input/suggestions.

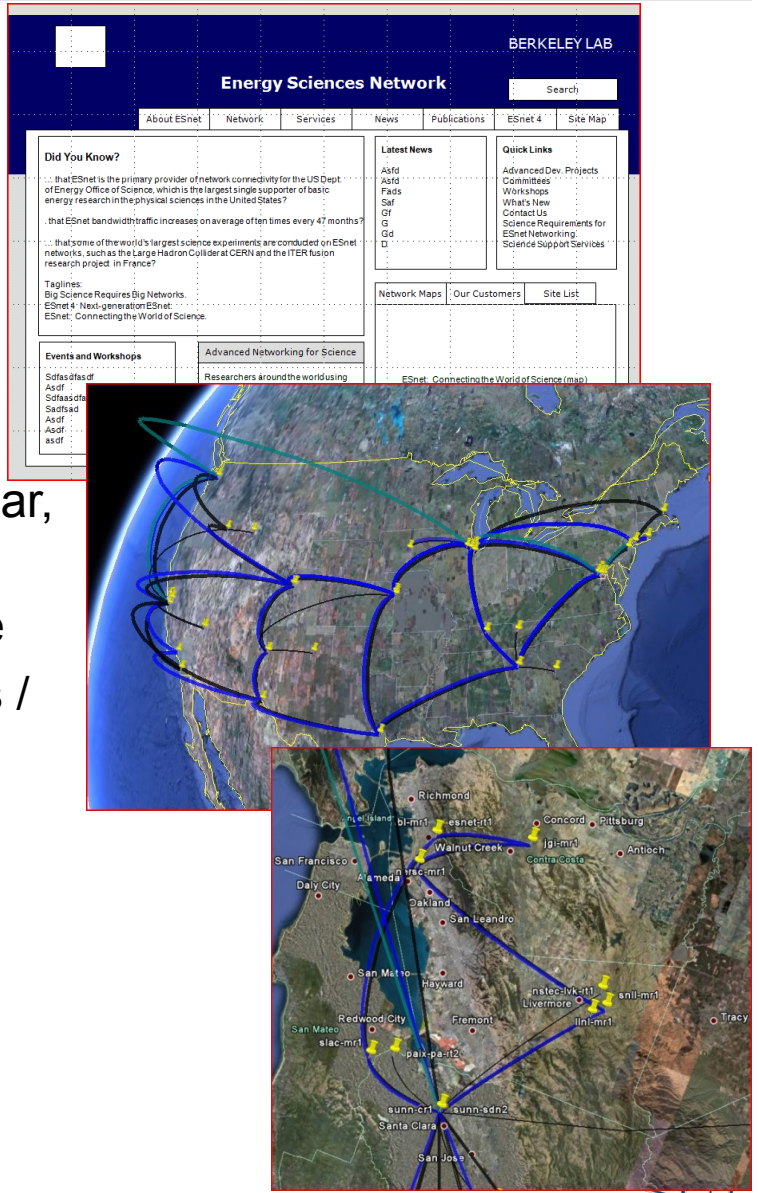

Agenda

- Network Update
- OSCARS
- perfSONAR
- Federated Trust

Multi-Domain Virtual Circuit Service
OSCARS

The OSCARS service requirements:

- Guaranteed bandwidth with resiliency
  - User specified bandwidth - requested and managed in a Web Services framework
  - Explicit backup paths can be requested
- Traffic isolation
  - Allows for high-performance, non-standard transport mechanisms that cannot co-exist with commodity TCP-based transport
- Traffic engineering (for ESnet operations)
  - Enables the engineering of explicit paths to meet specific requirements
    - e.g. bypass congested links; using higher bandwidth, lower latency paths; etc.
- Secure connections
  - The circuits are “secure” to the edges of the network (the site boundary) because they are managed by the control plane of the network which is highly secure and isolated from general traffic
- End-to-end, cross-domain connections between Labs and collaborating institutions

OSCARS Current (v0.5) Implementation

- Well defined inter-module interfaces
- Exchange of static topology information
- PCE integrated into OSCARS Core

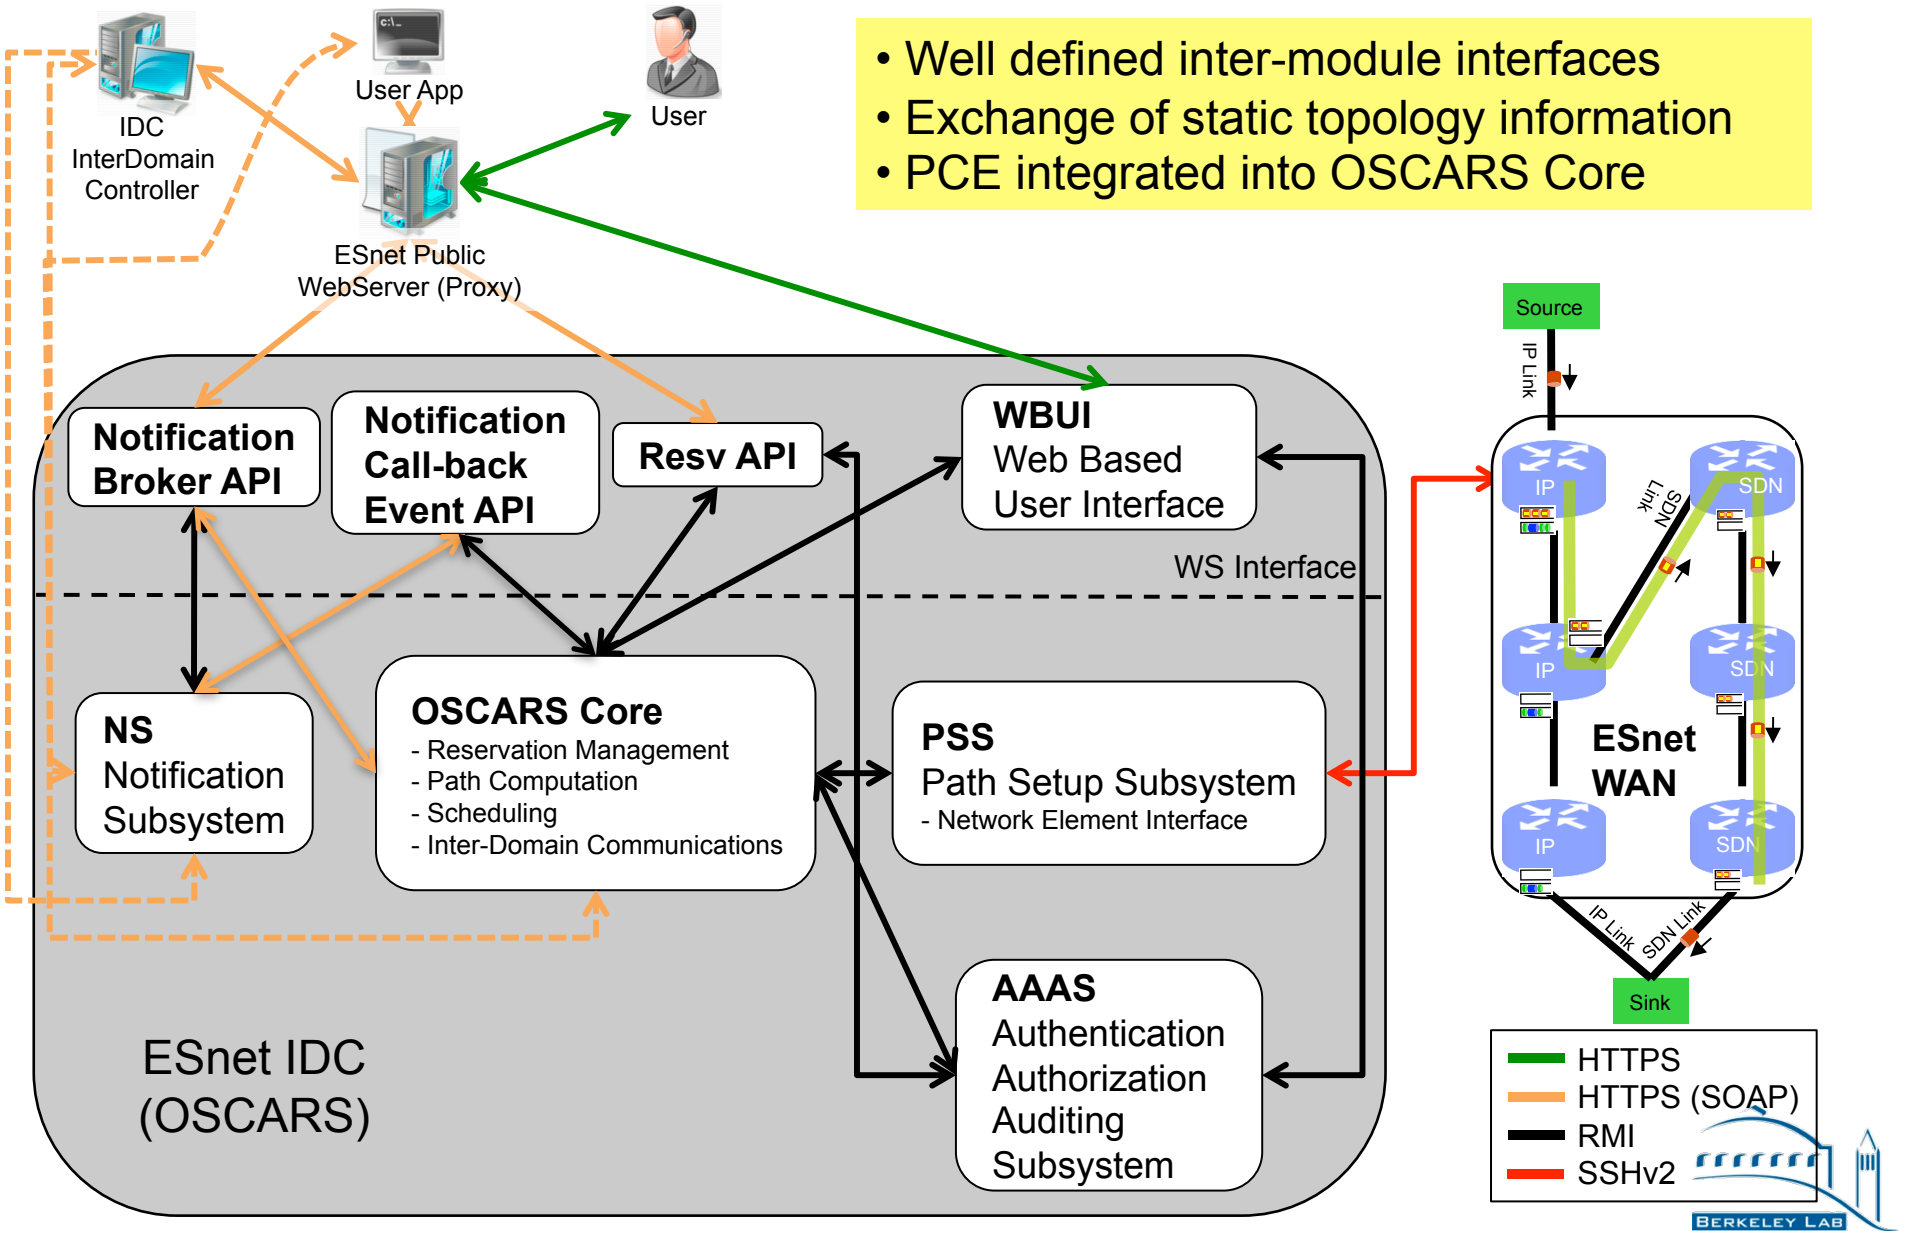

OSCARS Future Implementation

- Exchange of dynamic topology information
  - includes time dimension
- PCE separated from OSCARS Core
  - PCEs can be daisy changed
  - allows PCE to be pluggable
  - facilitates a research framework for collaboration

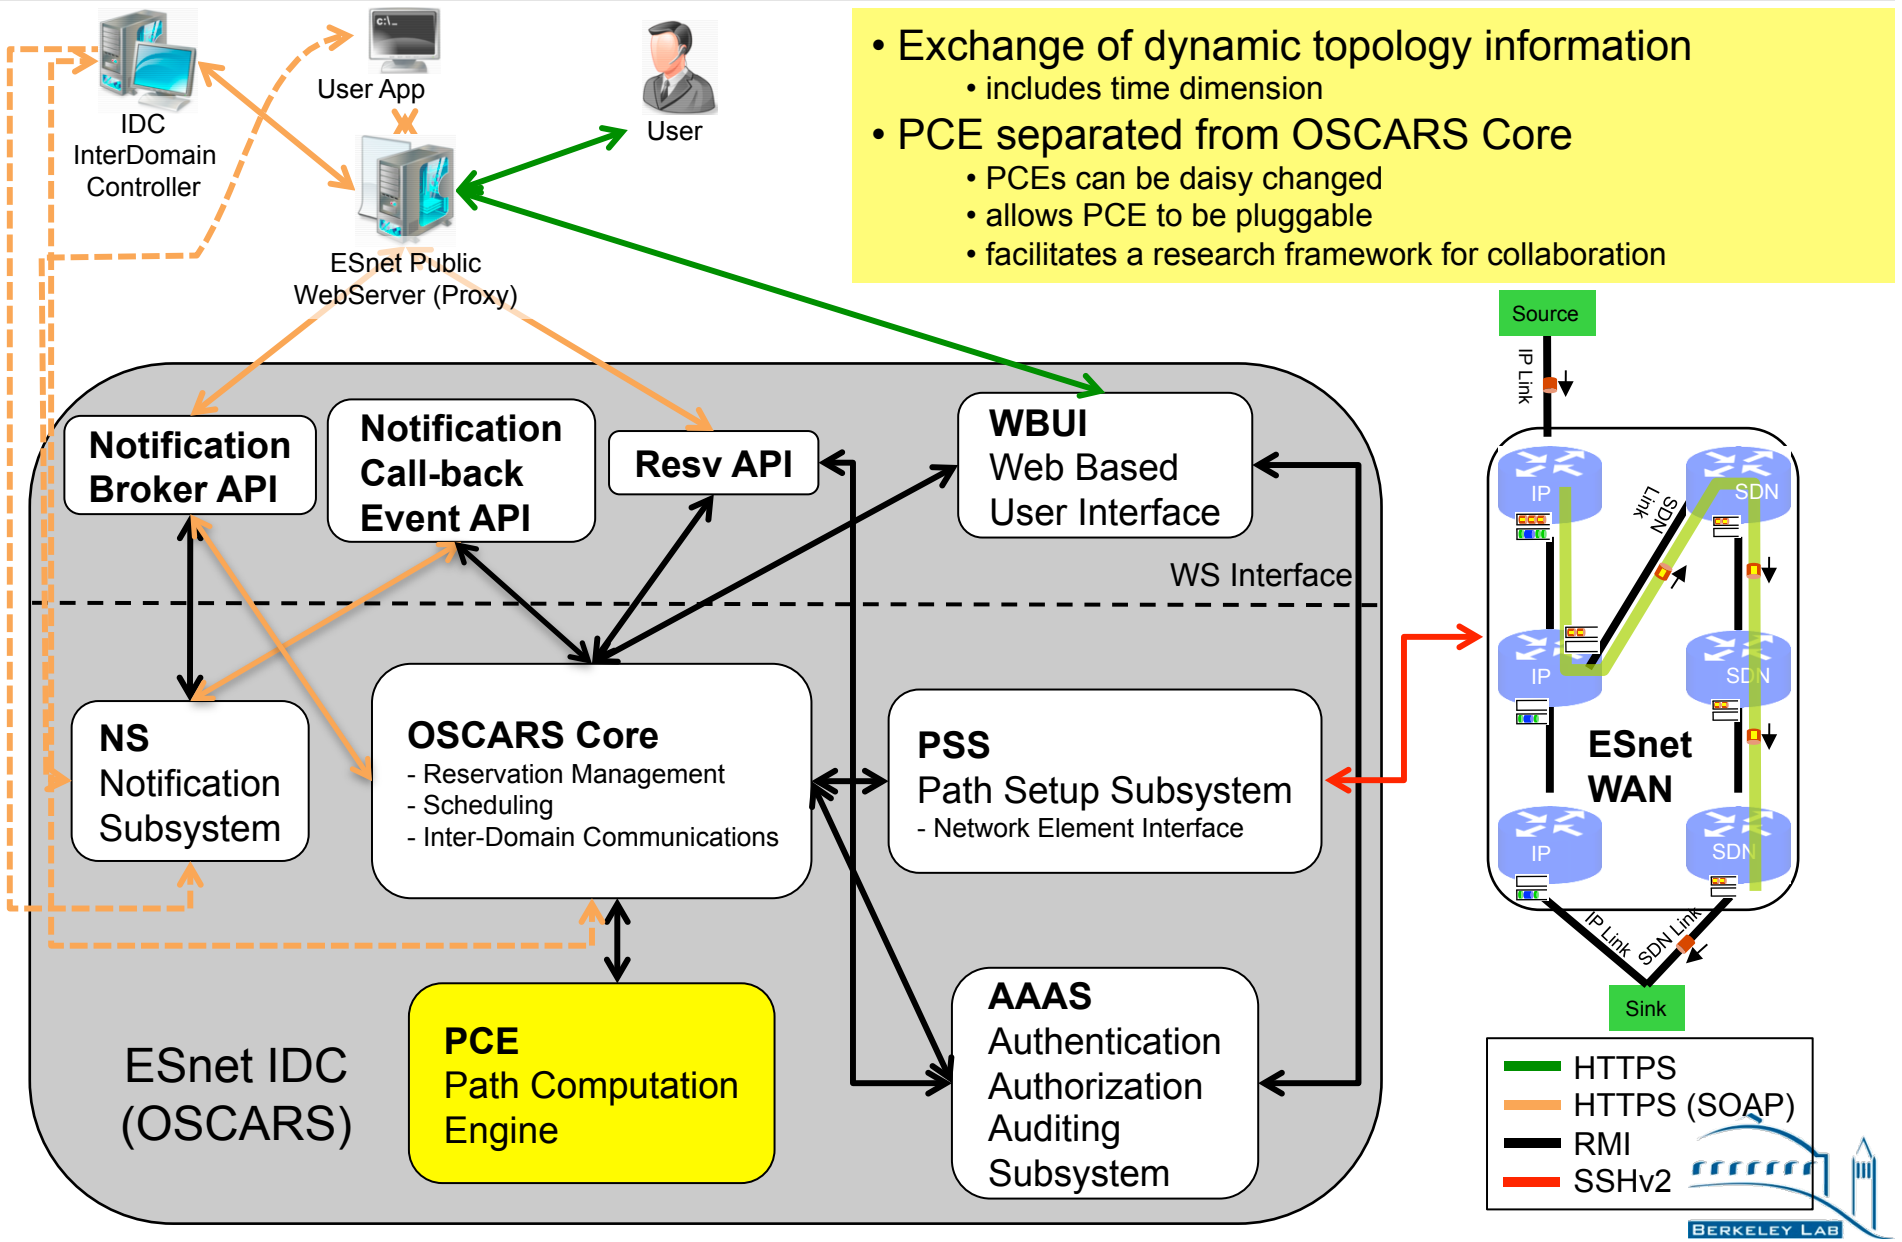

Production OSCARS

- Modifications needed by FNAL and BNL
  - Changed the reservation workflow, added a notification callback system, and added some parameters to the OSCARS API to improve interoperability with automated provisioning agents such as LambdaStation, Terapaths and Phoebus.
- Operational VC support
  - As of 12/2/08, there were 16 long-term production VCs instantiated, all of which support HEP
    - 4 VCs terminate at BNL
    - 2 VCs support LHC T0-T1 (primary and backup)
    - 12 VCs terminate at FNAL
    - 2 VCs support LHC T0-T1 (primary and backup)
    - For BNL and FNAL LHC T0-T1 VCs, except for the ESnet PE router at BNL (bnl-mr1.es.net) and FNAL (fnal-mr1-es.net), there are no other common nodes (router), ports (interfaces), or links between the primary and backup VC.
- Short-term dynamic VCs
  - Between 1/1/08 and 12/2/08, there were roughly 2650 successful HEP centric VCs reservations
    - 1950 reservations initiated by BNL using Terapaths
    - 1700 reservations initiated by FNAL using LambdaStation

OSCARS is a Production Service

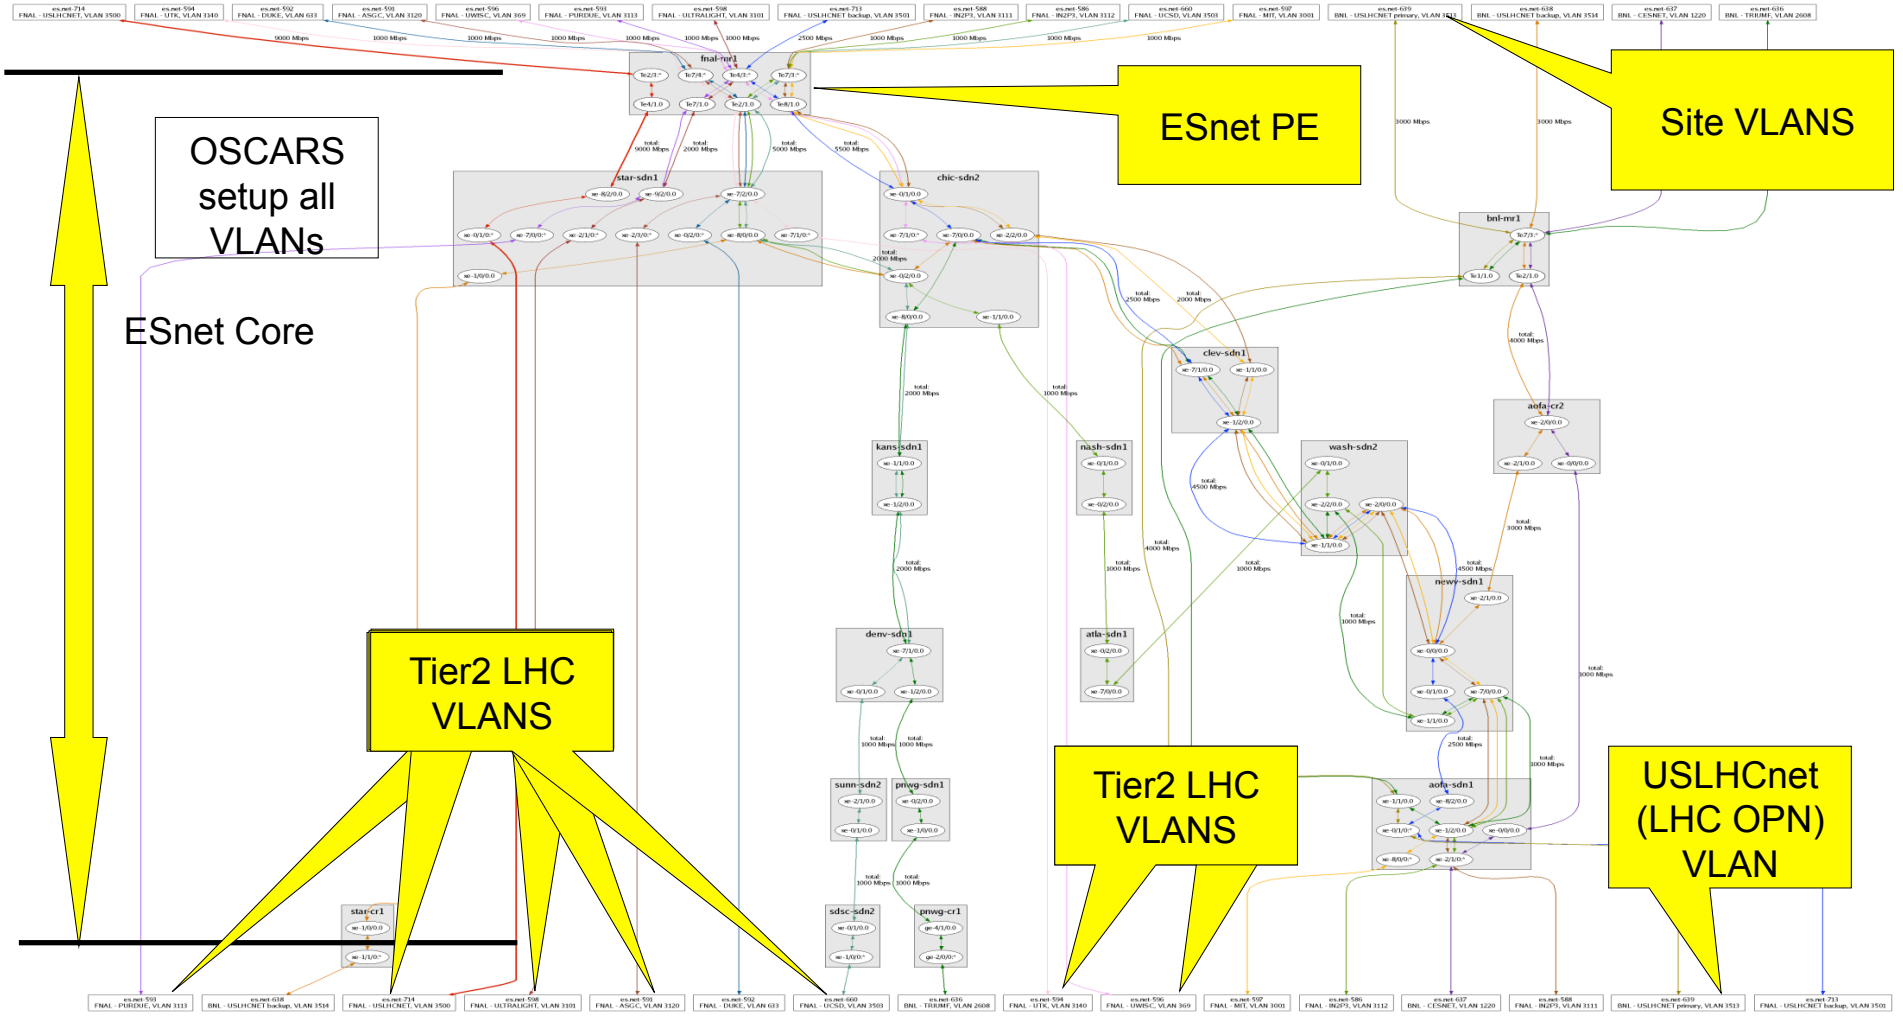

OSCARS generated and managed virtual circuits at FNAL – one of the US LHC Tier 1 data centers. This circuit map (minus the yellow callouts that explain the diagram) is automatically generated by an OSCARS tool and assists the connected sites with keeping track of what circuits exist and where they terminate.

Spectrum Now Monitors OSCARS Circuits

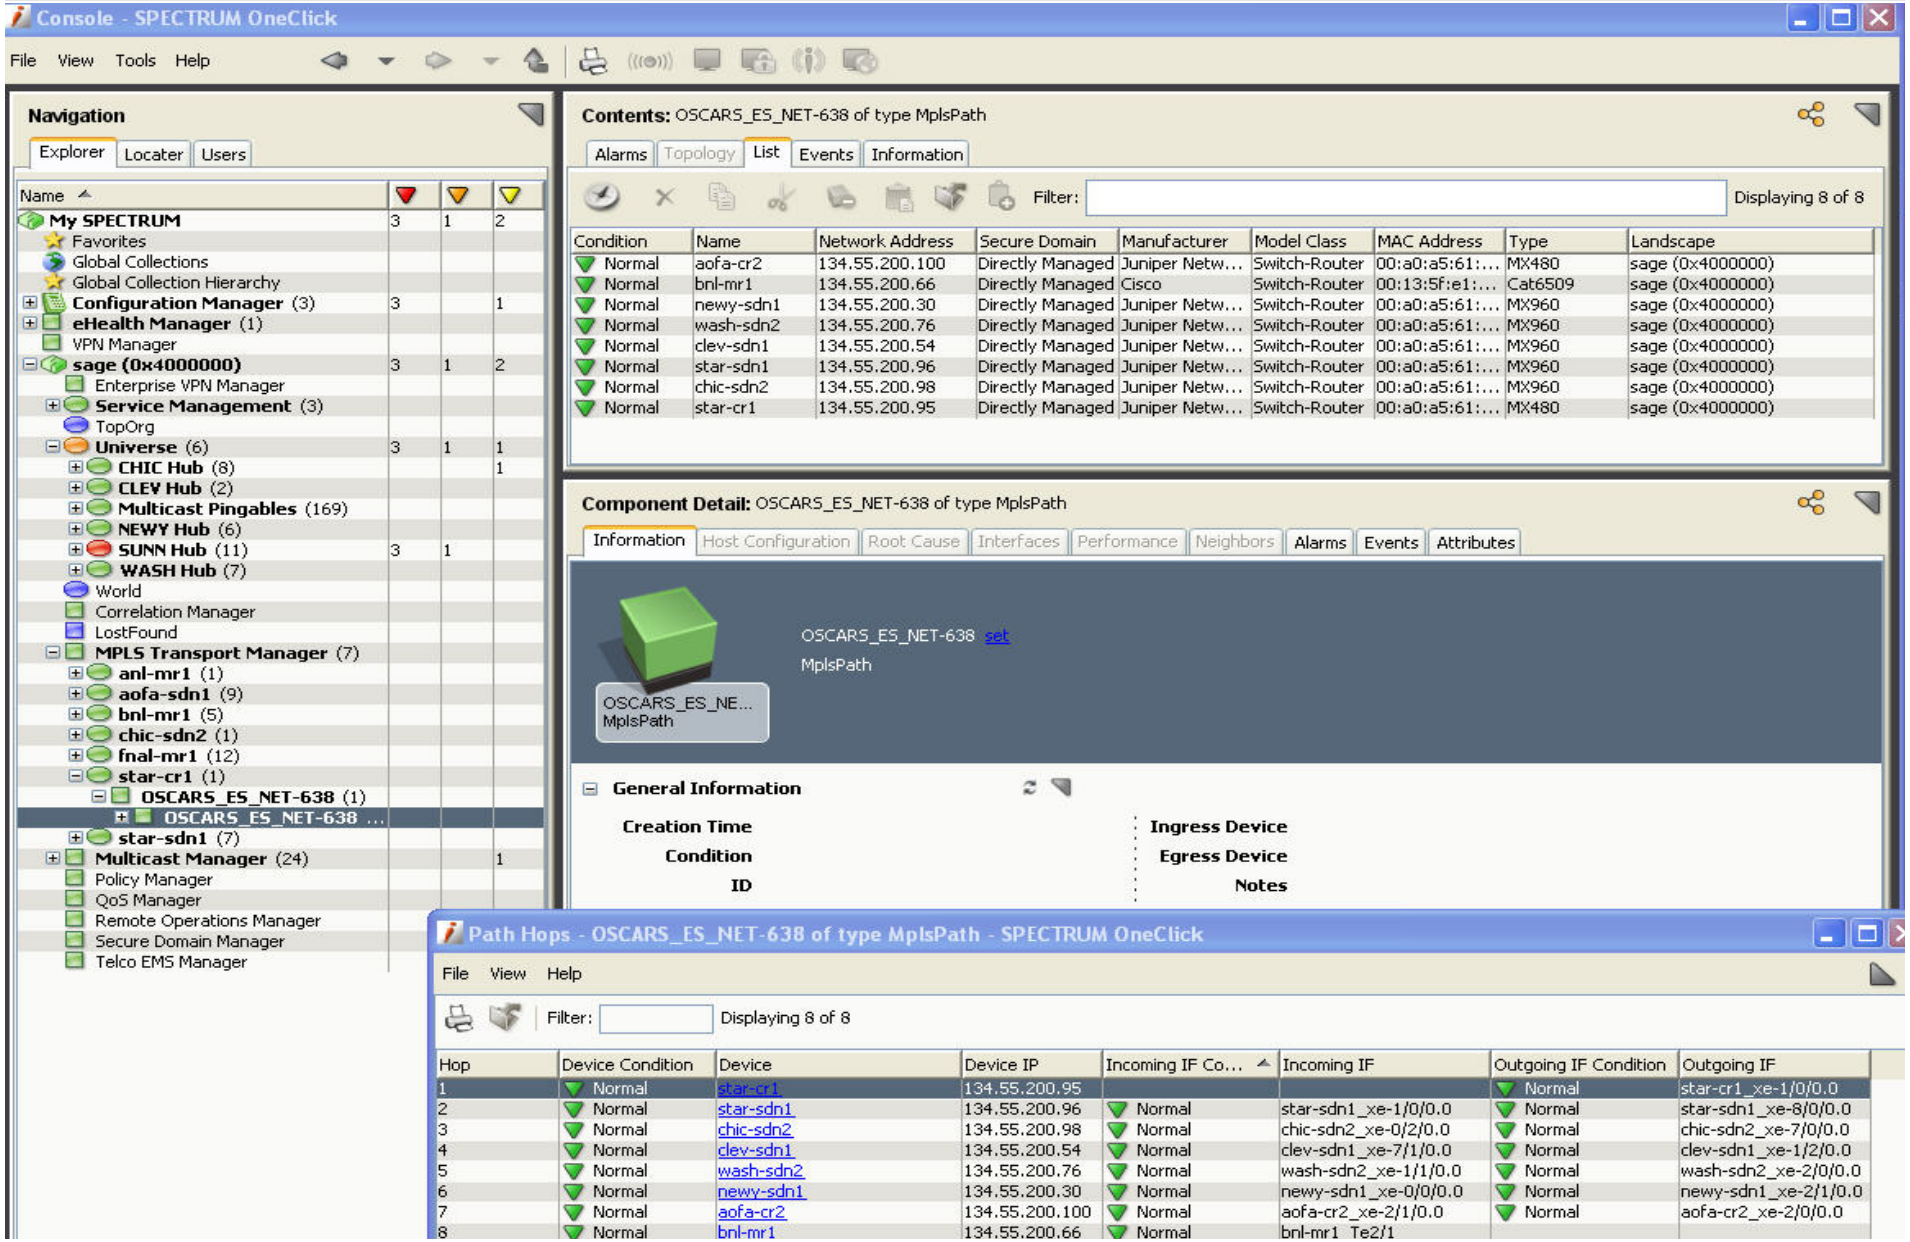

Agenda

- Network Update
- OSCARS
- perfSONAR
- Federated Trust

perfSONAR Services

- End-to-end monitoring service: providing useful, comprehensive, and meaningful ***information on the state of end-to-end paths.*** Supports regularly scheduled tests & archiving of results, acting as an intermediate layer, between the performance measurement tools and the diagnostic or visualization applications.
- Tools in the perfSONAR software suite:
  - SNMP Measurement Archive
  - Lookup Service
  - Topology Service
  - Circuit Status Measurement Archive
  - Status Measurement Archive
  - perfSONAR-BUOY
  - PingER Services
- Visualization
  - Allow ESnet user community to better understand our network & its capabilities.
  - Allow ESnet users to understand how their use impacts the backbone.
- Alarming
  - Automated analysis of regularly scheduled measurements to raise alerts.

ESnet Deployment Activities

- Currently deploying the hardware across the network to support adhoc measurements for debugging
  - OWAMP Servers
  - BWCTL Servers
  - Topology Service
  - Utilization Service
- perfSONAR Buoy Deployment
  - Between ESnet systems
  - To Internet2 & GEANT
  - To/From ESnet Sites
- Hardens the infrastructure
  - Continuous monitoring of servers & services
  - Centralized management of OS & Services configuration
  - Performance tuning & verifying everything is working as designed

perfSONAR R&D Activities

- Scaling & robustness enhancements
- Visualization Tools
  - Single Domain Tools
    - Utilization Browser
    - Topology Browser
    - Latency & Bandwidth Browser
  - Advanced Tools
    - Looking across multiple domains
    - Looking at correlations between different types of measurements
    - Application or user community specific views
- Alarming
- Integrating OSCARS circuits
  - Topology
  - Utilization
  - Active measurements across them

Agenda

- Network Update
- OSCARS
- perfSONAR
- Federated Trust Services

Federated Trust Services

- DOEGrids Certification Authority
  - New Logo and ID Mark
  - Operations
  - DOEGrids Audit progress
  - Cloning and Geographical Dispersion
- OpenID and Shibboleth
- Authorization Services Profile Document

DOEGrids CA - Operations

- Vista – IE browser support in development
  - Also beginning testing IE 8 browser
- ESnet 2-factor
  - Support ESnet 2-factor authentication token project
  - Add ESnet RA to list of official RAs in DOEGrids CA
- Recent problems – Dec 2008
  - CA not reading own Cert Revocation Lists
  - CA automatically certifying customers from a peer, highly trusted CA (CERN CA)
  - ➢ These problems have been corrected
    - All certifications since June 2007 were audited
    - No fraudulent certifications were discovered
    - By agreement with registration authorities, affected subscribers will undergo direct reverification at next renewal
      (RA's are free to require this at any time)
    - (See auditing slide)

DOEGrids CA (one of several CAs) Usage Statistics

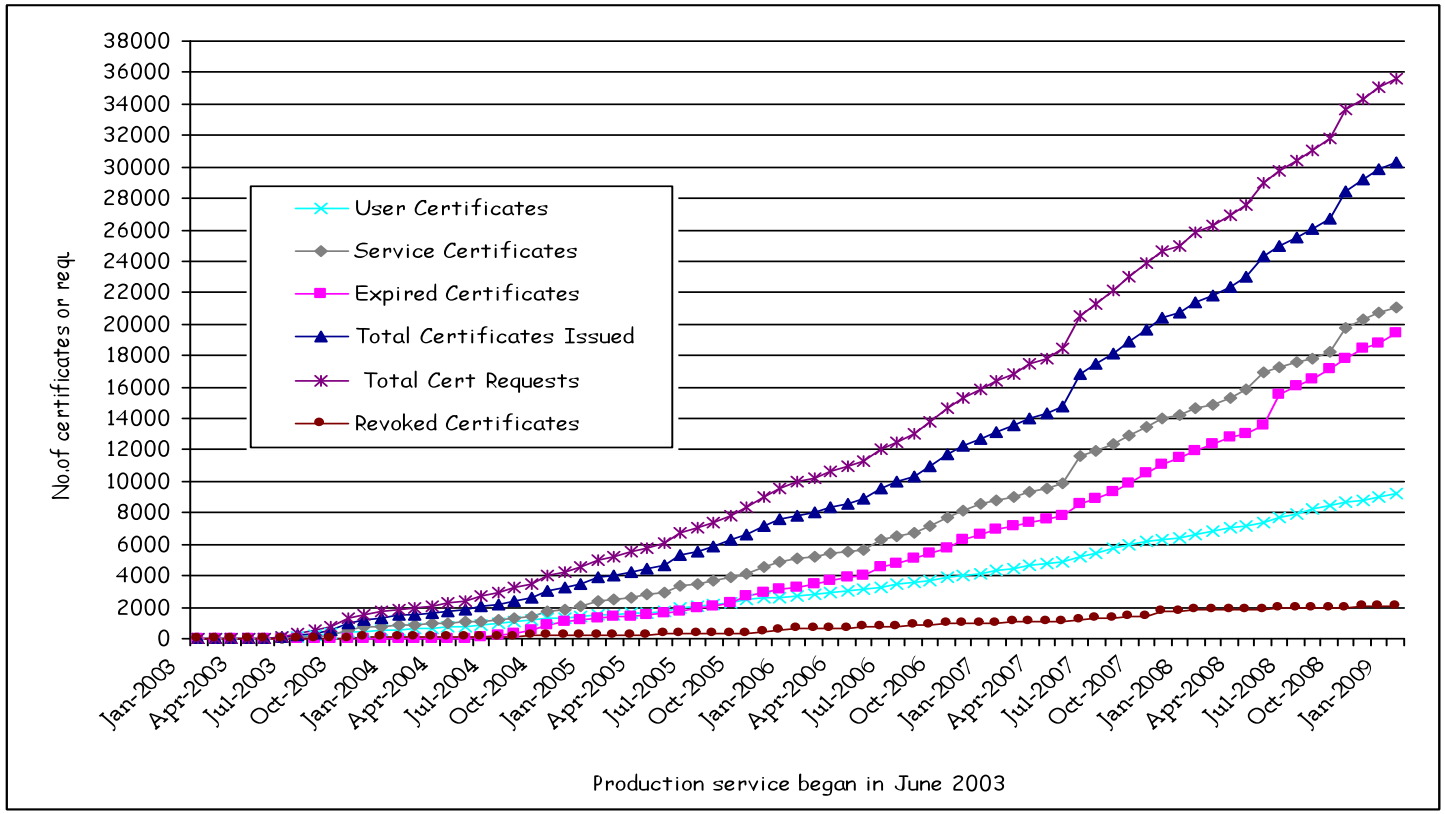

User Certificates	9259	Total No. of Revoked Certificates	2056
Host & Service Certificates	21043	Total No. of Expired Certificates	19452
Total No. of Requests	35629	Total No. of Certificates Issued	30331
		Total No. of Active Certificates	8823
ESnet SSL Server CA Certificates			50
FusionGRID CA certificates		* Report as of Jan 29, 2009	113

DOEGrids CA (Active Certificates) Usage Statistics

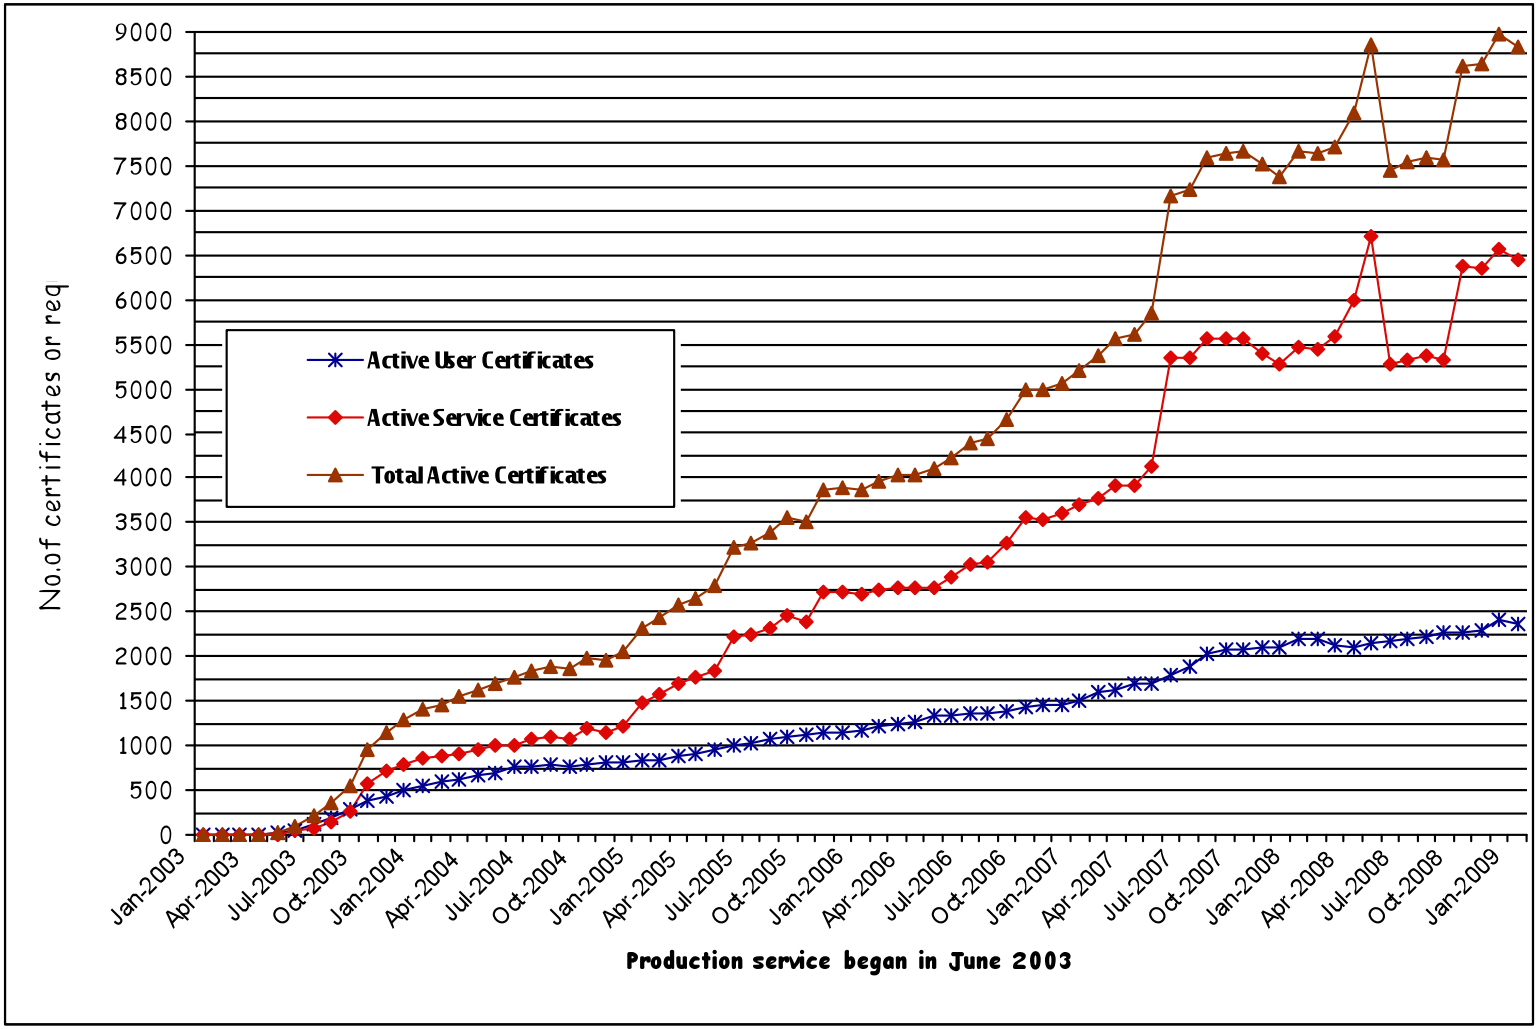

* Report as of Jan 29, 2009

Active DOEGrids CA Breakdown

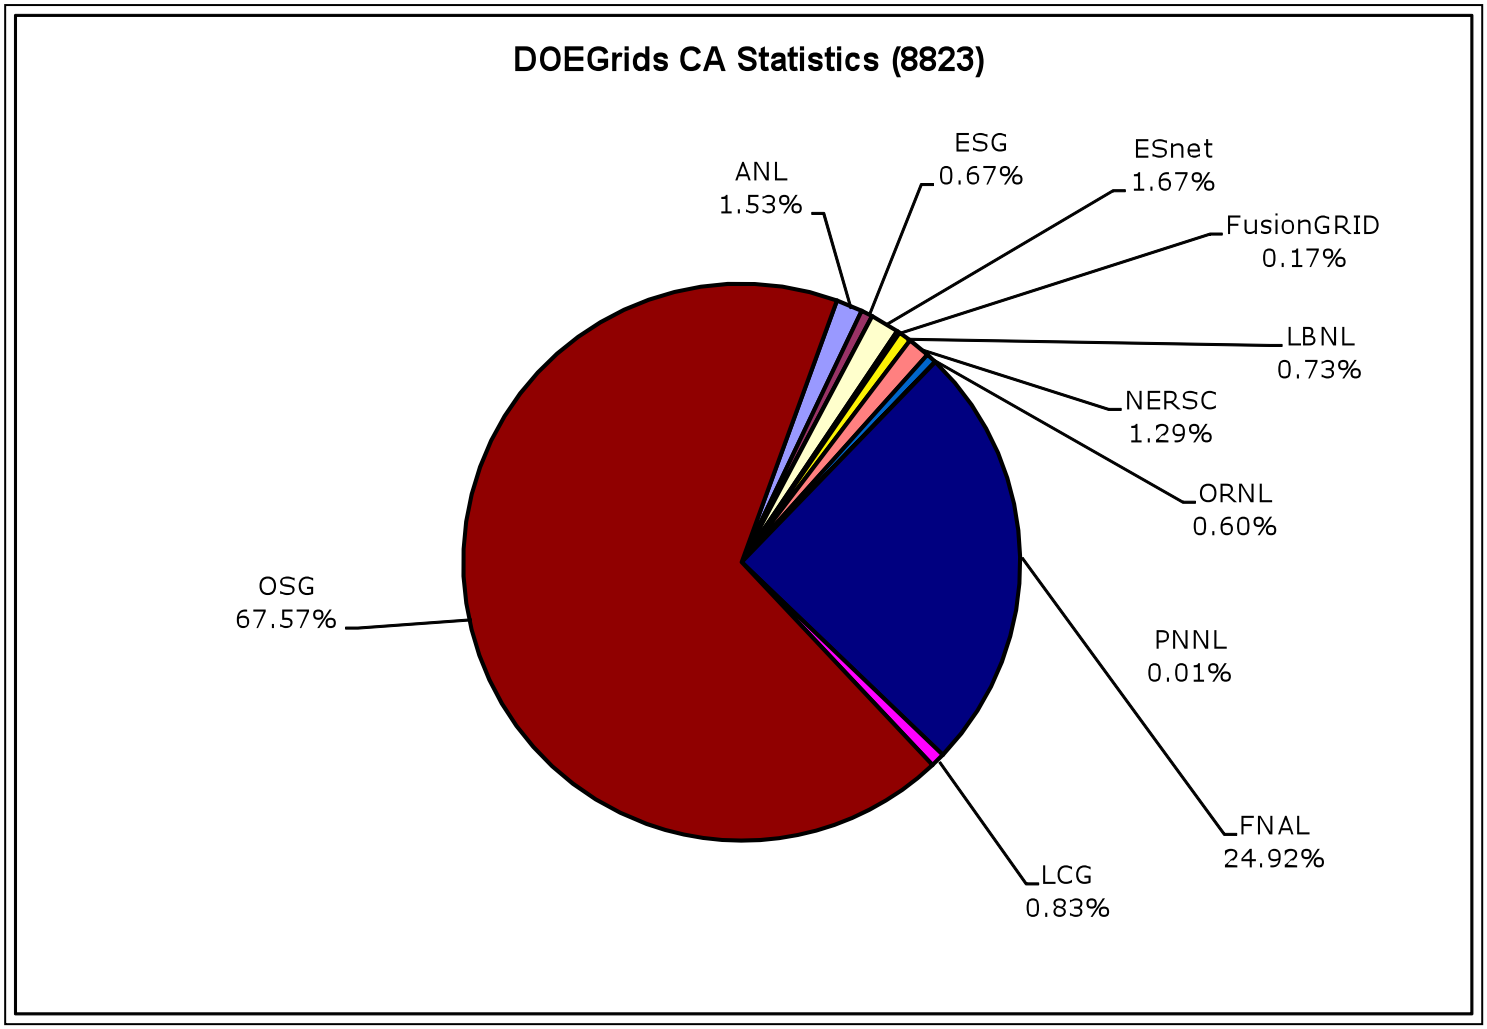

** OSG Includes (BNL, CDF, CIGI,CMS, CompBioGrid, DES, DOSAR, DZero, Engage, Fermilab, fMRI, GADU, geant4, GLOW, GPN, GRASE, GridEx, GUGrid, i2u2, ILC, JLAB, LIGO, mariachi, MIS, nanoHUB, NWICG, NYSGrid, OSG, OSGEDU, SBGrid, SDSS, SLAC, STAR & USATLAS)

DOEGrids CA - Audits

- The Certification Practices Statement (CPS) is being “translated” to the RFC 3647 format
  - Audit finding – requirement
  - Appropriate format for interoperation
- Next step will be to correct all documentation errors identified in the audit
- Scheduling an audit of configurations, modules, and operational scripts (see Dec 2008 problems) – Feb/Mar 2009

DOEGrids CA Cloning & Geographical Dispersion

- DOEGrids CA and its key management hardware will be cloned and dispersed around the US
  - Improve Continuity of Operations and disaster recovery issues (ESnet requirements)
  - Improve availability to customers
  - Provision for future, robust services
  - Current status: Testing and configuration of netHSM hardware, and project planning

OpenID and Shibboleth

- Continue efforts to promote this technology in DOE Laboratory community – won't you join us?
- OpenID: Summer project testing OpenID provider (mostly) with simple server
  - Objective: Use DOEGrids CA as source of identity
  - Objective: Test simple application (custom, and later simplified wiki)
  - See http://www.doegrids.org/OpenID/
  - Roadmap for phase 2: Robust version of summer project, with more SSO and addition of other OpenID consumers as opportunities appear
- Shibboleth: Similar roadmap as for OpenID
- Many security issues to consider
- WAYF/Discovery a problem for both services – perhaps this is an opportunity for a 3rd service, CardSpace

Identity and Federation Technology

- Shibboleth.
  - SAML 2.0
  - InCommon Federation

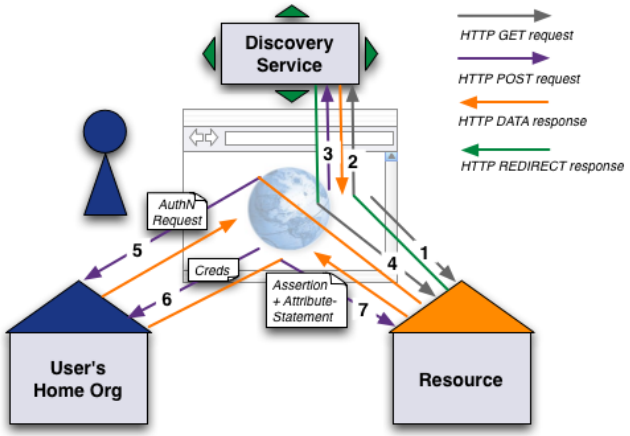

Graphics from SWITCH

- OpenID
  - OP and demo Consumer

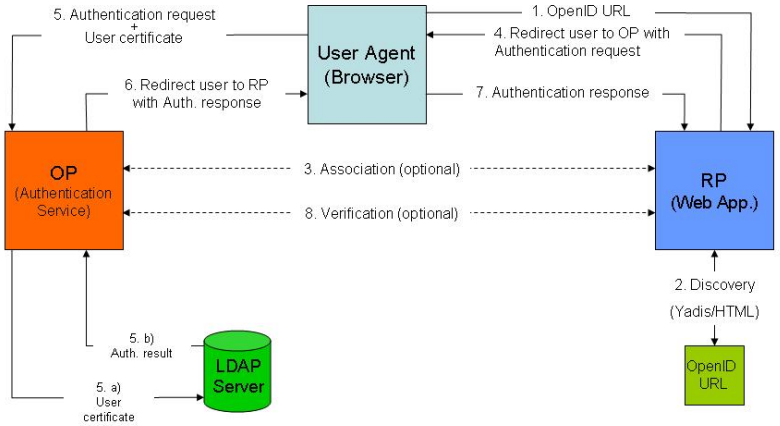

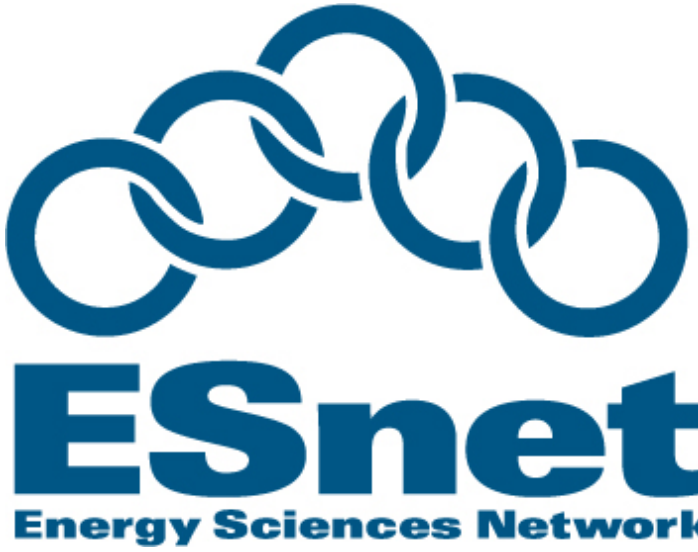
ESnet
Energy Sciences Network


U.S. DEPARTMENT OF
ENERGY
Office of
Science