ESnet Status Update

ESCC
January 23, 2008 (Aloha!)

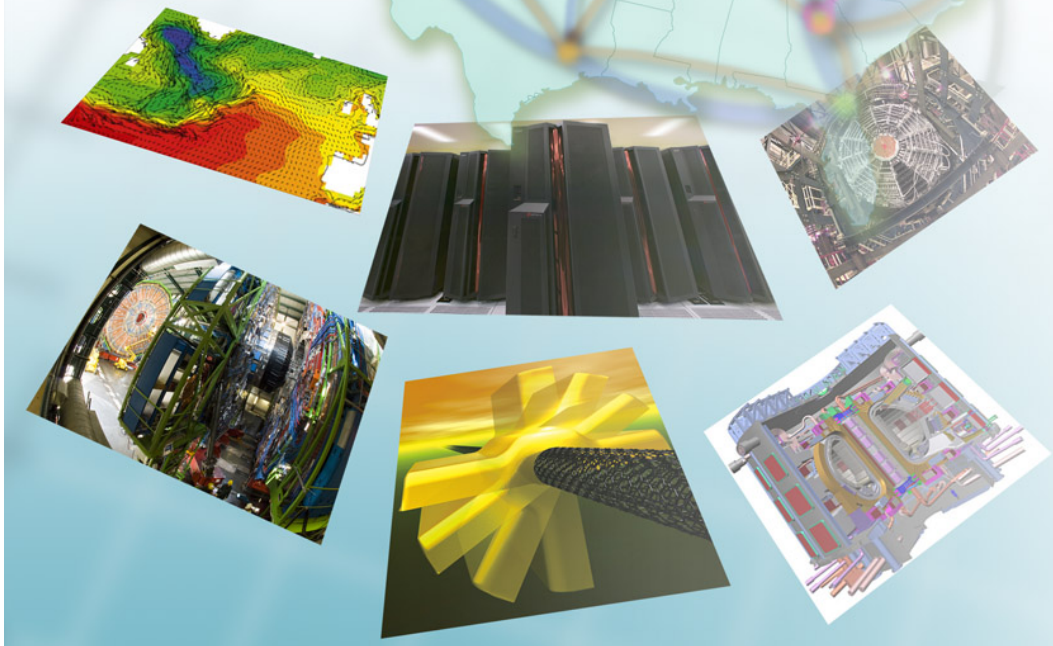

William E. Johnston
ESnet Department Head and Senior Scientist

Energy Sciences Network
Lawrence Berkeley National Laboratory

wej@es.net, www.es.net

This talk is available at www.es.net/ESnet4

Networking for the Future of Science

DOE Office of Science and ESnet – the ESnet Mission

- **ESnet's primary mission is to enable the large-scale science that is the mission of the Office of Science (SC) and that depends on:**
  - Sharing of massive amounts of data
  - Supporting thousands of collaborators world-wide
  - Distributed data processing
  - Distributed data management
  - Distributed simulation, visualization, and computational steering
  - Collaboration with the US and International Research and Education community
- ESnet provides network and collaboration services to Office of Science laboratories and many other DOE programs in order to accomplish its mission

ESnet Stakeholders and their Role in ESnet

- DOE Office of Science Oversight ("SC") of ESnet
  - The SC provides high-level oversight through the budgeting process
  - Near term input is provided by weekly teleconferences between SC and ESnet
  - Indirect long term input is through the process of ESnet observing and projecting network utilization of its large-scale users
  - Direct long term input is through the SC Program Offices Requirements Workshops (more later)
- SC Labs input to ESnet
  - Short term input through many daily (mostly) email interactions
  - Long term input through ESCC

ESnet Stakeholders and the Role in ESnet

- SC science collaborators input
  - Through numerous meeting, primarily with the networks that serve the science collaborators

Talk Outline

I. Building ESnet4

- **Ia.** Network Infrastructure
- **Ib.** Network Services
- **Ic.** Network Monitoring

II. Requirements

III. Science Collaboration Services

- **IIIa.** Federated Trust
- **IIIb.** Audio, Video, Data Teleconferencing

Ia. Building ESnet4 - Starting Point

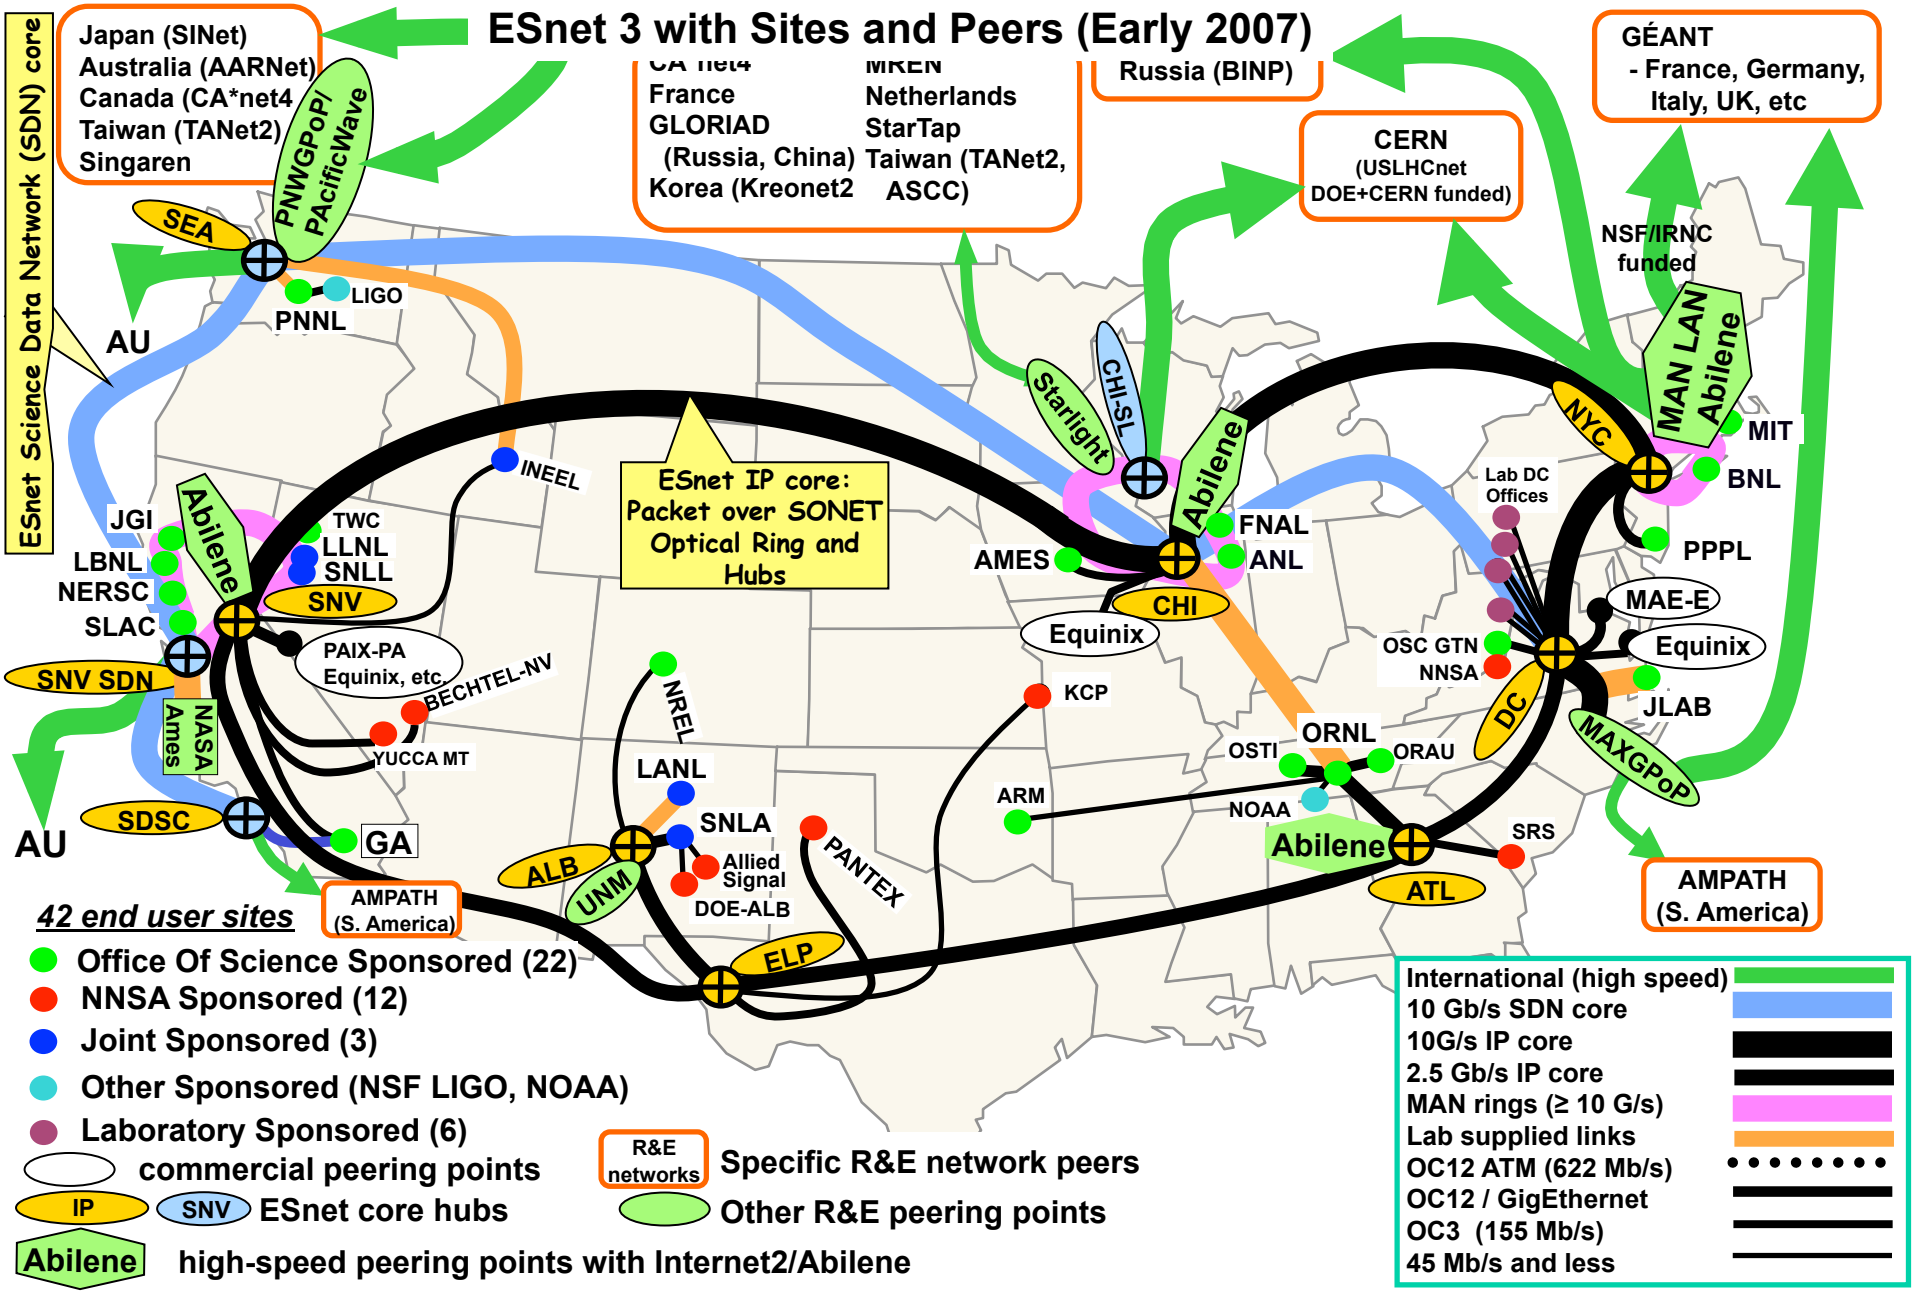

ESnet 3 Backbone as of January 1, 2007

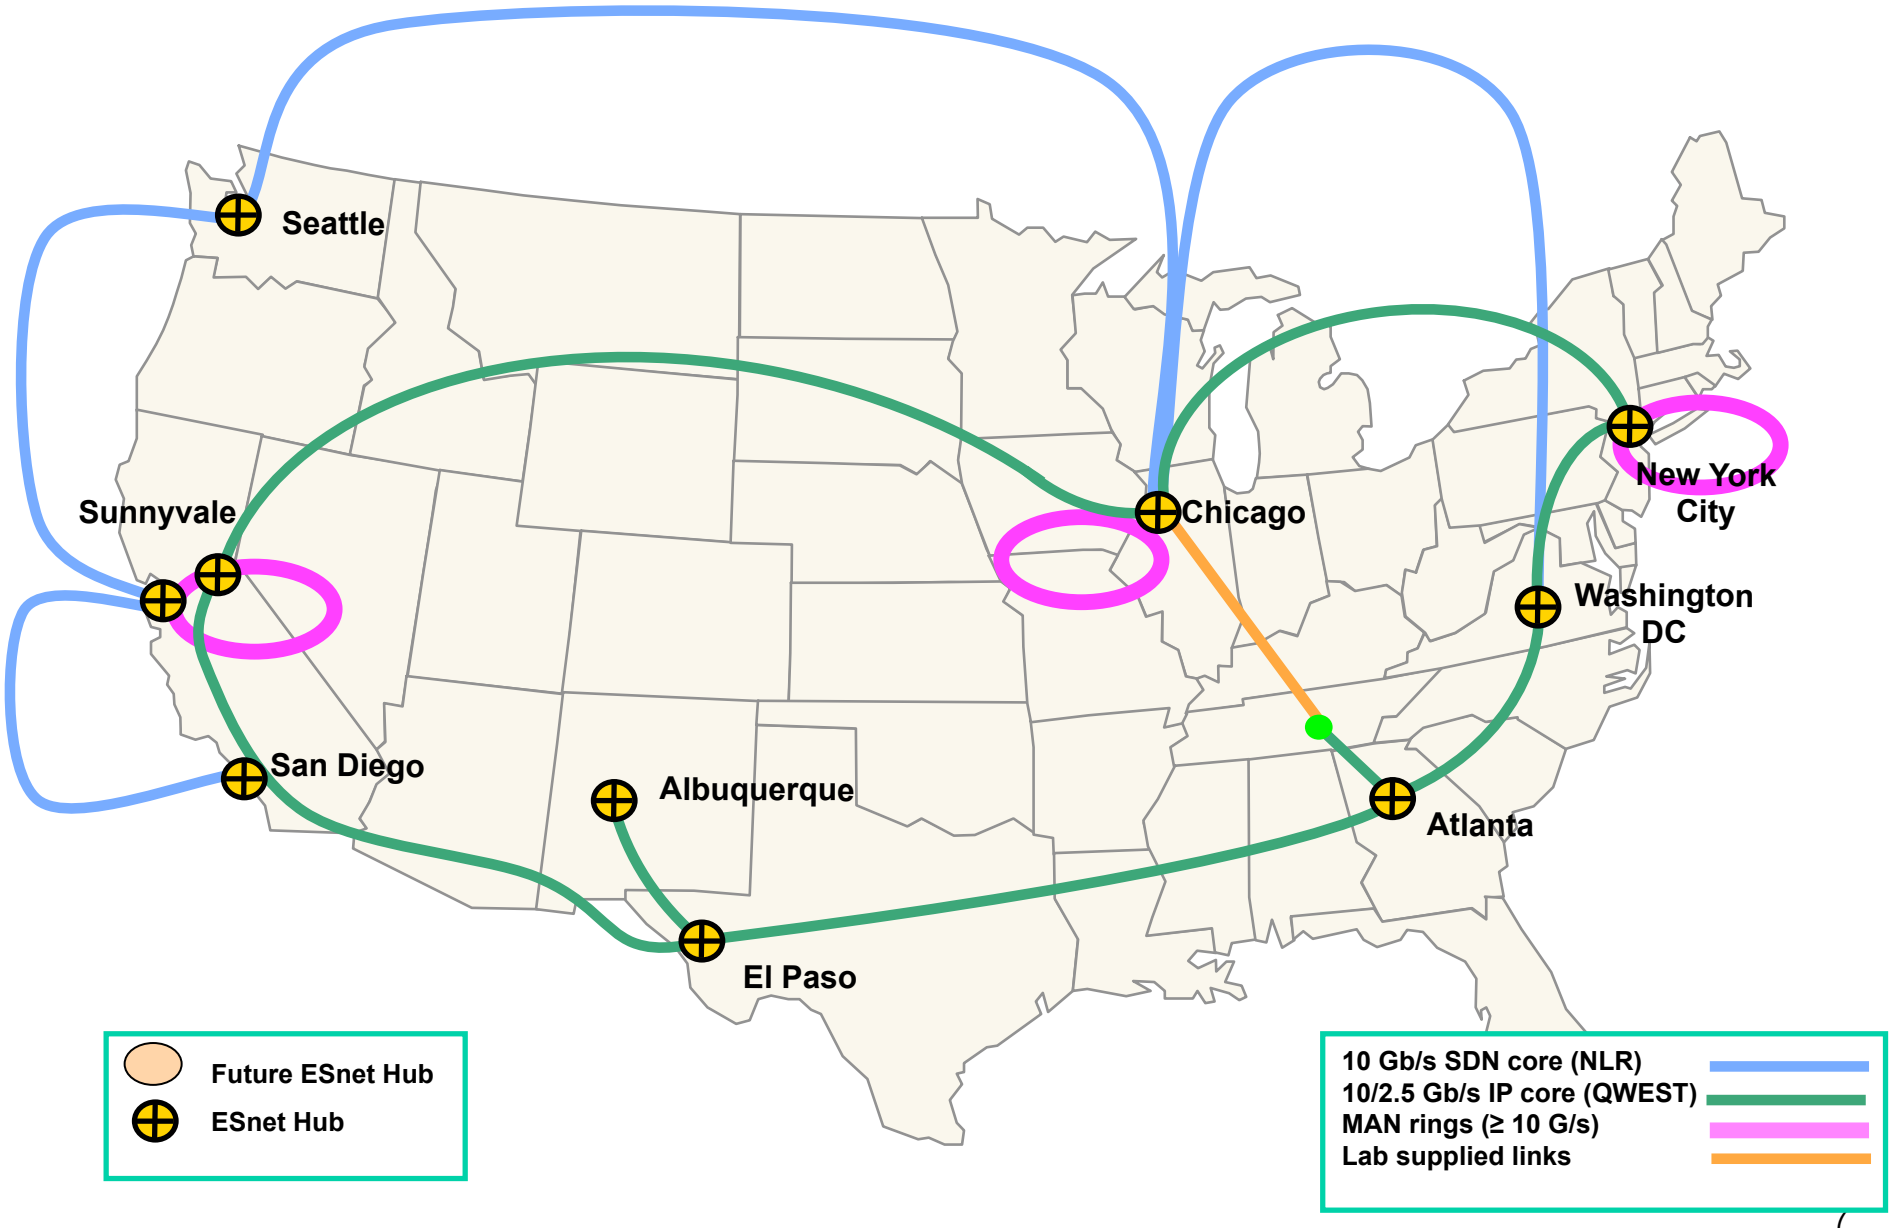

ESnet 4 Backbone as of April 15, 2007

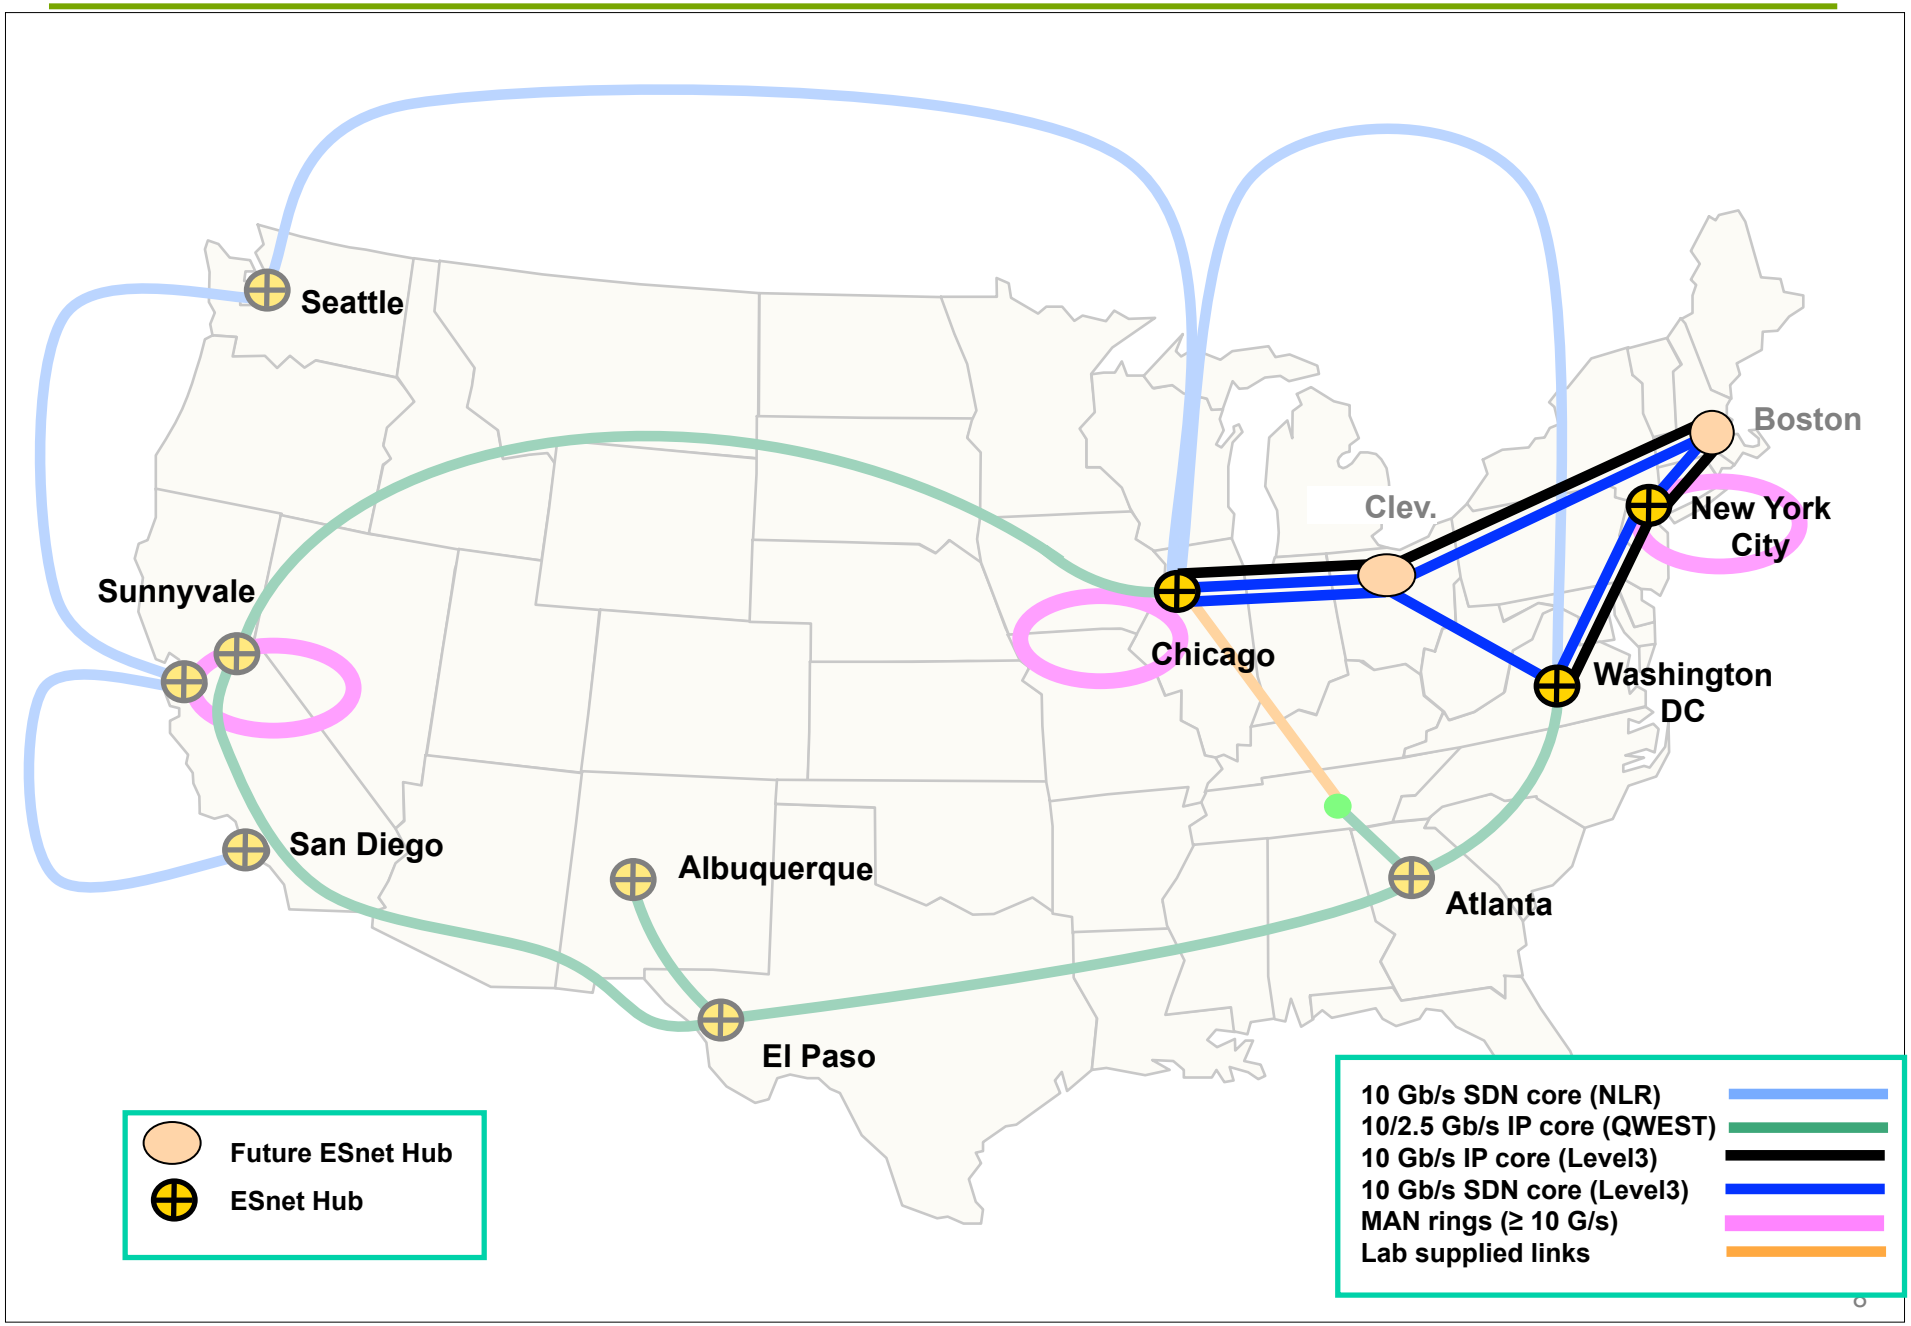

ESnet 4 Backbone as of May 15, 2007

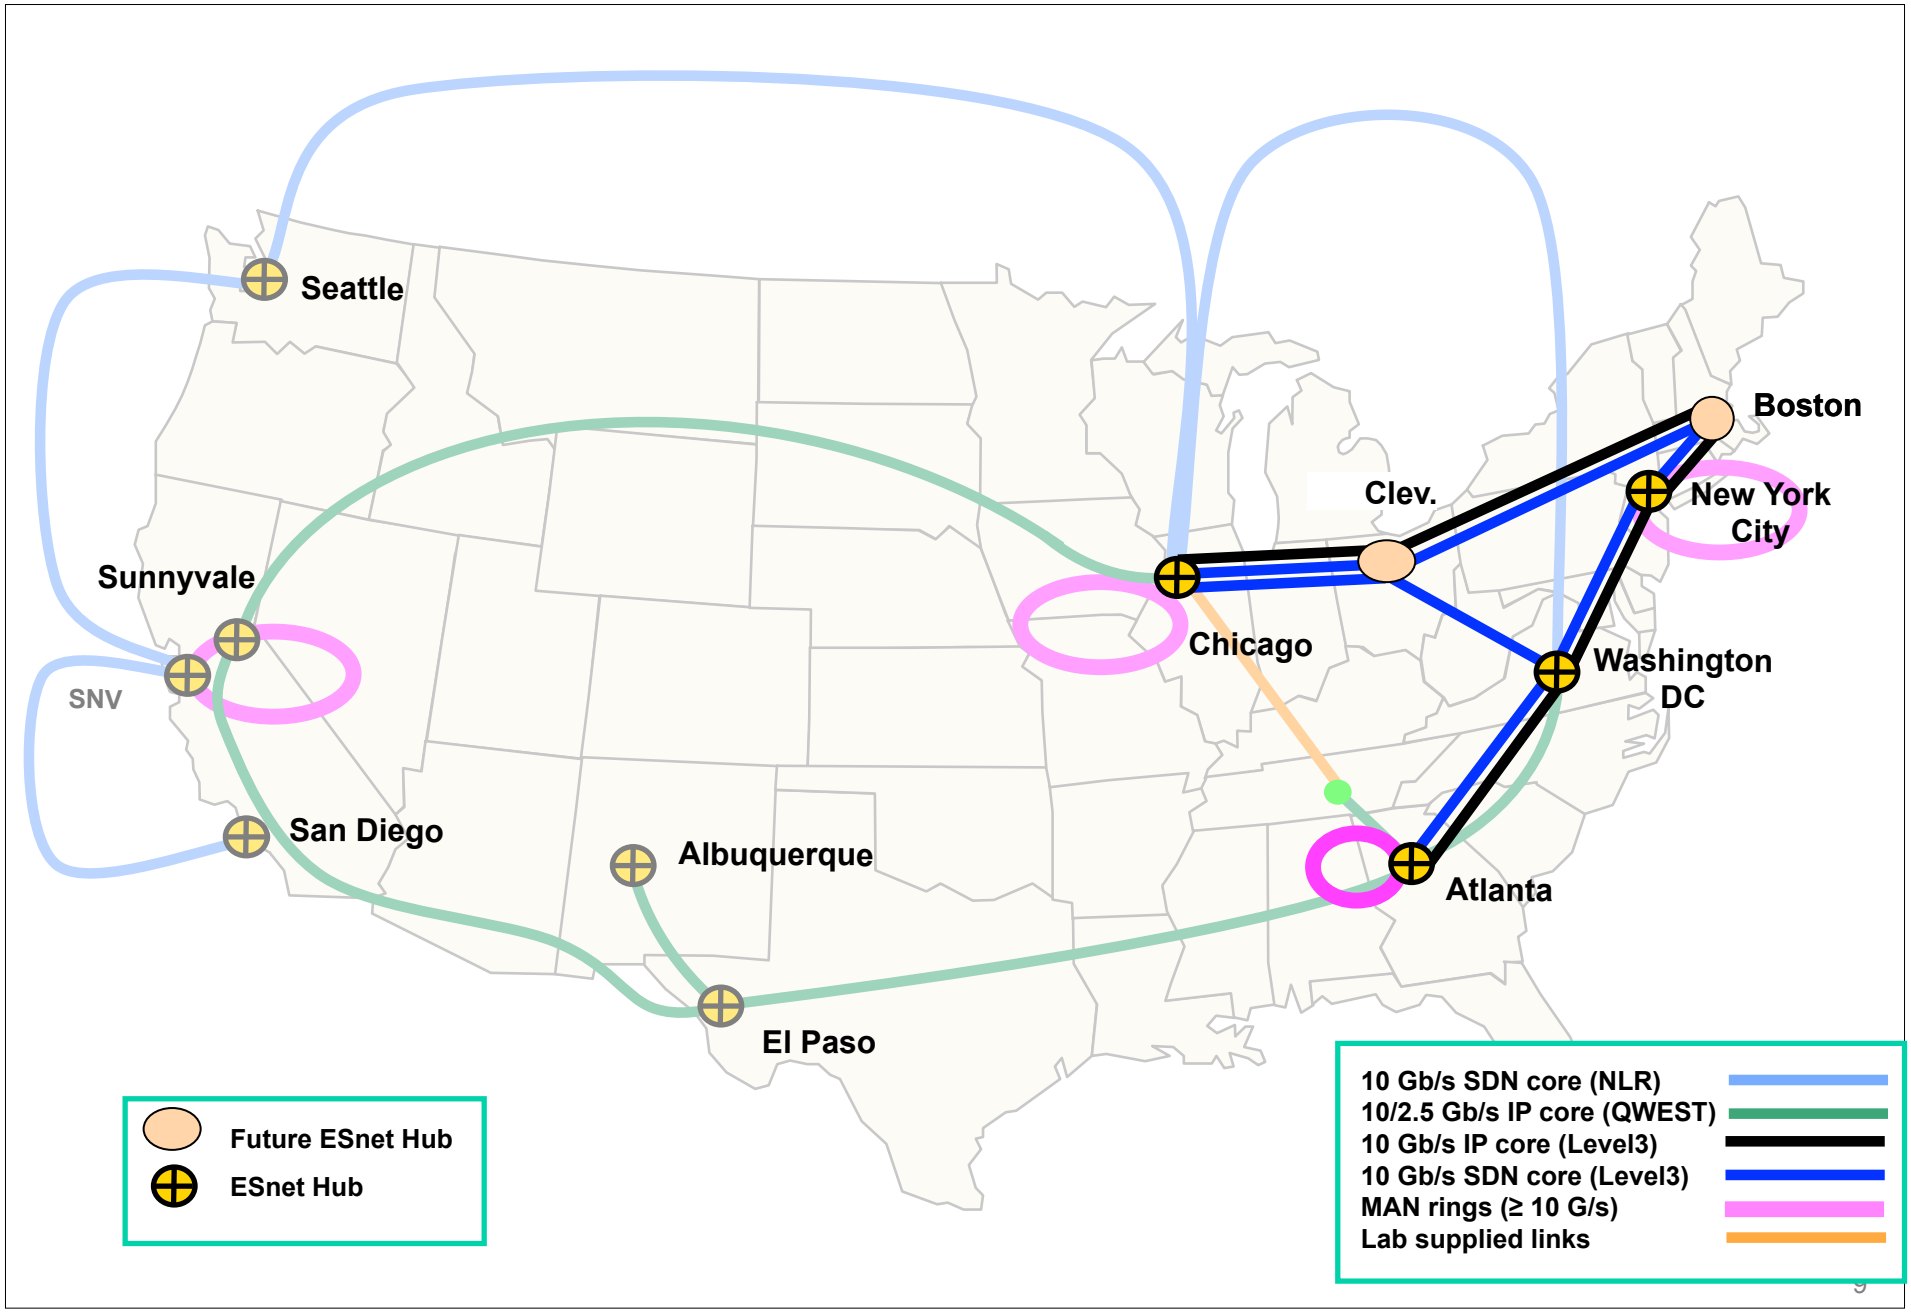

ESnet 4 Backbone as of June 20, 2007

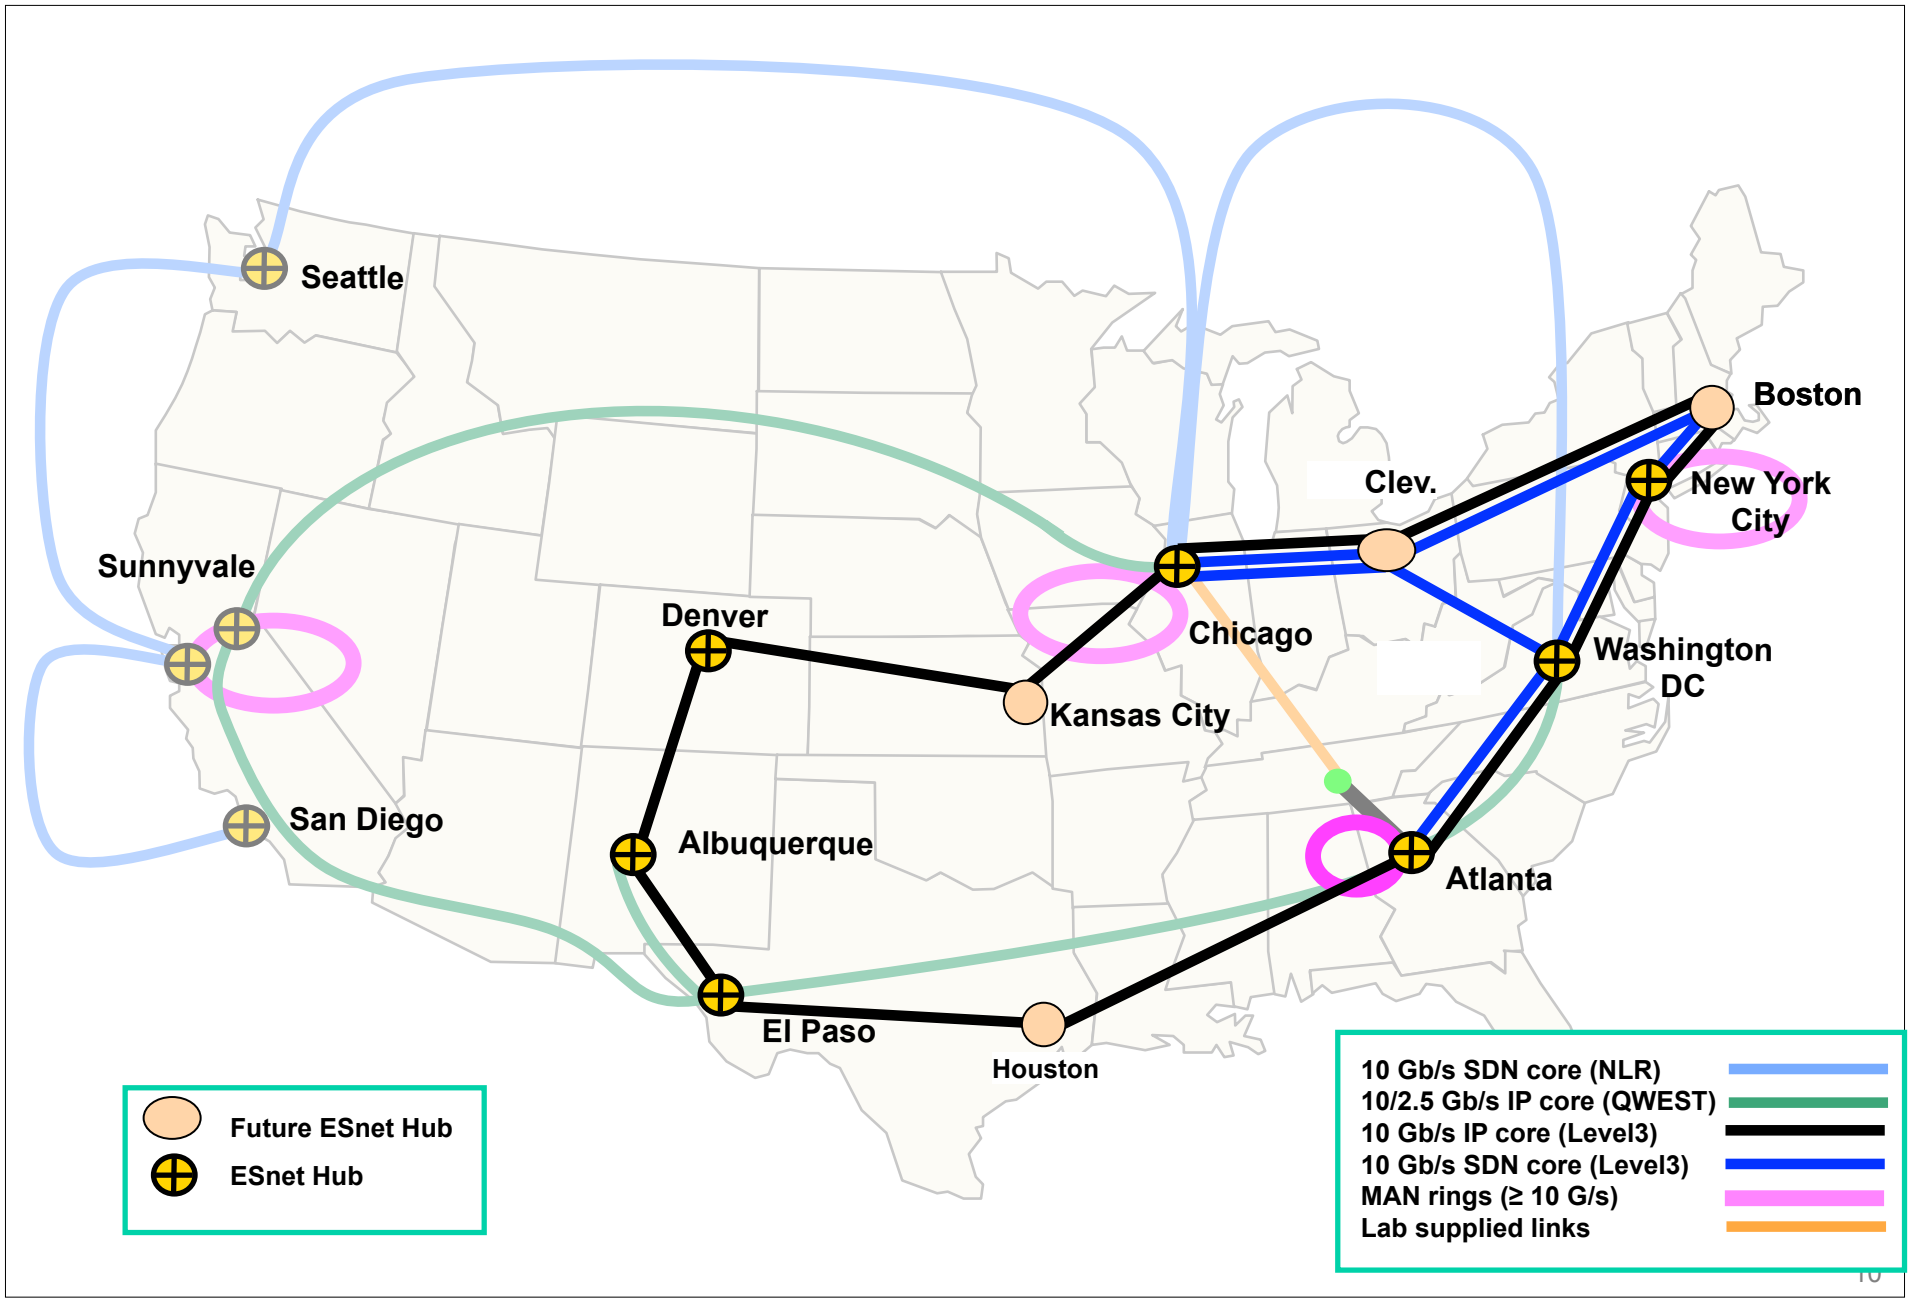

ESnet 4 Backbone August 1, 2007 (Last JT meeting at FNAL)

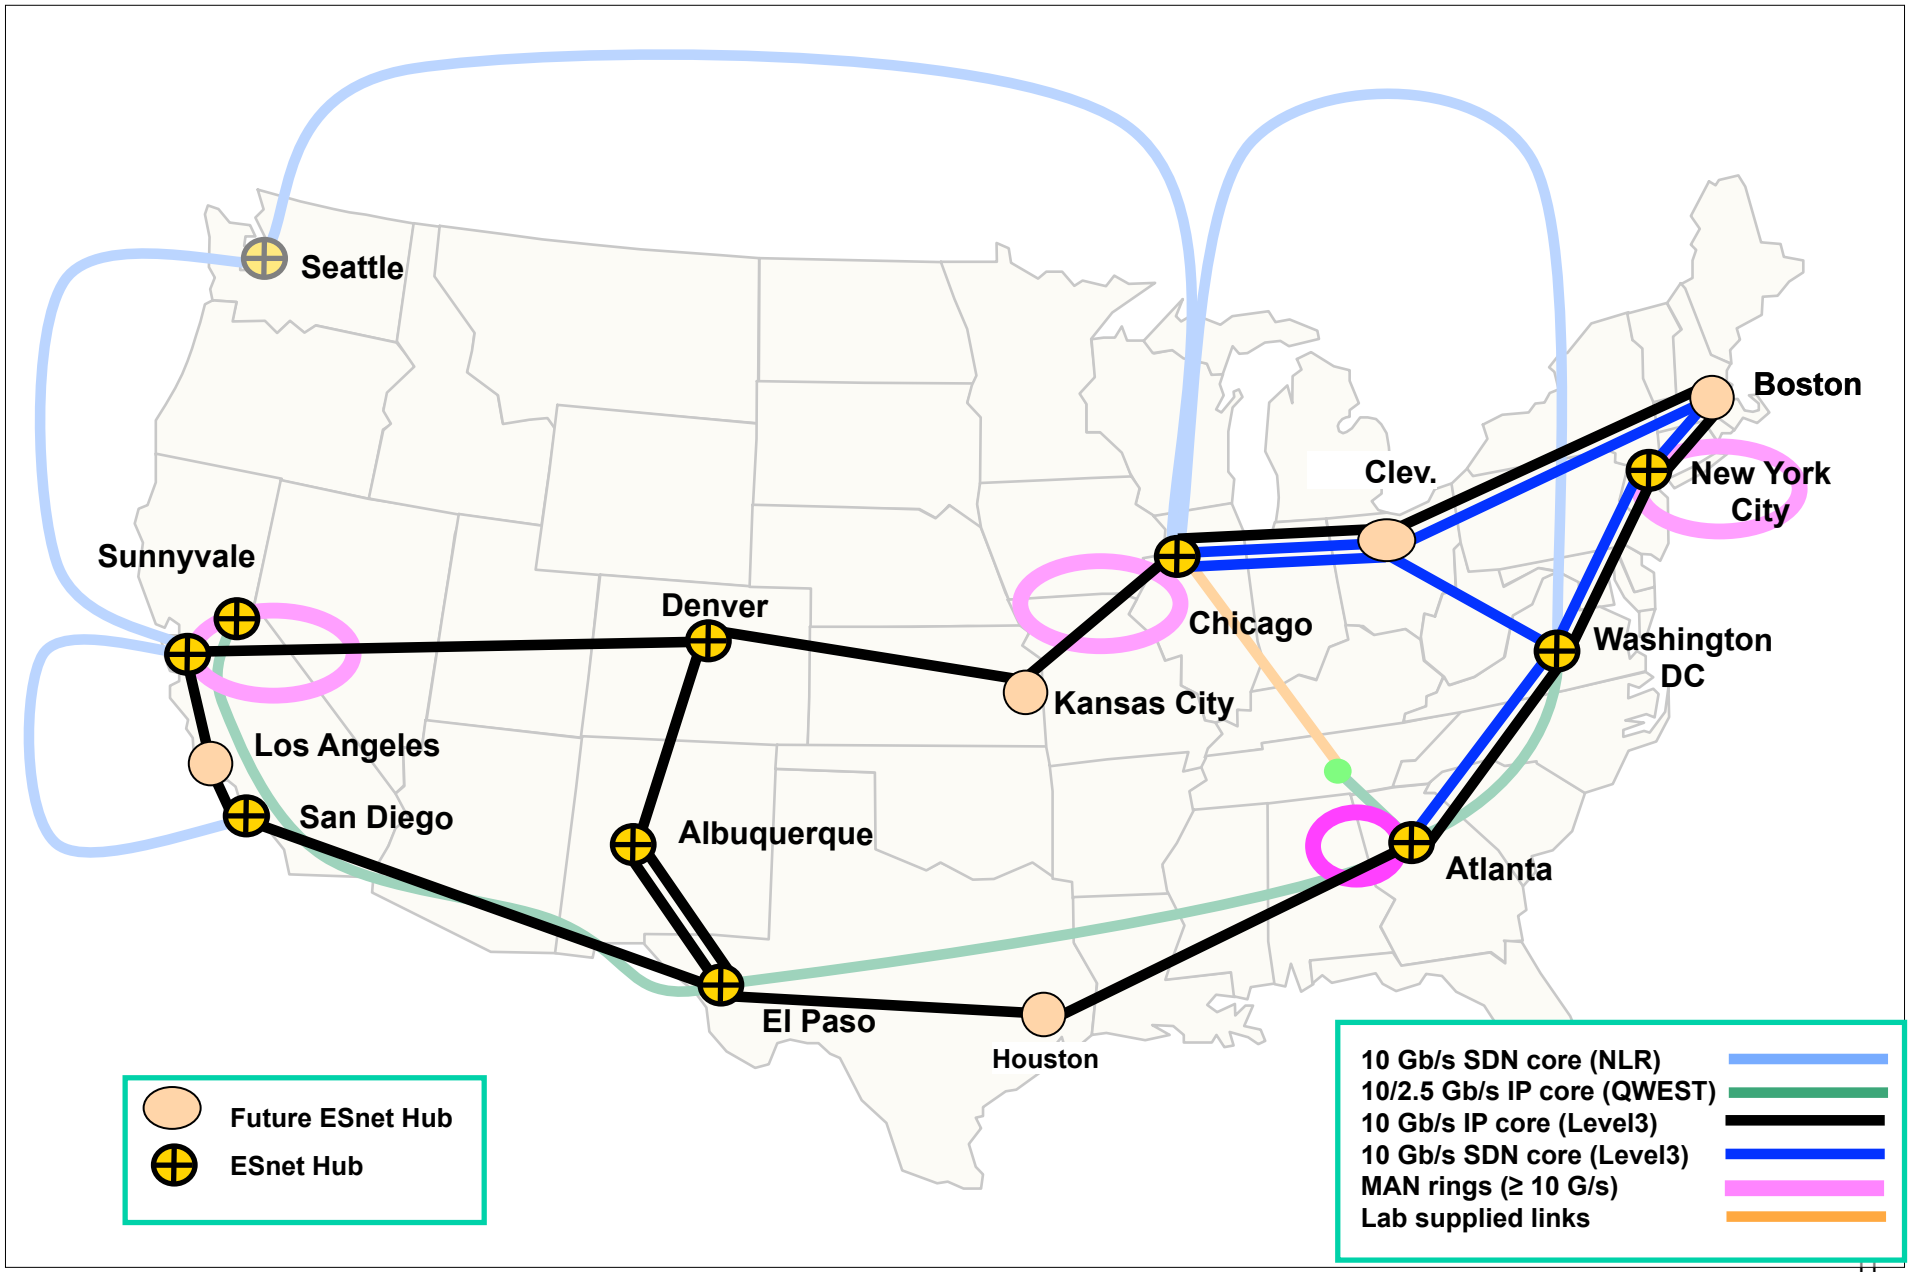

ESnet 4 Backbone September 30, 2007

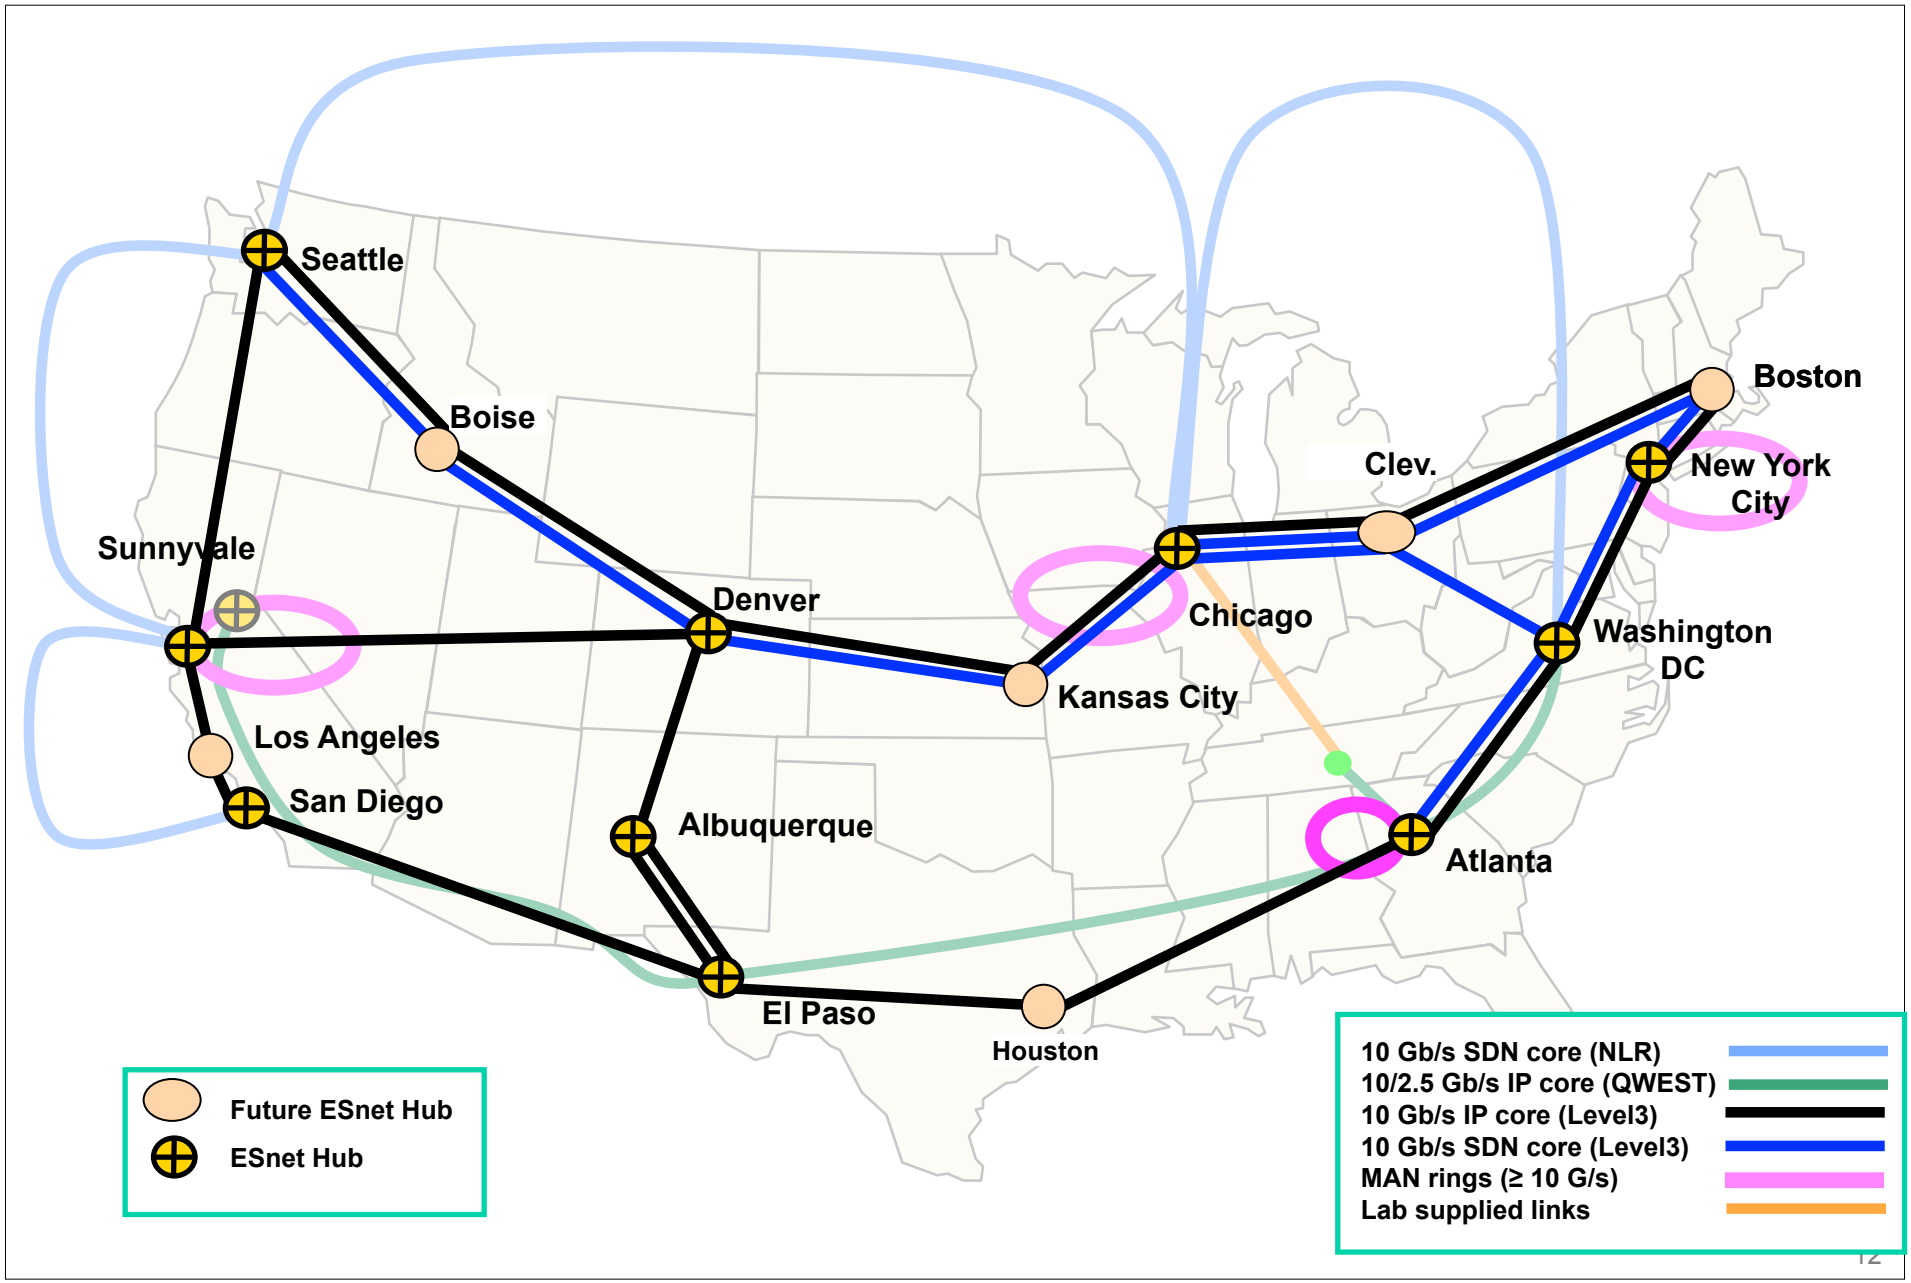

ESnet 4 Backbone December 2007

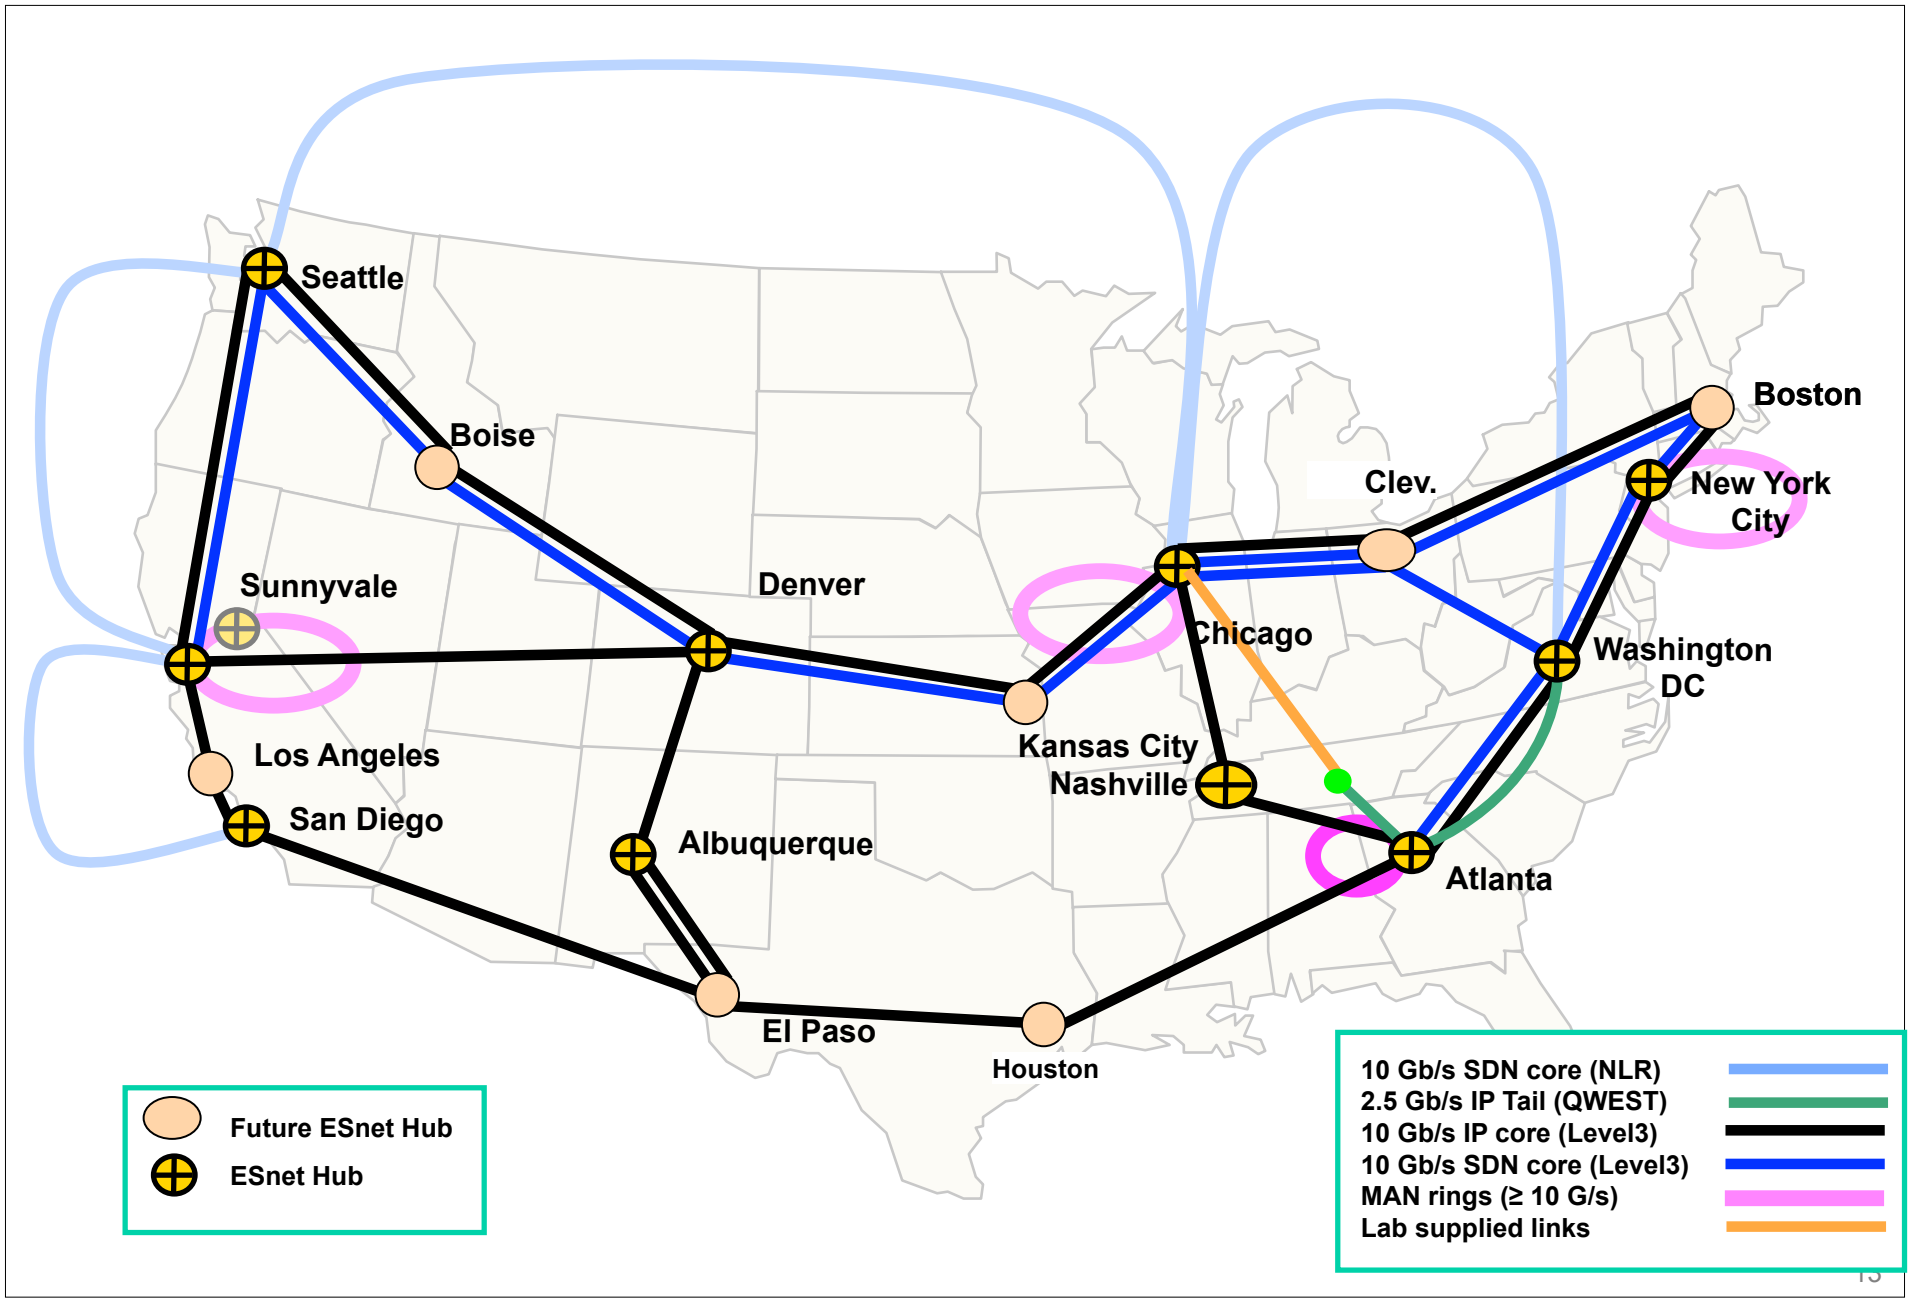

ESnet 4 Backbone Projected for December, 2008

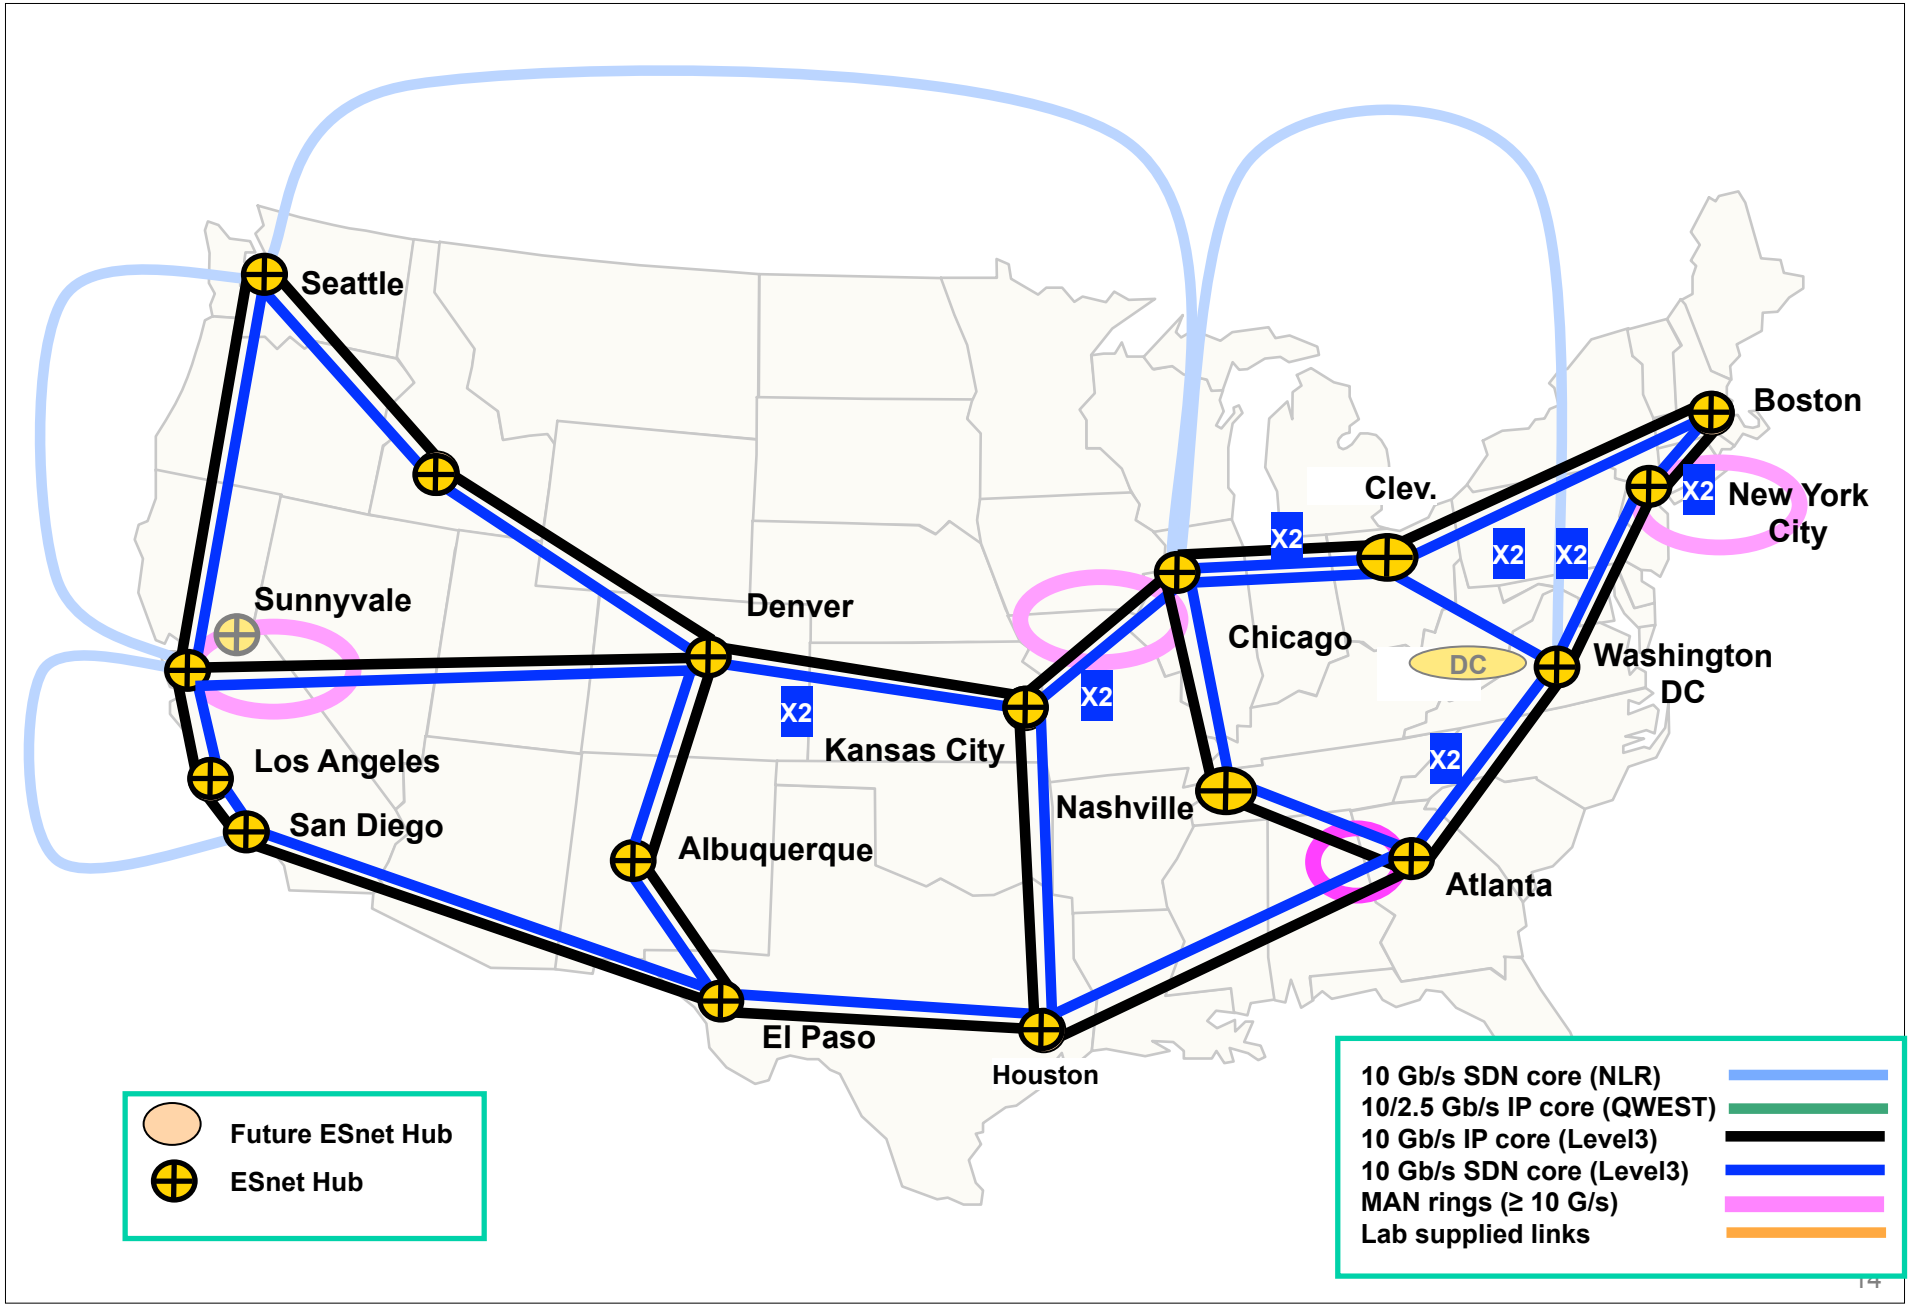

ESnet Provides Global High-Speed Internet Connectivity for DOE Facilities and Collaborators (12/2007)

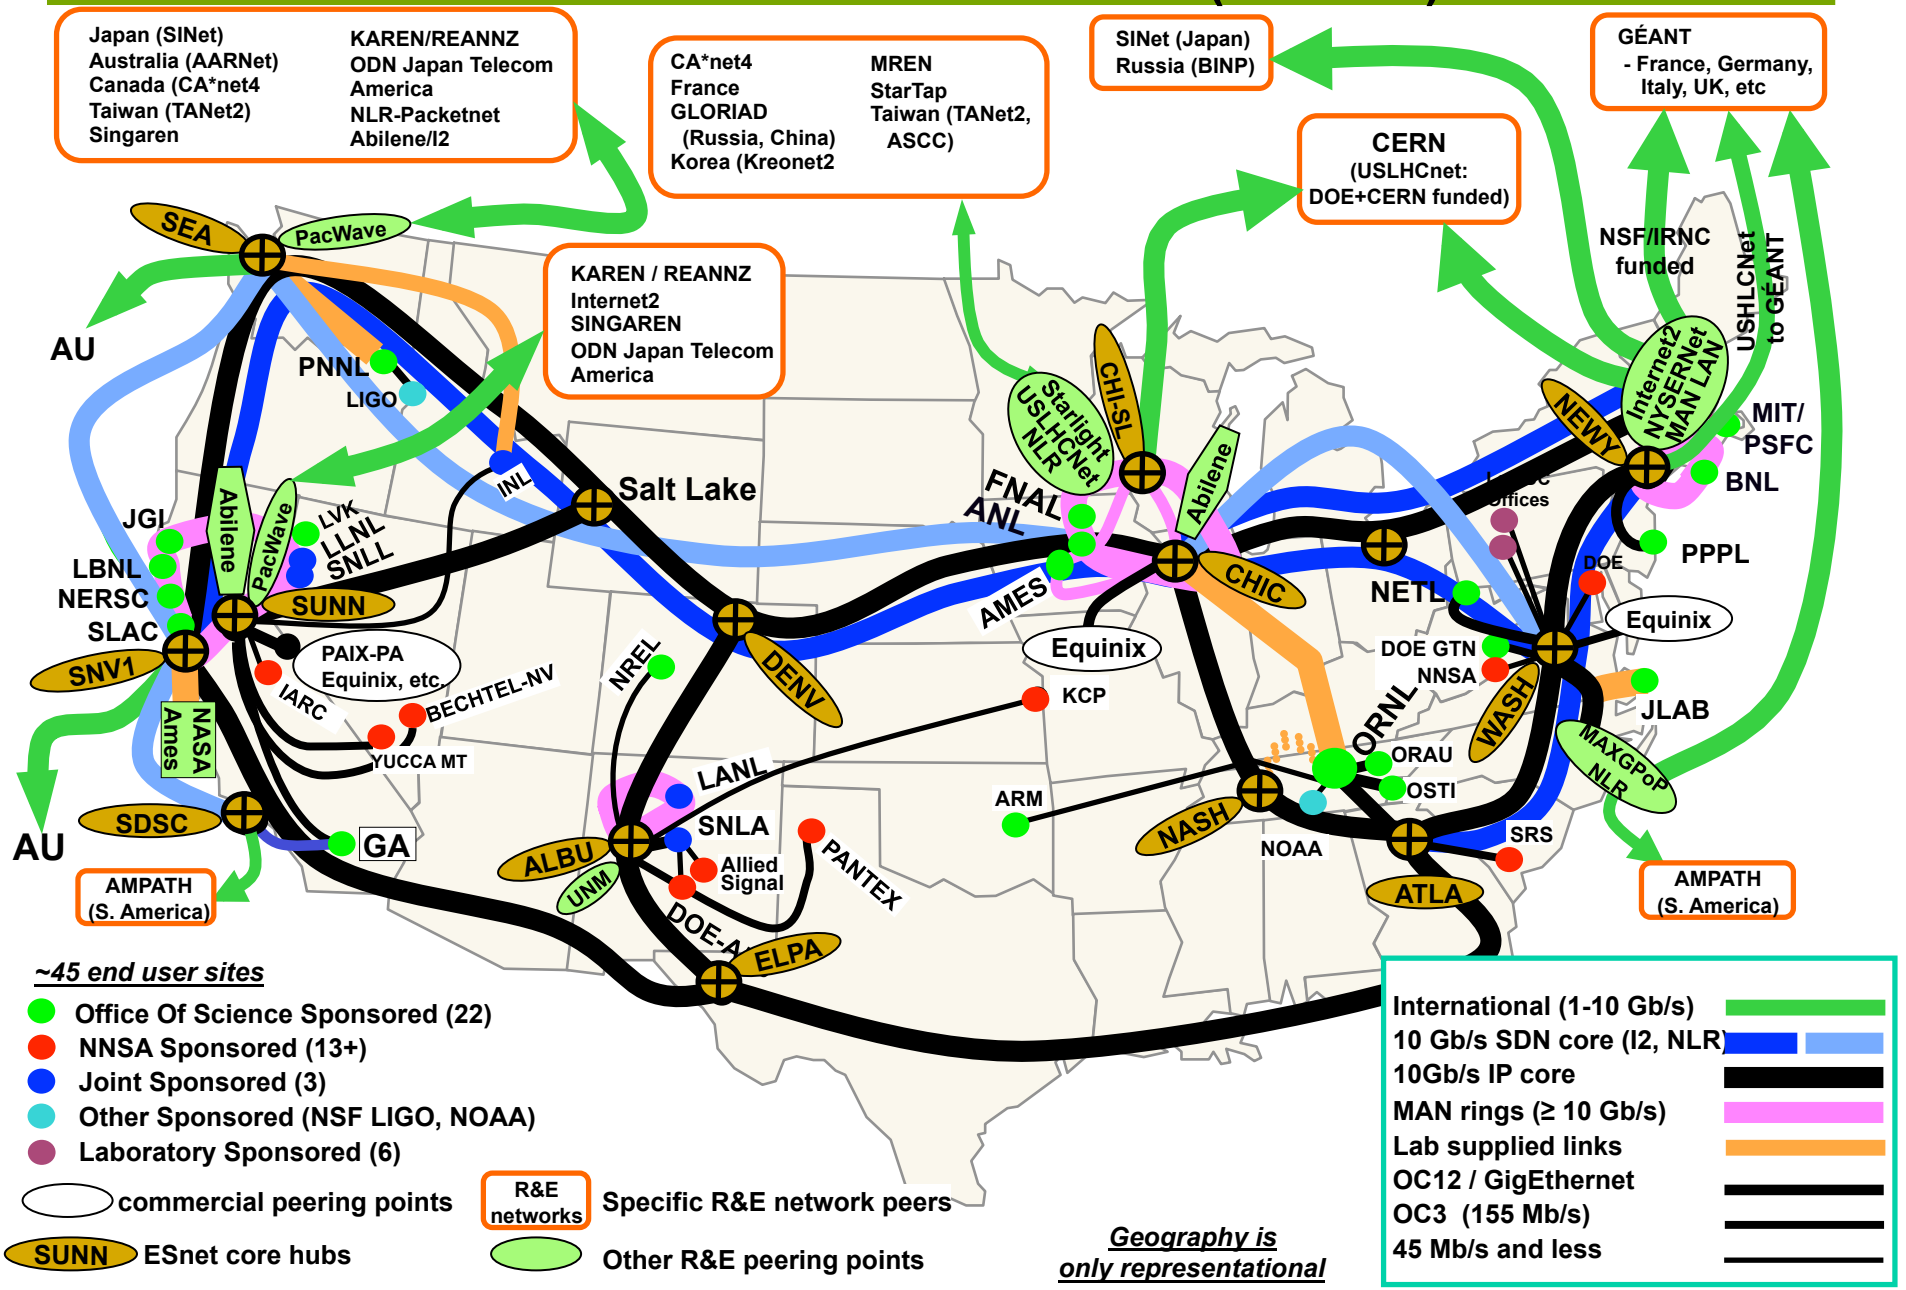

ESnet4 End-Game

Core networks 50-60 Gbps by 2009-2010 (10Gb/s circuits),
500-600 Gbps by 2011-2012 (100 Gb/s circuits)

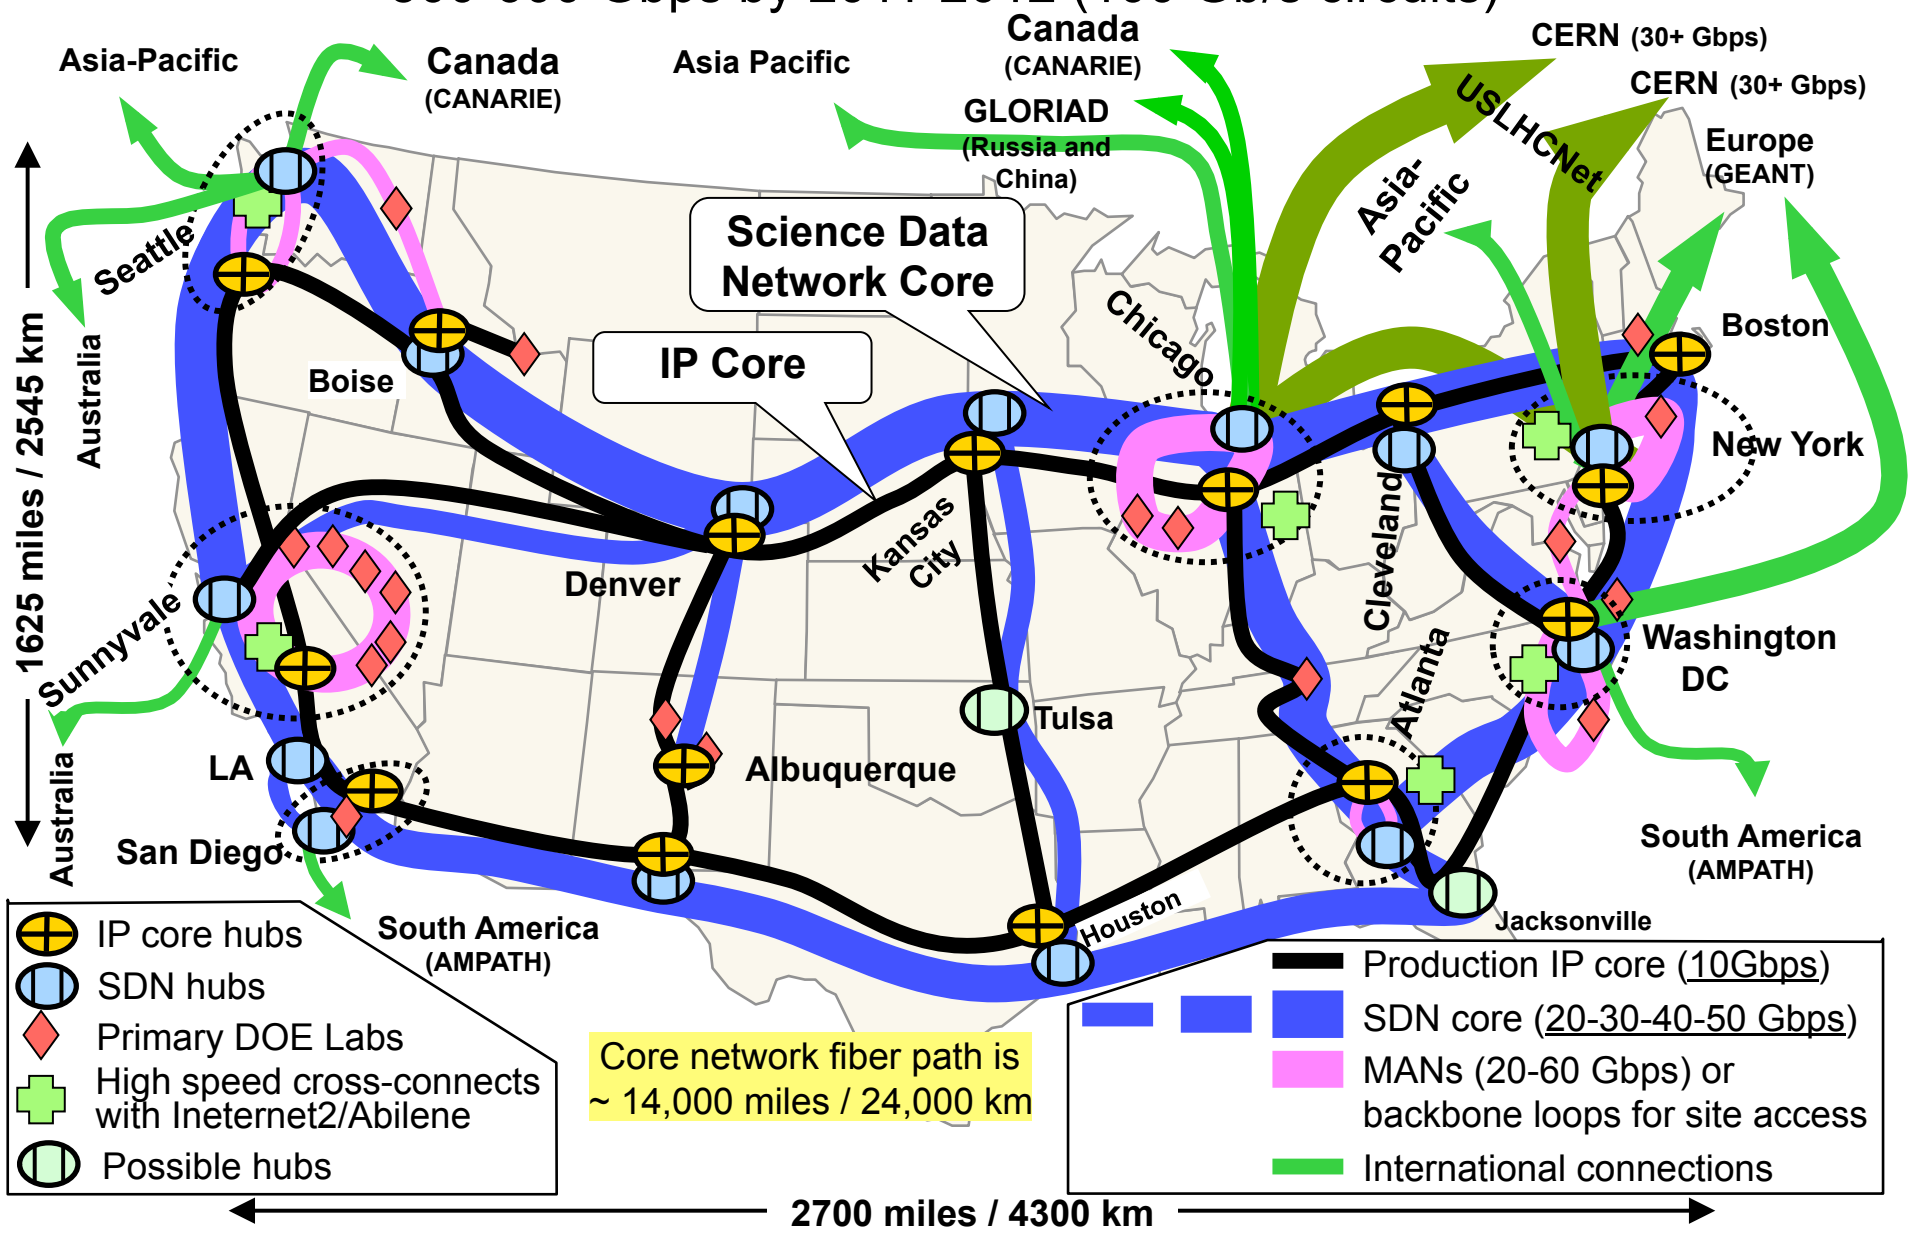

A Tail of Two ESnet4 Hubs

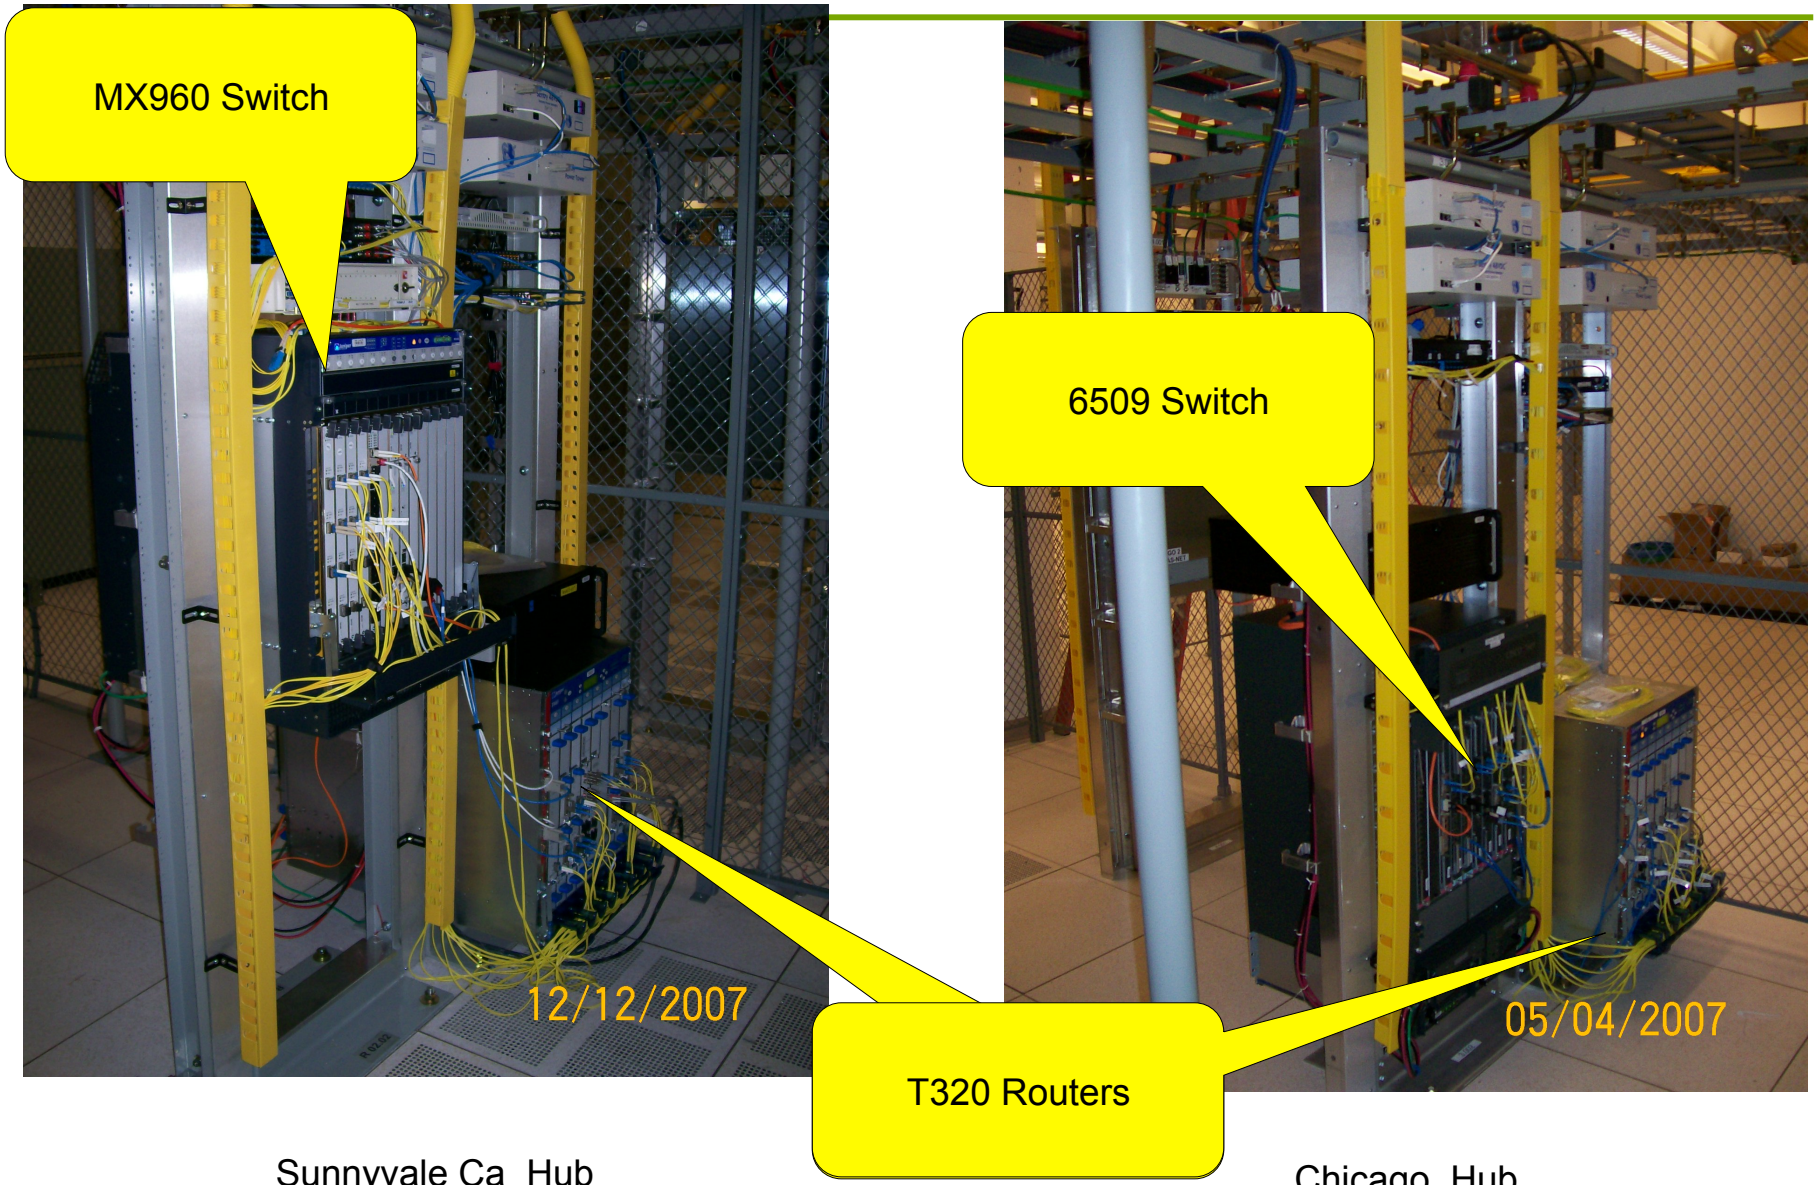

Sunnyvale Ca Hub

Chicago Hub

ESnet's SDN backbone is implemented with Layer2 switches; Cisco 6509s and Juniper MX960s - Each present their own unique challenges.

ESnet 4 Factoids as of January 21, 2008

- ESnet4 installation to date:
  - 32 new 10Gb/s backbone circuits
    - Over 3 times the number from last JT meeting
  - 20,284 10Gb/s backbone Route Miles
    - More than doubled from last JT meeting
  - 10 new hubs
    - Since last meeting
      - Seattle
      - Sunnyvale
      - Nashville
  - 7 new routers 4 new switches
  - Chicago MAN now connected to Level3 POP
    - 2 x 10GE to ANL
    - 2 x 10GE to FNAL
    - 3 x 10GE to Starlight

ESnet Traffic Continues to Exceed 2 Petabytes/Month

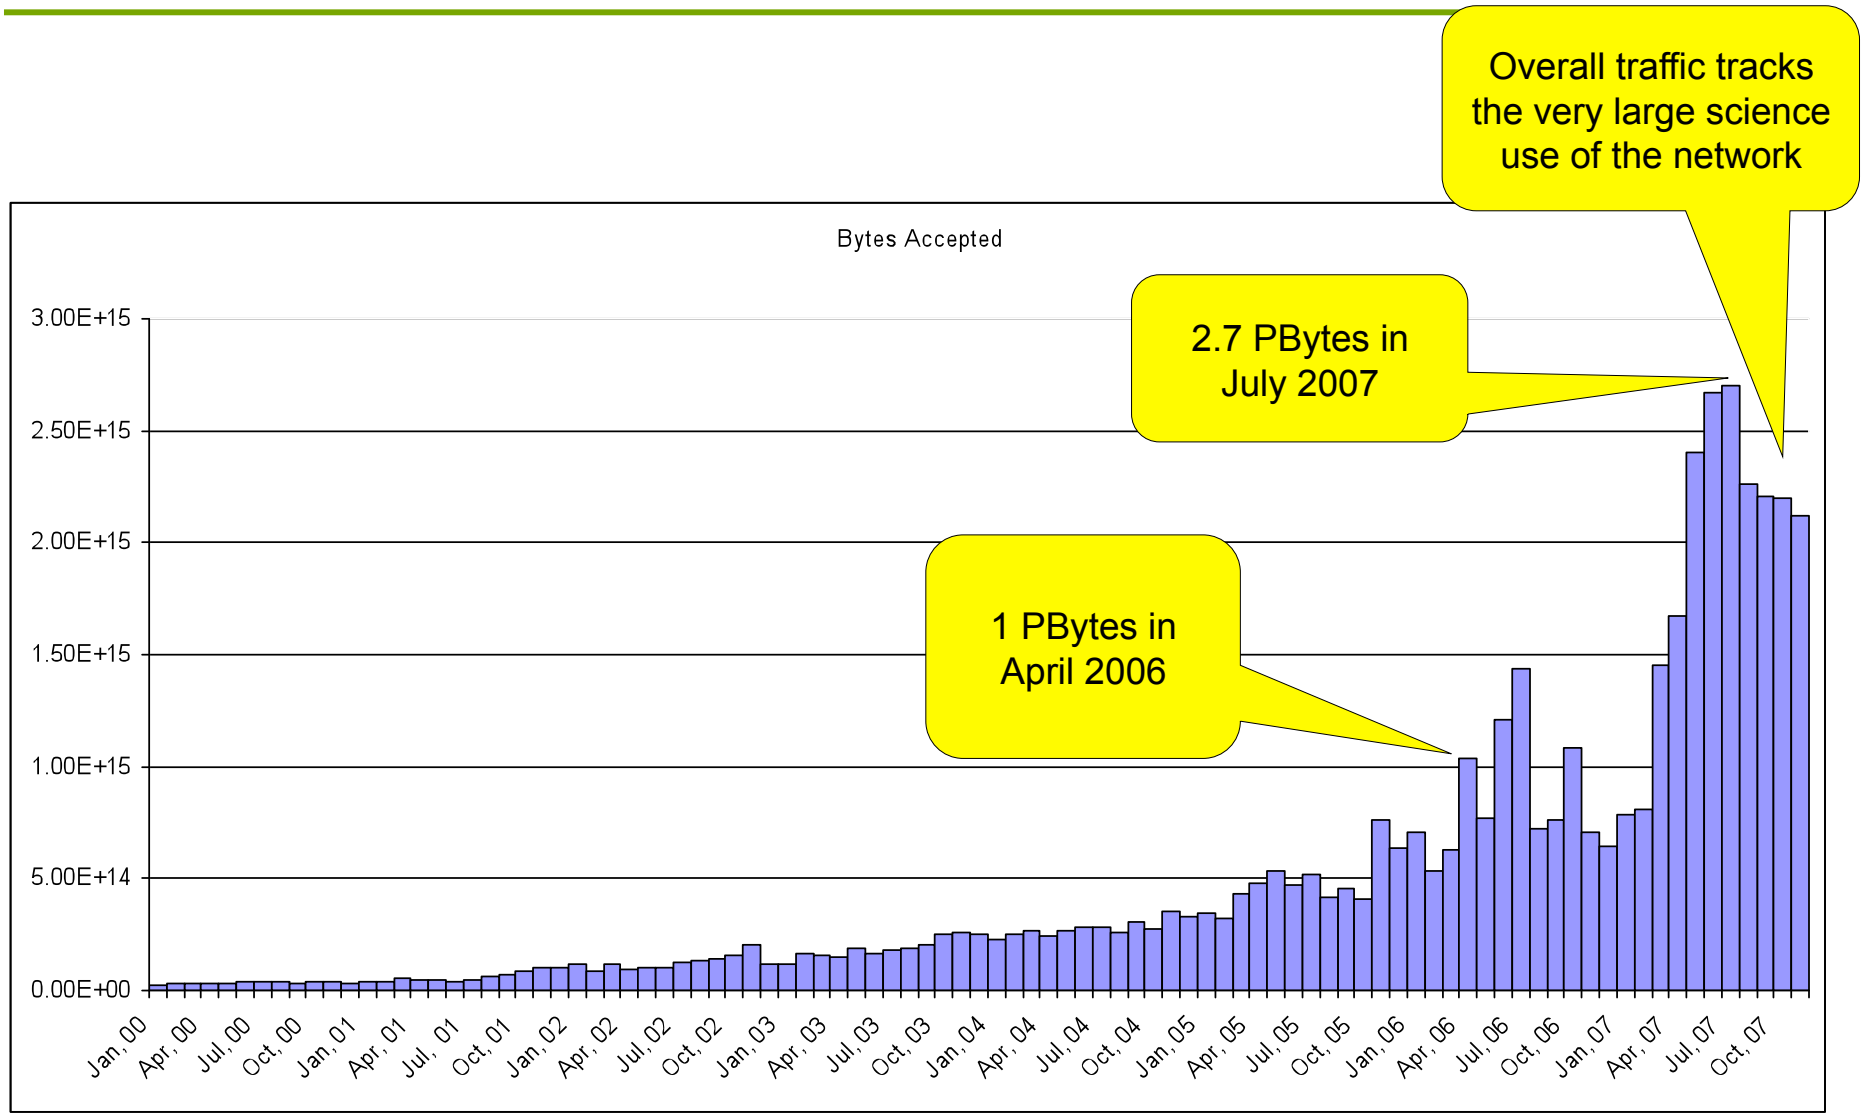

ESnet traffic historically has increased 10x every 47 months

When A Few Large Data Sources/Sinks Dominate Traffic it is Not Surprising that Overall Network Usage Follows the Patterns of the Very Large Users - This Trend Will Reverse in the Next Few Weeks as the Next Round of LHC Data Challenges Kicks Off

Fermilab

Fermi LightPath

FNAL Outbound Traffic

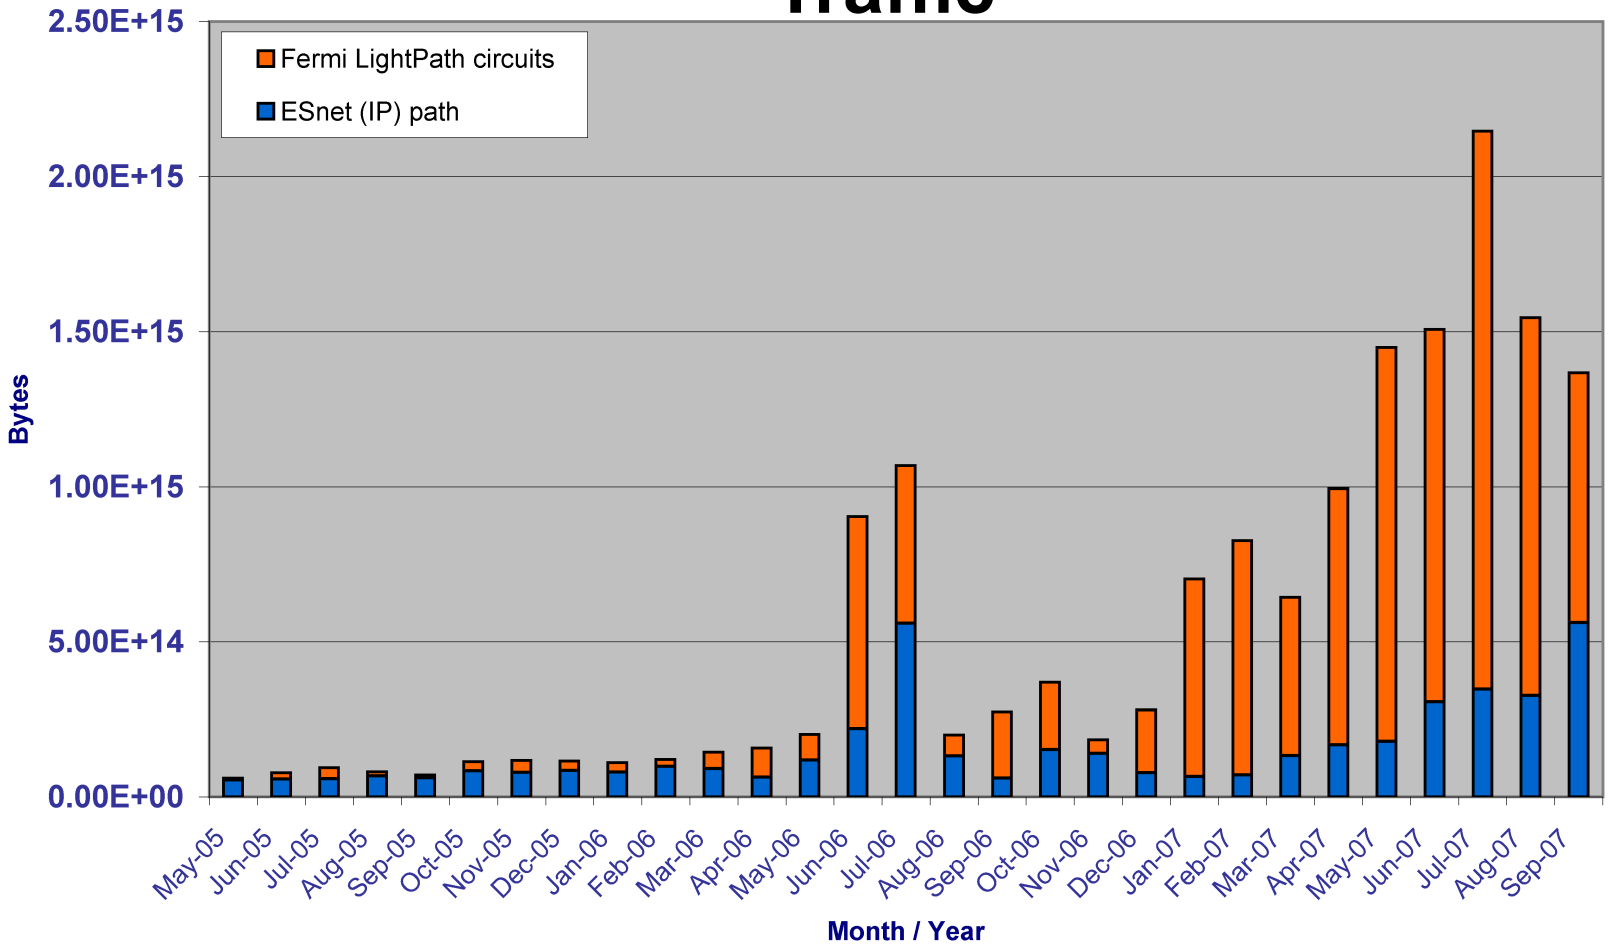

FNAL Traffic is Representative of all CMS Traffic

Accumulated data (Terabytes) received by CMS Data Centers ("tier1" sites) and many analysis centers ("tier2" sites) during the past 12 months (15 petabytes of data) [LHC/CMS]

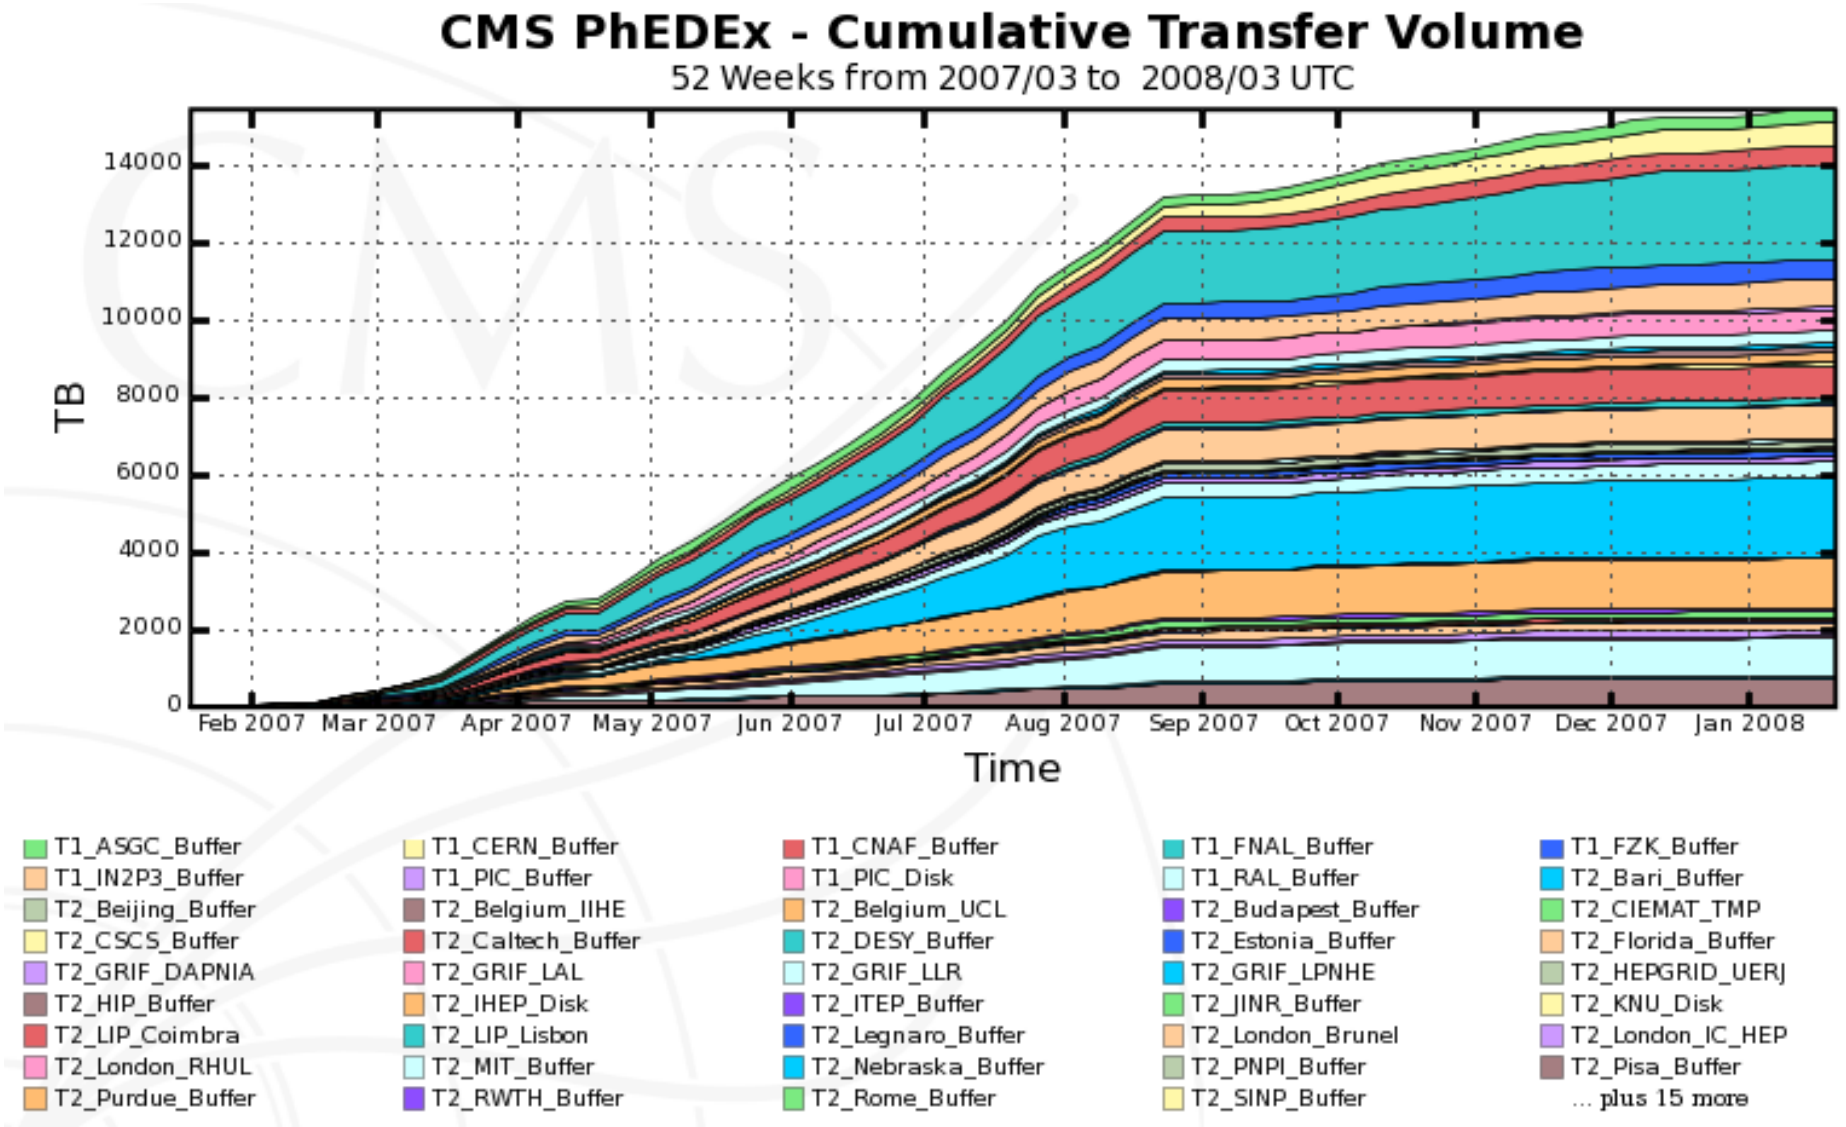

ESnet Continues to be Highly Reliable; Even During the Transition

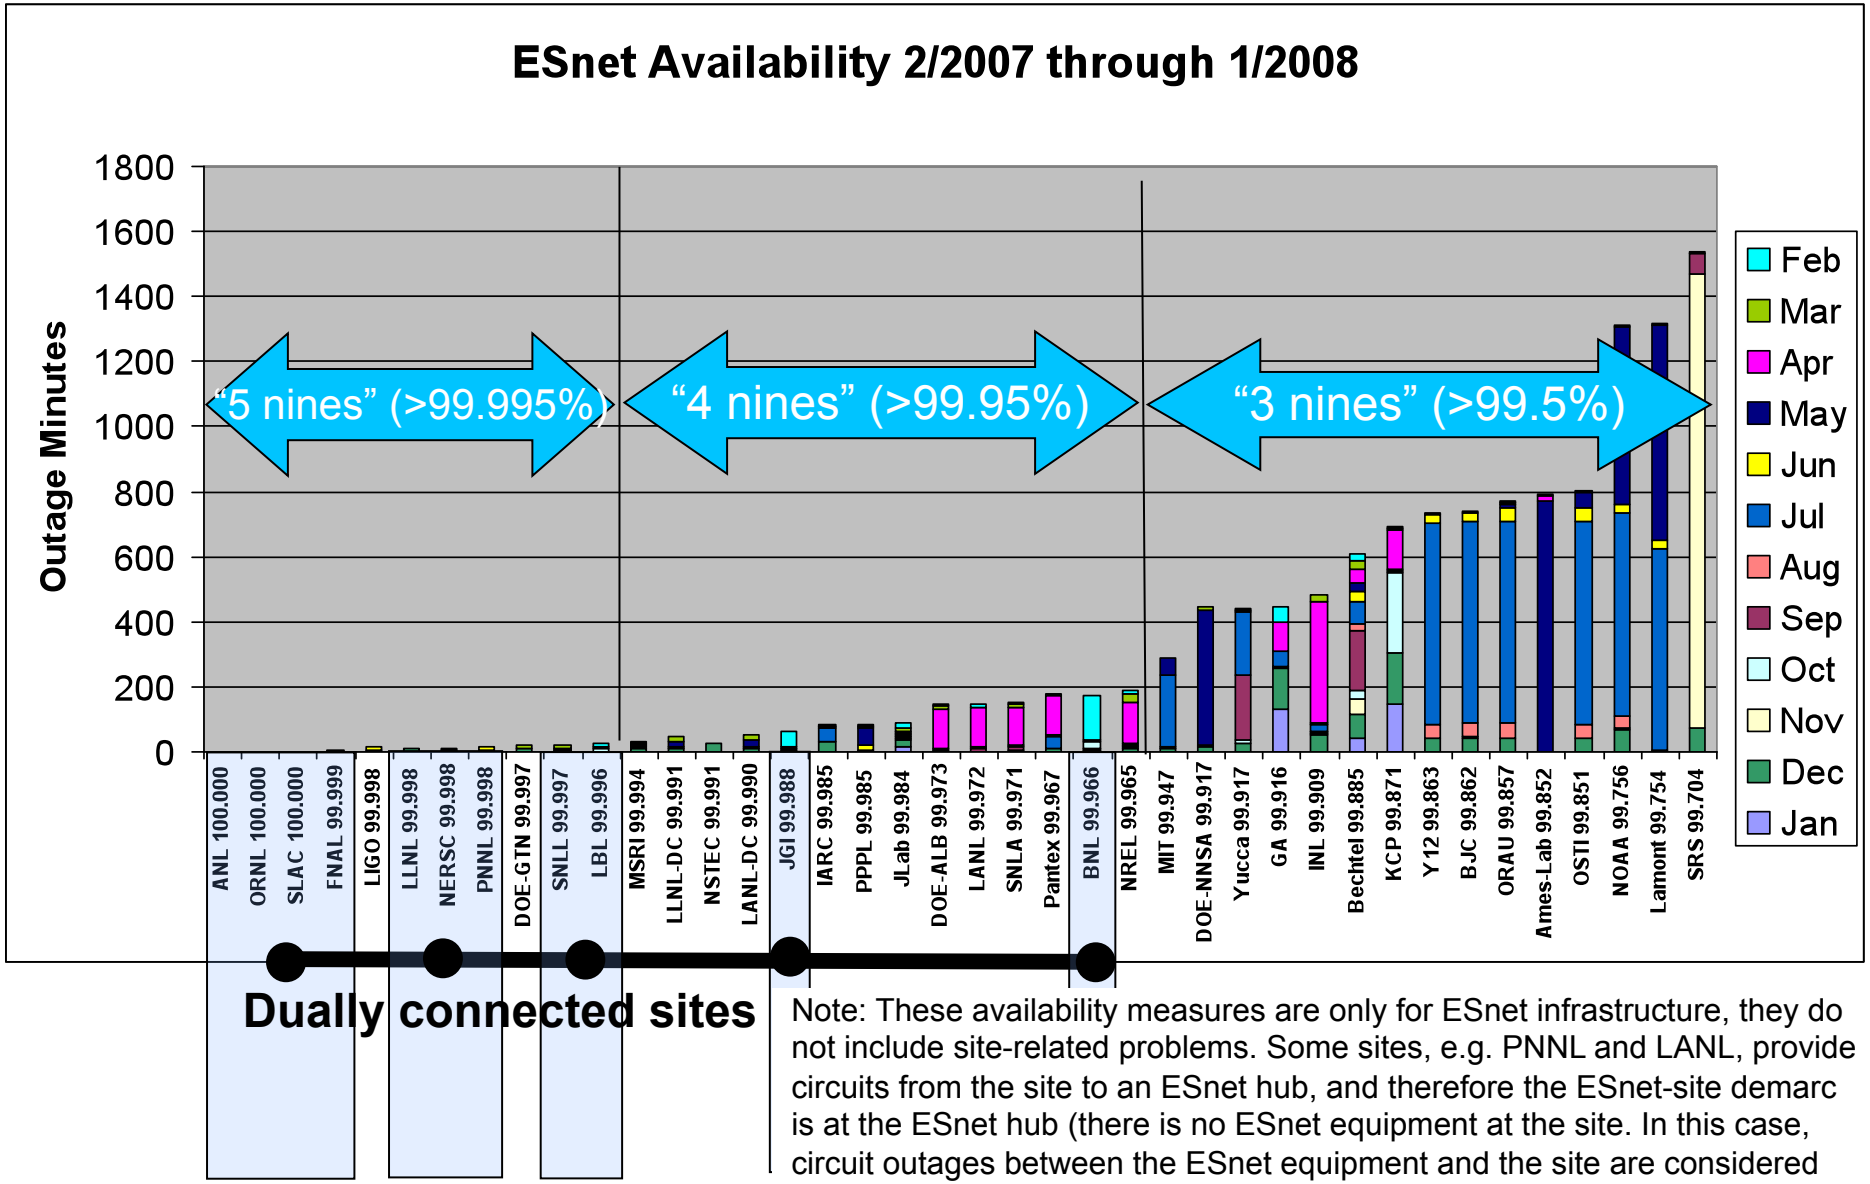

Note: These availability measures are only for ESnet infrastructure, they do not include site-related problems. Some sites, e.g. PNNL and LANL, provide circuits from the site to an ESnet hub, and therefore the ESnet-site demarc is at the ESnet hub (there is no ESnet equipment at the site. In this case, circuit outages between the ESnet equipment and the site are considered site issues and are not included in the ESnet availability metric.

Ib. Network Services for Large-Scale Science

- Large-scale science uses distributed system in order to:
  - Couple existing pockets of code, data, and expertise into a "system of systems"
  - Break up the task of massive data analysis into elements that are physically located where the data, compute, and storage resources are located - these elements are combined into a system using a "Service Oriented Architecture" approach
- Such ***systems***
  - ***are data intensive and high-performance***, typically moving terabytes a day for months at a time
  - ***are high duty-cycle***, operating most of the day for months at a time in order to meet the requirements for data movement
  - ***are widely distributed*** – typically spread over continental or inter-continental distances
  - ***depend on network performance and availability***, but these characteristics cannot be taken for granted, even in well run networks, when the multi-domain network path is considered
- The system elements ***must be able to get guarantees from the network*** that there is adequate bandwidth to accomplish the task at hand
- The systems ***must be able to get information from the network*** that allows graceful failure and auto-recovery and adaptation to unexpected network conditions that are short of outright failure

See, e.g., [ICFA SCIC]

Enabling Large-Scale Science

- These requirements are generally true for systems with widely distributed components to be reliable and consistent in performing the sustained, complex tasks of large-scale science

➢ **Networks must provide communication capability as a service that can participate in SOA**

- **configurable**
- **schedulable**
- **predictable**
- **reliable**
- **informative**
- **and the network and its services must be scalable and geographically comprehensive**

Networks Must Provide Communication Capability that is Service-Oriented

- Configurable
  - Must be able to provide multiple, specific “paths” (specified by the user as end points) with specific characteristics
- Schedulable
  - Premium service such as guaranteed bandwidth will be a scarce resource that is not always freely available, therefore time slots obtained through a resource allocation process must be schedulable
- Predictable
  - A committed time slot should be provided by a network service that is not brittle - reroute in the face of network failures is important
- Reliable
  - Reroutes should be largely transparent to the user
- Informative
  - When users do system planning they should be able to see average path characteristics, including capacity
  - When things do go wrong, the network should report back to the user in ways that are meaningful to the user so that informed decisions can about alternative approaches
- Scalable
  - The underlying network should be able to manage its resources to provide the appearance of scalability to the user
- Geographically comprehensive
  - The R&E network community must act in a coordinated fashion to provide this environment end-to-end

The ESnet Approach

- Provide *configurability, schedulability, predictability*, and reliability with a flexible virtual circuit service - OSCARS
  - User* specifies end points, bandwidth, and schedule
  - OSCARS can do fast reroute of the underlying MPLS paths
- Provide useful, comprehensive, and meaningful *information on the state of the paths*, or potential paths, to the user
  - perfSONAR, and associated tools, provide real time information in a form that is useful to the user (via appropriate network abstractions) and that is delivered through standard interfaces that can be incorporated in to SOA type applications
  - Techniques need to be developed to monitor virtual circuits based on the approaches of the various R&E nets - e.g. MPLS in ESnet, VLANs, TDM/grooming devices (e.g. Ciena Core Directors), etc., and then integrate this into a perfSONAR framework

* User = human or system component (process)

The ESnet Approach

- *Scalability* will be provided by new network services that, e.g., provide dynamic wave allocation at the optical layer of the network
  - Currently an R&D project
- *Geographic ubiquity* of the services can only be accomplished through active collaborations in the global R&E network community so that all sites of interest to the science community can provide compatible services for forming end-to-end virtual circuits
  - Active and productive collaborations exist among numerous R&E networks: ESnet, Internet2, CANARIE, DANTE/GÉANT, some European NRENs, some US regionals, etc.

OSCARS Overview

On-demand Secure Circuits and Advance Reservation System

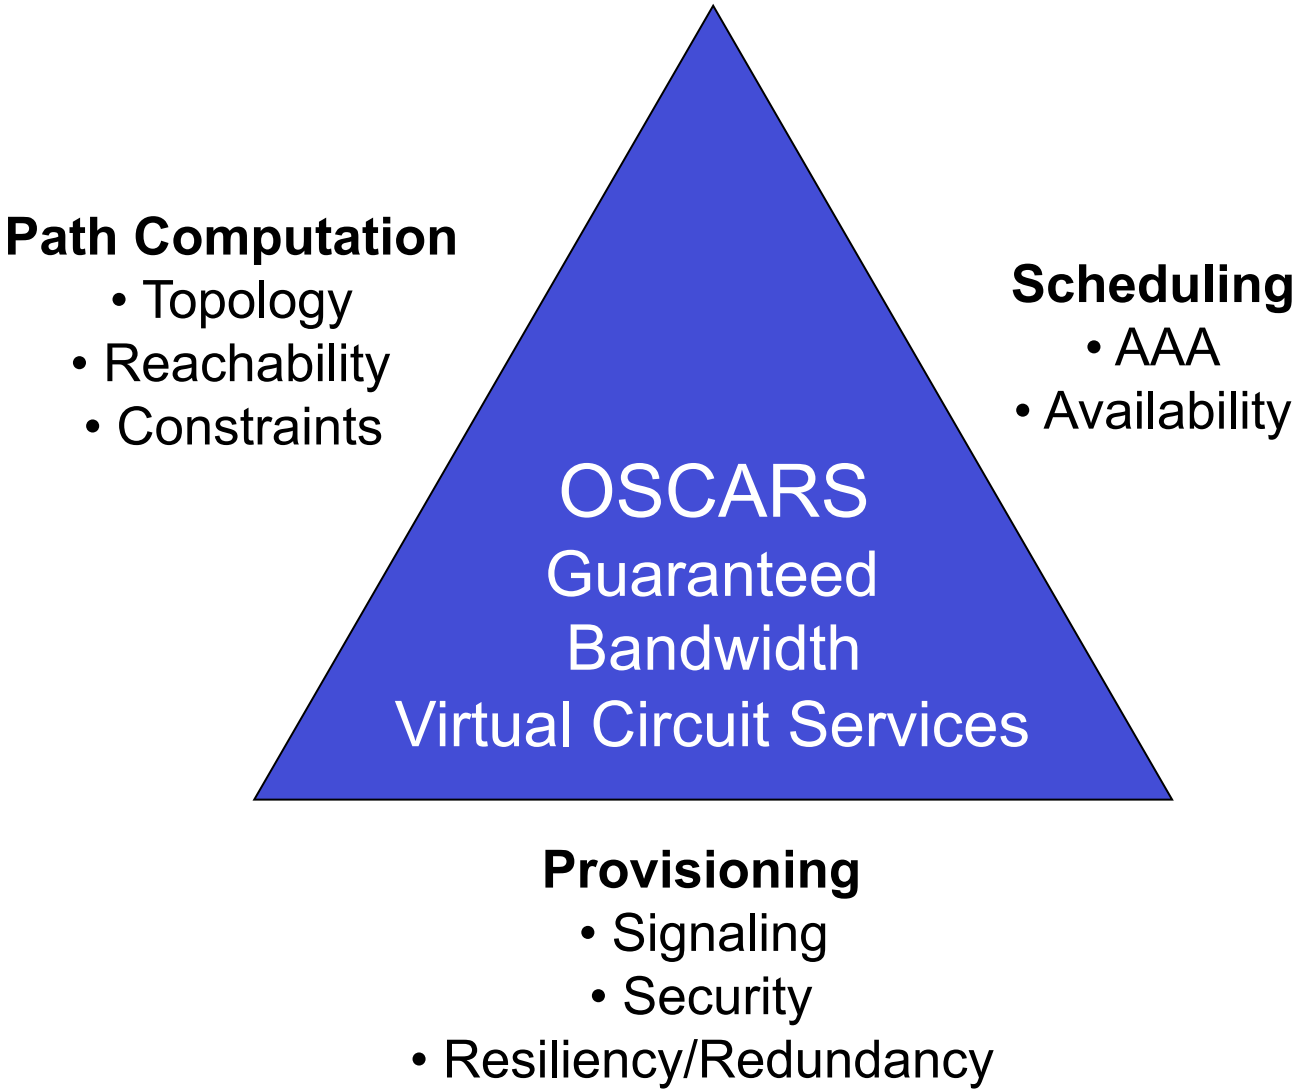

OSCARS Status Update

- ESnet Centric Deployment
  - Prototype layer 3 (IP) guaranteed bandwidth virtual circuit service deployed in ESnet (1Q05)
  - Prototype layer 2 (Ethernet VLAN) virtual circuit service deployed in ESnet (3Q07)
- Inter-Domain Collaborative Efforts
  - Terapaths (BNL)
    - Inter-domain interoperability for layer 3 virtual circuits demonstrated (3Q06)
    - Inter-domain interoperability for layer 2 virtual circuits demonstrated at SC07 (4Q07)
  - LambdaStation (FNAL)
    - Inter-domain interoperability for layer 2 virtual circuits demonstrated at SC07 (4Q07)
  - HOPI/DRAGON
    - Inter-domain exchange of control messages demonstrated (1Q07)
    - Integration of OSCARS and DRAGON has been successful (1Q07)
  - DICE
    - First draft of topology exchange schema has been formalized (in collaboration with NMWG) (2Q07), interoperability test demonstrated 3Q07
    - Initial implementation of reservation and signaling messages demonstrated at SC07 (4Q07)
  - UVA
    - Integration of Token based authorization in OSCARS under testing
  - Nortel
    - Topology exchange demonstrated successfully 3Q07
    - Inter-domain interoperability for layer 2 virtual circuits demonstrated at SC07 (4Q07)

Ic. Network Measurement Update

- Deploy network test platforms at all hubs and major sites
  - About 1/3 of the 10GE bandwidth test platforms & 1/2 of the latency test platforms for ESnet 4 have been deployed.
  - 10GE test systems are being used extensively for acceptance testing and debugging
  - Structured & ad-hoc external testing capabilities have not been enabled yet.
  - Clocking issues at a couple POPS are not resolved.
- Work is progressing on revamping the ESnet statistics collection, management & publication systems
  - ESxSNMP & TSDB & PerfSONAR Measurement Archive (MA)
  - PerfSONAR TS & OSCARS Topology DB
  - NetInfo being restructured to be PerfSONAR based

Network Measurement Update

- PerfSONAR provides a service element oriented approach to monitoring that has the potential to integrate into SOA
  - See Joe Metzger's talk

II. SC Program Network Requirements Workshops

- The Workshops are part of DOE's governance of ESnet
  - The ASCR Program Office owns the requirements workshops, not ESnet
  - The Workshops replaced the ESnet Steering Committee
  - The workshops are fully controlled by DOE....all that ESnet does is to support DOE in putting on the workshops
    - The content and logistics of the workshops is determined by an SC Program Manager from the Program Office that is the subject of the each workshop
      - SC Program Office sets the timing, location (almost always Washington so that DOE Program Office people can attend), and participants

Network Requirements Workshops

- Collect requirements from two DOE/SC program offices per year
- DOE/SC Program Office workshops held in 2007
  - Basic Energy Sciences (BES) – June 2007
  - Biological and Environmental Research (BER) – July 2007
- Workshops to be held in 2008
  - Fusion Energy Sciences (FES) – Coming in March 2008
  - Nuclear Physics (NP) – TBD 2008
- Future workshops
  - HEP and ASCR in 2009
  - BES and BER in 2010
  - And so on…

Network Requirements Workshops - Findings

- Virtual circuit services (traffic isolation, bandwidth guarantees, etc) continue to be requested by scientists
  - OSCARS service directly addresses these needs
    - http://www.es.net/OSCARS/index.html
    - Successfully deployed in early production today
    - ESnet will continue to develop and deploy OSCARS
- Some user communities have significant difficulties using the network for bulk data transfer
  - fasterdata.es.net – web site devoted to bulk data transfer, host tuning, etc. established
  - NERSC and ORNL have made significant progress on improving data transfer performance between supercomputer centers

Network Requirements Workshops - Findings

- Some data rate requirements are unknown at this time
  - Drivers are instrument upgrades that are subject to review, qualification and other decisions that are 6-12 months away
  - These will be revisited in the appropriate timeframe

BES Workshop Bandwidth Matrix as of June 2007

Project	Primary Site	Primary Partner Sites	Primary ESnet	2007 Bandwidth	2012 Bandwidth
ALS	LBNL	Distributed	Sunnyvale	3 Gbps	10 Gbps
APS, CNM, SAMM, ARM	ANL	FNAL, BNL, UCLA, and CERN	Chicago	10 Gbps	20 Gbps
Nano Center	BNL	Distributed	NYC	1 Gbps	5 Gbps
CRF	SNL/CA	NERSC, ORNL	Sunnyvale	5 Gbps	10 Gbps
Molecular Foundry	LBNL	Distributed	Sunnyvale	1 Gbps	5 Gbps
NCEM	LBNL	Distributed	Sunnyvale	1 Gbps	5 Gbps
LCLF	SLAC	Distributed	Sunnyvale	2 Gbps	4 Gbps
NSLS	BNL	Distributed	NYC	1 Gbps	5 Gbps
SNS	ORNL	LANL, NIST, ANL, U. Indiana	Nashville	1 Gbps	10 Gbps
Total				25 Gbps	74 Gbps

BER Workshop Bandwidth Matrix as of Dec 2007

Project	Primary Site	Primary Partner Sites	Primary ESnet	2007 Bandwidth	2012 Bandwidth
ARM	**BNL, ORNL, PNNL**	**NOAA, NASA, ECMWF (Europe), Climate Science**	**NYC, Nashville, Seattle**	**1 Gbps**	**5 Gbps**
Bioinformatics	**PNNL**	**Distributed**	**Seattle**	**.5 Gbps**	**3 Gbps**
EMSL	**PNNL**	**Distributed**	**Seattle**	**10 Gbps**	**50 Gbps**
Climate	**LLNL, NCAR, ORNL**	**NCAR, LANL, NERSC, LLNL, International**	**Sunnyvale, Denver, Nashville**	**1 Gbps**	**5 Gbps**
JGI	**JGI**	**NERSC**	**Sunnyvale**	**1 Gbps**	**5 Gbps**
Total				**13.5 Gbps**	**68 Gbps**

ESnet Site Network Requirements Surveys

- Surveys given to ESnet sites through ESCC
- Many sites responded, many did not
- Survey was lacking in several key areas
  - Did not provide sufficient focus to enable consistent data collection
  - Sites vary widely in network usage, size, science/business, etc… very difficult to make one survey fit all
  - In many cases, data provided not quantitative enough (this appears to be primarily due to the way in which the questions were asked)
- Surveys were successful in some key ways
  - It is clear that there are many significant projects/programs that cannot be captured in the DOE/SC Program Office workshops
  - DP, industry, other non-SC projects
  - Need better process to capture this information
- New model for site requirements collection needs to be developed

IIIa. Federated Trust Services

- Remote, multi-institutional, identity authentication is critical for distributed, collaborative science in order to permit sharing widely distributed computing and data resources, and other Grid services
- Public Key Infrastructure (PKI) is used to formalize the existing web of trust within science collaborations and to extend that trust into cyber space
  - The function, form, and policy of the ESnet trust services are driven entirely by the requirements of the science community and by direct input from the science community
  - International scope trust agreements that encompass many organizations are crucial for large-scale collaborations
- The service (and community) has matured to the point where it is revisiting old practices and updating and formalizing them

DOEGrids CA Audit

- "Request" by EUGridPMA
  - EUGridPMA is auditing all "old" CAs
- OGF Audit Framework
  - Developed from WebTrust for CAs & al
- Partial review of NIST 800-53
- Audit Day 11 Dec 2007 – Auditors:

Robert Cowles (SLAC)	Dan Peterson (ESnet)	Mary Thompson (ex-LBL)
John Volmer (ANL)	Scott Rea (HEBCA*)(obsrv)	

*** Higher Education Bridge Certification Authority**

The goal of the Higher Education Bridge Certification Authority (HEBCA) is to facilitate trusted electronic communications within and between institutions of higher education as well as with federal and state governments.

DOEGrids CA Audit – Results

- Final report in progress
- Generally good – many documentation errors need to be addressed
- EUGridPMA is satisfied

➢ EUGridPMA has agreed to recognize US research science ID verification as acceptable for initial issuance of certificate

  – This is a BIG step forward

- The ESnet CA projects have begun a year-long effort to converge security documents and controls with NIST 800-53

DOEGrids CA Audit – Issues

- ID verification – no face to face/ID doc check
  - We have collectively agreed to drop this issue – US science culture is what it is, and has a method for verifying identity
- Renewals – we must address the need to re-verify our subscribers after 5 years
- Auditors recommend we update the format of our Certification Practices Statement (for interoperability and understandability)
- Continue efforts to improve reliability & disaster recovery
- We need to update our certificate formats again (minor errors)
- There are many undocumented or incompletely documented security practices (a problem both in the CPS and NIST 800-53)

DOEGrids CA (one of several CAs) Usage Statistics

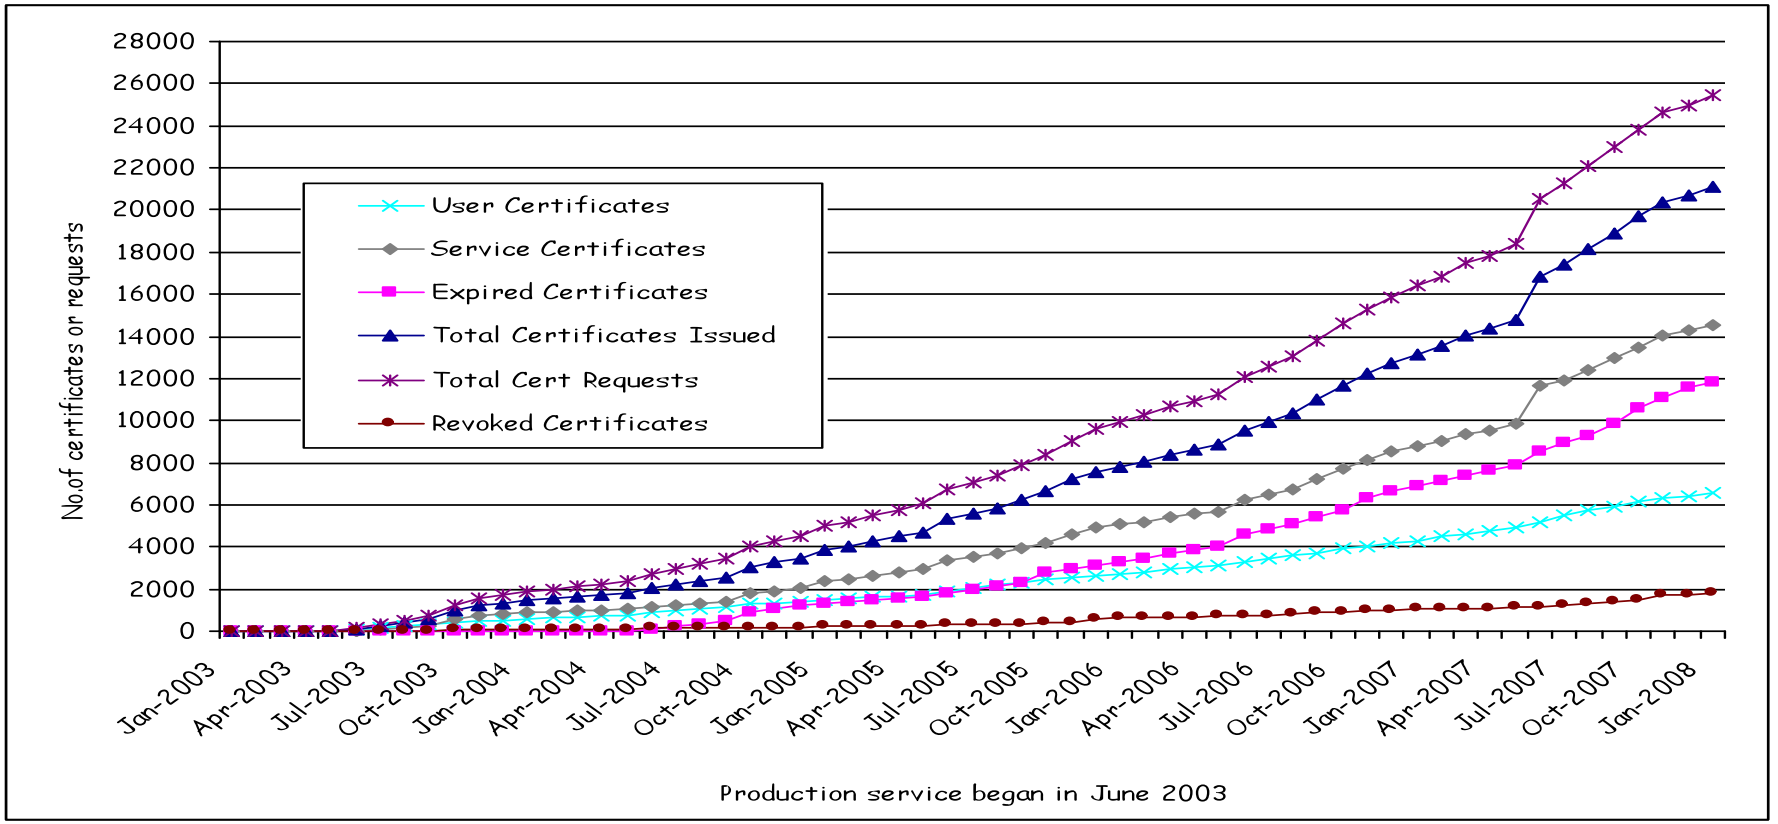

User Certificates	6549	Total No. of Revoked Certificates	1776
Host & Service Certificates	14545	Total No. of Expired Certificates	11797
Total No. of Requests	25470	Total No. of Certificates Issued	21095
		Total No. of Active Certificates	7547
ESnet SSL Server CA Certificates			49
FusionGRID CA certificates			113

* Report as of Jan 17, 2008

DOEGrids CA (Active Certificates) Usage Statistics

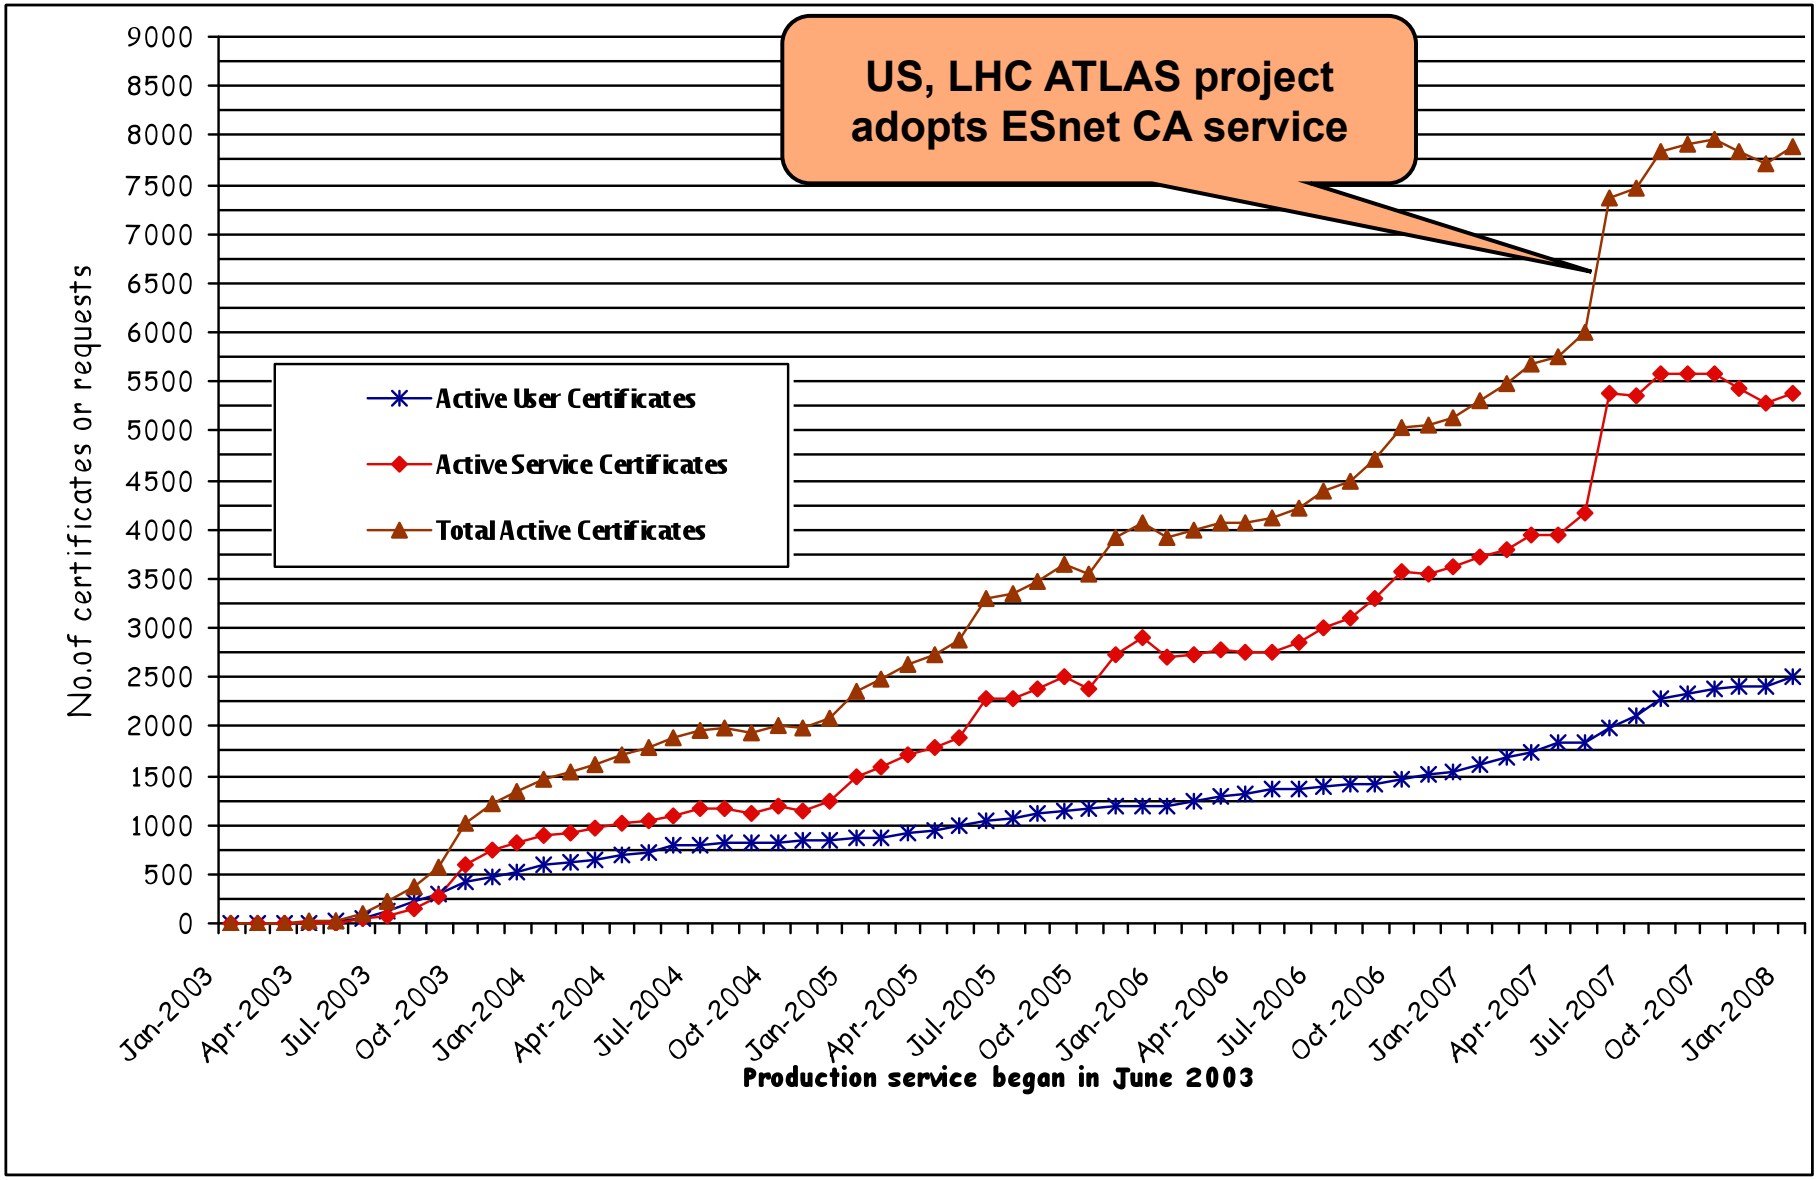

* Report as of Jan 17, 2008

DOEGrids CA Usage - Virtual Organization Breakdown

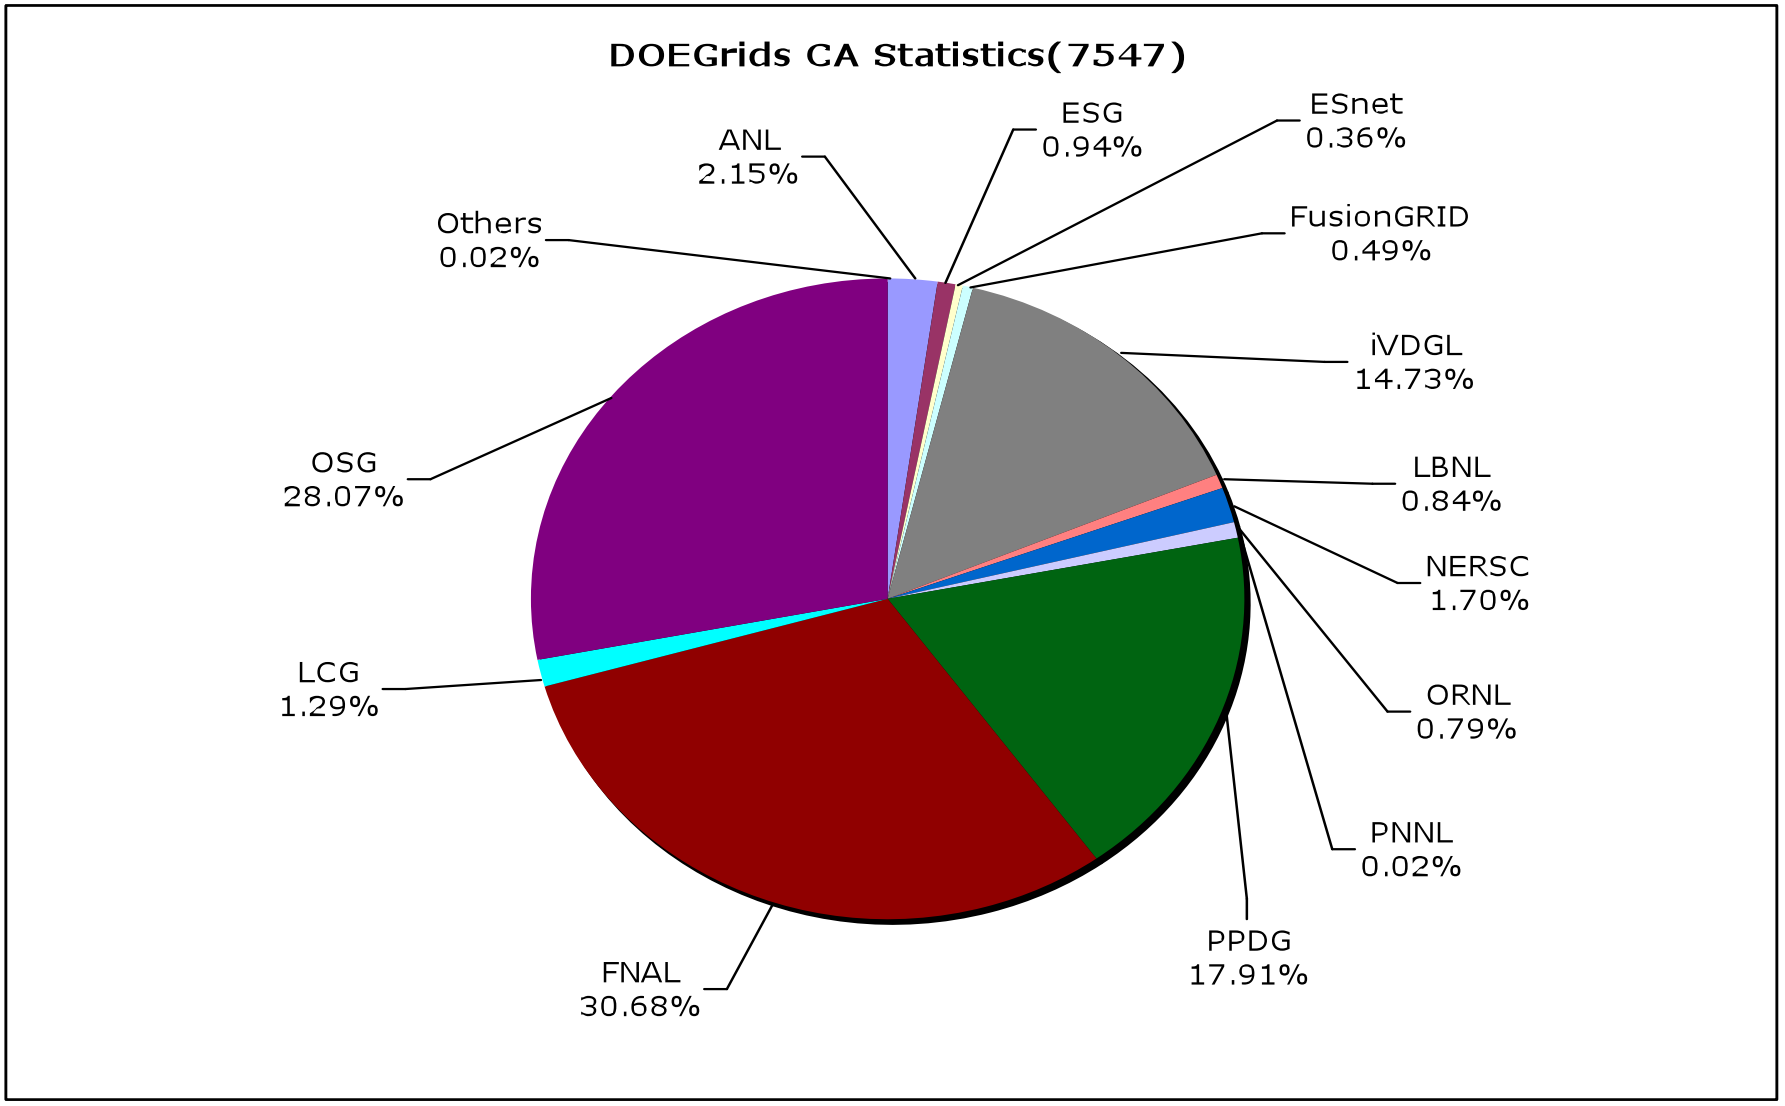

* DOE-NSF collab. & Auto renewals

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

DOEGrids CA Usage - Virtual Organization Breakdown

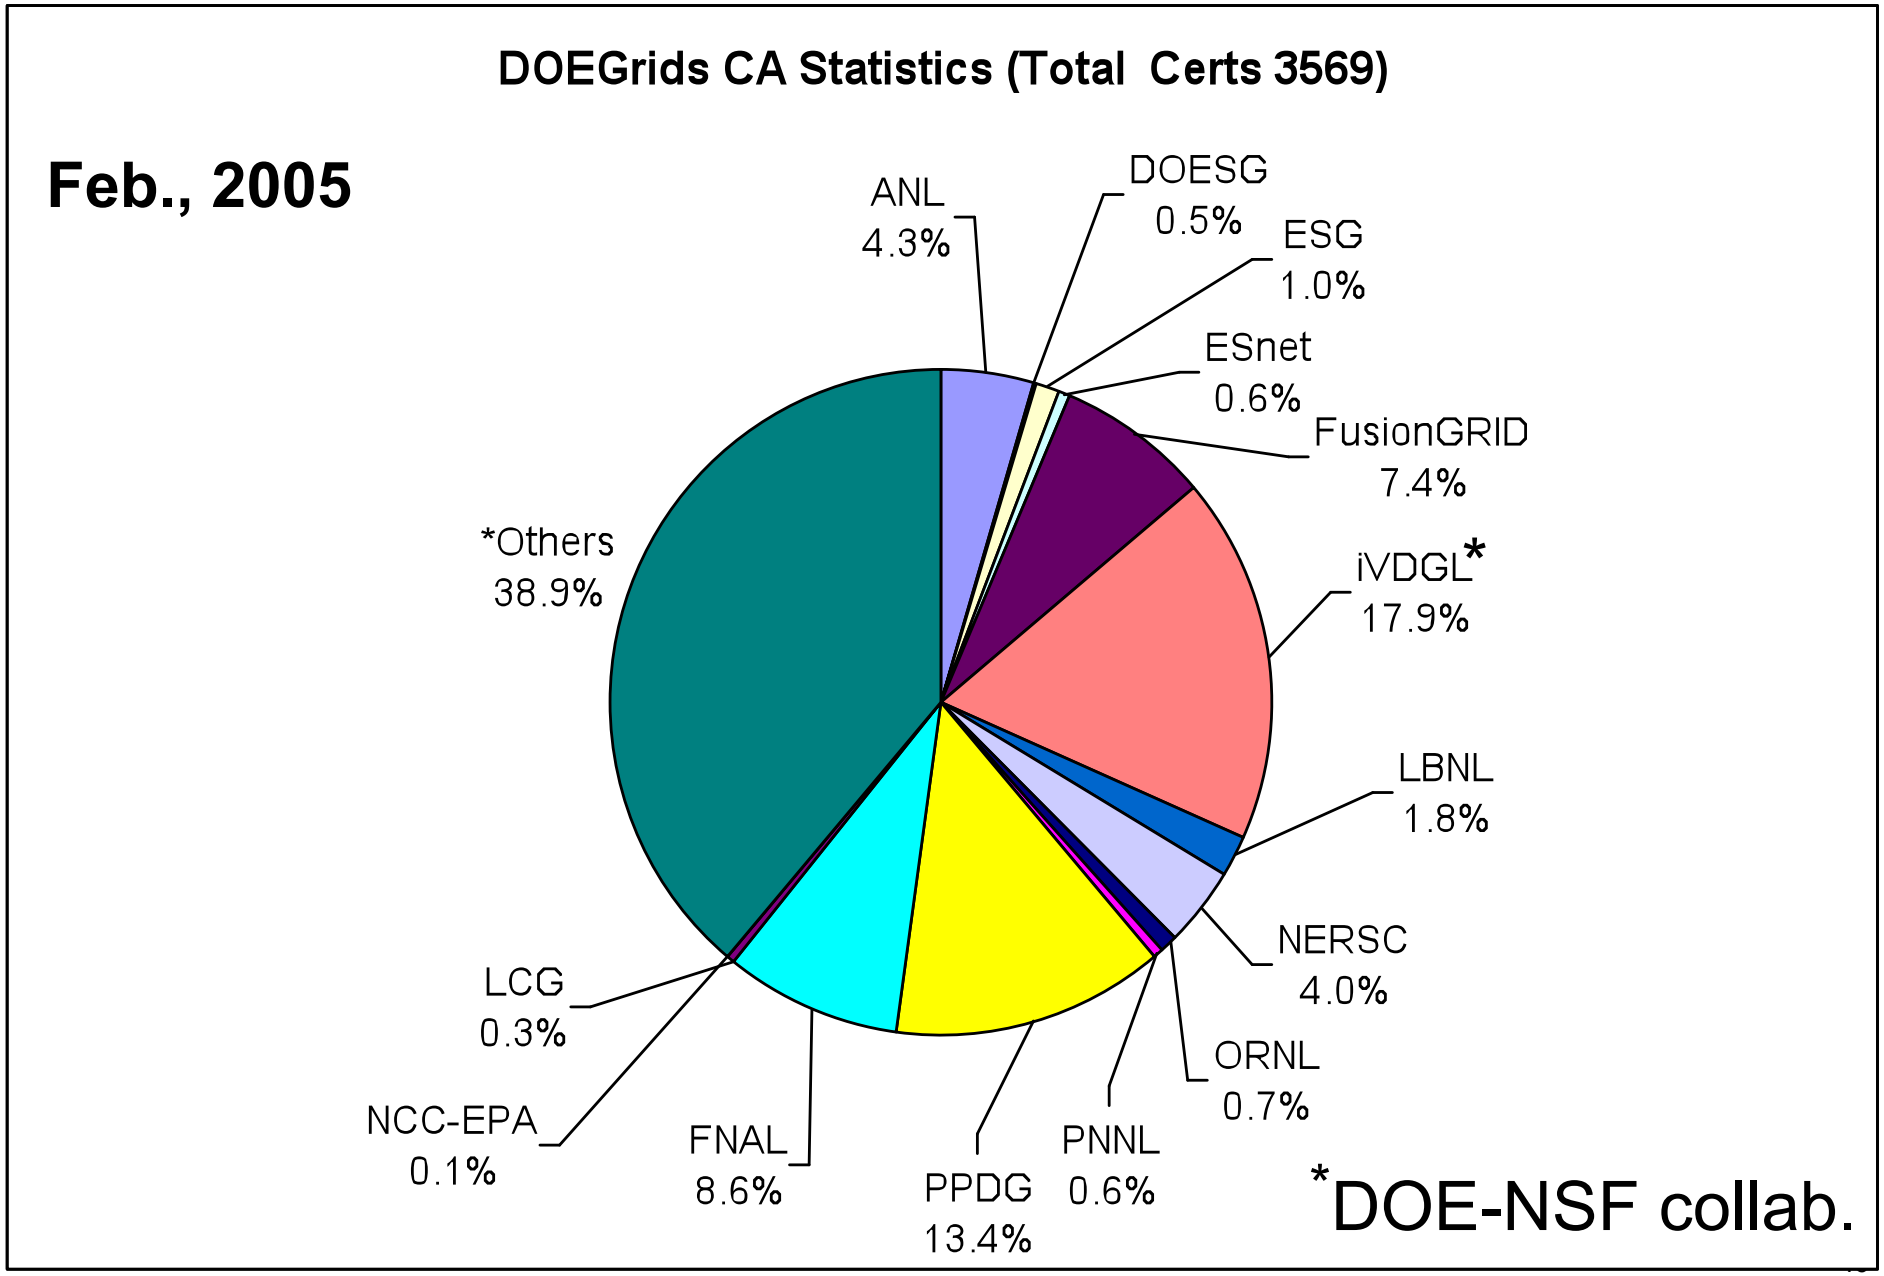

DOEGrids Disaster Recovery

- Recent upgrades and electrical incidents showed some unexpected vulnerabilities
- Remedies:
  - Update ESnet battery backup control system @LBL to protect ESnet PKI servers better
  - “Clone” CAs and distribute copies around the country
    - A lot of engineering
    - A lot of security work and risk assessment
    - A lot of politics
  - Clone and distribute CRL distribution machines

Policy Management Authority

- DOEGrids PMA needs re-vitalization
  - Audit finding
  - Will transition to (t)wiki format web site
  - Unclear how to re-energize
- ESnet owns the IGTF domains, and now the TAGPMA.org domain
  - 2 of the important domains in research science Grids
- TAGPMA.org
  - CANARIE needed to give up ownership
  - Currently finishing the transfer
  - Developing Twiki for PMA use
- IGTF.NET
  - Acquired in 2007
  - Will replace “gridpma.org” as the home domain for IGTF
  - Will focus on the wiki foundation used in TAGPMA, when it stabilizes

Possible Use of Grid Certs. For Wiki Access

- Experimenting with Wiki and client cert authentication
  - Motivation – no manual registration, large community, make PKI more useful
- Twiki – popular in science; upload of documents; many modules; some modest access control
  - Hasn't behaved well with client certs; the interaction of Apache <-> Twiki <-> TLS client is very difficult
- Some alternatives:
  - GridSite (but uses Media Wiki)
  - OpenID

Possible Use of Federation for ECS Authentication

- The Federated Trust / DOEGrids approach to managing authentication has successfully scaled to about 8000 users
  - This is possible because of the Registration Agent approach that puts initial authorization and certificate issuance in the hands of community representatives rahter than ESnet
  - Such an approach, in theory, could also work for ECS authentication and maybe first-level problems (e.g. “I have forgotten my password”)
- Upcoming ECS technology refresh includes authentication & authorization improvements.

Possible Use of Federation for ECS Authentication

- Exploring:
  - Full integration with DOEGrids – use its registration directly, and its credentials
  - Service Provider in federation architecture (Shibboleth, maybe openID)
  - Indico – this conference/room scheduler has become popular. Authentication/authorization services support needed
  - Some initial discussions with Tom Barton @ U Chicago (Internet2) on federation approaches took place in December, more to come soon
- Questions to Mike Helm and Stan Kluz

IIIb. ESnet Conferencing Service (ECS)

- An ESnet Science Service that provides audio, video, and data teleconferencing service to support human collaboration of DOE science
  - Seamless voice, video, and data teleconferencing is important for geographically dispersed scientific collaborators
  - Provides the central scheduling essential for global collaborations
  - ECS serves about 1600 DOE researchers and collaborators worldwide at 260 institutions
    - Videoconferences - about 3500 port hours per month
    - Audio conferencing - about 2300 port hours per month
    - Data conferencing - about 220 port hours per month Web-based, automated registration and scheduling for all of these services

ESnet Collaboration Services (ECS)

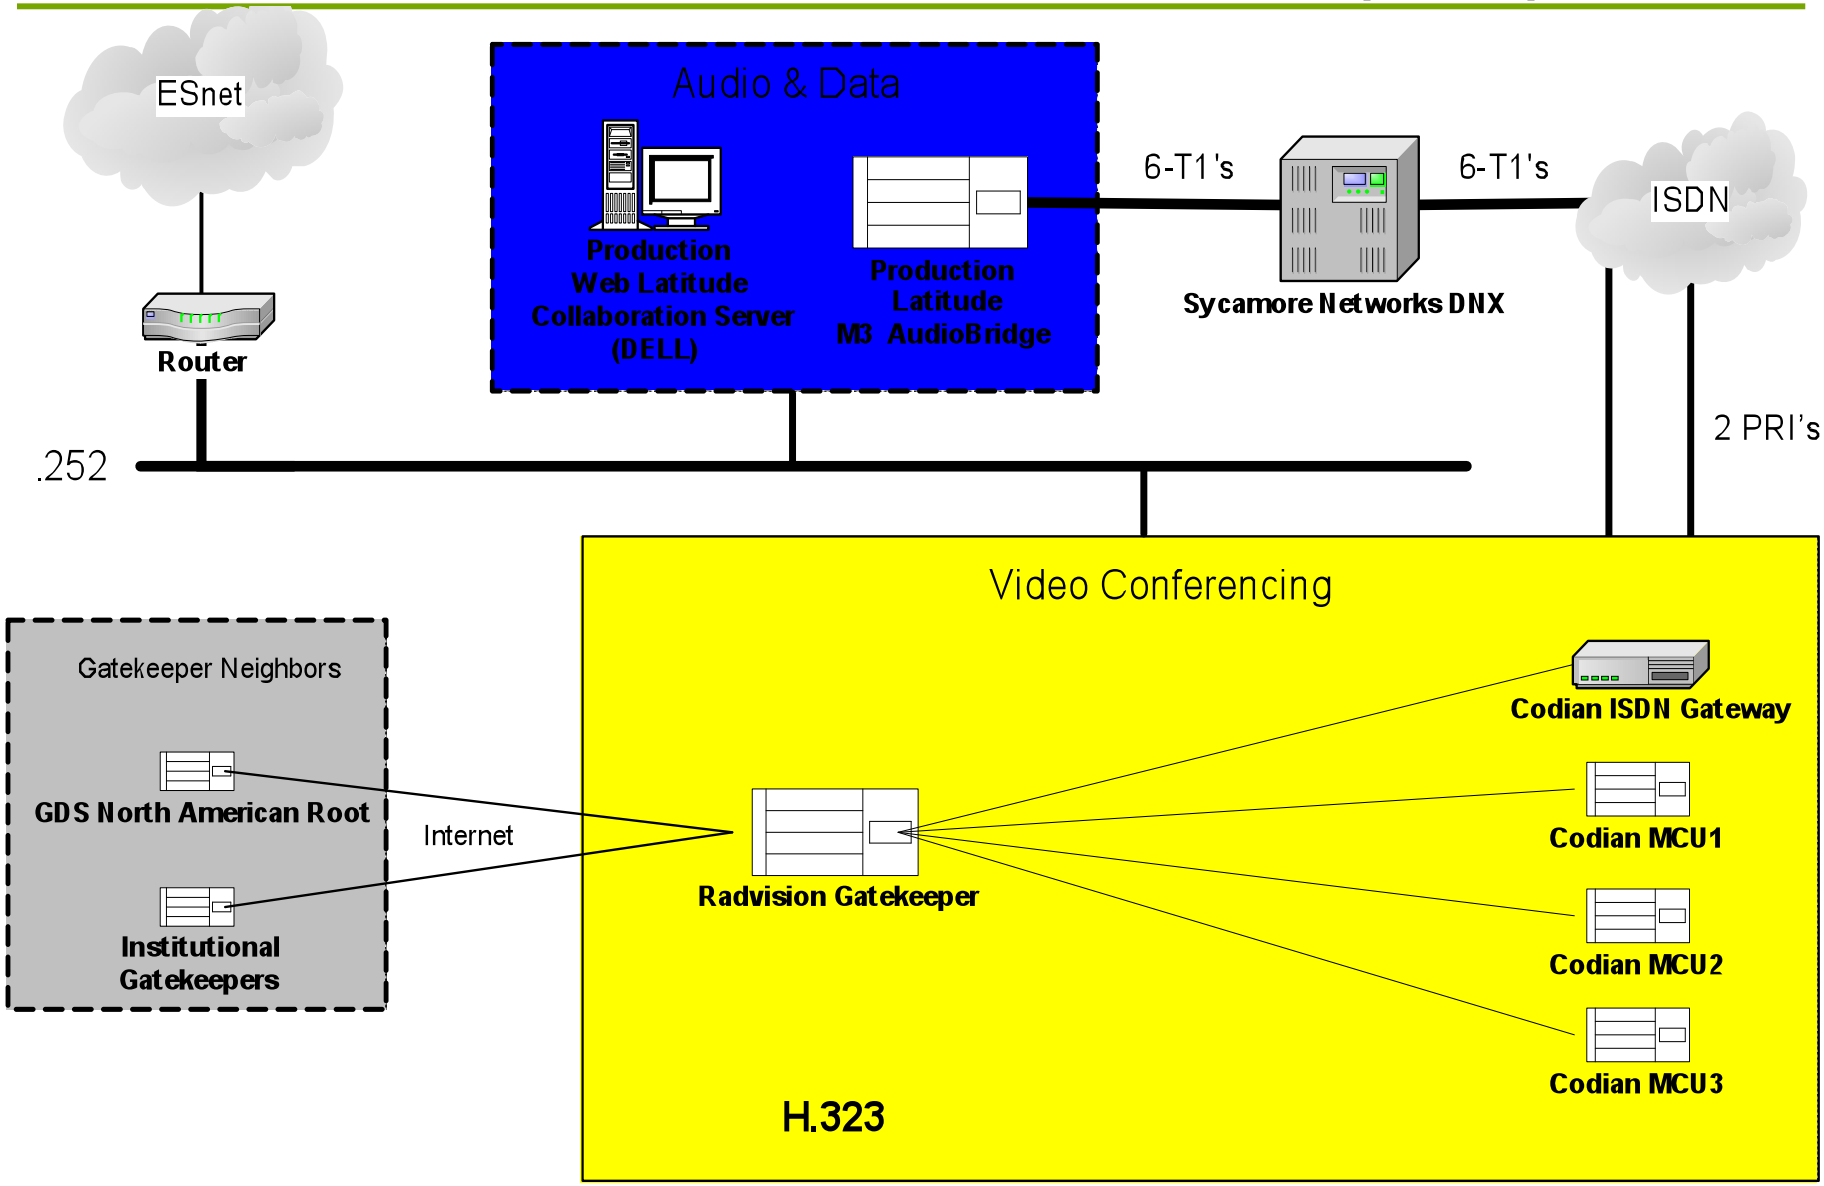

ECS Video Collaboration Service

- High Quality videoconferencing over IP and ISDN
- Reliable, appliance based architecture
- Ad-Hoc H.323 and H.320 multipoint meeting creation
- Web Streaming options on 3 Codian MCU's using Quicktime or Real
- 3 Codian MCUs with Web Conferencing Options
- 120 total ports of video conferencing on each MCU (40 ports per MCU)
- 384k access for video conferencing systems using ISDN protocol
- Access to audio portion of video conferences through the Codian ISDN Gateway

ECS Voice and Data Collaboration

- 144 usable ports
  - Actual conference ports readily available on the system.
- 144 overbook ports
  - Number of ports reserved to allow for scheduling beyond the number of conference ports readily available on the system.
- 108 Floater Ports
  - Designated for unexpected port needs.
  - Floater ports can float between meetings, taking up the slack when an extra person attends a meeting that is already full and when ports that can be scheduled in advance are not available.
- Audio Conferencing and Data Collaboration using Cisco MeetingPlace
- Data Collaboration = WebEx style desktop sharing and remote viewing of content
- Web-based user registration
- Web-based scheduling of audio / data conferences
- Email notifications of conferences and conference changes
- 650+ users registered to schedule meetings (not including guests)

ECS Futures

- ESnet is still on-track to replicate the teleconferencing hardware currently located at LBNL in a Central US or Eastern US location
  - have about come to the conclusion that the ESnet hub in NYC is not the right place to site the new equipment
- The new equipment is intended to provide at least comparable service to the current (upgraded) ECS system
  - Also intended to provide some level of backup to the current system
  - A new Web based registration and scheduling portal may also come out of this

ECS Service Level

- ESnet Operations Center is open for service 24x7x365.
- A trouble ticket is opened within15 to 30 minutes and assigned to the appropriate group for investigation.
- Trouble ticket is closed when the problem is resolved.
- ***ECS support*** is provided Monday to Friday, 8AM to 5 PM Pacific Time excluding LBNL holidays
  - Reported problems are addressed within 1 hour from receiving a trouble ticket during ECS support period
  - ESnet does NOT provide a real time (during-conference) support service

Real Time ECS Support

- A number of user groups have requested “real-time” conference support (monitoring of conferences while in-session)
- Limited Human and Financial resources currently prohibit ESnet from:

  A) Making real time information available to the public on the systems status (network, ECS, etc) This information is available only on some systems to our support personnel

  B) 24x7x365 real-time support

  C) Addressing simultaneous trouble calls as in a real time support environment.

  - This would require several people addressing multiple problems simultaneously

Real Time ECS Support

- Solution
  - A fee-for-service arrangement for real-time conference support
  - Available from TKO Video Communications, ESnet's ECS service contractor
  - Service offering could provide:
    - Testing and configuration assistance prior to your conference
    - Creation and scheduling of your conferences on ECS Hardware
    - Preferred port reservations on ECS video and voice systems
    - Connection assistance and coordination with participants
    - Endpoint troubleshooting
    - Live phone support during conferences
    - Seasoned staff and years of experience in the video conferencing industry
    - ESnet community pricing

ECS Impact from LBNL Power Outage, January 9th 2008

- Heavy rains caused LBNL sub-station one of two 12Kv buss to fail
  - 50% of LBNL lost power
  - LBNL estimates 48 hr before power restored
  - ESnet lost power to data center
  - Backup generator for ESnet data center failed to start due to a failed starter battery
  - ESnet staff kept MAN Router functioning by swapping batteries in UPS.
  - ESnet services, ECS, PKI, etc.. were shut down to protect systems and reduce heat load in room
  - Internal ESnet router lost UPS power and shut down
- After ~25 min generator was started by “jump” starting.
  - ESnet site router returned to service
  - No A/C in data center when running on generator
  - Mission critical services brought back on line
- After ~ 2 hours house power was restored
  - Power reliability still questionable
  - LBNL strapped buss one to feed buss two
- After 24 hrs remaining services restored to normal operation
- Customer Impact
  - ~ 2 Hrs instability of ESnet services to customers

Power Outage Lessons Learned

- As of Jan 22, 2008
  - Normal building power feed has still not been restored
- EPA rules restrict operation of generator in non-emergency mode.
  - However, monthly running of generator will resume
- Current critical systems list to be evaluated and priorities adjusted.
- Internal ESnet router relocated to bigger UPS or removed from the ESnet services critical path.
- ESnet staff need more flashlights!

➢Summary

- **Transition to ESnet4 is going smoothly**
  - **New network services to support large-scale science are progressing**
  - **Measurement infrastructure is rapidly becoming widely enough deployed to be very useful**
- **New ESC hardware and service contract are working well**
  - **Plans to deploy replicate service are on-track**
- **Federated trust - PKI policy and Certification Authorities**
  - **Service continues to pick up users at a pretty steady rate**
  - **Maturing of service - and PKI use in the science community generally - is maturing**

References

[OSCARS]

For more information contact Chin Guok (chin@es.net). Also see

- http://www.es.net/oscars

[LHC/CMS]

http://cmsdoc.cern.ch/cms/aprom/phedex/prod/Activity::RatePlots?view=global

[ICFA SCIC] "Networking for High Energy Physics." International Committee for Future Accelerators (ICFA), Standing Committee on Inter-Regional Connectivity (SCIC), Professor Harvey Newman, Caltech, Chairperson.

- http://monalisa.caltech.edu:8080/Slides/ICFASCIC2007/

[E2EMON] Geant2 E2E Monitoring System –developed and operated by JRA4/WI3, with implementation done at DFN
http://cnmdev.lrz-muenchen.de/e2e/html/G2_E2E_index.html
http://cnmdev.lrz-muenchen.de/e2e/lhc/G2_E2E_index.html

[TrViz] ESnet PerfSONAR Traceroute Visualizer
https://performance.es.net/cgi-bin/level0/perfsonar-trace.cgi