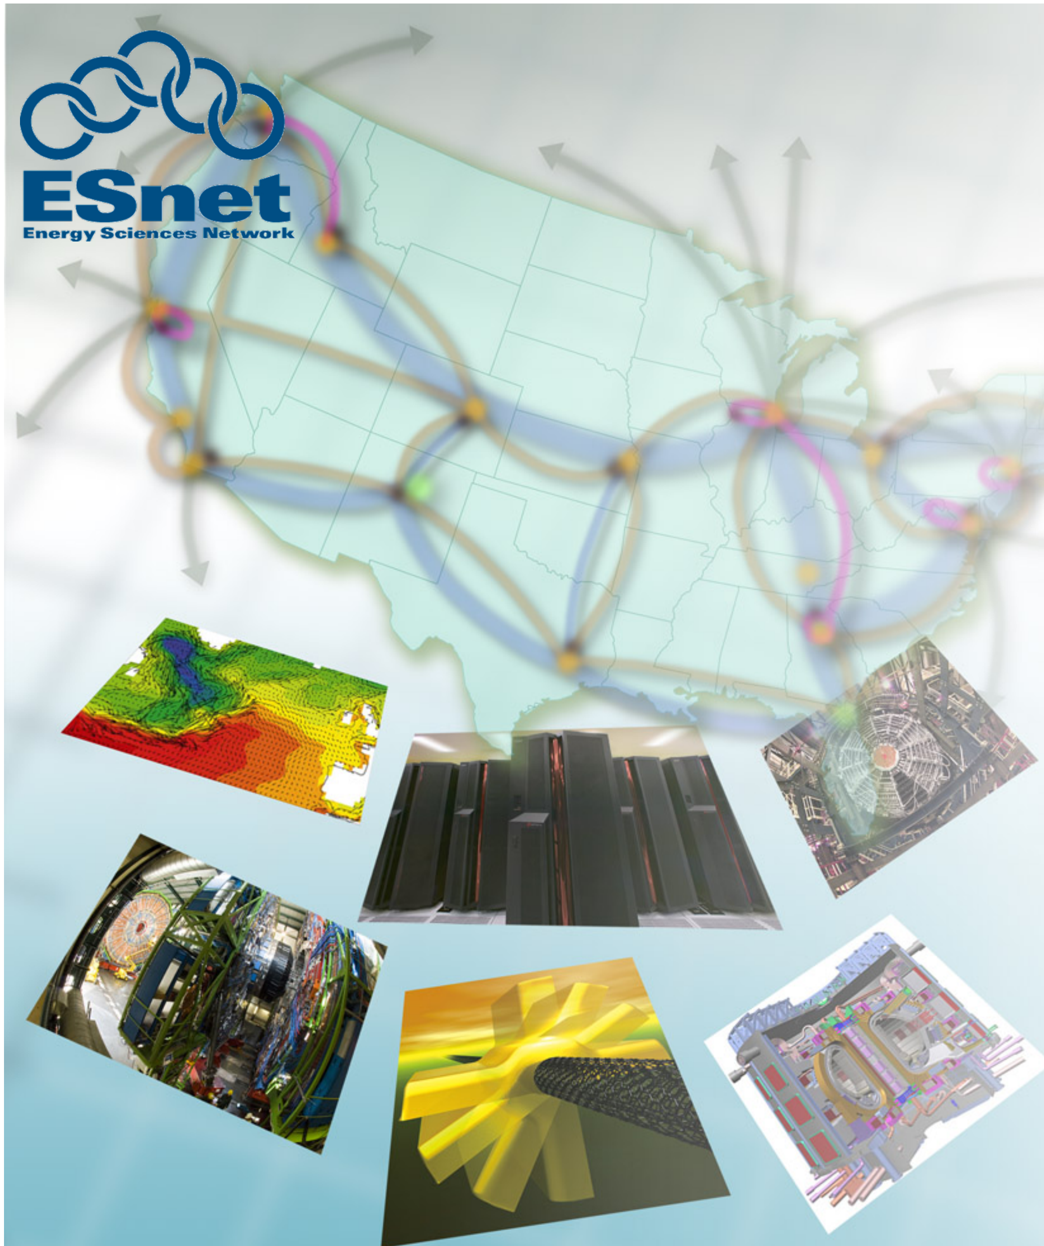

ESnet4: Networking for the Future of DOE Science

Office of Science, Science Programs Requirements Workshops: Nuclear Physics
May, 2008

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's Office of Science: Enabling Large-Scale Science

- The Office of Science (SC) is the single ***largest supporter of basic research in the physical sciences in the United States***, … providing more than 40 percent of total funding … for the Nation's research programs in high-energy physics, nuclear physics, and fusion energy sciences. (http://www.science.doe.gov) – SC funds 25,000 PhDs and PostDocs
- A primary mission of SC's National Labs is to ***build and operate very large scientific instruments*** - particle accelerators, synchrotron light sources, very large supercomputers - that generate massive amounts of data and involve very large, distributed collaborations
- **ESnet - the Energy Sciences Network - is an SC program whose *primary mission is to enable the large-scale science* of the Office of Science 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
- In addition to the National Labs, ESnet serves much of the rest of DOE, including NNSA – about 75,000-100,000 users in total

Office of Science, US Community

Drives ESnet Design for Domestic Connectivity

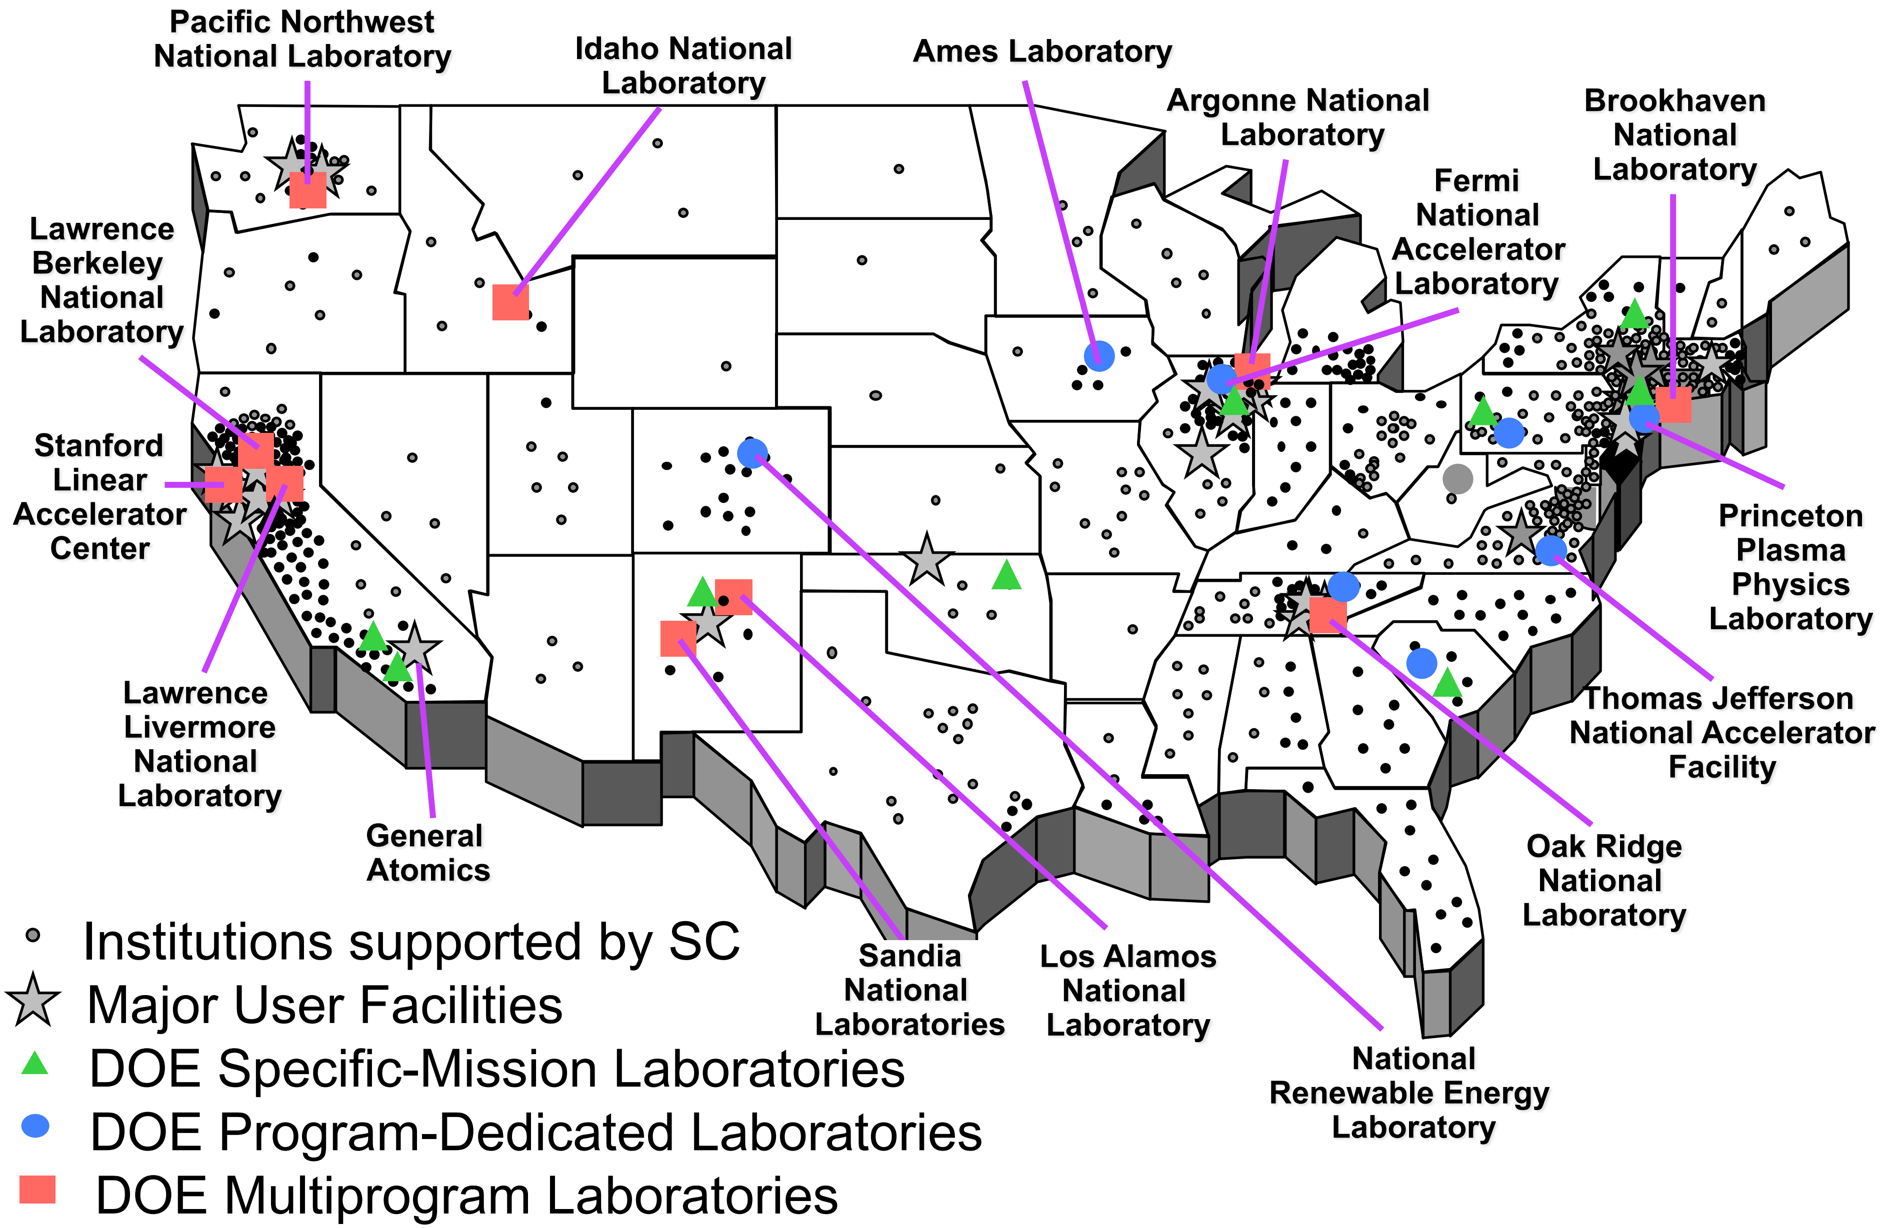

ESnet as a Large-Scale Science Data Transport Network

A few hundred of the 25,000,000 hosts that the Labs connect to every month account for 50% of all ESnet traffic (which are primarily high energy physics, nuclear physics, and climate data at this point)

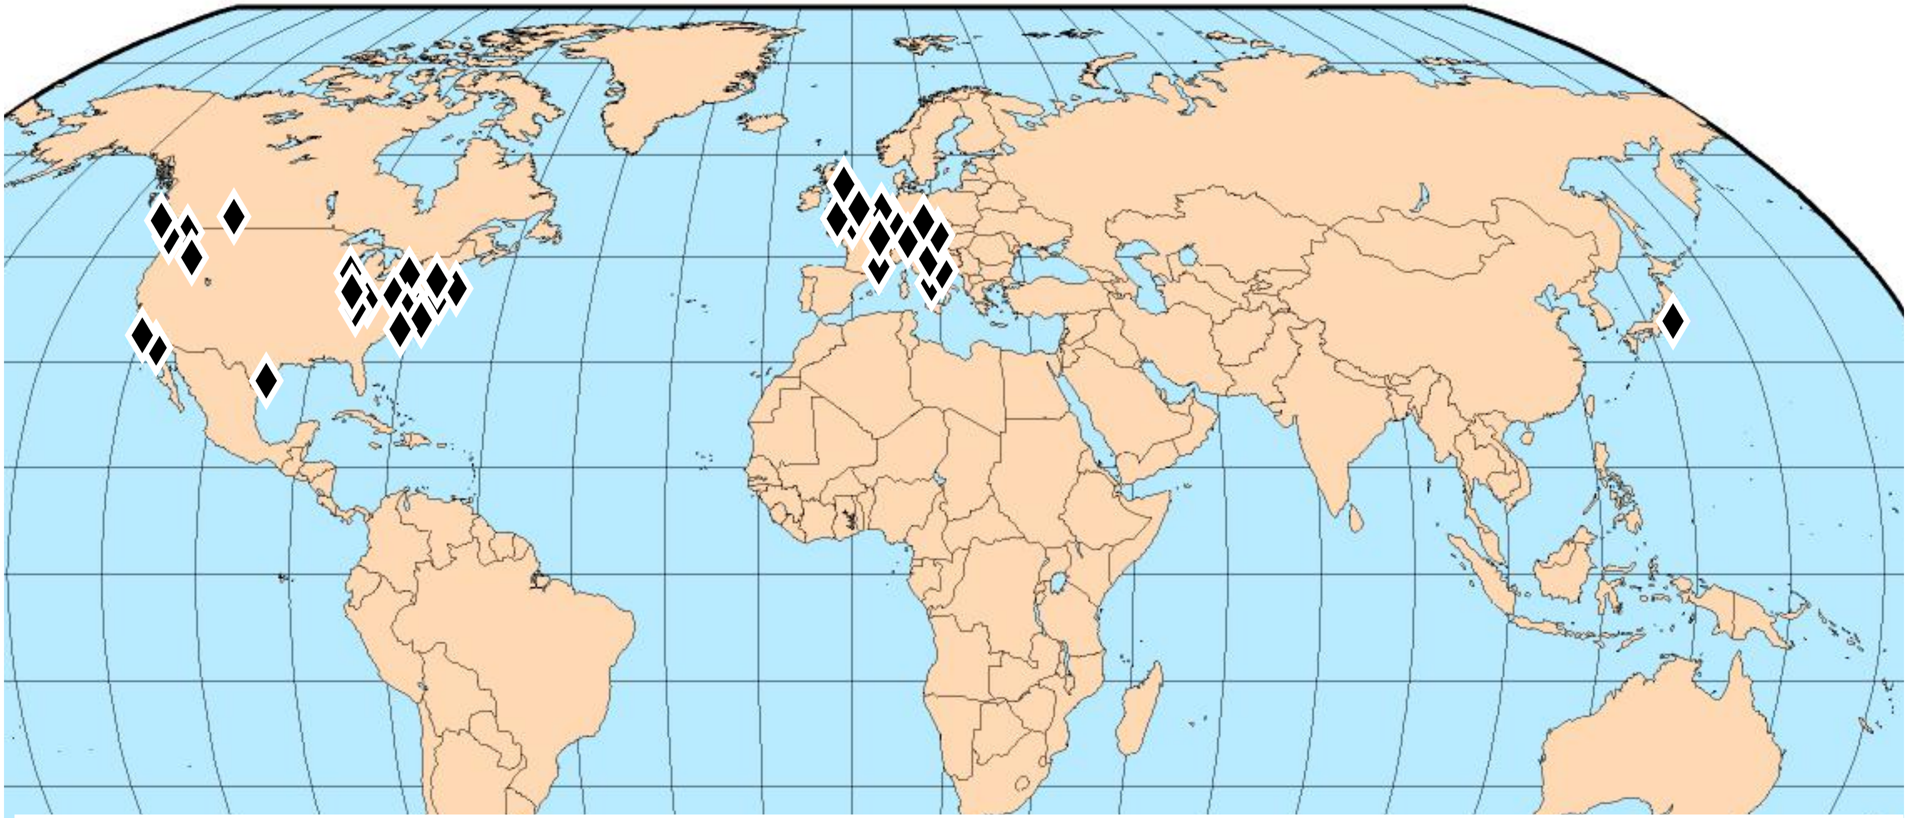

Most of ESnet's traffic (>85%) goes to and comes from outside of ESnet. This reflects the highly collaborative nature of large-scale science (which is one of the main focuses of DOE's Office of Science).

♦ = the R&E source or destination of ESnet's top 100 sites (all R&E) (the DOE Lab destination or source of each flow is not shown)

High-Level View of all ESnet Traffic

ESnet Inter-Sector Traffic Summary, Jan. 2007

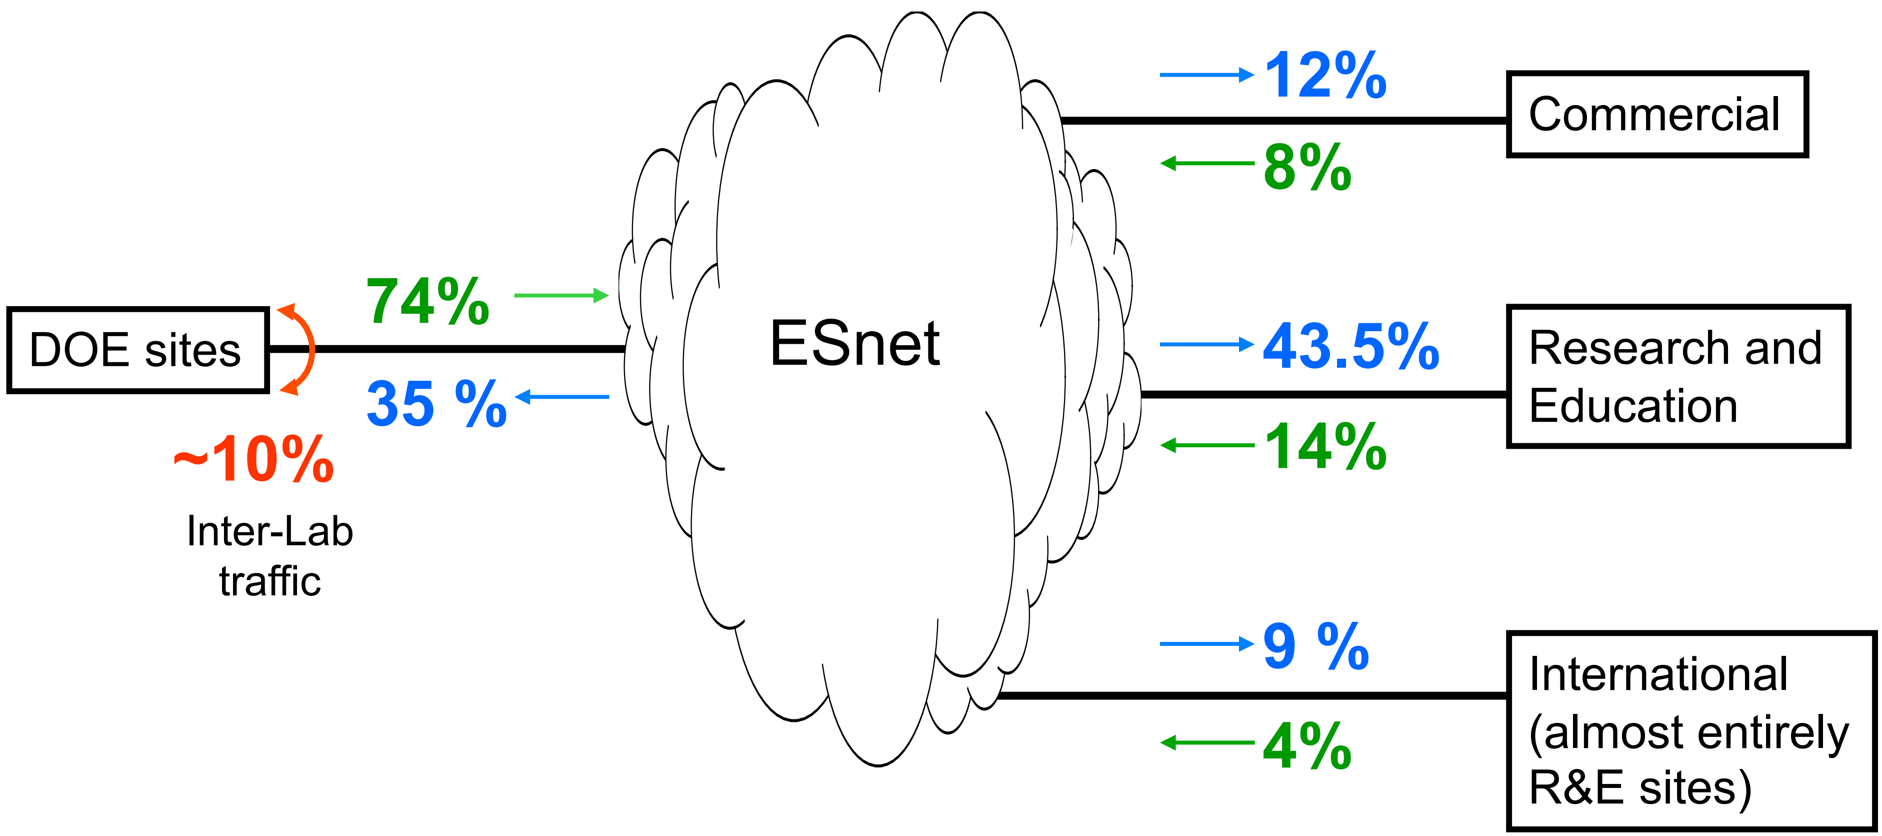

Traffic notes

• ***more than 90% of all traffic is Office of Science***

Traffic coming into ESnet = Green
Traffic leaving ESnet = Blue
Traffic between ESnet sites

ESnet Defined

- A national optical circuit infrastructure
  - ESnet shares an optical network on a dedicated national fiber infrastructure with Internet2 (US national research and education (R&E) network)
    - ESnet has exclusive use of a group of 10Gb/s optical channels on this infrastructure
- A large-scale IP network
  - A tier 1 Internet Service Provider (ISP) (direct connections with all major commercial networks providers)
- A large-scale science data transport network
  - With multiple high-speed connections to all major US and international research and education (R&E) networks in order to enable large-scale science
  - Providing circuit services specialized to carry the massive science data flows of the National Labs

ESnet Defined

- The primary DOE network provider
  - Provides production Internet service to all of the major DOE Labs* and most other DOE sites
  - Based on DOE Lab populations, it is estimated that between 50,000 -100,000 users depend on ESnet for global Internet access
    - additionally, each year more than 18,000 non-DOE researchers from universities, other government agencies, and private industry use Office of Science facilities
- An organization of 30 professionals structured for the service
  - The ESnet organization designs, builds, and operates the network
- An operating entity with an FY06 budget of about $26.5M

* PNNL supplements its ESnet service with commercial service

ESnet as Tier 1 ISP – Providing DOE with Access to the Global Internet

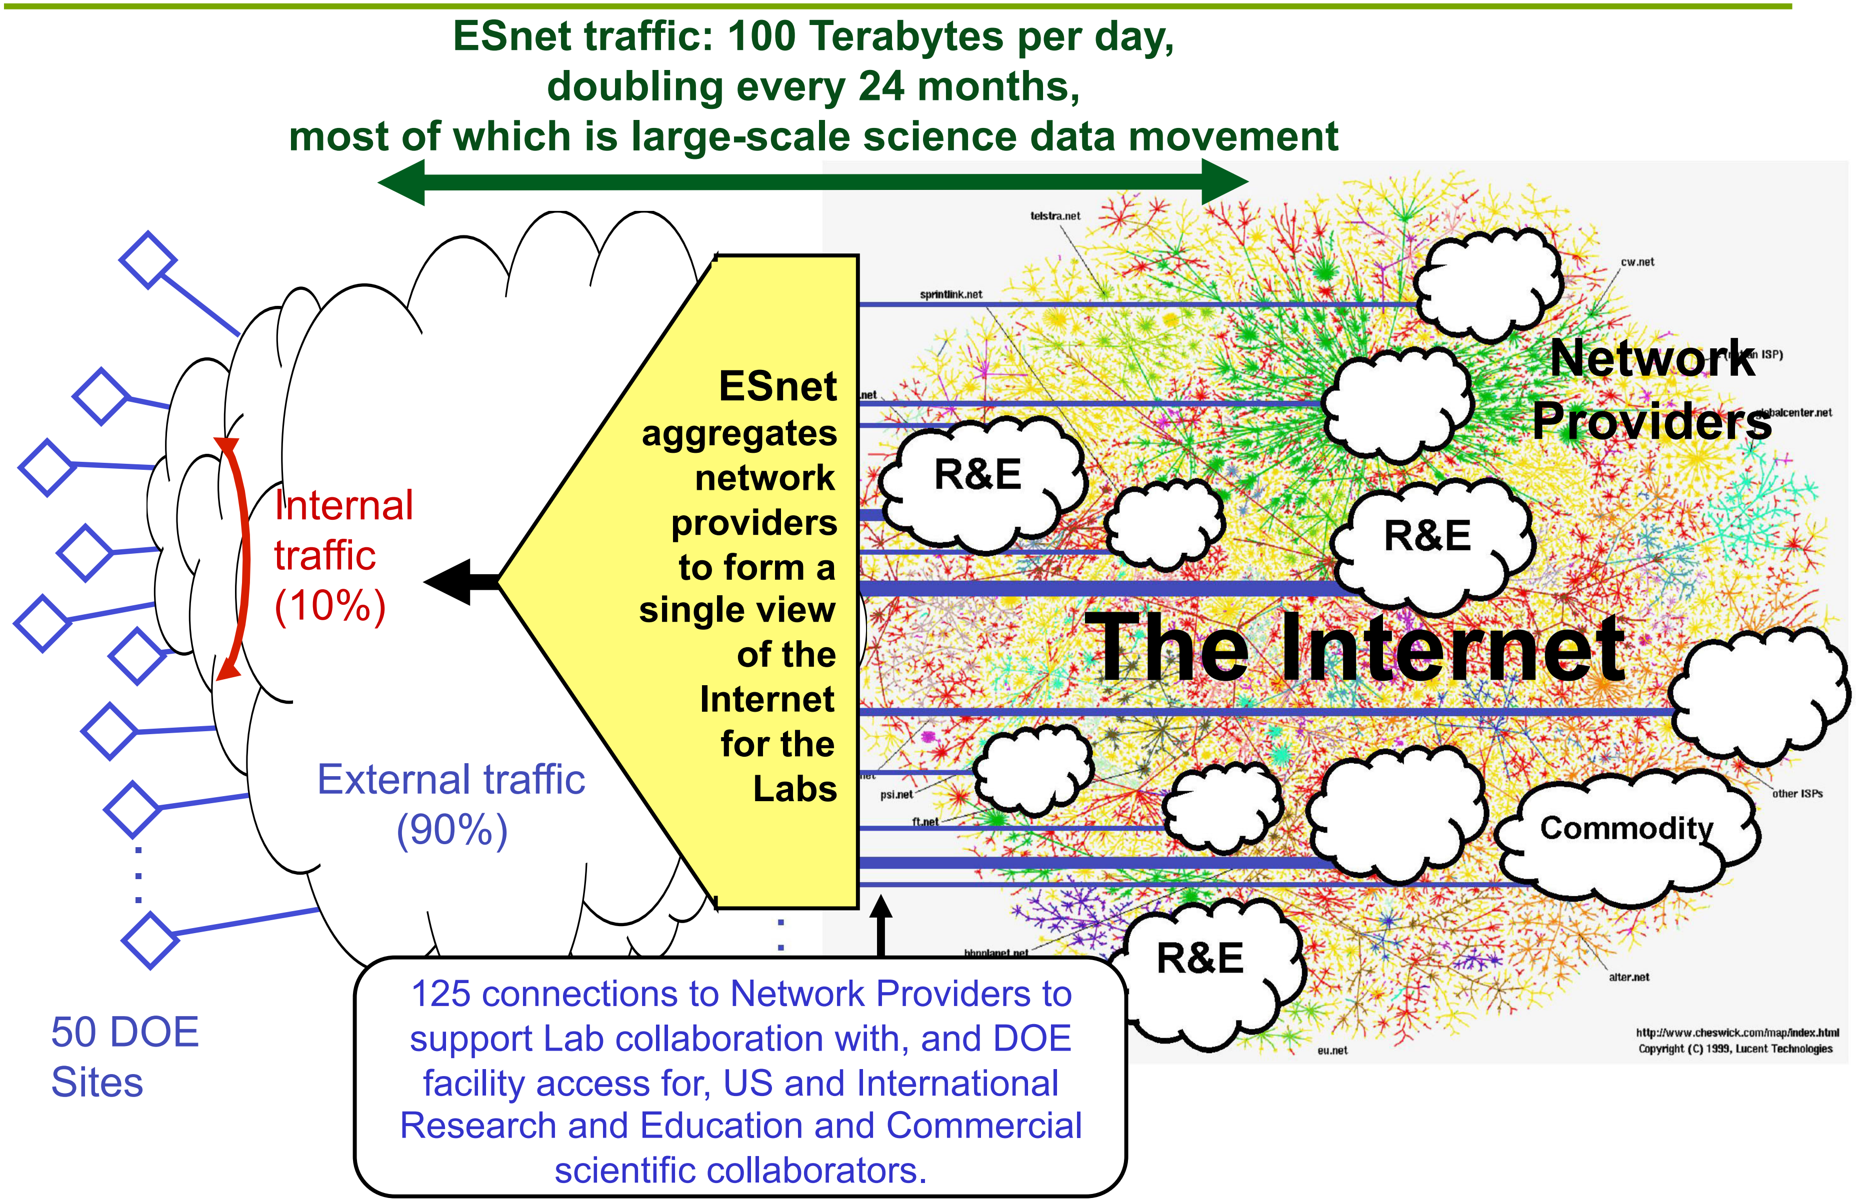

No single external network provider offers an aggregation service with connectivity or capacity that would meet DOE Science requirements.

What Does ESnet Provide? - 1

- An architecture tailored to accommodate DOE's large-scale science
  - Move huge amounts of data between a small number of sites that are scattered all over the world
- Comprehensive connectivity
  - High bandwidth access to DOE sites and DOE's primary science collaborators: Research and Education institutions in the US, Europe, Asia Pacific, and elsewhere
  - Full access to the global commercial Internet for DOE Labs
- Highly reliable transit networking
  - Fundamental goal is to deliver every packet that is received to the destination location

What Does ESnet Provide? - 2

- A full suite of network services
  - Routing and address space management for IPv4 and IPv6 unicast and multicast
  - Primary DNS services
  - Circuit services (layer 2 e.g. Ethernet VLANs), MPLS overlay networks (to isolate traffic from the general Internet)
  - Scavenger service so that certain types of bulk traffic can use all available bandwidth, but will give priority to any other traffic when it shows up

What Does ESnet Provide? - 3

- New network services
  - Guaranteed bandwidth “virtual circuits”
- Services supporting collaborative science
  - Federated trust services / PKI Certification Authorities with science oriented policy
  - Audio-video-data teleconferencing
- Highly reliable and secure operation of the network
  - Extensive disaster recovery infrastructure
  - Comprehensive internal security
  - Cyberdefense for the WAN

What Does ESnet Provide? - 4

- Comprehensive user support, including “owning” all trouble tickets involving ESnet users (including problems at the far end of an ESnet connection) until they are resolved
  - 24 hr x 365 d/yr coverage
  - ESnet’s mission to enable the network based aspects of SC science involves troubleshooting network problems wherever they occur
- A highly collaborative and interactive relationship with the DOE Labs and scientists for planning, configuration, and operation of the network
  - ESnet and its services evolve continuously in direct response to SC science needs
  - Engineering services for special requirements

ESnet is An Organization Structured for the Service

Network operations and user support (24x7x365, end-to-end problem resolution	5.5
Network engineering, routing and network services, WAN security	8.3
Deployment and WAN maintenance	
Internal infrastructure, disaster recovery, security	7.4
Applied R&D for new network services (Circuit services and end-to-end monitoring)	3.0
Science collaboration services (Public Key Infrastructure certification authorities, AV conferencing, email lists, Web services)	5.5
Management, accounting, compliance	3.5

ESnet FY06 Budget is Approximately $26.6M

Approximate Budget Categories

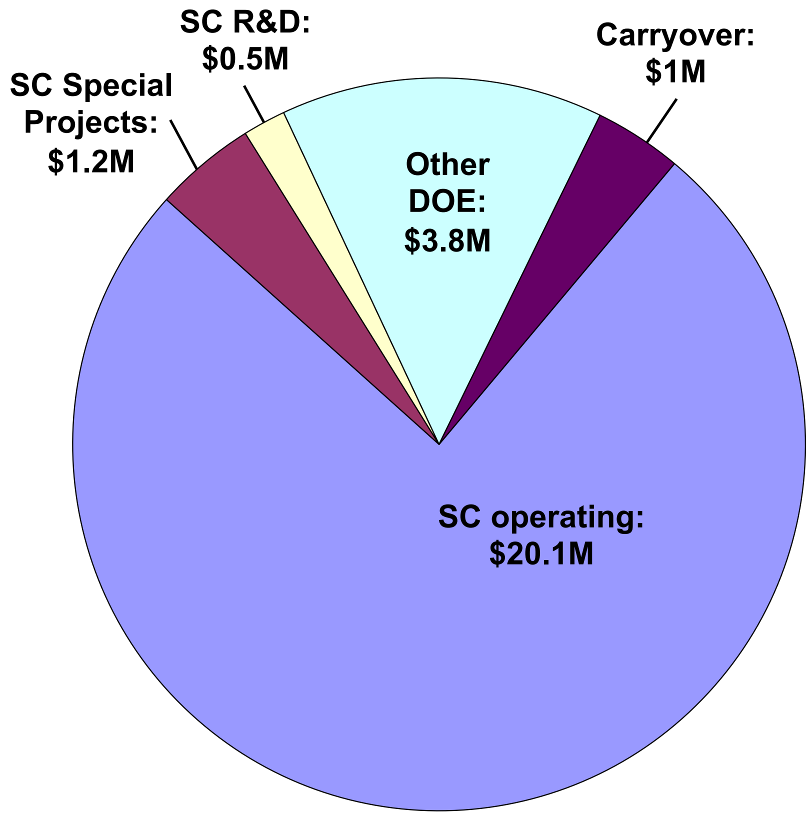

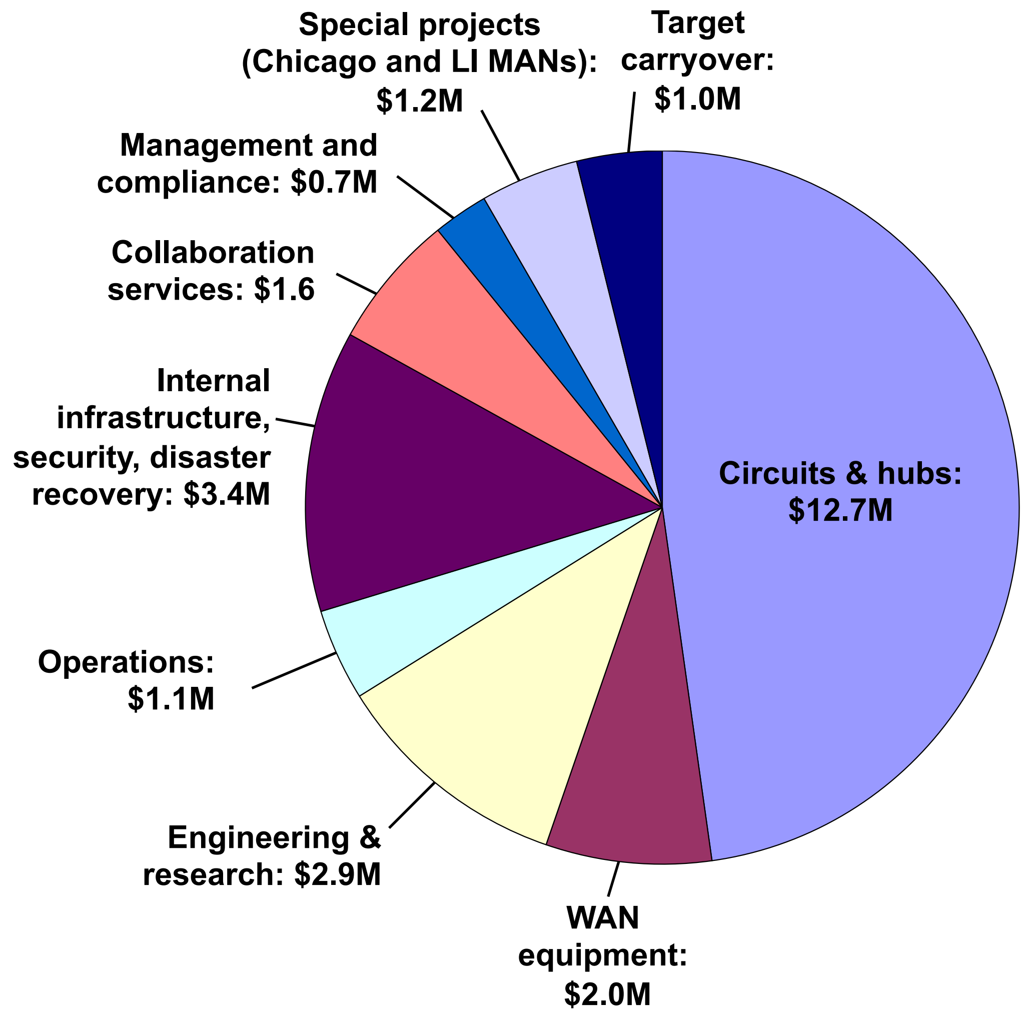

Total funds: $26.6M

Total expenses: $26.6M

ESnet History

ESnet0/MFENet mid-1970s-1986	ESnet0/MFENet	56 Kbps microwave and satellite links
ESnet1 1986-1995	ESnet formed to serve the Office of Science	56 Kbps to 45 Mbps X.25
ESnet2 1995-2000	Partnership with Sprint to build the first national footprint ATM network	IP over 155 Mbps ATM net
ESnet3 2000-2007	Partnership with Qwest to build a national Packet over SONET network and optical channel Metropolitan Area fiber	IP over 2.5 and 10Gbps SONET
ESnet4 2007-2012	Partnership with Internet2 and US Research & Education community to build a dedicated national optical network	IP and virtual circuits on a configurable optical infrastructure with at least 5-6 optical channels of 10-100 Gbps each

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

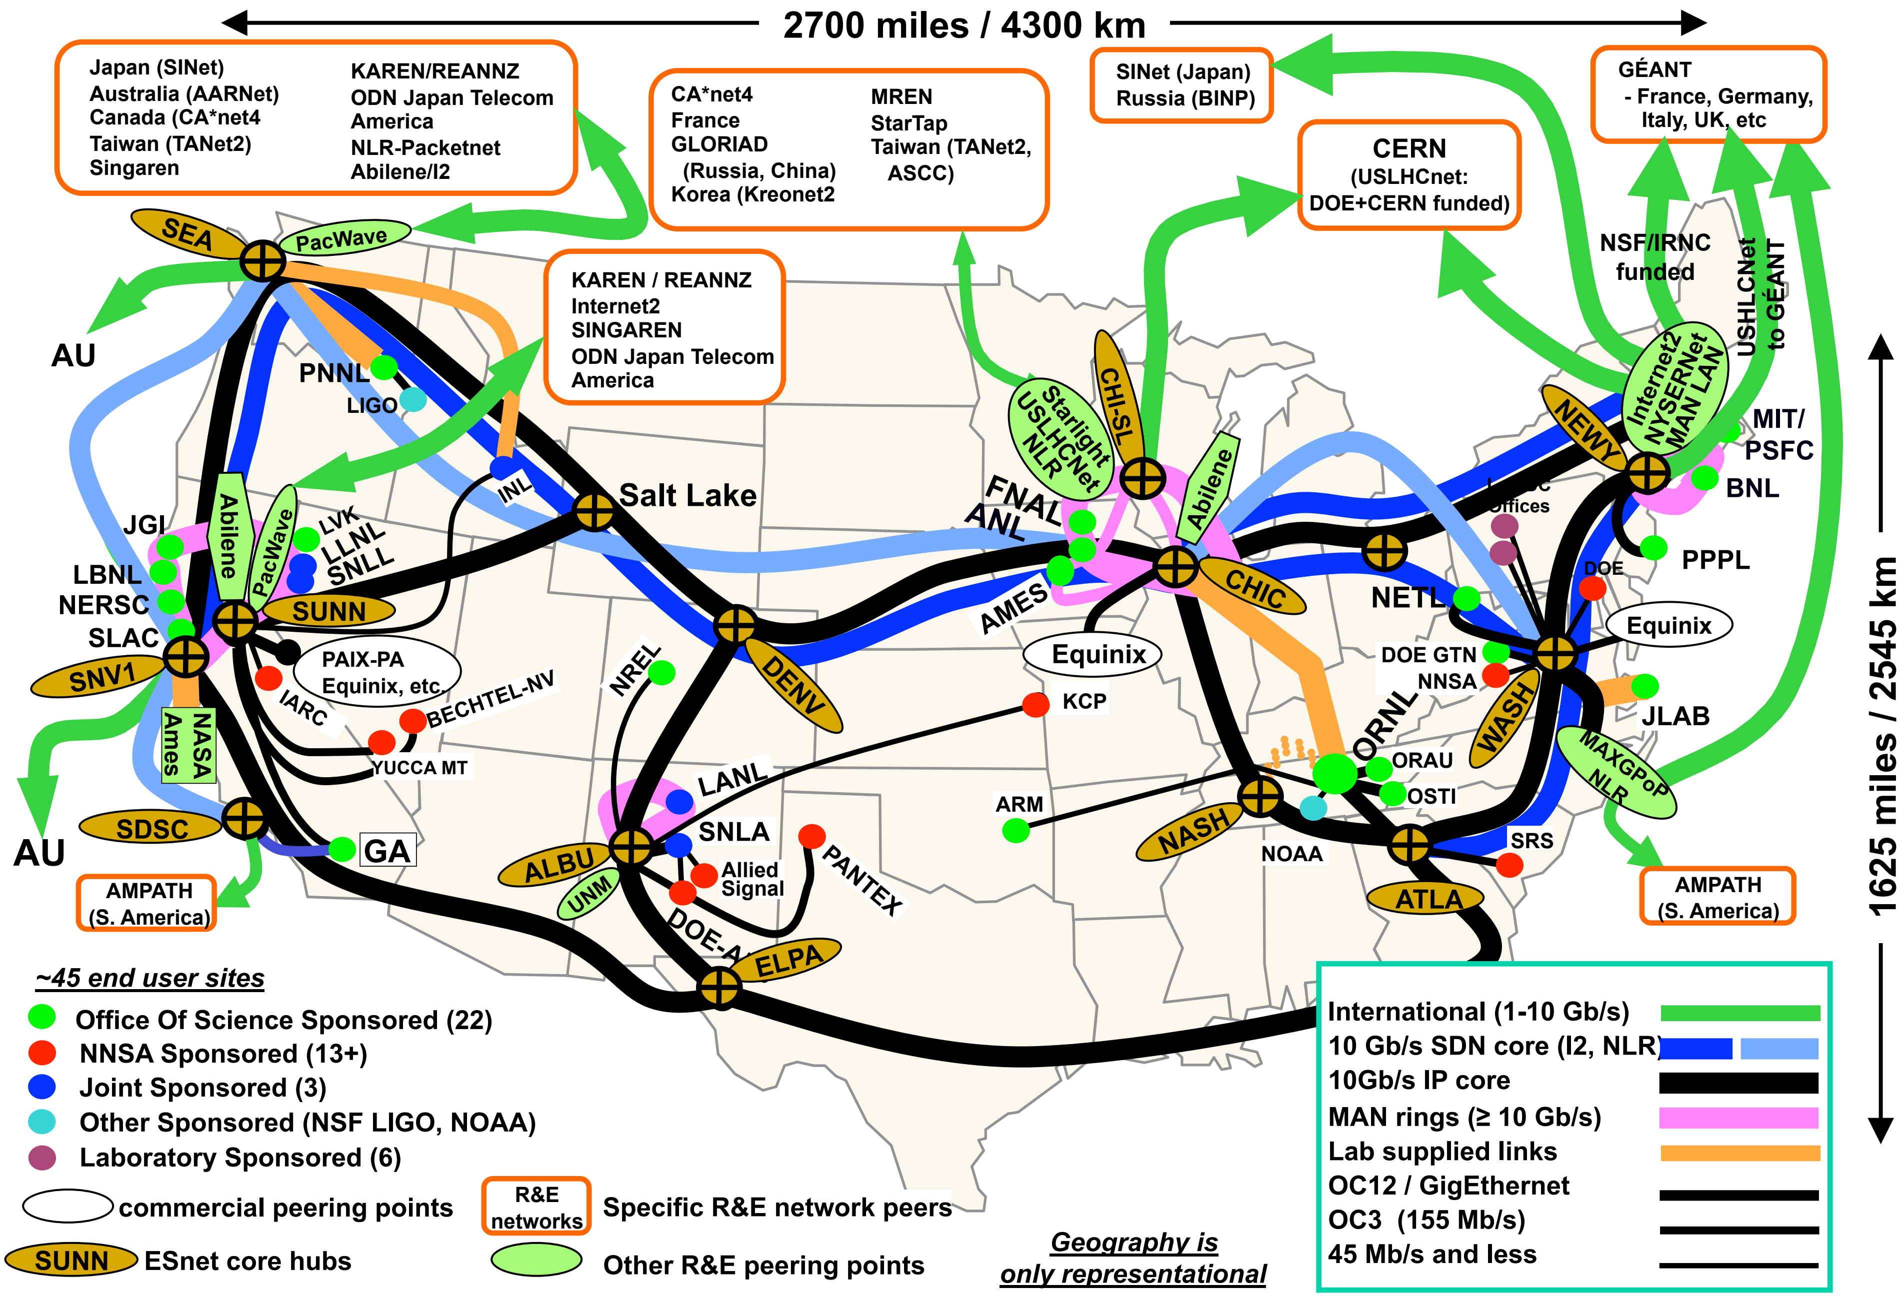

ESnet's Place in U. S. and International Science

- ESnet, Internet2, and National Lambda Rail (NLR) provide most of the nation's transit networking for basic science
  - Internet2 provides national transit networking for most of the US universities by interconnecting the regional networks (mostly via the GigaPoPs)
  - ESnet provides national transit networking and ISP service for the DOE Labs
  - NLR provides various 10Gb/s circuits
- DANTE/GÉANT plays a role in Europe similar to Internet2 and ESnet in the US – it interconnects the European National Research and Education Networks (NRENs), to which the European R&E sites connect
  - GÉANT and NSF each provide a 10G transatlantic circuit for the R&E community and these currently carry all non-LHC ESnet traffic to Europe - this DOE to Europe traffic is a significant fraction of all ESnet traffic

How the Network is Operated

- Research and education networks like ESnet typically operate with a small, dedicated staff
- The key to everything that ESnet provides is scalability
  - How do you manage a huge infrastructure with a small number of people?
  - This issue dominates all others when looking at whether to support new services (e.g. Sciences Services / Grid middleware)
    - Can the service be structured so that its operational aspects do not increase as a function of increasing user population?
      - if not, then ESnet cannot offer it as a service

Scalable Operation is Essential

The entire ESnet network is operated by about 16 people

Management, resource management, circuit accounting, group leads (4 FTE)

Infrastructure (7 FTE)

Core Engineering Group (5 FTE)

7X24 On-Call Engineers (7 FTE)

7X24 Operations Desk (4 FTE)

Automated Alarm Analysis System

Science Services (middleware and collaboration tools) (5 FTE)

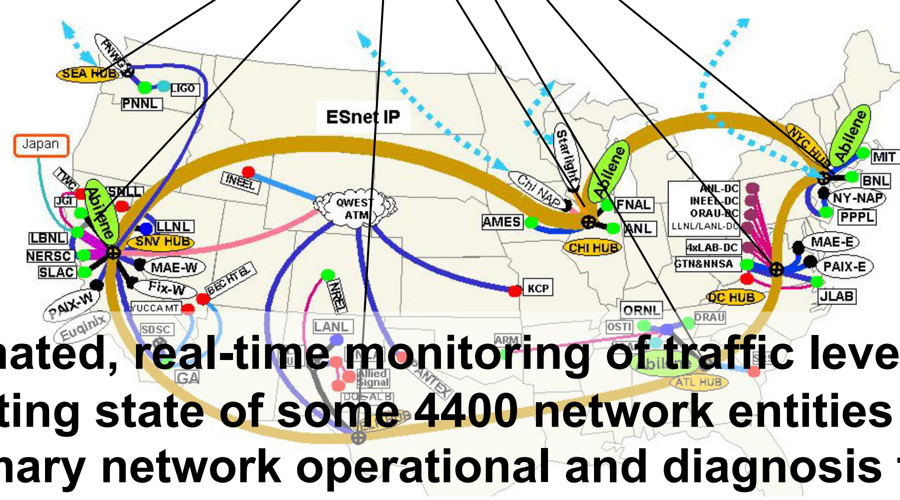

Automated, real-time monitoring of traffic levels and operating state of some 4400 network entities is the primary network operational and diagnosis tool

Operating *Science Mission* Critical Infrastructure

- ESnet is a visible and critical piece of DOE science infrastructure
  - if ESnet fails,10s of thousands of DOE and University users know it within minutes if not seconds
- Requires high reliability and high operational security in the systems that are integral to the operation and management of the network
  - Secure and redundant mail and Web systems are central to the operation and security of ESnet
    - trouble tickets are by email
    - engineering communication by email
    - engineering database interfaces are via Web
  - Secure network access to Hub routers
  - Backup secure telephone modem access to Hub equipment
  - 24x7 help desk and 24x7 on-call network engineer

    trouble@es.net (end-to-end problem resolution)

ESnet is a Highly Reliable Infrastructure

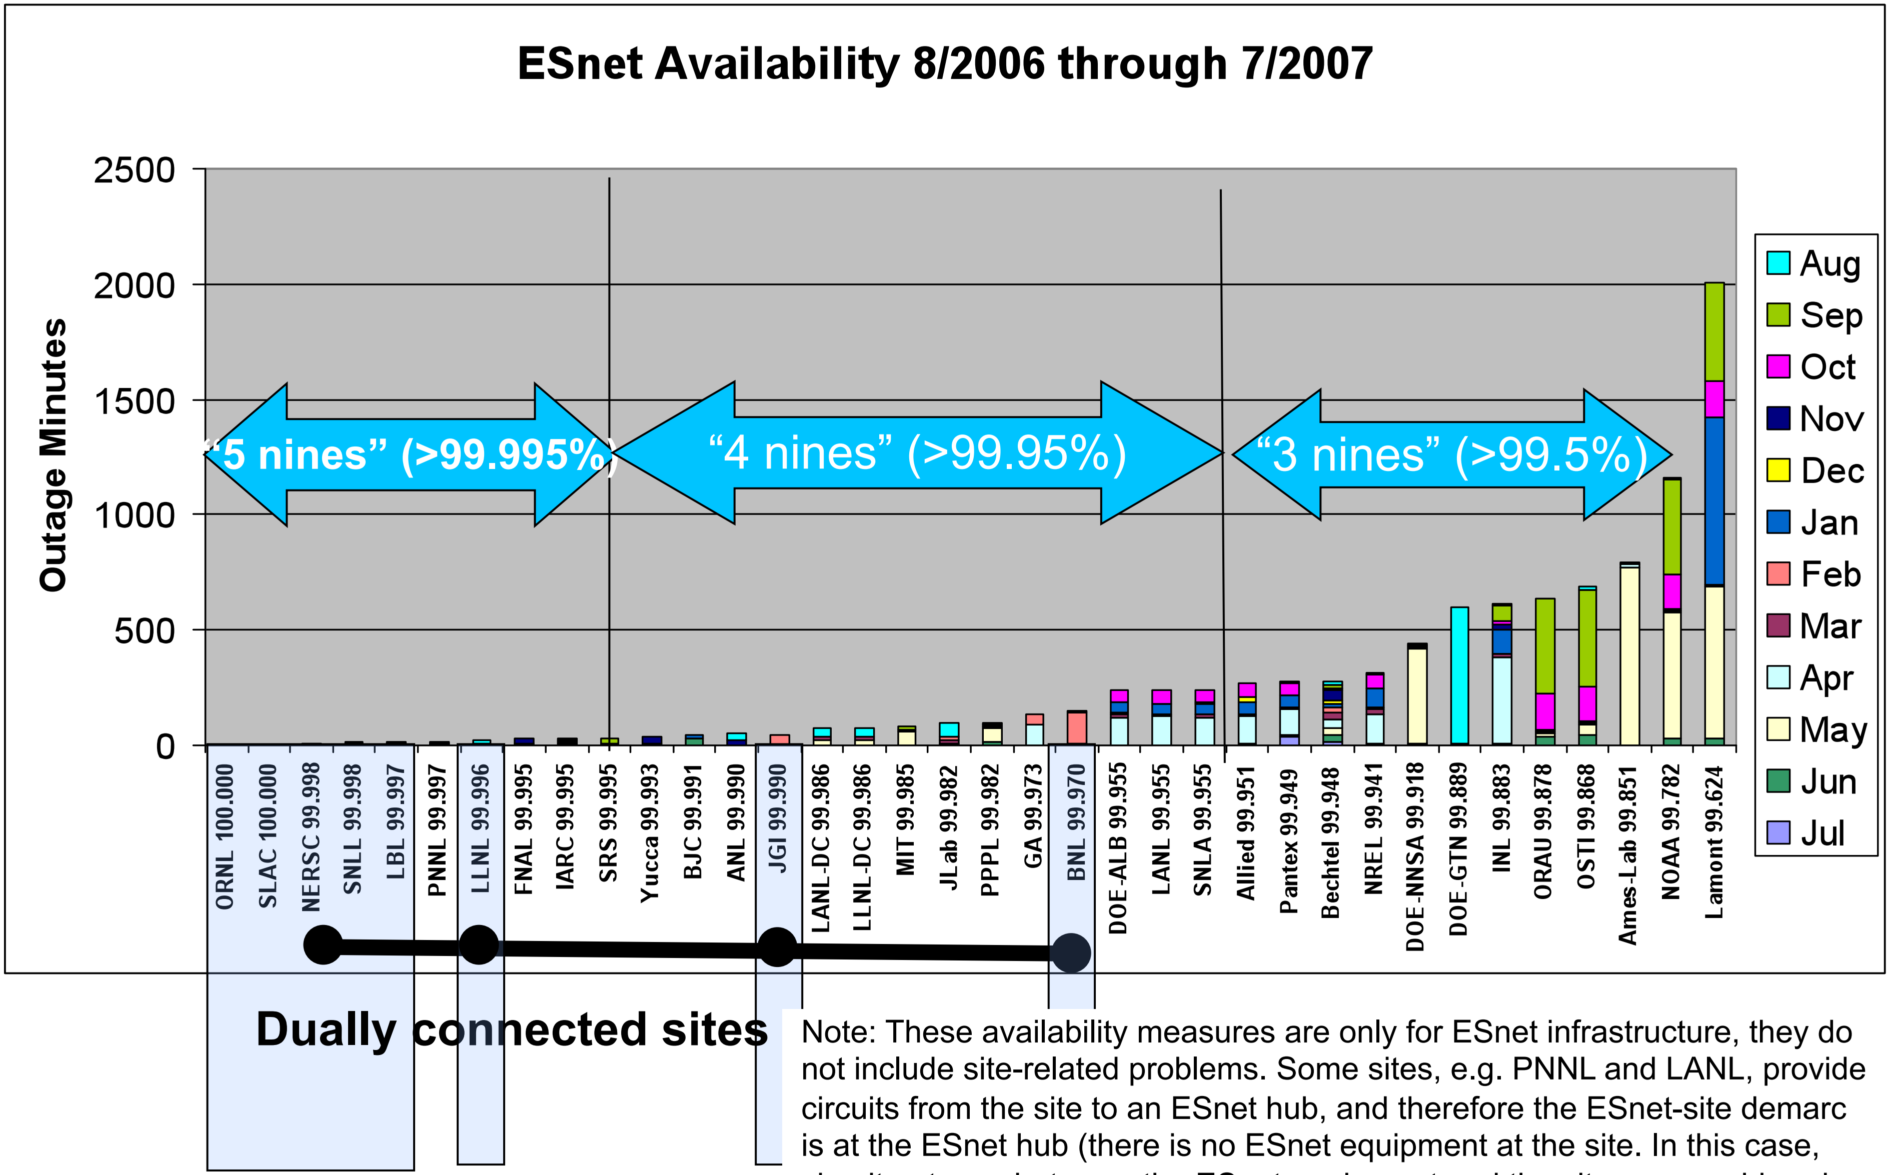

Note: These availability measures are only for ESnet infrastructure, they do not include site-related problems. Some sites, e.g. PNNL and LANL, provide circuits from the site to an ESnet hub, and therefore the ESnet-site demarc is at the ESnet hub (there is no ESnet equipment at the site. In this case,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

Disaster Recovery and Stability

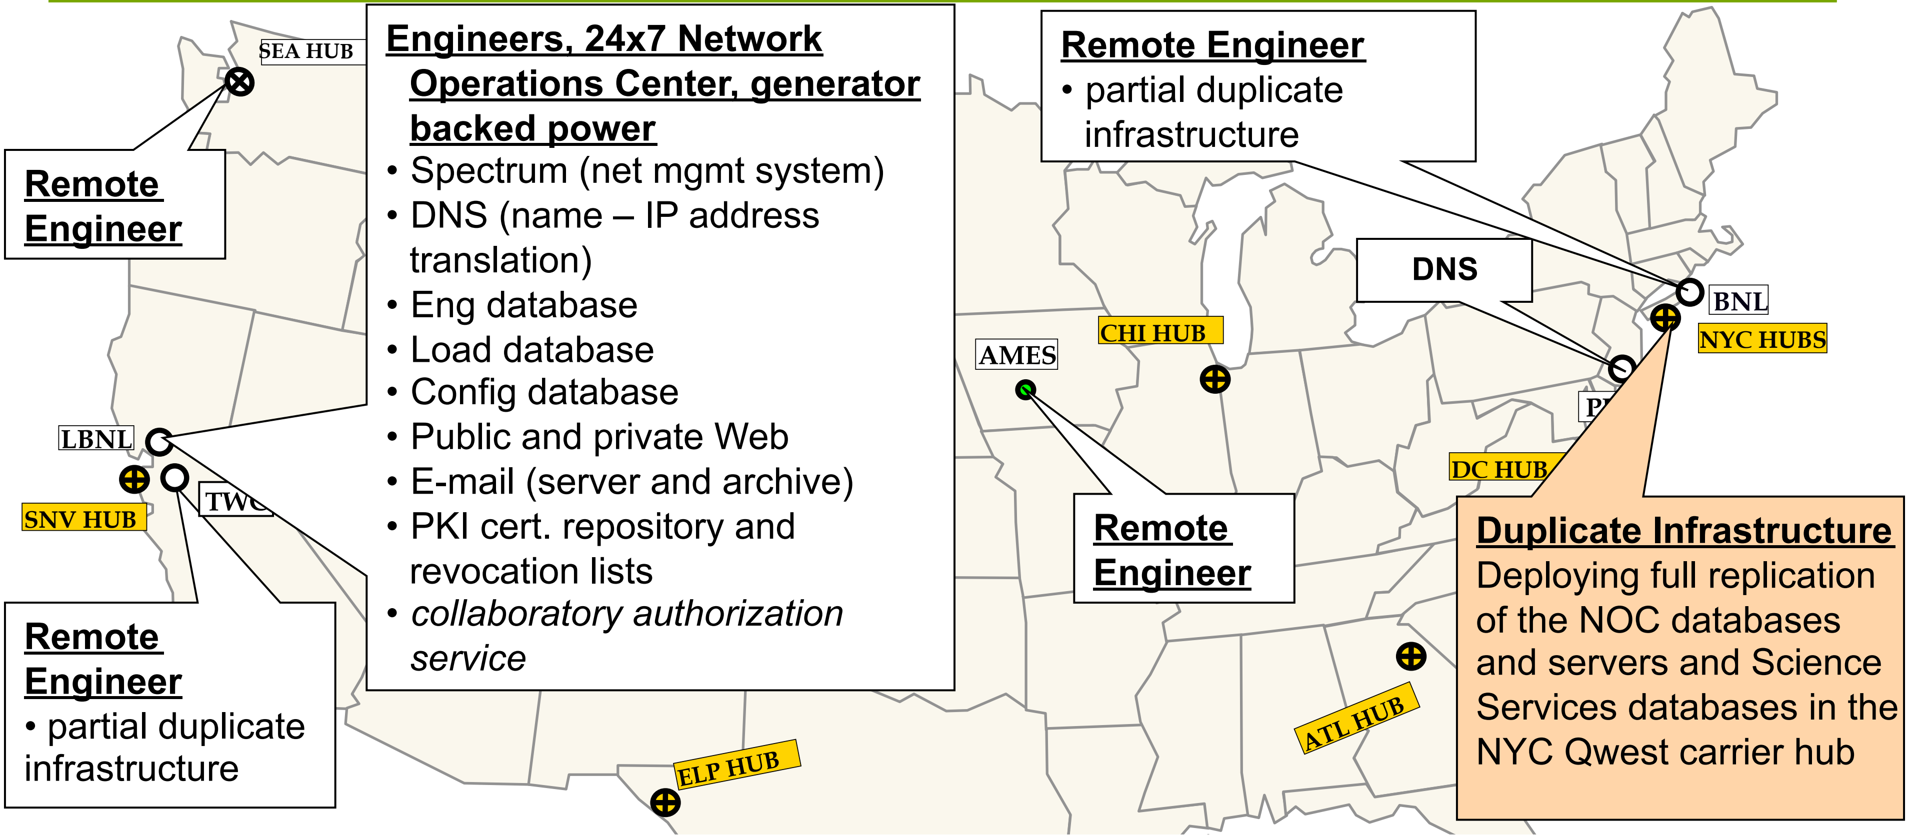

- The network must be kept available even if, e.g., the West Coast is disabled by a massive earthquake, etc.

Reliable operation of the network involves

- remote Network Operation Centers (4)
- replicated support infrastructure
- generator backed UPS power at all critical network and infrastructure locations

- high physical security for all equipment
- non-interruptible core - **<u>ESnet core operated without interruption</u>** through
  - N. Calif. Power blackout of 2000
  - the 9/11/2001 attacks, and
  - the Sept., 2003 NE States power blackout

ESnet Maintaining Science Mission Critical Infrastructure in the Face of a Cyberattack

- ESnet has tools in place to allow ESnet engineers to respond to extraordinary situations
- The tools involve things like
  - Separating ESnet core routing functionality from external commercial Internet connections by means of a “peering” router that can have a policy different from the core routers
  - Reducing the priority of traffic from the Internet so that it will not compete with site-to-site traffic
  - Applying filter based signatures to incoming Internet traffic depending on the nature of the attack
- One example of an “extraordinary situation” is the large-scale denial-of-service attack that occurred in early 2003
- ESnet engineers used a phased response that ranged from blocking the “bad” traffic at the ingress points to blocking infected sites as they (autonomously) start to attack other ESnet sites
  - A complete isolation of ESnet from the commercial Internet is possible

Historical Example of ESnet Response to DDOS Attack

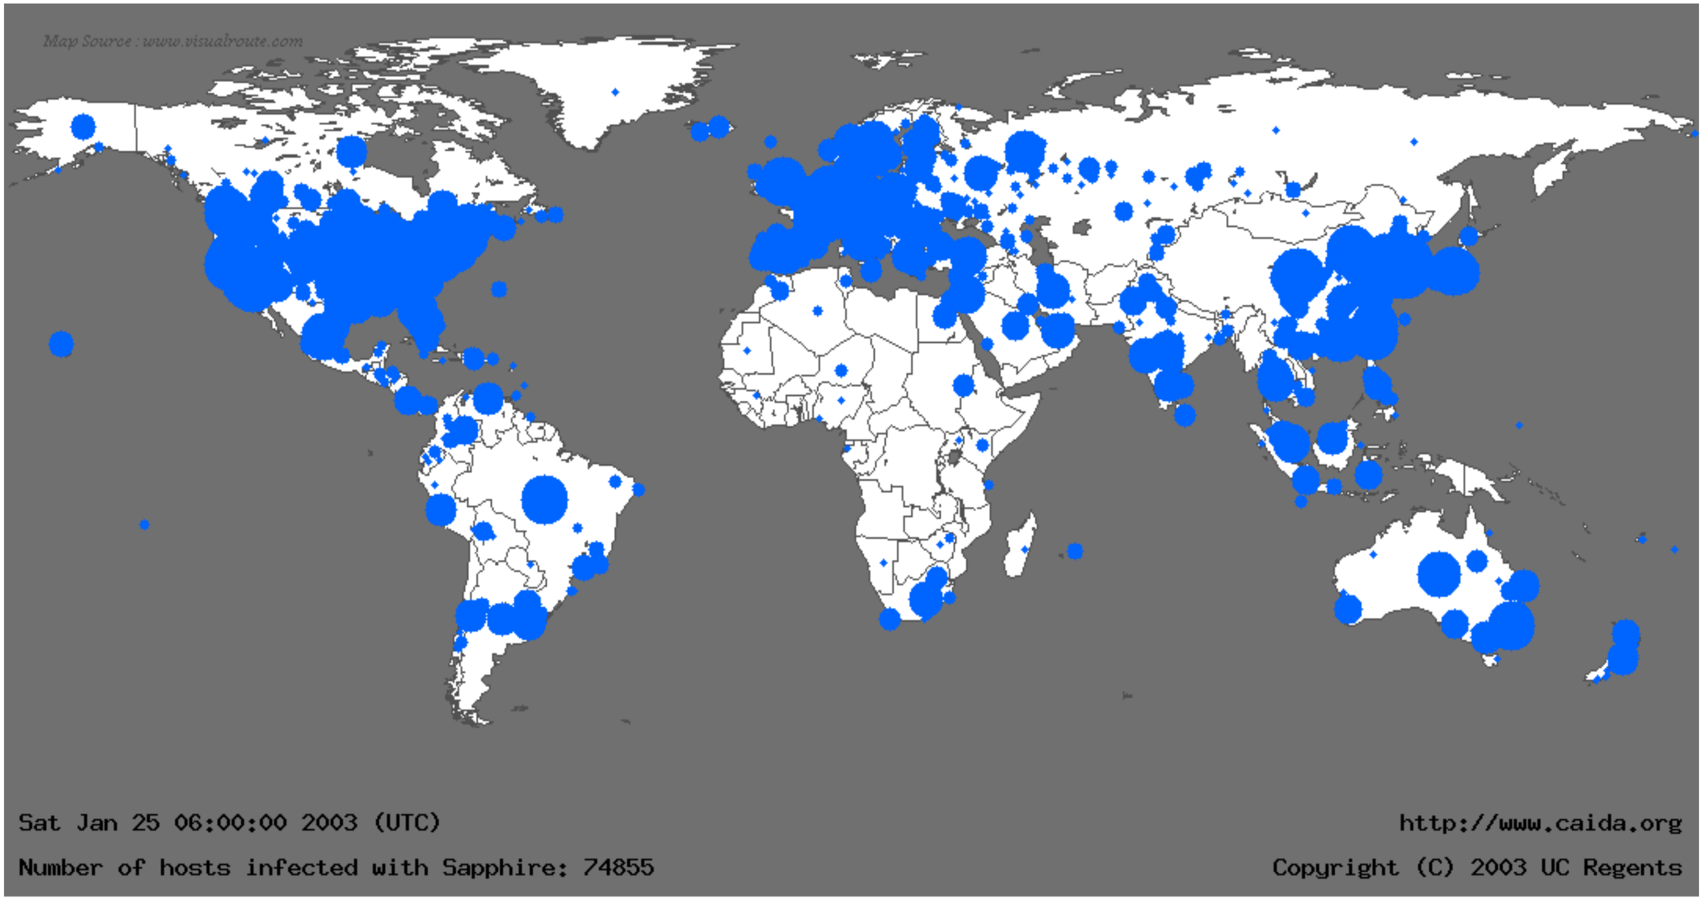

30 minutes after the Sapphire/Slammer worm was released, 75,000 hosts running Microsoft's SQL Server (port 1434) were infected.

(From "The Spread of the Sapphire/Slammer Worm," David Moore (CAIDA & UCSD CSE), Vern Paxson (ICIR & LBNL), Stefan Savage (UCSD CSE), Colleen Shannon (CAIDA), Stuart Staniford (Silicon Defense), and Nicholas Weaver (Silicon Defense & UC Berkeley EECS). Jan., 2003. (http://www.caida.org/publications/papers/2003/sapphire/))

Historical Example of ESnet Response to DDOS Attack

Sapphire/Slammer worm infection hits creating almost a full Gb/s (1000 megabit/sec.) traffic spike on the ESnet backbone

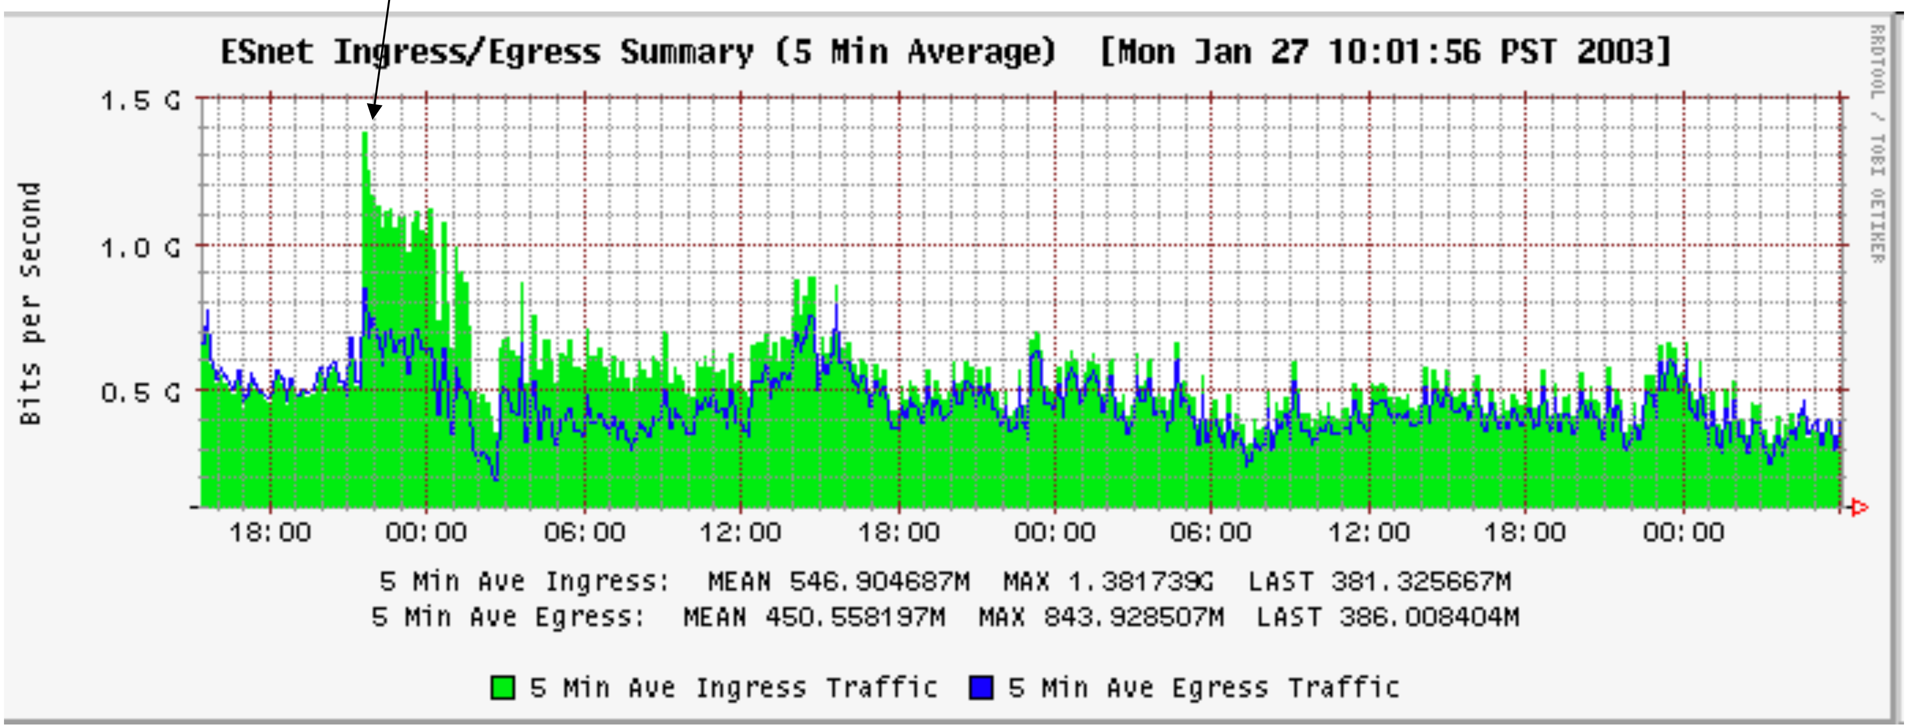

ESnet Cyberattack Defense

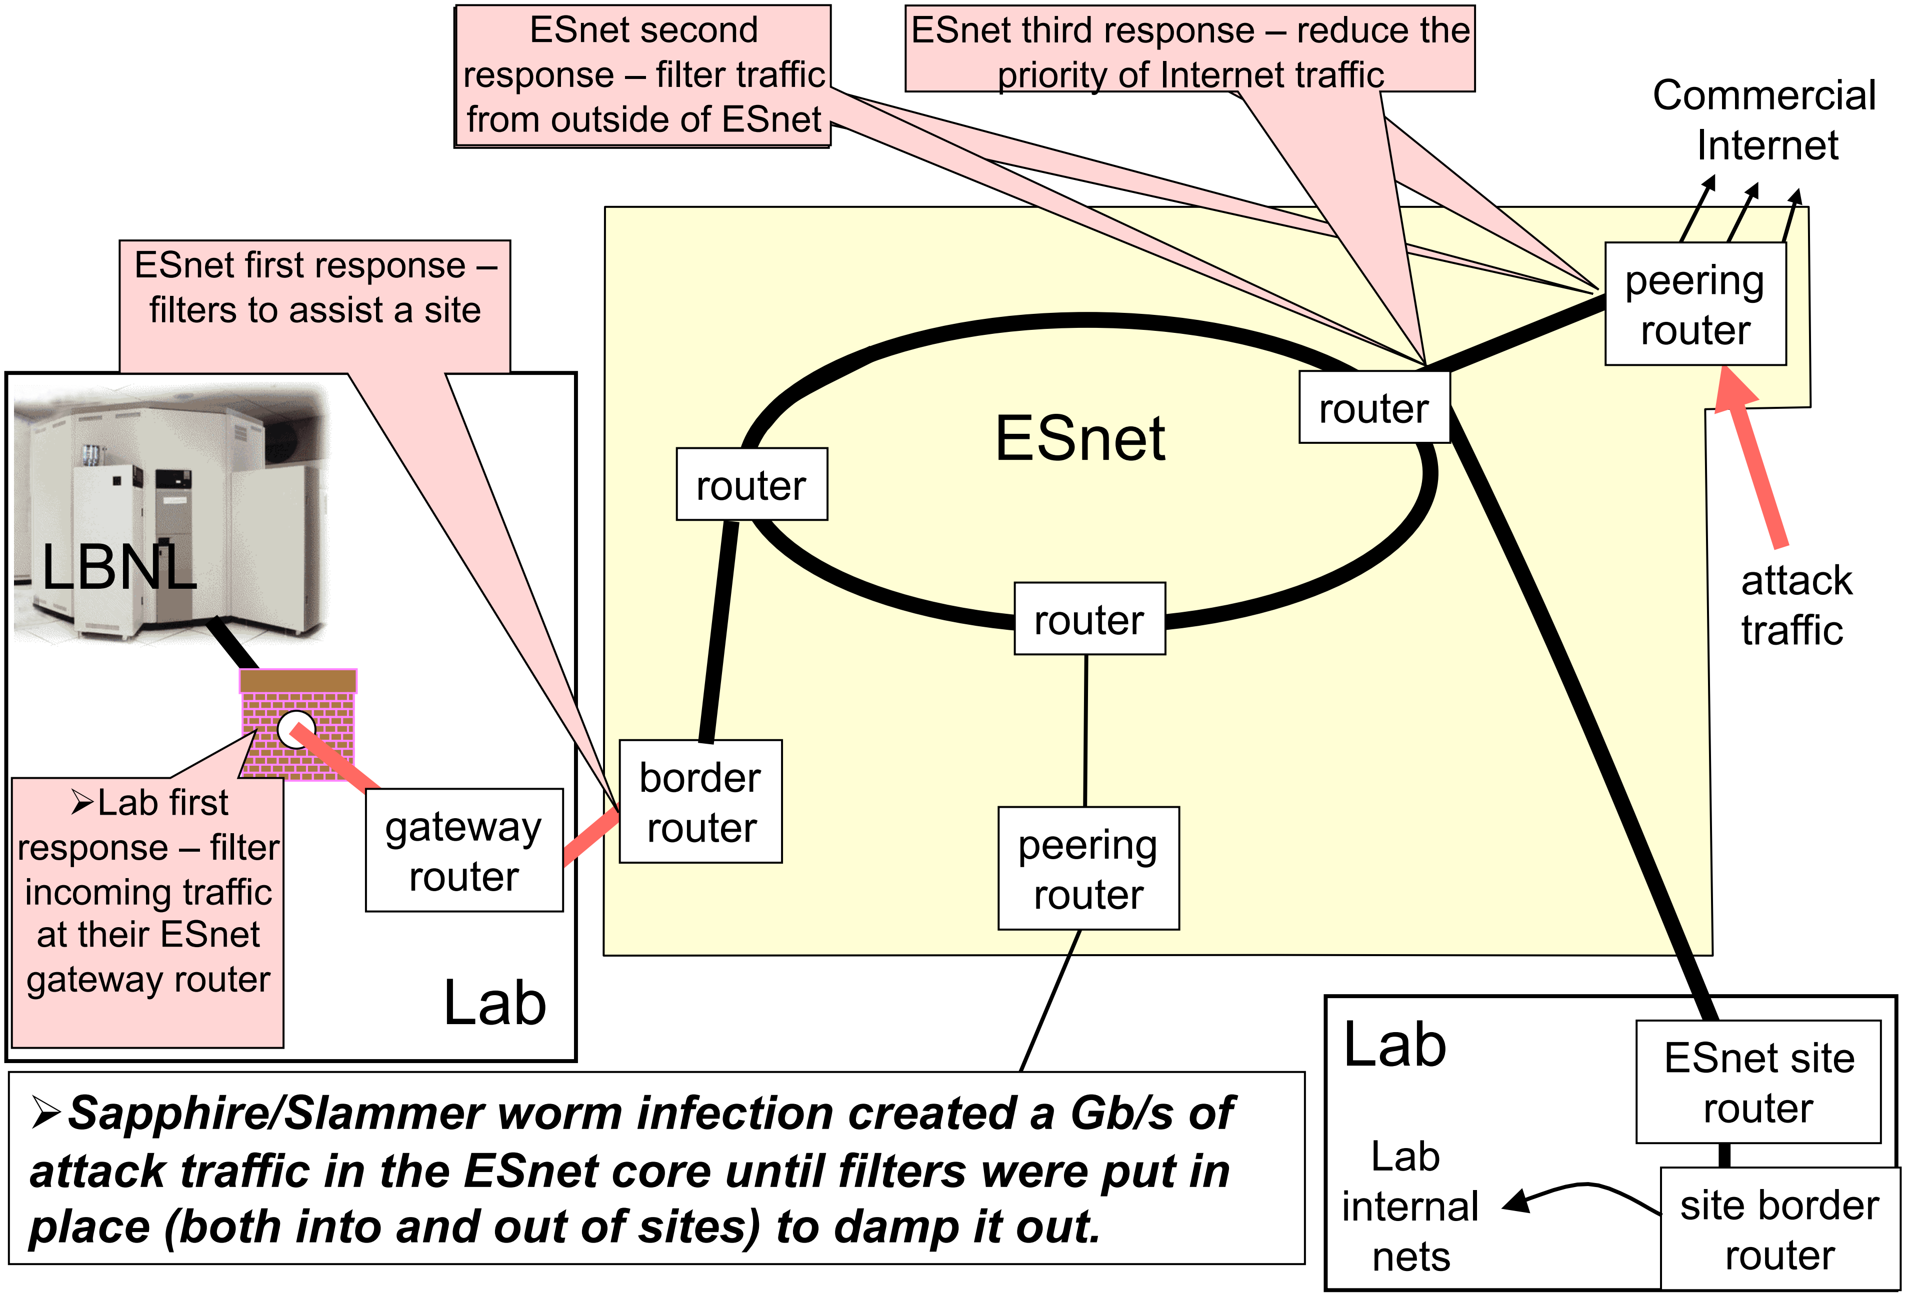

- ***Sapphire/Slammer worm infection created a Gb/s of attack traffic in the ESnet core until filters were put in place (both into and out of sites) to damp it out.***

➢A Changing Science Environment is the Key Driver of the Next Generation ESnet

- Large-scale collaborative science – big facilities, massive data, thousands of collaborators – is now a significant, if not dominate, aspect of the Office of Science ("SC") program
- SC science community is almost equally split between Labs and universities
  - SC facilities have users worldwide
- Very large non-US international facilities (e.g. LHC and ITER) and international collaborators are now a key element of SC science
- Distributed systems - e.g. Grids - for data analysis, simulations, instrument operation, etc., are essential and are now common (in fact dominate the data analysis that now generates 50% of all ESnet traffic)

Planning and Building the Future Network - ESnet4

- Requirements are primary drivers for ESnet – science focused
- Sources of Requirements
  1. Office of Science (SC) Program Managers
     - The Program Offices Requirements Workshops
       - BES completed
       - BER in July, 2007
       - FES in March, 2008
       - Others to follow at the rate of 3 a year
  2. Direct gathering through interaction with science users of the network
     - Example case studies (updated 2005/2006)
       - Magnetic Fusion
       - Large Hadron Collider (LHC)
       - Climate Modeling
       - Spallation Neutron Source
  3. Observation of traffic trends in the network
- Requirements aggregation
  - Convergence on a complete set of network requirements

1. Network Requirements Workshop Example: Basic Energy Sciences (BES)

- Input from BES facilities, science programs and sites
  - Light Sources
  - SNS at ORNL, Neutron Science program
  - Nanoscience Centers
  - Combustion Research
  - Computational Chemistry
  - Other existing facilities (e.g. National Center for Electron Microscopy at LBL)
  - Facilities currently undergoing construction (e.g. LCLS at SLAC)

Workshop Process

- Three inputs
  - Discussion of Program Office – goals, future projects, and science portfolio
  - Discussions with representatives of individual programs and facilities
  - Group discussion about common issues, future technologies (e.g. detector upgrades), etc.
- Additional discussion – ESnet4
  - Architecture
  - Deployment schedule
  - Future services

BES Workshop Findings (1)

- BES facilities are unlikely to provide the magnitude of load that we expect from the LHC
  - However, significant detector upgrades are coming in the next 3 years
  - LCLS may provide significant load
  - SNS data repositories may provide significant load
  - Theory and simulation efforts may provide significant load
- Broad user base
  - Makes it difficult to model facilities as anything other than point sources of traffic load
  - Requires wide connectivity
- Most facilities and disciplines expect significant increases in PKI service usage

2. Case Studies For Requirements

- Advanced Scientific Computing Research (ASCR)
  - NERSC
  - NLCF
- Basic Energy Sciences
  - Advanced Light Source
    - Macromolecular Crystallography
  - Chemistry/Combustion
  - Spallation Neutron Source
- Biological and Environmental
  - Bioinformatics/Genomics
  - Climate Science
- Fusion Energy Sciences
  - Magnetic Fusion Energy/ITER
- High Energy Physics
  - LHC
- Nuclear Physics
  - RHIC

Science Networking Requirements Aggregation Summary

Science Drivers Science Areas / Facilities	End2End Reliability	Connectivity	Today End2End Band width	5 years End2End Band width	Traffic Characteristics	Network Services
Magnetic Fusion Energy	99.999% (Impossible without full redundancy)	• DOE sites • US Universities • Industry	200+ Mbps	1 Gbps	• Bulk data • Remote control	• Guaranteed bandwidth • Guaranteed QoS • Deadline scheduling
NERSC (LBNL) and ACLF (ANL)	-	• DOE sites • US Universities • International • Other ASCR supercomputers	10 Gbps	20 to 40 Gbps	• Bulk data • Remote control • Remote file system sharing	• Guaranteed bandwidth • Guaranteed QoS • Deadline Scheduling • PKI / Grid
NLCF (ORNL)	-	• DOE sites • US Universities • Industry • International	Backbone Band width parity	Backbone band width parity	• Bulk data • Remote file system sharing	
Nuclear Physics (RHIC - BNL)	-	• DOE sites • US Universities • International	12 Gbps	70 Gbps	• Bulk data	• Guaranteed bandwidth • PKI / Grid
Spallation Neutron Source (ORNL)	High (24x7 operation)	• DOE sites	640 Mbps	2 Gbps	• Bulk data	

Science Network Requirements Aggregation Summary

Science Drivers Science Areas / Facilities	End2End Reliability	Connectivity	Today End2End Band width	5 years End2End Band width	Traffic Characteristics	Network Services
Advanced Light Source	-	• DOE sites • US Universities • Industry	1 TB/day 300 Mbps	5 TB/day 1.5 Gbps	• Bulk data • Remote control	• Guaranteed bandwidth • PKI / Grid
Bioinformatics	-	• DOE sites • US Universities	625 Mbps 12.5 Gbps in two years		• Bulk data • Remote control • Point-to-multipoint	• Guaranteed bandwidth • High-speed multicast
Chemistry / Combustion	-	• DOE sites • US Universities • Industry	-	10s of Gigabits per second	• Bulk data	• Guaranteed bandwidth • PKI / Grid
Climate Science	-	• DOE sites • US Universities • International	-	5 PB per year 5 Gbps	• Bulk data • Remote control	• Guaranteed bandwidth • PKI / Grid
		Immediate Requirements and Drivers				
High Energy Physics (LHC)	99.95+% (Less than 4 hrs/year)	• US Tier1 (FNAL, BNL) • US Tier2 (Universities) • International (Europe, Canada)	10 Gbps	60 to 80 Gbps (30-40 Gbps per US Tier1)	• Bulk data • Coupled data analysis processes	• Guaranteed bandwidth • Traffic isolation • PKI / Grid

LHC will be the largest scientific experiment and generate the most data that the scientific community has ever tried to manage.
The data management model involves a world-wide collection of data centers that store, manage, and analyze the data and that are integrated through network connections with typical speeds in the 10+ Gbps range.

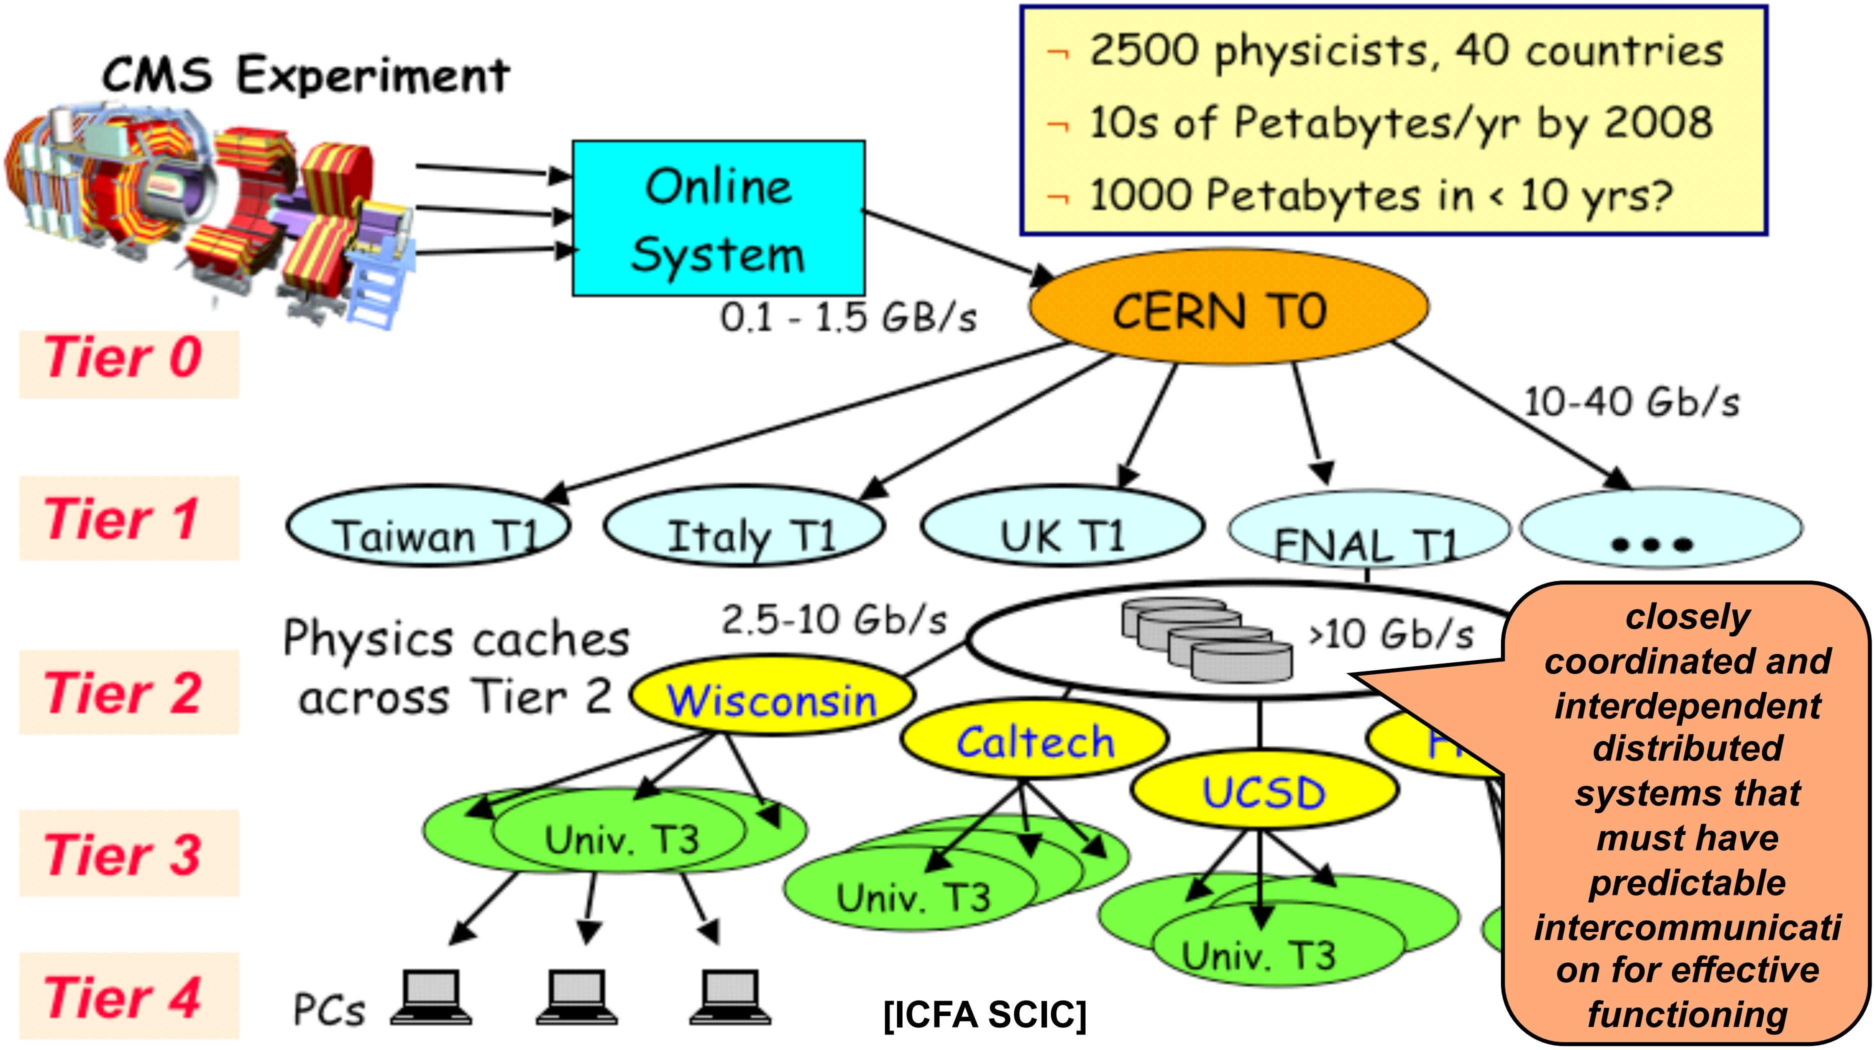

[ICFA SCIC]

The Next Level of Detail In Network Requirements

- Consider the LHC networking as an example, because their requirements are fairly well characterized

The Next Level of Detail: LHC Tier 0, 1, and 2 Connectivity Requirements Summary

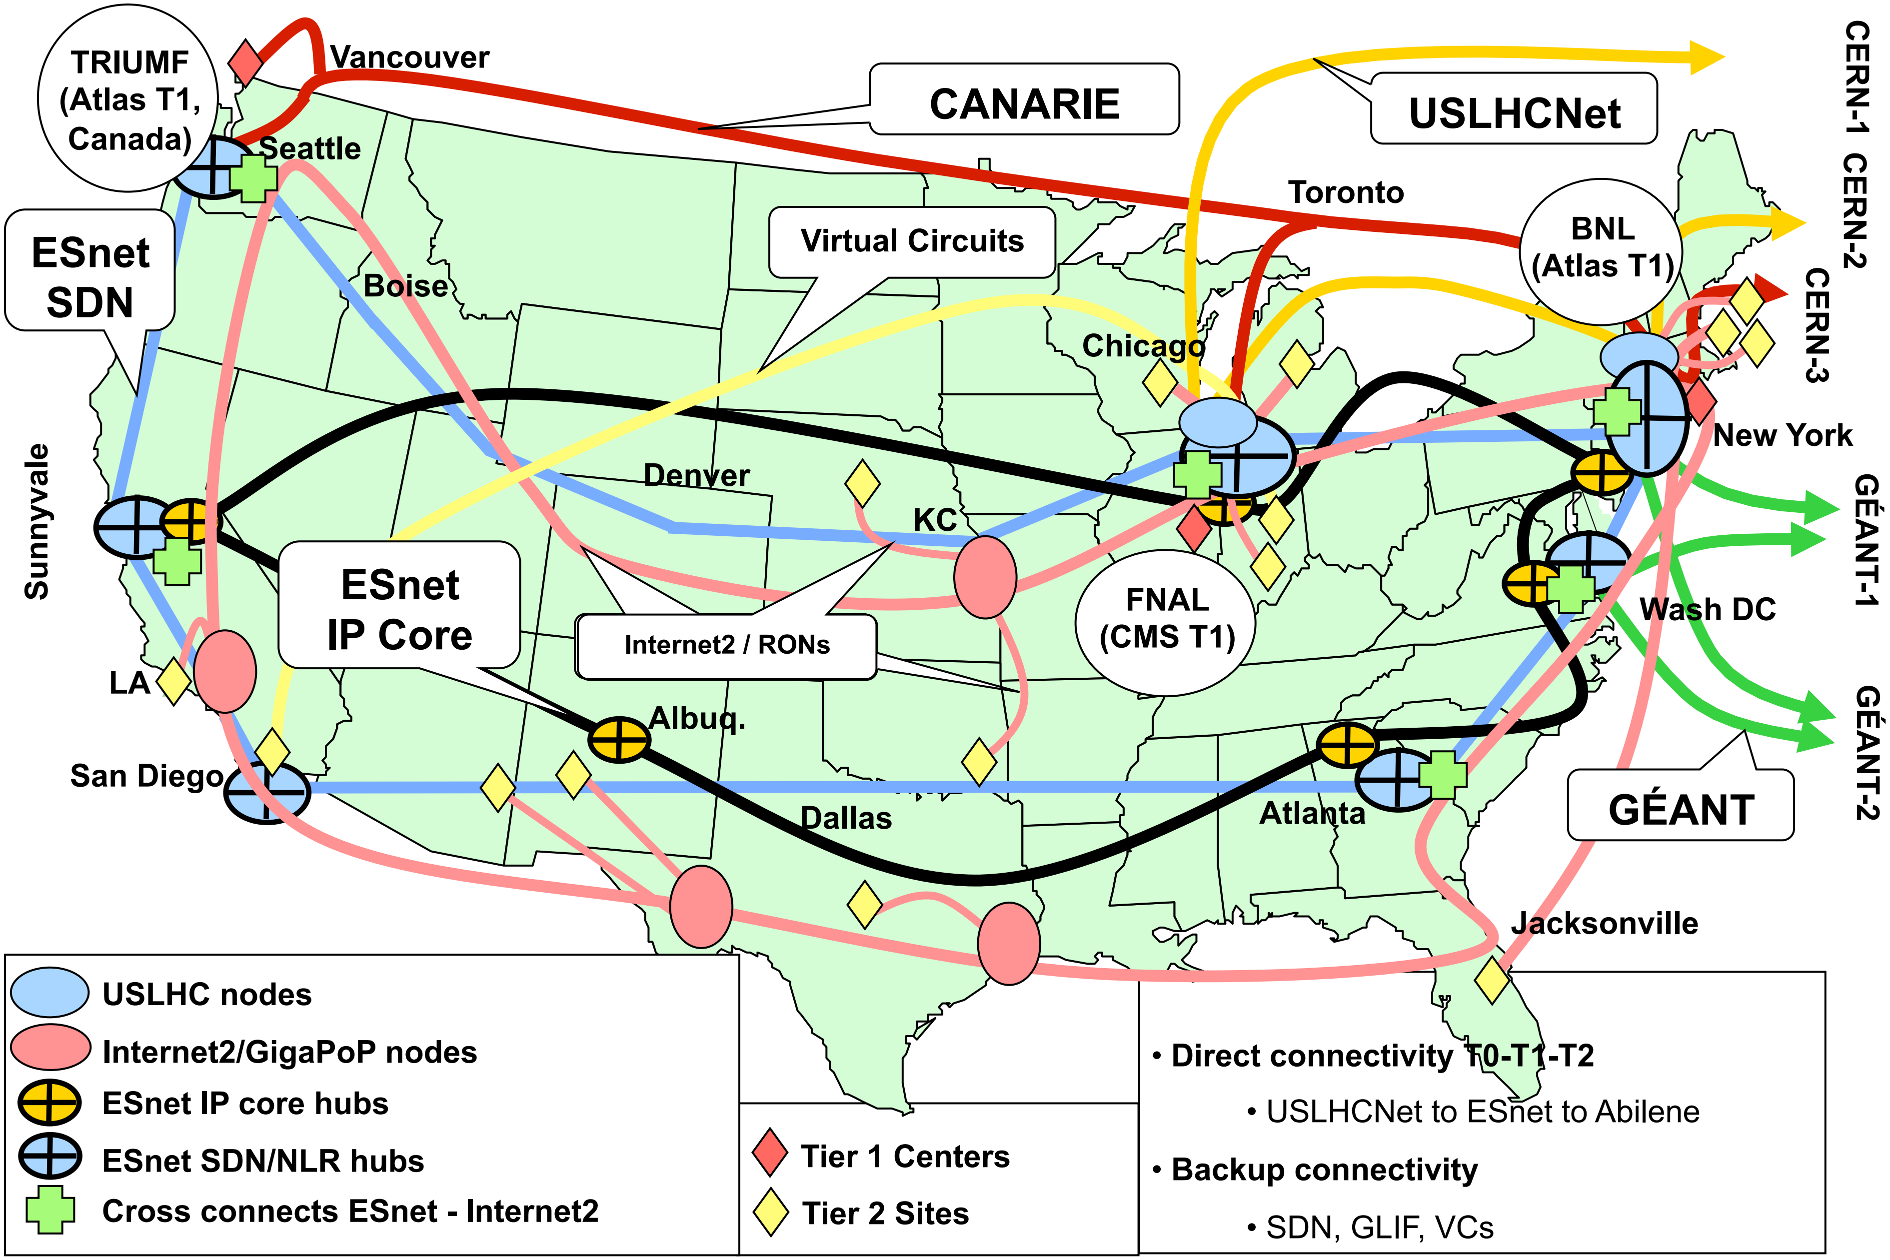

- Direct connectivity T0-T1-T2
    - USLHCNet to ESnet to Abilene
- Backup connectivity
    - SDN, GLIF, VCs

The Next Level of Detail:
LHC ATLAS Bandwidth Matrix as of April 2007

Site A	Site Z	ESnet A	ESnet Z	A-Z 2007 Bandwidth	A-Z 2010 Bandwidth
CERN	BNL	AofA (NYC)	BNL	10Gbps	20-40Gbps
BNL	U. of Michigan (Calibration)	BNL (LIMAN)	Starlight (CHIMAN)	3Gbps	10Gbps
BNL	Boston University	BNL (LIMAN)	Internet2 / NLR Peerings	3Gbps (Northeastern Tier2 Center)	10Gbps (Northeastern Tier2 Center)
BNL	Harvard University				
BNL	Indiana U. at Bloomington	BNL (LIMAN)	Internet2 / NLR Peerings	3Gbps (Midwestern Tier2 Center)	10Gbps (Midwestern Tier2 Center)
BNL	U. of Chicago				
BNL	Langston University	BNL (LIMAN)	Internet2 / NLR Peerings	3Gbps (Southwestern Tier2 Center)	10Gbps (Southwestern Tier2 Center)
BNL	U. Oklahoma Norman				
BNL	U. of Texas Arlington				
BNL	Tier3 Aggregate	BNL (LIMAN)	Internet2 / NLR Peerings	5Gbps	20Gbps
BNL	TRIUMF (Canadian ATLAS Tier1)	BNL (LIMAN)	Seattle	1Gbps	5Gbps

LHC CMS Bandwidth Matrix as of April 2007

Site A	Site Z	ESnet A	ESnet Z	A-Z 2007 Bandwidth	A-Z 2010 Bandwidth
CERN	FNAL	Starlight (CHIMAN)	FNAL (CHIMAN)	10Gbps	20-40Gbps
FNAL	U. of Michigan (Calibration)	FNAL (CHIMAN)	Starlight (CHIMAN)	3Gbps	10Gbps
FNAL	Caltech	FNAL (CHIMAN)	Starlight (CHIMAN)	3Gbps	10Gbps
FNAL	MIT	FNAL (CHIMAN)	AofA (NYC)/ Boston	3Gbps	10Gbps
FNAL	Purdue University	FNAL (CHIMAN)	Starlight (CHIMAN)	3Gbps	10Gbps
FNAL	U. of California at San Diego	FNAL (CHIMAN)	San Diego	3Gbps	10Gbps
FNAL	U. of Florida at Gainesville	FNAL (CHIMAN)	SOX / Ultralight at Starlight	3Gbps	10Gbps
FNAL	U. of Nebraska at Lincoln	FNAL (CHIMAN)	Starlight (CHIMAN)	3Gbps	10Gbps
FNAL	U. of Wisconsin at Madison	FNAL (CHIMAN)	Starlight (CHIMAN)	3Gbps	10Gbps
FNAL	Tier3 Aggregate	FNAL (CHIMAN)	Internet2 / NLR Peerings	5Gbps	20Gbps

Estimated Aggregate Link Loadings, 2007-08

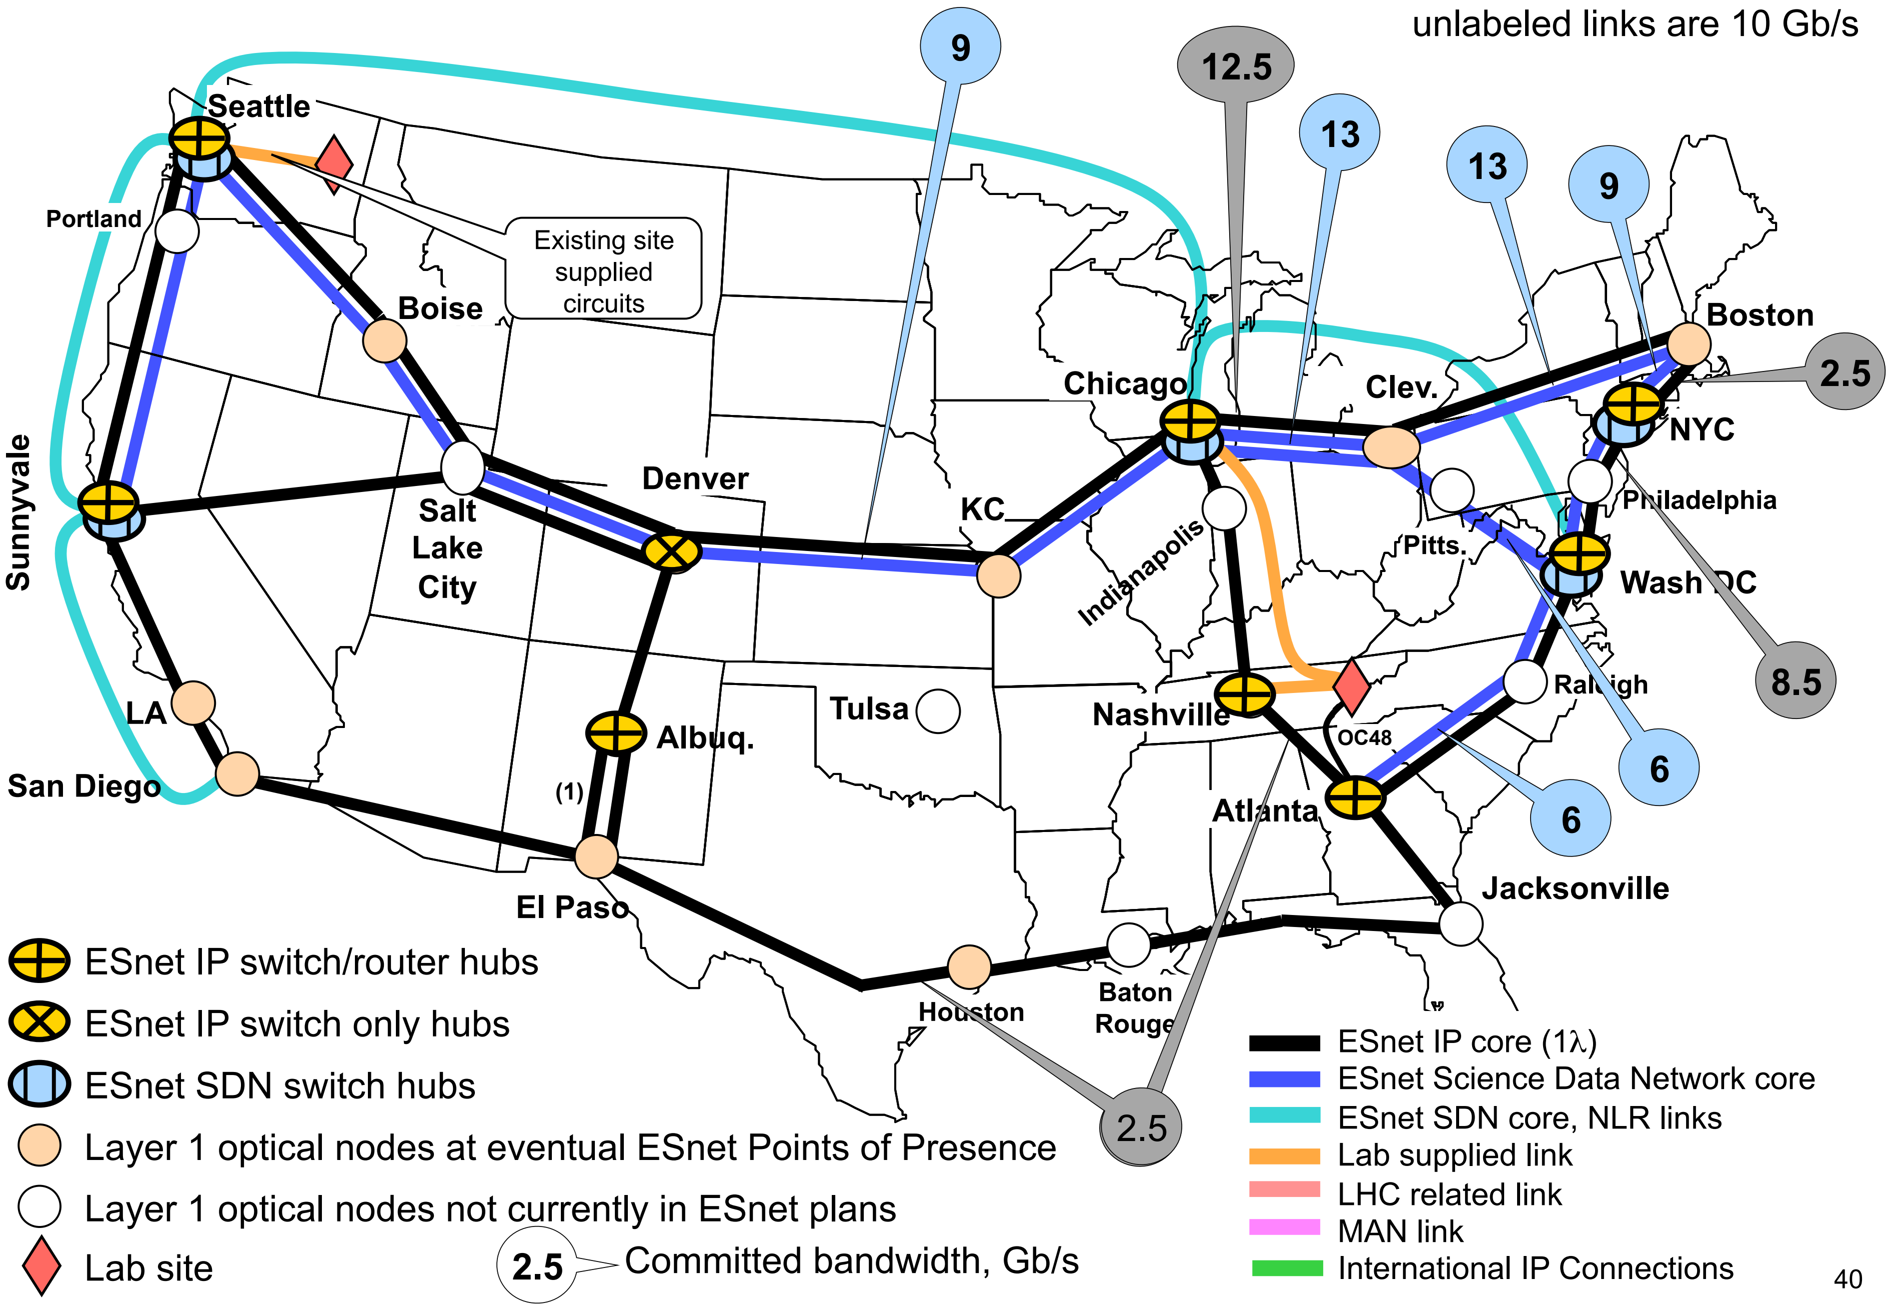

ESnet4 2007-8 Estimated Bandwidth Commitments

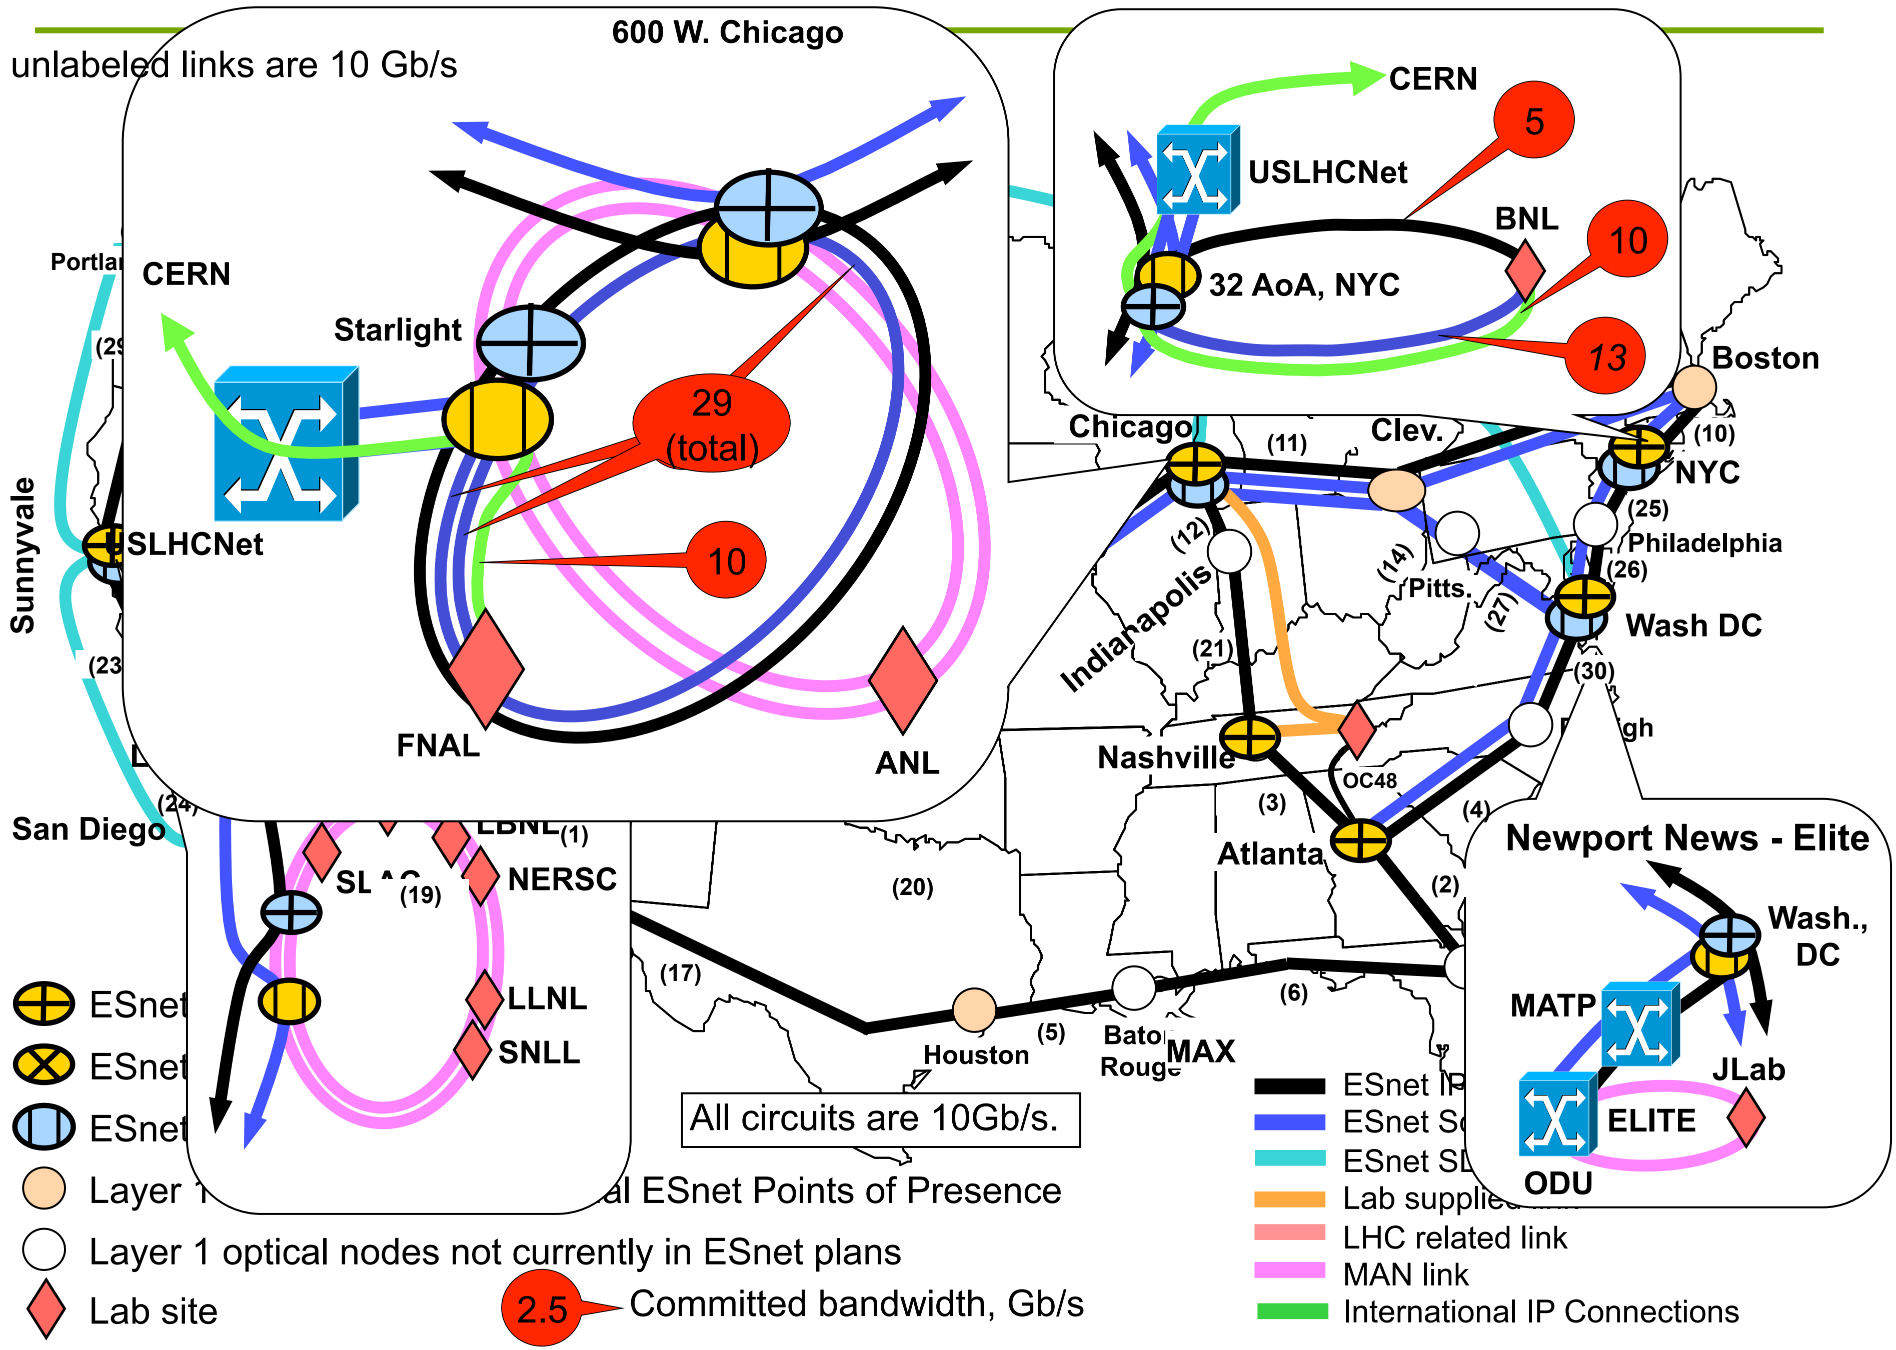

Are These Estimates Realistic? YES!

FNAL outbound CMS traffic for 4 months, to Sept. 1, 2007

<u>Max= 8.9 Gb/s</u> (1064 MBy/s of data), <u>Average = 4.1 Gb/s</u> (493 MBy/s of data)

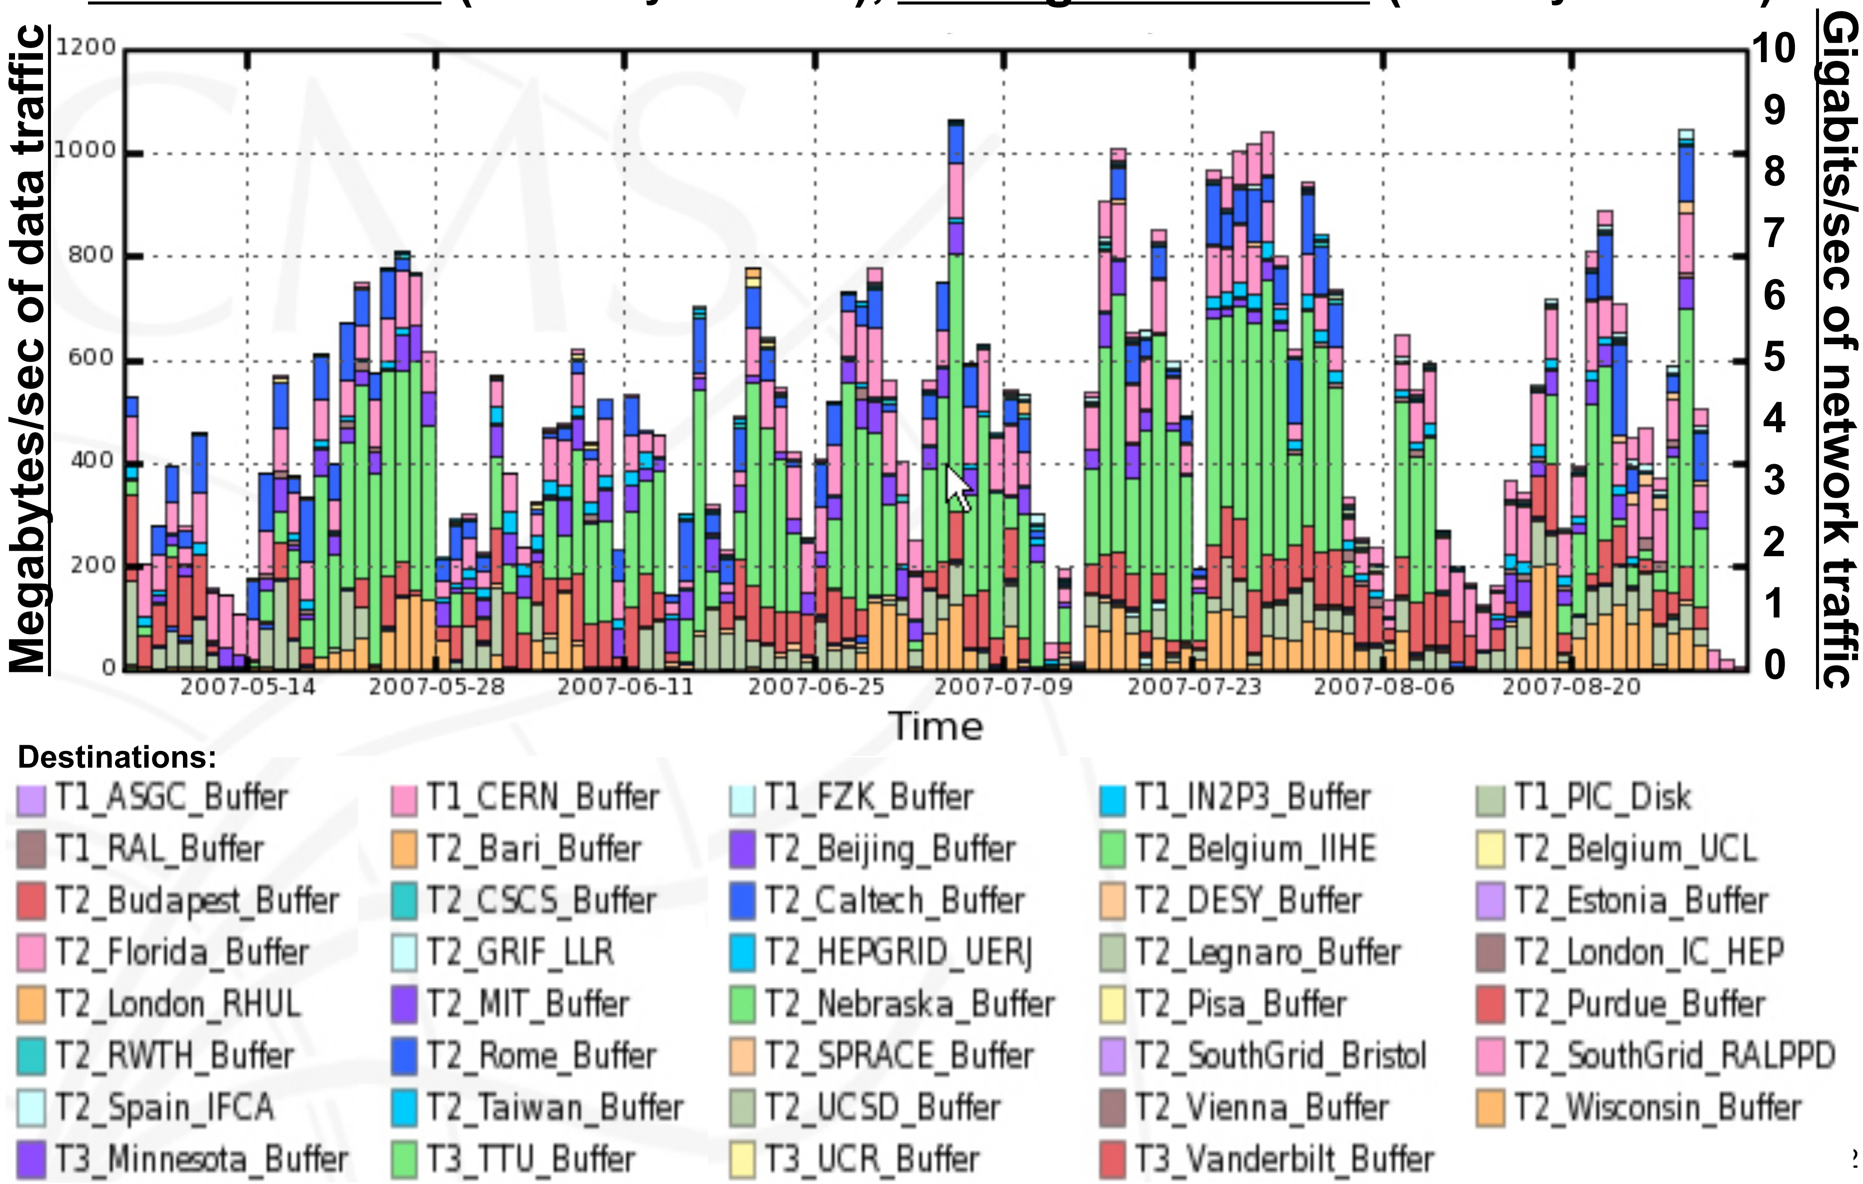

Estimated Aggregate Link Loadings, 2010-11

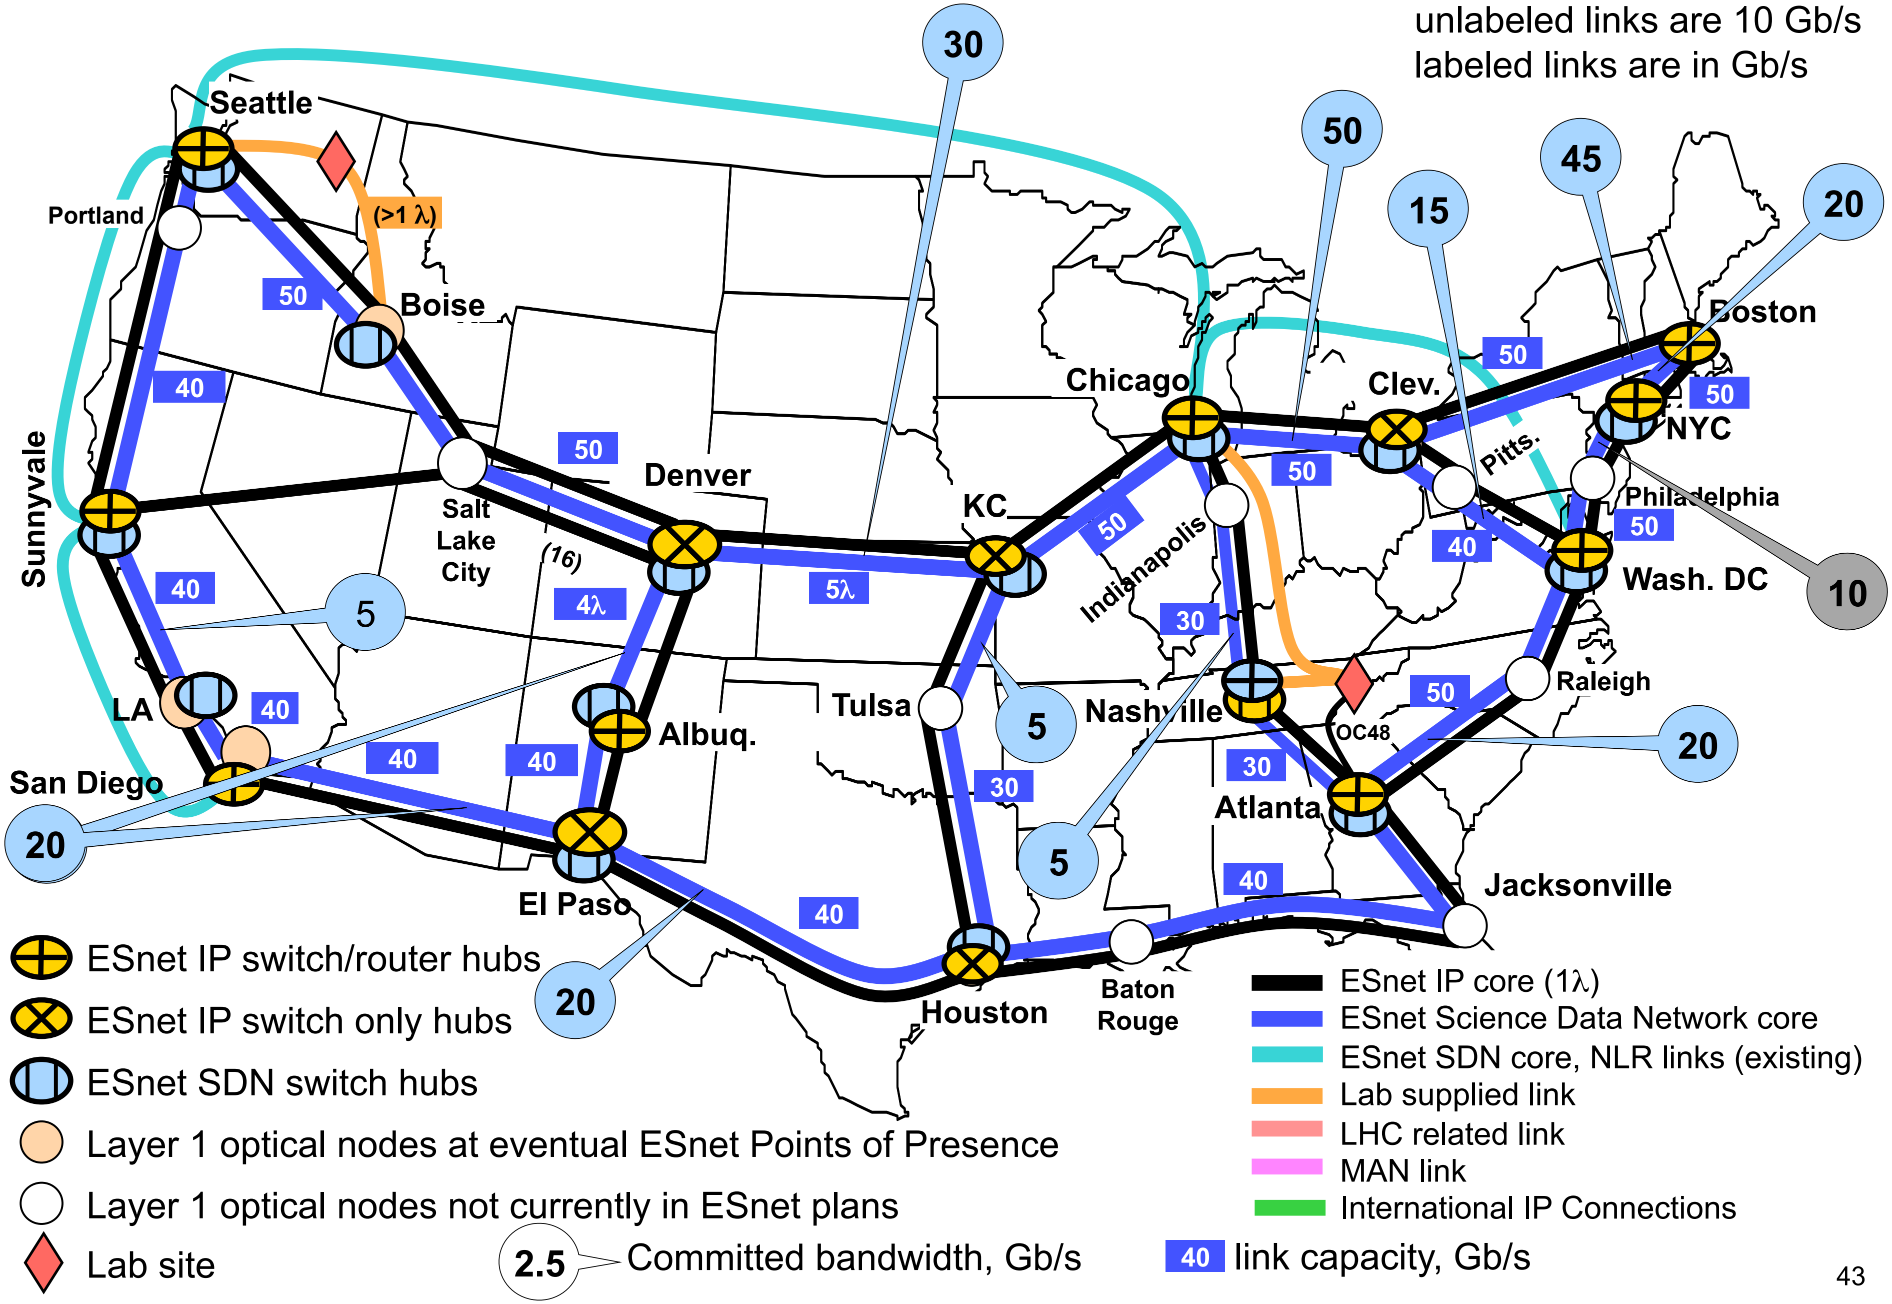

ESnet4 2010-11 Estimated Bandwidth Commitments

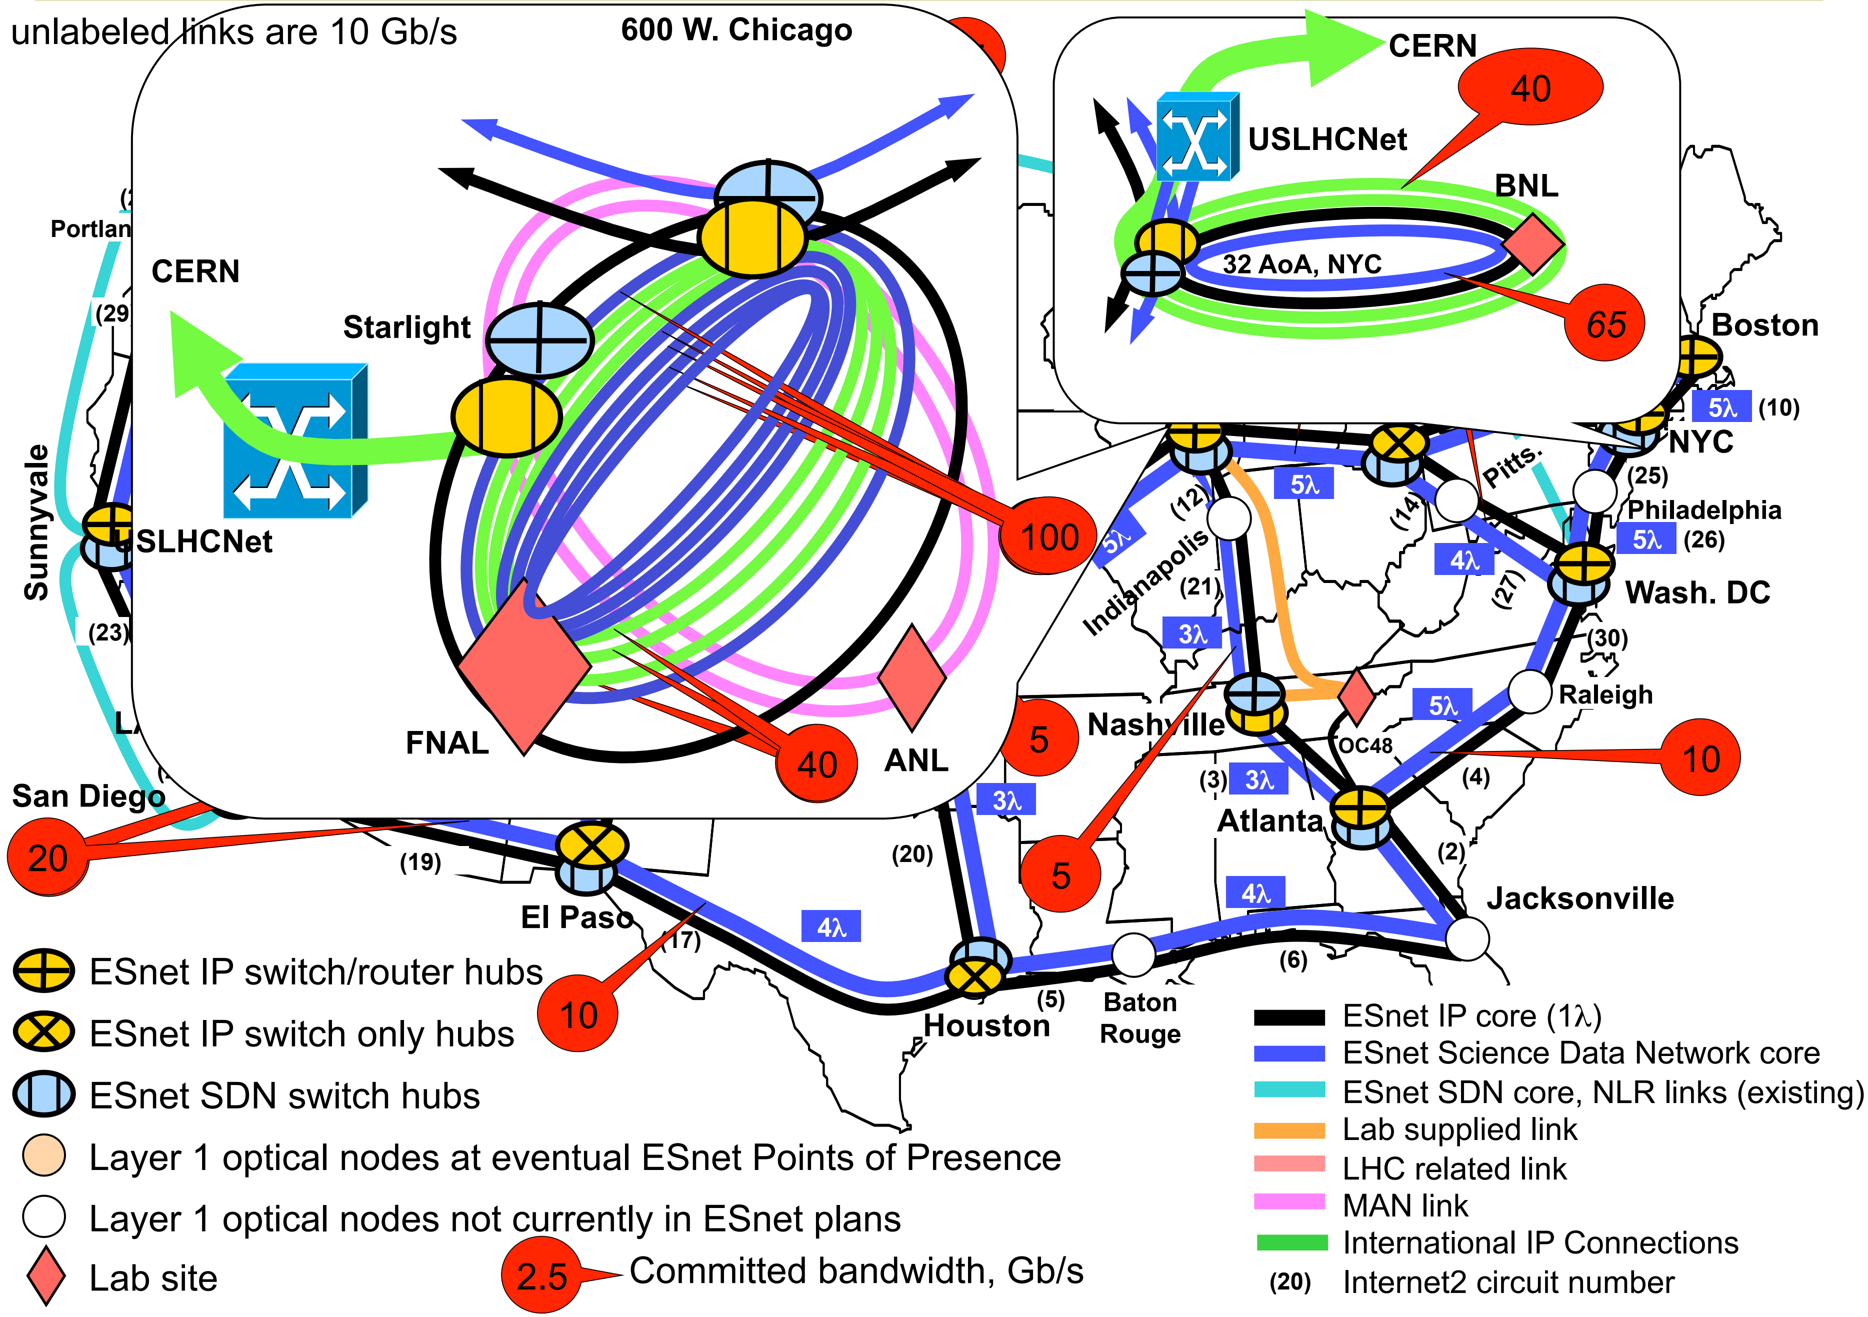

Large-Scale Data Analysis Systems (Typified by the LHC) have Several Characteristics that Result in Requirements for the Network and its Services

- The ***systems are data intensive and high-performance***, typically moving terabytes a day for months at a time
- The ***system are high duty-cycle***, operating most of the day for months at a time in order to meet the requirements for data movement
- The ***systems are widely distributed*** – typically spread over continental or inter-continental distances
- Such ***systems 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 applications ***must be able to get guarantees from the network*** that there is adequate bandwidth to accomplish the task at hand
- The applications ***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
- *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*

This slide drawn from [ICFA SCIC]

Example: Multidisciplinary Simulation

A "complete" approach to climate modeling involves many interacting models and data that are <u>provided by different groups at different locations</u>

(Tim Killeen, NCAR)

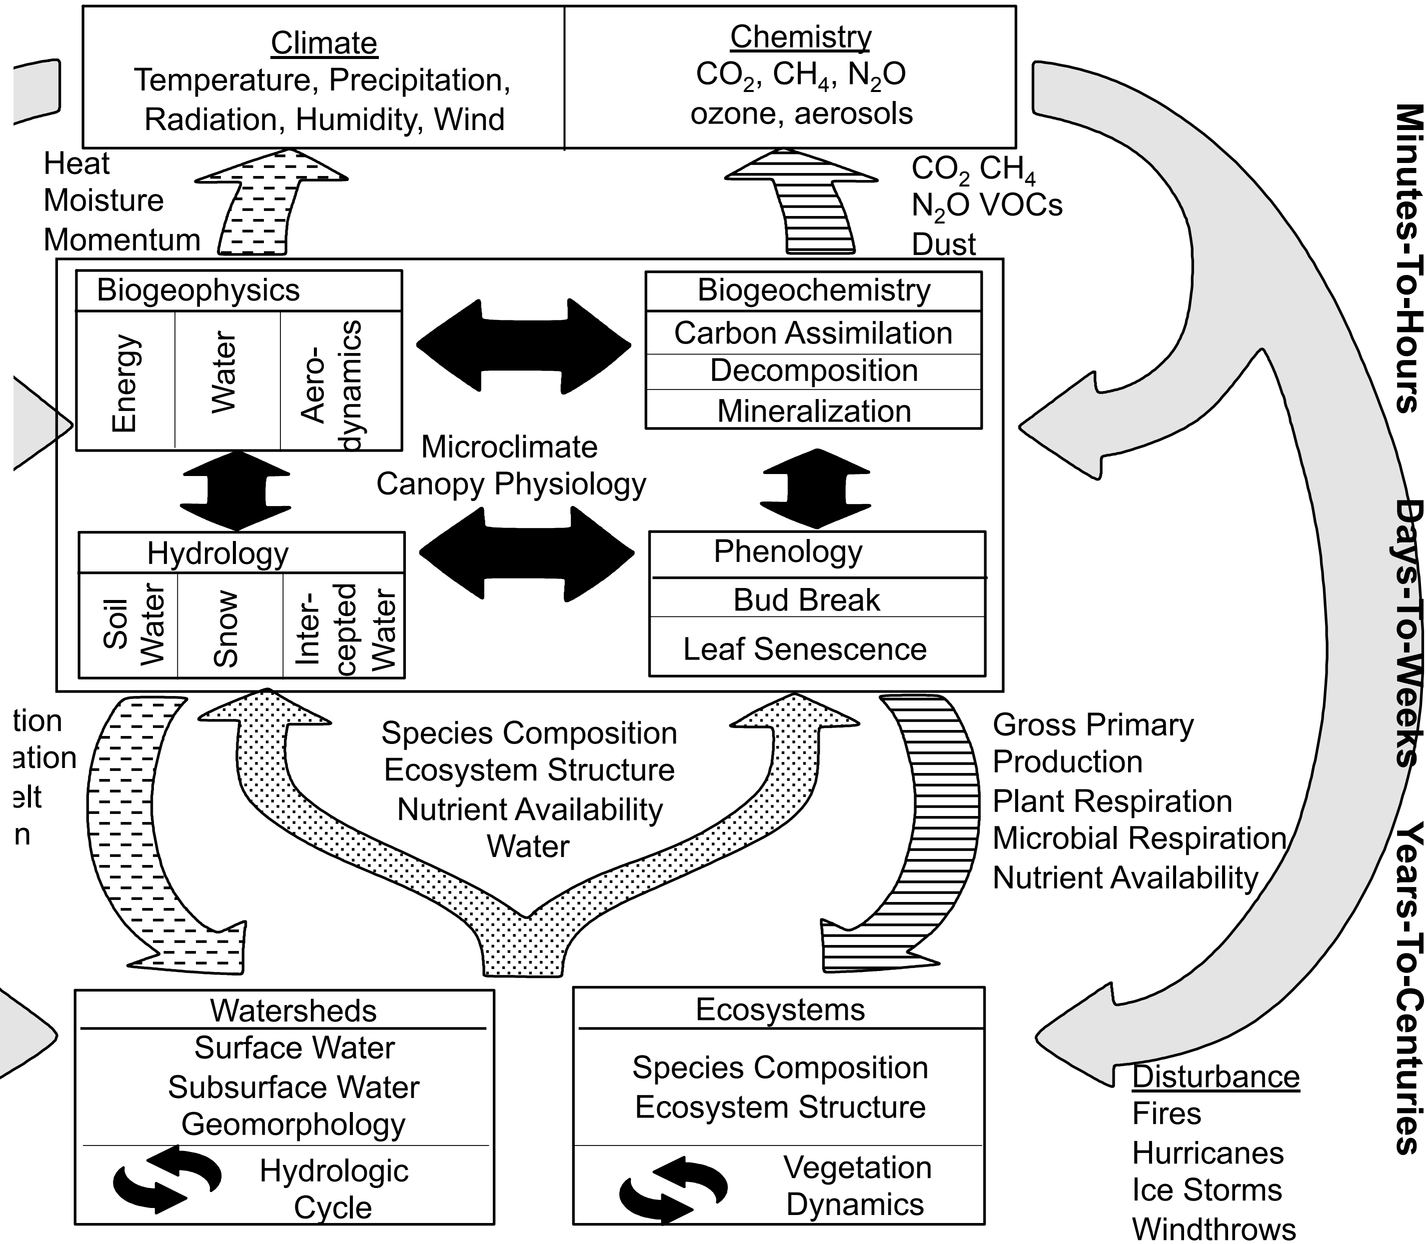

(Courtesy Gordon Bonan, NCAR: *Ecological Climatology: Concepts and Applications*. Cambridge University Press, Cambridge, 2002.)

Distributed Multidisciplinary Simulation

- Distributed multidisciplinary simulation involves integrating computing elements at several remote locations
  - *Requires co-scheduling of computing, data storage, and network elements*
  - Also Quality of Service (e.g. bandwidth guarantees)
  - There is not a lot of experience with large-scale versions of this scenario yet, but it is coming (e.g. the new Office of Science supercomputing facility at Oak Ridge National Lab has a distributed computing elements model)

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 that is service-oriented: configurable, schedulable, predictable, reliable, and informative – and the network and its services must be scalable

Capacity Planning by Observation

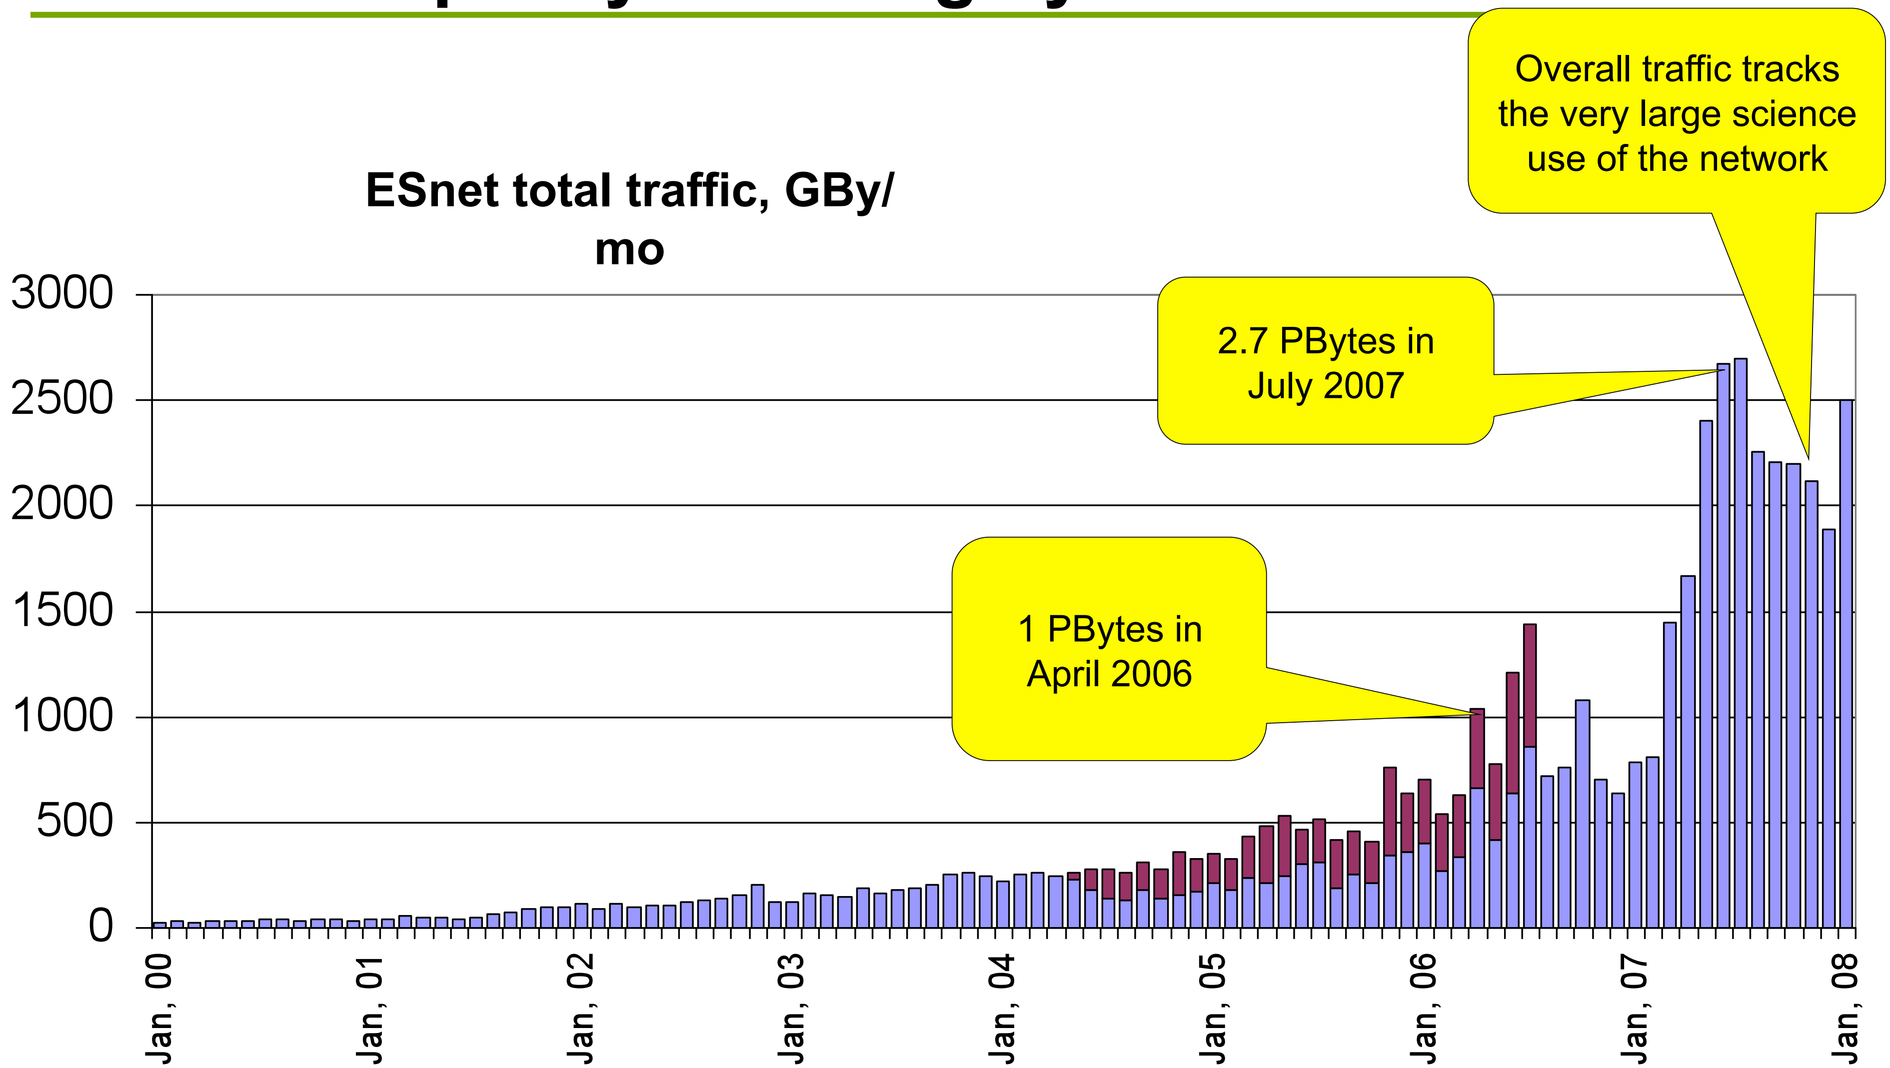

ESnet traffic historically has increased 10x every 47 months. However with large science usage now dominating the traffic the growth will track the science project usage of the network.

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

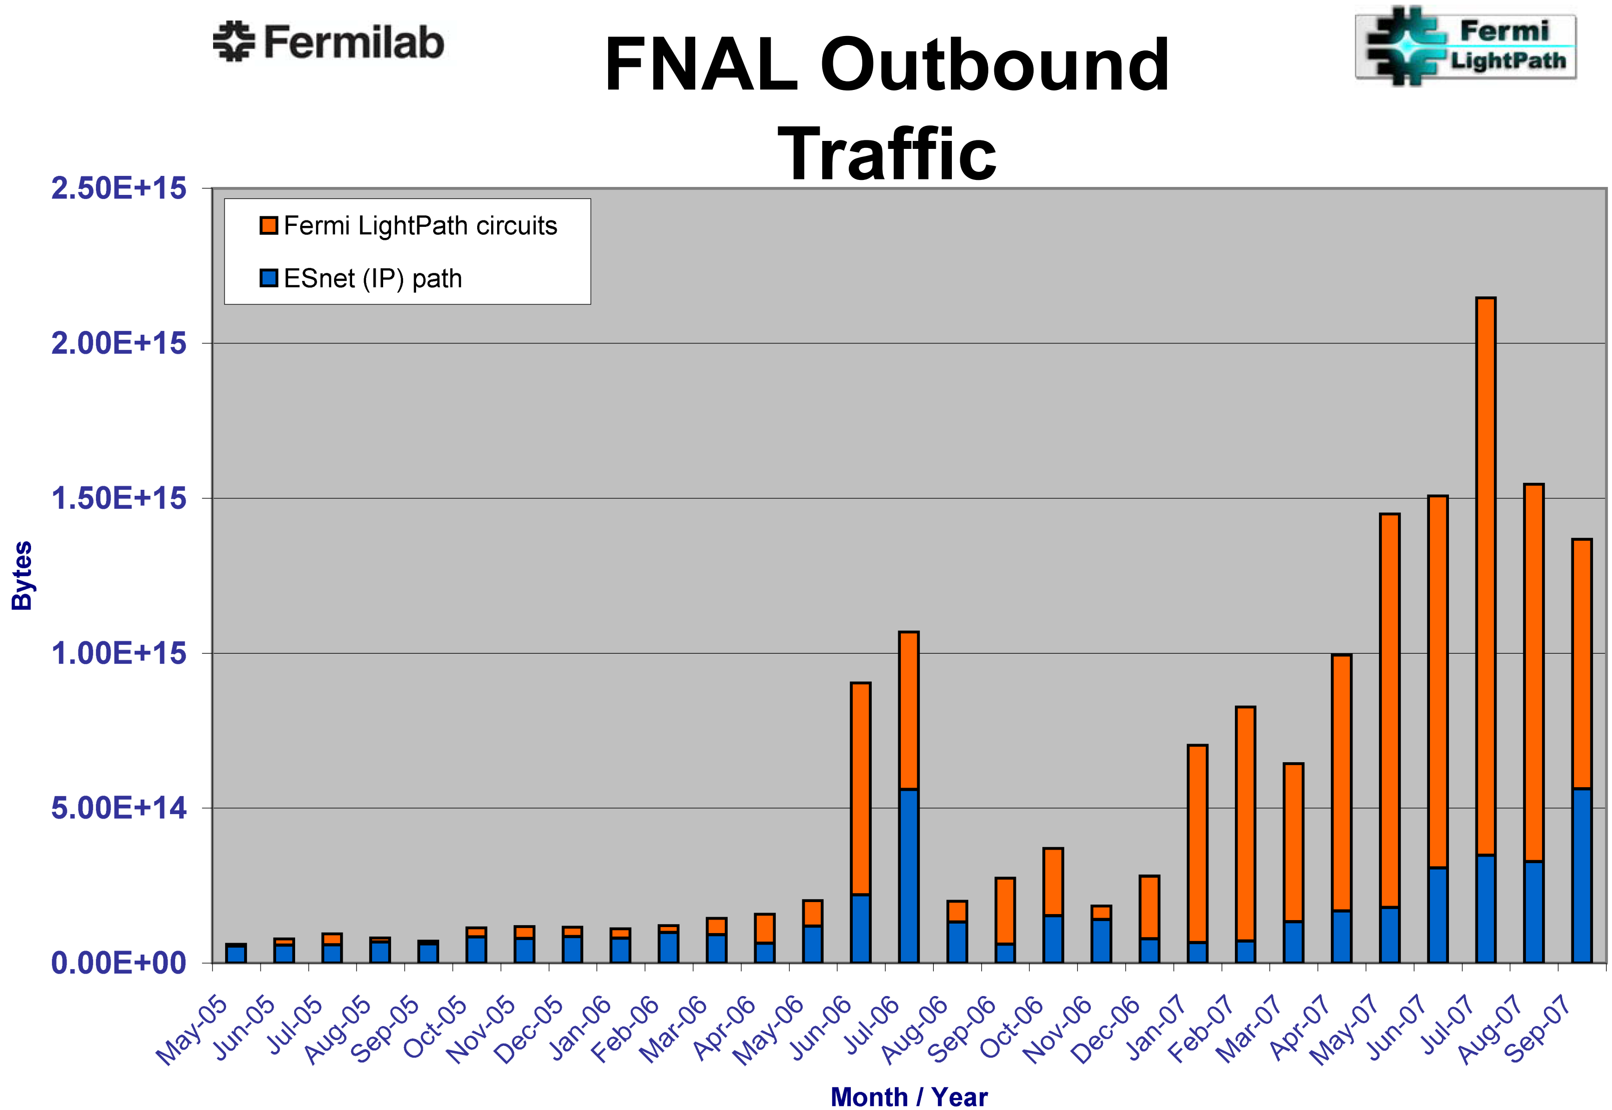

ESnet Traffic has Increased by 10X Every 47 Months, on Average, Since 1990

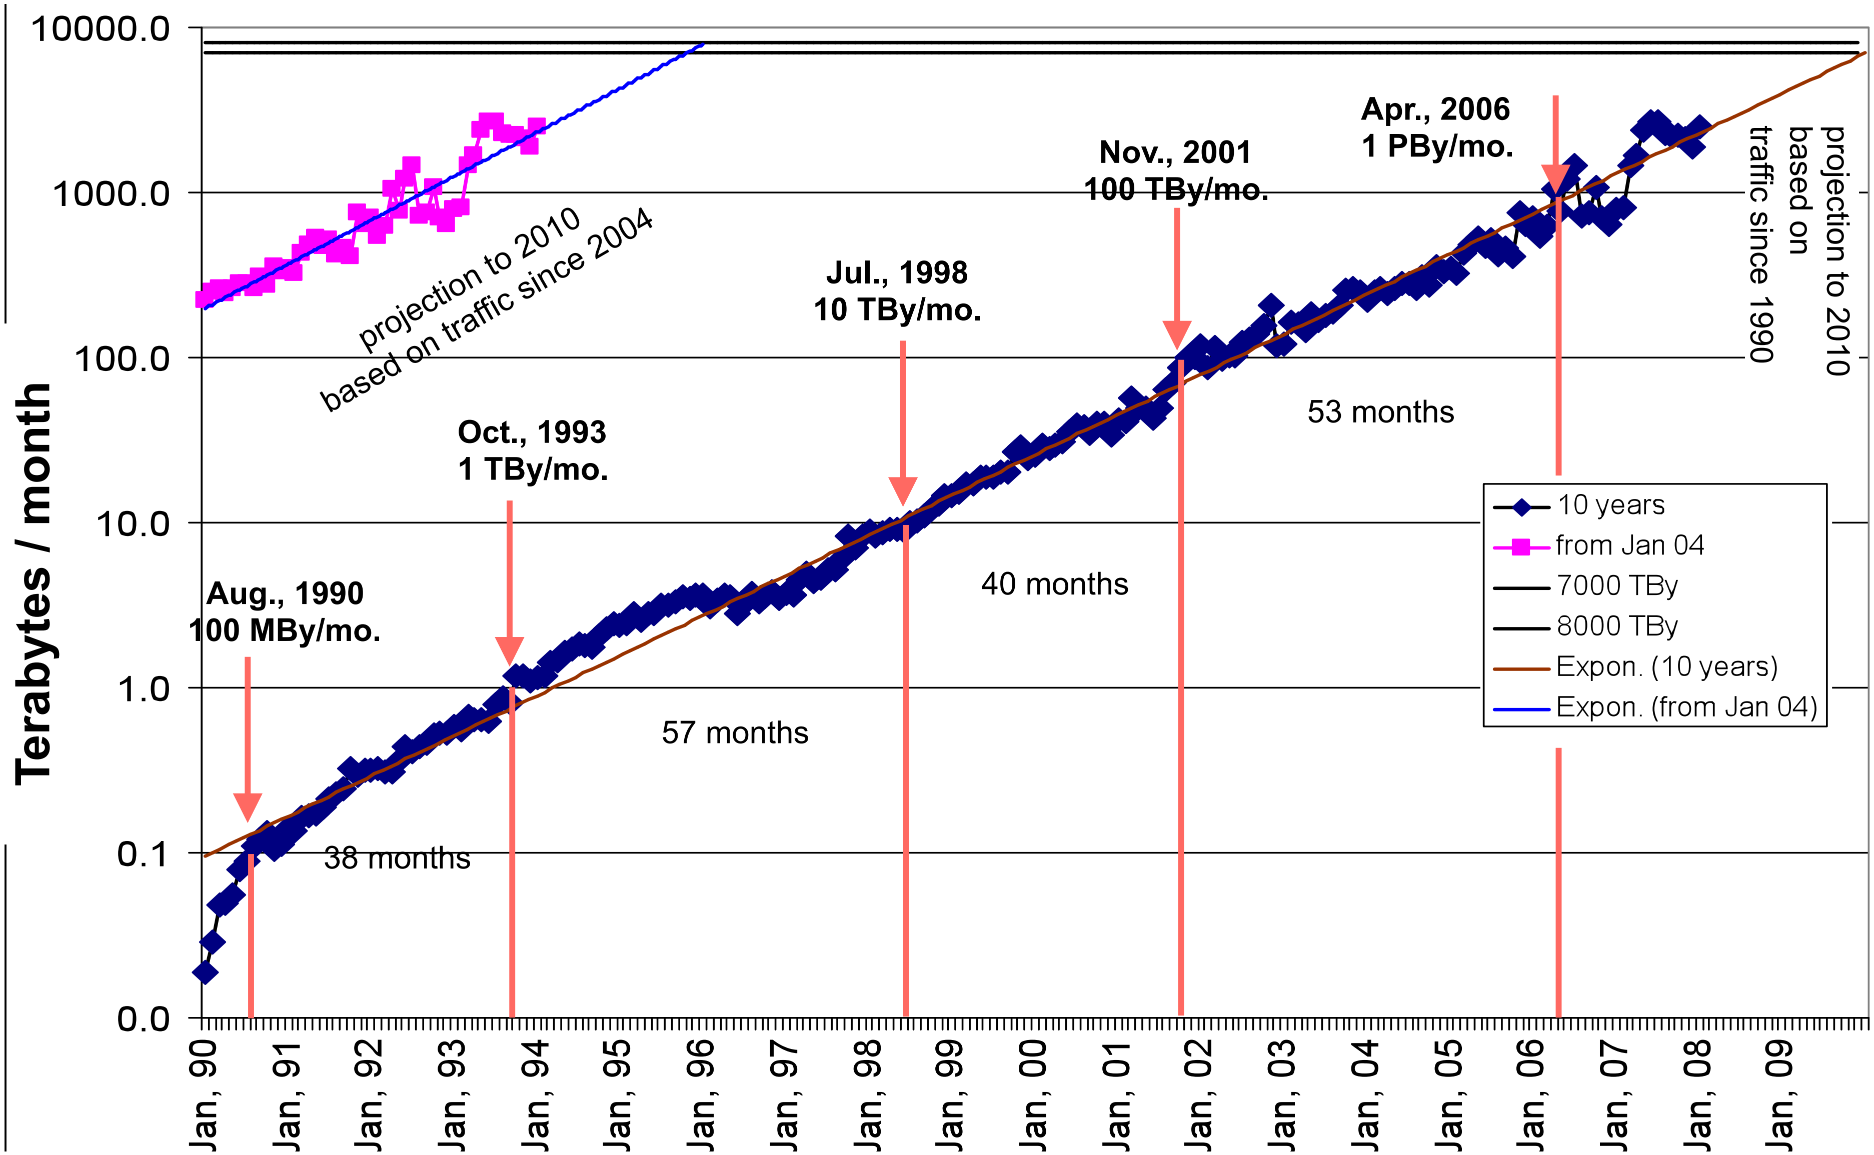

Log Plot of ESnet Monthly Accepted Traffic, January, 1990 – June, 2007

Large-Scale Flow Trends, June 2006
Subtitle: “Onslaught of the LHC”)

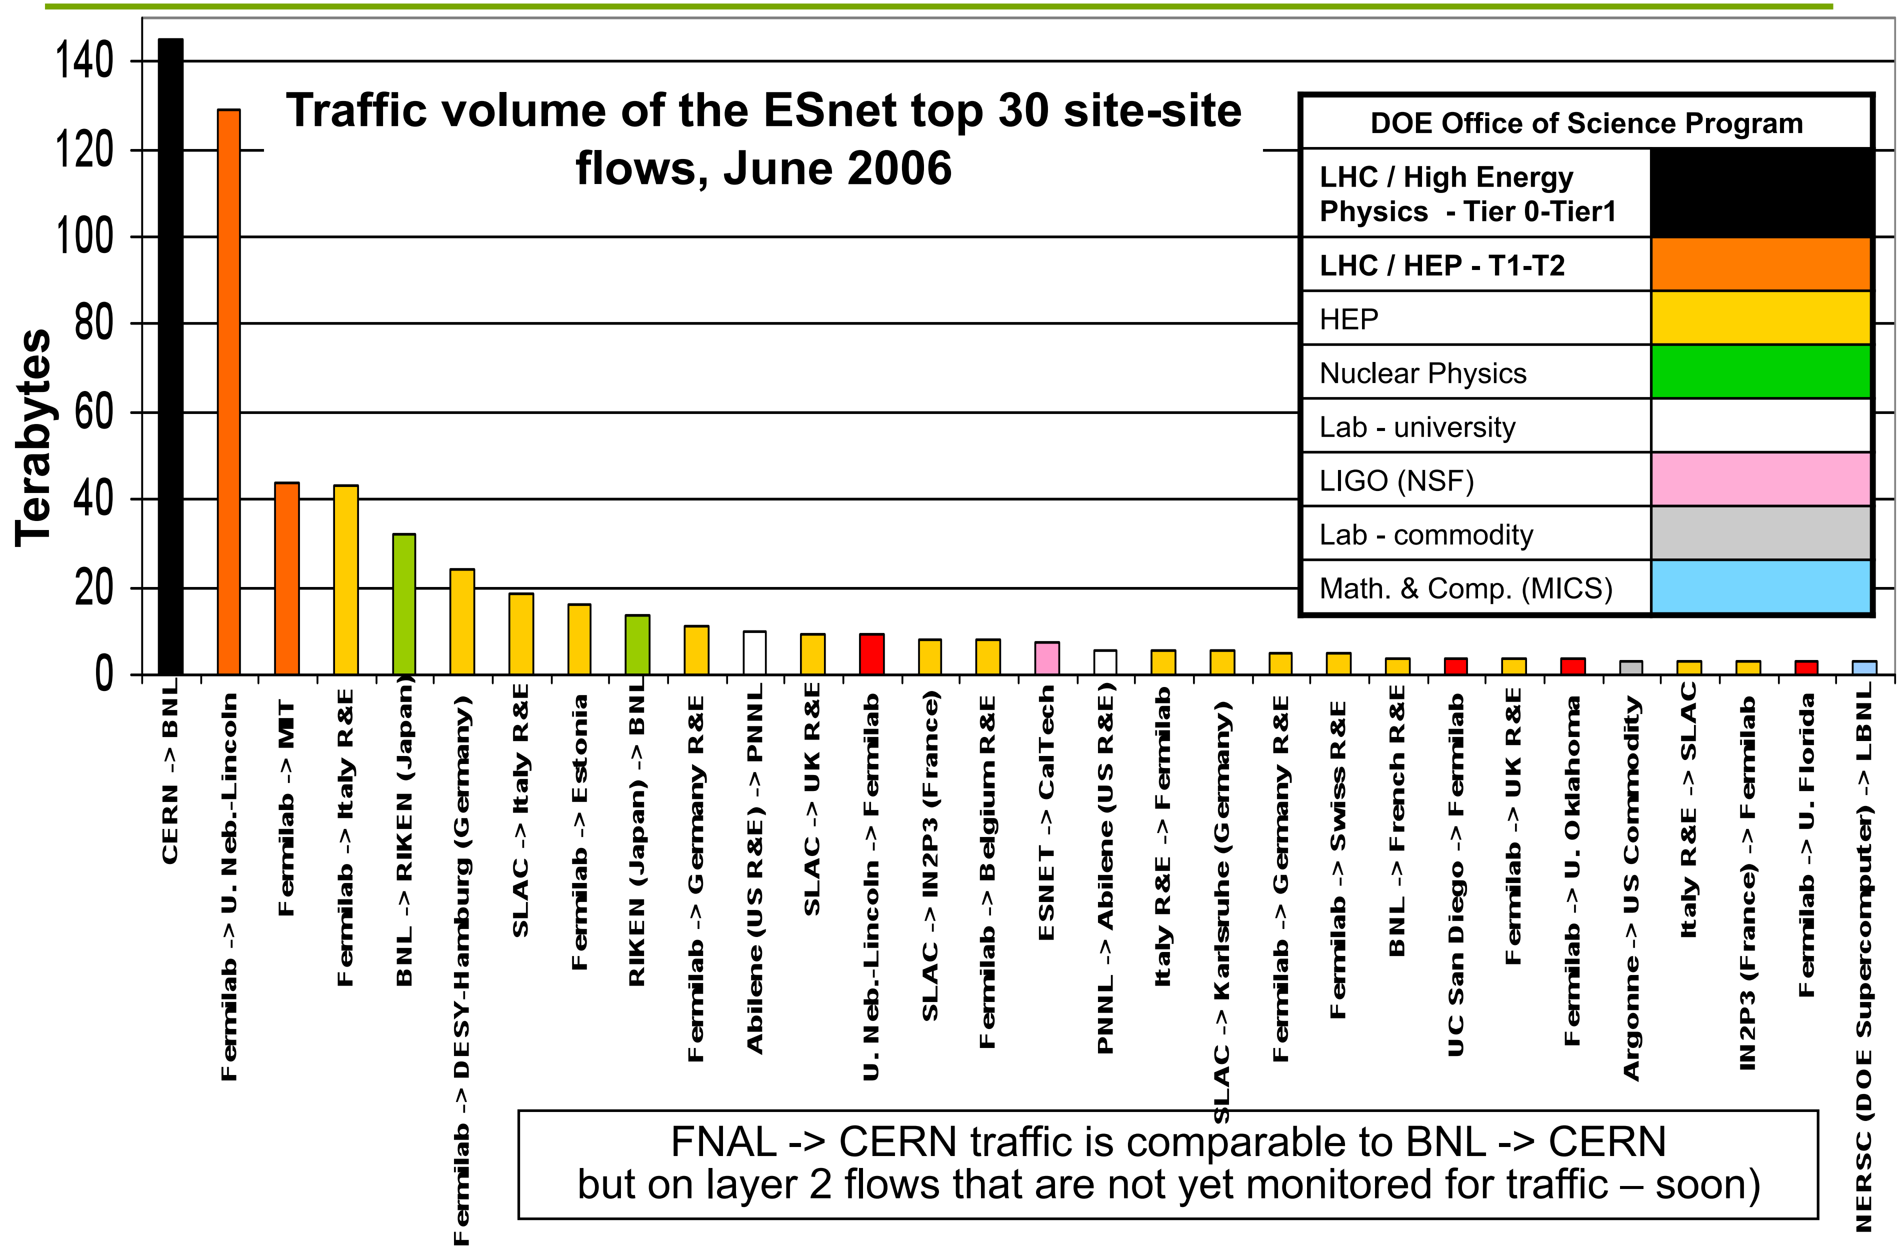

BNL Has Traffic Similar to FNAL

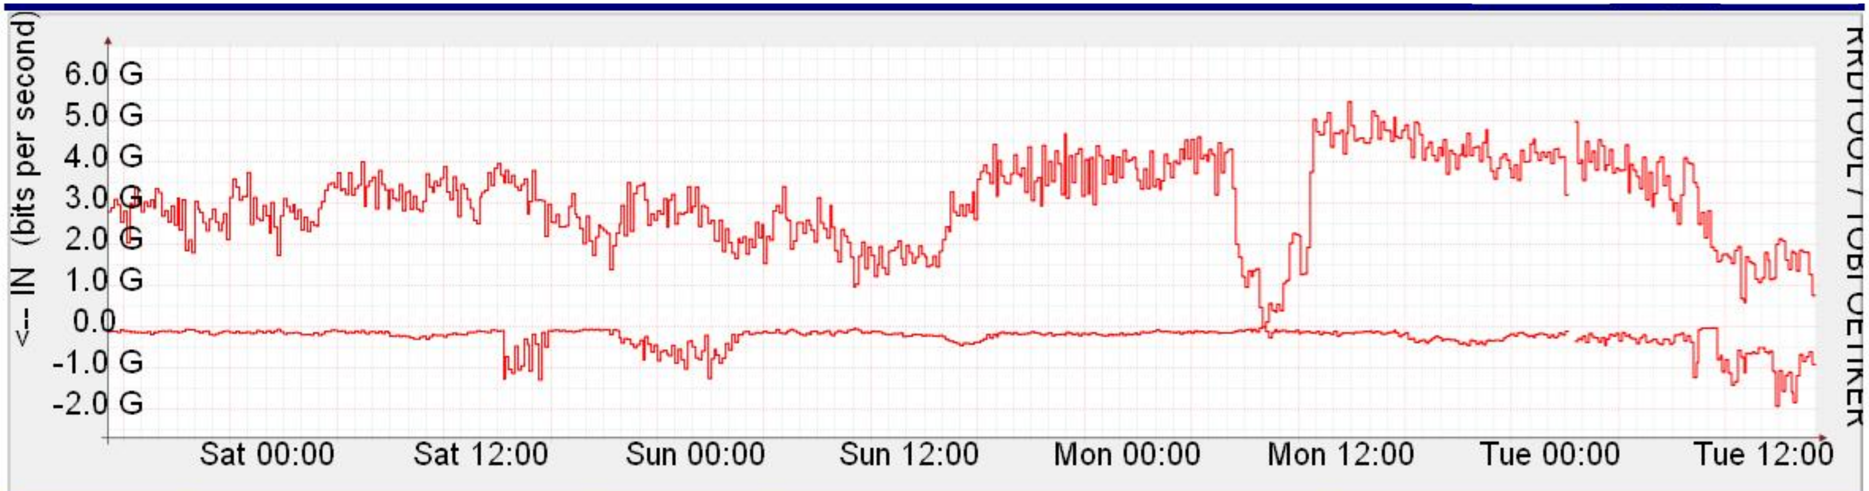

Color	Site(s)	Label	Average Rate (bps)		Volume (Bytes)	
			In	Out	In	Out
	BNL LHCOPN	src	261.573 M	3.069 G	10.224 T	119.966 T

BNL Has Traffic Similar to FNAL

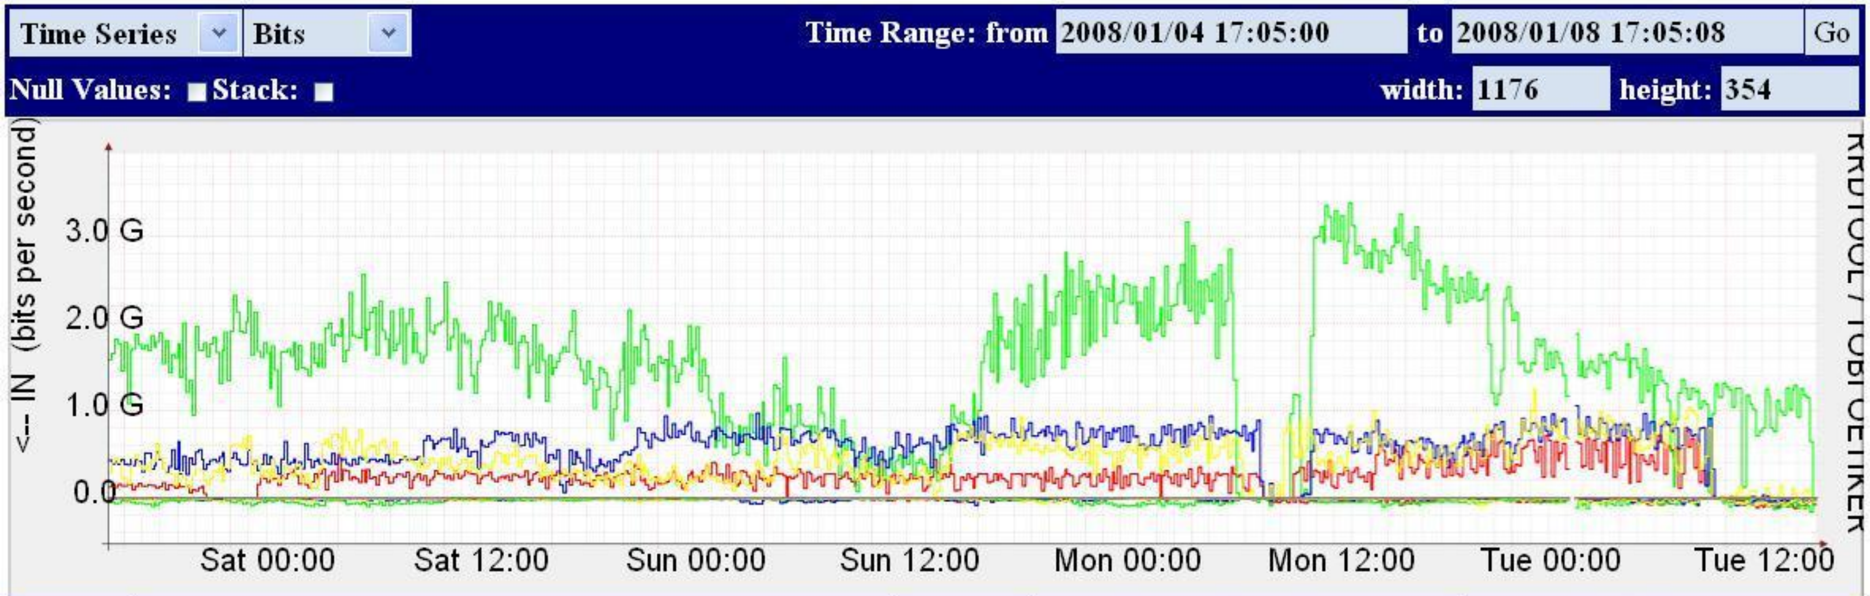

Color	Site(s)	Label	Average Rate (bps)		Volume (Bytes)	
			In	Out	In	Out
	/BNL LHCOPN - Indiana University	tcp	15.907 M	567.953 M	637.783 G	22.237 T
	/BNL LHCOPN - SLAC	tcp	14.302 M	246.064 M	573.398 G	9.634 T
	/BNL LHCOPN - University Chicago	tcp	15.447 M	436.613 M	619.339 G	17.095 T
	/BNL LHCOPN - University Michigan /Michigan States	tcp	37.854 M	1.562 G	1.482 T	61.158 T

Large-Scale Science Now Dominates ESnet

- ESnet carries about 2 x 10^9 flows (connections) per month which account for all of the ESnet traffic
- Large-scale science – LHC, RHIC, climate data, etc. now generate a few thousand flows/month that account for about 90% of all ESnet traffic
- The traffic on ESnet today is completely dominated by DOE's large-scale science

The Onslaught of Grids

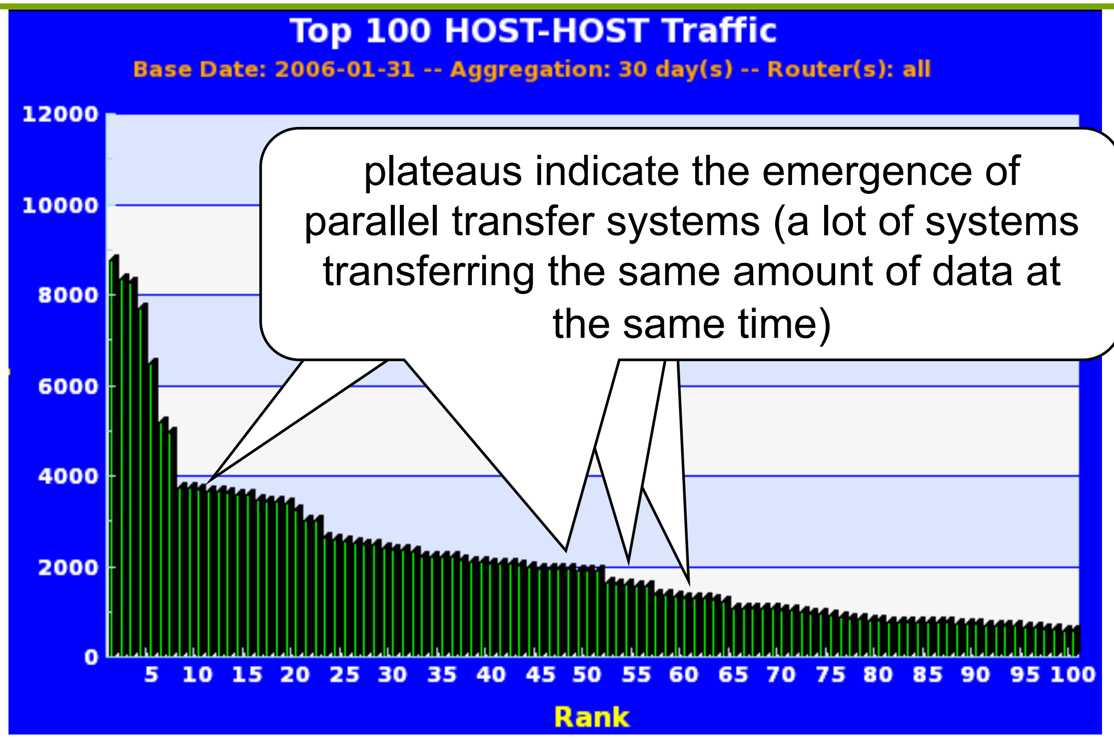

Most large data transfers are now done by parallel / Grid data movers

- In June, 2006 ***72%*** of the hosts generating the top 1000 flows were involved in parallel data movers (Grid applications)
- ***This is the most significant traffic pattern change in the history of ESnet***
- This has implications for the network architecture that favor path multiplicity and route diversity

Requirements from Network Utilization Observation

- In 4 years, we can expect a 10x increase in traffic over current levels *without the addition of production LHC traffic*
  - Nominal average load on busiest backbone links is ~1.5 Gbps today (June, 2006)
  - In 4 years that figure will be ~15 Gbps based on current trends
- Measurements of this type are science-agnostic
  - It doesn’t matter who the users are, the traffic load is increasing exponentially
  - Predictions based on this sort of forward projection tend to be conservative estimates of future requirements because they cannot predict new uses

Requirements from Traffic Flow Observations

- Observation: Most of ESnet science traffic has a source or sink outside of ESnet
  - Requirements:
  - High-bandwidth peering with R&E community
  - Reliability and bandwidth requirements demand that peering be redundant
  - Multiple 10 Gbps peerings today, must be able to add more bandwidth flexibly and cost-effectively
  - Bandwidth and service guarantees must traverse R&E peerings
    - Collaboration with other R&E networks on a common framework is critical
    - Seamless fabric
- Observation: Large-scale science is now the dominant user of the network - traffic patterns are different than commodity Internet
  - Requirements:
  - A network that is purpose-built, with a scalable architecture to satisfy the demands of large-scale science traffic into the future

Summary of All Requirements To-Date

Requirements from SC Programs:

- Provide "consulting" on system / application network tuning

Requirements from science case studies:

- Build the ESnet core up to 100 Gb/s within 5 years
- Deploy network to accommodate LHC collaborator footprint
- Implement network to provide for LHC data path loadings
- Provide the network as a service-oriented capability

Requirements from observing traffic growth and change trends in the network:

- Provide 15 Gb/s core by 2008 and 150 Gb/s core by 2010
- Provide a rich diversity and high bandwidth for R&E peerings
- Economically accommodate a very large volume of circuit-like traffic

Where Performance Problems Arise and Who Has to Address Them

- “The network” is really a collection of autonomous domains that have very specific boundaries of operation and responsibility
- Almost all end-to-end (system-to-system) performance problems involve many autonomous domains and therefore have to be solved cooperatively

Architecture of the Internet

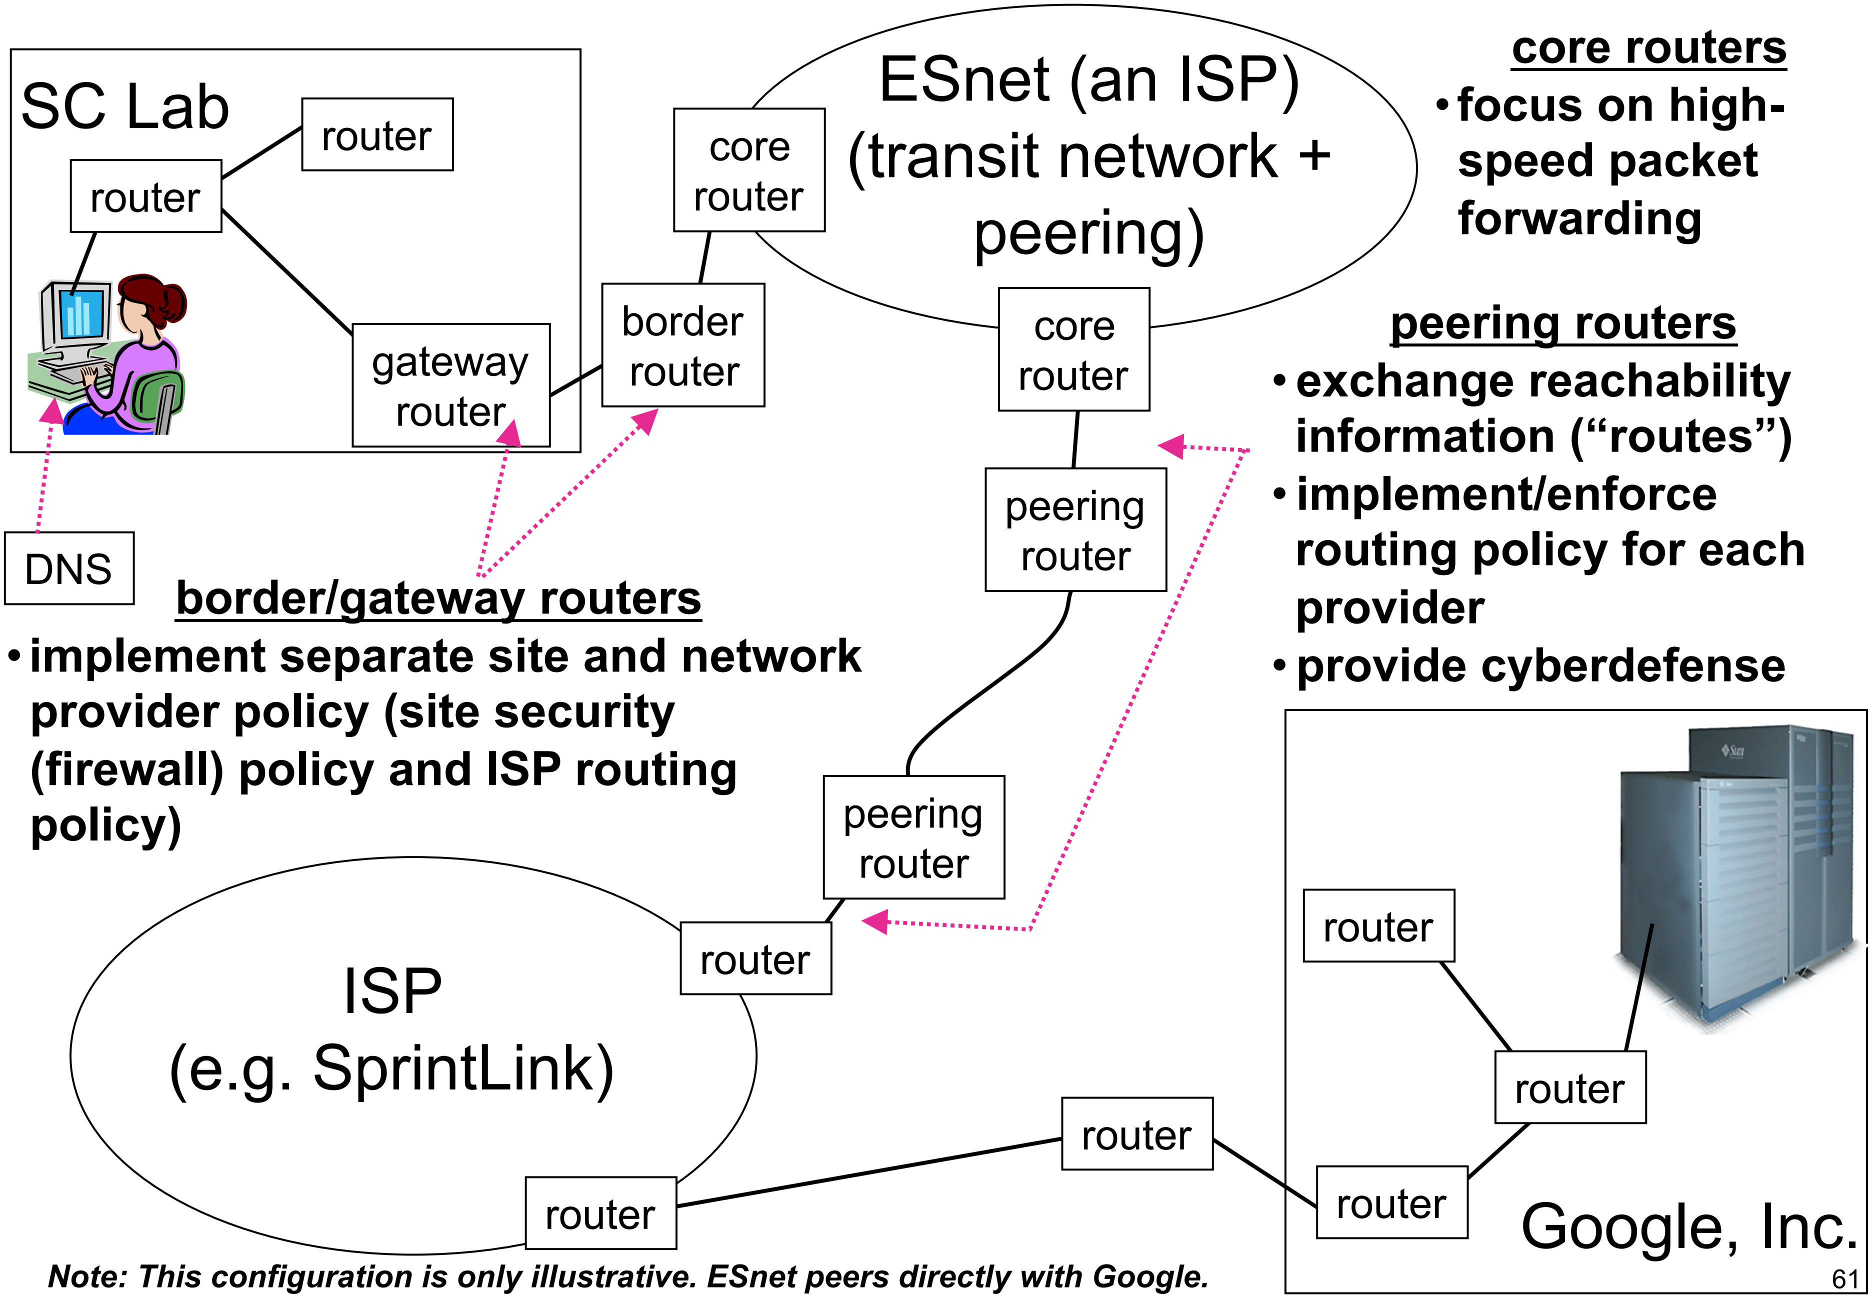

core routers

- focus on high-speed packet forwarding

peering routers

- exchange reachability information ("routes")
- implement/enforce routing policy for each provider
- provide cyberdefense

border/gateway routers

- implement separate site and network provider policy (site security (firewall) policy and ISP routing policy)

Note: This configuration is only illustrative. ESnet peers directly with Google.

Architecture of the R&E Environment

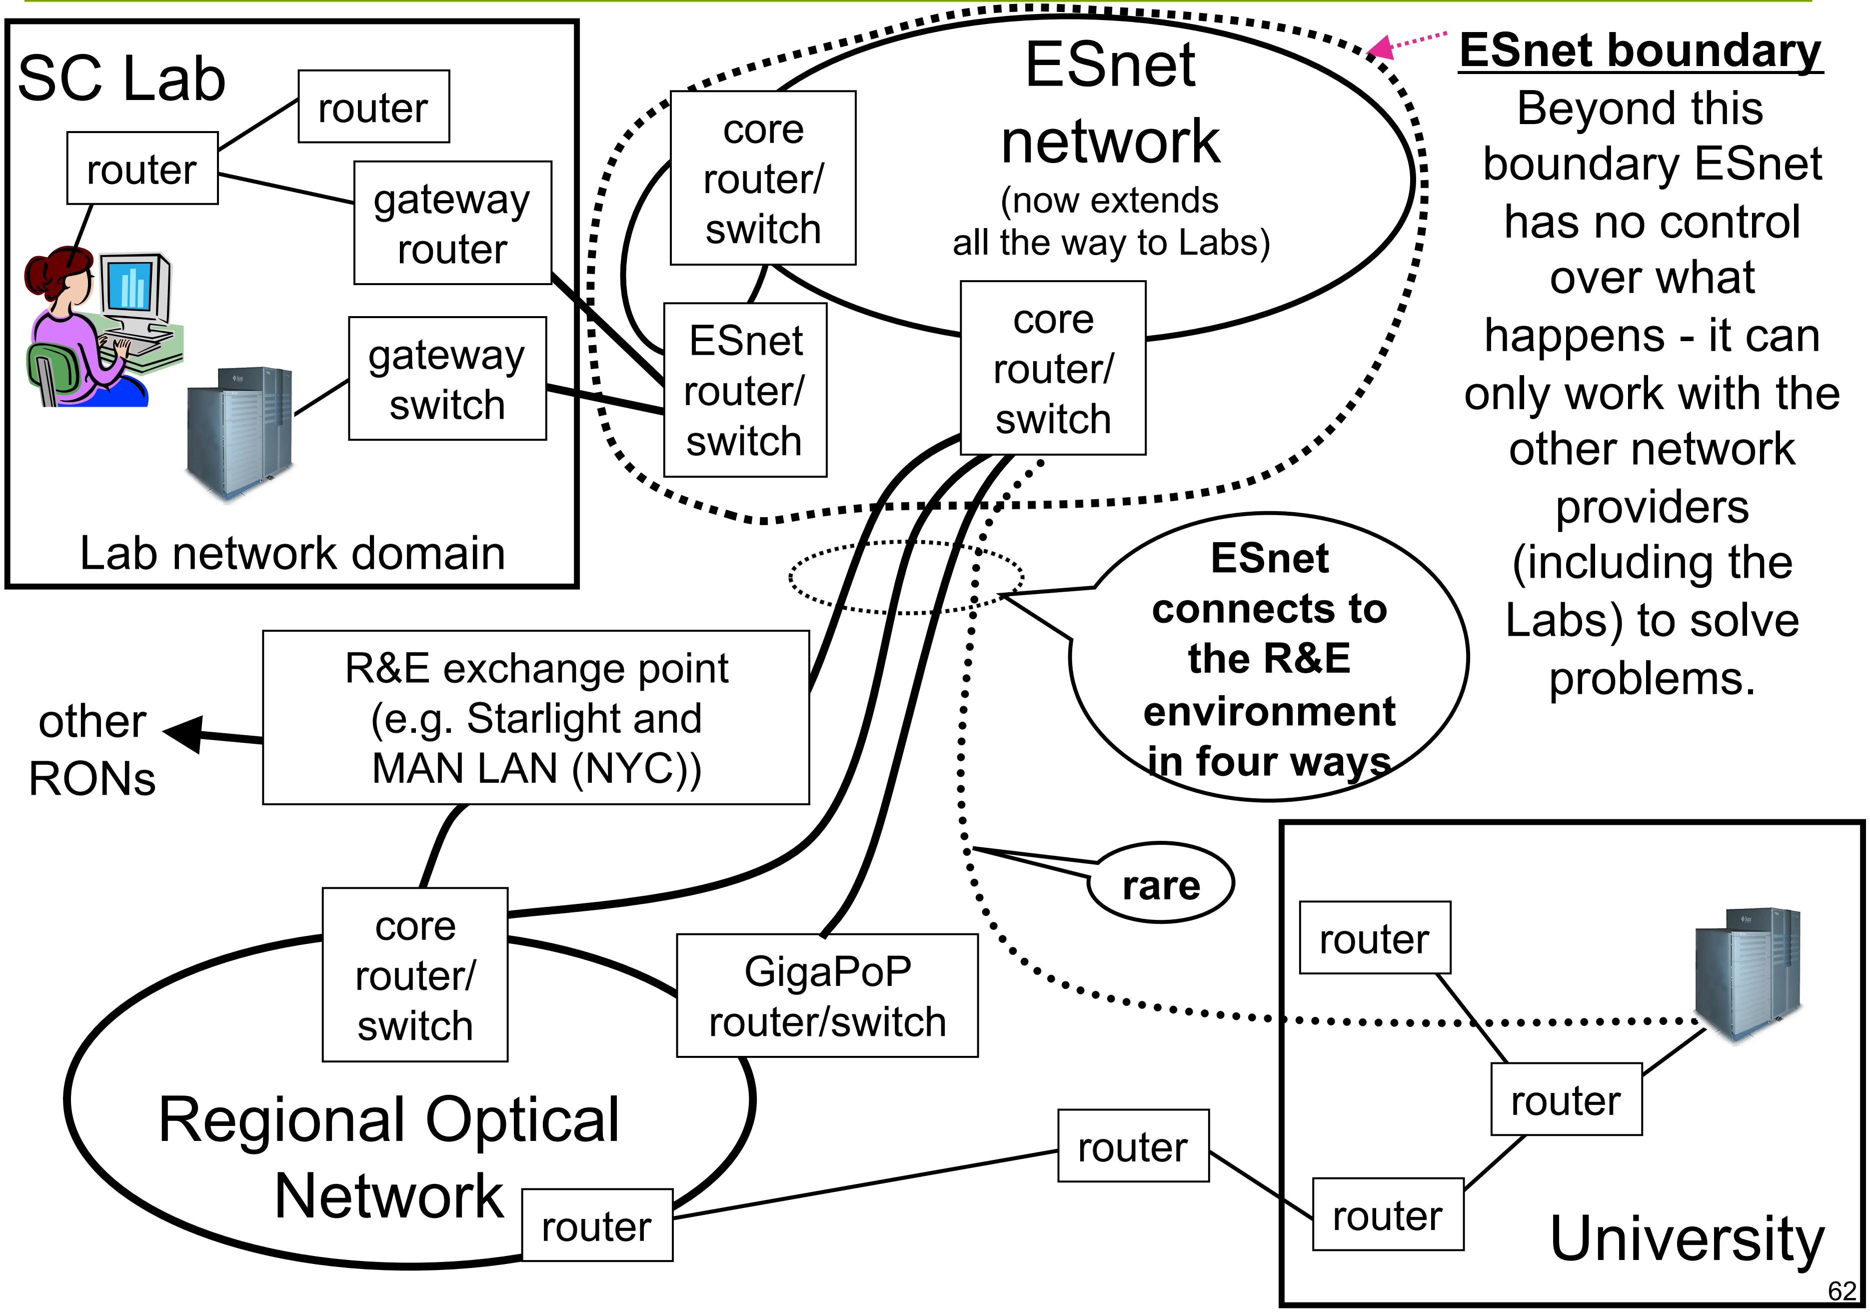

Relationship of ESnet to Sites/Labs and Facilities

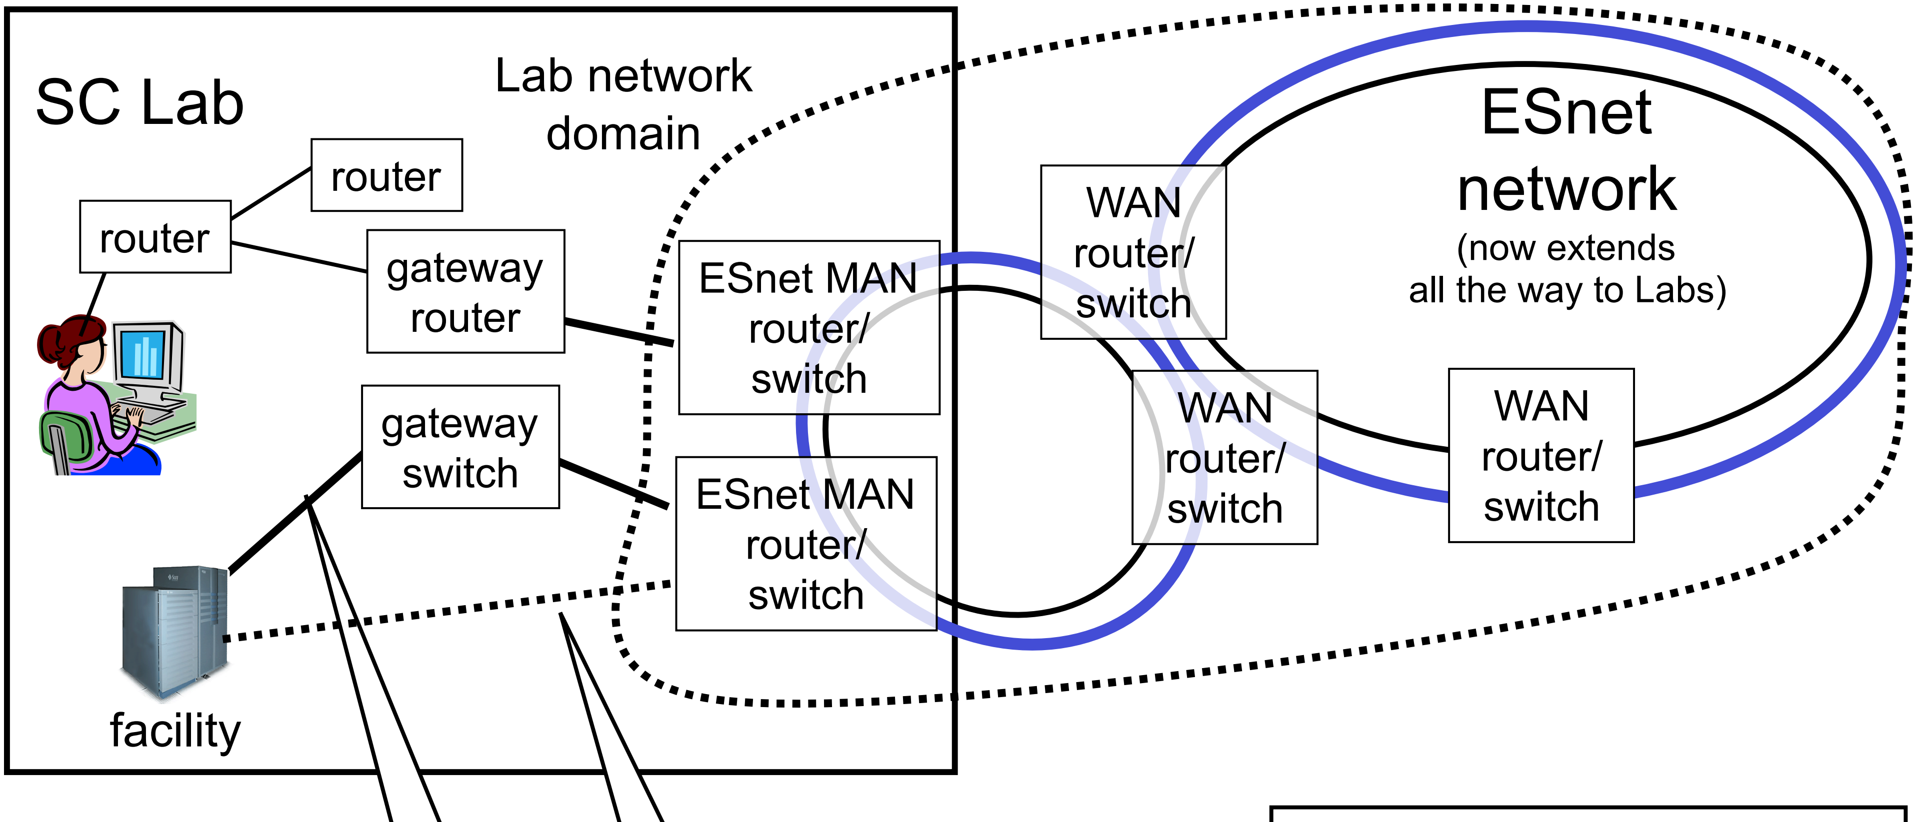

- Connection from a facility to ESnet can be either through a site switch/router or directly to the ESnet switch/router at the site.
- In the case of a direct connection, the site must provide the physical path (fiber) from the ESnet equipment to the facility equipment - and the Site Coordinator (Lab primary contact for ESnet) must agree to the connection.

ESnet boundary

Beyond this boundary ESnet has no control over what happens - it can only work with the other network providers (including the Labs) to solve problems.

Most Common Network Performance Bottlenecks

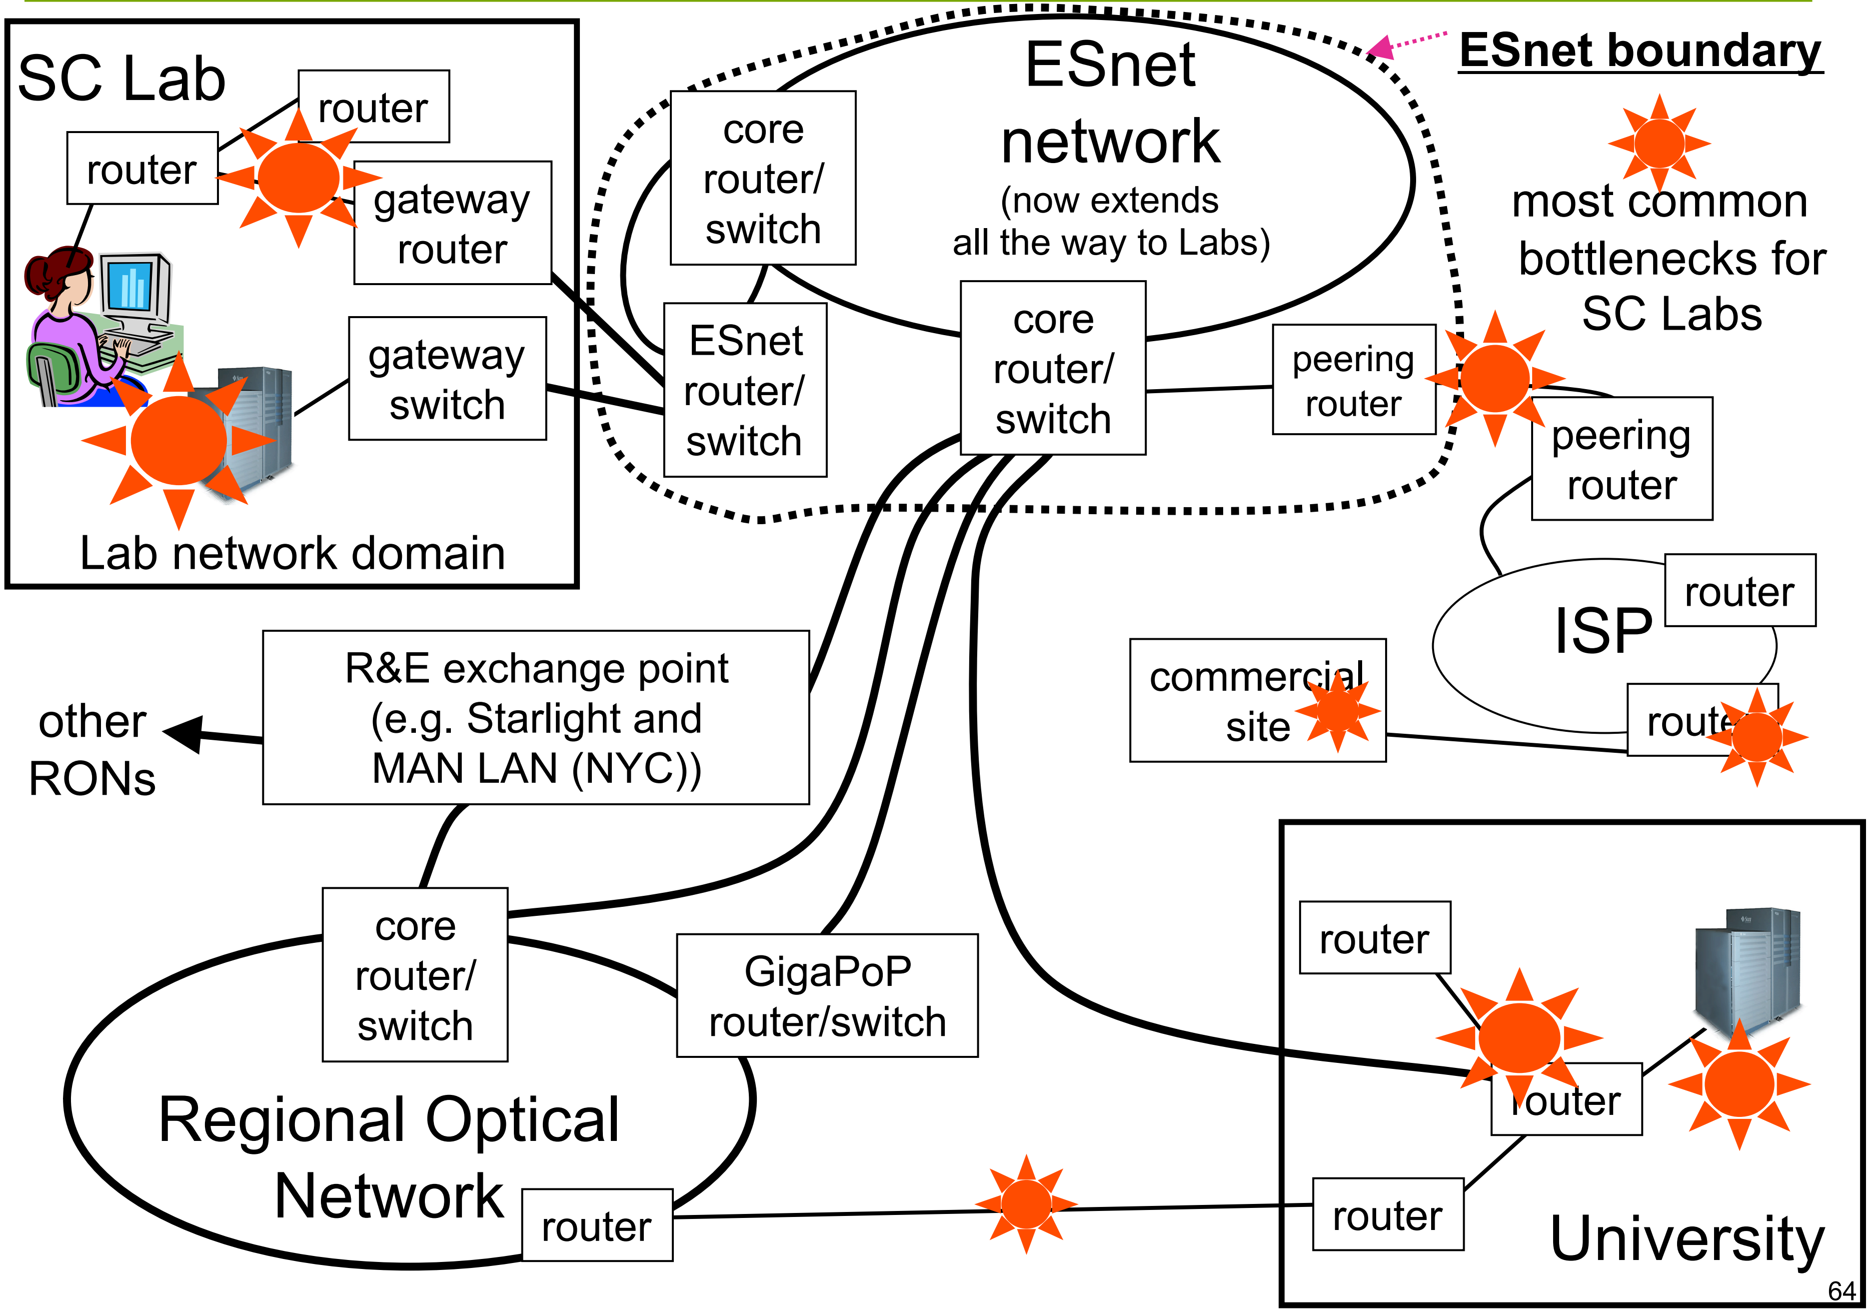

➢ ESnet4 - The Response to the Requirements

I) A new network architecture and implementation strategy

- Provide two networks: IP and circuit-oriented Science Data Netework
  - Reduces cost of handling high bandwidth data flows
    - Highly capable routers are not necessary when every packet goes to the same place
    - Use lower cost (factor of 5x) switches to relatively route the packets
- Rich and diverse network topology for flexible management and high reliability
- Dual connectivity at every level for all large-scale science sources and sinks
- A partnership with the US research and education community to build a shared, large-scale, R&E managed optical infrastructure
  - a scalable approach to adding bandwidth to the network
  - dynamic allocation and management of optical circuits

II) Development and deployment of a virtual circuit service

- Develop the service cooperatively with the networks that are intermediate between DOE Labs and major collaborators to ensure and-to-end interoperability

III) Develop and deploy service-oriented, user accessable network monitoring systems

IV) Provide “consulting” on system / application network performance tuning

ESnet4

- Internet2 has partnered with Level 3 Communications Co. and Infinera Corp. for a dedicated optical fiber infrastructure with a national footprint and a rich topology - the "Internet2 Network"
  - The fiber will be provisioned with Infinera Dense Wave Division Multiplexing equipment that uses an advanced, integrated optical-electrical design
  - Level 3 will maintain the fiber and the DWDM equipment
  - The DWDM equipment will initially be provisioned to provide10 optical circuits (lambdas - λs) across the entire fiber footprint (40/80 λs is max.)
- ESnet has partnered with Internet2 to:
  - Share the optical infrastructure
  - Develop new circuit-oriented network services
  - Explore mechanisms that could be used for the ESnet Network Operations Center (NOC) and the Internet2/Indiana University NOC to back each other up for disaster recovery purposes

ESnet4

- ESnet has built its next generation IP network and its new circuit-oriented Science Data Network primarily on the Internet2 optical circuits that are dedicated to ESnet, together with a few National Lambda Rail and other circuits
  - ESnet will provision and operate its own routing and switching hardware that is installed in various commercial telecom hubs around the country, as it has done for the past 20 years
  - ESnet's peering relationships with the commercial Internet, various US research and education networks, and numerous international networks will continue and evolve as they have for the past 20 years

Typical Internet2 and ESnet Optical Node

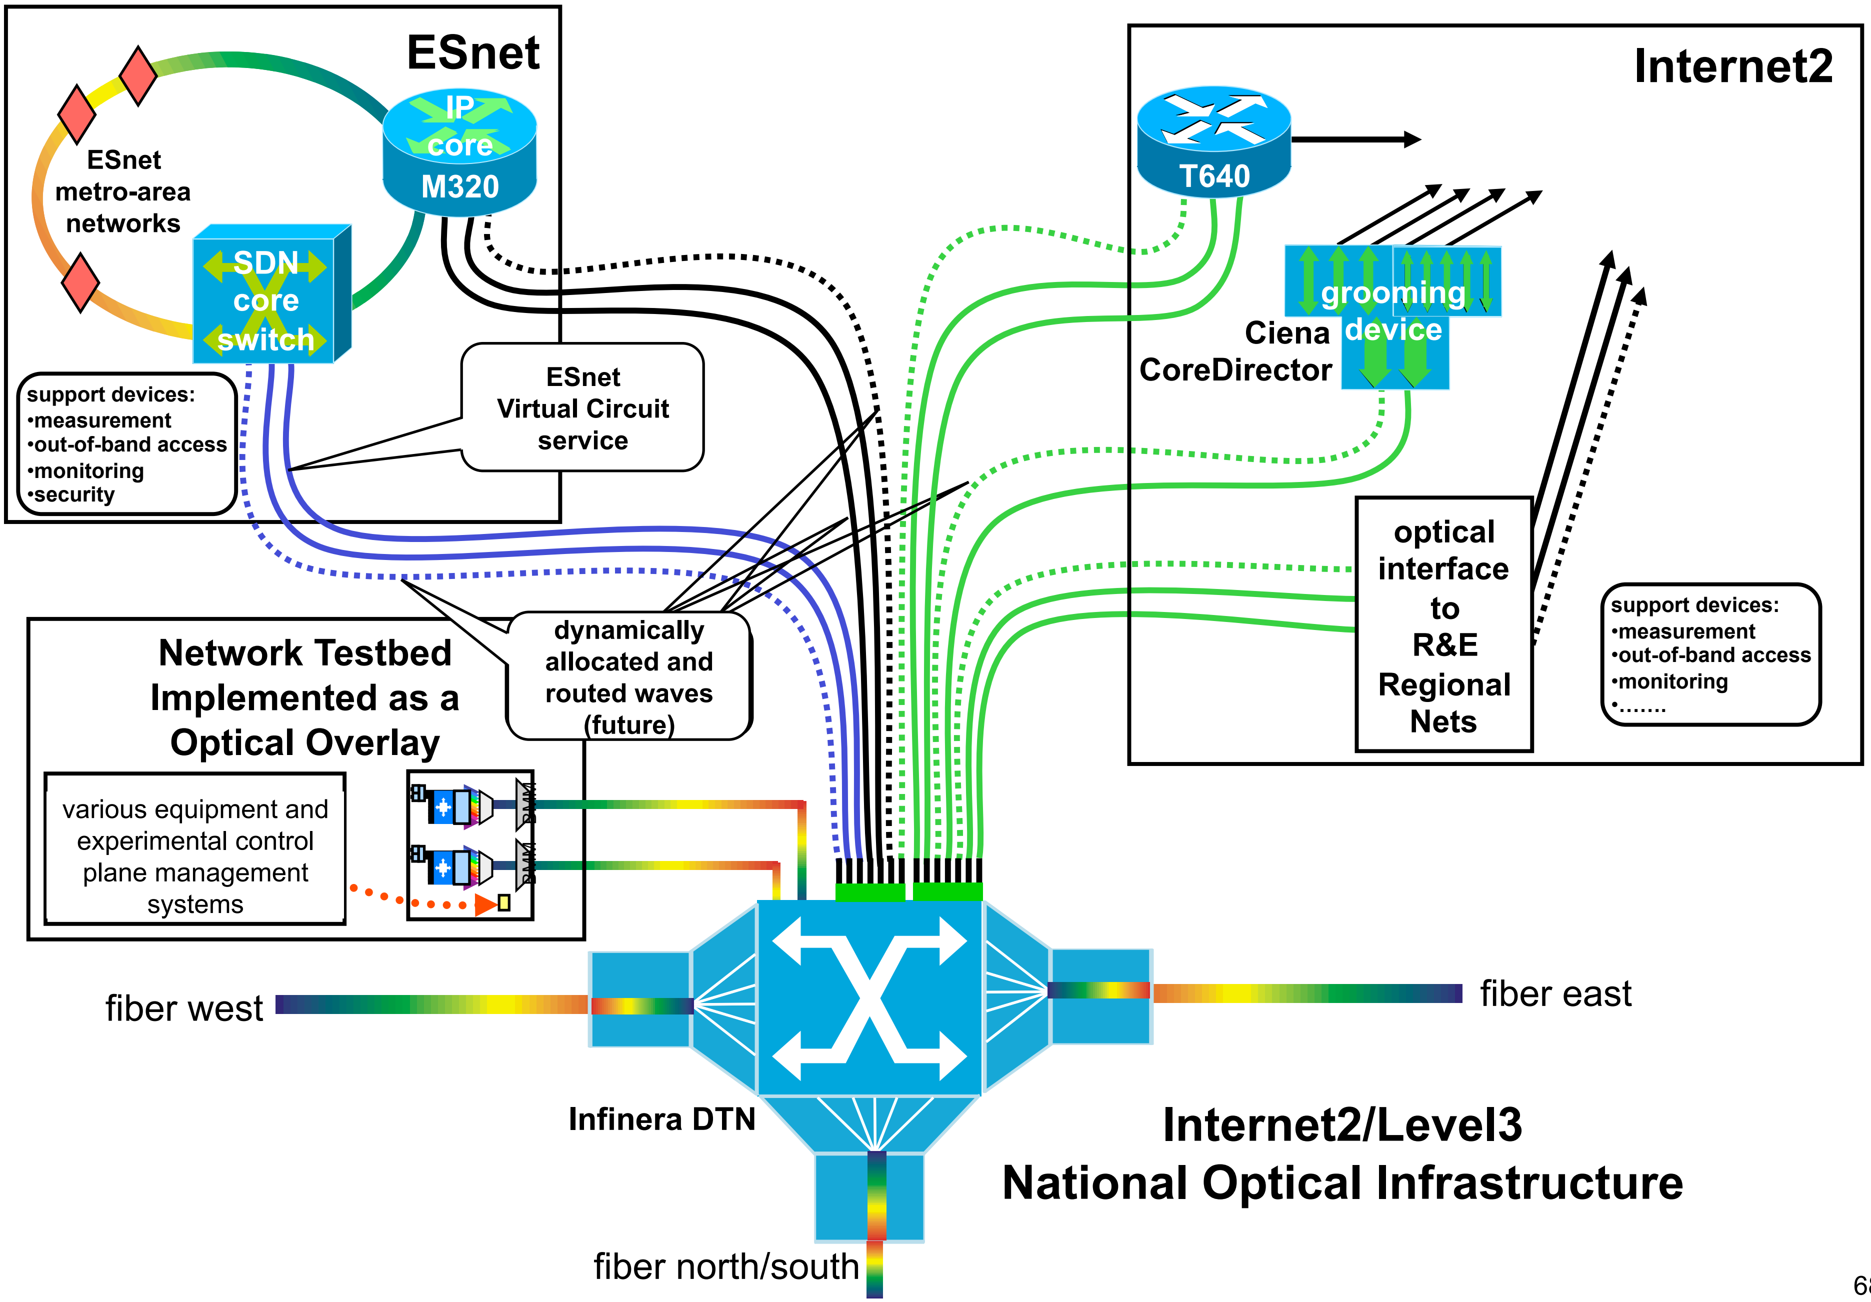

Conceptual Architecture of the Infinera DTN Digital Terminal

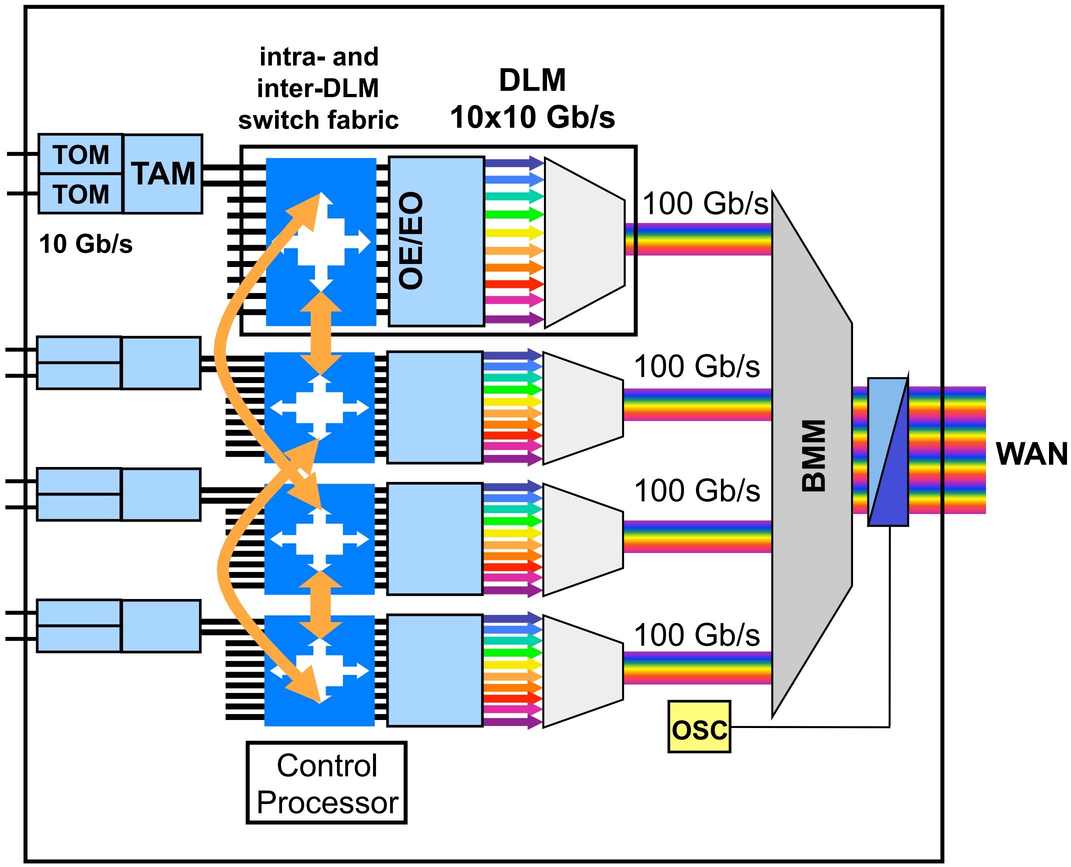

Notes:
1) This conceptual architecture is based on W. Johnston's understanding of the Infinera DTN
2) All channels are full duplex, which is not made explicit here
3) The switch fabric operates within a DLM and also interconnects the DLMs
4) The Management Control Module is not shown

- Band Mux Module (BMM)
  - Multiplexes 100Gb/s bands onto 400 Gb/s or 800 Gb/s line
- Digital Line Module (DLM)
  - 100Gb/s DWDM line module (10 λ x 10Gb/s)
  - Integrated digital switch enables add/drop & grooming at 2.5Gb/s (ODU-1) granularity
- Tributary Adapter Module (TAM)
  - Provides customer facing optical interface
  - Supports variety of data rates and service types
  - Up to 5 TAMs per DLM
- Tributary Optical Module (TOM)
  - Pluggable client side optical module (XFP, SFP)
  - Supports 10GE, OC48&OC192 SONET
- Optical Supervisory Channel (OSC)
  - Management Plane Traffic—Includes traffic from the remote management systems to access network elements for the purpose of managing them
  - Control Plane Traffic—GMPLS routing and signaling control protocol traffic
  - Datawire Traffic—Customer management traffic by interconnecting customer's 10Mbps Ethernet LAN segments at various sites through AUX port interfaces

ESnet4

- ESnet4 will also involve an expansion of the multi-10Gb/s Metropolitan Area Rings in the San Francisco Bay Area, Chicago, Long Island, Newport News (VA/Washington, DC area), and Atlanta
  - provide multiple, independent connections for ESnet sites to the ESnet core network
  - expandable
- Several 10Gb/s links provided by the Labs that will be used to establish multiple, independent connections to the ESnet core
  - currently PNNL and ORNL

ESnet Metropolitan Area Network Ring Architecture for High Reliability Sites

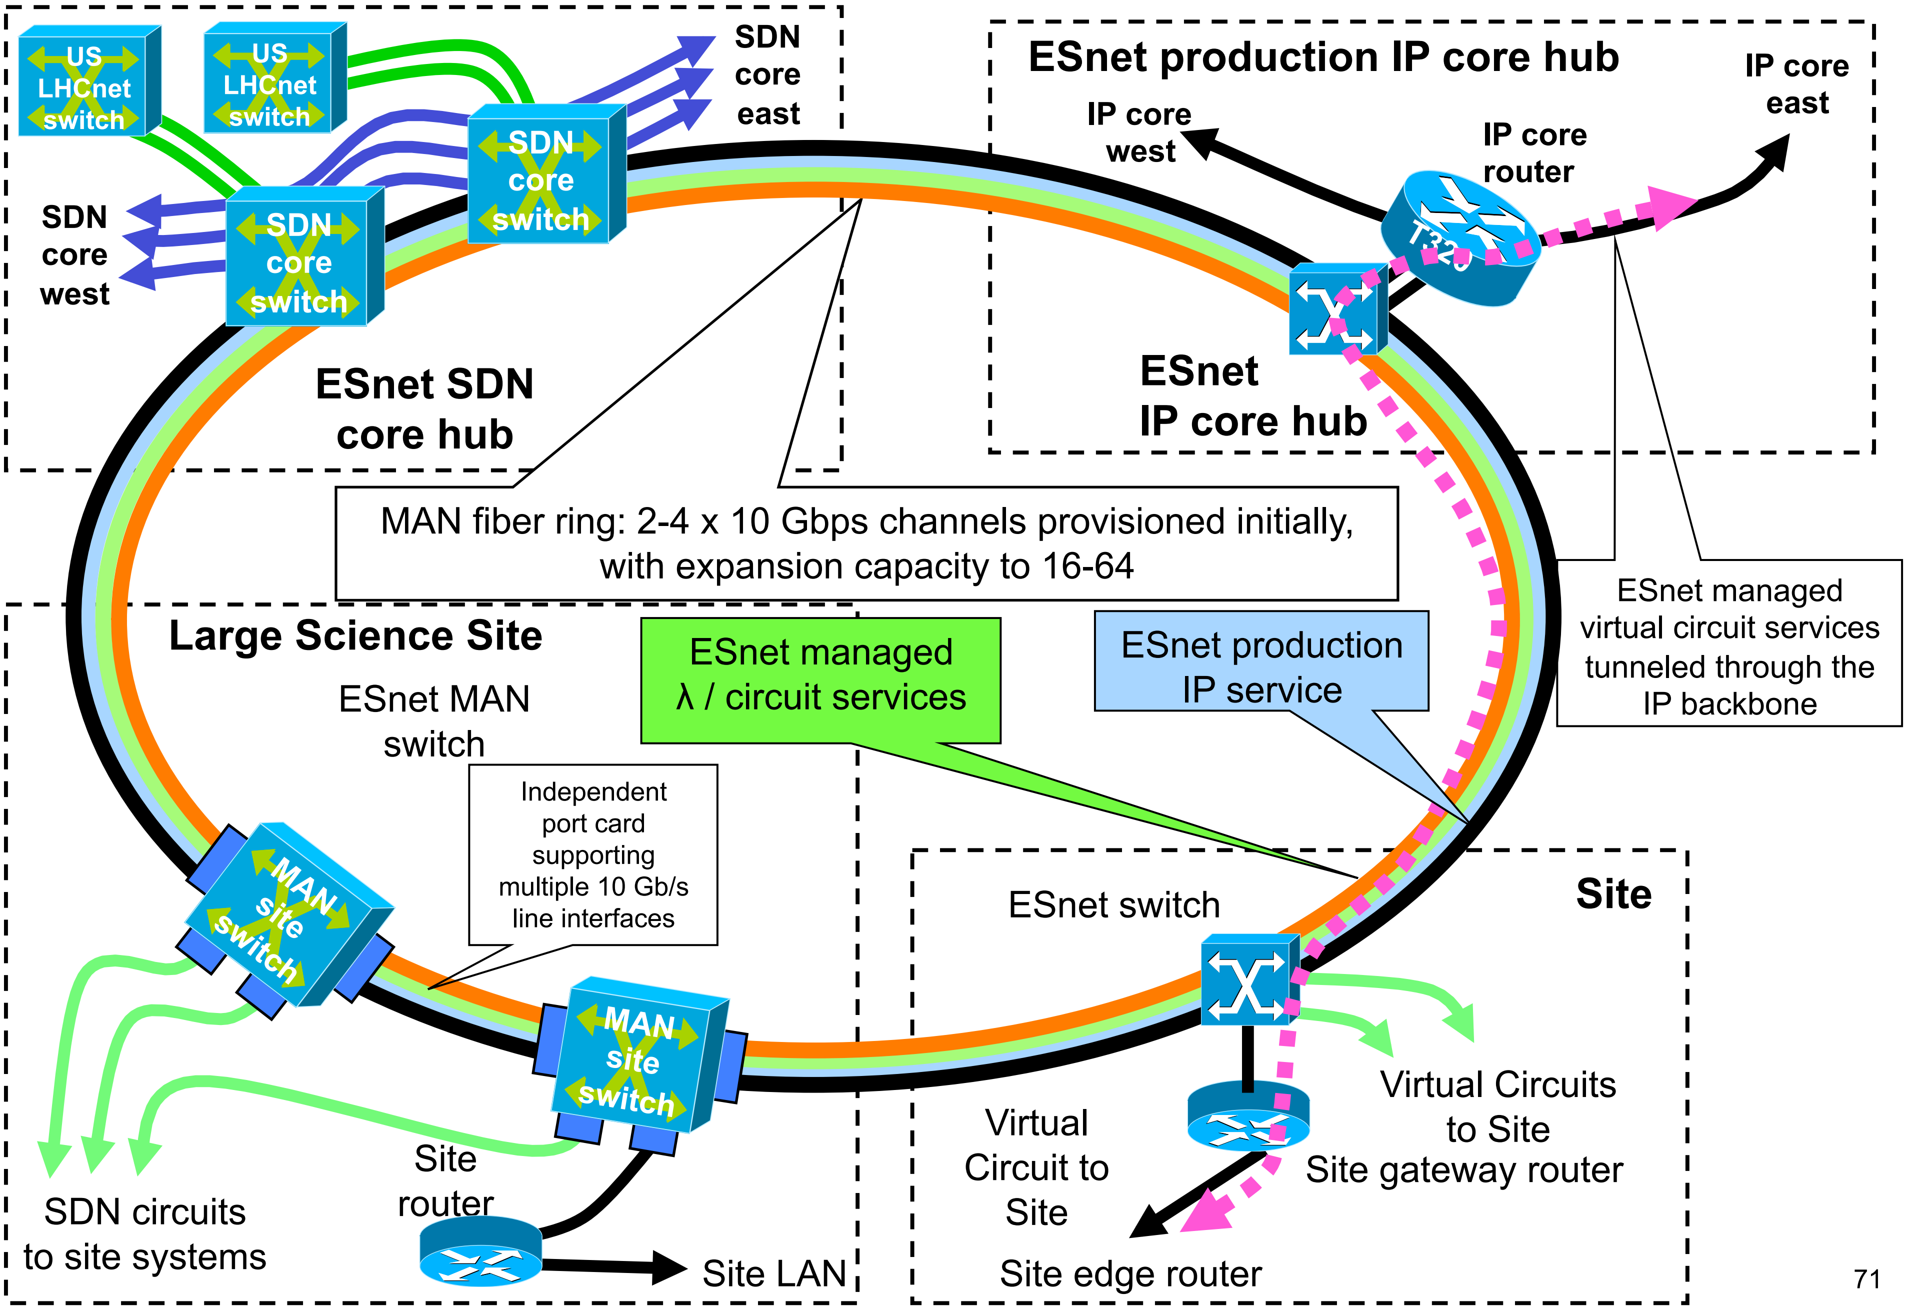

The Evolution of ESnet Architecture

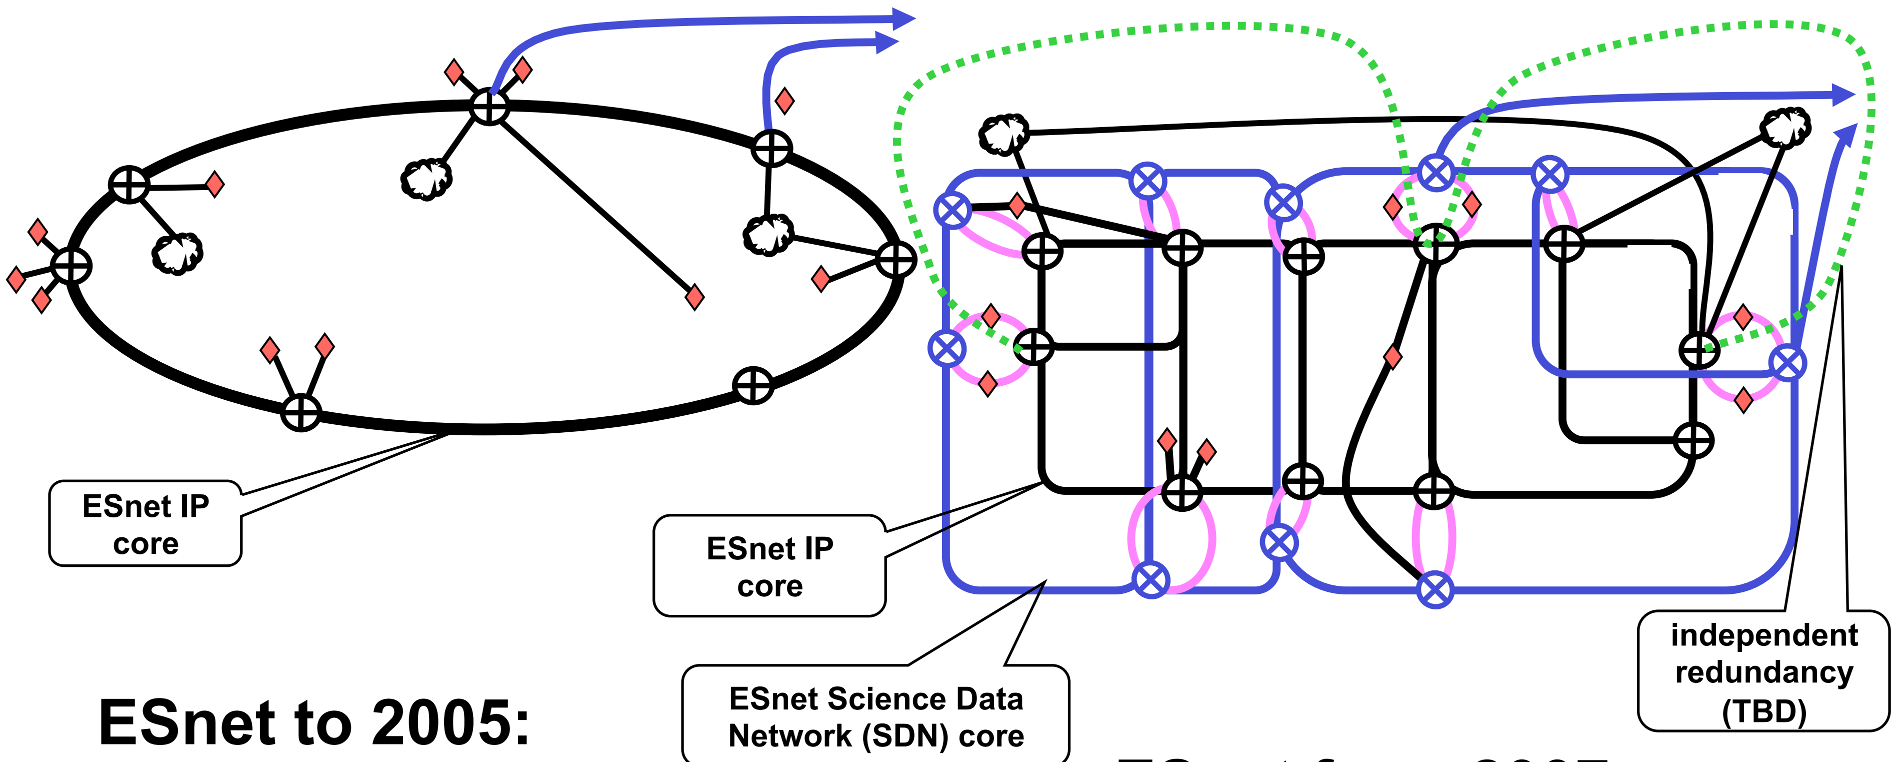

ESnet to 2005:

- A routed IP network with sites singly attached to a national core ring

ESnet sites

ESnet hubs / core network connection points

Metro area rings (MANs)

Other IP networks

Circuit connections to other science networks (e.g. USLHCNet)

ESnet from 2007:

- A routed IP network with sites dually connected on metro area rings or dually connected directly to core ring
- A switched network providing virtual circuit services for data-intensive science
- Rich topology offsets the lack of dual, independent national cores

ESnet 3 Backbone as of January 1, 2007

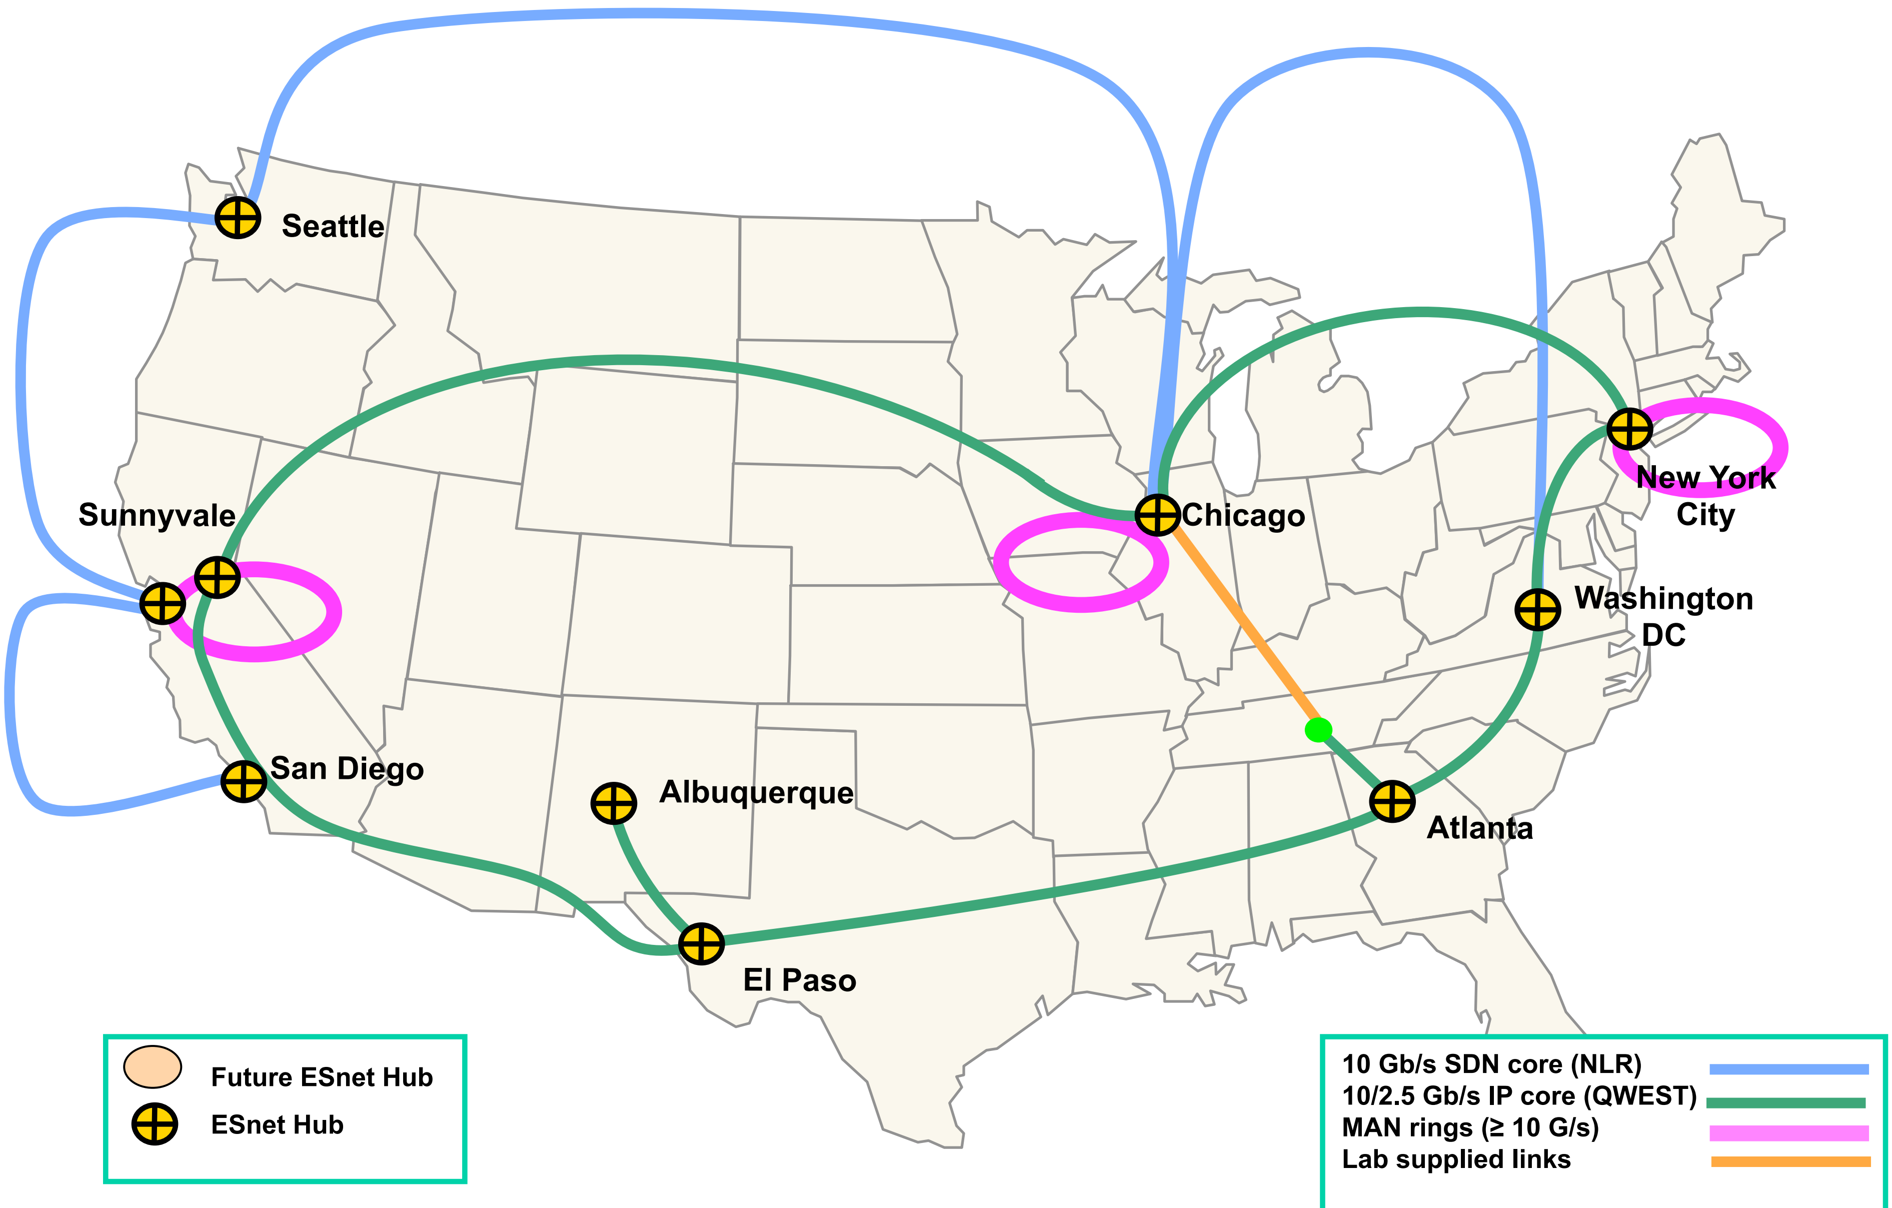

ESnet 4 Backbone as of April 15, 2007

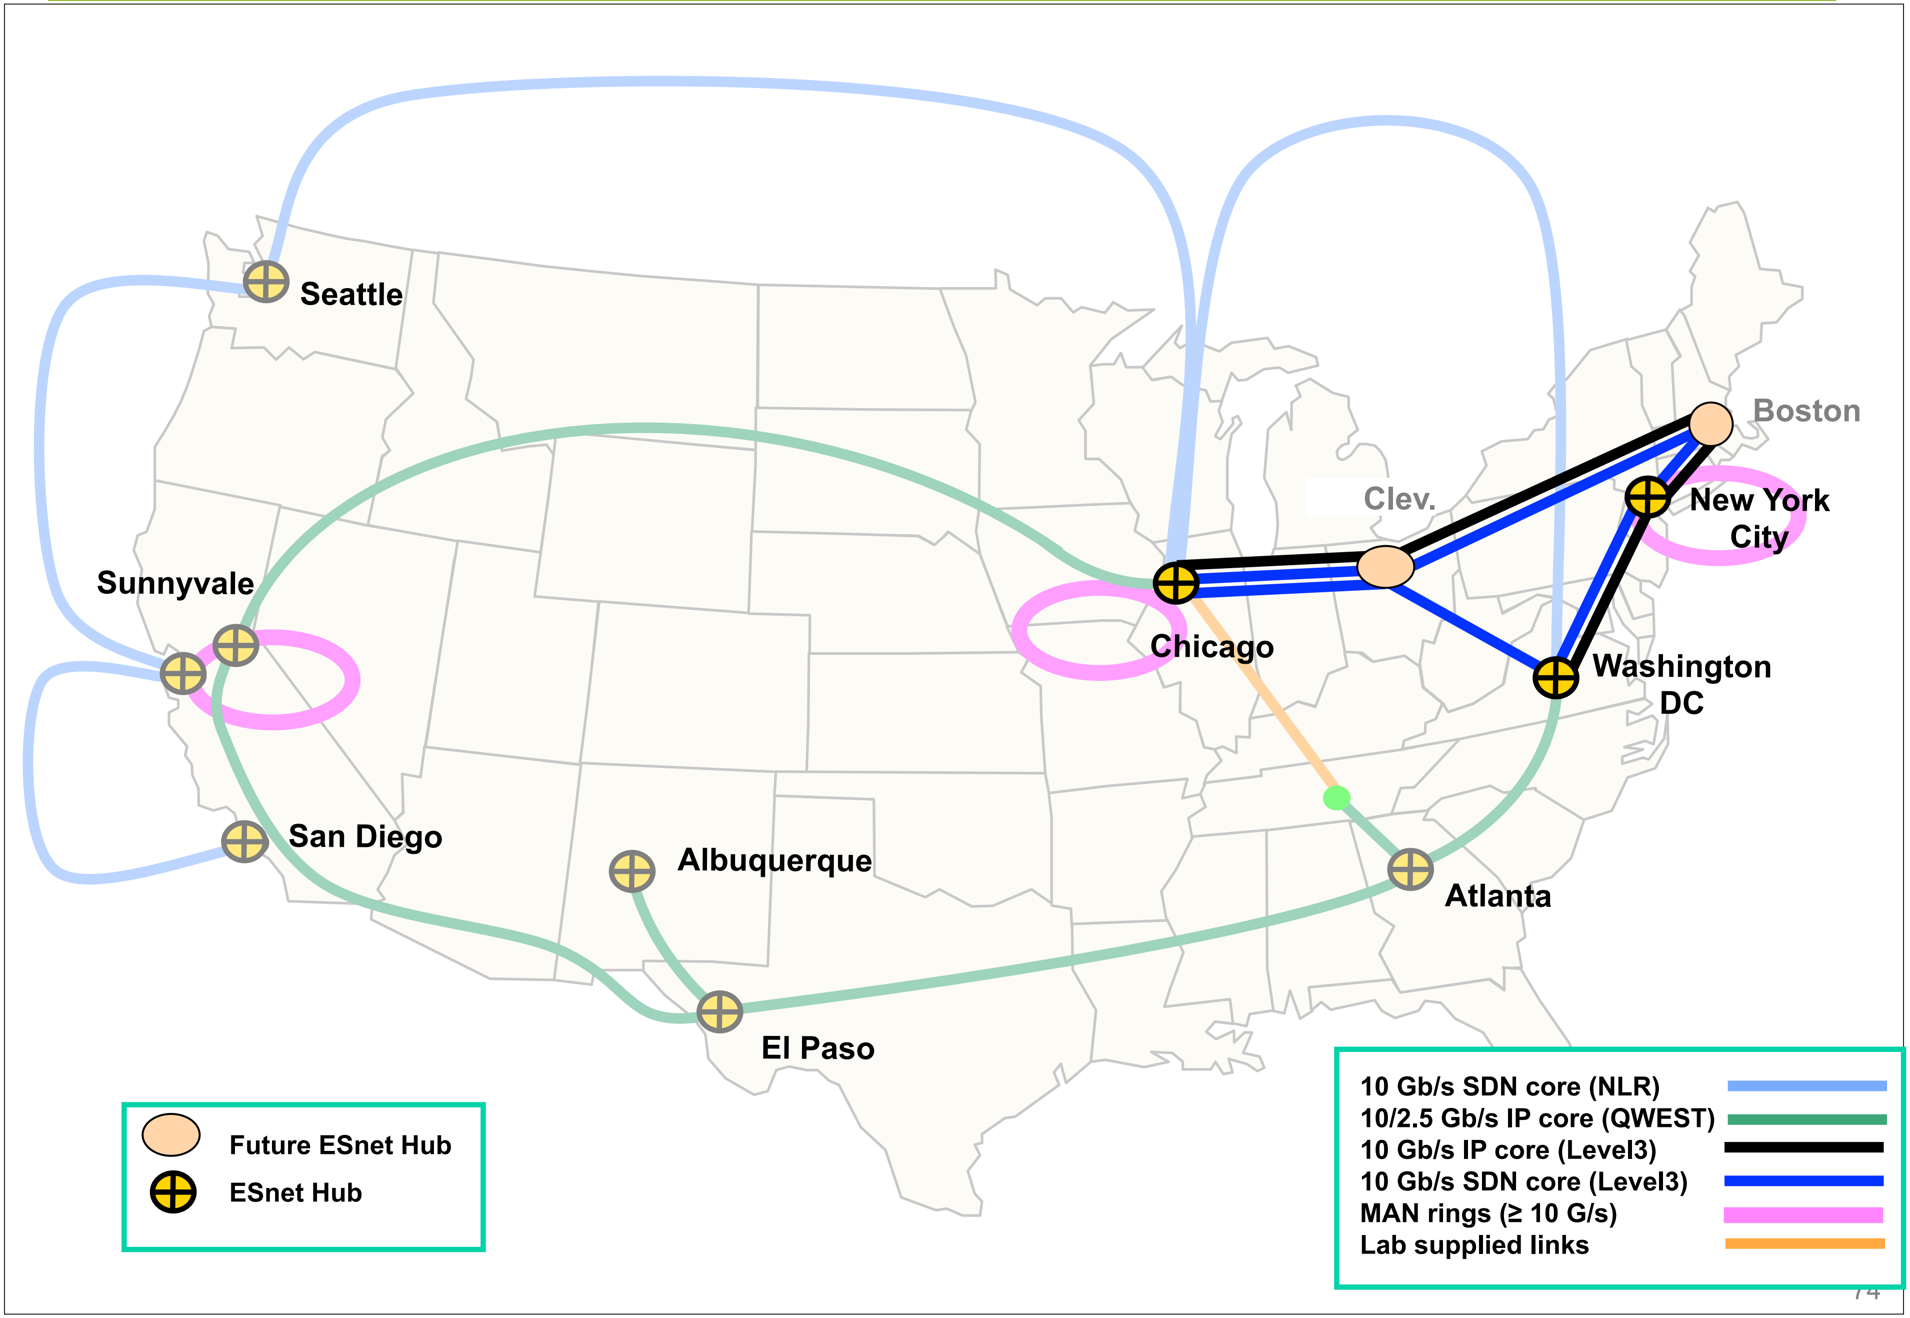

ESnet 4 Backbone as of May 15, 2007

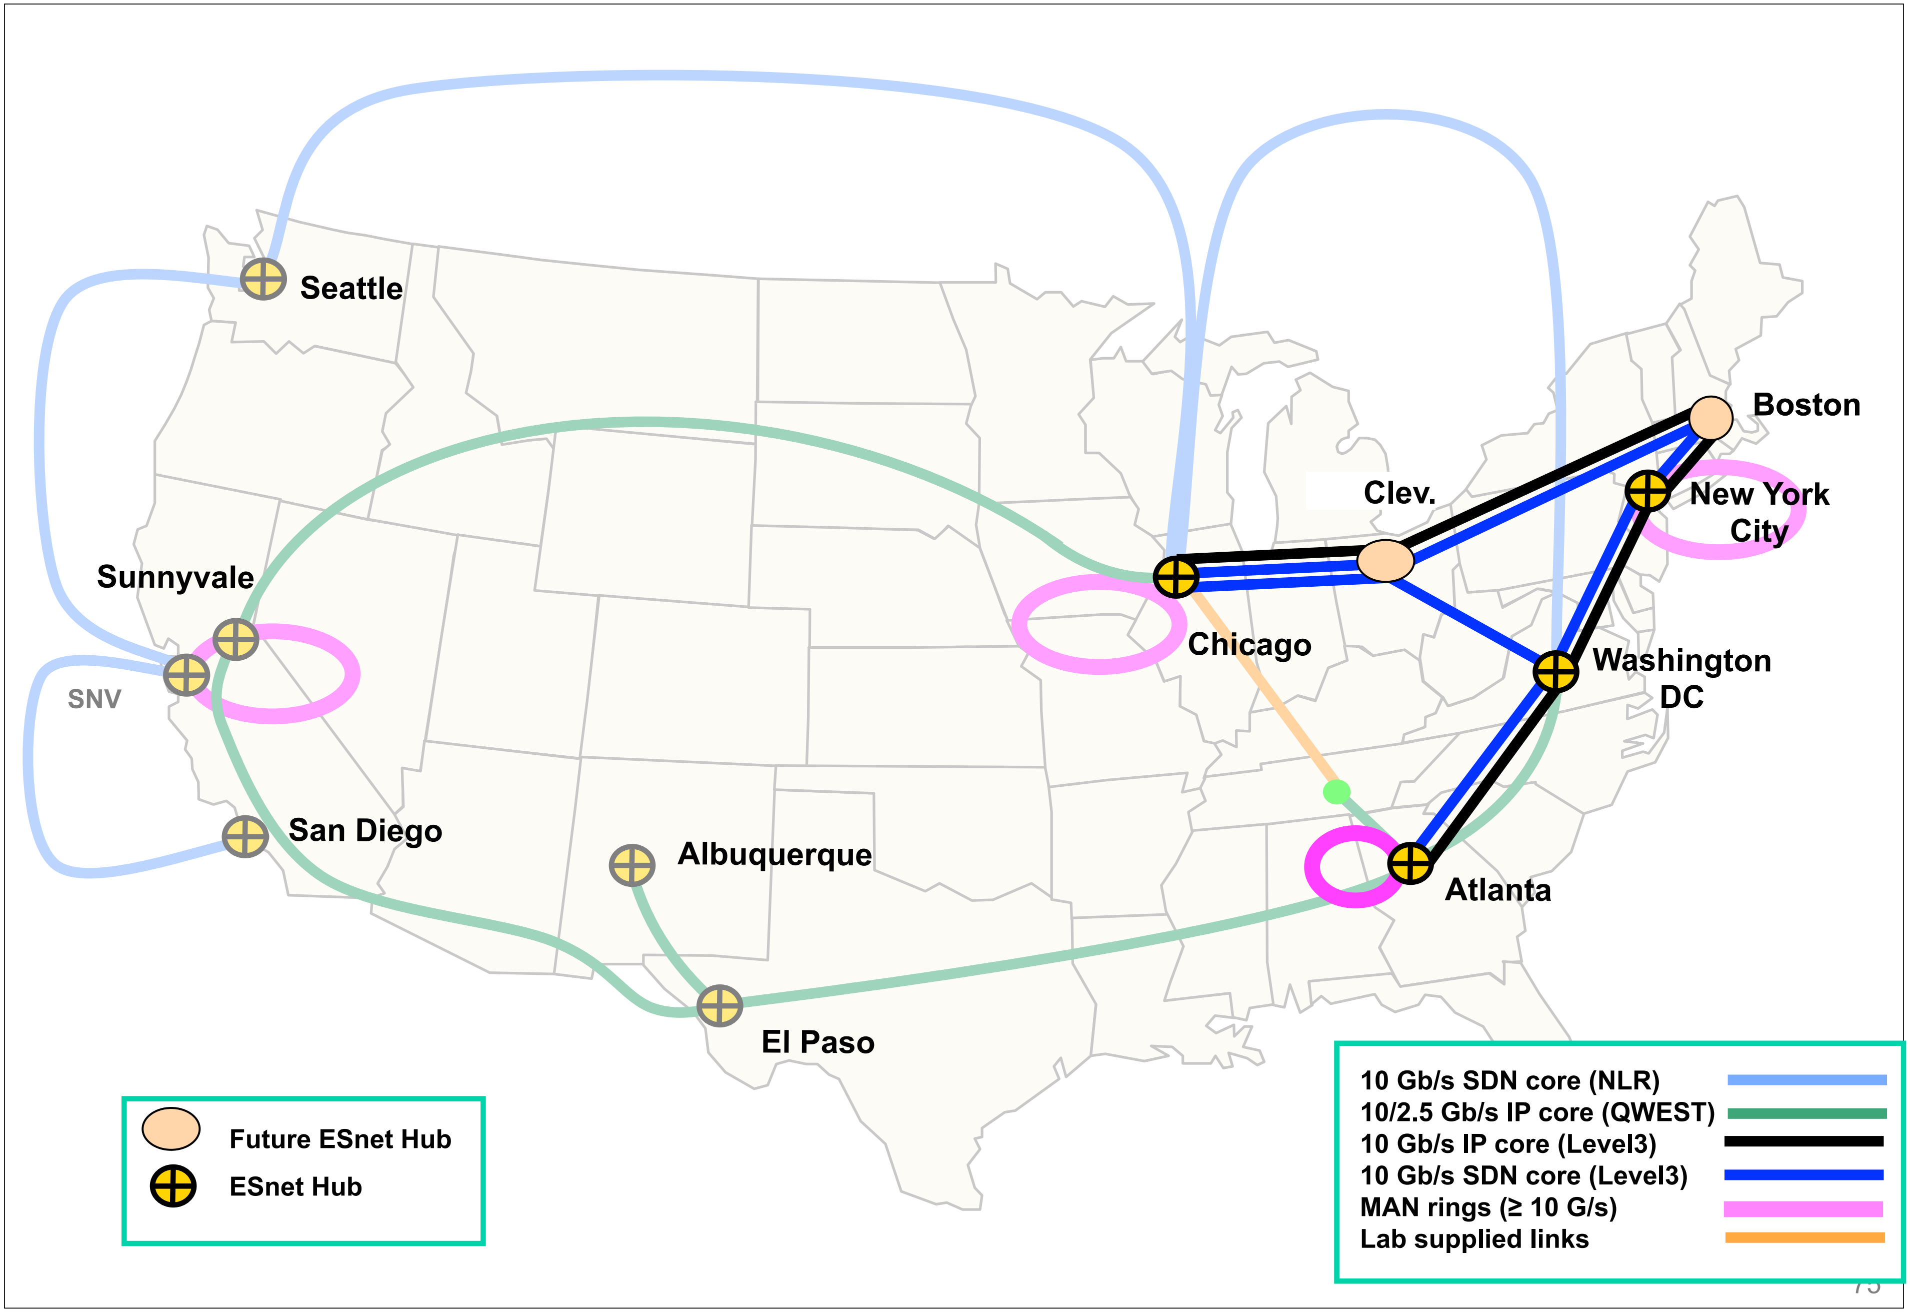

ESnet 4 Backbone as of June 20, 2007

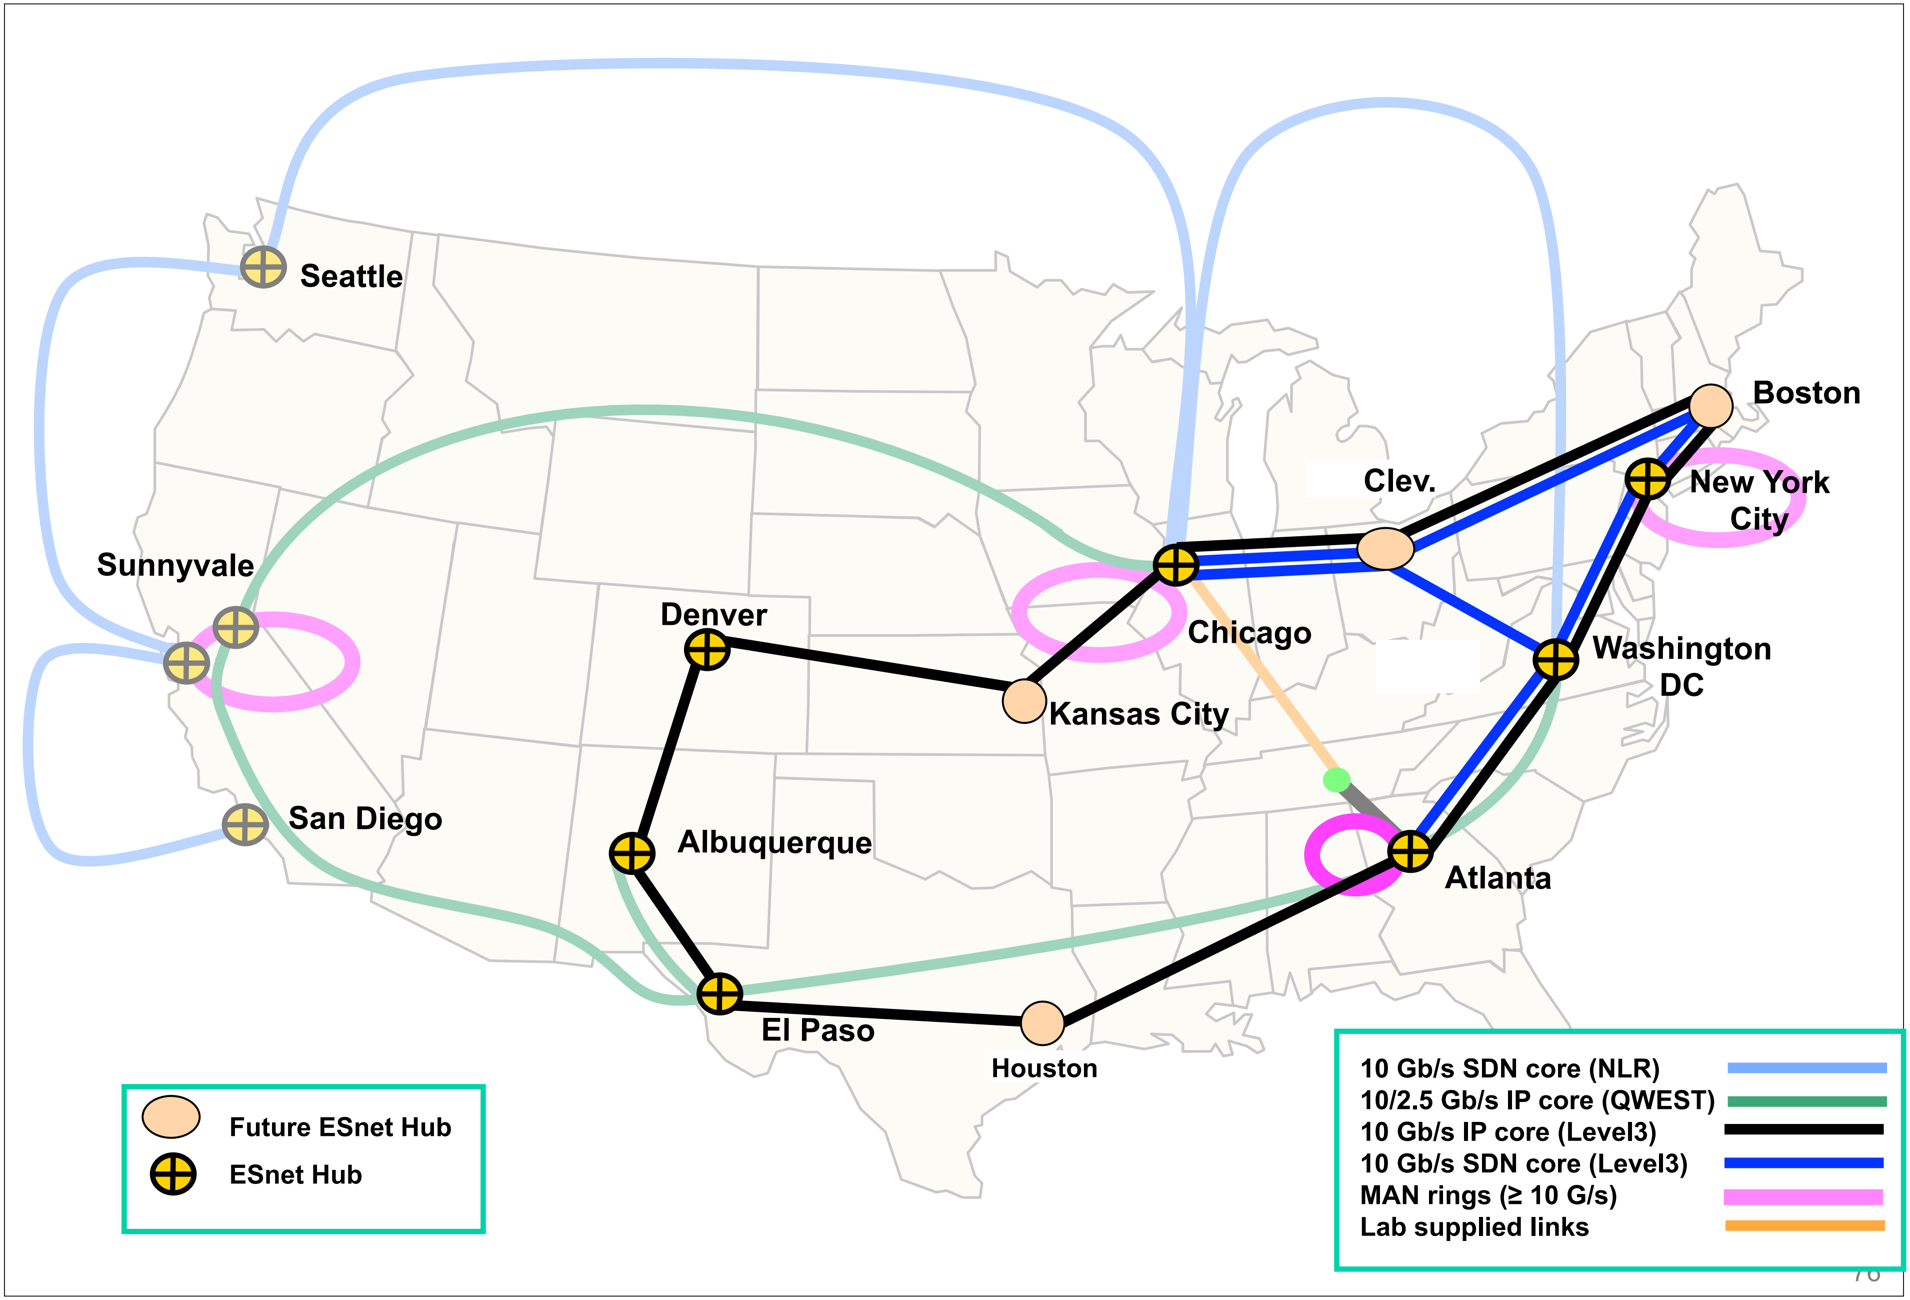

ESnet 4 Backbone August 1, 2007 (Last JT meeting at FNAL)

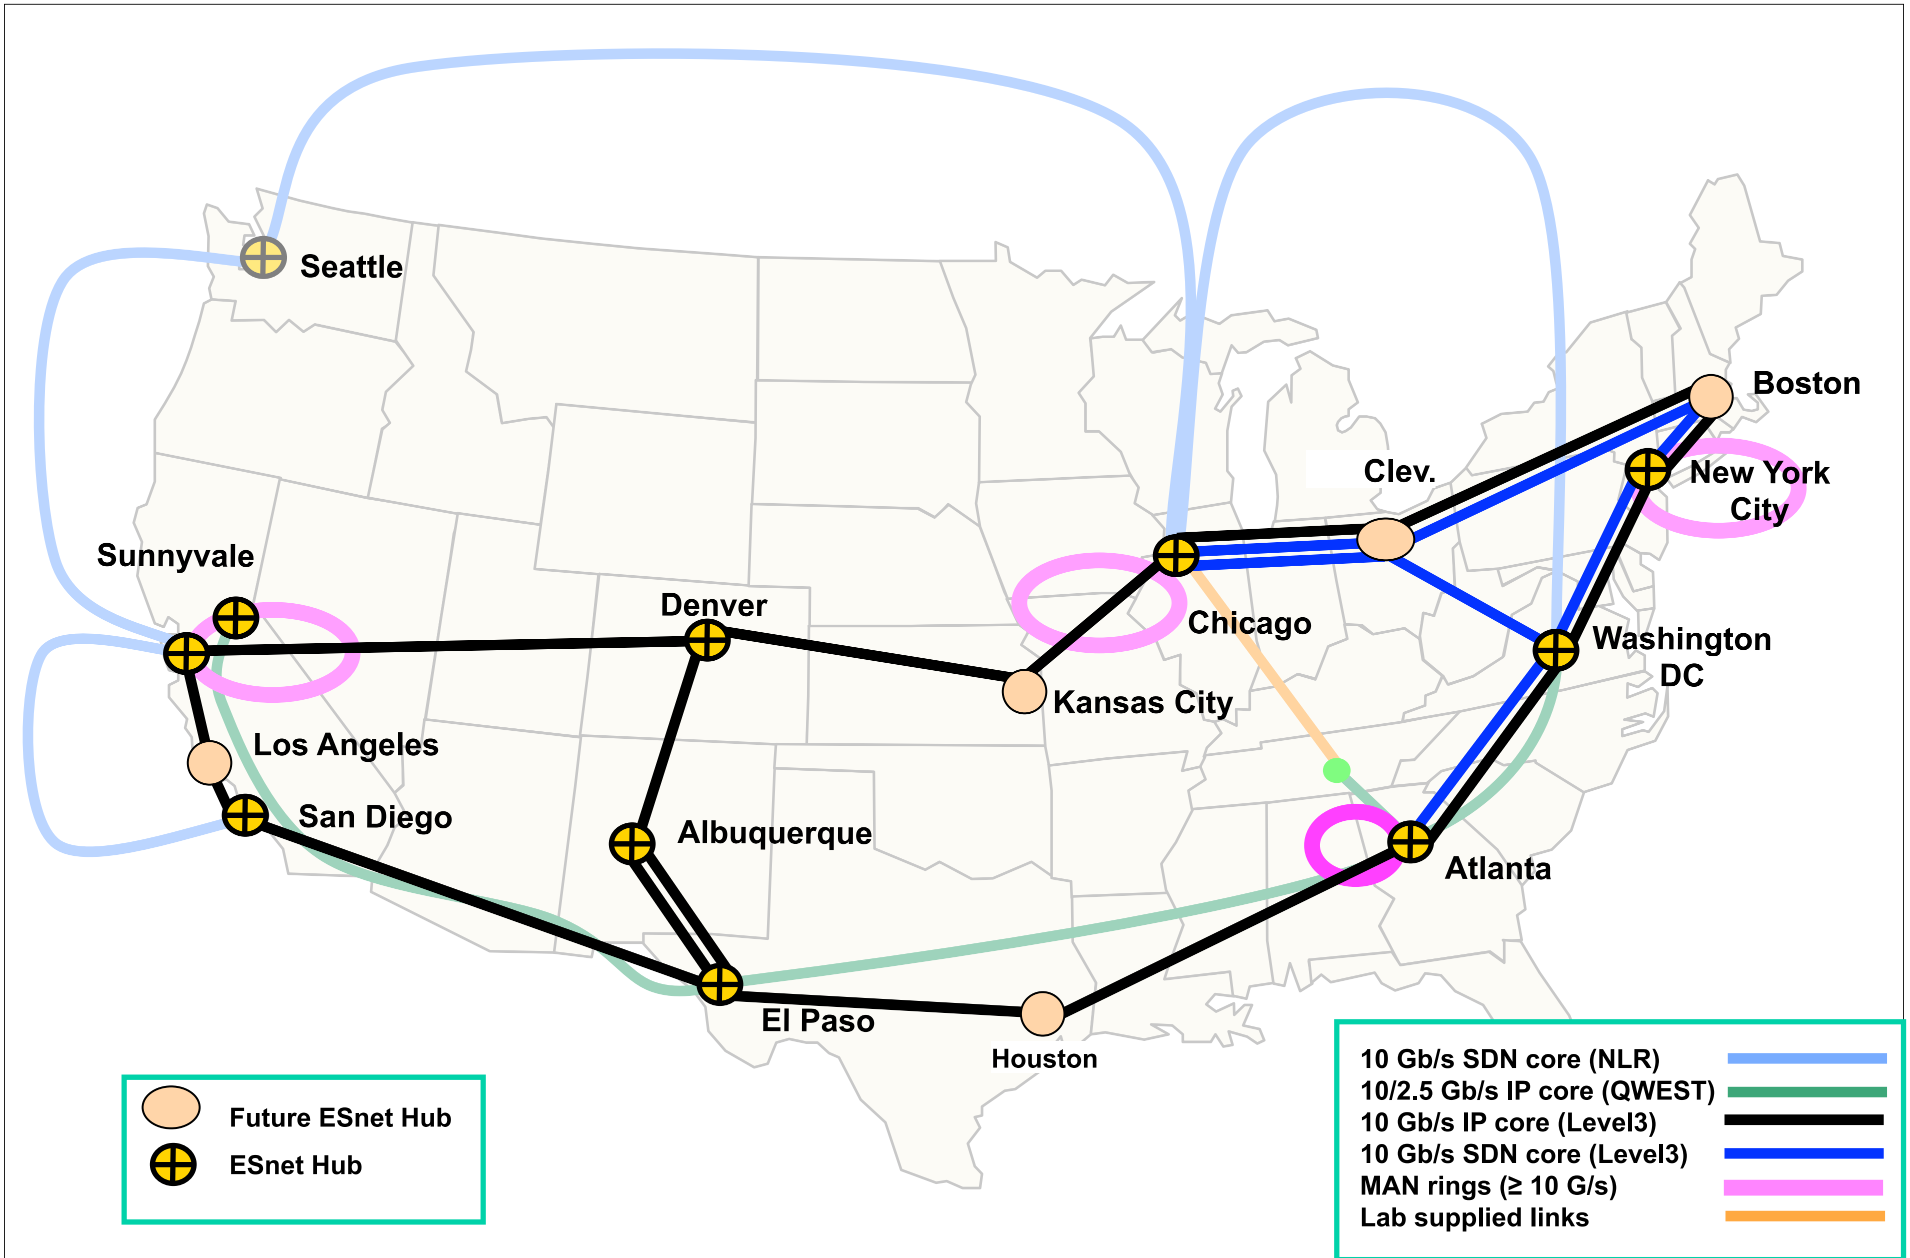

ESnet 4 Backbone September 30, 2007

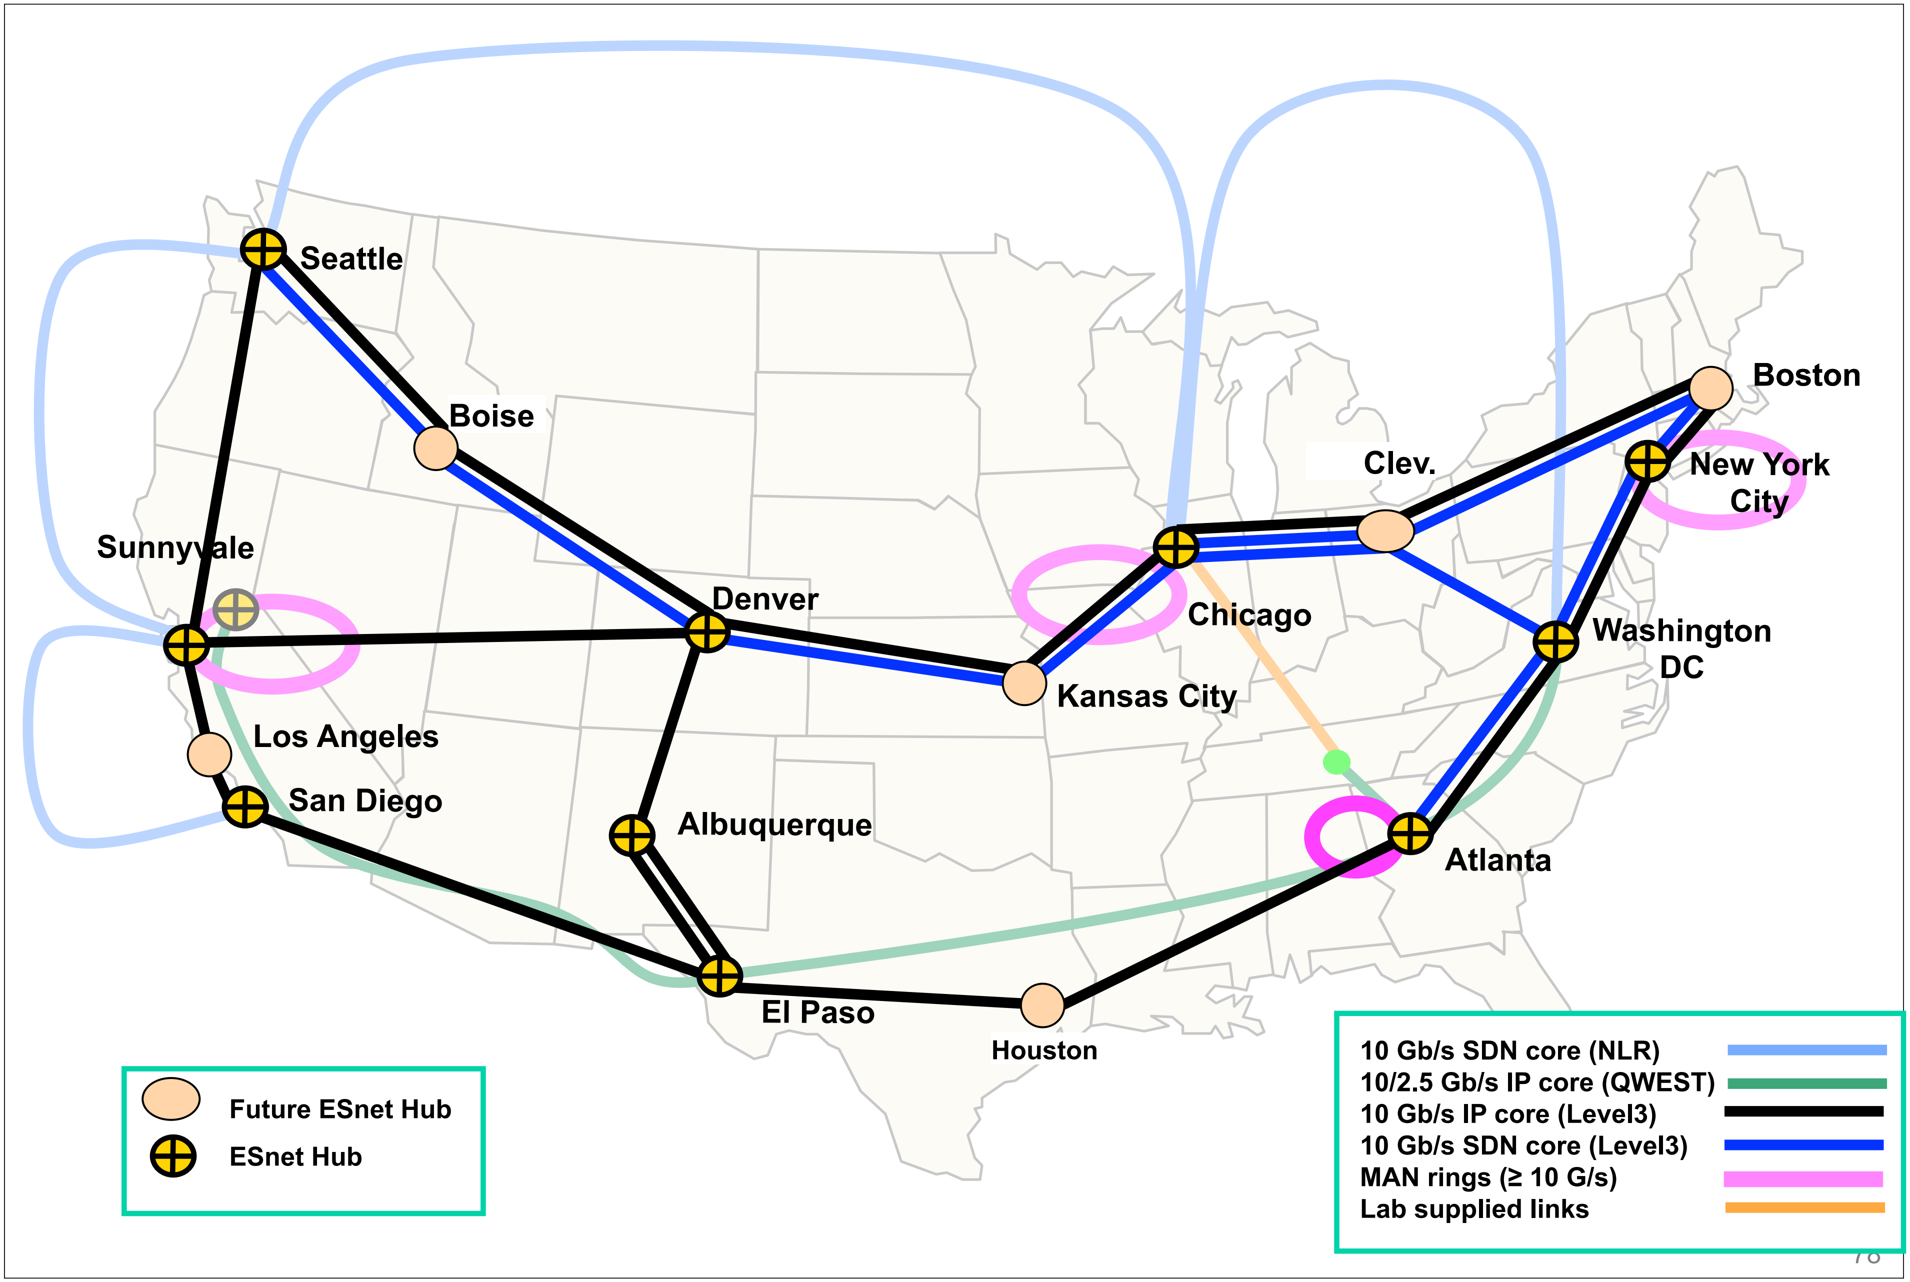

ESnet 4 Backbone December 2007

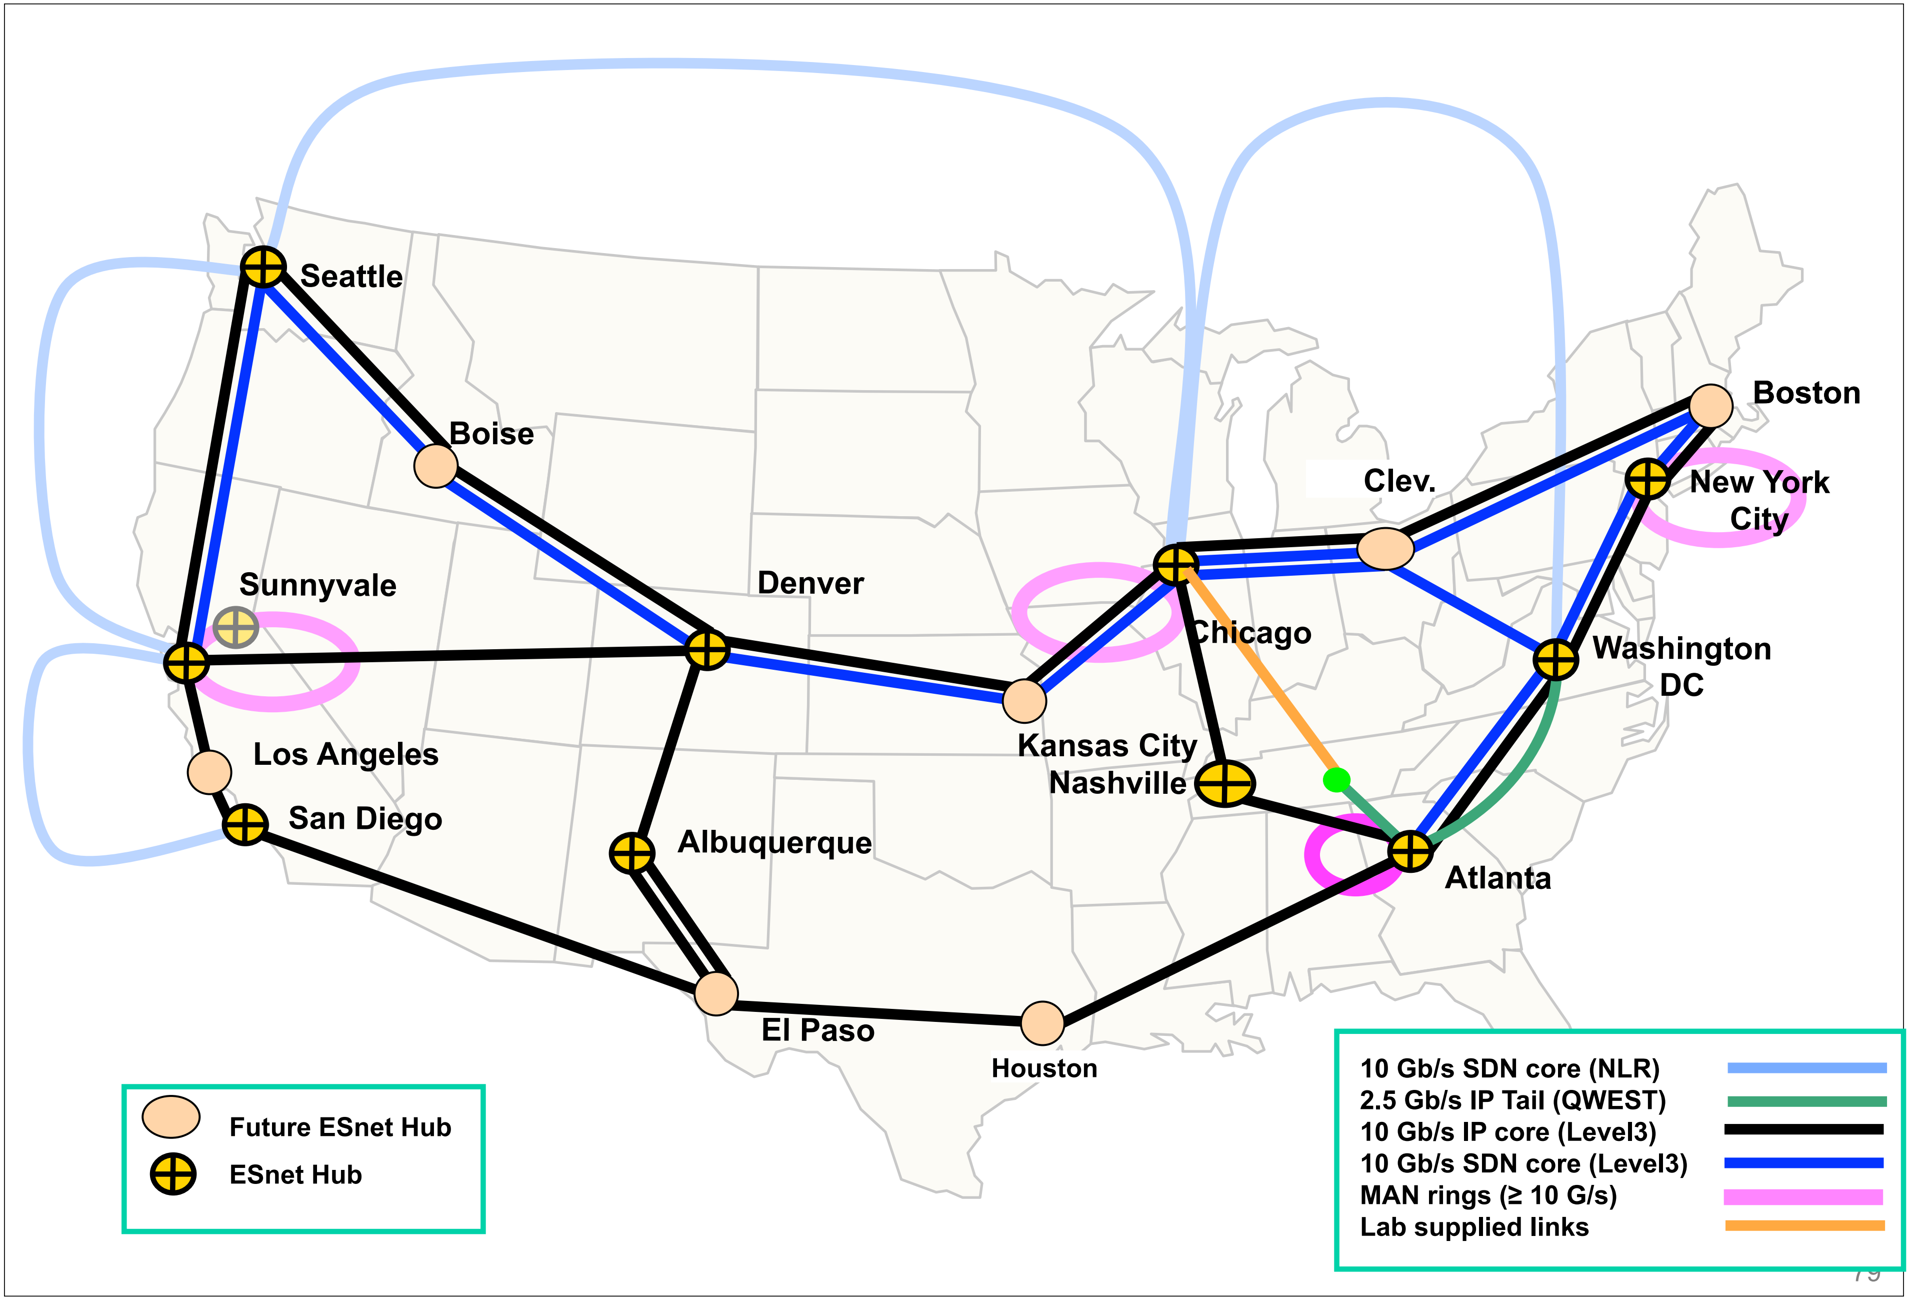

ESnet 4 Backbone December, 2008

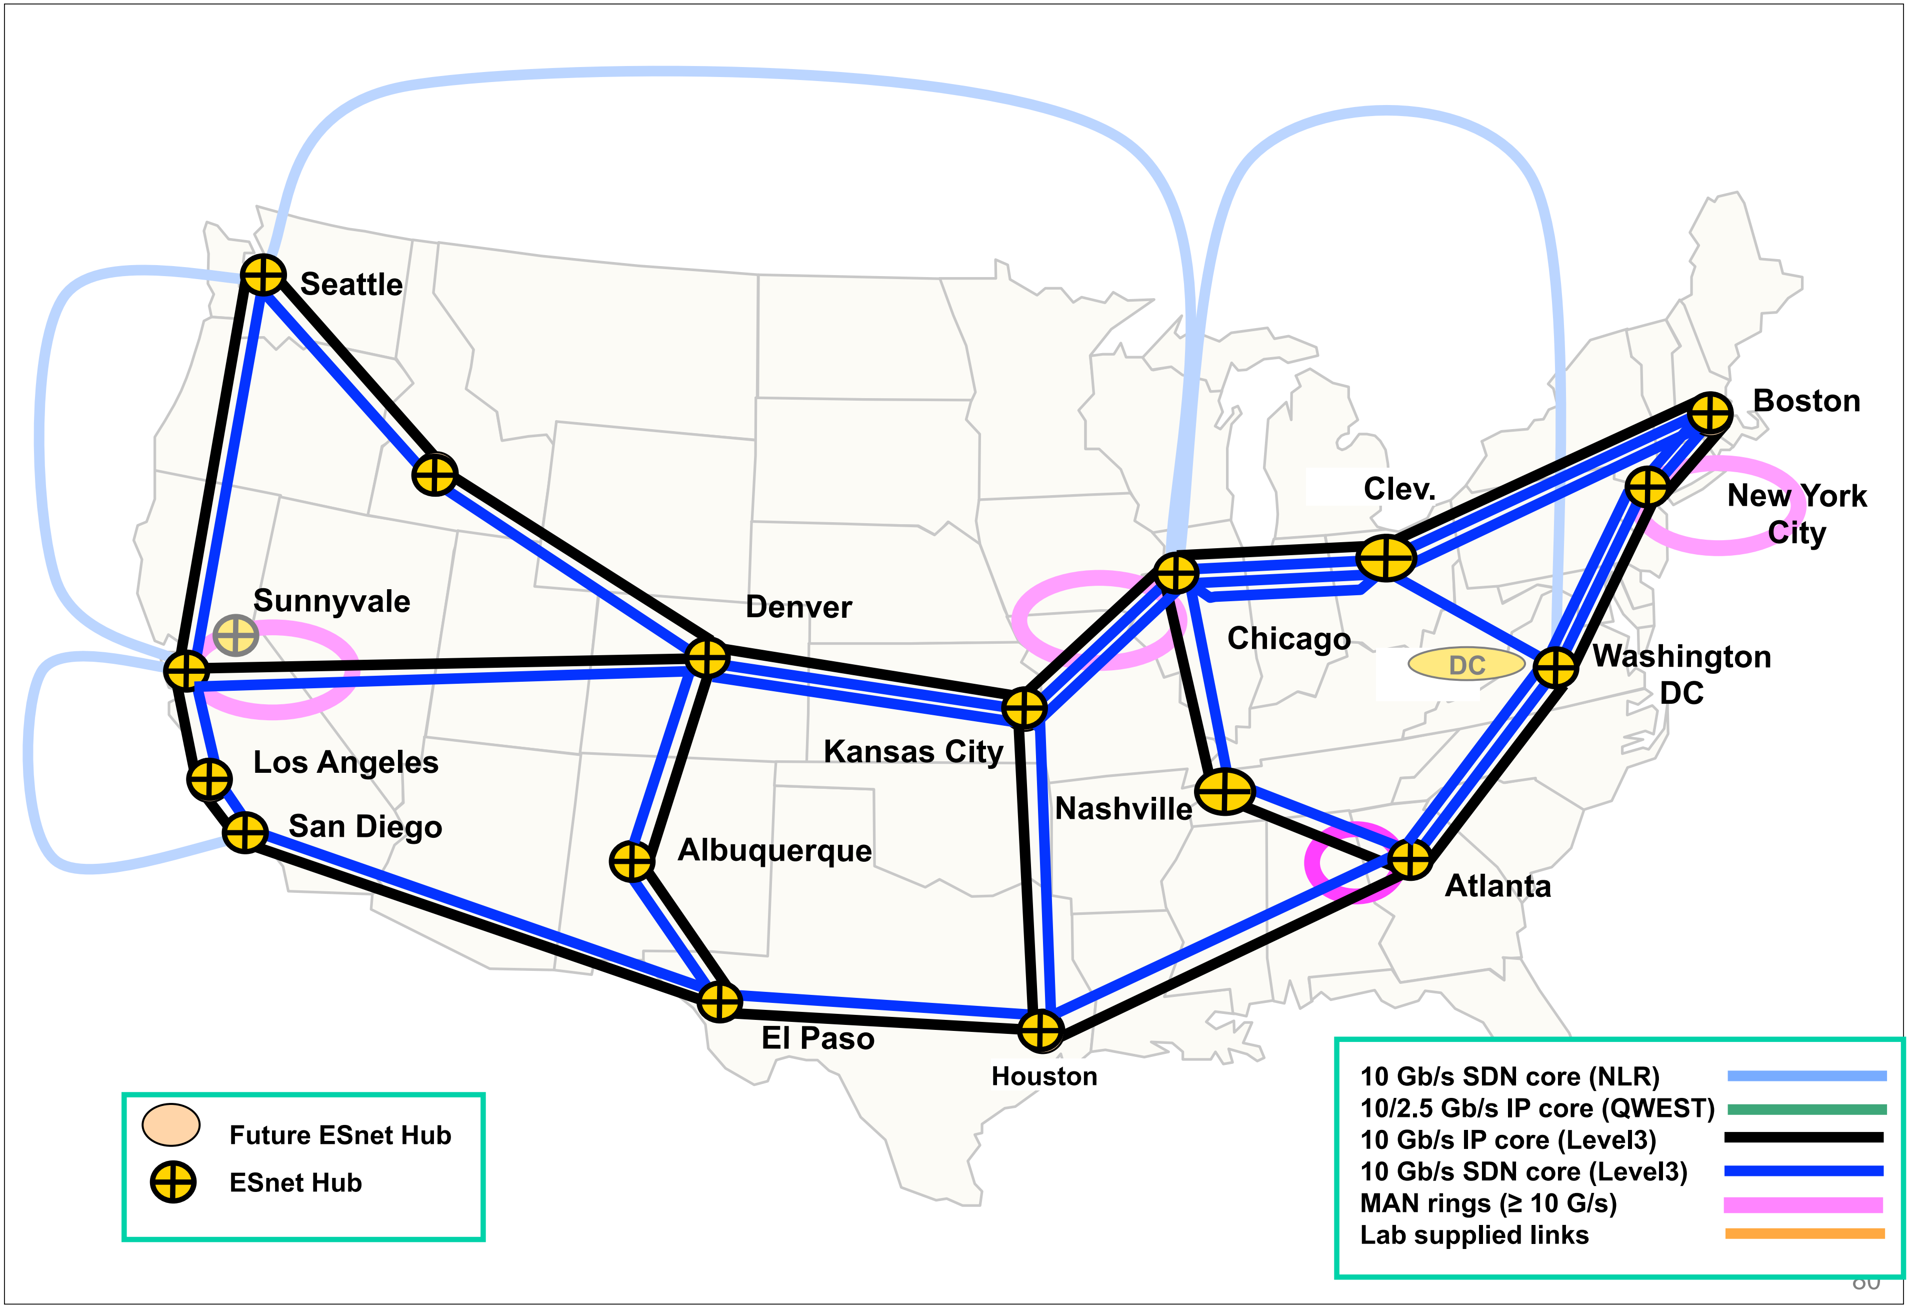

Connections at a Typical ESnet4 Wide Area Network Hub

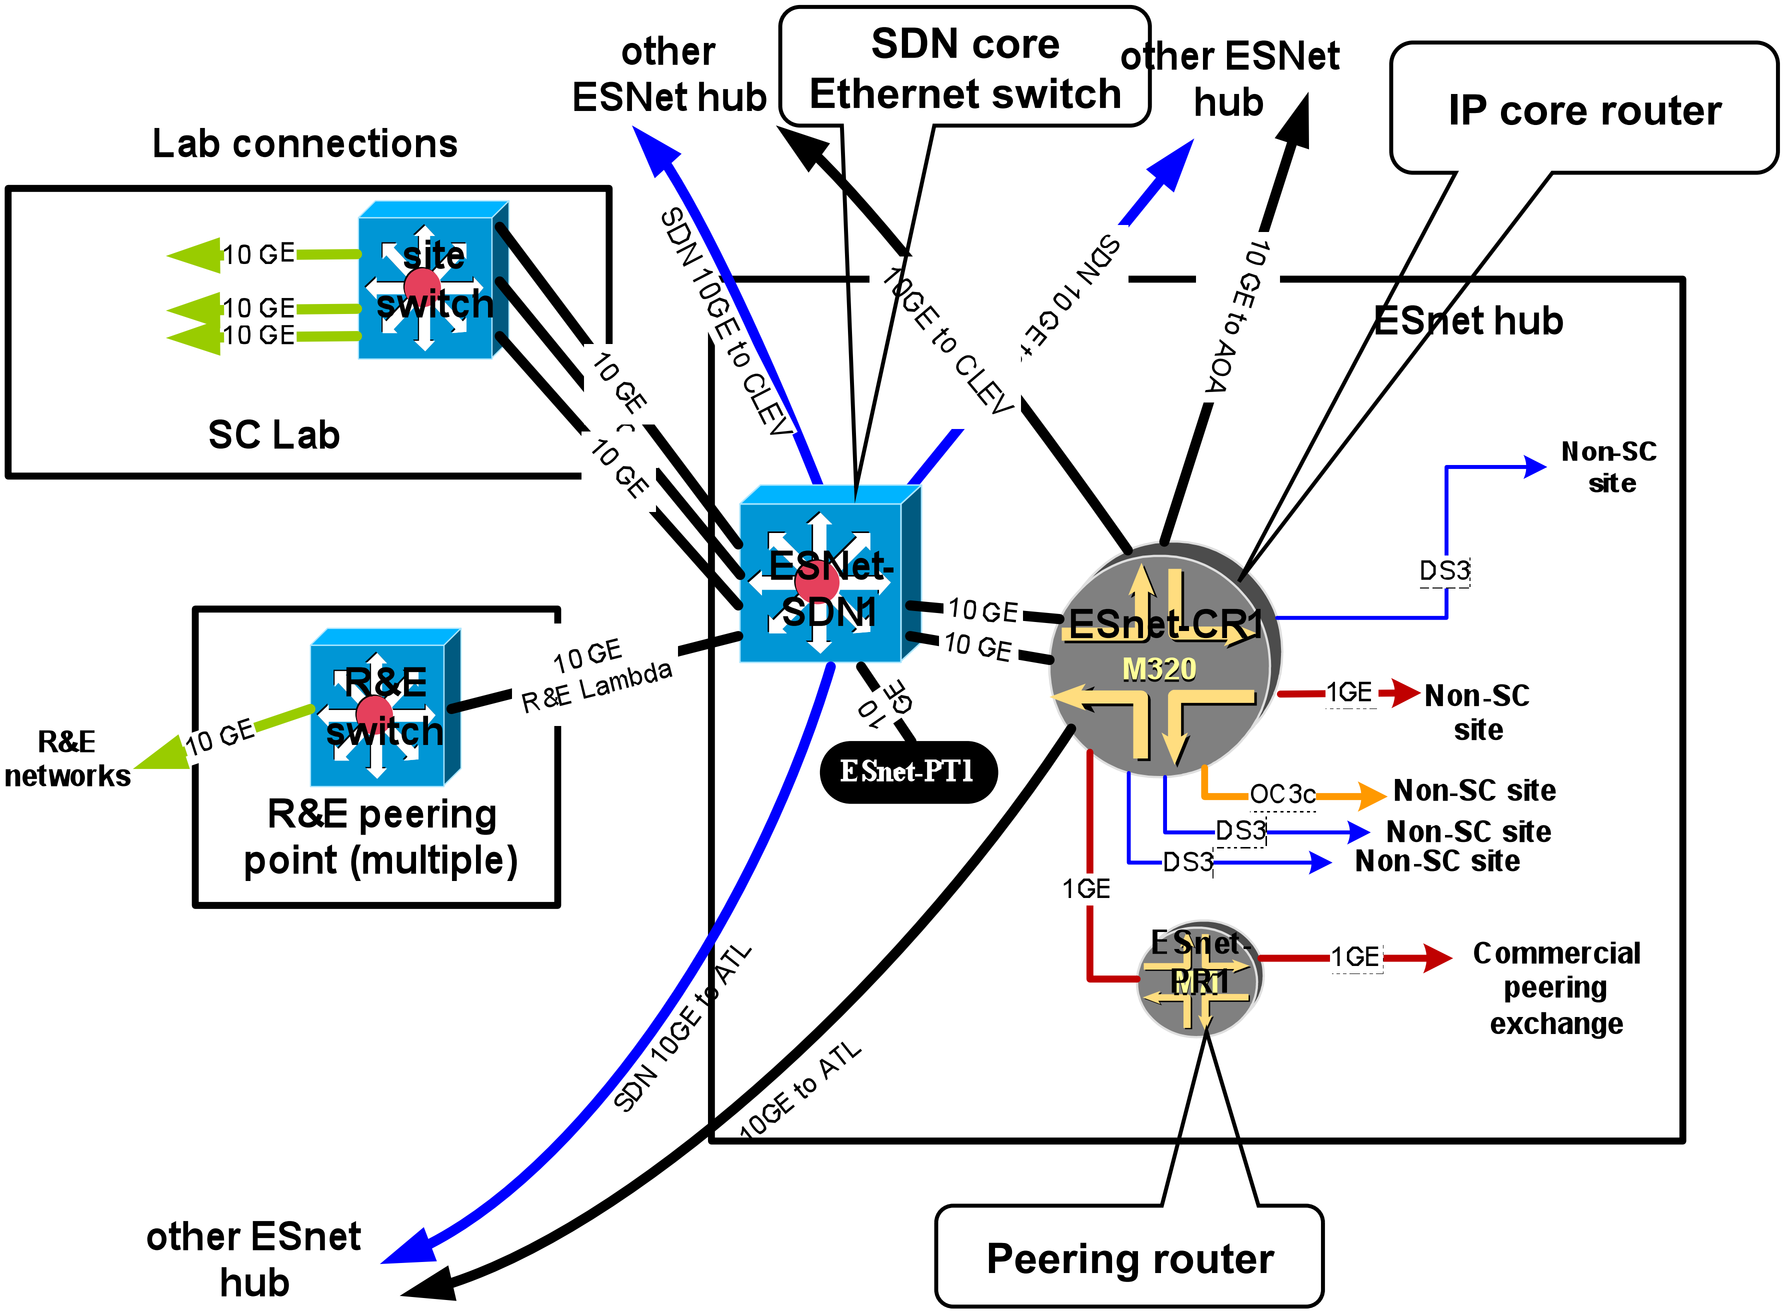

Typical Medium Size ESnet4 Hub

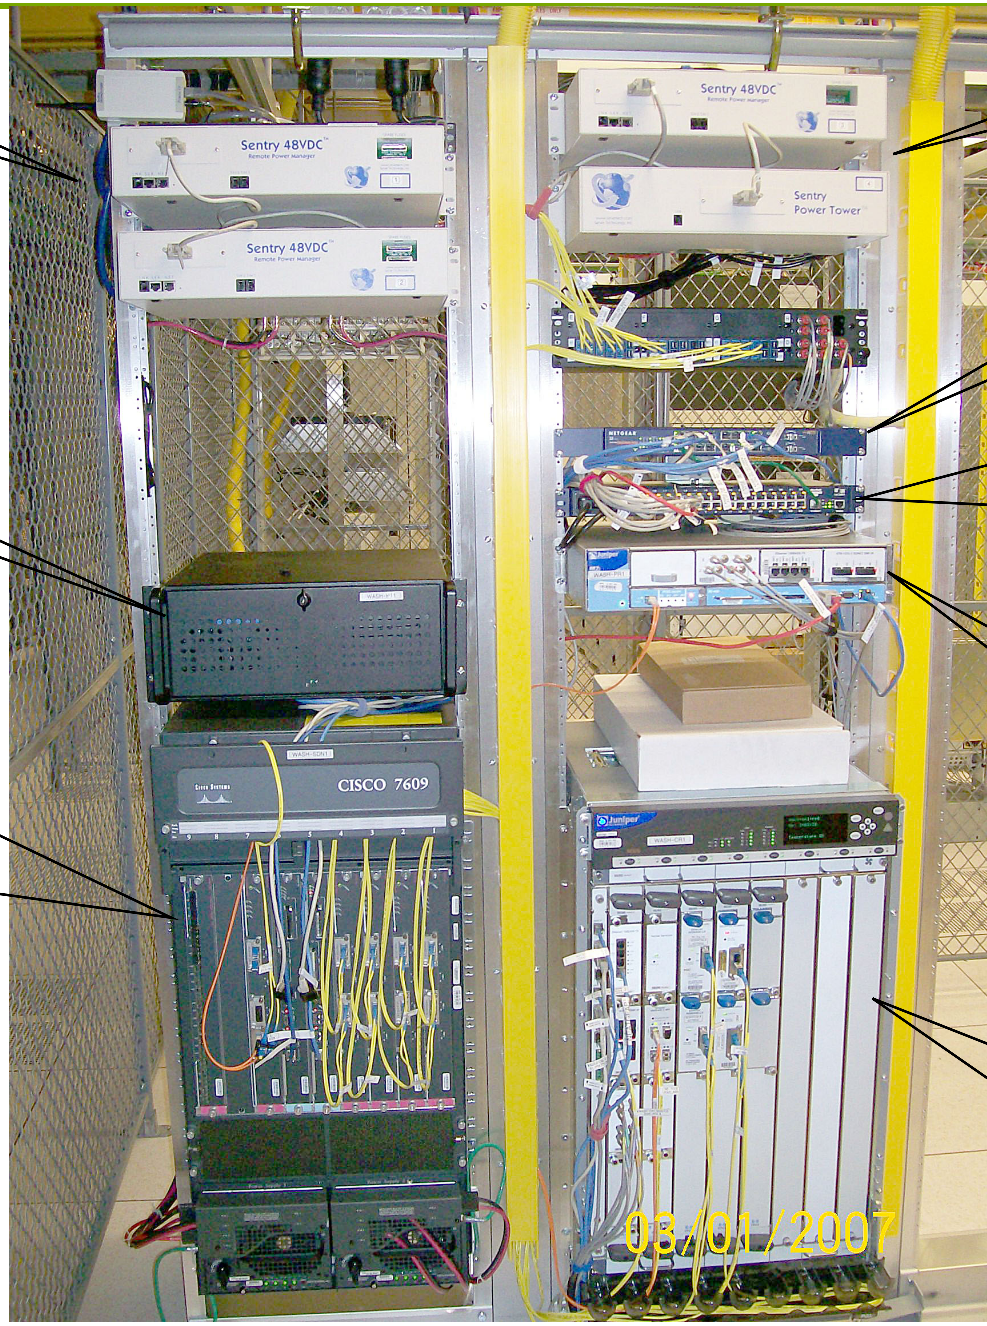

DC power controllers

10 Gbps network tester

Cisco 7609 or Juniper MX960 Science Data Network switch (about $365K, list)

DC power controllers

local Ethernet

secure terminal server (telephone modem access)

Juniper M7i, IP peering router, (about $60K, list)

Juniper M320, core IP router, (about $1.65M, list)

ESnet4 IP + SDN, 2008 Configuration

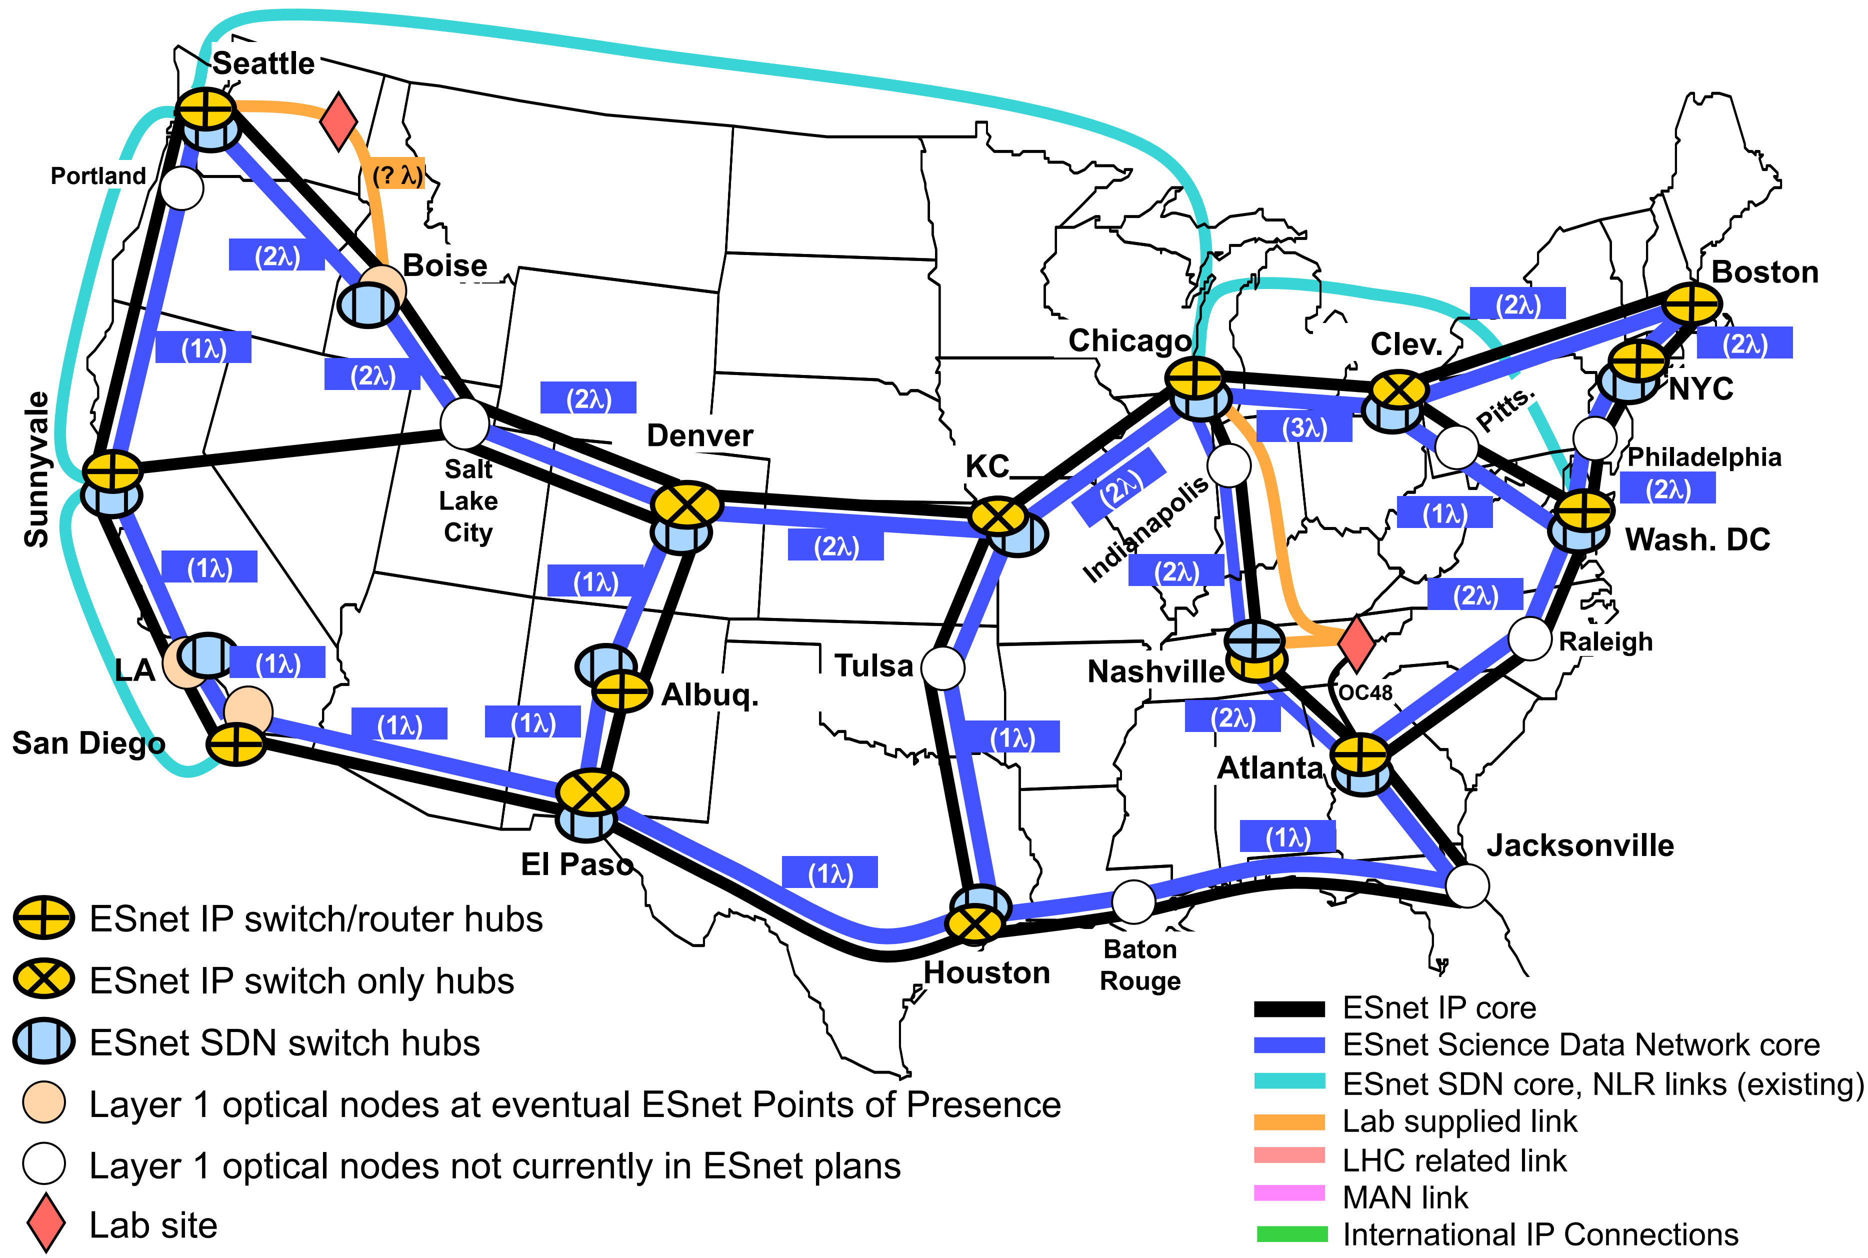

ESnet4 Metro Area Rings, 2007 Configurations

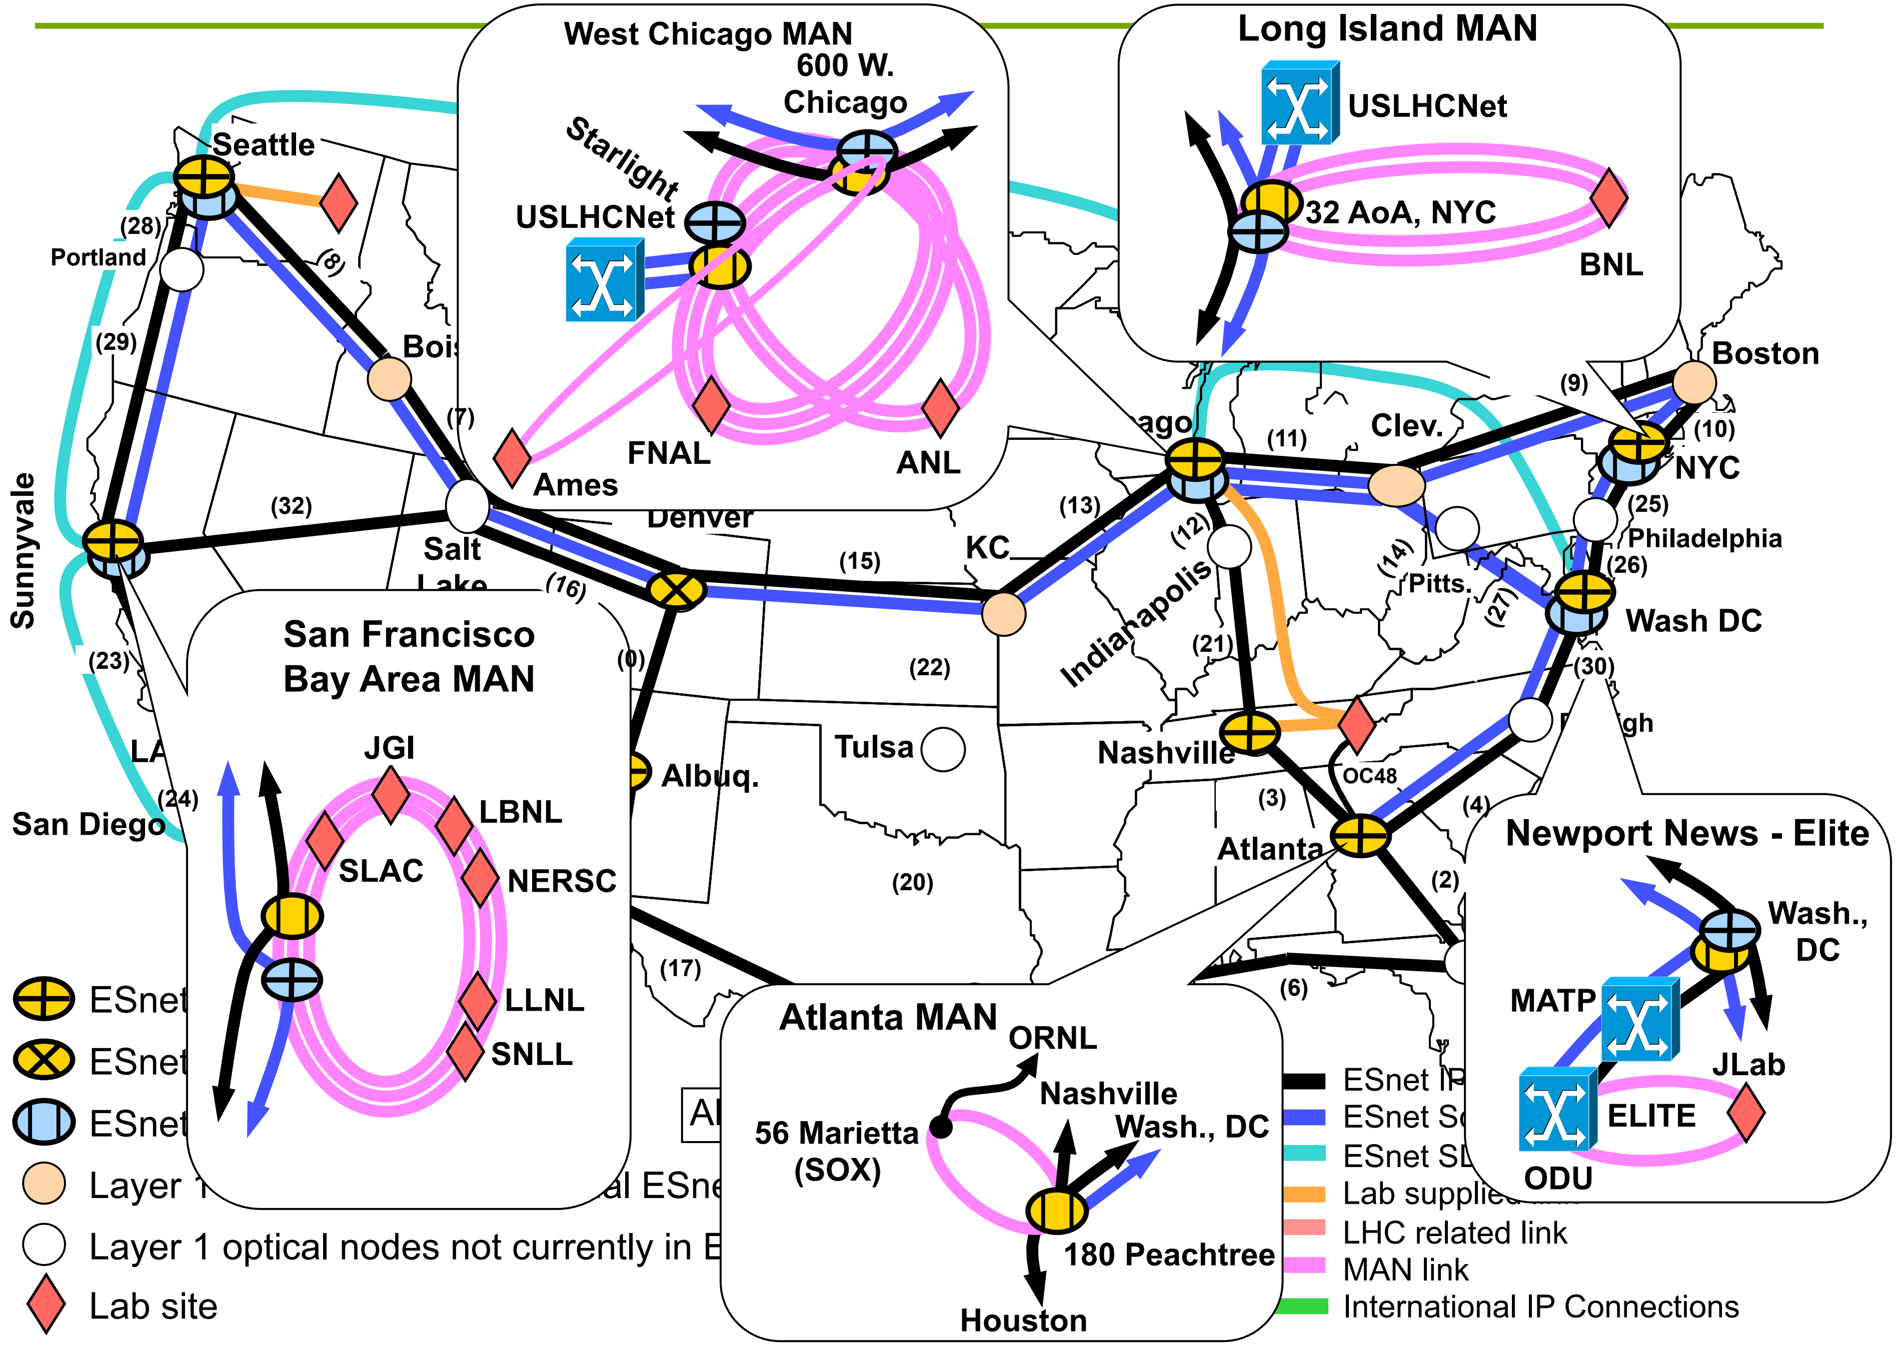

➢ Note that the major ESnet sites are now directly on the ESnet “core” network

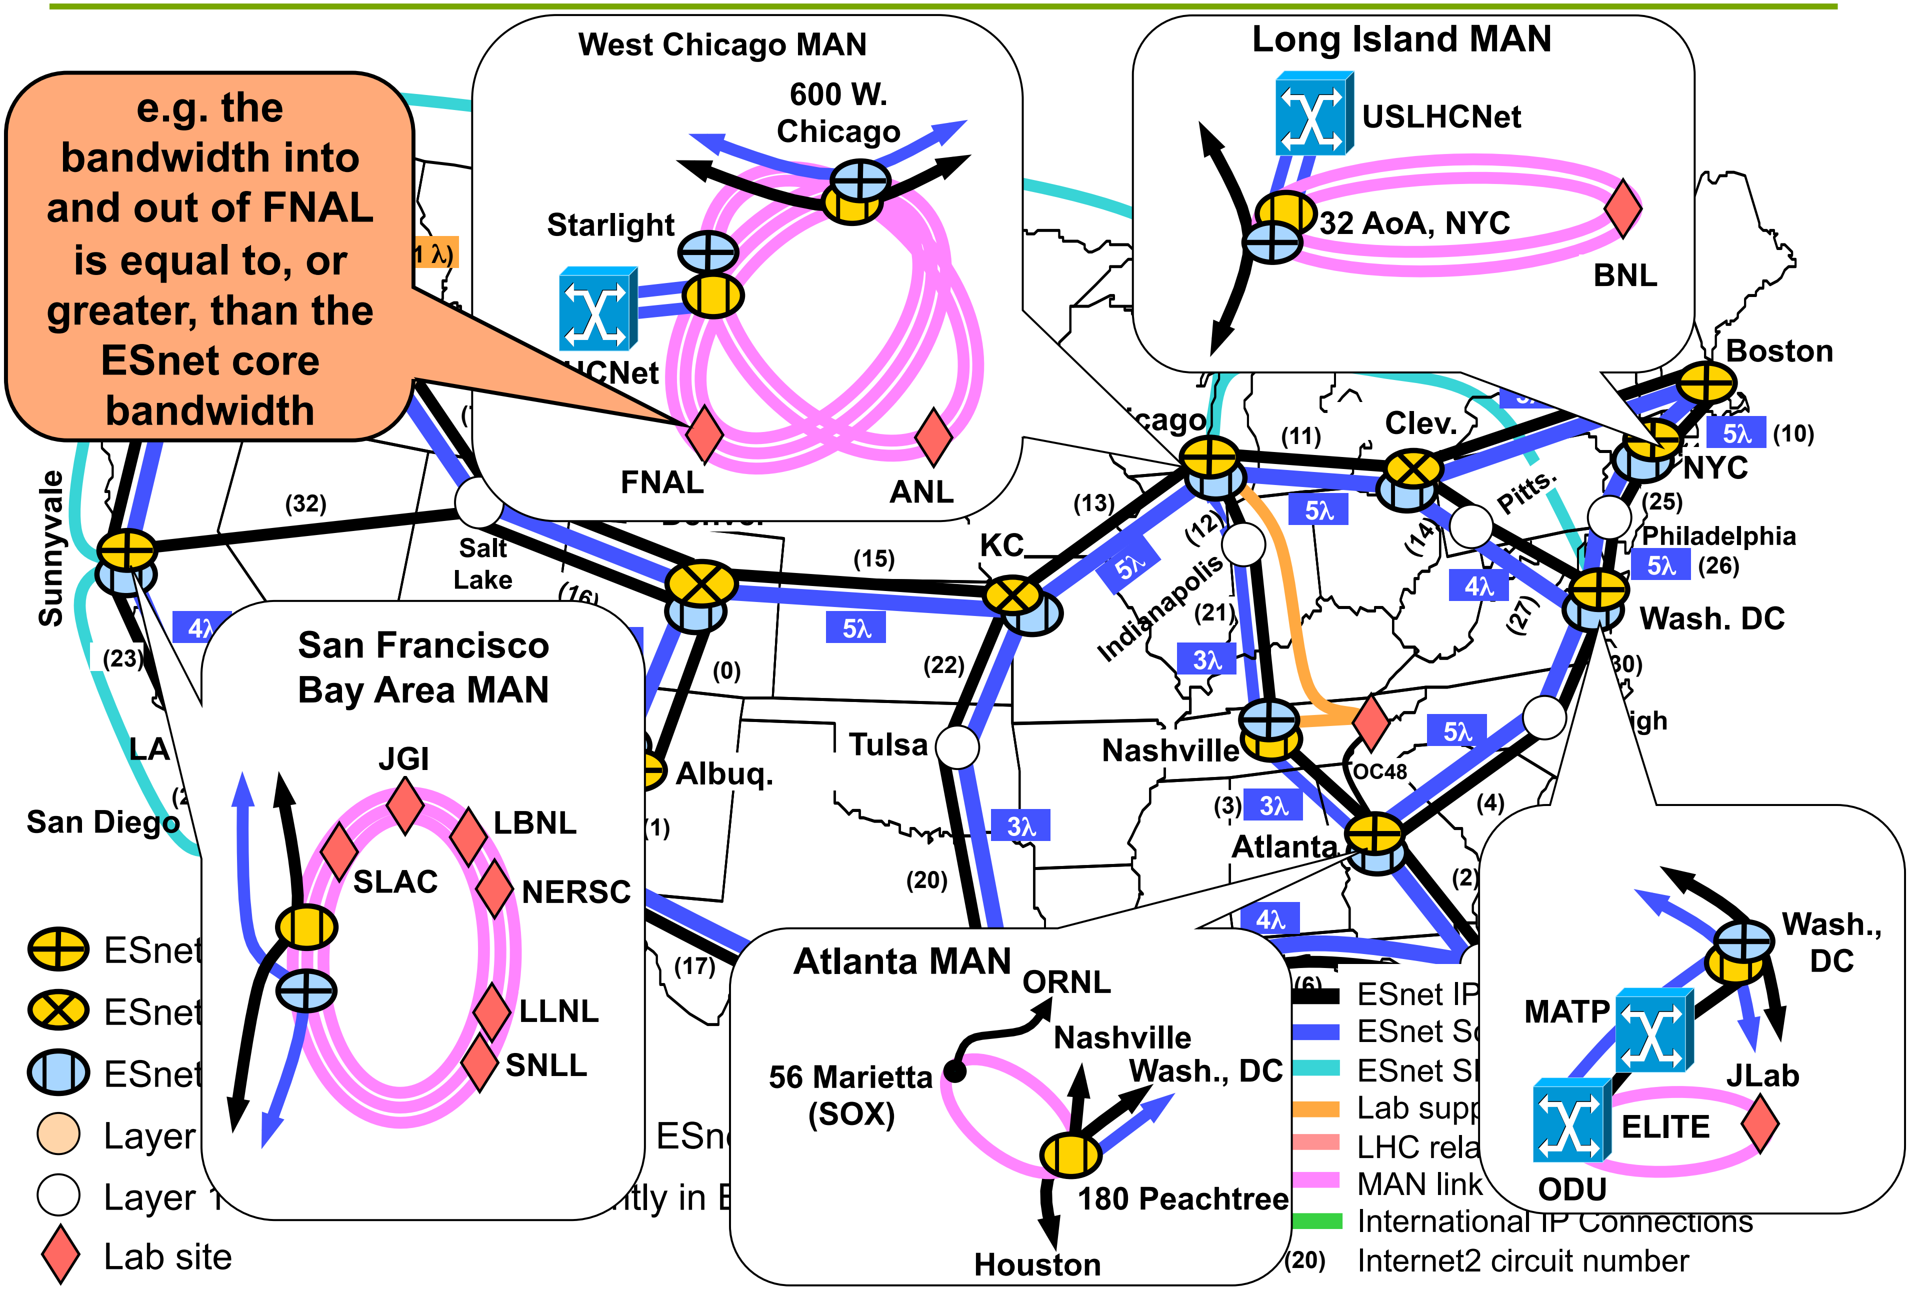

Estimated ESnet4 2009 Configuration

(Some of the circuits may be allocated dynamically from shared a pool.)

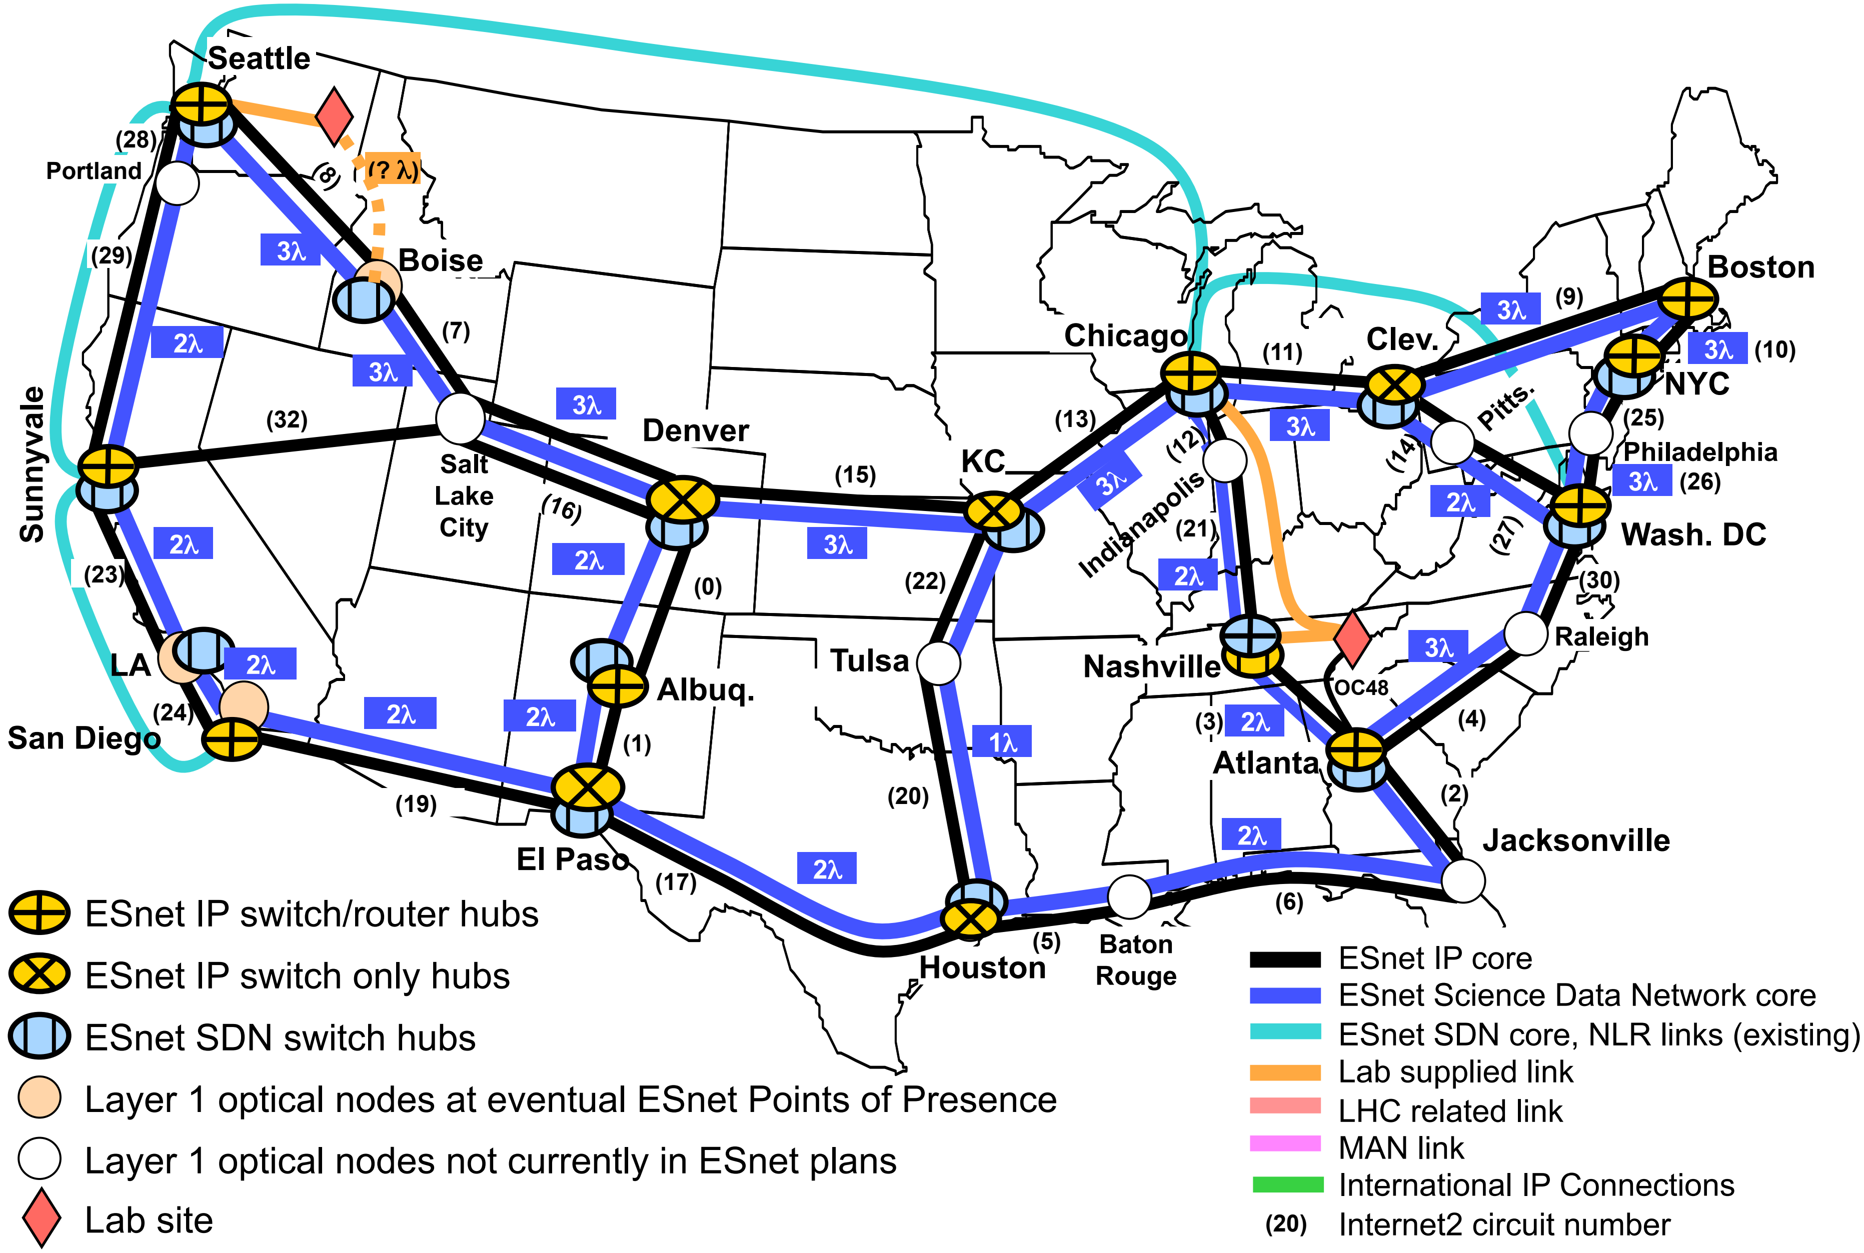

Aggregate Estimated Link Loadings, 2010-11

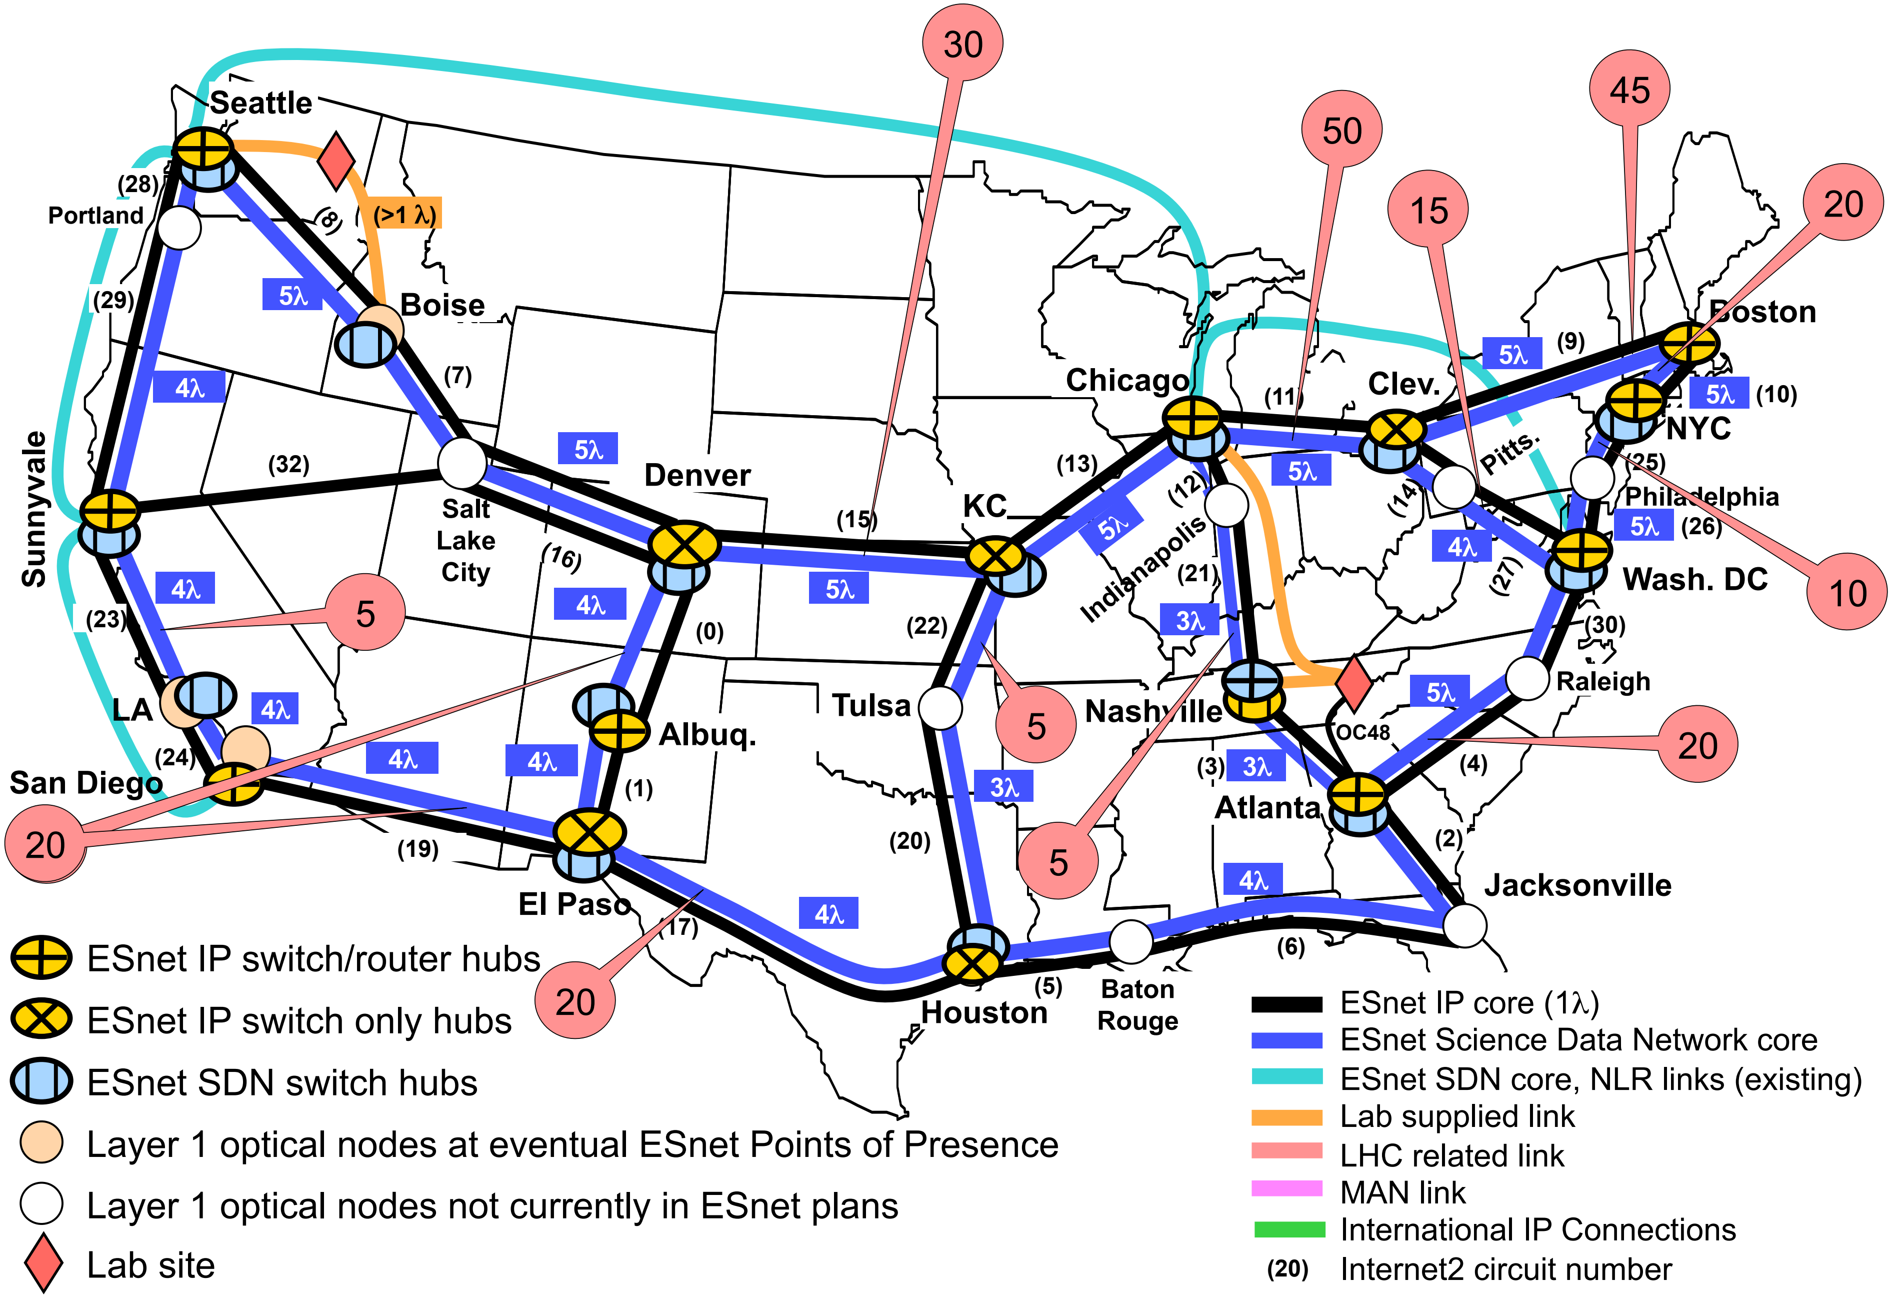

(2C) ESnet4 2010-11 Estimated Bandwidth Commitments

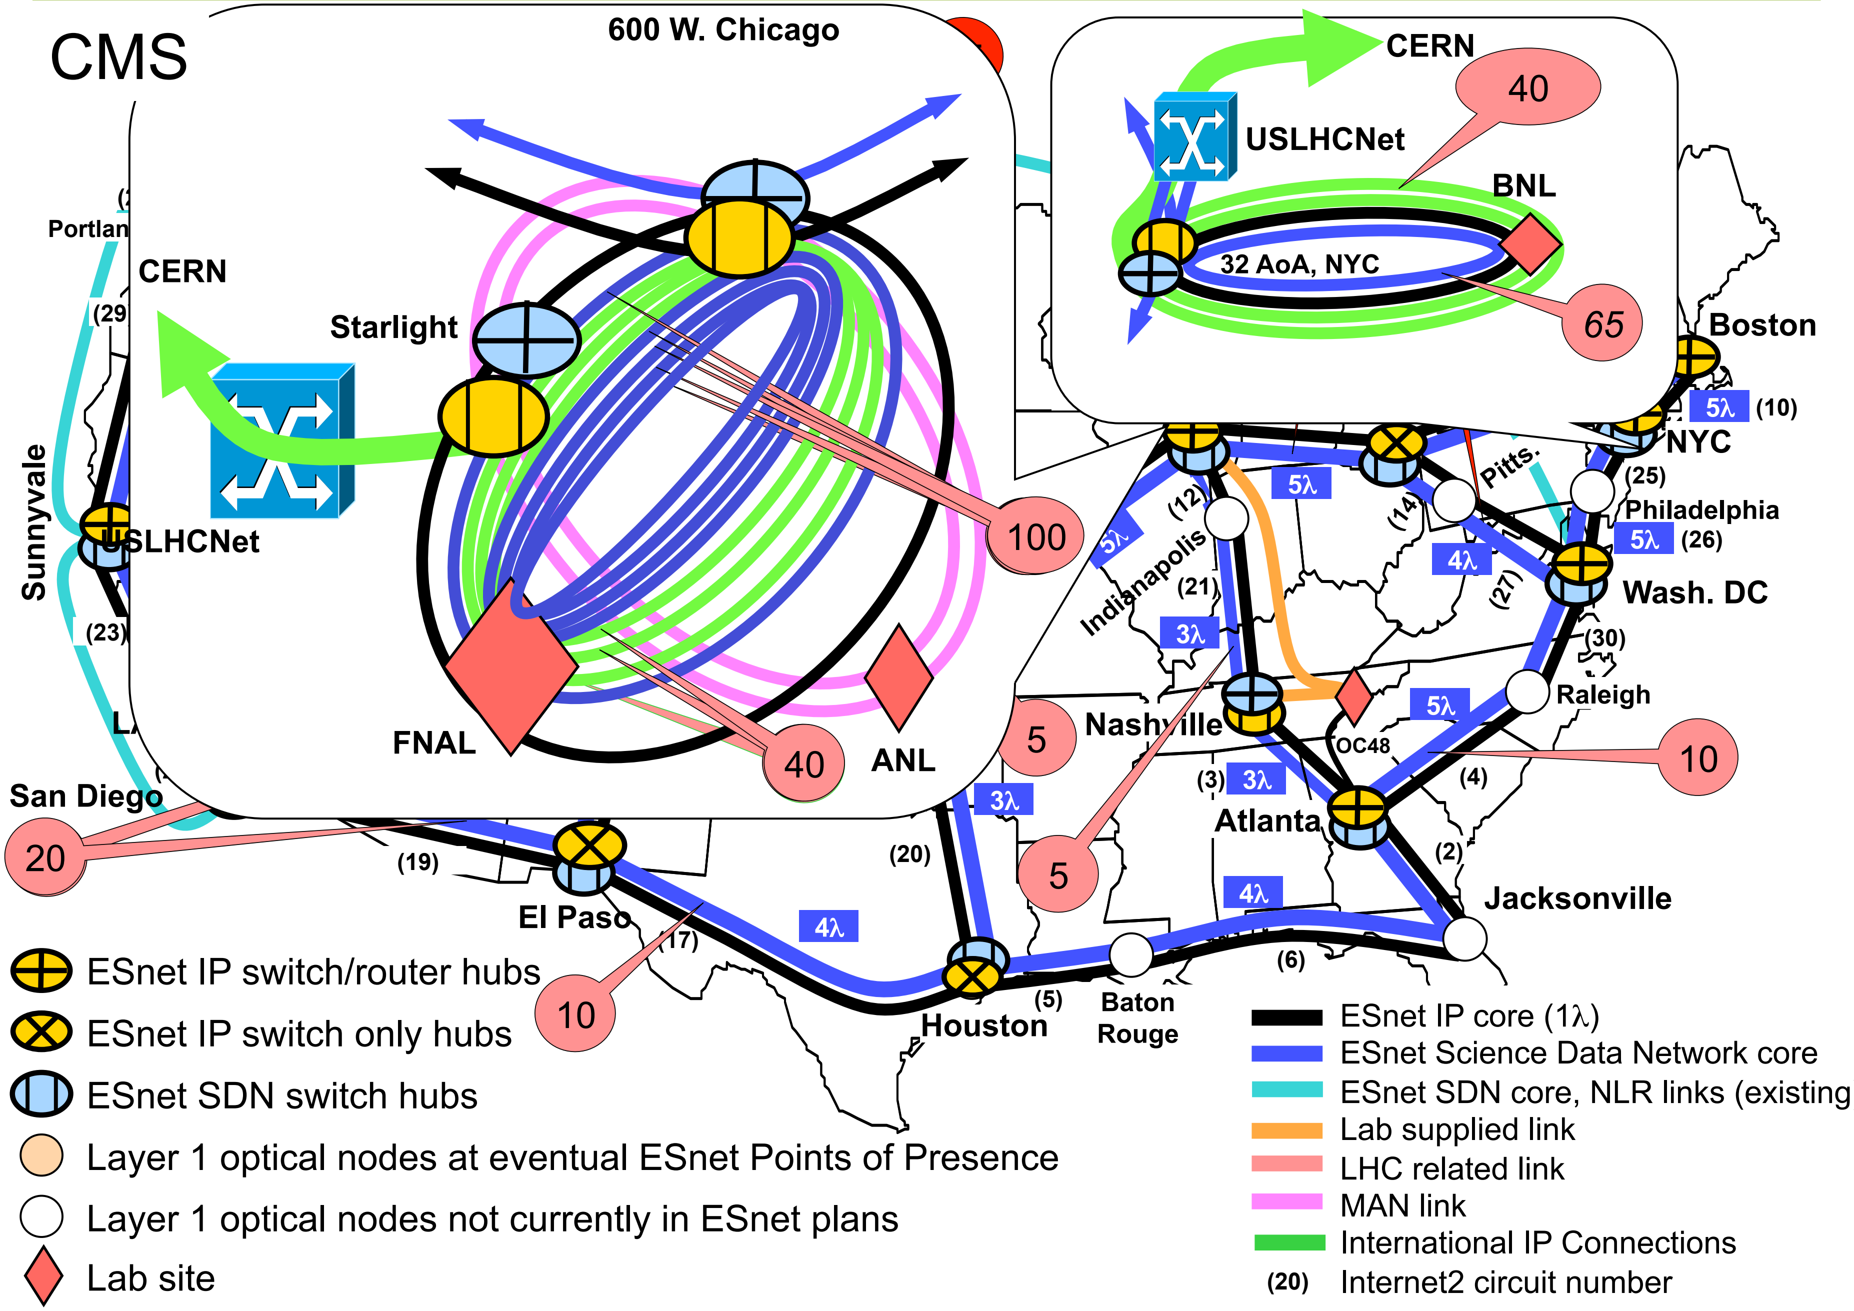

ESnet4 IP + SDN, 2011 Configuration

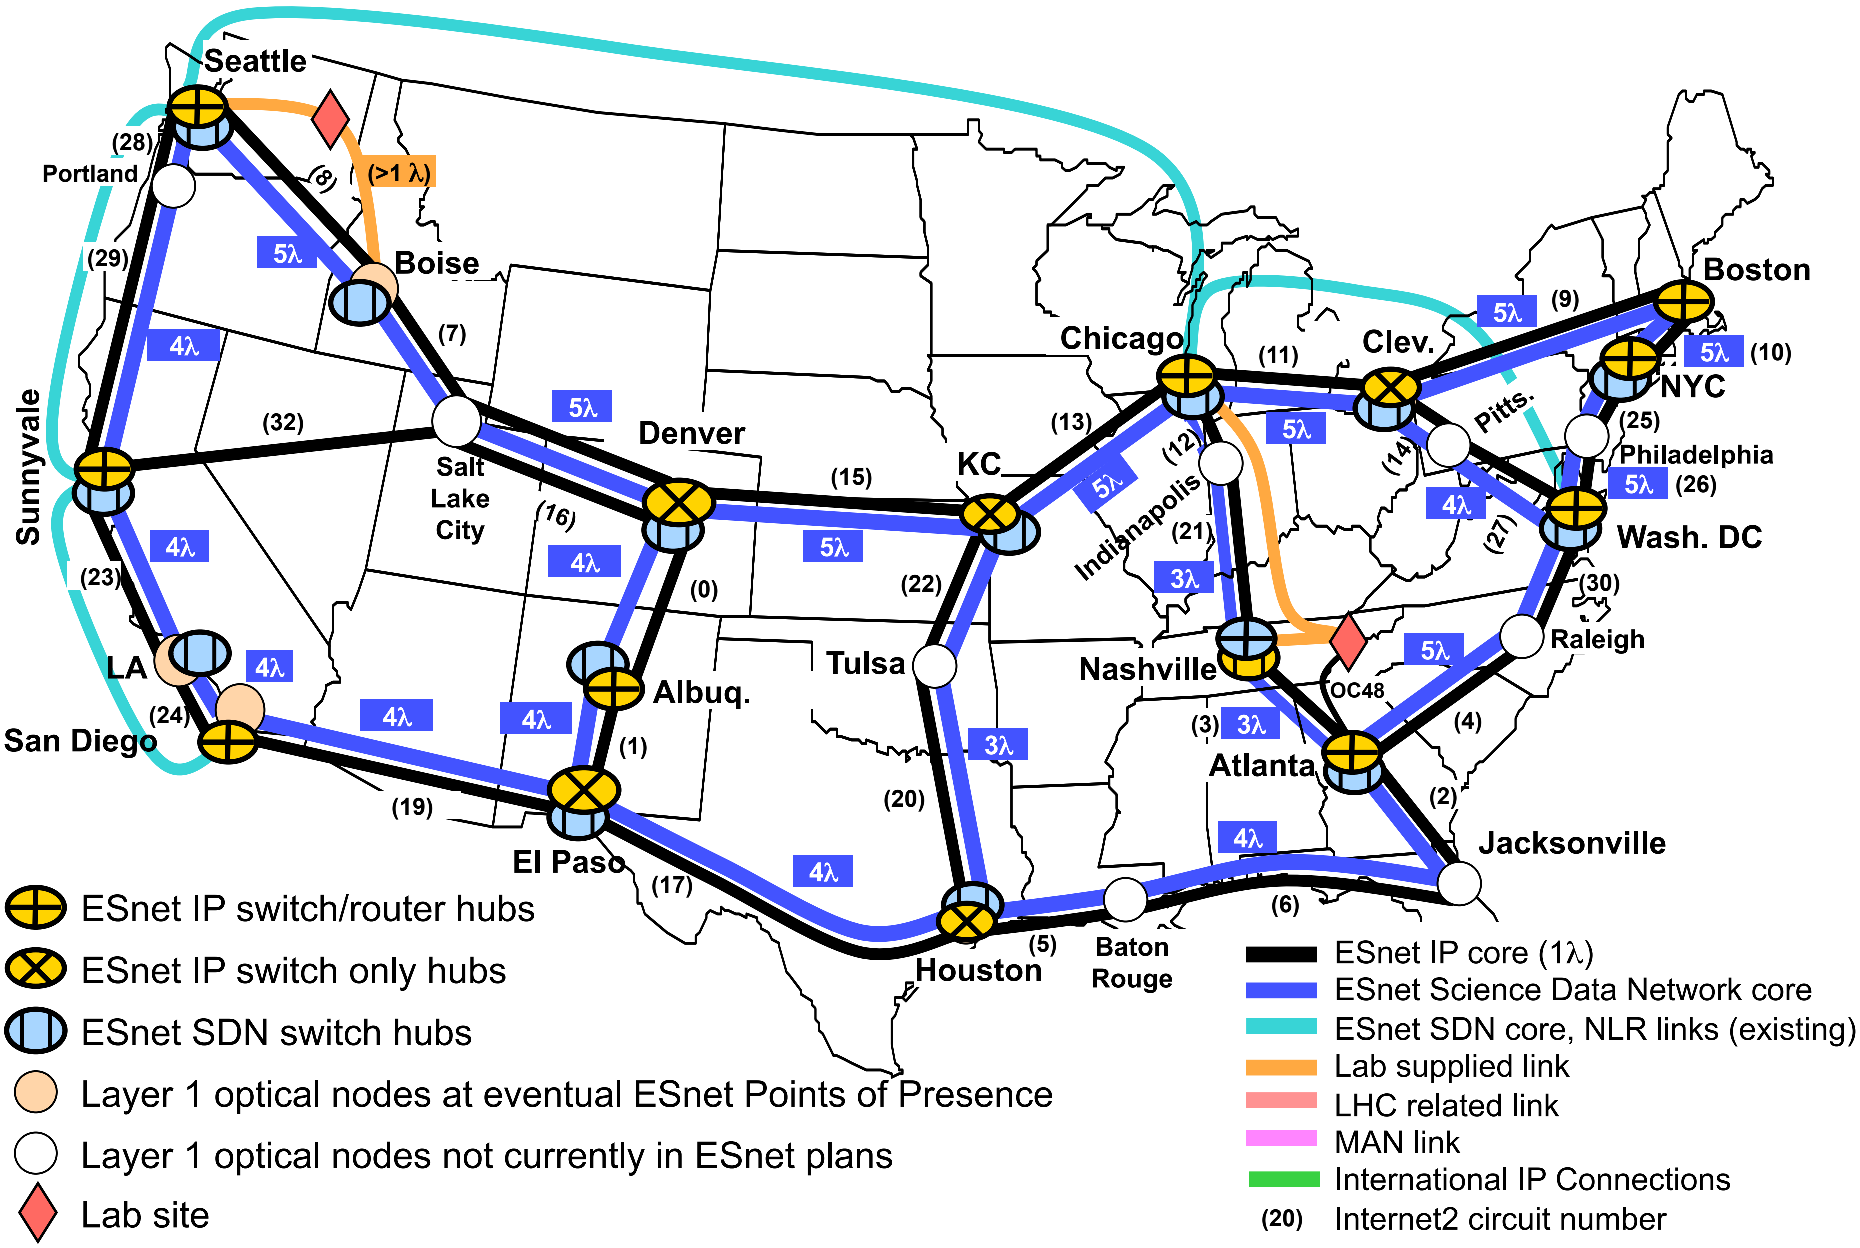

ESnet4 Planed Configuration

40-50 Gbps in 2009-2010, 160-400 Gbps in 2011-2012

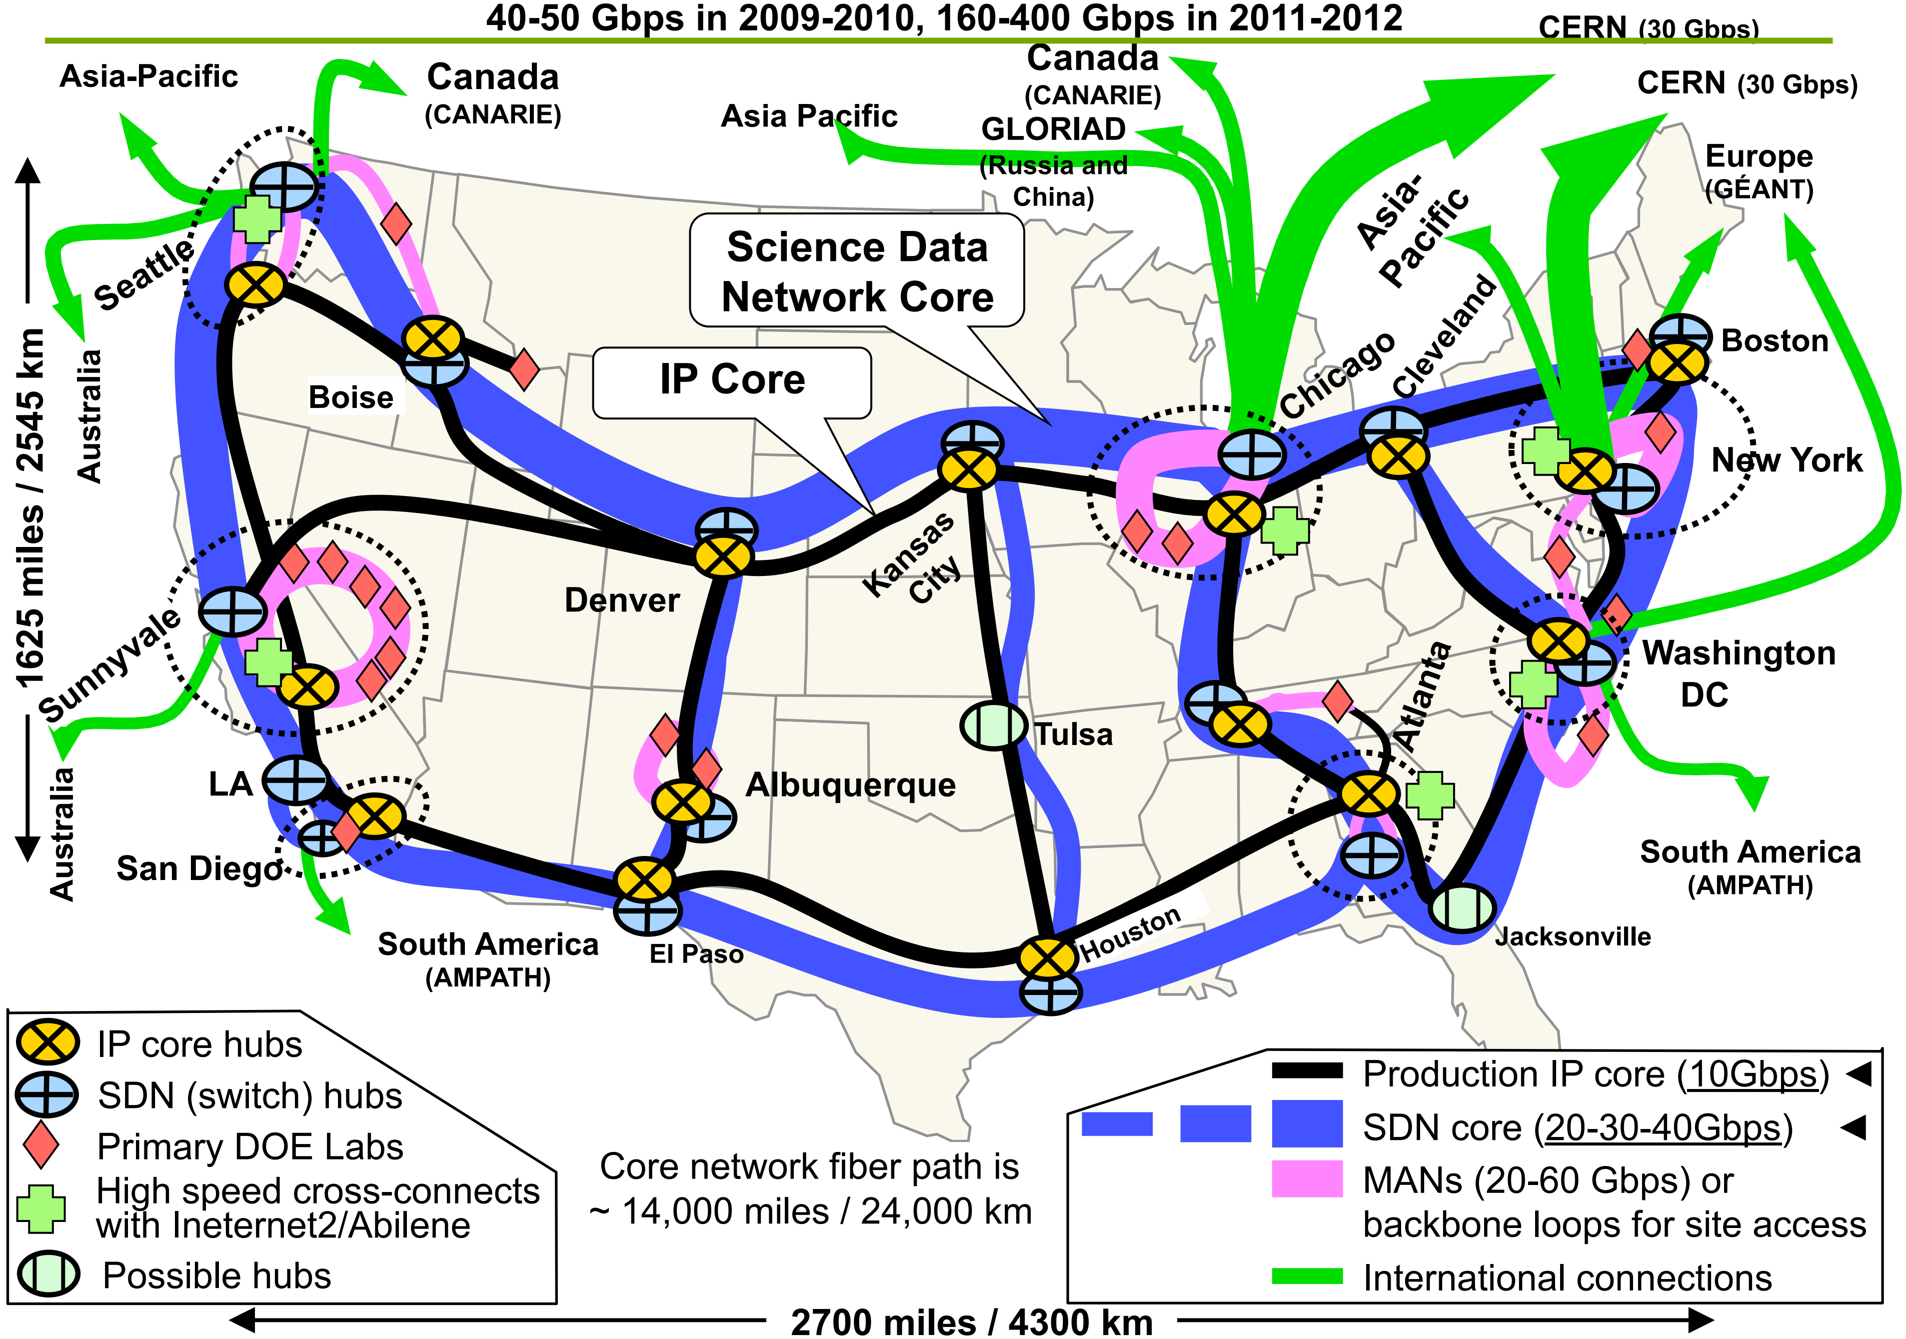

➢ Multi-Domain Virtual Circuits as a Service

- Guaranteed bandwidth service
  - User specified bandwidth - requested and managed in a Web Services framework
- Traffic isolation and traffic engineering
  - Provides for high-performance, non-standard transport mechanisms that cannot co-exist with commodity TCP-based transport
  - Enables the engineering of explicit paths to meet specific requirements
    - e.g. bypass congested links, using lower bandwidth, lower latency paths
- 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 isolated from the general traffic
- End-to-end (cross-domain) connections between Labs and collaborating institutions
- Reduced cost of handling high bandwidth data flows
  - Highly capable routers are not necessary when every packet goes to the same place
  - Use lower cost (factor of 5x) switches to relatively route the packets

Virtual Circuit Service Functional Requirements

- Support user/application VC reservation requests
  - Source and destination of the VC
  - Bandwidth, latency, start time, and duration of the VC
  - Traffic characteristics (e.g. flow specs) to identify traffic designated for the VC
- Manage allocations of scarce, shared resources
  - Authentication to prevent unauthorized access to this service
  - Authorization to enforce policy on reservation/provisioning
  - Gathering of usage data for accounting
- Provide virtual circuit setup and teardown mechanisms and security
  - Widely adopted and standard protocols (such as MPLS and GMPLS) are well understood within a single domain
  - Cross domain interoperability is the subject of ongoing, collaborative development
  - secure and-to-end connection setup is provided by the network control plane
  - accommodate heterogeneous circuit abstraction (e.g.. MPLS, GMPLS, VLANs, VCAT/LCAS)
- Enable the claiming of reservations
  - Traffic destined for the VC must be differentiated from “regular” traffic
- Enforce usage limits
  - Per VC admission control polices usage, which in turn facilitates guaranteed bandwidth
  - Consistent per-hop QoS throughout the network for transport predictability

ESnet Virtual Circuit Service: OSCARS
(On-demand Secured Circuits and Advanced Reservation System)

Software Architecture (see ref. 9)

- Web-Based User Interface (WBUI) will prompt the user for a username/password (or Grid proxy cert.) and forward it to the AAAS.
- Authentication, Authorization, and Auditing Subsystem (AAAS) will handle access, enforce policy, and generate usage records.
- Bandwidth Scheduler Subsystem (BSS) will track reservations and map the state of the network (present and future).
- Path Setup Subsystem (PSS) will setup and teardown the on-demand paths (LSPs).

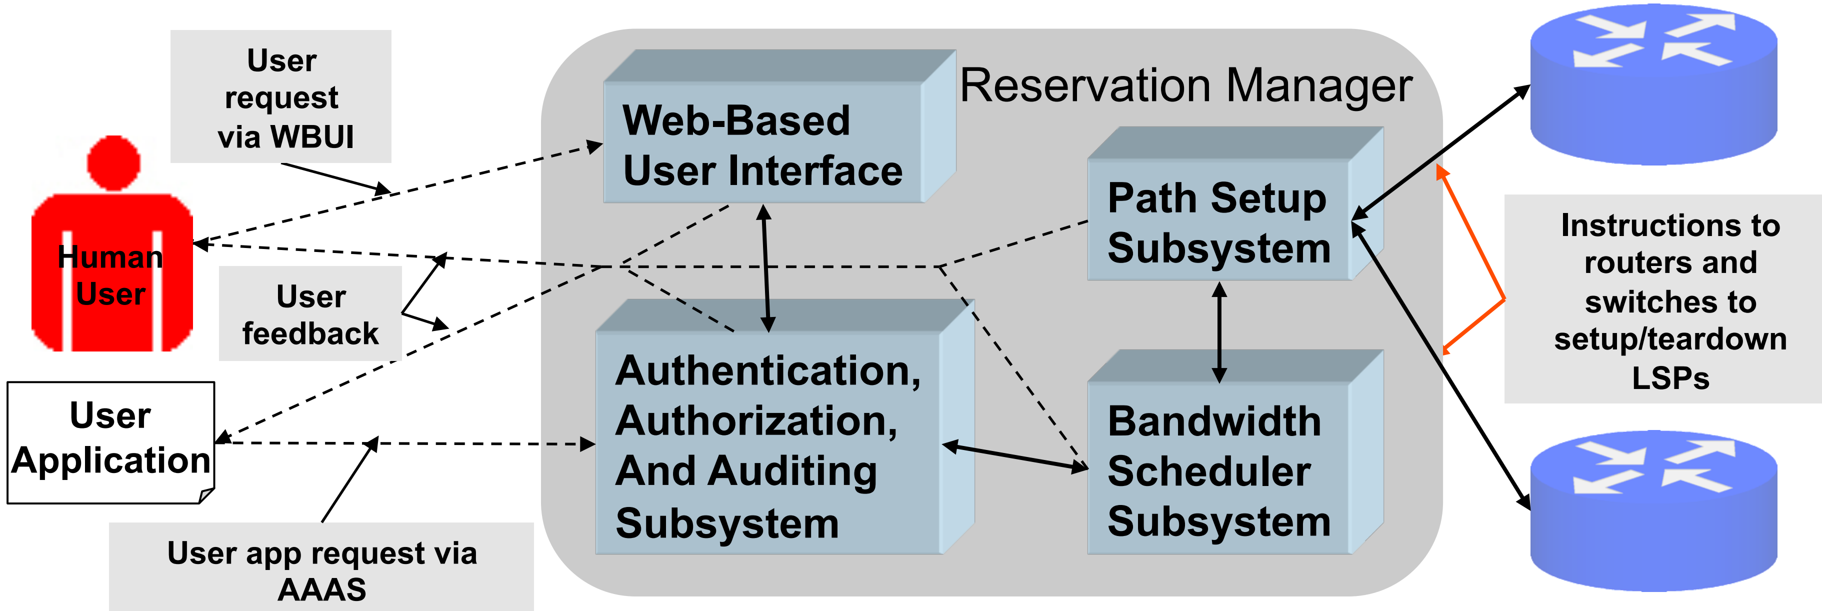

Environment of Science is Inherently Multi-Domain

- End points will be at independent institutions – campuses or research institutes - that are served by ESnet, Abilene, GÉANT, and their regional networks
  - Complex inter-domain issues – typical circuit will involve five or more domains - of necessity this involves collaboration with other networks
  - For example, a connection between FNAL and DESY involves five domains, traverses four countries, and crosses seven time zones

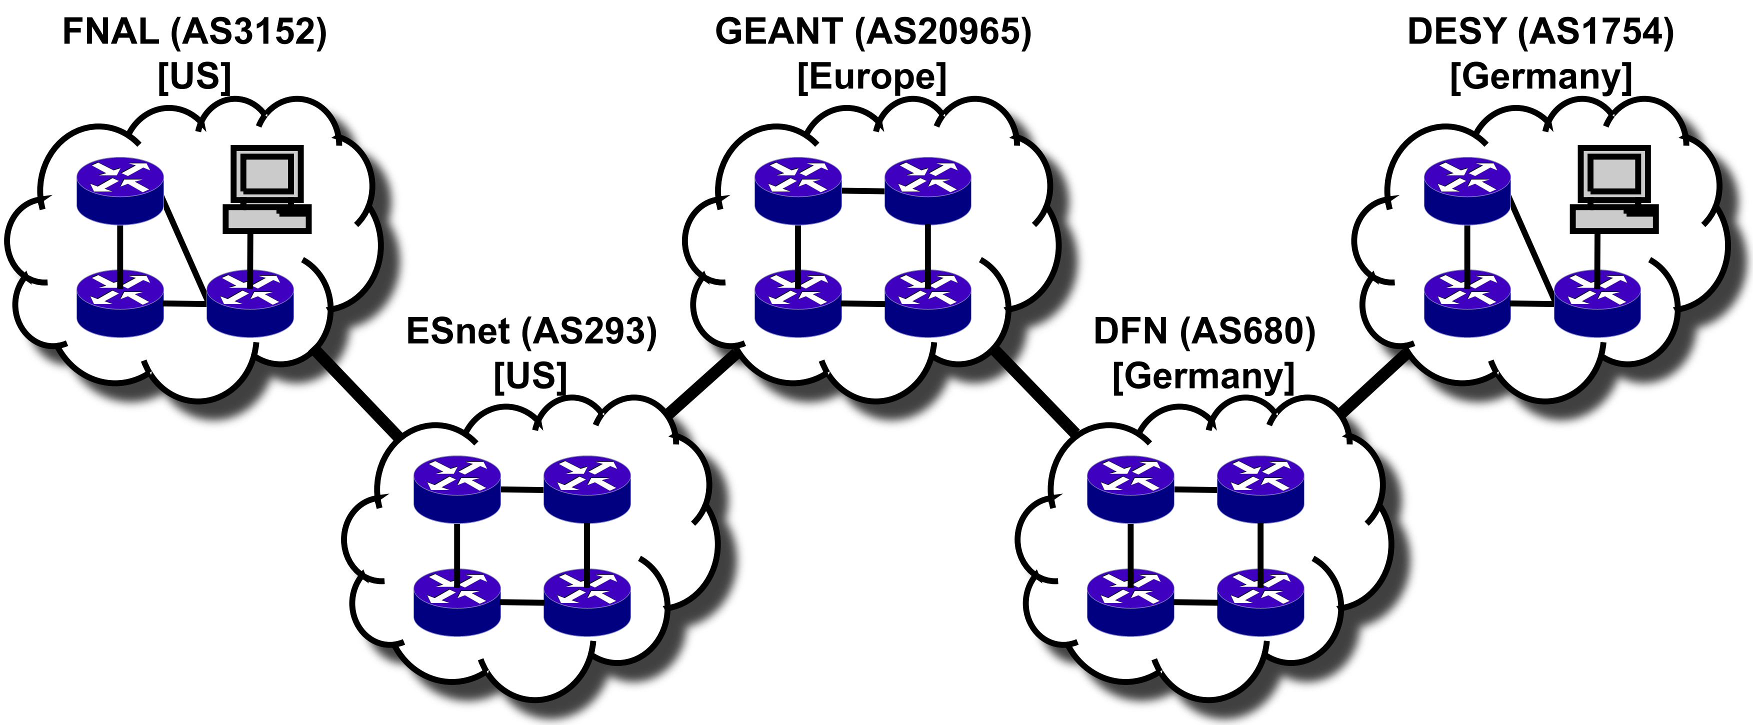

Inter-domain Reservations: A Hard Problem

- **Motivation**:
  - For a virtual circuit service to be successful, it must
    - Be end-to-end, potentially crossing several administrative domains
    - Have consistent network service guarantees throughout the circuit
- **Observation:**
  - Setting up an intra-domain circuit is easy compared with coordinating an inter-domain circuit
- **Issues:**
  - Cross domain authentication *and* authorization
    - A mechanism to authenticate and authorize a bandwidth on-demand (BoD) circuit request must be agreed upon in order to automate the process
  - Multi-domain Acceptable Use Policies (AUPs)
    - Domains may have very specific AUPs dictating what the BoD circuits can be used for and where they can transit/terminate
  - Domain specific service offerings
    - Domains must have way to guarantee a certain level of service for BoD circuits
  - Security concerns
    - Are there mechanisms for a domain to protect itself (e.g. RSVP filtering)
      - ➢ Mostly “no” - so there will not be signaling across domain boundaries

Environment of Science is Inherently Multi-Domain

- In order to set up end-to-end circuits across multiple domains without violating security or allocation management policy of any of the domains, the process of setting up end-to-end circuits is handled by a broker that
    1) Negotiates for bandwidth that is available for this purpose
    2) Requests that the domains each establish the “meet-me” state so that the circuit in one domain can connect to a circuit in another domain (thus maintaining domain control over the circuits)

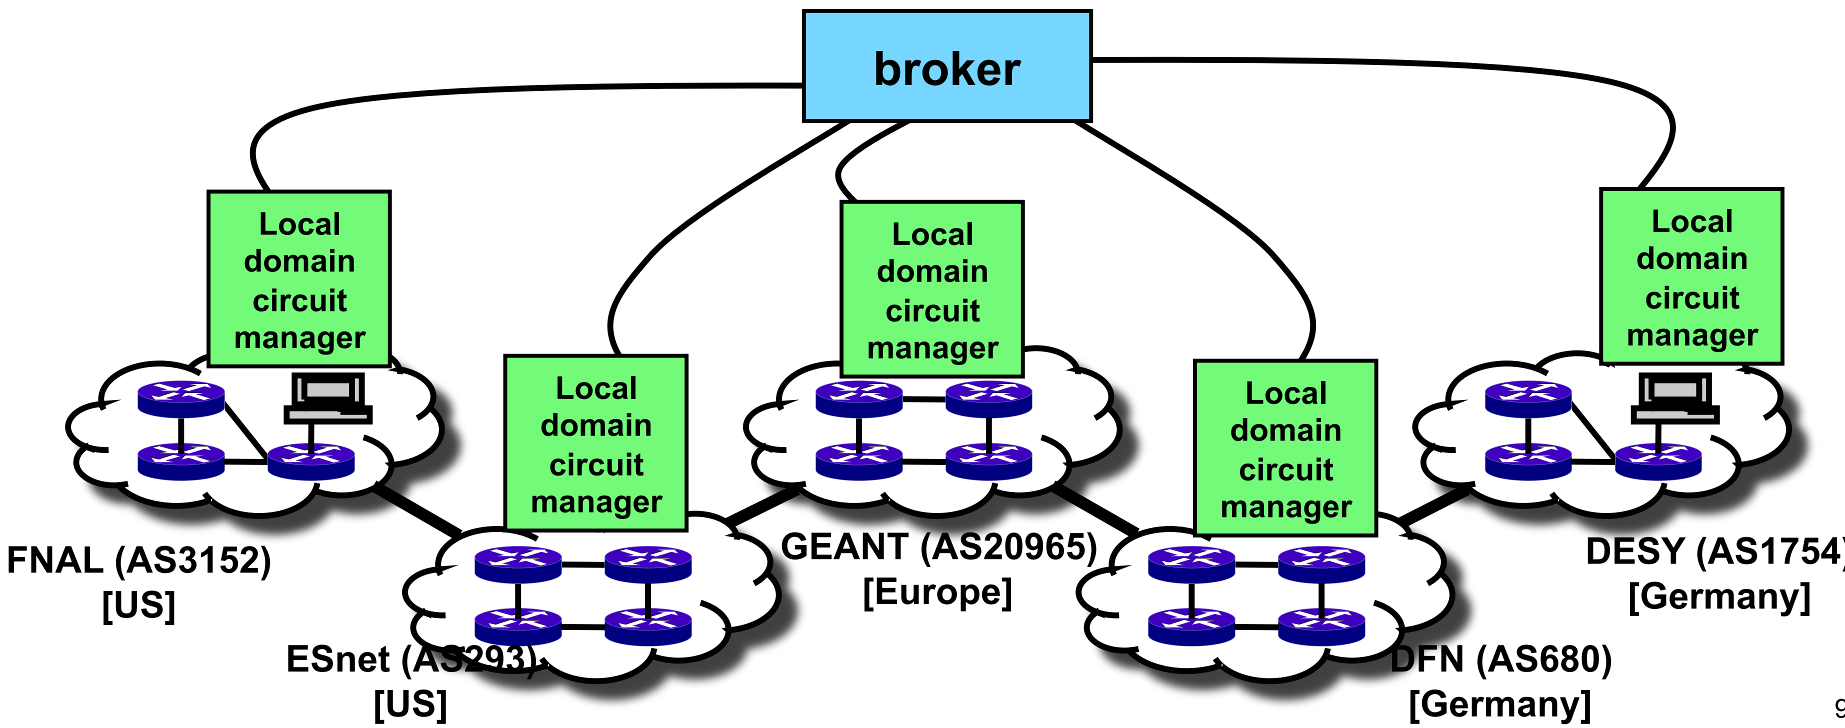

OSCARS Update

- OSCARS is the ESnet control plane mechanism
  - Establishes the circuit path
  - Transport is then via MPLS and VLANs
- Completed porting OSCARS from Perl to Java to better support web-services
  - This is now the common code base for OSCARS and I2's BRUW
- Paper on OSCARS was accepted by the IEEE GridNets
- Collaborative efforts
  - Working with I2 and DRAGON to support interoperability between OSCARS/ BRUW and DRAGON
    - currently in the process of installing an instance of DRAGON in ESnet
  - Working with I2, DRAGON, and TeraPaths (Brookhaven Lab) to determine an appropriate interoperable AAI (authentication and authorization infrastructure) framework (this is in conjunction with GEANT2's JRA5)
  - Working with DICE Control Plane group to determine schema and methods of distributing topology and reachability information
    - DICE=Internet2, ESnet, GEANT, CANARIE/UCLP; see http://www.garr.it/dice/presentation.htm for presentations from the last meeting
  - Working with Tom Lehman (DRAGON), Nagi Rao (USN), Nasir Ghani (Tennessee Tech) on multi-level, multi-domain hybrid network performance measurements (funded by OASCR)

➢ Monitoring as a Service-Oriented Communications Service

- perfSONAR is a community effort to define network management data exchange protocols, and standardized measurement data gathering and archiving
- Path performance monitoring is an example of a perfSONAR application
  - provide users/applications with the end-to-end, multi-domain traffic and bandwidth availability
  - provide real-time performance such as path utilization and/or packet drop
- Multi-domain path performance monitoring tools are in development based on perfSONAR protocols and infrastructure
  - One example – Traceroute Visualizer [TrViz] – has been deployed in about 10 R&E networks in the US and Europe that have deployed at least some of the required perfSONAR measurement archives to support the tool

Traceroute Visualizer

- Forward direction bandwidth utilization on application path from LBNL to INFN-Frascati (Italy)
  - traffic shown as bars on those network device interfaces that have an associated MP services (the first 4 graphs are normalized to 2000 Mb/s, the last to 500 Mb/s)

1 ir1000gw (131.243.2.1)
2 er1kgw
3 lbl2-ge-lbnl.es.net

link capacity is also provided

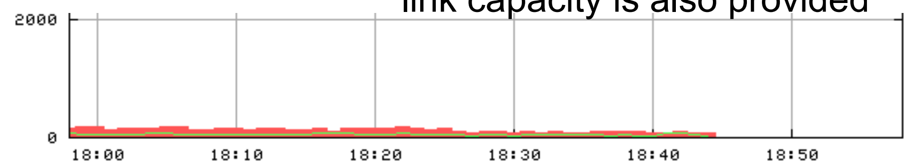

4 slacmr1-sdn-lblmr1.es.net (GRAPH OMITTED)
5 snv2mr1-slacmr1.es.net (GRAPH OMITTED)
6 snv2sdn1-snv2mr1.es.net

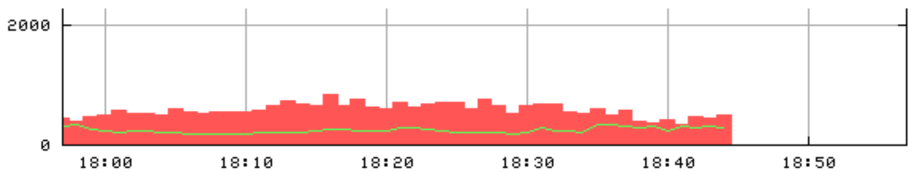

7 chislsdn1-oc192-snv2sdn1.es.net (GRAPH OMITTED)
8 chiccr1-chislsdn1.es.net

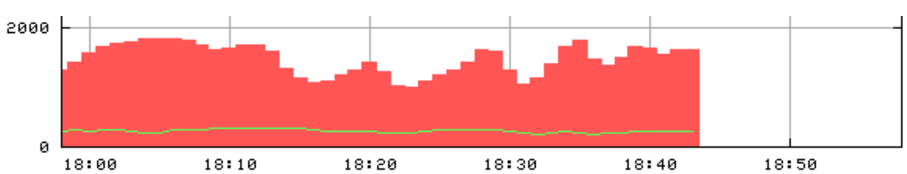

9 aofacr1-chicsdn1.es.net (GRAPH OMITTED)

10 esnet.rt1.nyc.us.geant2.net (NO DATA)
11 so-7-0-0.rt1.ams.nl.geant2.net (NO DATA)
12 so-6-2-0.rt1.fra.de.geant2.net (NO DATA)
13 so-6-2-0.rt1.gen.ch.geant2.net (NO DATA)
14 so-2-0-0.rt1.mil.it.geant2.net (NO DATA)
15 garr-gw.rt1.mil.it.geant2.net (NO DATA)
16 rt1-mi1-rt-mi2.mi2.garr.net

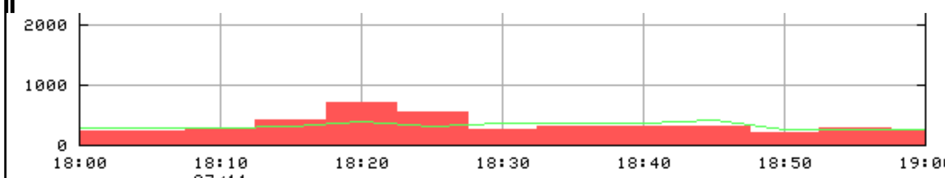

17 rt-mi2-rt-rm2.rm2.garr.net (GRAPH OMITTED)
18 rt-rm2-rc-fra.fra.garr.net (GRAPH OMITTED)
19 rc-fra-ru-lnf.fra.garr.net (GRAPH OMITTED)

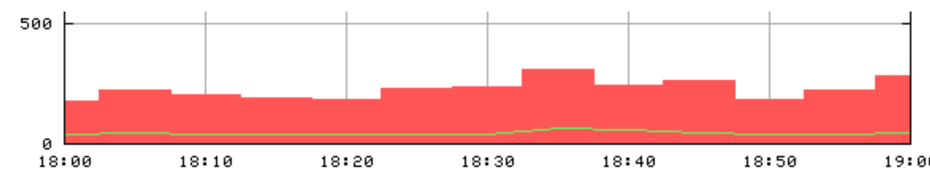

20
21 www6.lnf.infn.it (193.206.84.223) 189.908 ms 189.596 ms 189.684 ms

➢ 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

Federated Trust Services – Support for Large-Scale Collaboration

- CAs are provided with different policies as required by the science community
  - DOEGrids CA has a policy tailored to accommodate international science collaboration
  - NERSC CA policy integrates CA and certificate issuance with NIM (NERSC user accounts management services)
  - FusionGrid CA supports the FusionGrid roaming authentication and authorization services, providing complete key lifecycle management

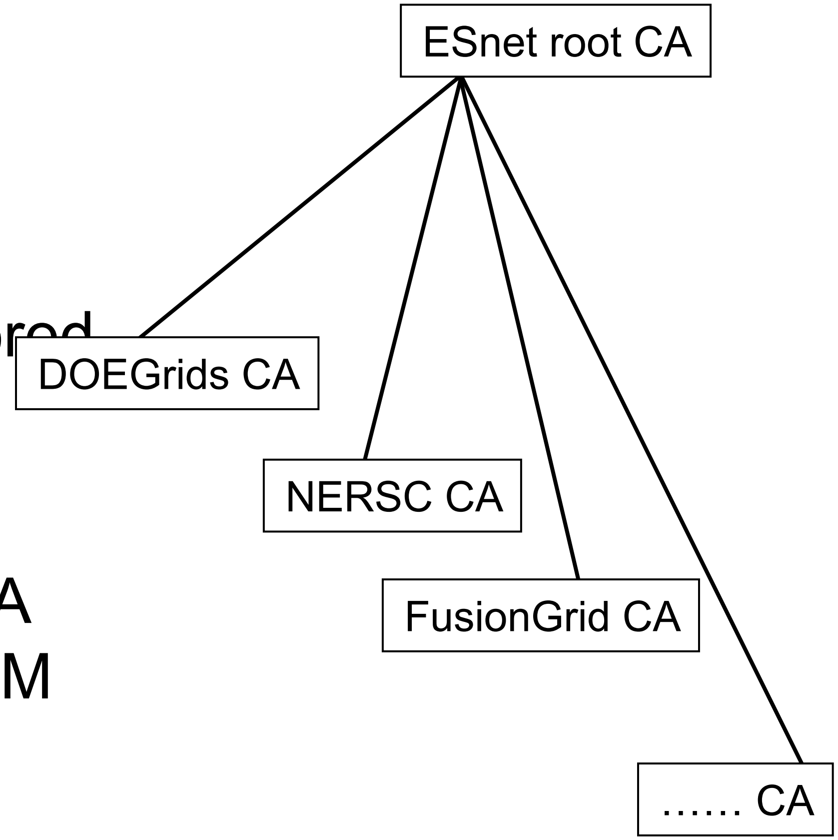

See www.doegrids.org

DOEGrids CA (Active Certificates) Usage Statistics

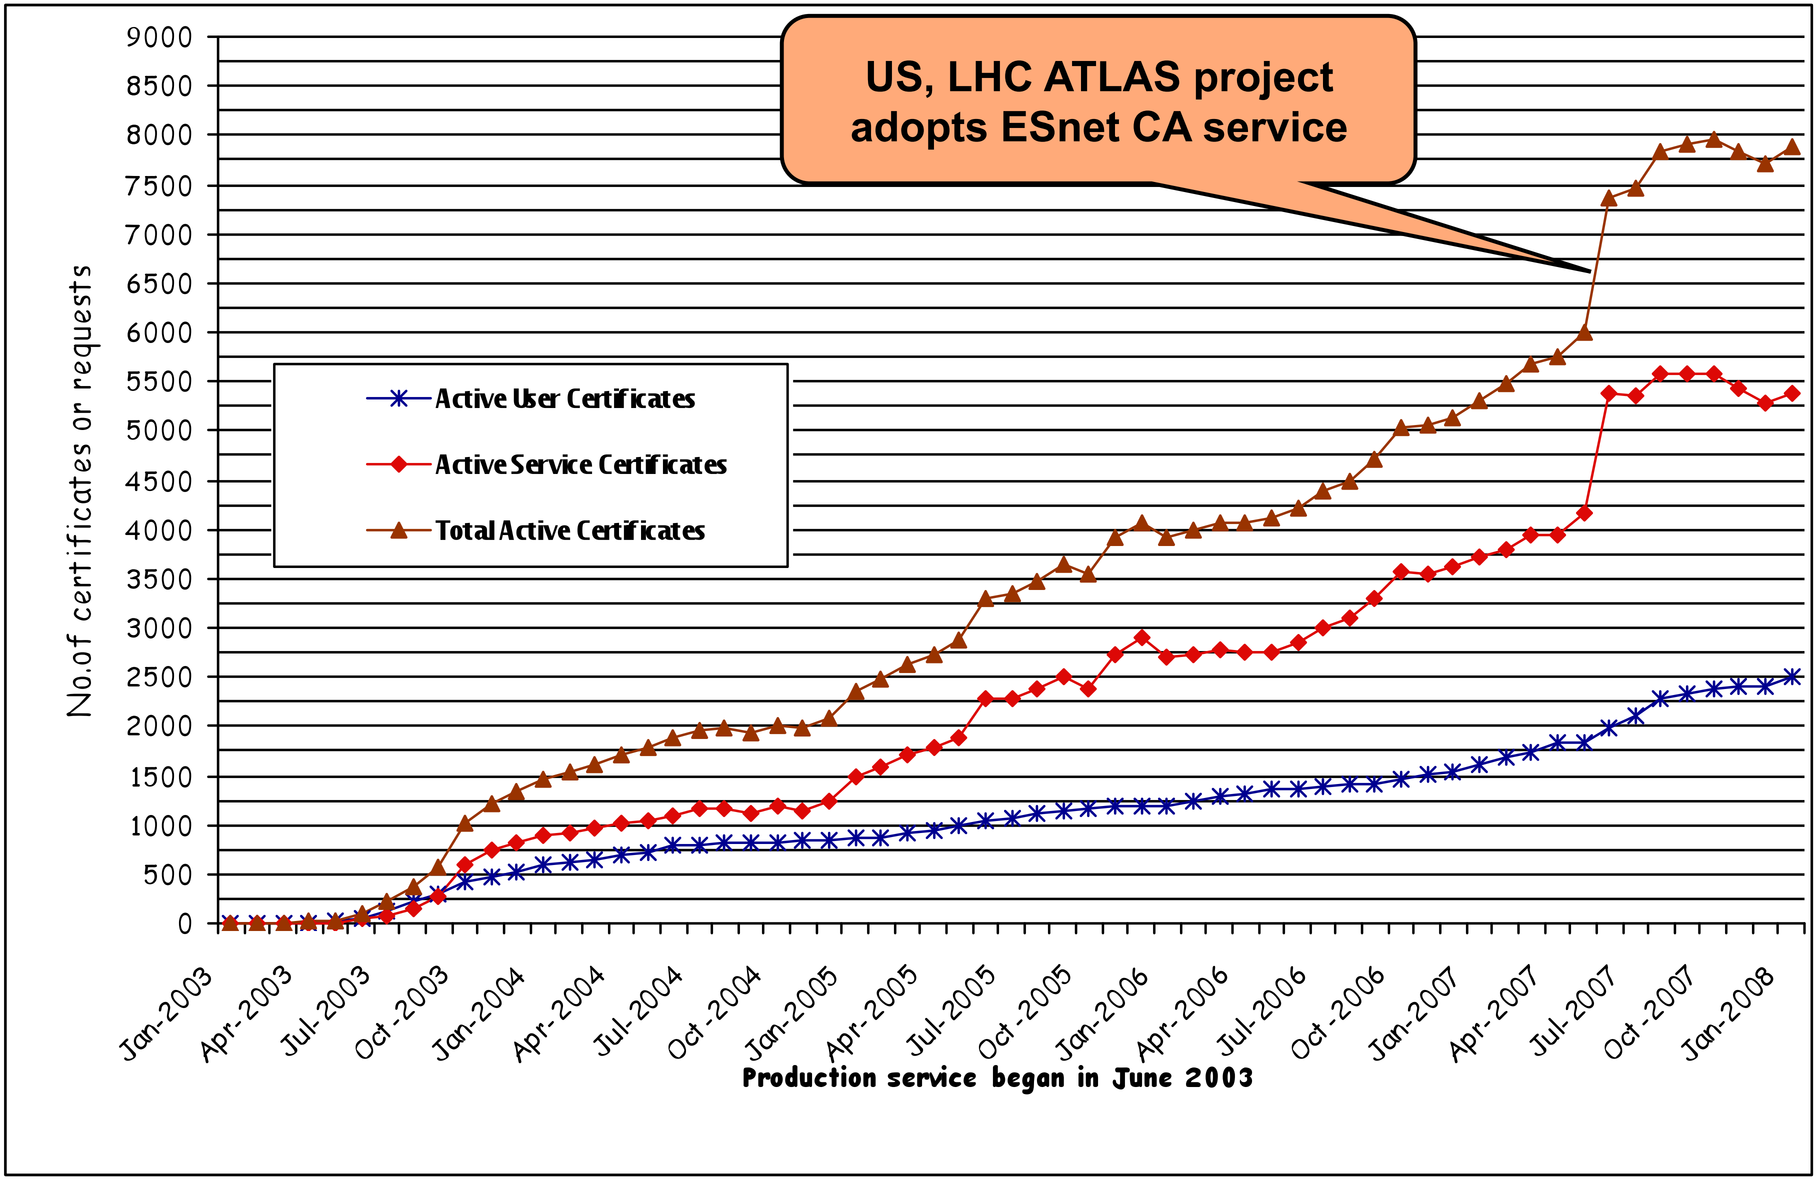

* Report as of Jan 17, 2008

DOEGrids CA Usage - Virtual Organization Breakdown

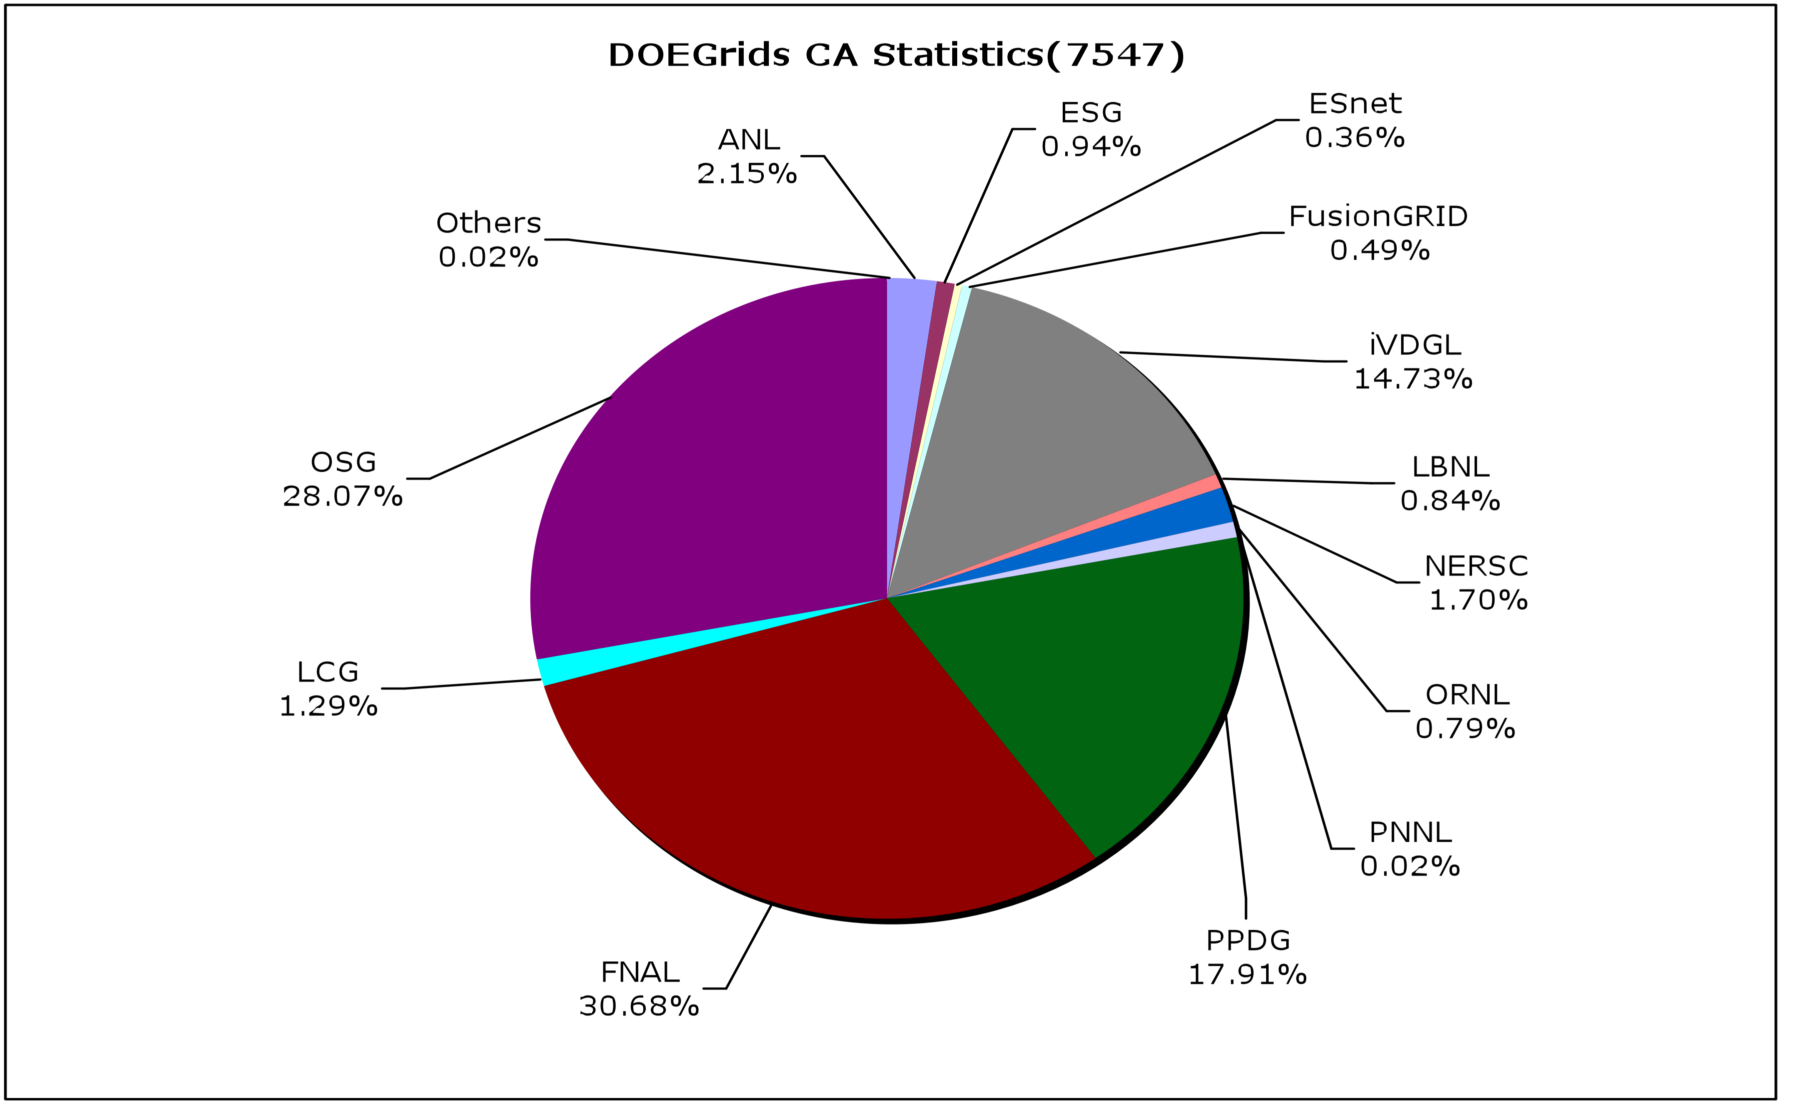

* DOE-NSF collab. & Auto renewals

** 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)

➢ 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)

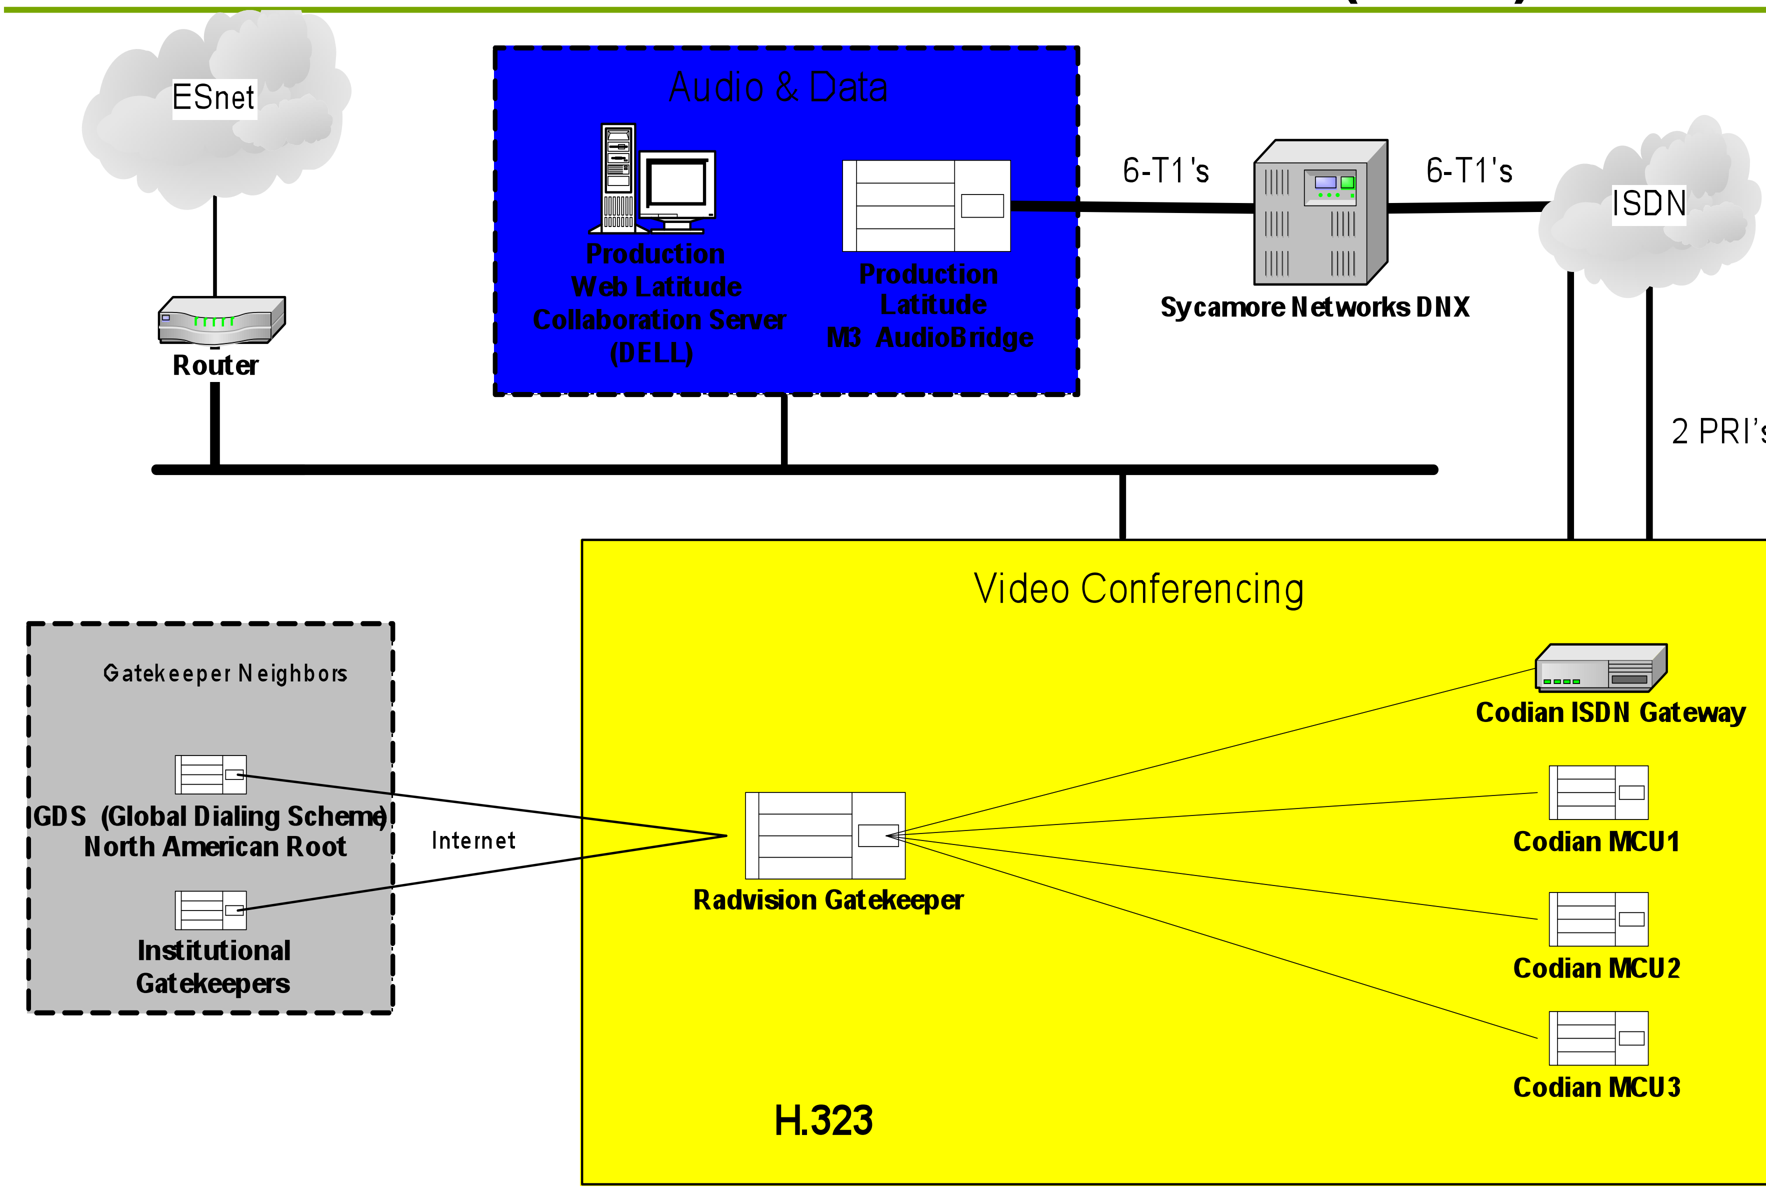

ECS Video Collaboration Service

- 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 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

Typical Problems Reported to ECS Support

Video Conferencing

- User E.164 look up
- Gatekeeper registration problems – forgotten IP address or user network problems
- Gateway Capacity for ISDN service expanded to 2 full PRI's = 46 x 64kbps chs
- For the most part, problems are with user-side network and systems configuration.

Voice and Data Collaboration

- Scheduling Conflicts and Scheduling Capacity has been addressed by expanding overbooking capacity to 100% of actual capacity
    - Future equipment plans will allow for optimal configuration of scheduling parameters.
- Browser compatibility with Java based data sharing client – users are advised to test before meetings
- Lost UserID and/or passwords

Words of Wisdom

We advise users that at least two actions must be taken in advance of conferences to reduce the likelihood of problems:

A) testing of the configuration to be used for the audio, video and data conference.

B) appropriate setup time must be allocated BEFORE the conference to ensure punctuality and correct local configuration. (at least 15 min recommended)

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

- Proposed solution
  - A fee-for-service arrangement for real-time conference support
  - Such an arrangement could be made by ***contracting directly with TKO Video Communications, ESnet's ECS service provider***
  - Service offering would 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 at $xxx per hour (Commercial Price: $yyy/hr)

Evolution of ECS

- Video
  - Increased ports and redundant MCU at geographically diverse location
- Audio
  - Maybe outsource the whole service
  - New Federal telecom contract (Networx - successor to FTS2000) has the potential for providing much more capacity at substantially less per port cost than current system
    - Testing for functionality
    - Not clear that appropriate management and fiduciary controls are possible

➢Summary

- **ESnet is currently satisfying its mission by enabling SC science that is dependant on networking and distributed, large-scale collaboration**

References

1. High Performance Network Planning Workshop, August 2002
   – http://www.doecollaboratory.org/meetings/hpnpw
2. Science Case Studies Update, 2006 (contact eli@es.net)
3. DOE Science Networking Roadmap Meeting, June 2003
   – http://www.es.net/hypertext/welcome/pr/Roadmap/index.html
4. DOE Workshop on Ultra High-Speed Transport Protocols and Network Provisioning for Large-Scale Science Applications, April 2003
   – http://www.csm.ornl.gov/ghpn/wk2003
5. Science Case for Large Scale Simulation, June 2003
   – http://www.pnl.gov/scales/
6. Workshop on the Road Map for the Revitalization of High End Computing, June 2003
   – http://www.cra.org/Activities/workshops/nitrd
   – http://www.sc.doe.gov/ascr/20040510_hecrtf.pdf (public report)
7. ASCR Strategic Planning Workshop, July 2003
   – http://www.fp-mcs.anl.gov/ascr-july03spw
8. Planning Workshops-Office of Science Data-Management Strategy, March & May 2004
   – http://www-conf.slac.stanford.edu/dmw2004

17. 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