



ESnet4

IP Network and Science Data Network

JET Presentation

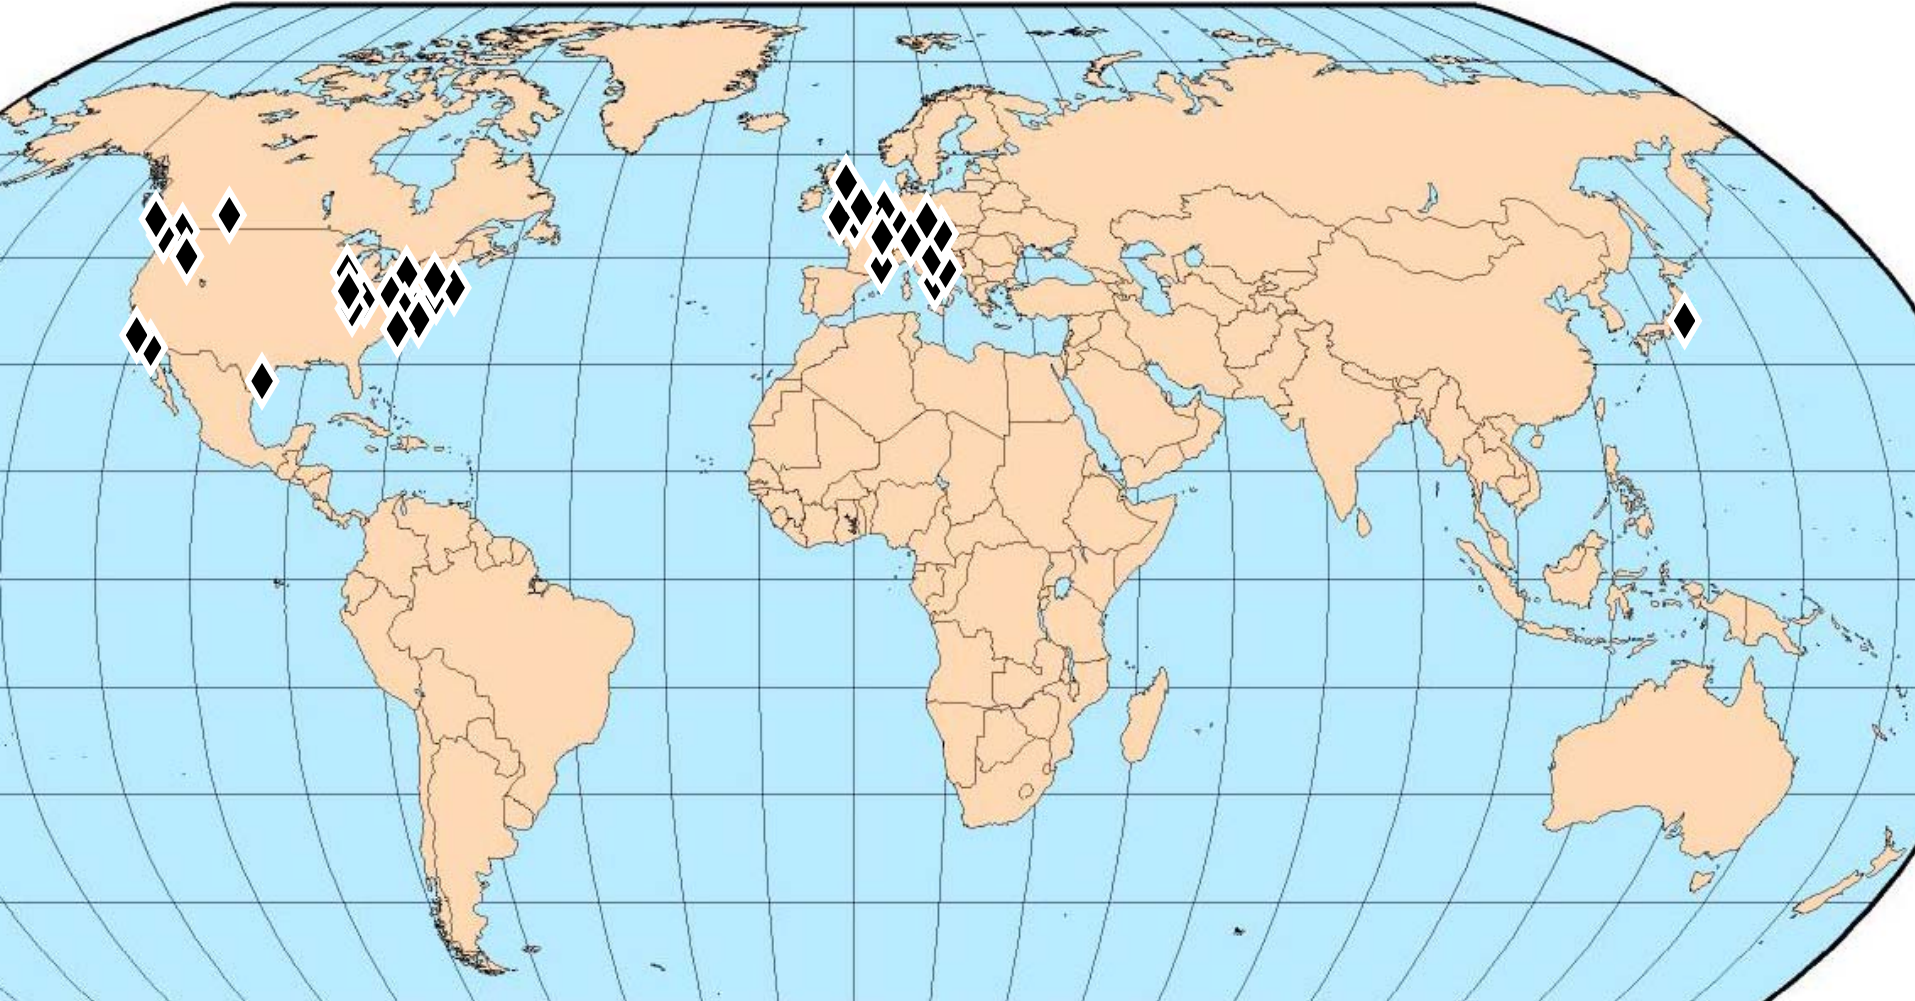
10/17/2006

For more information contact
William E. Johnston (wej@es.net), ESnet Dept. Head,
or
Joe Burrencia (joeb@es.net), ESnet General Manager

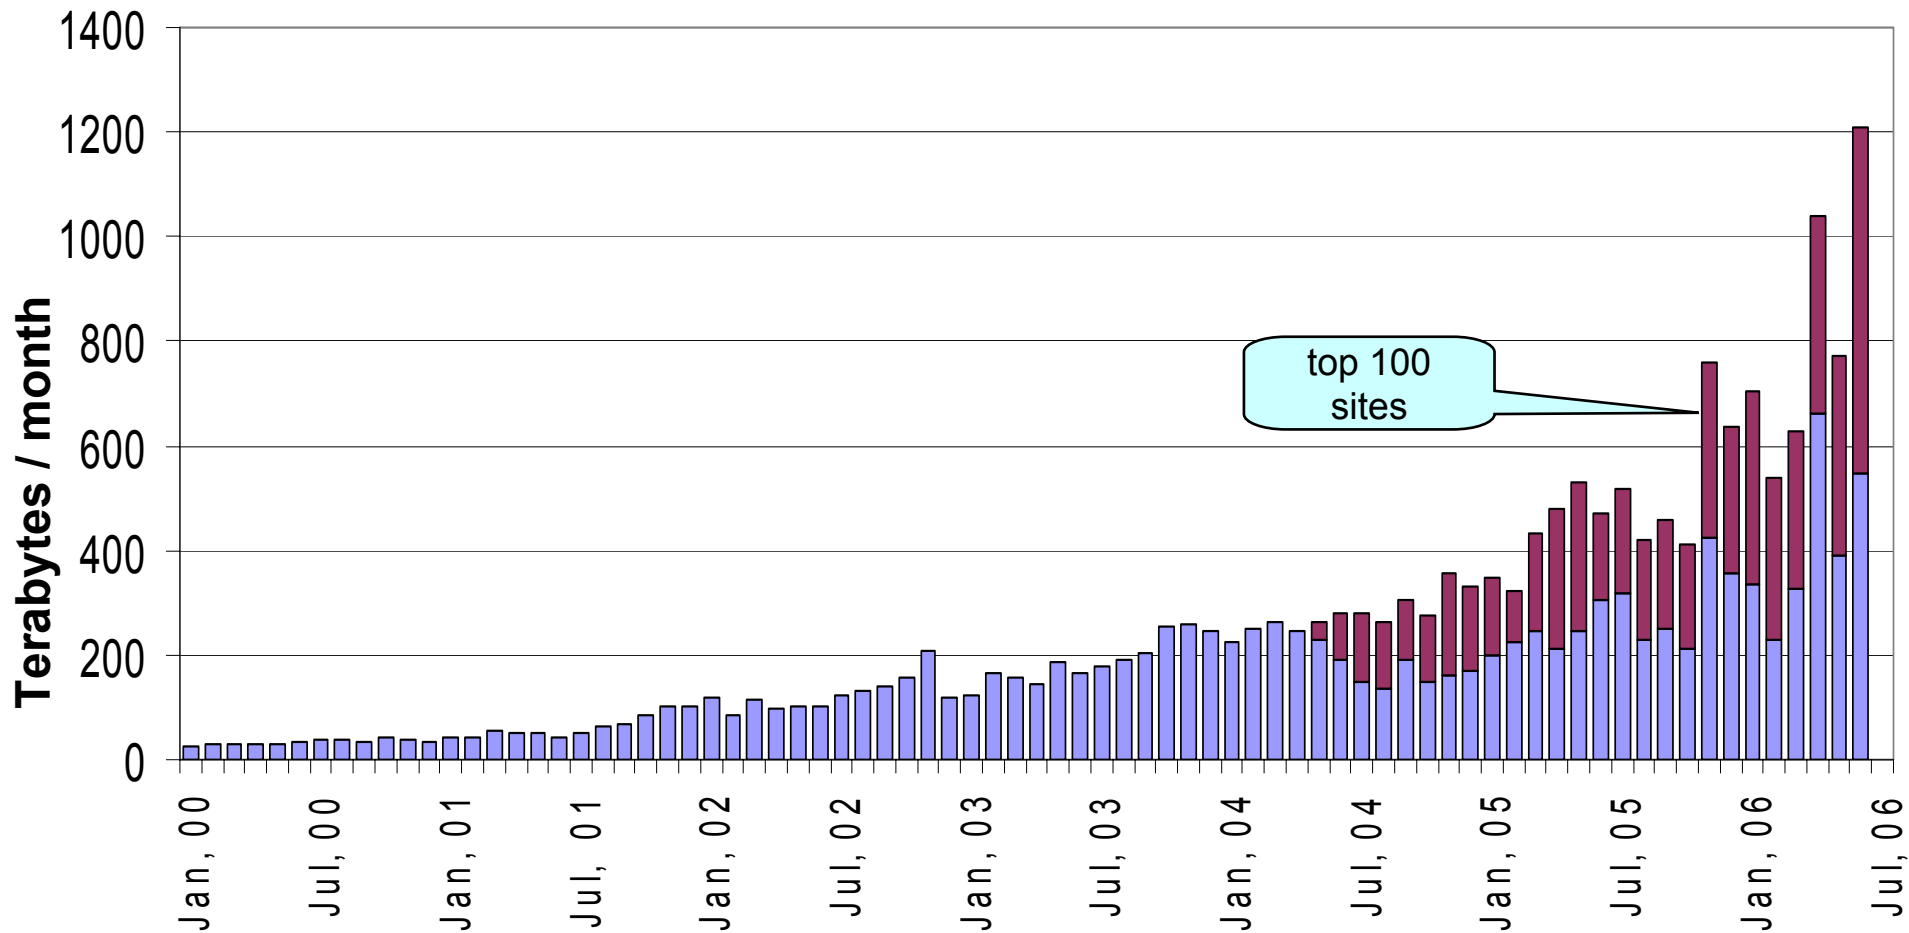


- **ESnet's primary mission is to enable the large-scale science that is the mission of DOE's Office of Science (SC):**
 - Sharing of massive amounts of data
 - Supporting thousands of collaborators world-wide
 - Distributed data processing
 - Distributed data management
 - Distributed simulation, visualization, and computational steering
- ESnet provides network and collaboration services to Office of Science laboratories and to sites of other DOE programs in cases where this increases cost effectiveness

Footprint of Largest SC Data Sharing Collaborators (50% of ESnet traffic) Drives the Footprint that ESnet Must Support



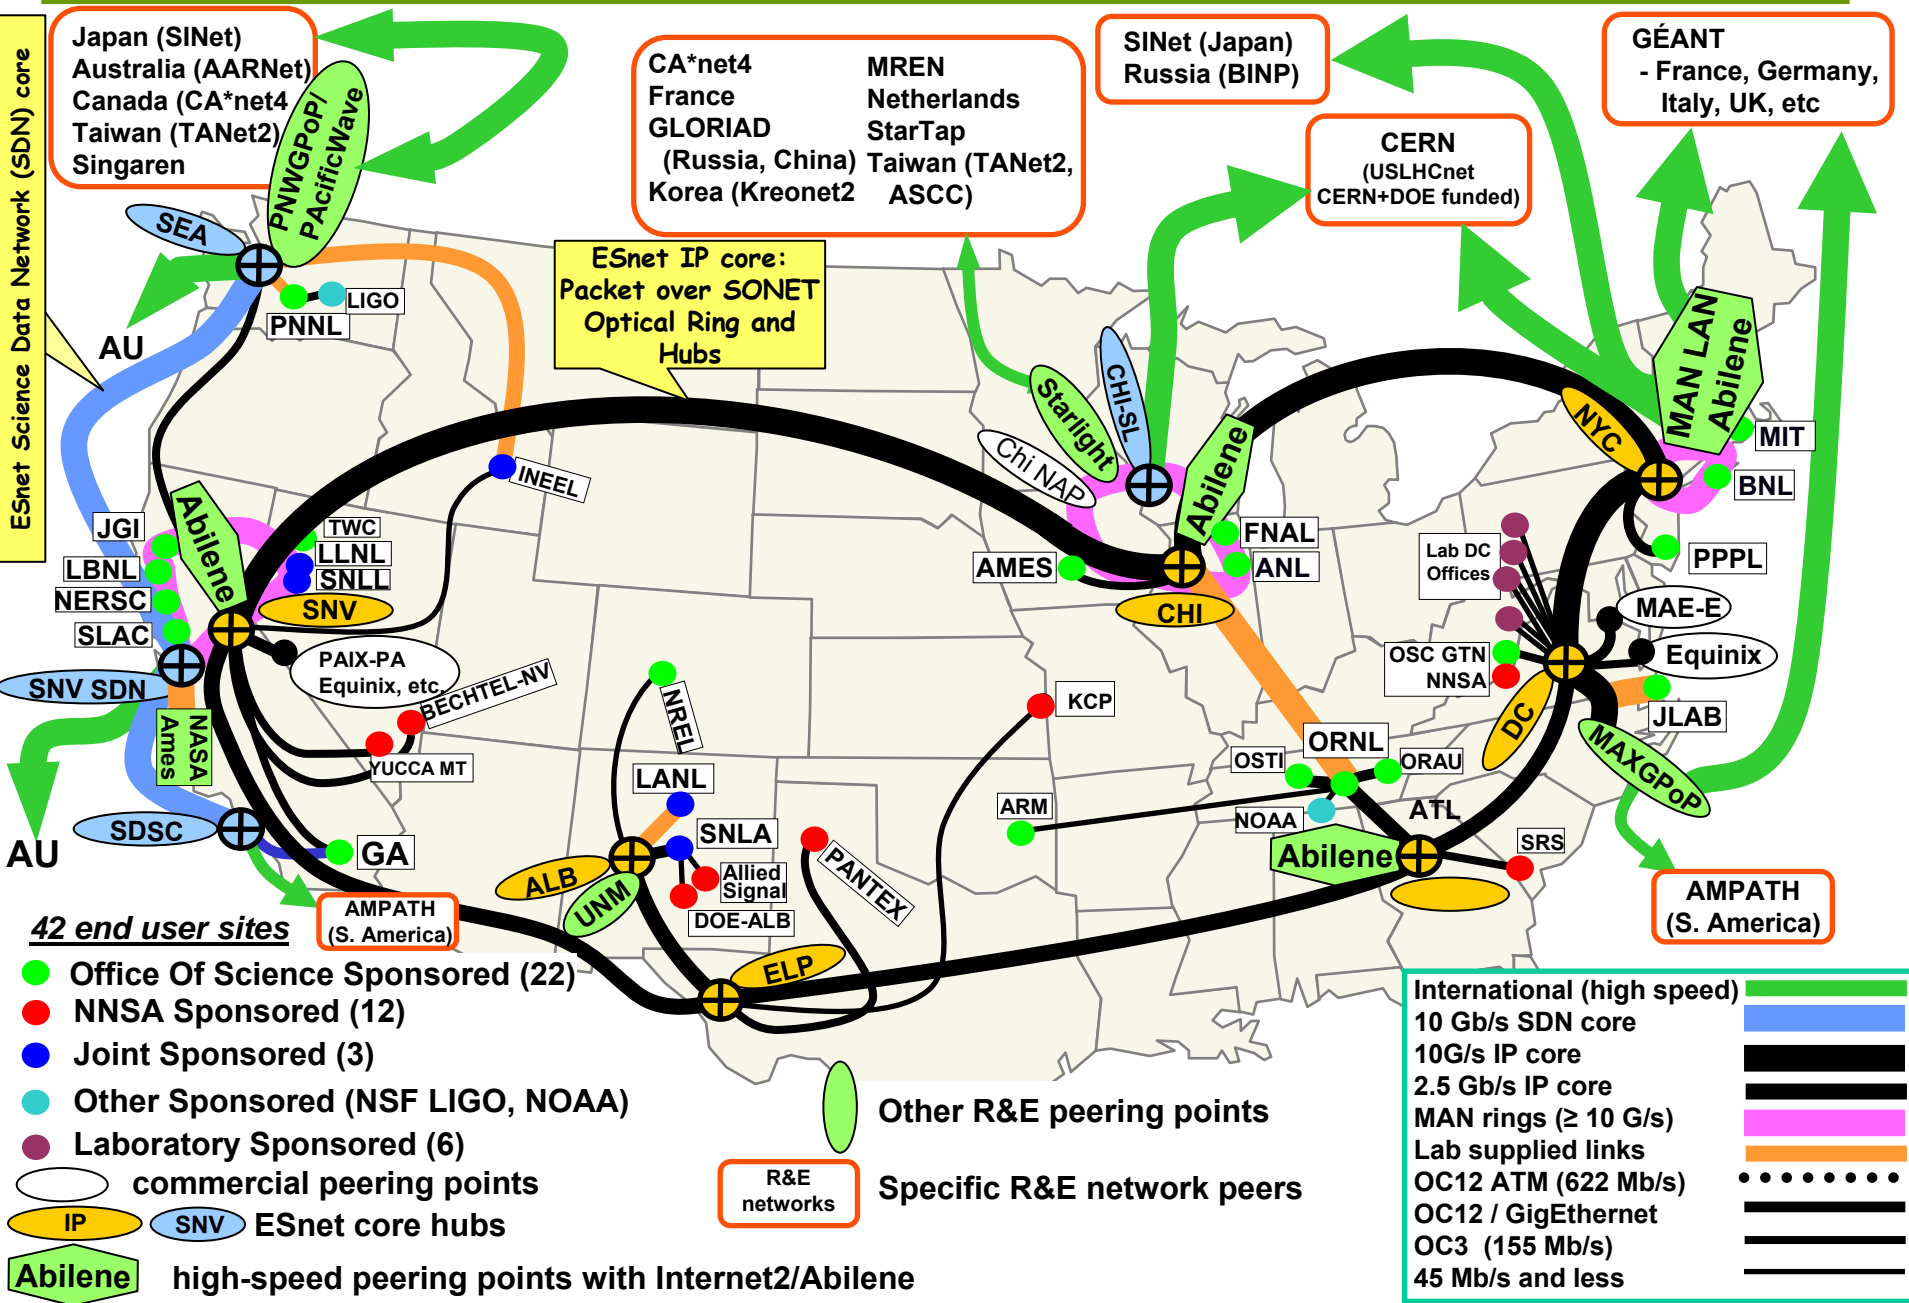
Evolution of ESnet Traffic Patterns




ESnet Monthly Accepted Traffic, January, 2000 – June, 2006

- ESnet is currently transporting more than 1 petabyte (1000 terabytes) per month
- More than 50% of the traffic is now generated by the top 100 work flows (system to system)

ESnet3 Today (Summer, 2006) Provides Global High-Speed Internet Connectivity for DOE Facilities and Collaborators



The Evolution of ESnet Architecture



ESnet to 2005:

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

ESnet from 2006-07:

- 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

- ◆ ESnet sites
- ⊕ ESnet hubs / core network connection points
- Metro area rings (MANs)
- ☁ Other IP networks
- ➡ Circuit connections to other science networks (e.g. USLHCNet)

- Internet2 has partnered with Level 3 Communications Co. 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 provide 10 optical circuits (λ s) across the entire fiber footprint (80 λ s 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

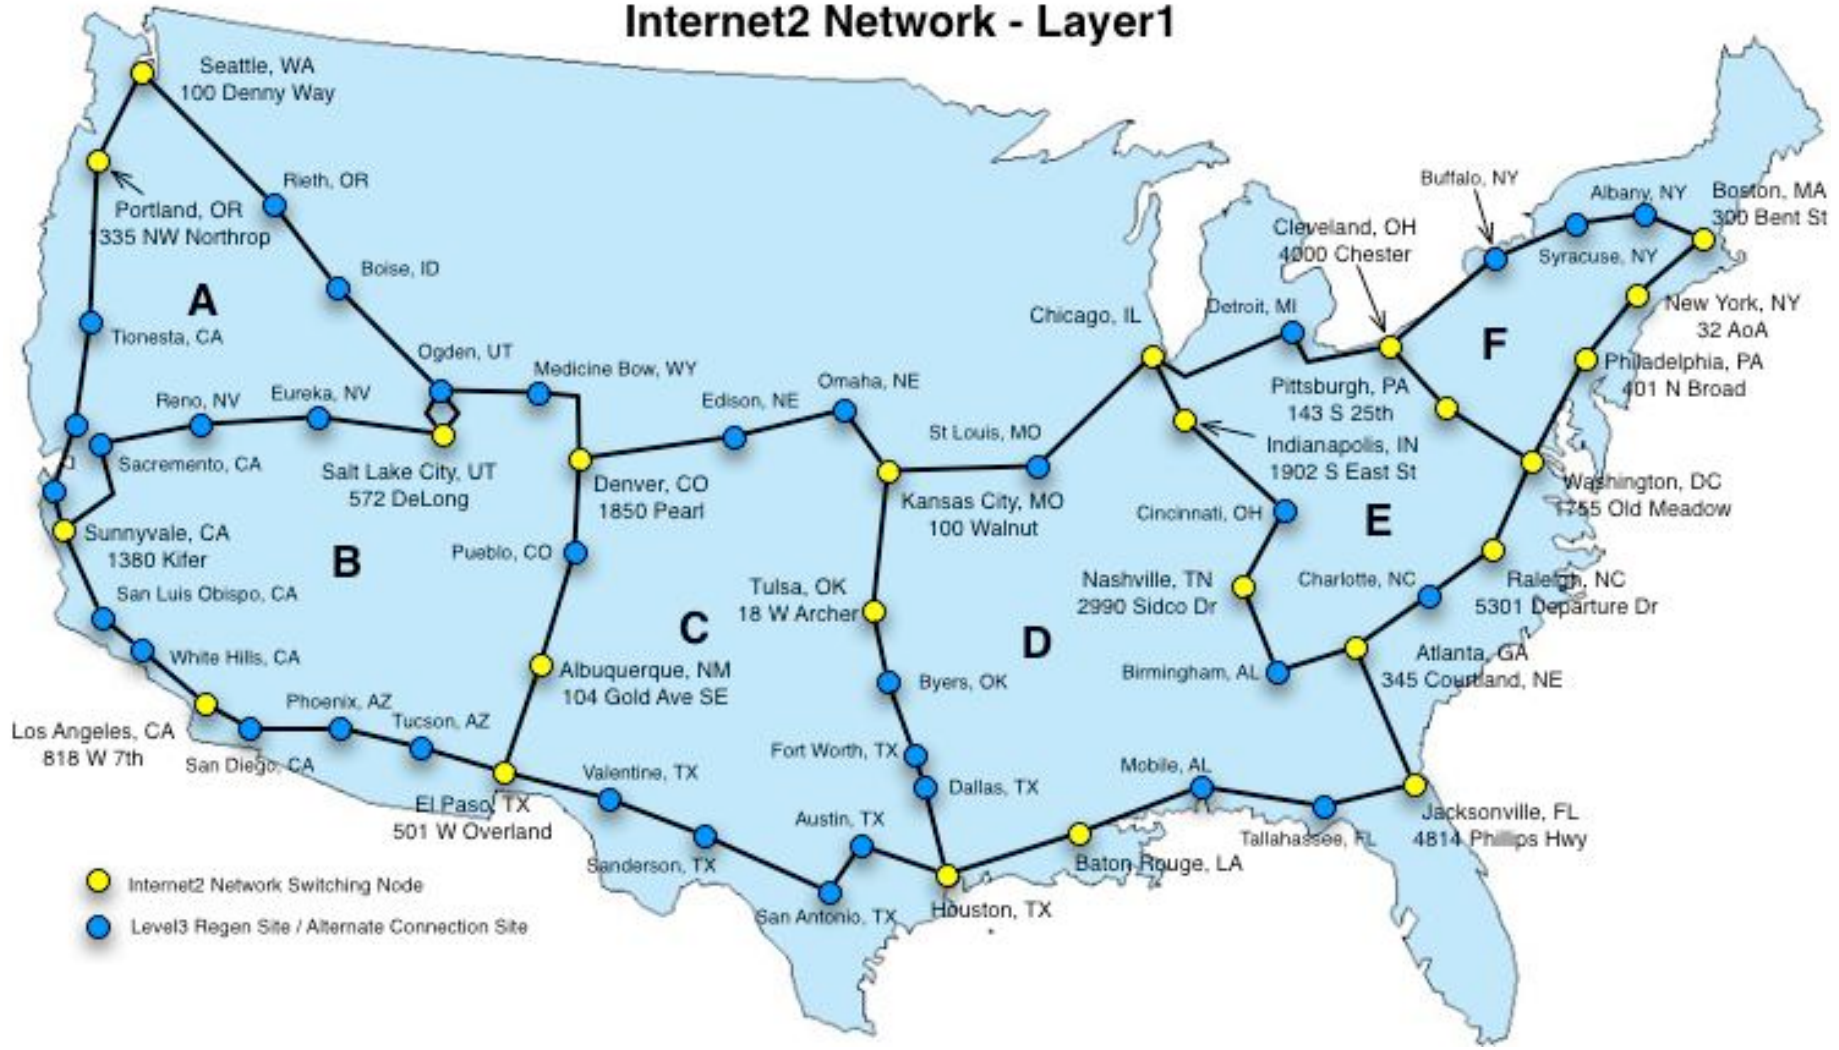
- ESnet will build its next generation IP network and its new circuit-oriented Science Data Network primarily on the Internet2 Network circuits (λ s) that are dedicated to ESnet, together with a few National Lambda Rail and other circuits
 - ESnet will provision and operate its own routing and switching 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

- ESnet4 will also involve an expansion of the multi-10Gb/s Metropolitan Area Rings in the San Francisco Bay Area, Chicago, Long Island, and Newport News, VA/Washington, DC area
 - 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

Internet2 Network Layer 1 Footprint

1625 miles / 2545 km

Internet2 Network - Layer1

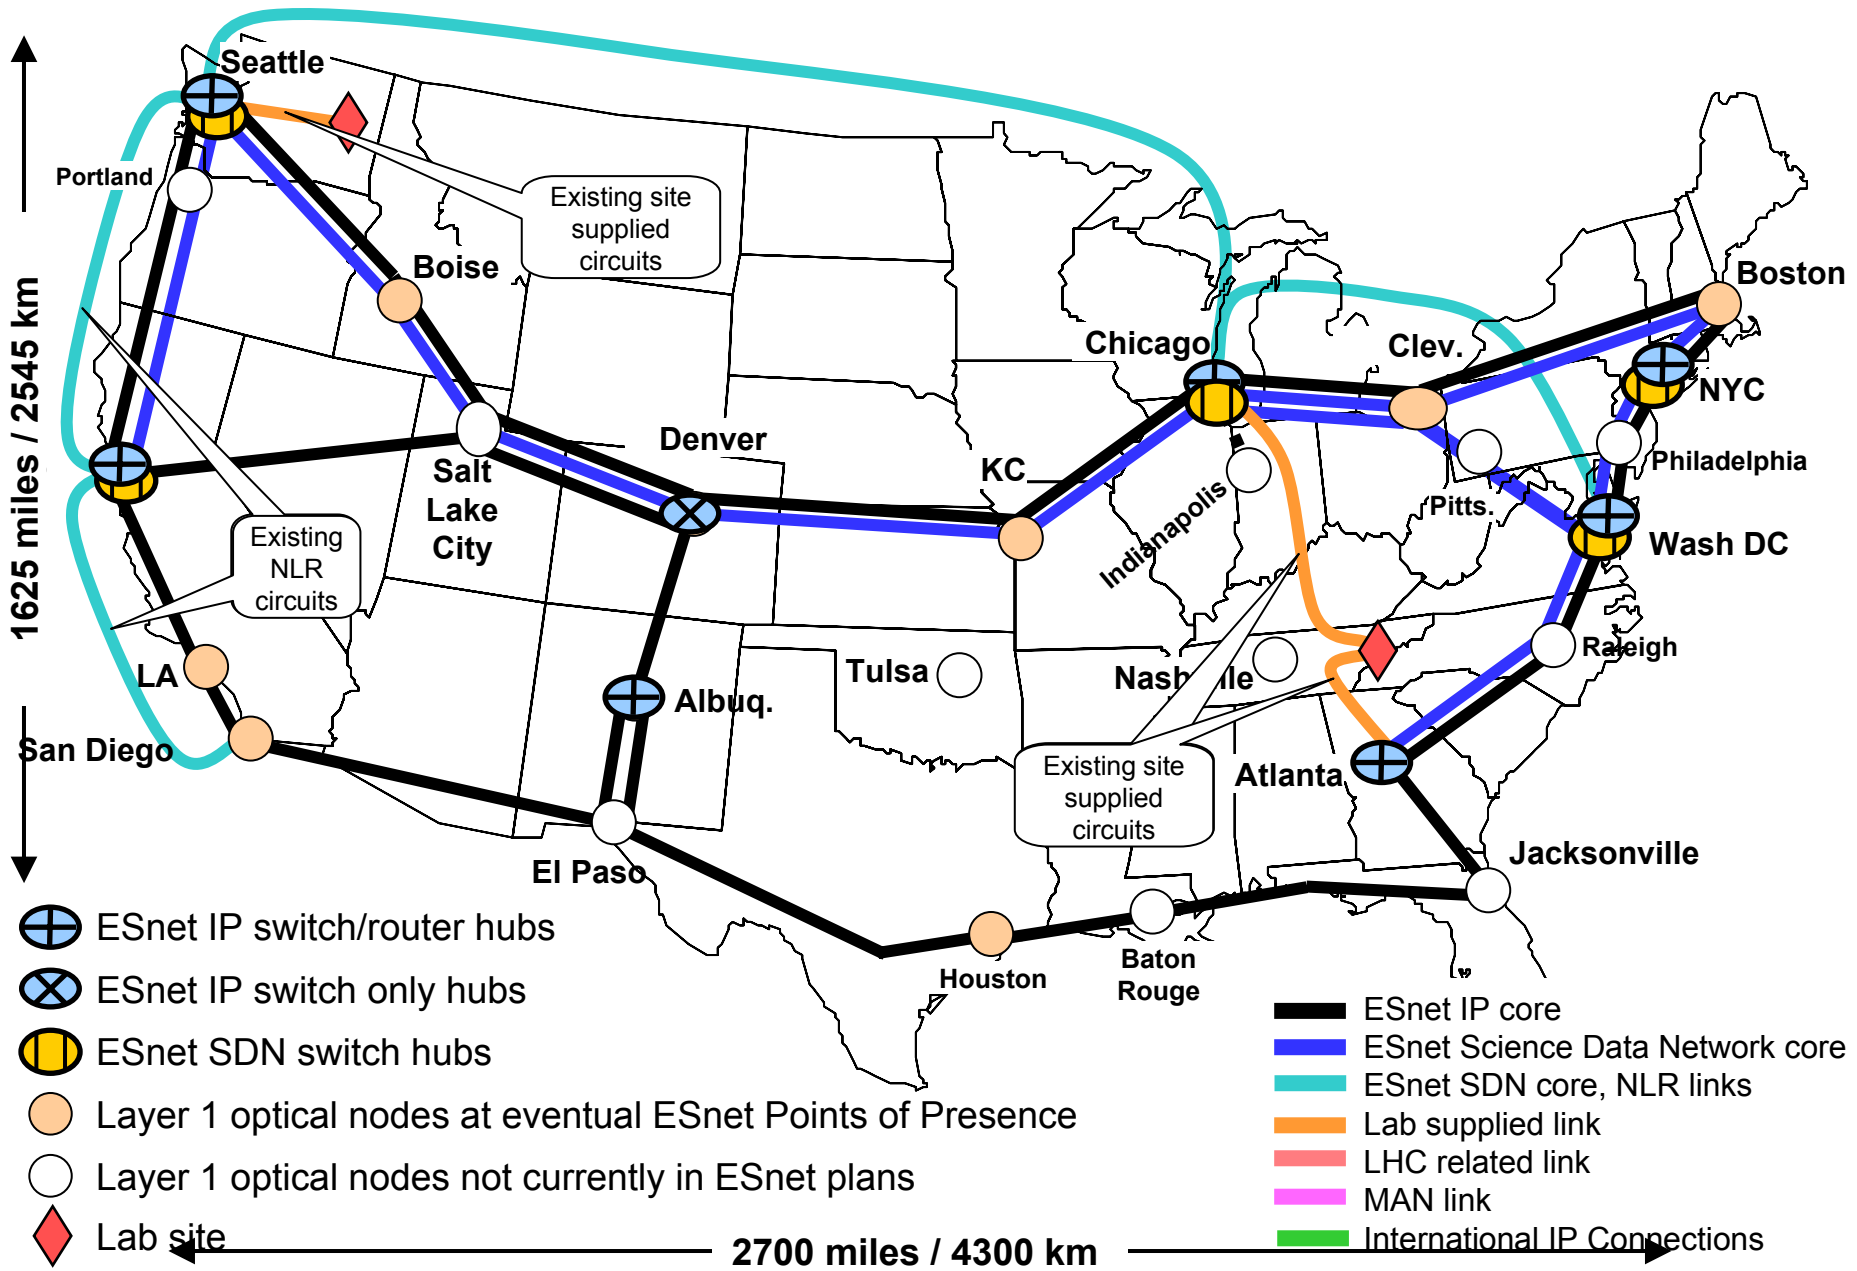


Core network fiber path is
~ 14,000 miles / 24,000 km

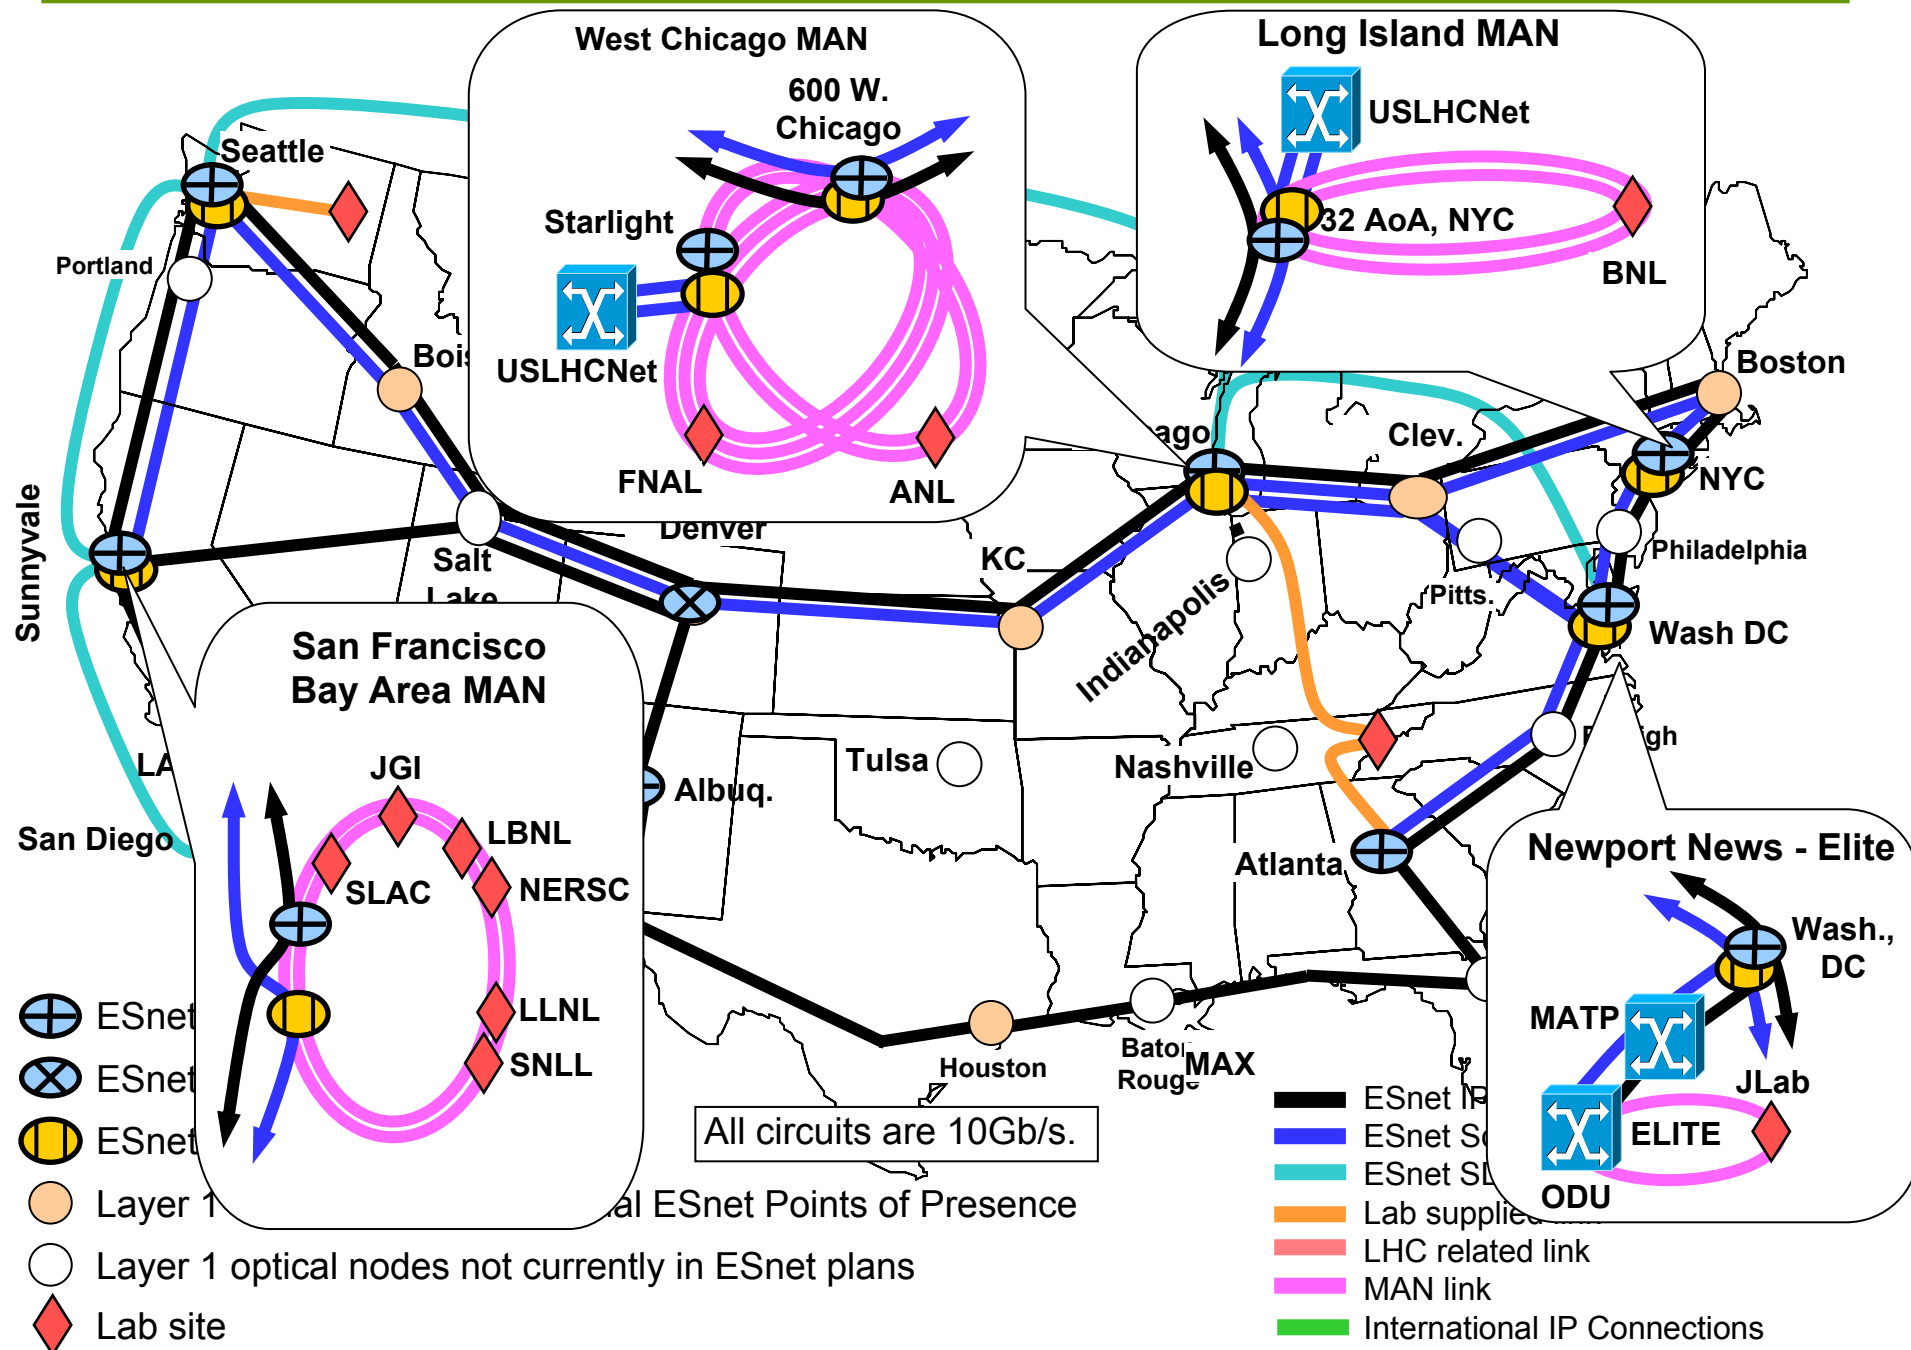
2700 miles / 4300 km

ESnet4 IP + SDN Configuration, mid-August, 2007

All circuits are 10Gb/s.

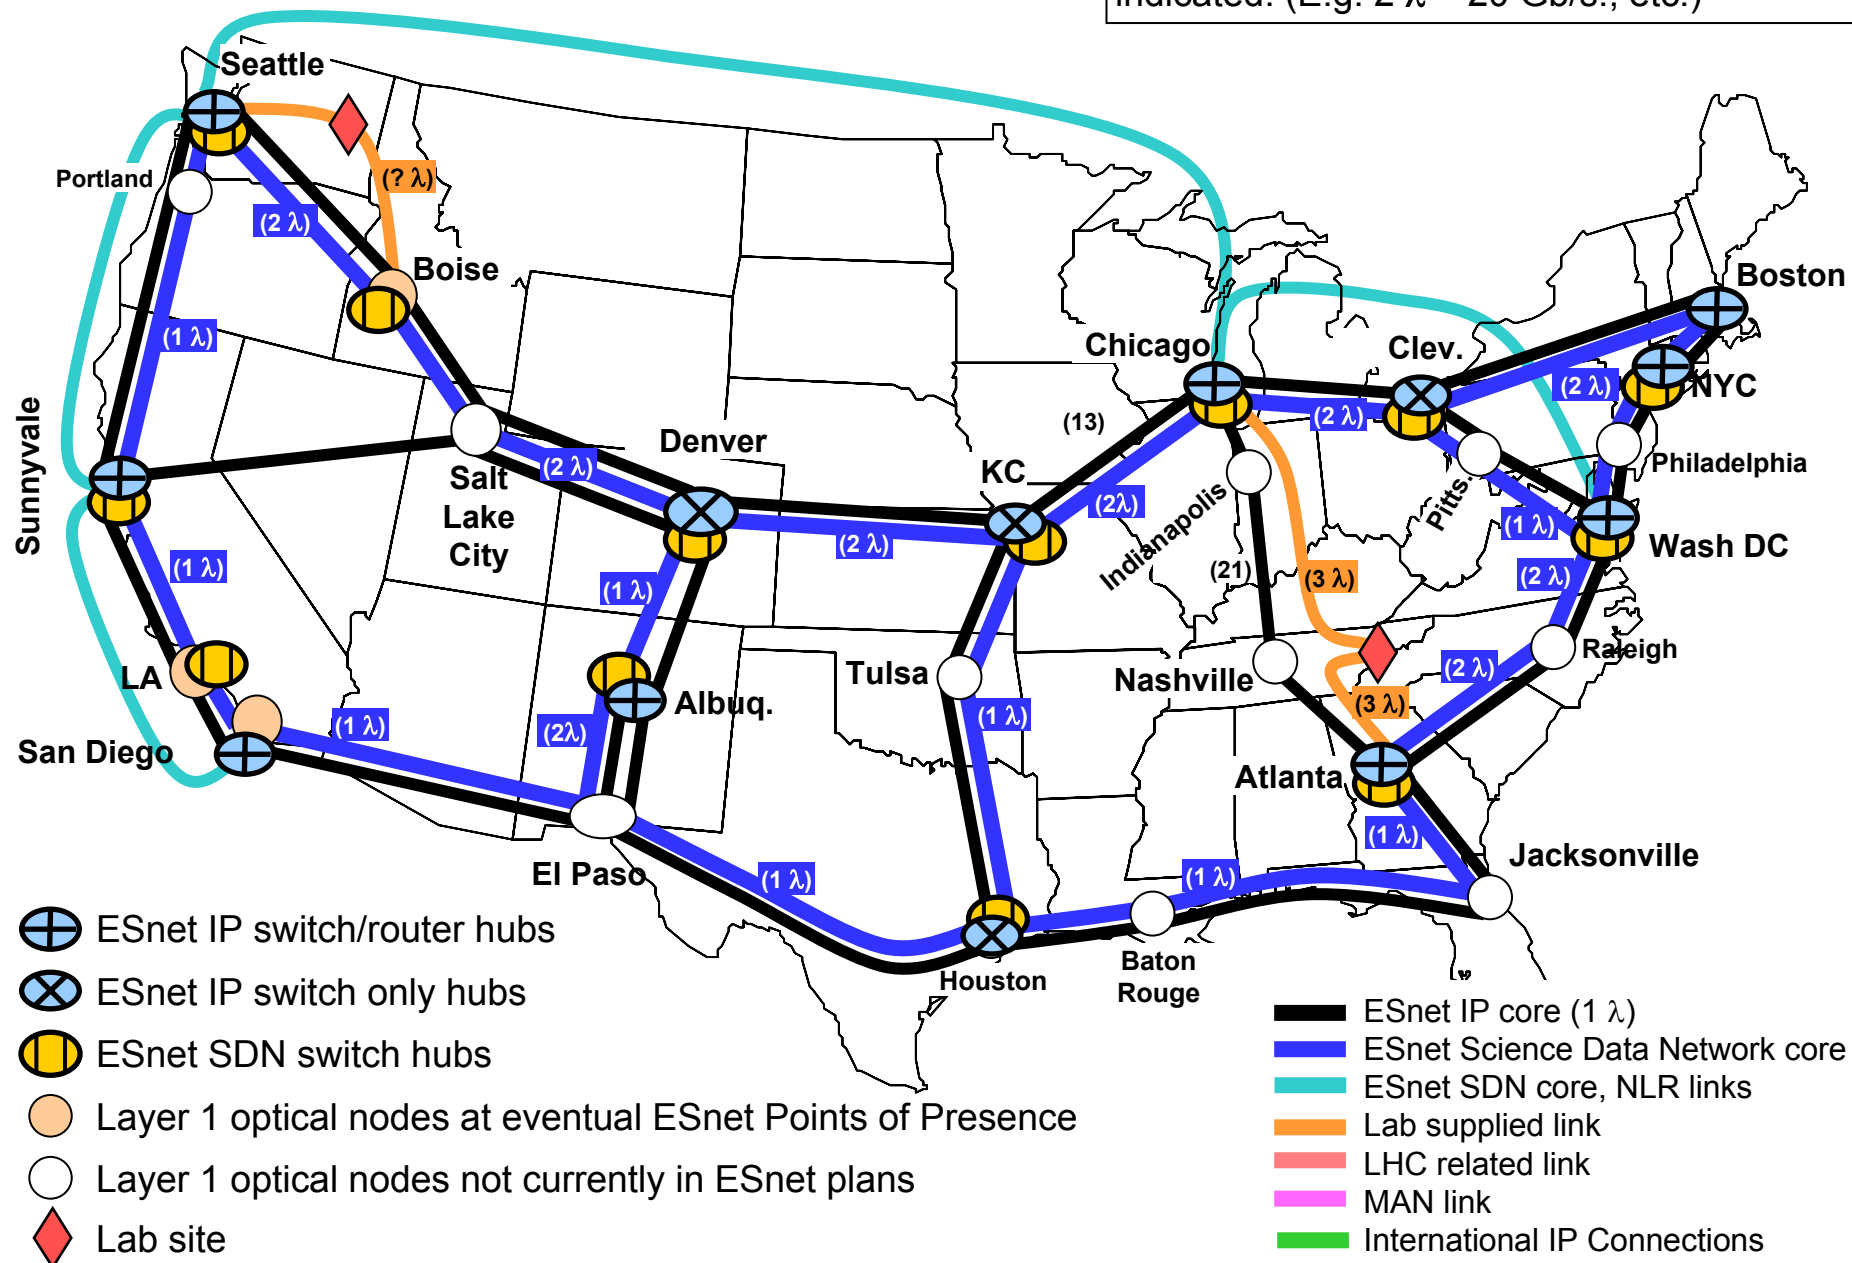


ESnet4 Metro Area Rings, 2007 Configurations

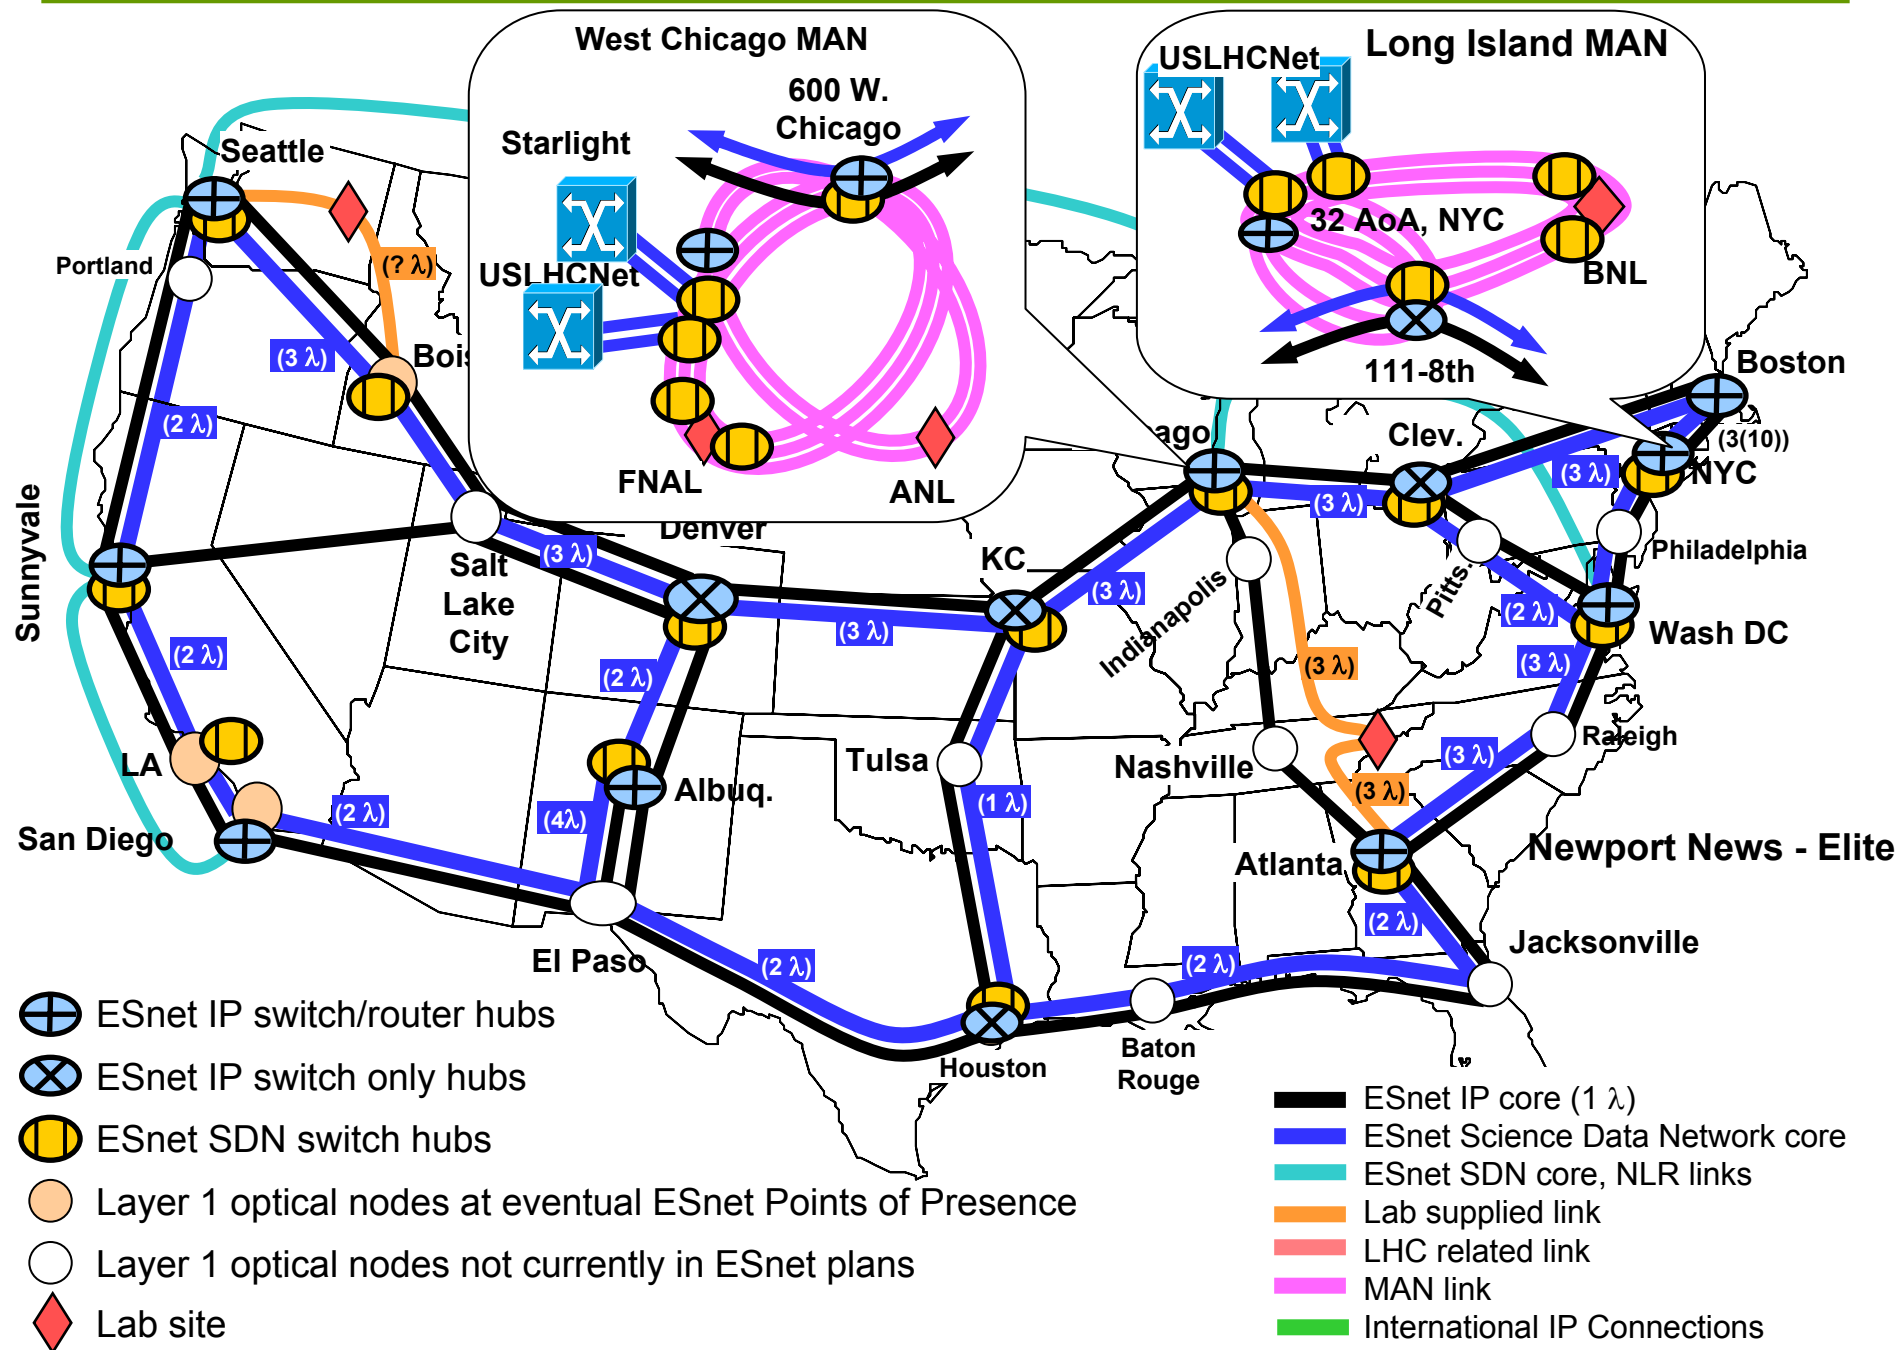


ESnet4 IP + SDN, 2008 Configuration

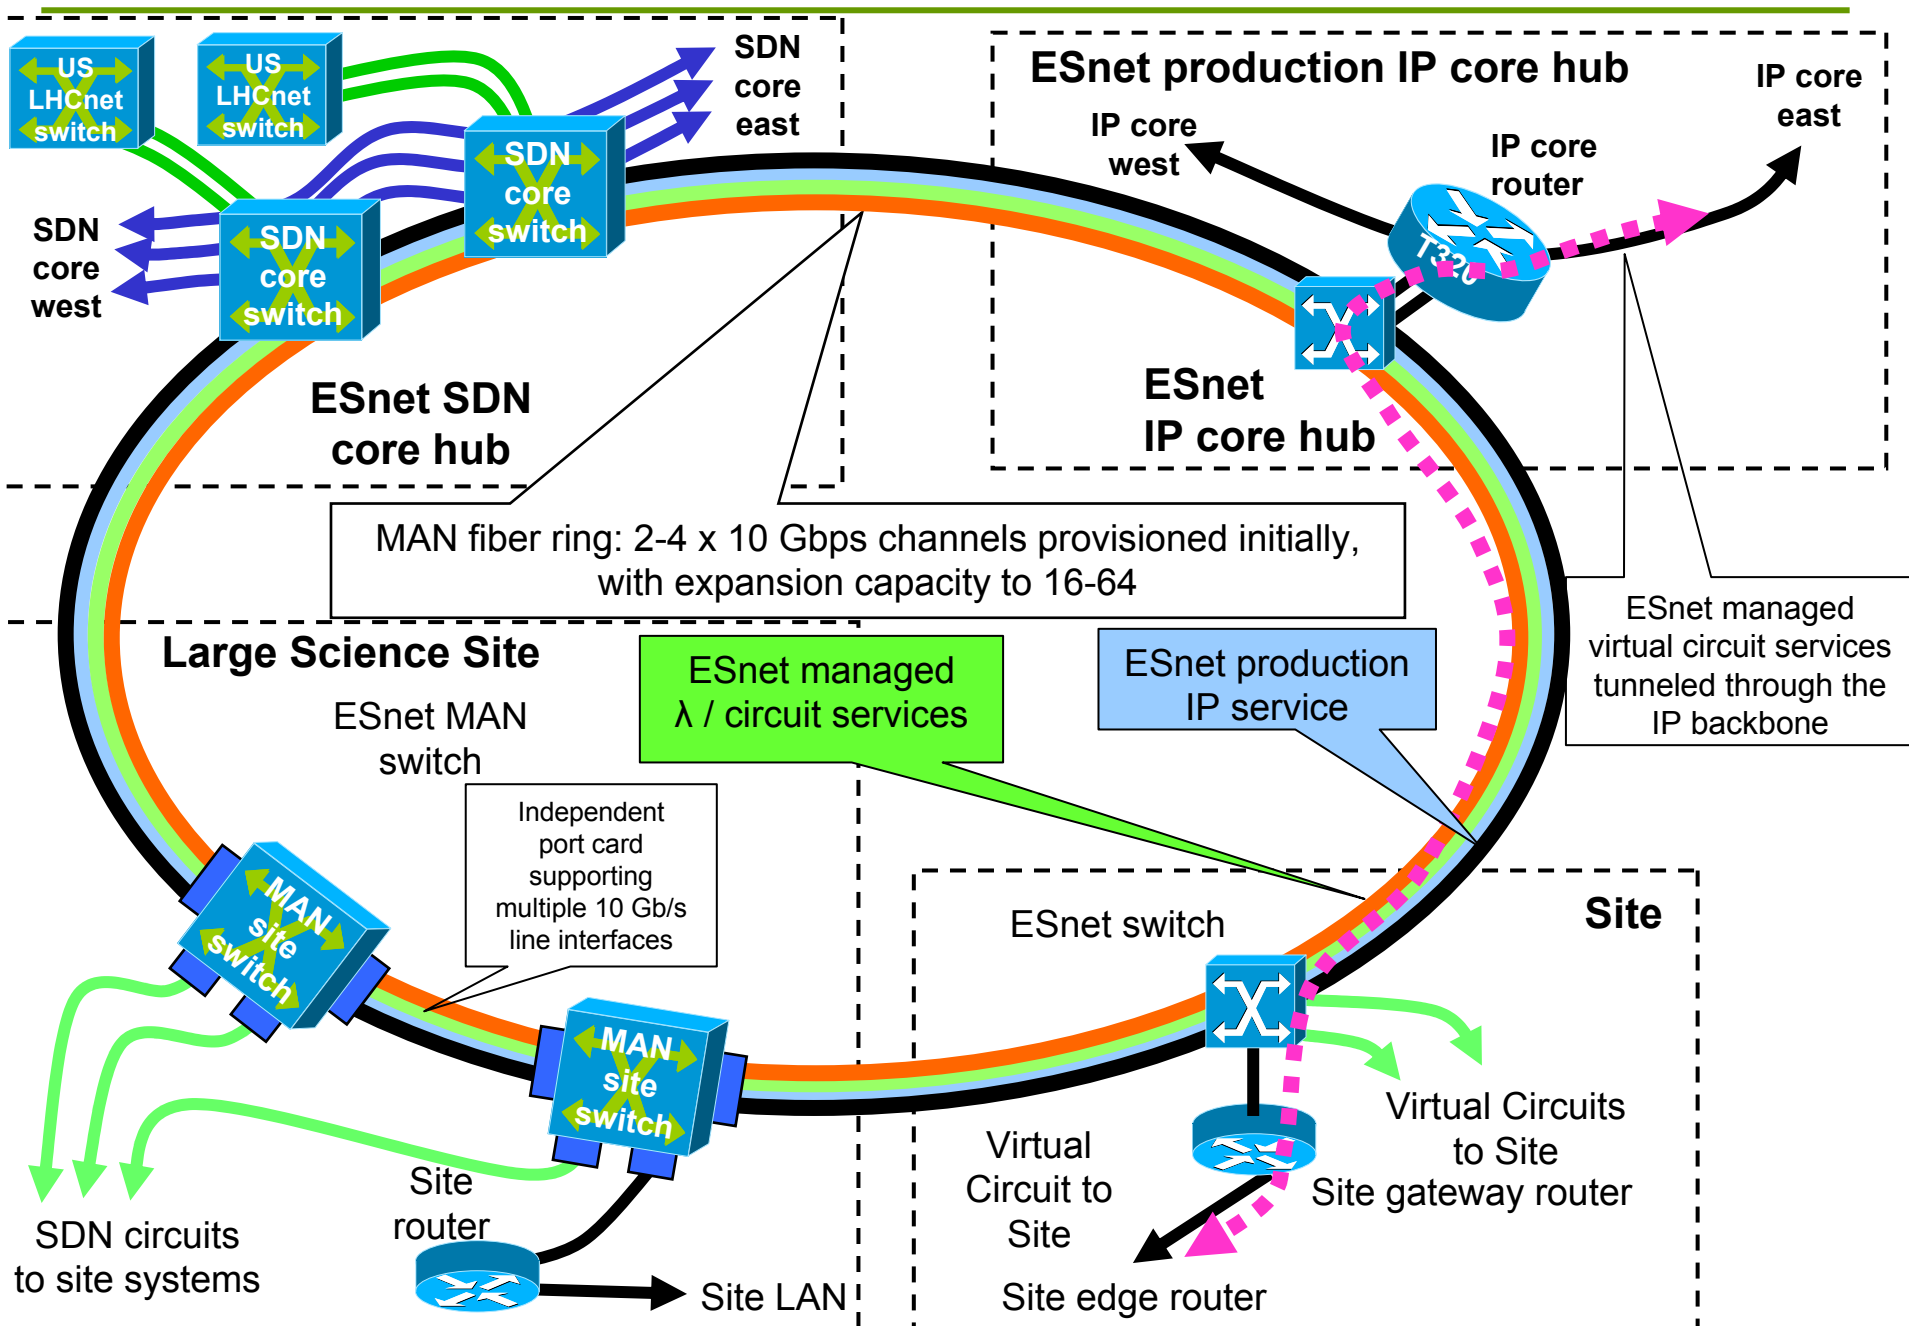
All circuits are 10Gb/s, or multiples thereof, as indicated. (E.g. $2 \lambda = 20 \text{ Gb/s.}$, etc.)



ESnet4 Metro Area Rings, 2009 Configurations



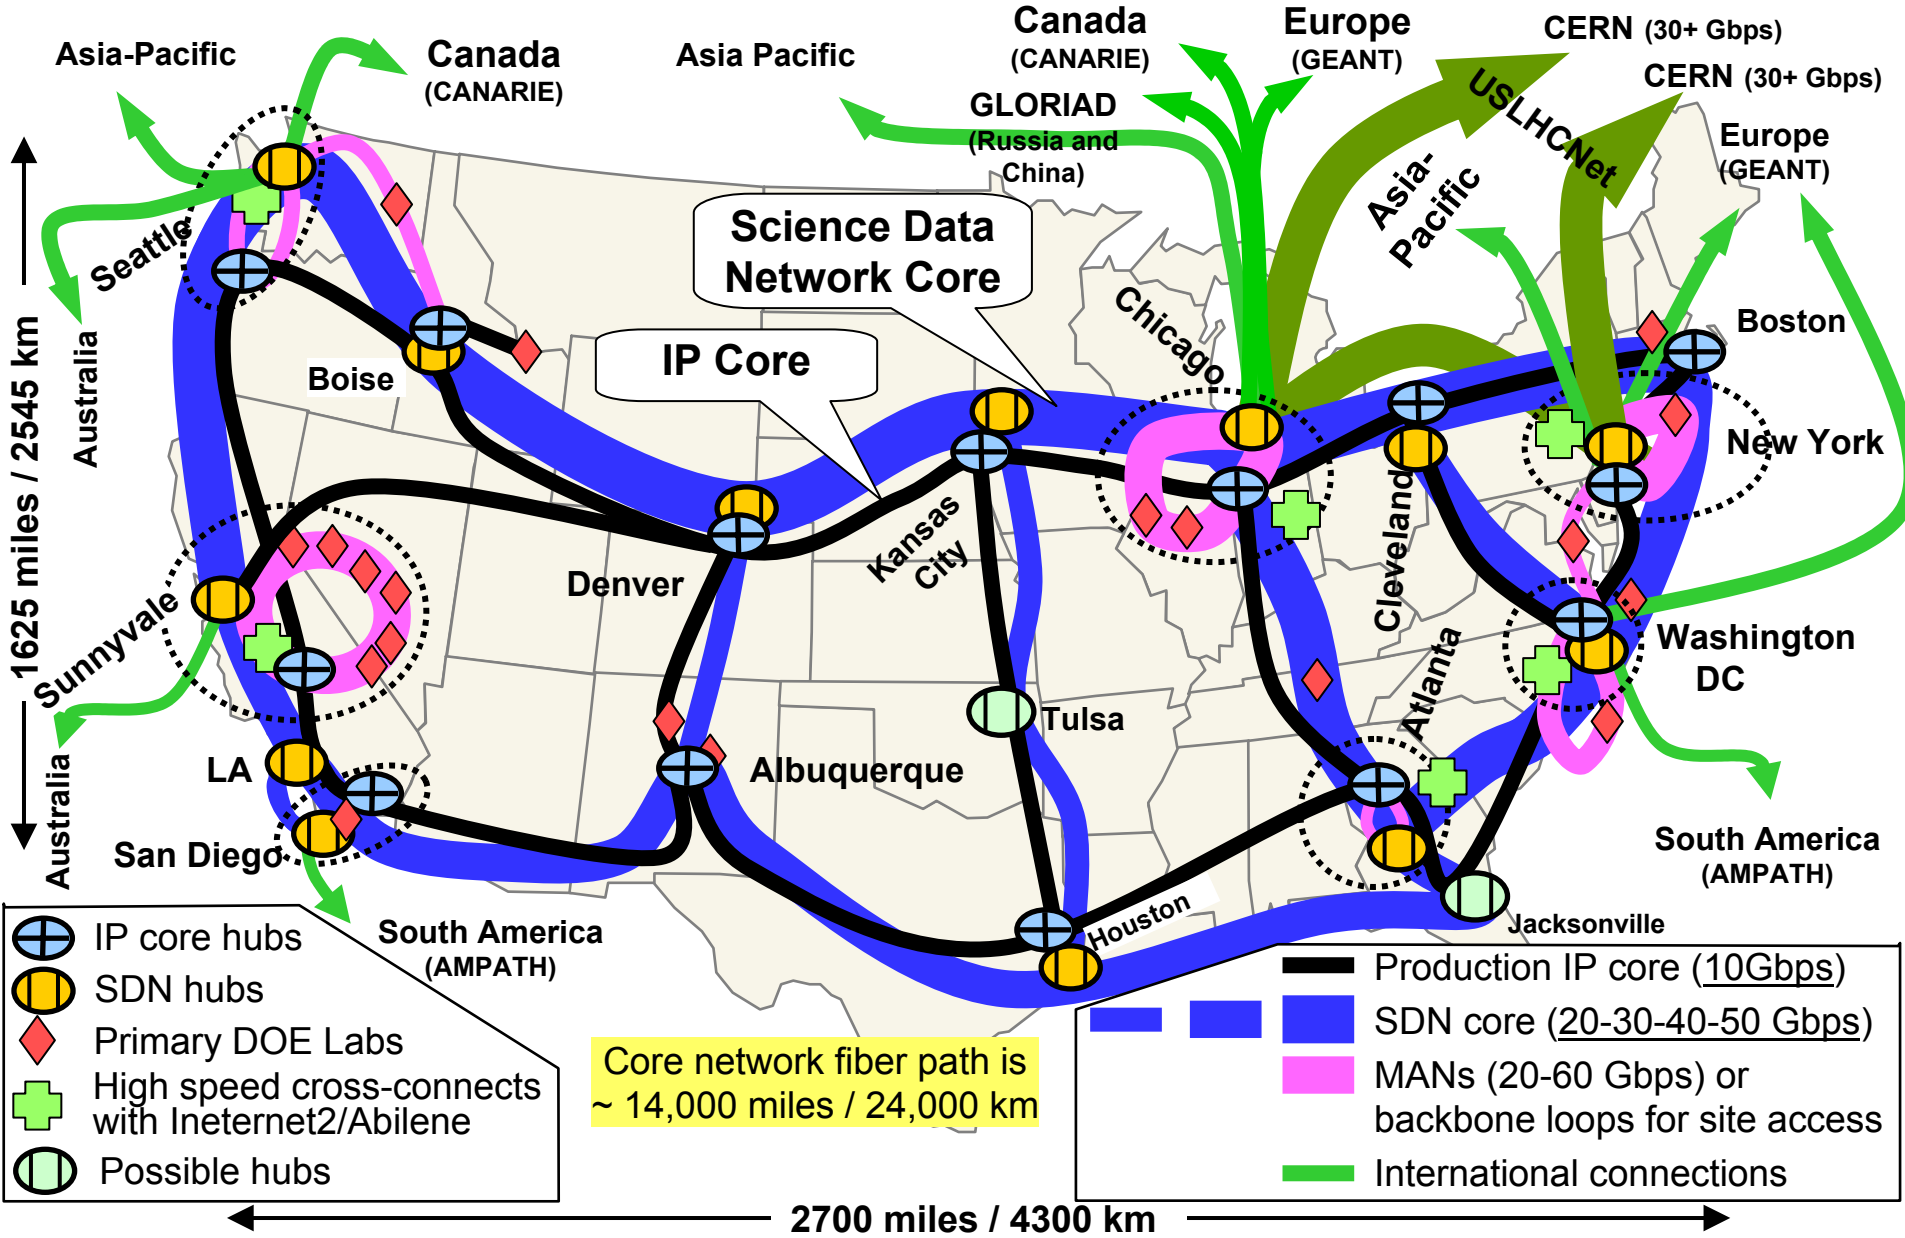
ESnet Metropolitan Area Network Ring Architecture for High Reliability Sites



ESnet4

Core networks: 50-60 Gbps by 2009-2010, 200-600 Gbps by 2011-2012

(IP peering connections not shown)



End

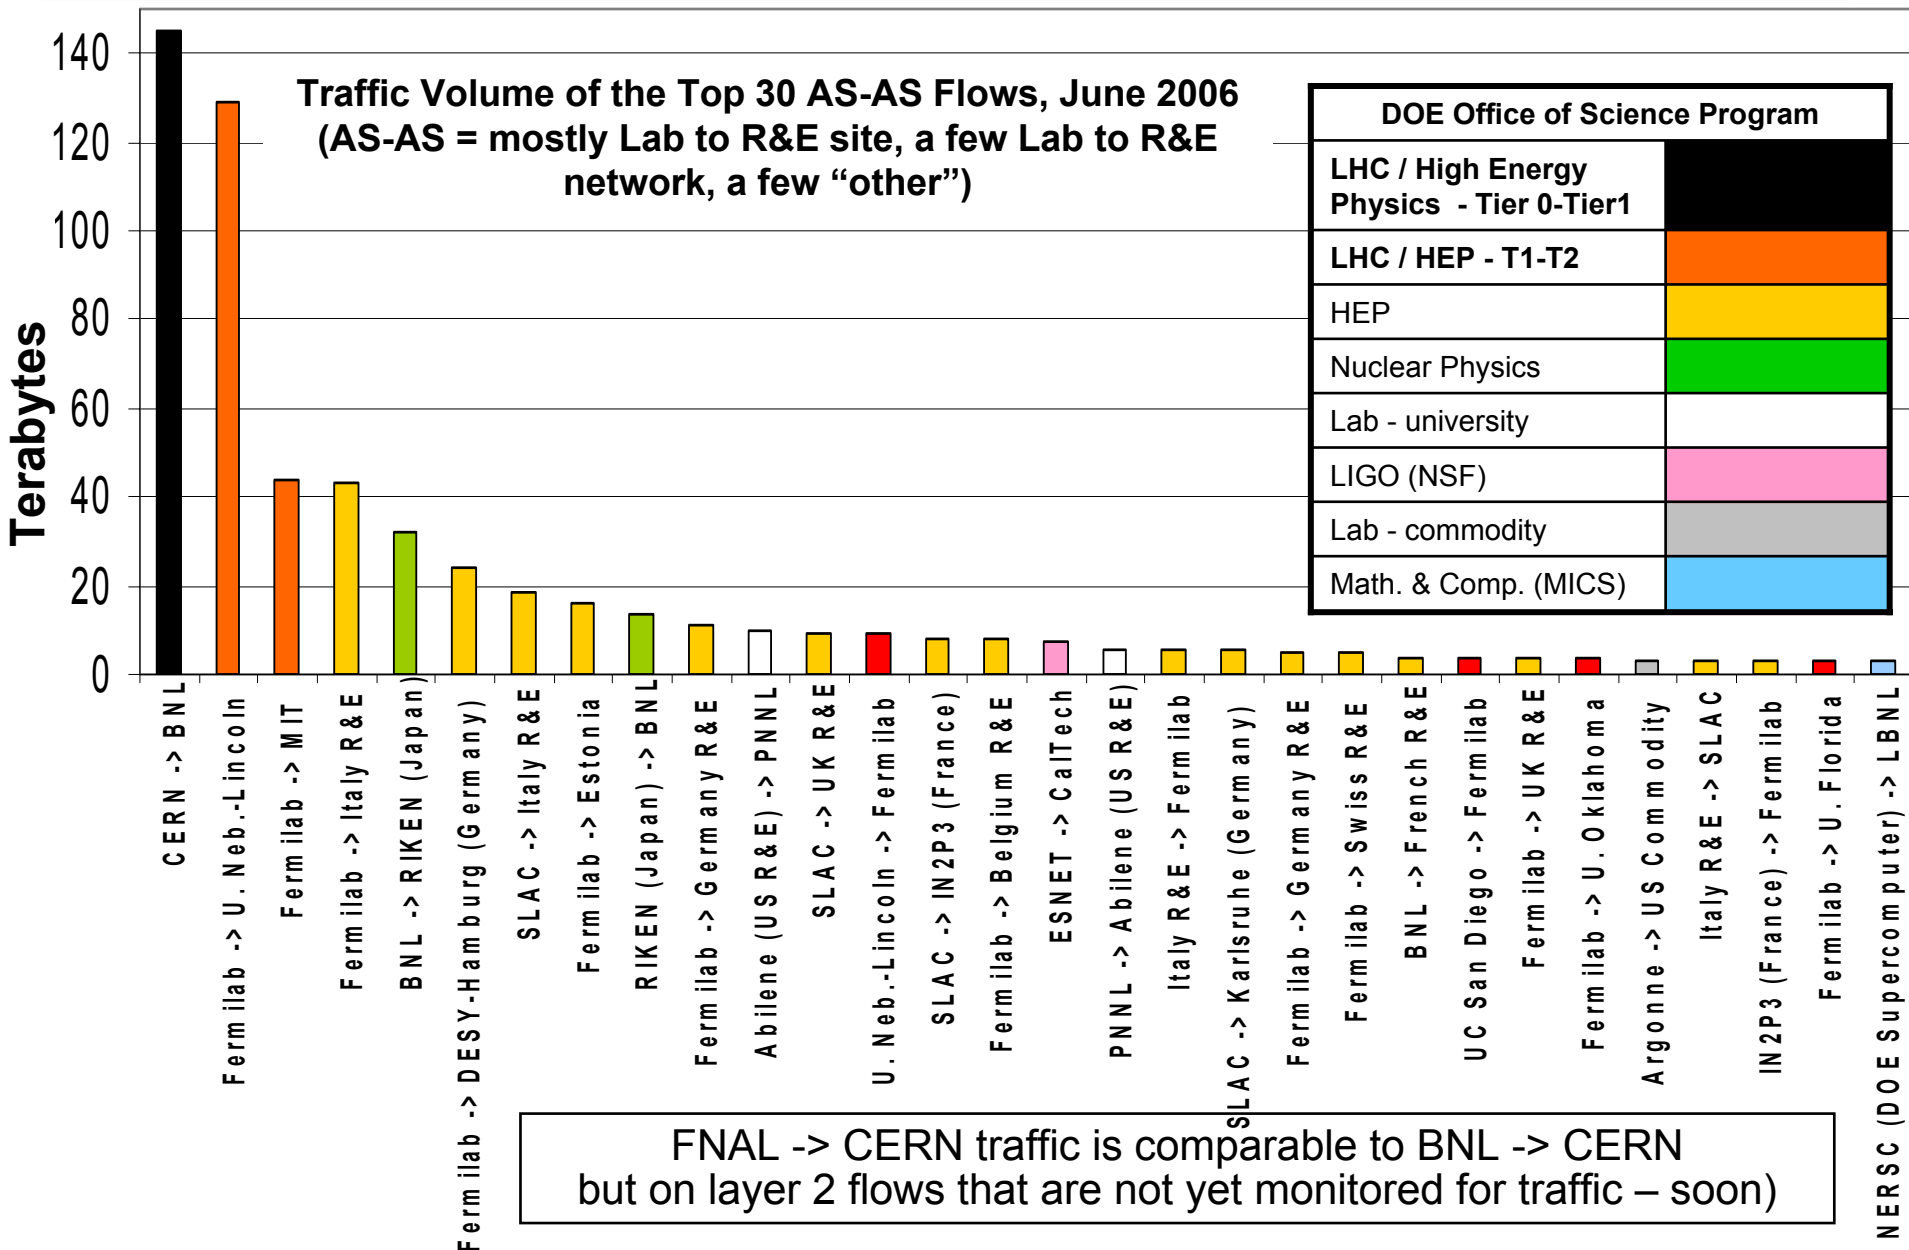
Additional Material Follows

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 dominate feature of the Office of Science (“SC”) programs
- Distributed systems for data analysis, simulations, instrument operation, etc., are essential and are now common
- These changes are supported by network traffic pattern observations

Large-Scale Flow Trends, June 2006


Subtitle: “Onslaught of the LHC”




Traffic Patterns are Changing Dramatically

Top 100 HOST-HOST Traffic

Base Date: 2005-01-31 -- Aggregation: 30 day(s) -- Router(s): all





total traffic, TBy





Top 100 HOST-HOST Traffic

Base Date: 2005-07-31 -- Aggregation: 30 day(s) -- Router(s): all




Top 100 HOST-HOST Traffic




Top 100 HOST-HOST Traffic

Base Date: 2006-07-14 -- Aggregation: 30 day(s) -- Router(s): all




total traffic, TBy



- While the total traffic is increasing exponentially
 - Peak flow – that is system-to-system – bandwidth is decreasing
 - The number of large flows is increasing

Question: Why is peak flow bandwidth decreasing while total traffic is increasing?



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

- In June, 2006 **72%** of the hosts generating the 1000 work flows were involved in parallel data movers (Grid applications)
- ***This, combined with the dramatic increase in the proportion of traffic due to large-scale science (now 50% of all traffic) represents the most significant traffic pattern change in the history of ESnet***
- This probably argues for a network architecture that favors path multiplicity and route diversity

Network Observation – Circuit-like Behavior

- For large-scale data handling projects (e.g. LIGO - Caltech) the work flows (system to system data transfers) exhibit circuit-like behavior
- This circuit has a duration of about 3 months (all of the top traffic generating work flows are similar to this) - this argues for a circuit-oriented element in the network architecture

