



***U.S. Department of Energy's
Office of Science***

Connecting Scientists

**...to collaborators, to facilities, to datastores, to
computational resources, to....**

Energy Sciences Network (ESnet)

**April 8, 2005
FESAC Meeting**

**Mary Anne Scott
Program Manager
Advanced Scientific Computing Research
Mathematical, Information, and Computational Research Division**
scott@er.doe.gov



What is ESnet?

- **Mission:**
 - Provide, interoperable, effective and reliable communications infrastructure and leading-edge network services that support missions of the Department of Energy, especially the Office of Science
- **Vision:**
 - Provide seamless and ubiquitous access, via shared collaborative information and computational environments, to the facilities, data, and colleagues needed to accomplish their goals.
- **Role:**
 - A component of the Office of Science infrastructure critical to the success of its research programs (funded through ASCR/MICS and managed and operated by ESnet staff at LBNL).

Essentially all of the national data traffic supporting US science is carried by two networks—ESnet and Internet-2/Abilene (which plays a similar role for the university community)



What is ESnet's user base?

- Between 10,000 and 100,000 researchers in the US (guesstimate)
- Mainly Office of Science programs—ASCR, BER, BES, FES, HEP, NP
- Also traffic for NNSA and others
- All the US national labs
- Hundreds of universities
- Hundreds of foreign institutions

Characteristics of the user base

- Many casual users
- Users from science disciplines that span SC interests
- Users concerned with any data intensive and computationally intensive tasks
- Collaborators distributed geographically, small to large groups

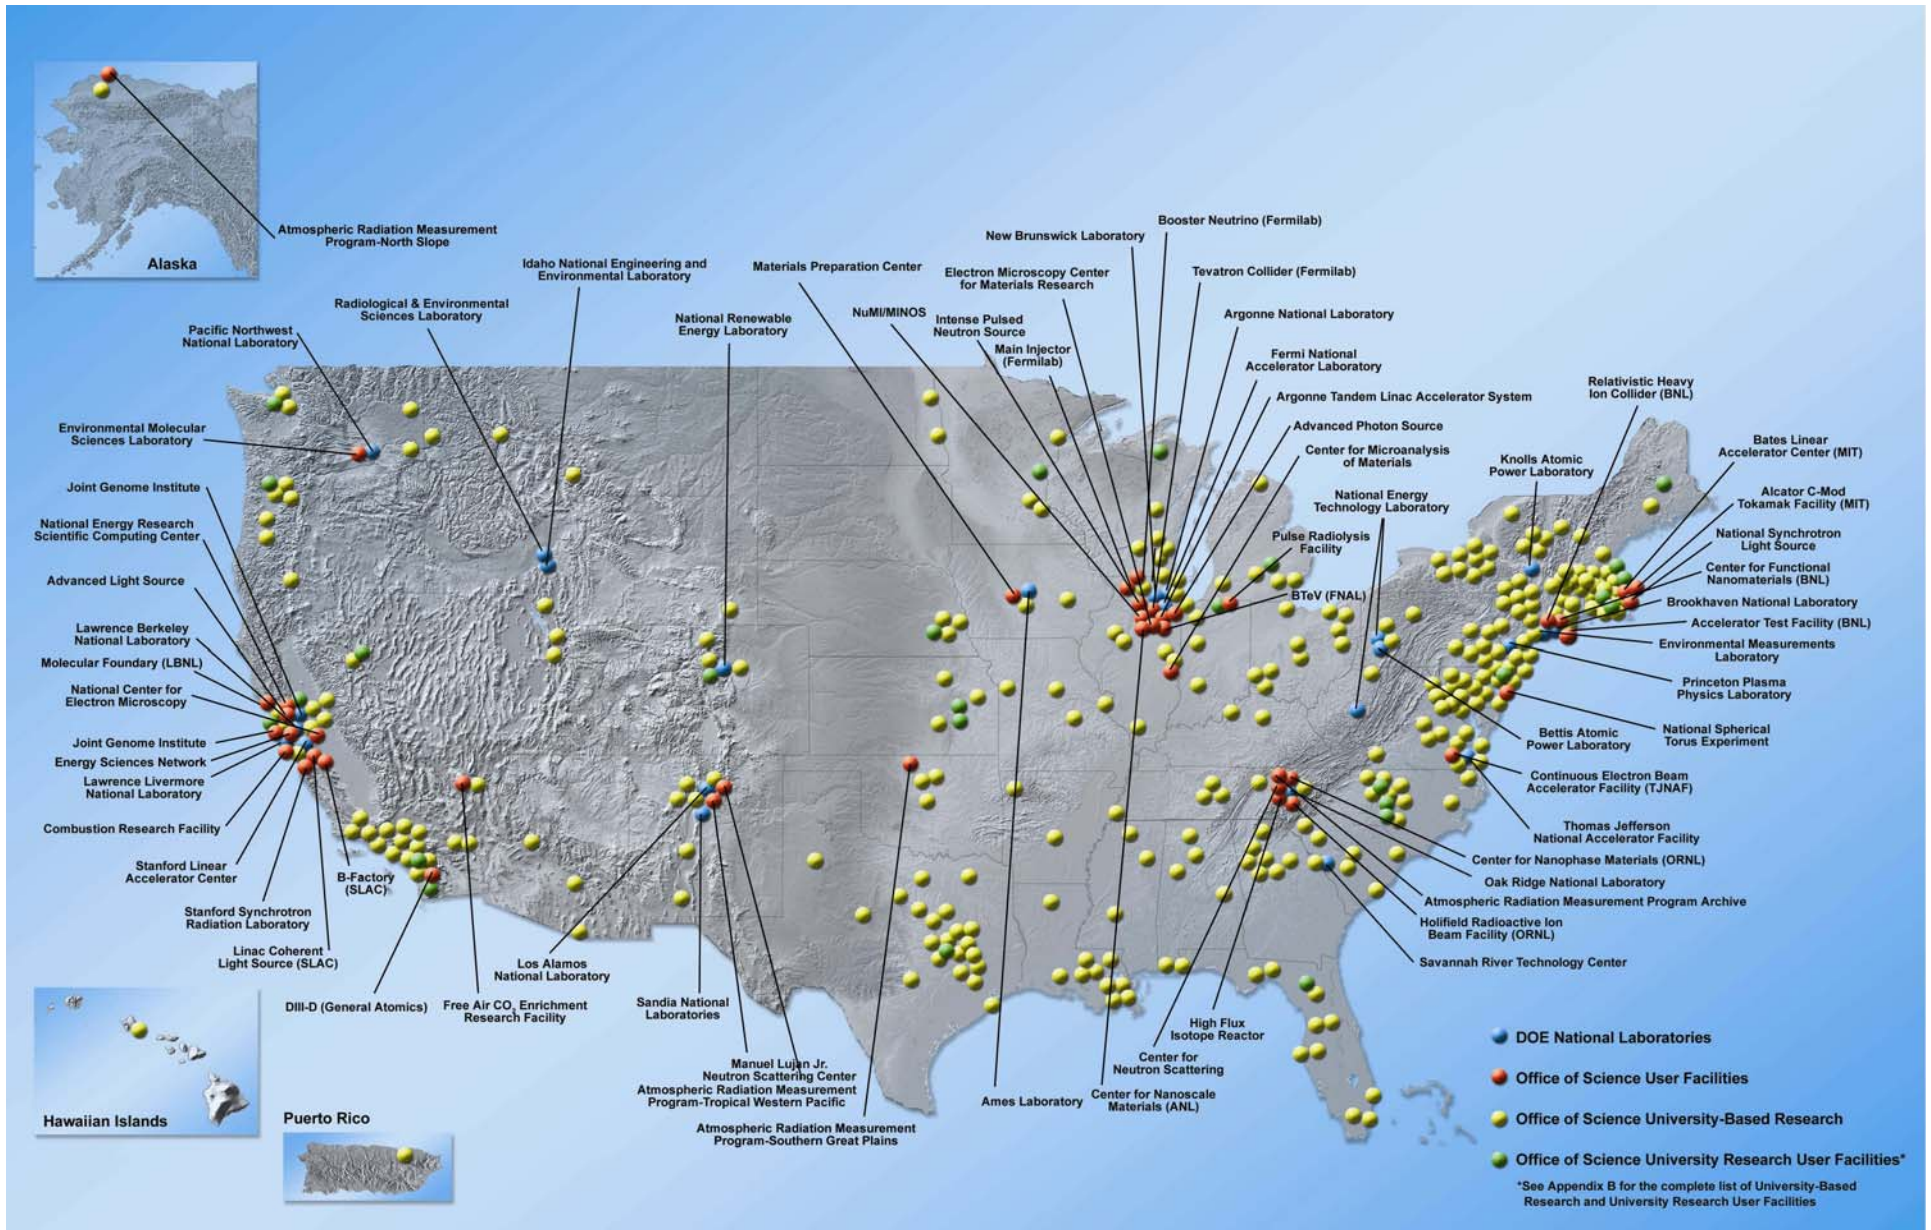


Why is ESnet important?

- Enables thousands of DOE, university and industry scientists and collaborators worldwide to make effective use of unique DOE research facilities and computing resources independent of time and geographic location
 - Direct connections to all major DOE sites
 - Access to the global Internet (managing 150,000 routes at 10 commercial peering points)
 - User demand has grown by a factor of more than 10,000 since its inception in the mid 1990's—a 100 percent increase every year since 1990
- Capabilities not available through commercial networks
 - Architected to move huge amounts of data between a small number of sites
 - High bandwidth peering to provide access to US, European, Asia-Pacific, and other research and education networks.

Objective: *Support scientific research* by providing seamless and ubiquitous access to the facilities, data, and colleagues

Office of Science – A Distributed Enterprise





Science Mission Critical Infrastructure

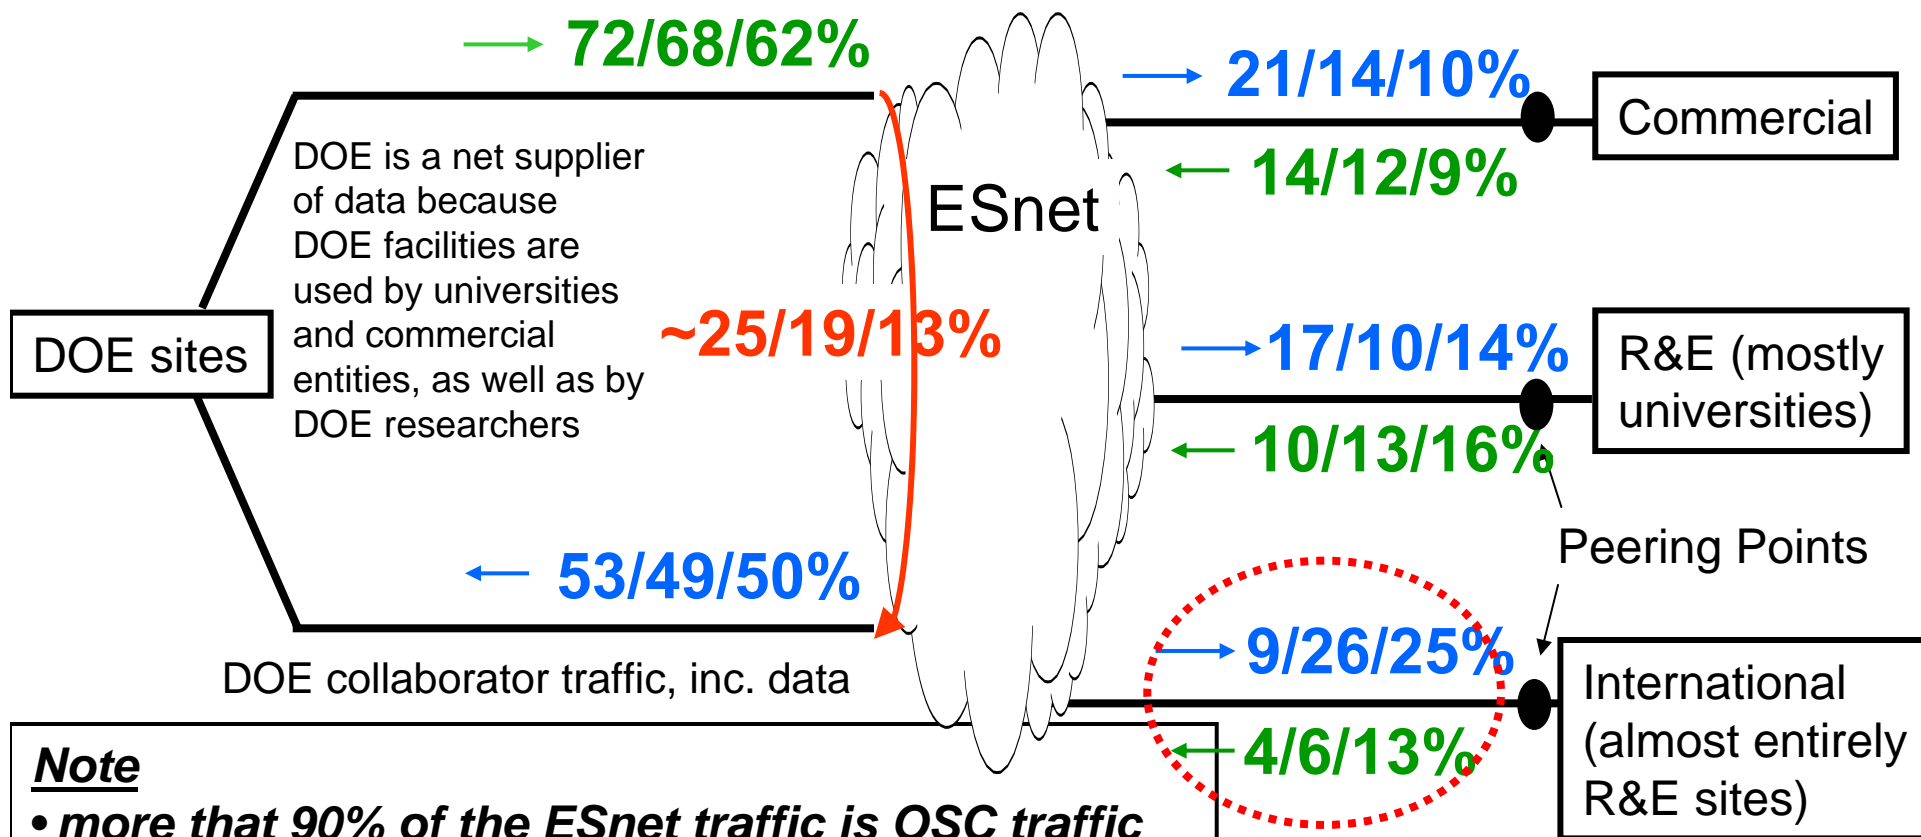
- ESnet is a visible and critical piece of DOE science infrastructure
 - If ESnet fails, tens of thousands of DOE and University users know it within minutes if not seconds
- Scalable operation is essential
 - R&E networks typically operate with a small staff
- Requires high reliability and high operational security in the network operations

Who Generates ESnet Traffic?

ESnet Inter-Sector Traffic Summary, Jan 03 / Feb 04/ Nov 04:

1.7X overall traffic increase, 1.9X Office of Science increase

(The overall traffic, and the fraction going to international sites, is increasing due to physics data sharing by FNAL, SLAC, and BNL)



Note

- more than 90% of the ESnet traffic is OSC traffic
- less than 20% of the traffic is inter-Lab

ESnet Basic Appropriate Use Policy (AUP):

All ESnet traffic must originate or terminate at an ESnet site (no public transit traffic is allowed)

Traffic coming into ESnet = Green
 Traffic leaving ESnet = Blue
 Traffic between ESnet sites
 % = of total ingress or egress traffic

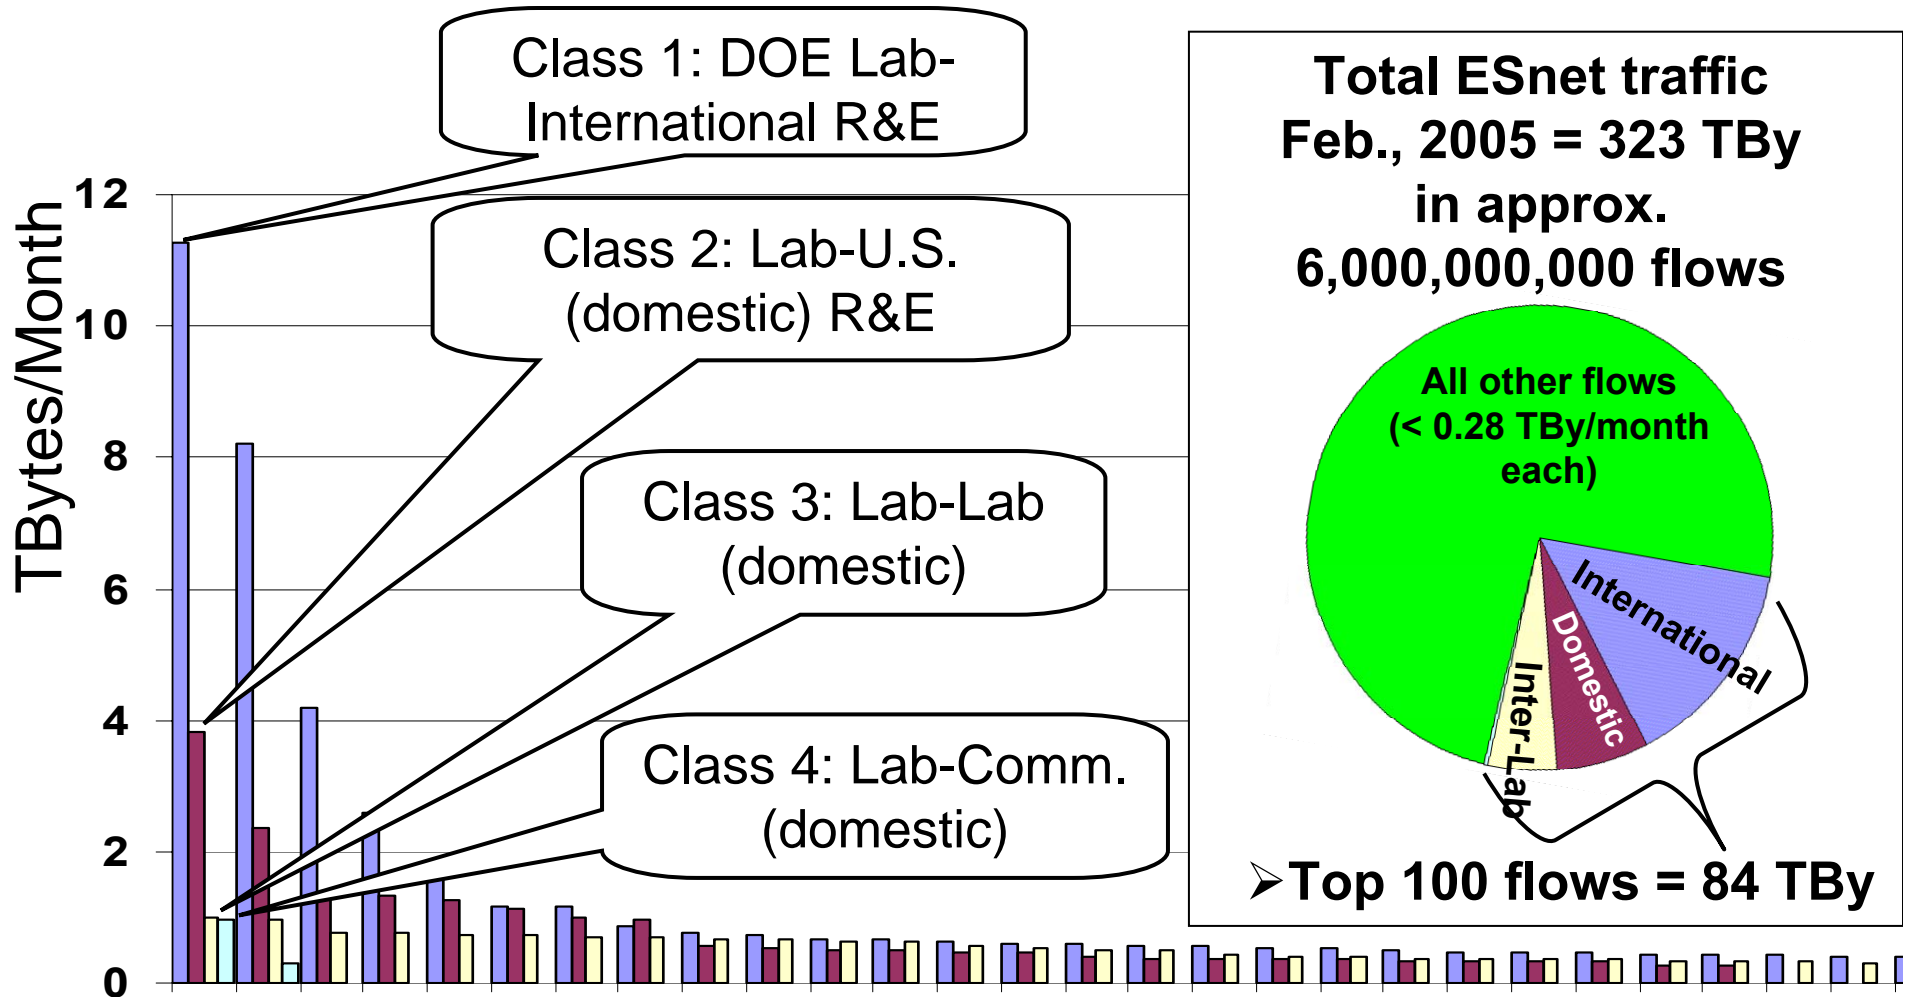


ESnet Traffic

- Since SLAC and FNAL based high energy physics experiment data analysis started, the top 100 ESnet flows have consistently accounted for 25% - 40% of ESnet's monthly total traffic
 - Sites in Europe analyze large amounts of this data
- As LHC (CERN high energy physics accelerator) data starts to move, the large science flows will increase a lot (200-2000 times)
 - Both LHC tier 1 (primary U.S. experiment data centers) are at DOE Labs – Fermilab and Brookhaven

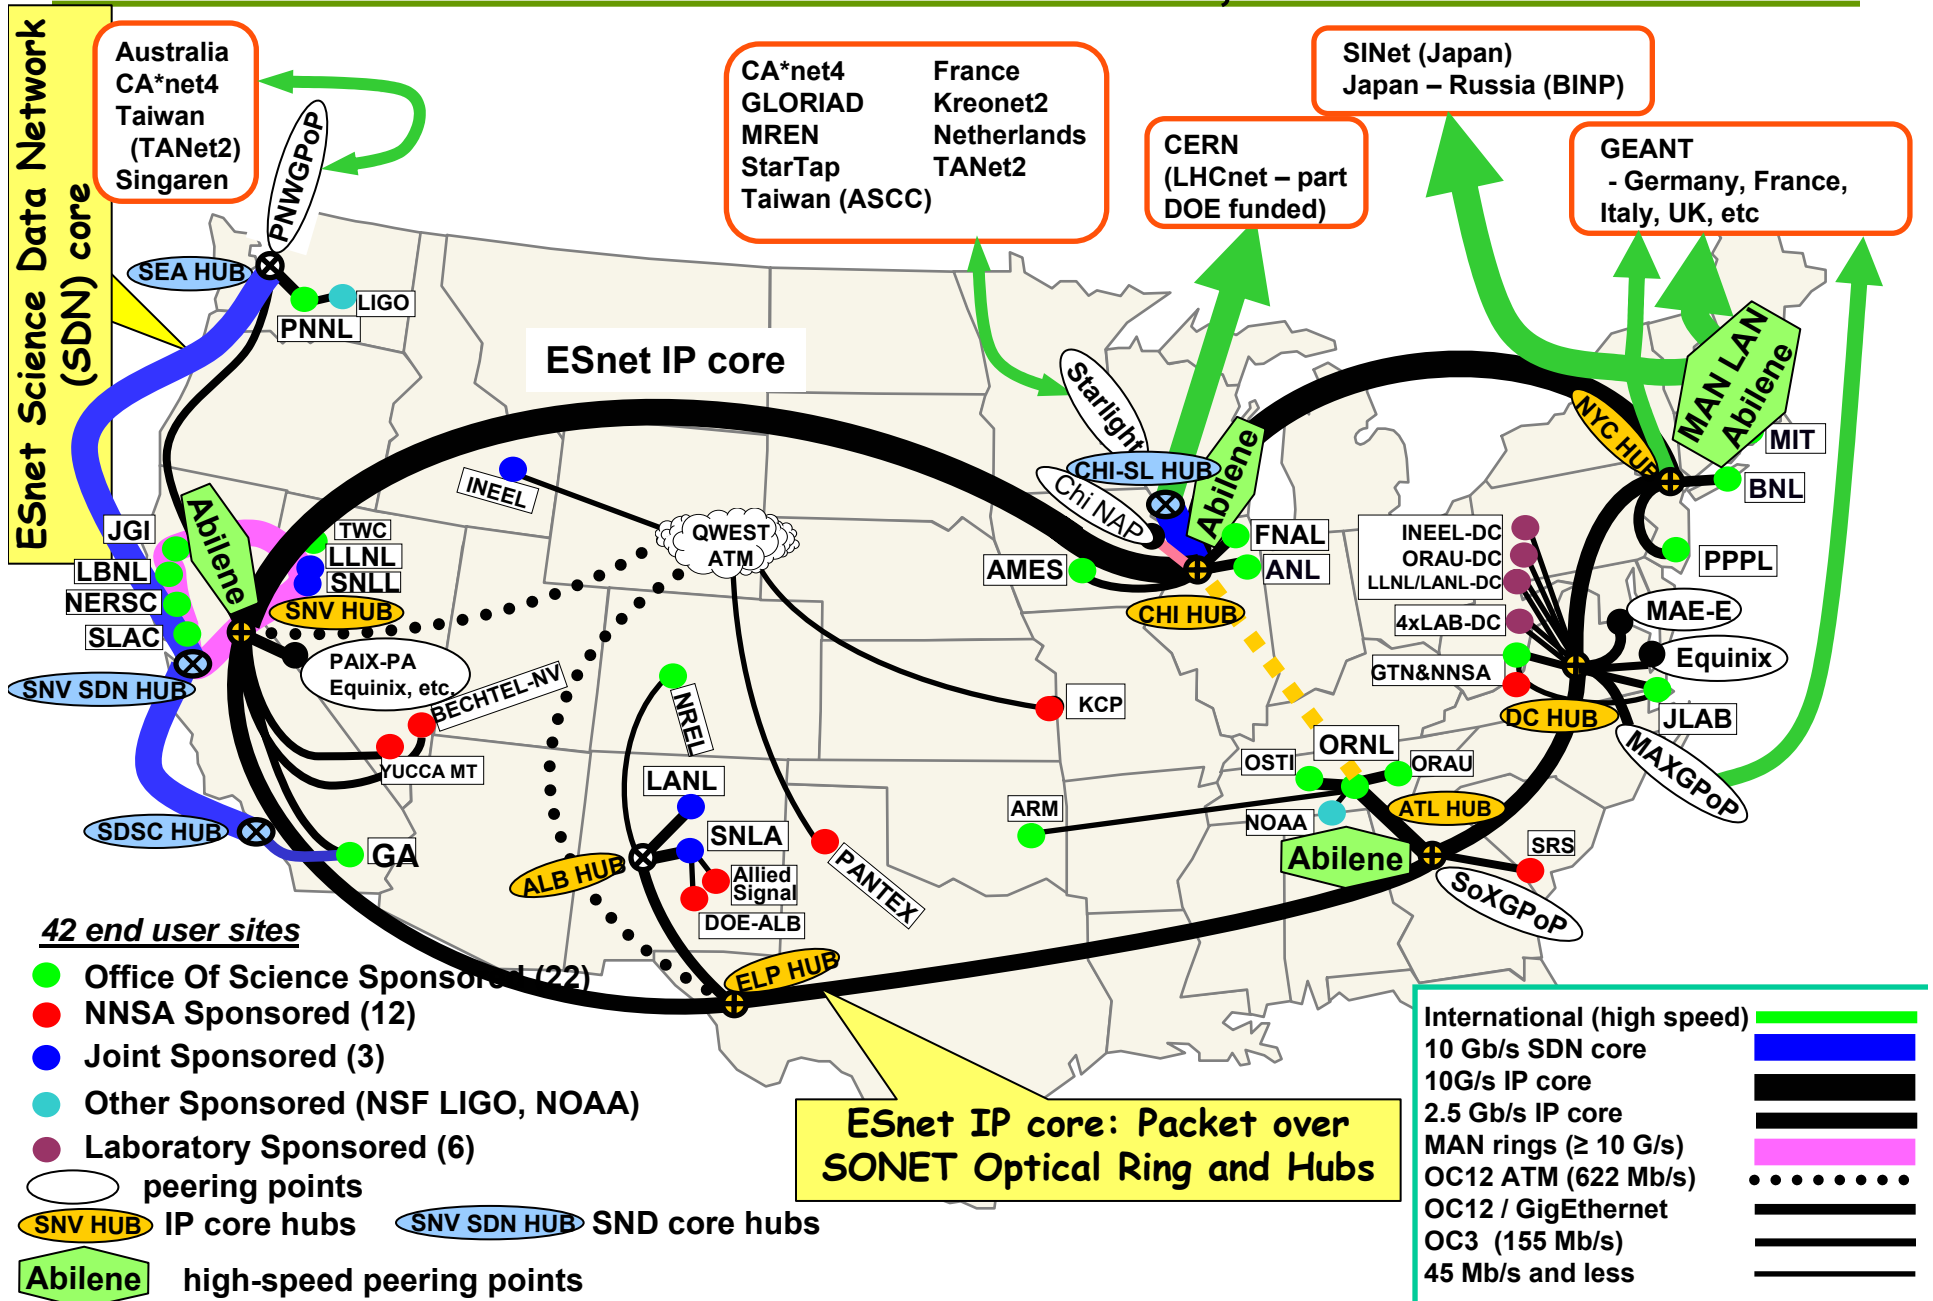
A Small Number of Science Users Account for a Significant Fraction of all ESnet Traffic

ESnet Top 100 *Host-to-Host* Flows, Feb., 2005

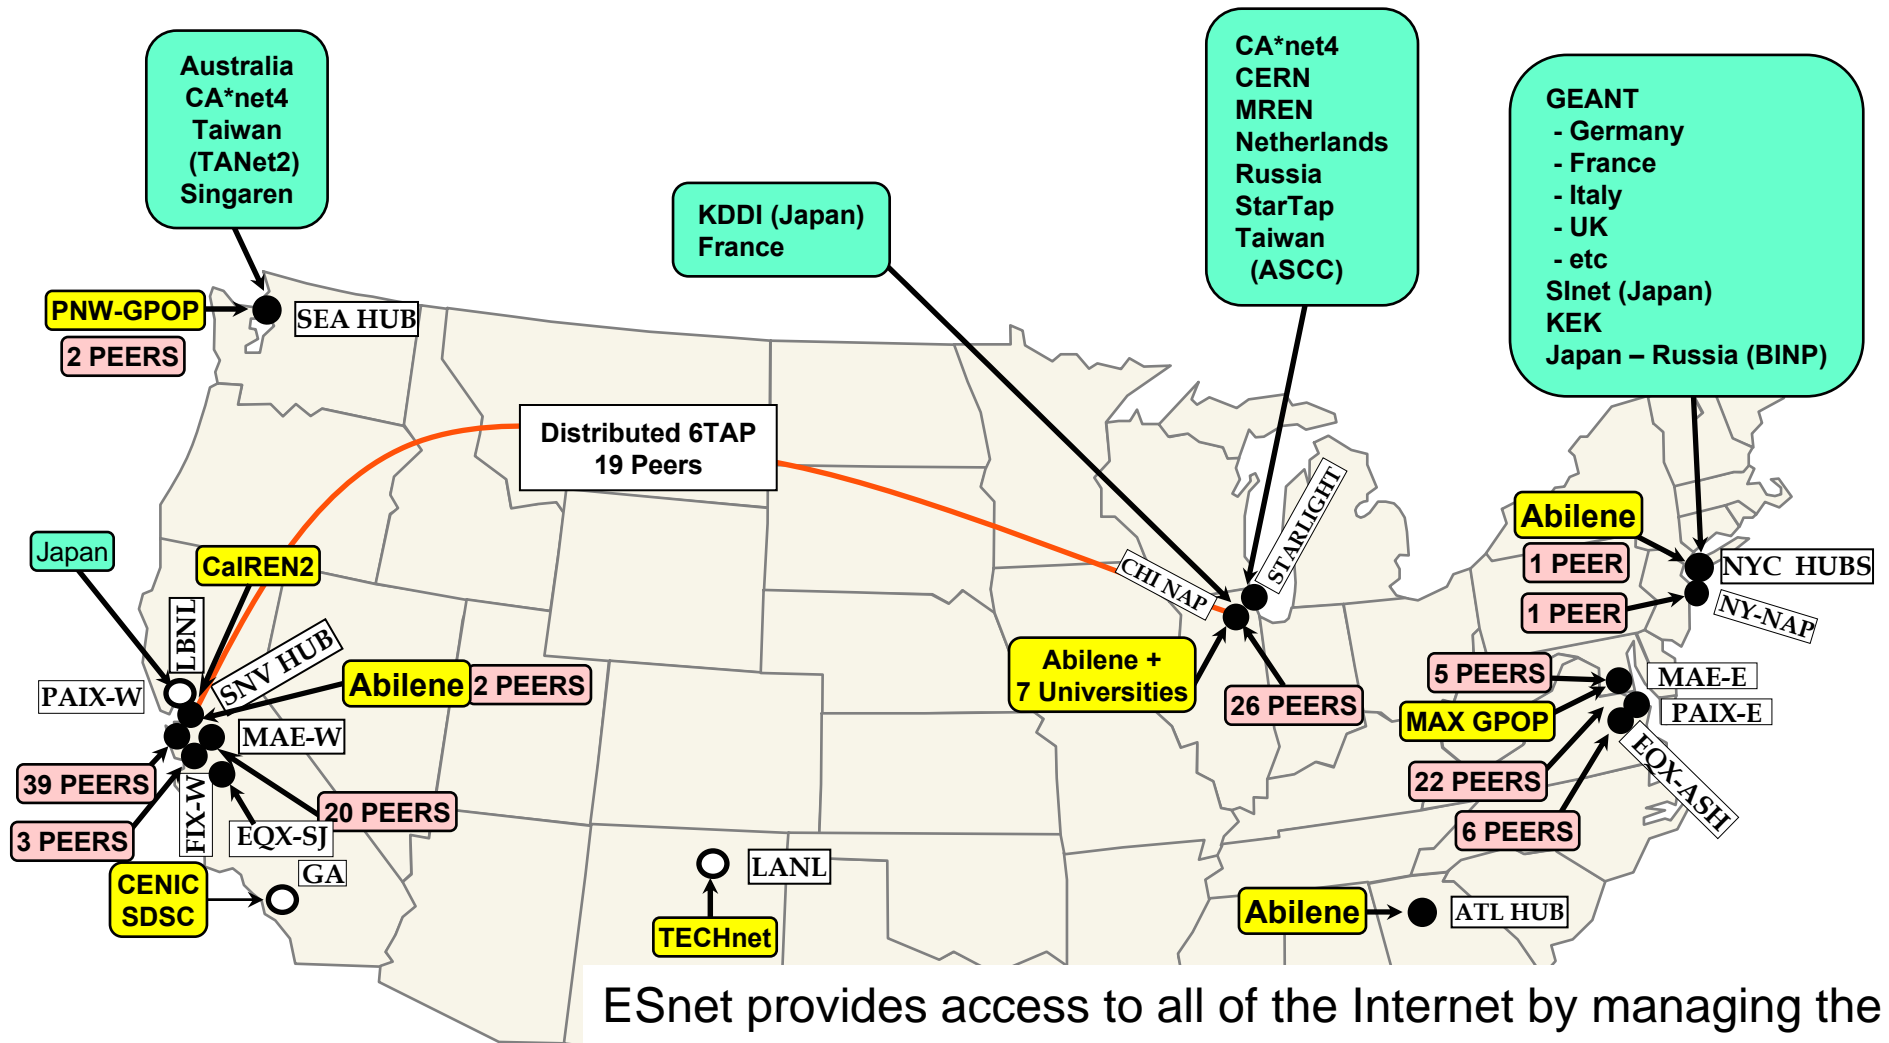


- Notes:**
- 1) This data does not include intra-Lab (LAN) traffic (ESnet ends at the Lab border routers, so science traffic on the Lab LANs is invisible to ESnet)
 - 2) Some Labs have private links that are not part of ESnet - that traffic is not represented here

ESnet High-Speed Physical Connectivity to DOE Facilities and Collaborators, Summer 2005



ESnet's Peering Infrastructure Connects the DOE Community With its Collaborators



ESnet Peering (connections to other networks)	University
	International
	Commercial

ESnet provides access to all of the Internet by managing the full complement of Global Internet routes (about 150,000) at 10 general/commercial peering points + high-speed peerings w/ Abilene and the international R&E networks. *This is a lot of work, and is very visible, but provides full access for DOE.*



How is ESnet Managed?

- ASCR/MICS provides core funding and program direction
- Engineered and operated by network staff at LBNL
- A community endeavor is essential for success
 - Strategic guidance from the Office of Science programs and other stakeholders
 - Energy Science Network Steering Committee (ESSC)
 - Network operation is a shared activity with the community
 - ESnet Site Coordinators Committee (ESCC)
 - Ensures the right operational “sociology” for success
 - Discussions underway to provide laboratories a voice
- Retrospective peer review (every 3-4 yr)
- Operational review (August, yearly)

Planning, configuration, and operation all have input and active participation from the DOE science community and the DOE Labs



Science Drivers for Networks

- The large-scale science that is the mission of the Office of Science is dependent on networks for
 - Sharing of massive amounts of data
 - Supporting thousands of collaborators world-wide
 - Distributed data processing
 - Distributed simulation, visualization, and computational steering
 - Distributed data management
- These issues were explored in two Office of Science workshops that formulated networking requirements to meet the needs of the science programs



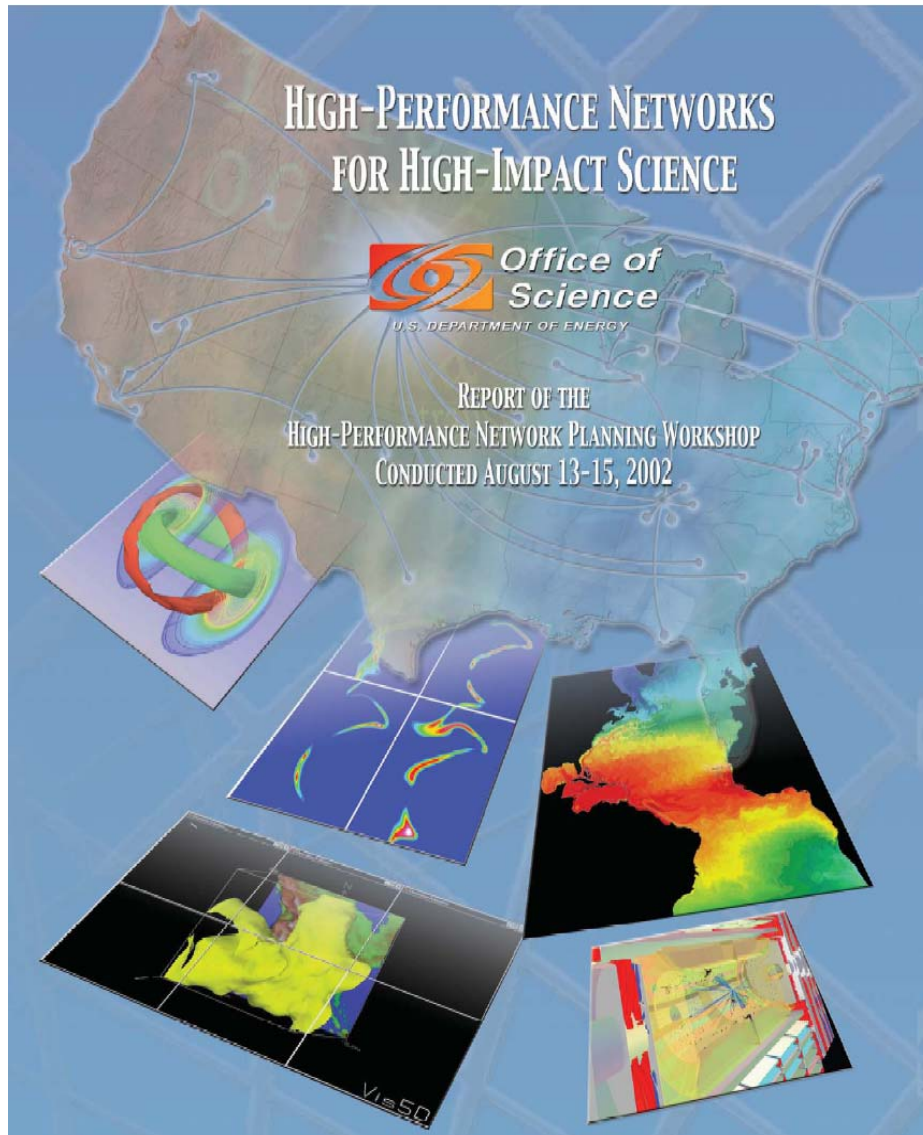
What's a possible future?

VISION –

- A seamless, high-performance network infrastructure in which science applications and advanced facilities are "n-way" interconnected to terascale computing, petascale storage, and high-end visualization capabilities.
- This advanced network facilitates collaborations among researchers and interactions between researchers and experimental and computational resources,
- Science, especially large-scale science, moves to a new regime that eliminates isolation, discourages redundancy, and promotes rapid scientific progress through the interplay of theory, simulation, and experiment.



Network and Middleware Needs of DOE Science



- **Focused on science requirements driving advanced network infrastructure**
 - *Middleware research*
 - *Network research*
 - *Network governance model*
- **Requirements for DOE science developed for a representative cross-section of the OSC scientific community**
 - *Climate simulation*
 - *Spallation Neutron Source facility*
 - *Macromolecular Crystallography*
 - *High Energy Physics experiments*
 - *Magnetic Fusion Energy Sciences*
 - *Chemical Sciences*
 - *Bioinformatics*

Evolving Quantitative Science Requirements for Networks

Science Areas considered in the Workshop (not Nuclear Physics and Supercomputing)	Today <i>End2End</i> Throughput	5 years End2End Documented Throughput Requirements	5-10 Years End2End <u>Estimated</u> Throughput Requirements	Remarks
High Energy Physics	0.5 Gb/s	100 Gb/s	1000 Gb/s	high bulk throughput
Climate (Data & Computation)	0.5 Gb/s	160-200 Gb/s	N x 1000 Gb/s	high bulk throughput
SNS NanoScience	Not yet started	1 Gb/s	1000 Gb/s + QoS for control channel	remote control and time critical throughput
Fusion Energy	0.066 Gb/s (500 MB/s burst)	0.198 Gb/s (500MB/ 20 sec. burst)	N x 1000 Gb/s	time critical throughput
Astrophysics	0.013 Gb/s (1 TBy/week)	N*N multicast	1000 Gb/s	computational steering and collaborations
Genomics Data & Computation	0.091 Gb/s (1 TBy/day)	100s of users	1000 Gb/s + QoS for control channel	high throughput and steering



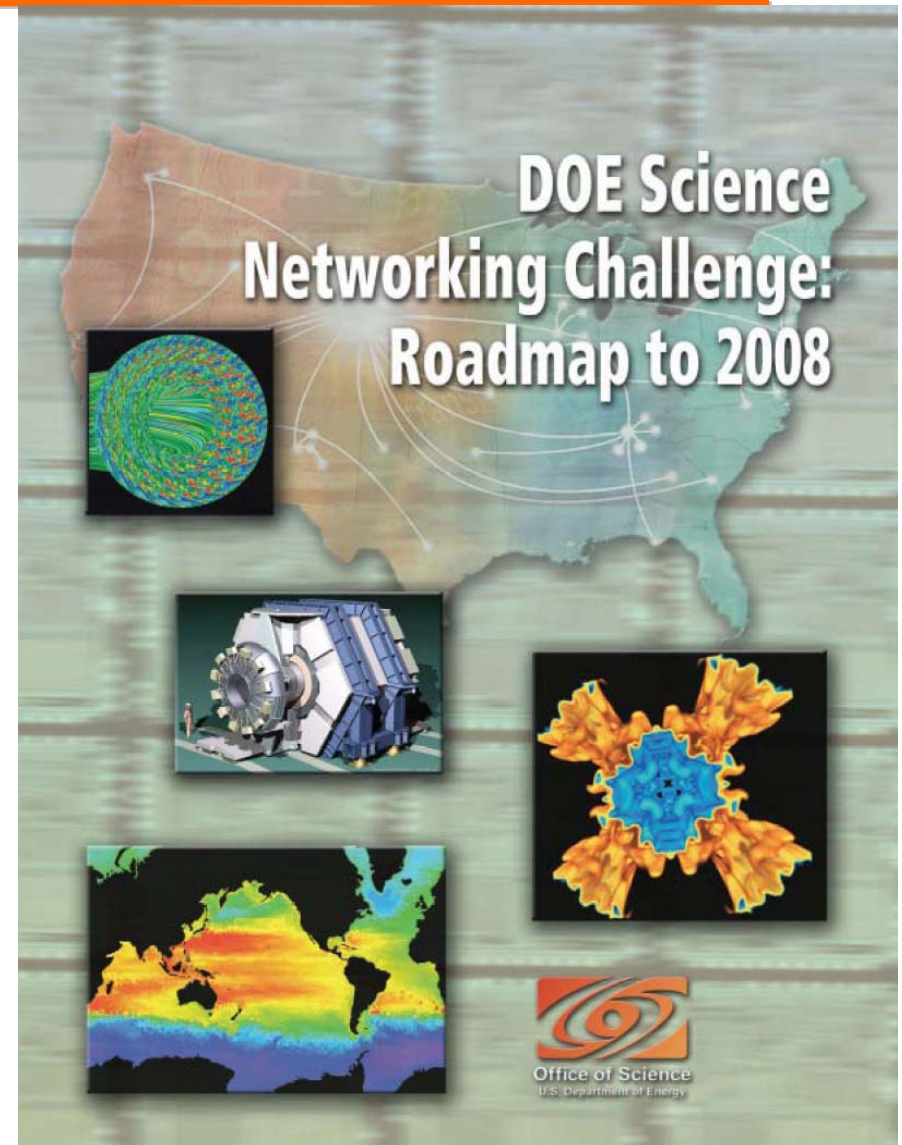
New Strategic Directions to Address Needs of DOE Science

Focused on what is needed to achieve the science driven network requirements of the previous workshop

- **THE #1 DRIVER for continuing advancements in networking and middleware** – Petabyte scale experimental and simulation data systems will be increasing to exabyte scale data systems.
 - Bioinformatics, Climate, LHC, etc.
- Computational systems that process or produce data continue to advance with Moore's Law
-

Organized by the ESSC

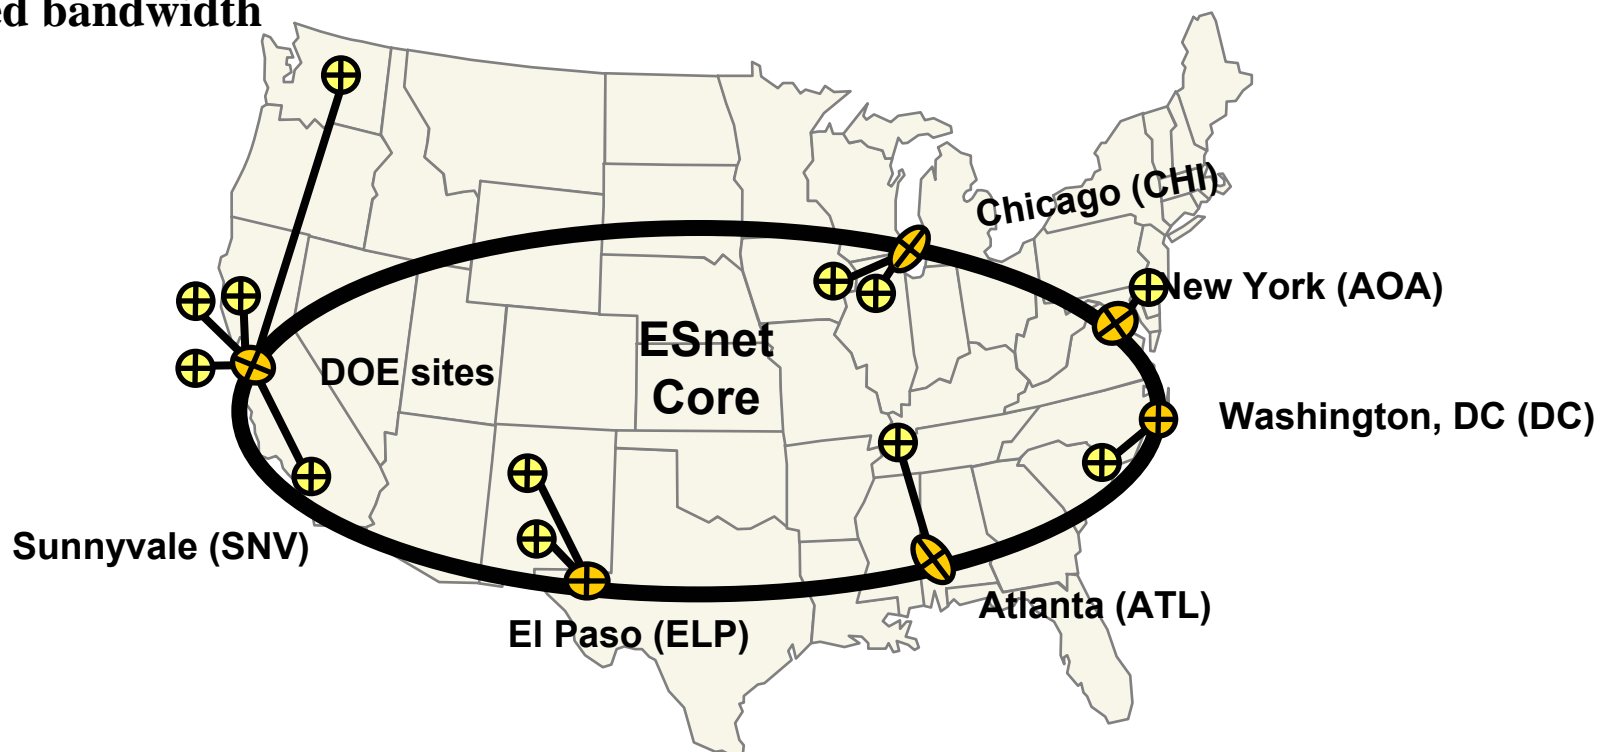
- Workshop Chair, Roy Whitney, JLAB
- Workshop Editors, Roy Whitney, JLAB; Larry Price, ANL





ESnet Evolution

- **With the current architecture ESnet cannot address**
 - Increasing reliability requirements (Labs and science experiments are insisting on network redundancy)
 - Long-term bandwidth needs (e.g. LHC will need dedicated 10/20/30/40 Gb/s into and out of FNAL & BNL)
- **The current core ring cannot handle the anticipated large science data flows at affordable cost**
- **The current point-to-point tail circuits are neither reliable nor scalable to the required bandwidth**





Evolution of ESnet

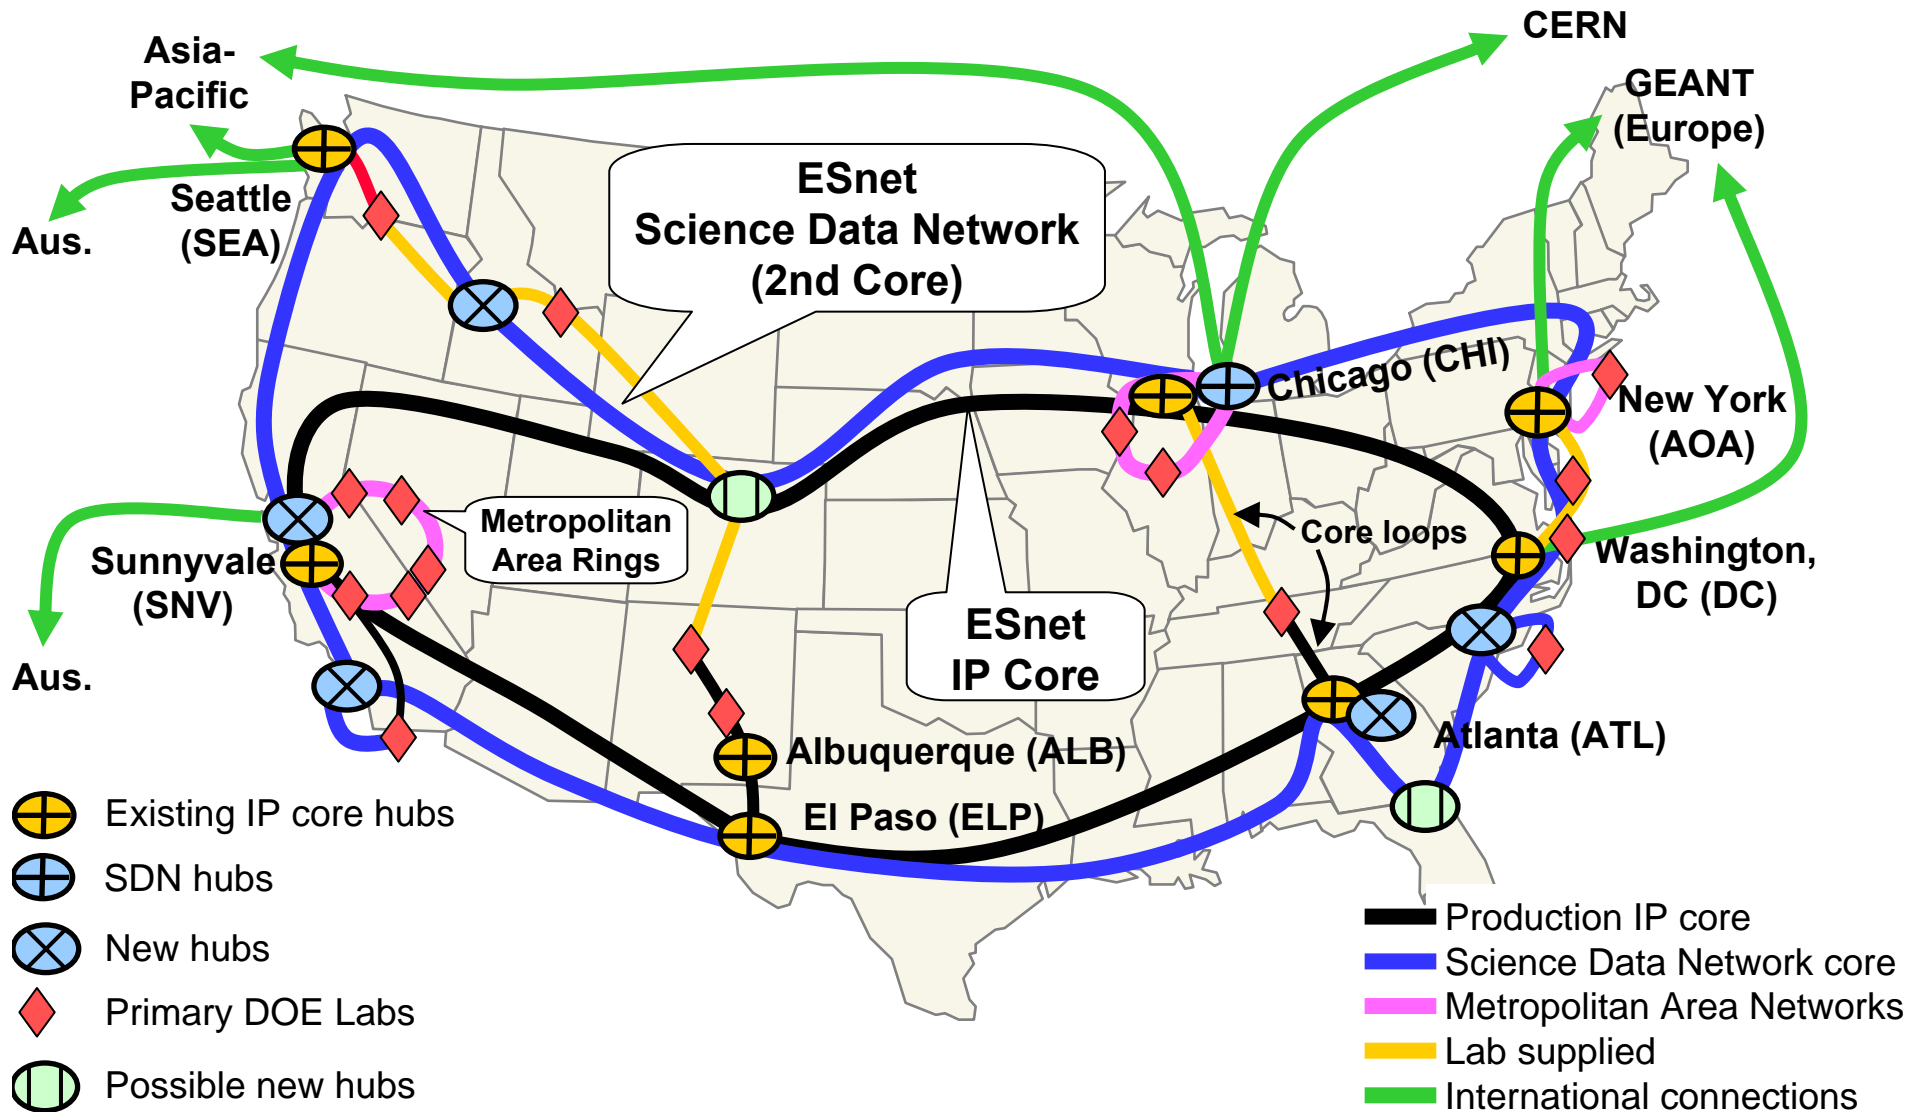
- Upgrading ESnet to accommodate the anticipated increase from the current 100%/yr traffic growth to 300%/yr over the next 5-10 years is priority number 7 out of 20 in DOE's "Facilities for the Future of Science – A Twenty Year Outlook"
- Based on the requirements, ESnet must address
 - I. Capable, scalable, and reliable production IP networking
 - University and international collaborator connectivity
 - Scalable, reliable, and high bandwidth site connectivity
 - II. Network support of high-impact science
 - provisioned circuits with guaranteed quality of service (e.g. dedicated bandwidth)
 - III. Evolution to optical switched networks
 - Partnership with UltraScienceNet
 - Close collaboration with the network R&D community
 - IV. Science Services to support Grids, laboratories, etc



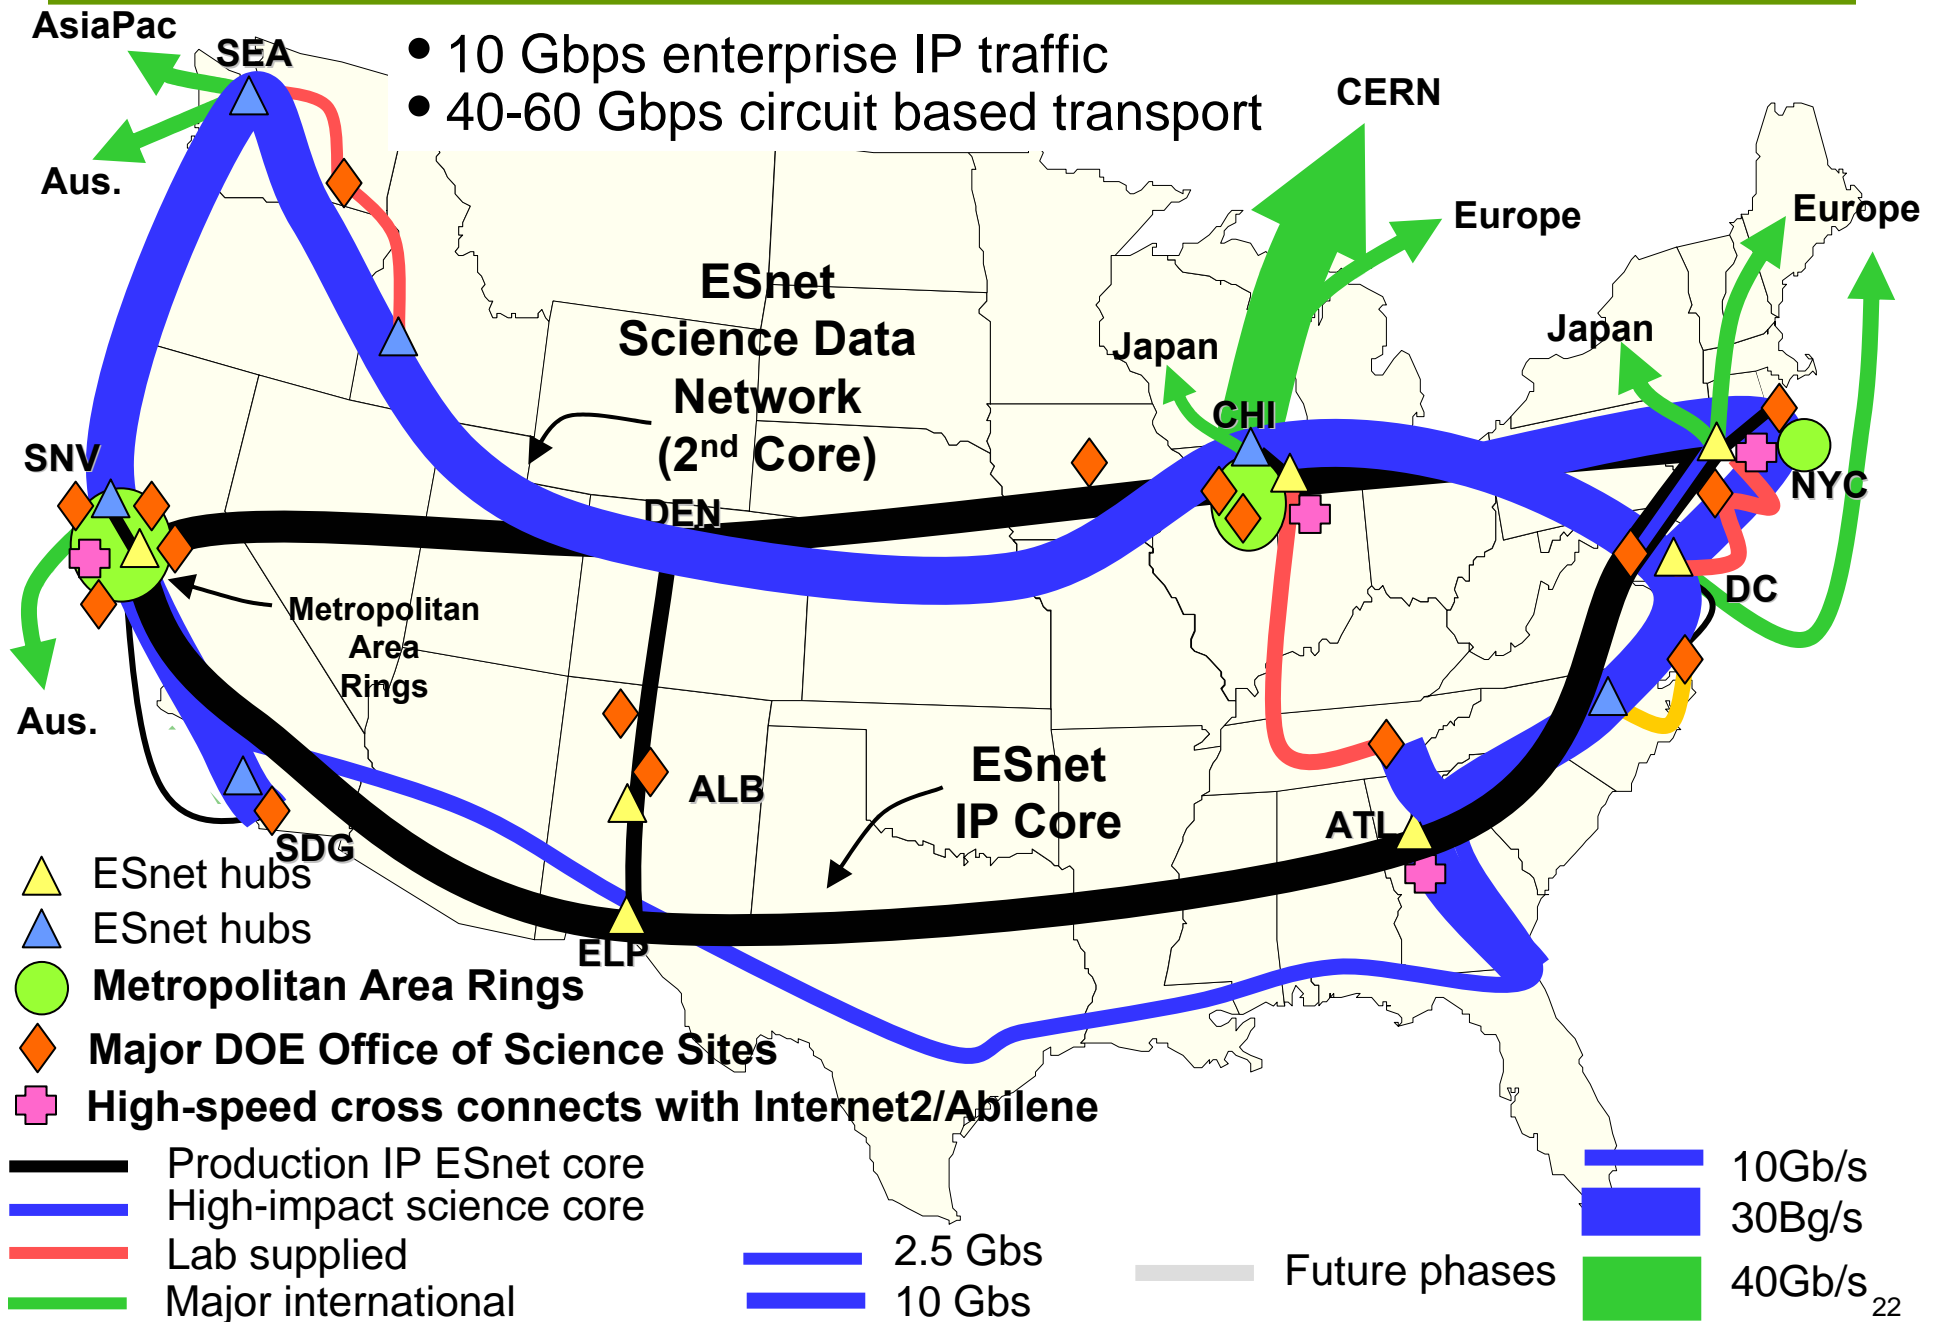
A New ESnet Architecture

- Goals
 - Fully redundant connectivity for every site
 - High-speed access for every site (at least 20 Gb/s)
- Three part strategy
 - 1) **Metropolitan Area Network** (MAN) rings provide dual site connectivity and much higher site-to-core bandwidth
 - 2) A **Science Data Network** core for
 - multiply connected MAN rings for protection against hub failure
 - expanded capacity for science data
 - a platform for provisioned, guaranteed bandwidth circuits
 - alternate path for production IP traffic
 - carrier circuit and fiber access neutral hubs
 - 3) **High-reliability IP core** (e.g. the current ESnet core)

New ESnet Strategy: Science Data Network + IP Core + MANs



ESnet Meeting Science Requirements – 2007/2008





Immediate Issues

- FNAL must increase its production traffic bandwidth to accommodate D0 data analysis
- ORNL's leadership class machine is "ASCR's (OSC's?) highest priority"
 - Current plans will provide 12.5 Gbps
 - Network requirements document is being prepared
- BNL must have 2.5 Gbps in 2005, 10 Gbps in 2006, and 40 Gbps in 2008
- PPPL must increase bandwidth to accommodate the new ITER Office
- JLab, PNNL, SLAC, NERSC,



Conclusions

- ESnet is an infrastructure that is critical to DOE's science mission
- Focused on the Office of Science Labs, but serves many other parts of DOE
- ESnet is working hard to meet the current and future networking need of DOE mission science in several ways:
 - Evolving a new high speed, high reliability, leveraged architecture
 - Championing several new initiatives which will keep ESnet's contributions relevant to the needs of our community
 - Grid middleware services for large numbers of users are hard – but they can be provided if careful attention is paid to scaling



Where do you come in?

- Early identification of requirements
 - Evolving programs
 - New facilities
- Participation in management activities



References – DOE Network Related Planning Workshops

- High Performance Network Planning Workshop, August 2002
<http://www.doecollaboratory.org/meetings/hpnpw>
- DOE Science Networking Roadmap Meeting, June 2003
<http://www.es.net/hypertext/welcome/pr/Roadmap/index.html>
- 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>
- Science Case for Large Scale Simulation, June 2003
<http://www.pnl.gov/scales/>
- Workshop on the Road Map for the Revitalization of High End Computing, June 2003
http://www.sc.doe.gov/ascr/20040510_hecrtf.pdf (public report)
- ASCR Strategic Planning Workshop, July 2003
<http://www.sc.doe.gov/ascr/ASCRstrategicplan07304final.pdf>
- Planning Workshops-Office of Science Data-Management Strategy, March & May 2004
 - <http://www-conf.slac.stanford.edu/dmw2004> (report coming soon)