


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

Energy Sciences Network Enabling Virtual Science

June 9, 2009

Steve Cotter steve@es.net
Dept. Head, Energy Sciences Network
Lawrence Berkeley National Lab






The Energy Sciences Network

The Department of Energy's Office of Science is one of the largest supporters of basic research in the physical sciences in the U.S.

- Directly supports the research of some 15,000 scientists, postdocs and graduate students at DOE laboratories, universities, other Federal agencies, and industry worldwide
- Operates major scientific facilities at DOE laboratories that that have participation by the US and international research and education (R&E) community

Established in 1985, ESnet is the Department of Energy's science networking program whose responsibility is to provide the network infrastructure supporting the missions of the Office of Science

- Enabling a new era in scientific discovery as we tackle global issues like climate change, alternative energy/fuels and understanding the origins of the universe.





ESnet: Driven by Science


Networking needs of researchers are far different than commercial users.

Therefore, ESnet regularly explores the plans and processes of major stakeholders to understand:


- The extreme data characteristics of instruments and facilities
 - How much data will be generated by instruments coming on-line over the next 5-10 years?
- The future process of science
 - How and where will the new data be analyzed and used – that is, how will the process of doing science change over 5-10 years?
- SC Networking Requirements Workshops
 - 2 workshops a year, rotating thru BES, BER, FES, NP, HEP, ASCR communities
 - Workshop reports: <http://www.es.net/hypertext/requirements.html>

Observing current and historical network traffic trends

- What do the trends in network patterns predict for future network needs?



Science: Driven by Data







Scientific data sets are growing exponentially

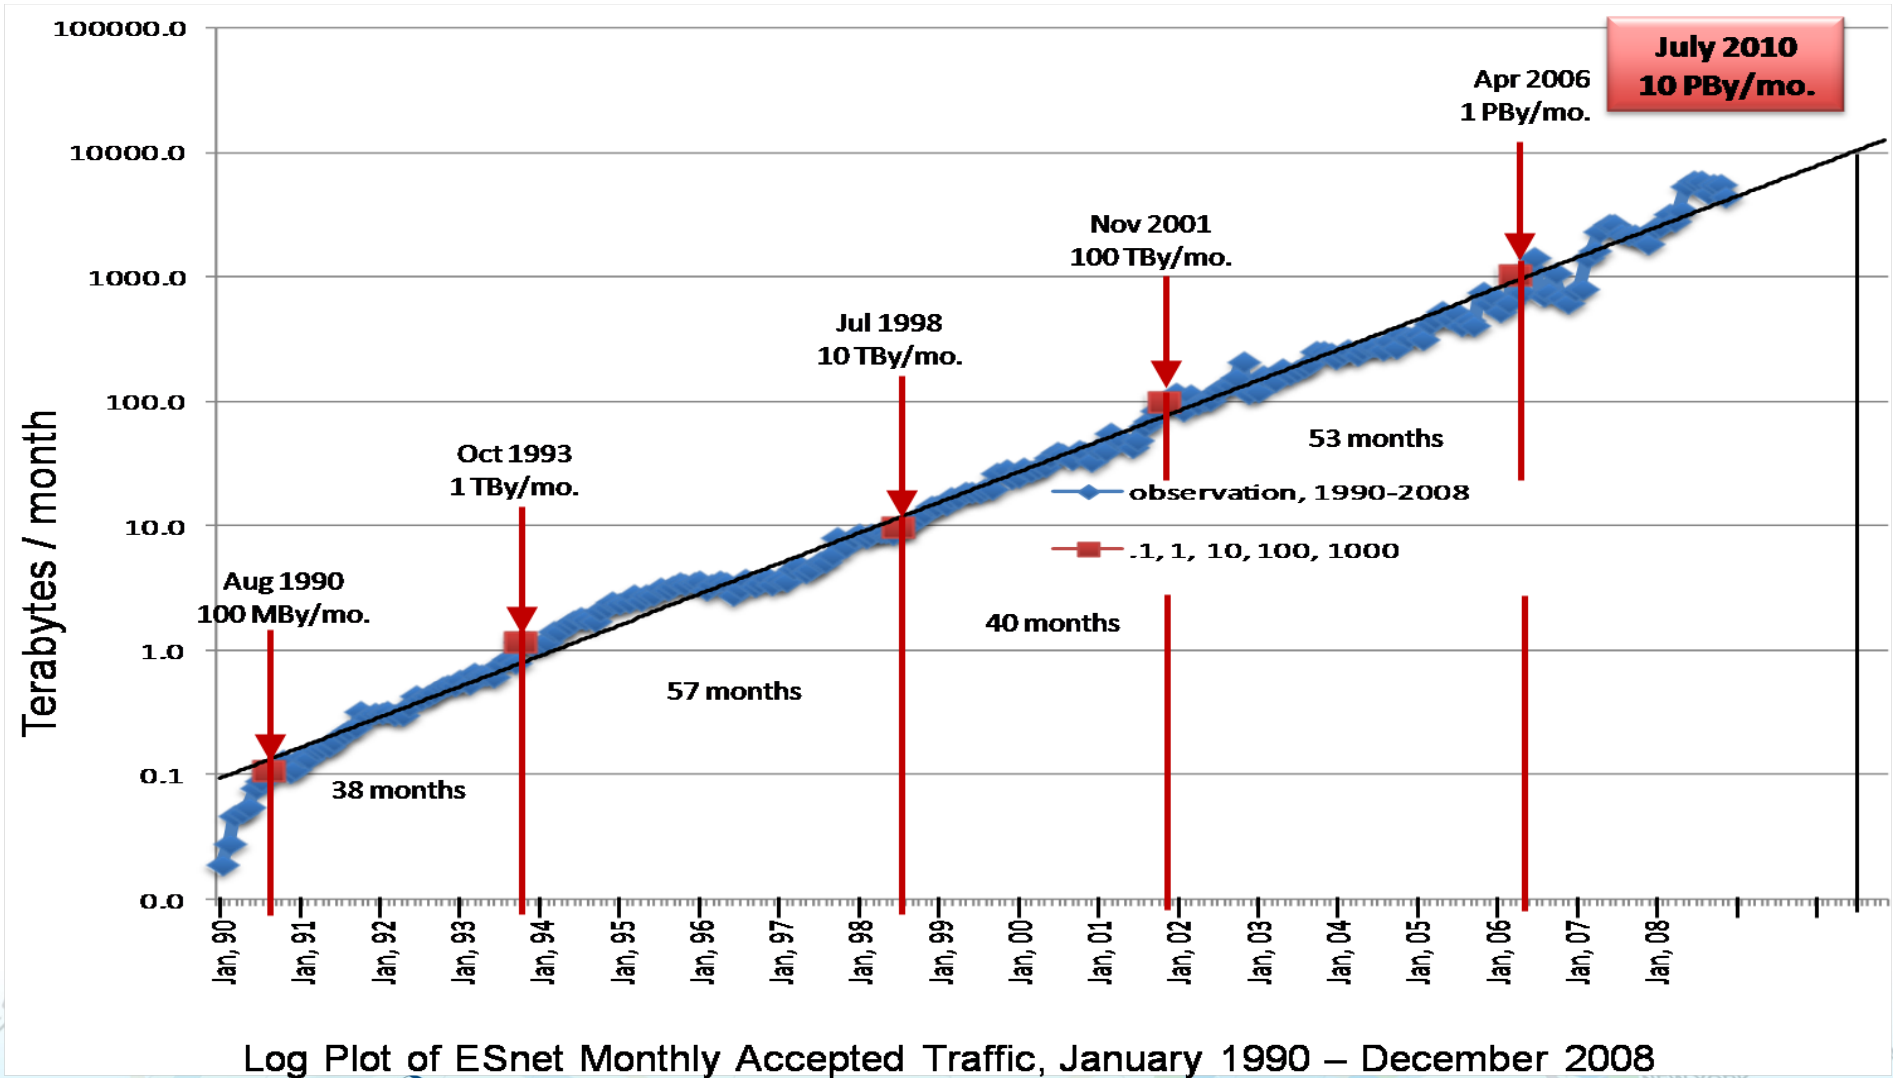
- Ability to generate data is exceeding our ability to store and analyze
- Simulation systems and some observational devices grow in capability with Moore's Law

Petabyte (PB) data sets will soon be common:

- Climate modeling: estimates of the next IPCC data is in 10s of petabytes
- Genome: JGI alone will have 0.5 petabyte of data this year and double each year
- Particle physics: LHC is projected to produce 16 petabytes of data per year
- Astrophysics: LSST and others will produce 5 petabytes/year



ESnet Challenges






Solution: ESnet4

ESnet4: a unique hybrid packet- & circuit-switched network infrastructure specifically designed to handle massive amounts of data

- Combines the flexibility and resiliency of IP routed networks with the deterministic, high-speed capability of a circuit-switched infrastructure

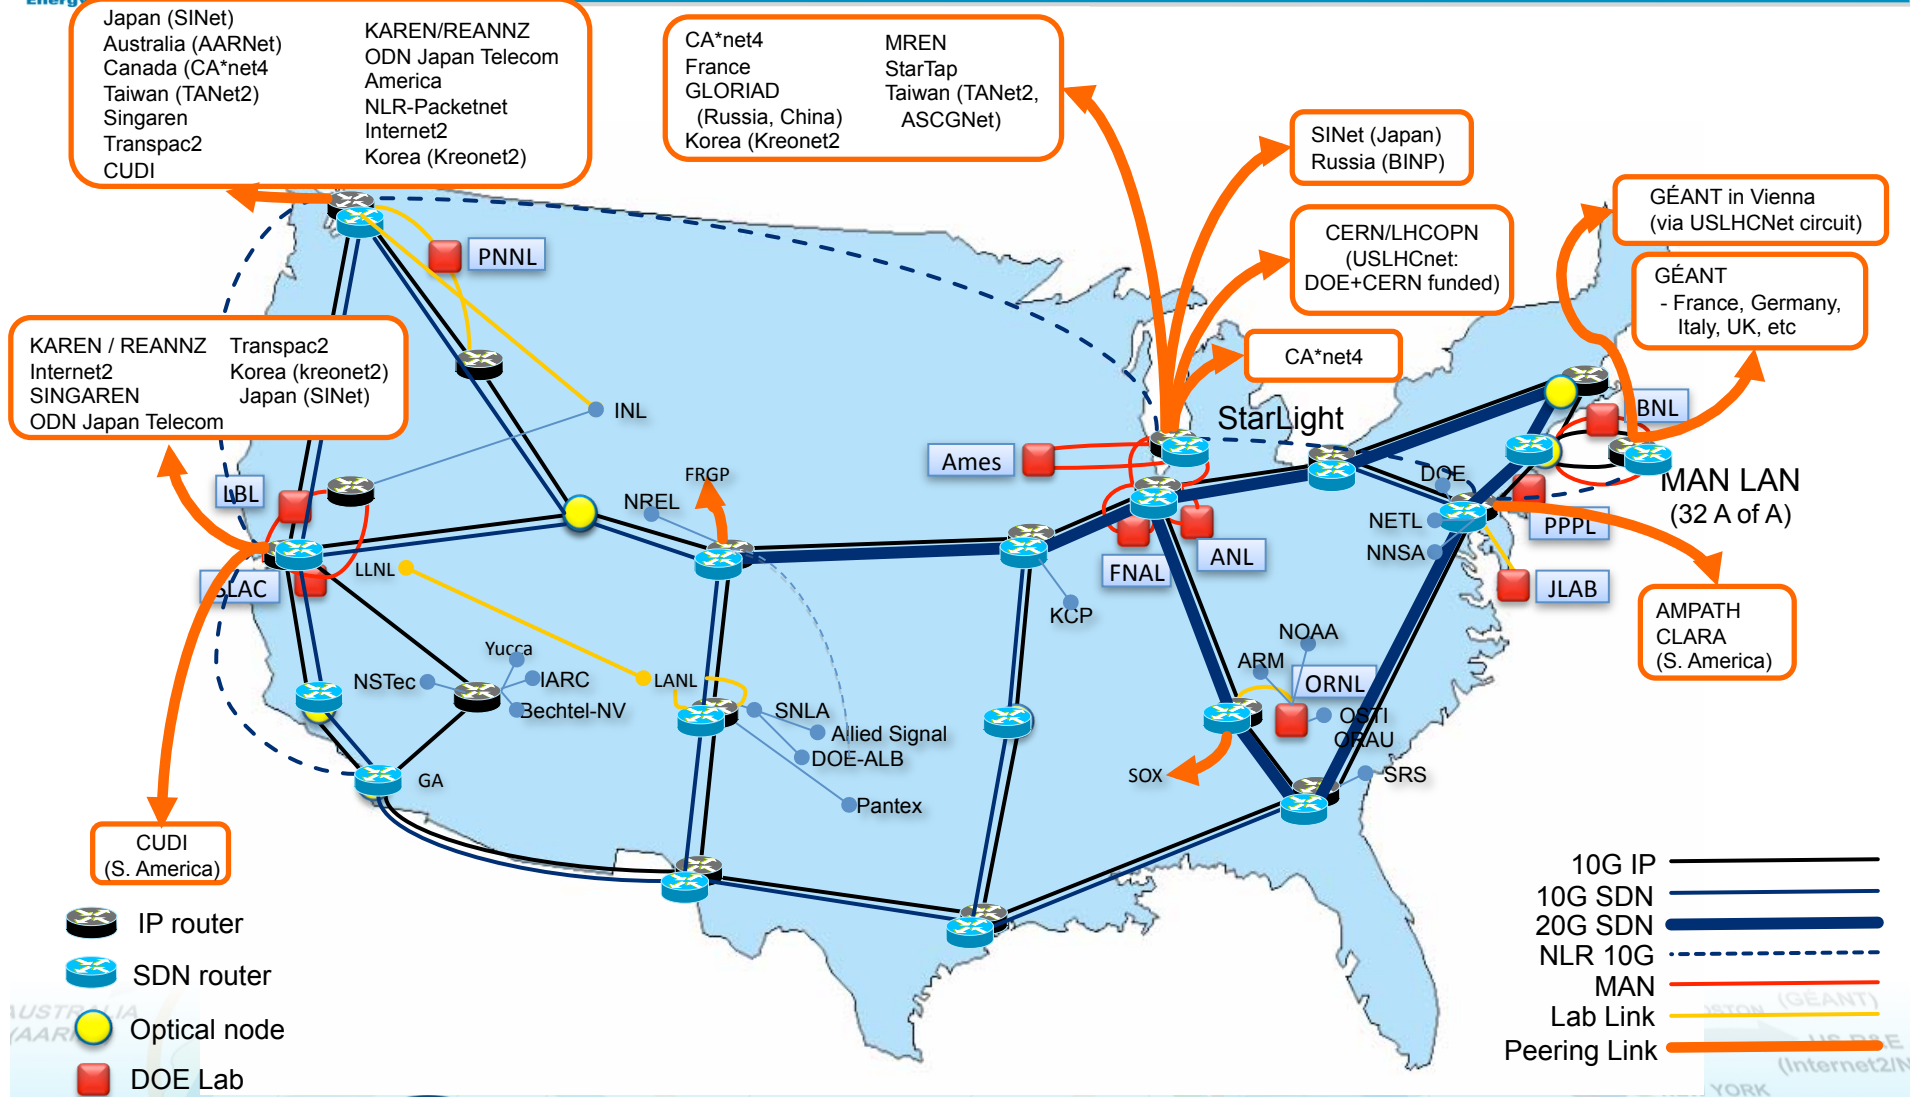
Used commercially available technologies to create two logically separate networks over which traffic seamlessly switches

- IP Network: One network for IP traffic using a single 10 Gbps circuit over which ESnet provides audio/videoconferencing and data collaboration tools
- Science Data Network: Circuit-switched core network consisting of multiple 10 Gbps circuits connecting directly with other high-speed R&E networks and utilizing Layer 2/3 switches





ESnet4 – June 2009






Solution: ESnet4

Key components tying the two networks together:

- ESnet developed and implemented the On-Demand Secure Circuits and Advance Reservation System (OSCARS) protocol which spans both networks.
 - Allows scientists to request dedicated bandwidth to move large amounts of data – up to terabytes at a time – across multiple network domains.
- Active participant in the perfSONAR consortium developing an open, modular infrastructure of services and applications that enables the gathering and sharing of network performance information, and facilitates troubleshooting of problems across network domains.






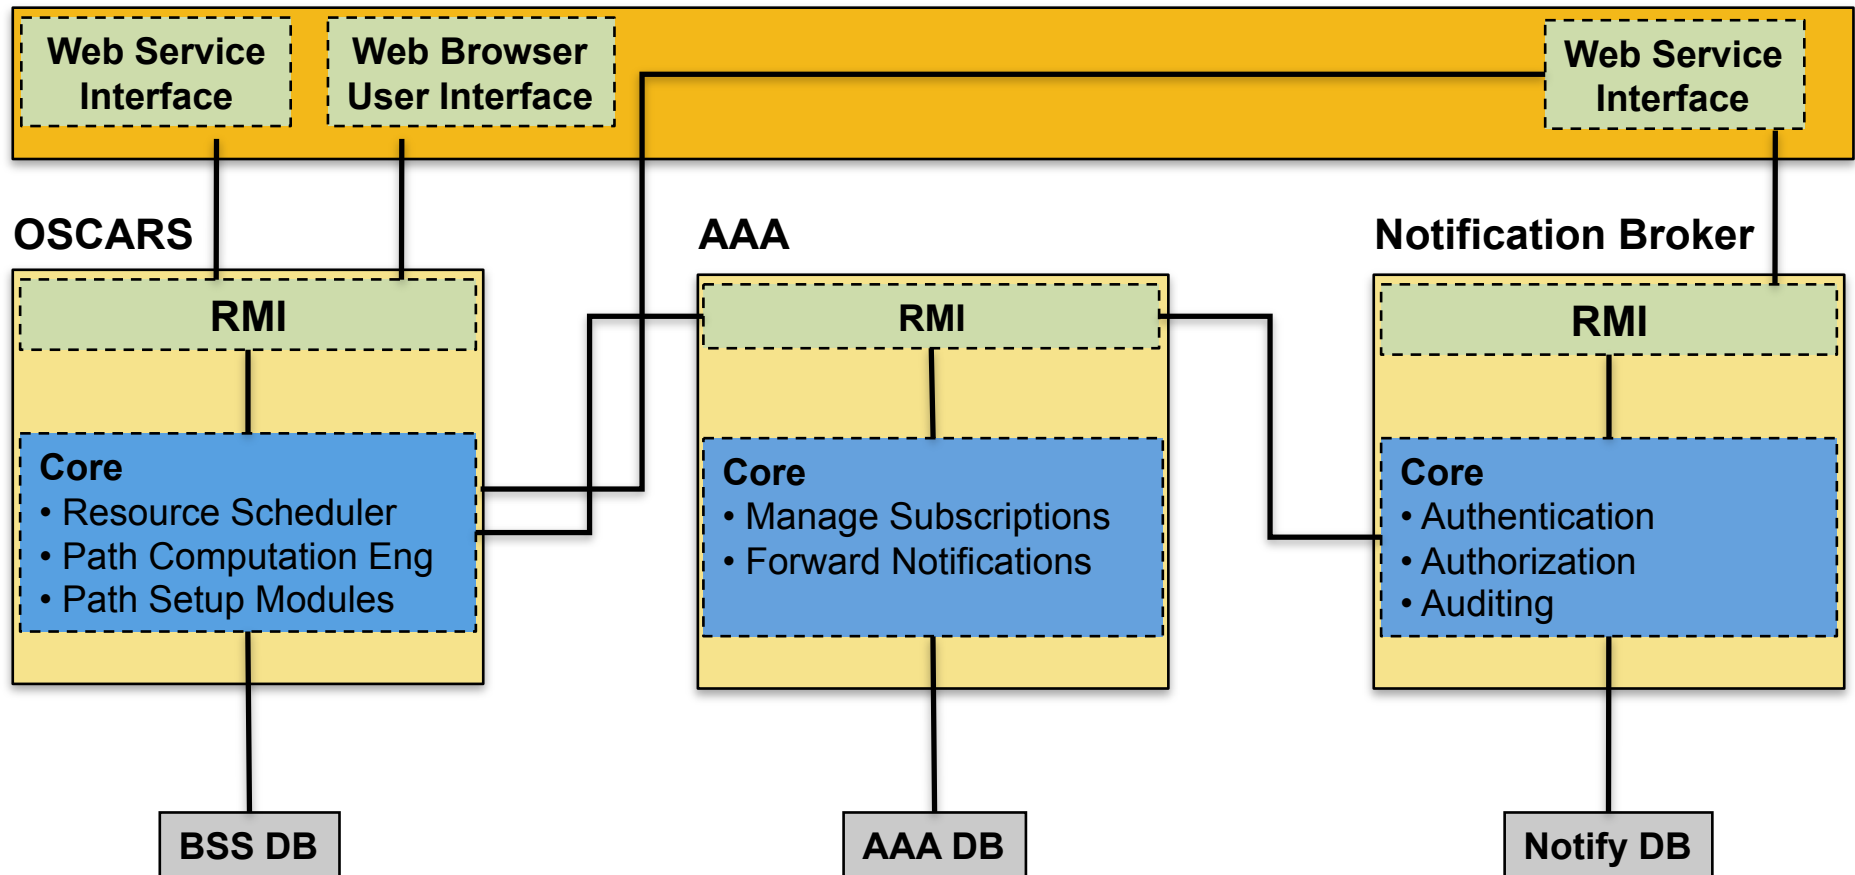
OSCARS: Multi-Domain Virtual Circuit Service

OSCARS Services

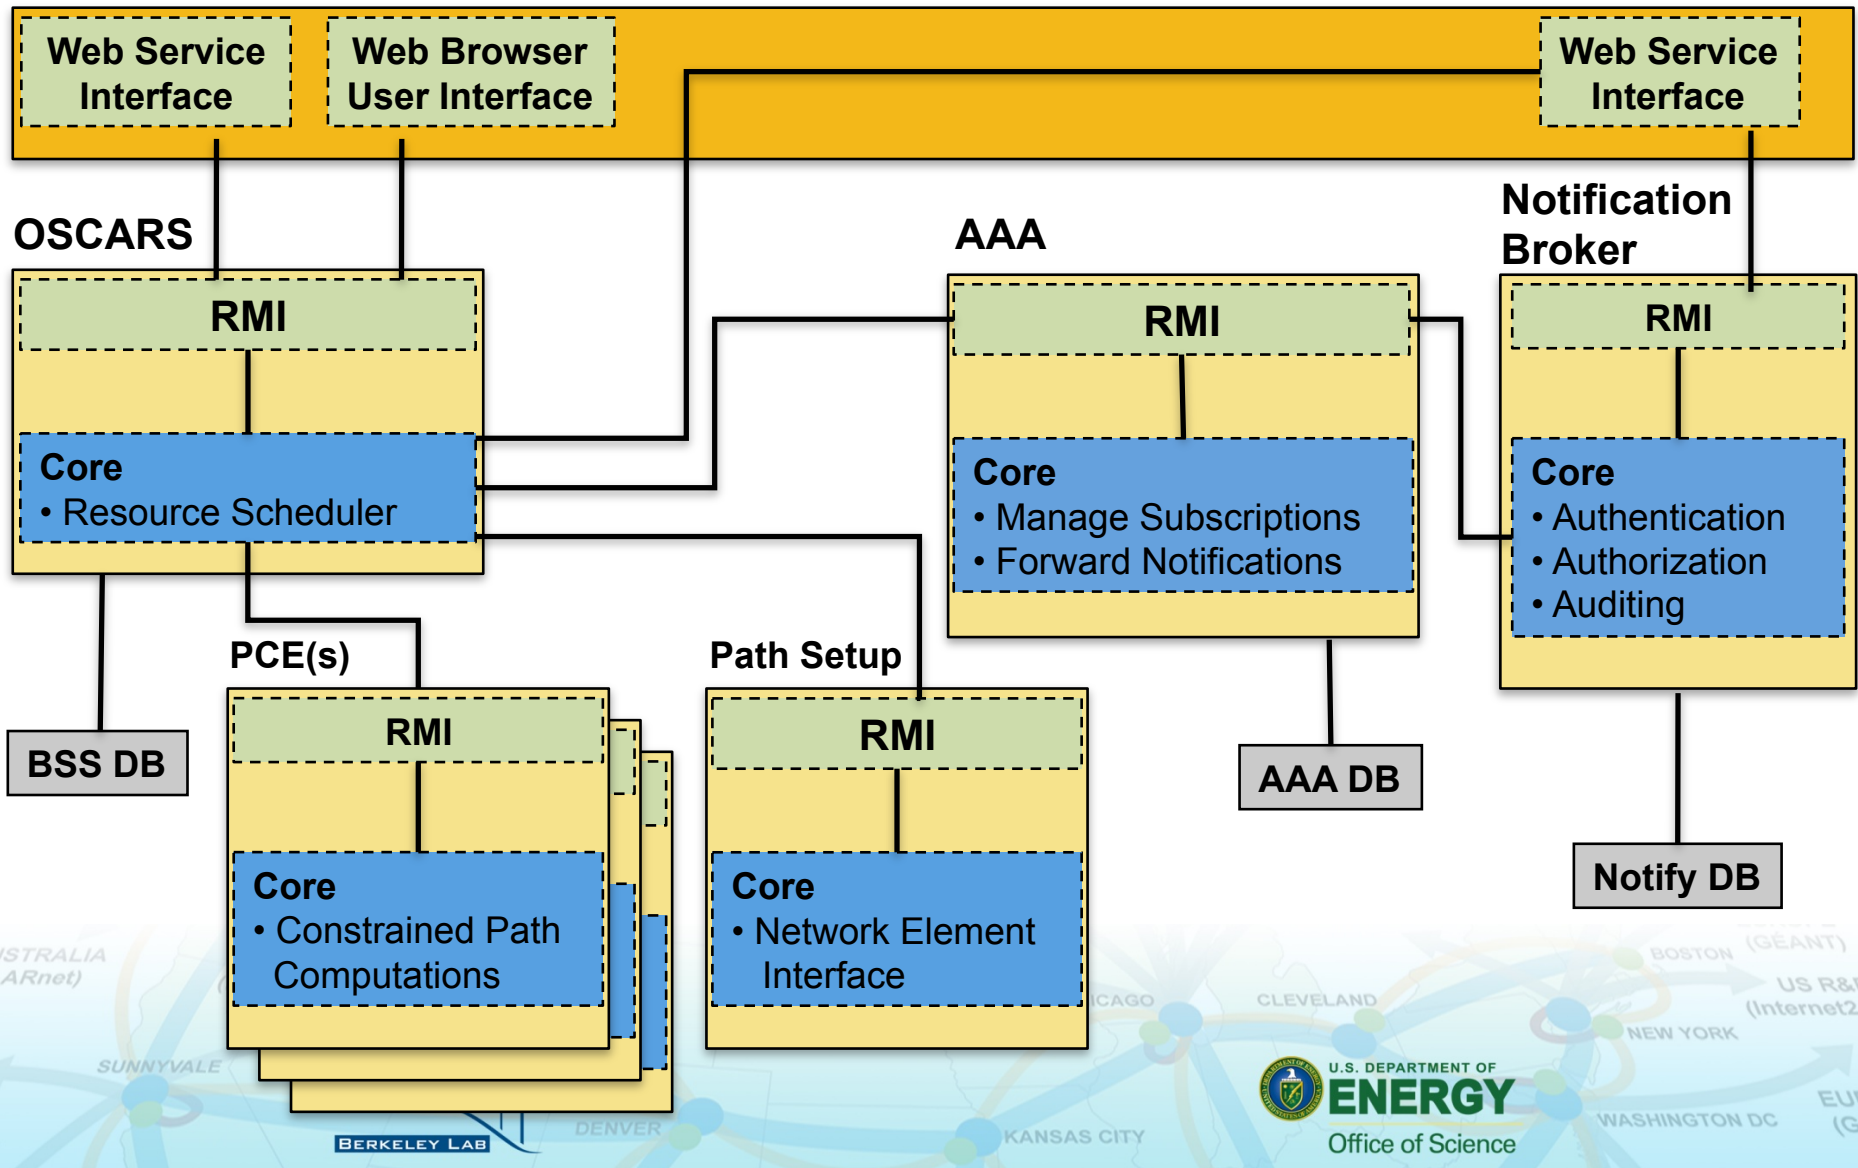
- Guaranteed bandwidth with resiliency: User specified bandwidth for primary and backup paths - requested and managed in a Web Services framework
- Traffic isolation: Allows for high-performance, non-standard transport mechanisms that cannot co-exist with commodity TCP-based transport
- Traffic engineering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



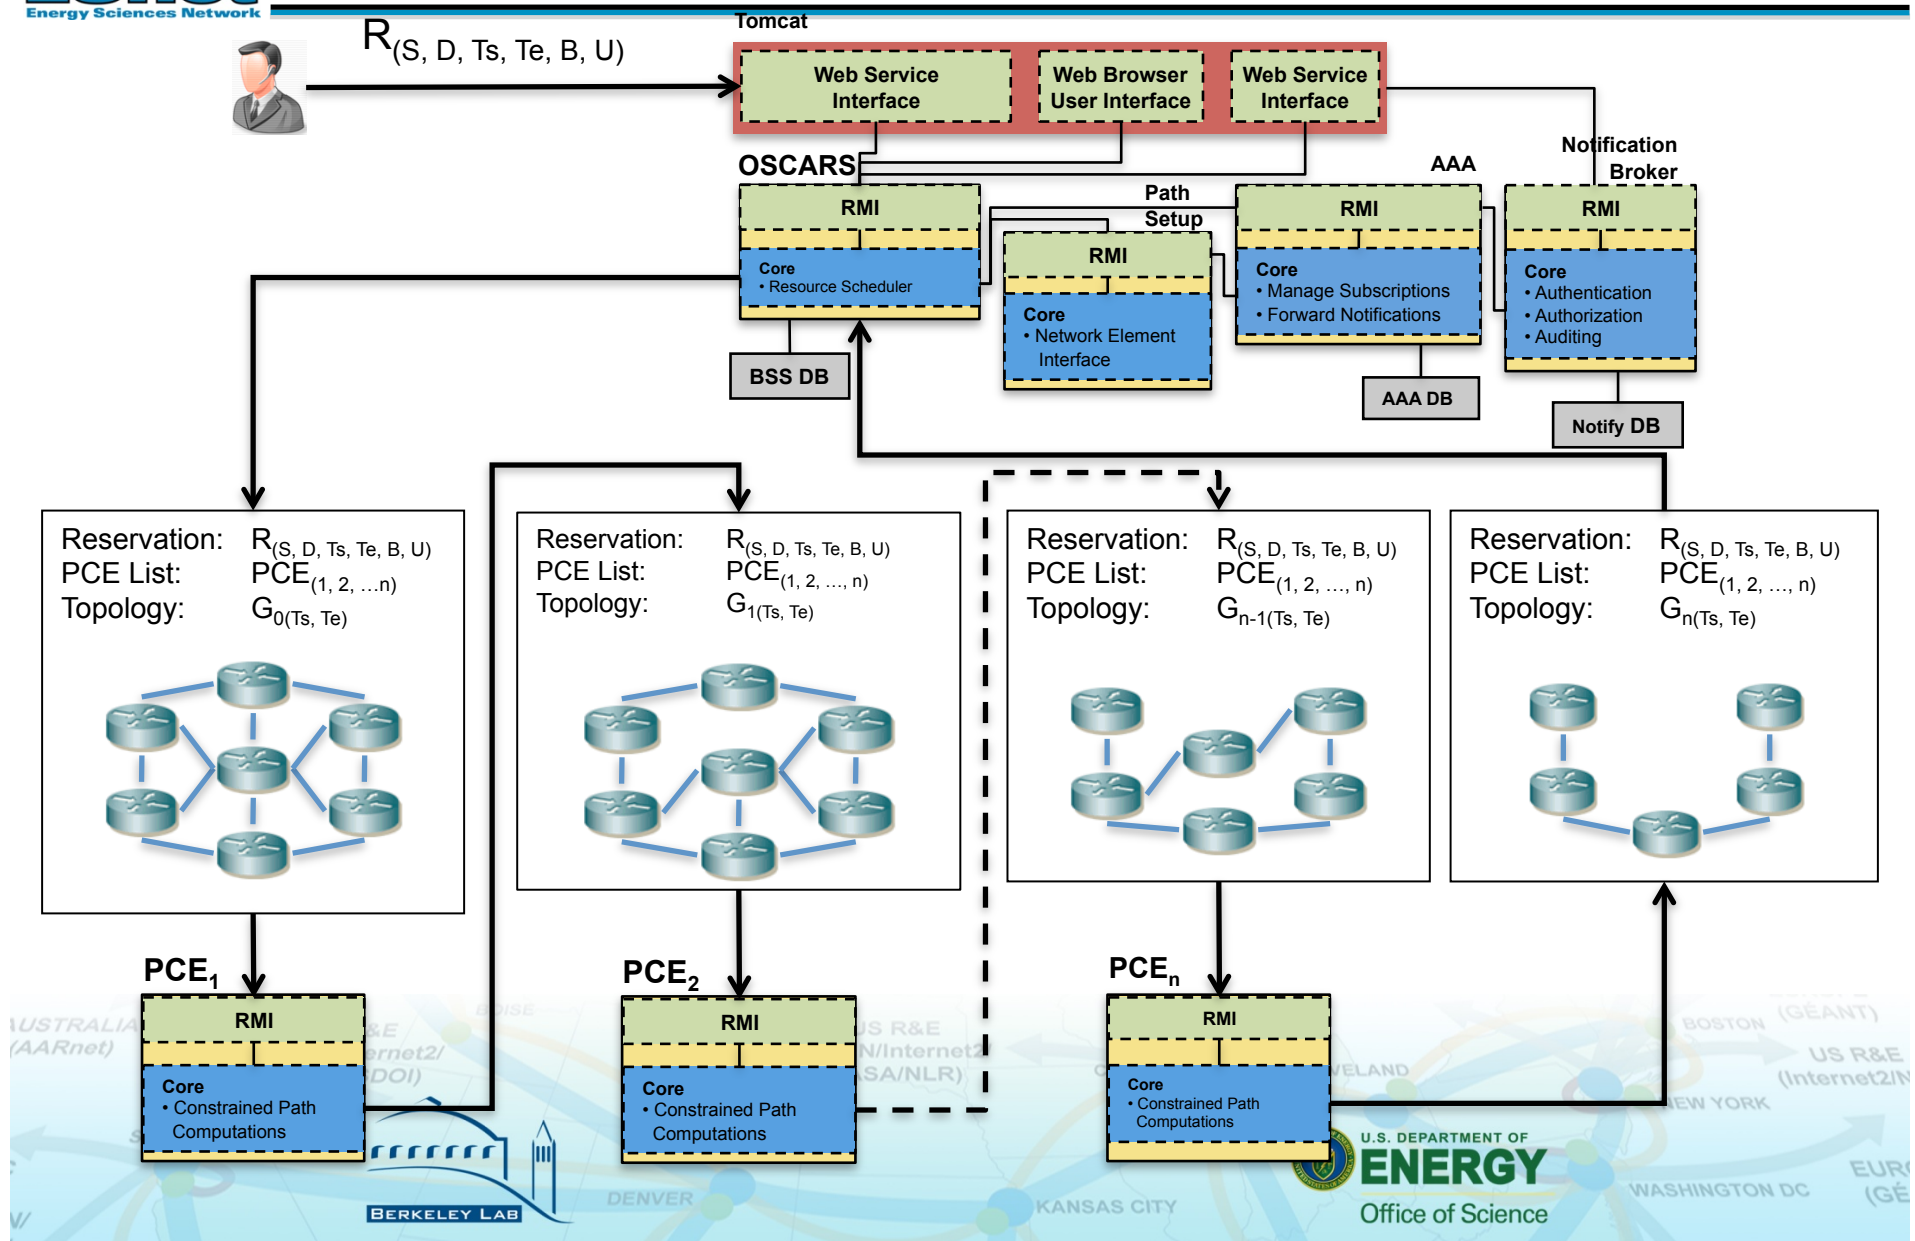
OSCARS 0.5 Architecture (1Q09)



OSCARS 0.6 Architecture (Target 3Q09)



Modular PCE Function






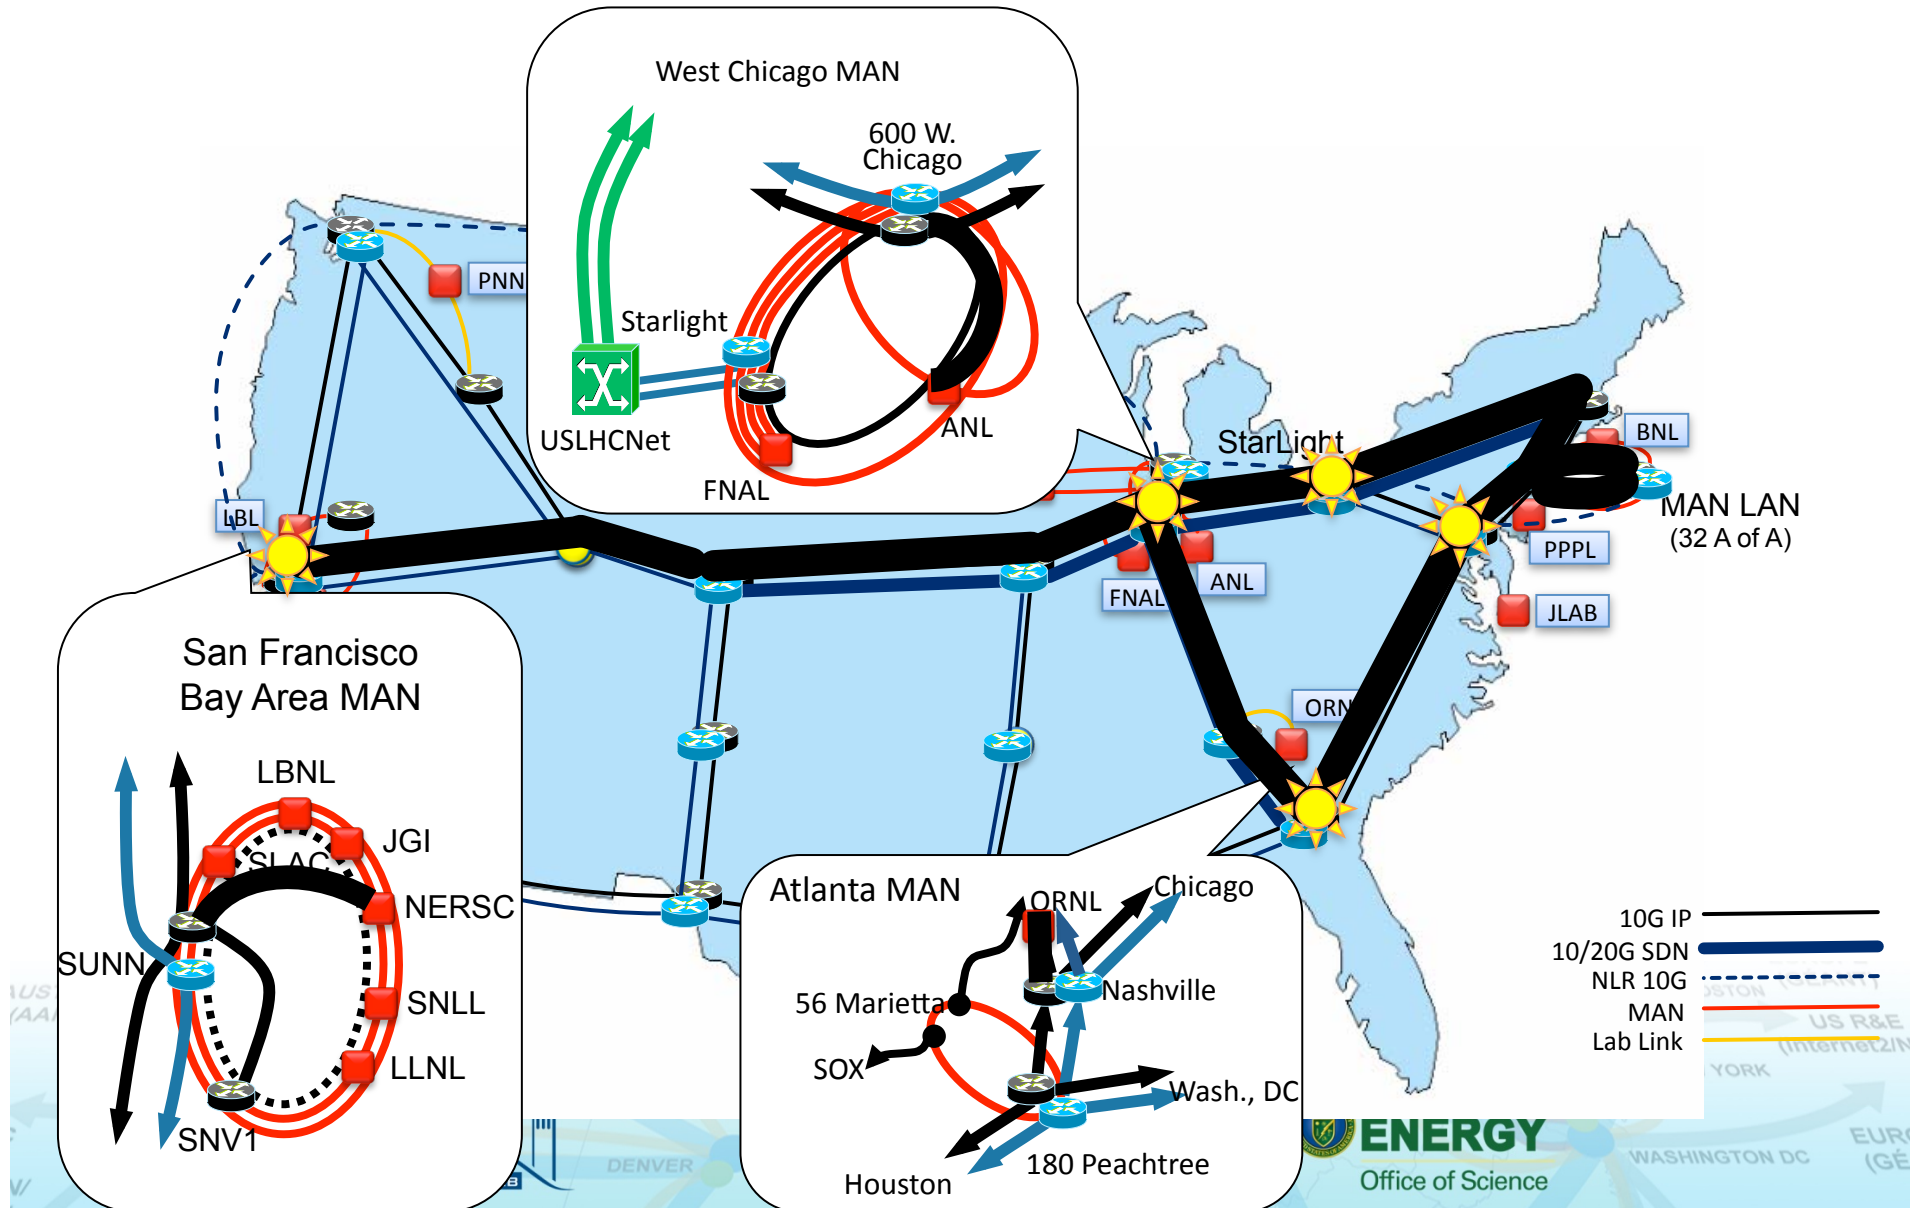
ESnet's Future Plans

ESnet recently designated to received ~\$70M in ARRA funds for an Advanced Networking Initiative

- Build a prototype wide area network to address our growing data needs while accelerating the development of 100 Gbps networking technologies
- Build a network testbed facility for researchers and industry
- Fund \$5M in network research with the goal of near term technology transfer to the production network




100 Gbps Prototype Network & Testbed





Experimental Optical Testbed

- Will consist of advanced network devices and components assembled to give network and middleware researchers the capabilities to prototype ESnet capabilities anticipated in the next decade.
- A community network R&D resource – the experimental facility will be open to researchers and industry to conduct research activities
- Multi-layer dynamic network technologies - that can support advanced services such as secure end-to-end on-demand bandwidth and circuits over Ethernet, SONET, and optical transport network technologies
- Ability to test the automatic classification of large bulk data flows and move them to a dedicated virtual circuit
- Network-aware application testing – provide opportunities for network researchers and application developers such as Grid-based middleware, cyber security services, and so on, to exploit advanced network capabilities in order to enhance end-to-end performance and security
- Technology transfer to production networks – ESnet, as host of the facility, will develop strategies to move mature technologies from testing mode to production service





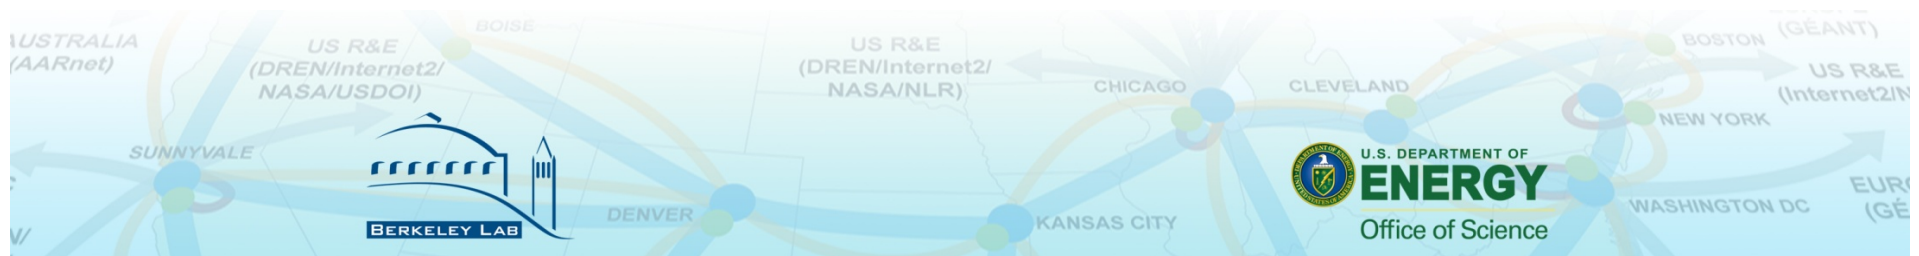
In Summary

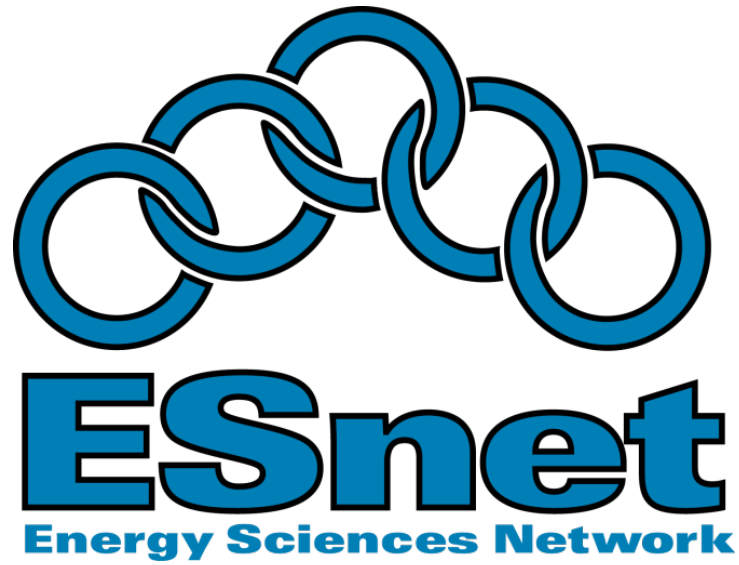
The deployment of 100 Gbps technologies is causing us to rethink our 'two network' strategy and the role for OSCARS

- Still believe that advanced reservations and end-to-end quality of service guarantees have a role

With these next generation networks, two opportunities exist:


- Ability to carry out distributed science at an unprecedented scale and with world-wide participation
- Unforeseen commercial applications that will develop using an innovative and reliable infrastructure that allows people around the world across multiple disciplines to exchange large datasets and analyses in an efficient way




A large graphic of the ESnet logo, consisting of a blue chain of four interlocking links above the text "ESnet" in a large, bold blue font, with "Energy Sciences Network" in a smaller blue font below it.

U.S. DEPARTMENT OF
ENERGY

Office of
Science






Extra Slide - OSCARS Status

Community approach to supporting end-to-end virtual circuits in the R&E environment is coordinated by the DICE (Dante, Internet2, Caltech, ESnet) working group

- Each organization potentially has their own InterDomain Controller approach (though the ESnet/Internet2 OSCARS code base is used by several organizations - flagged OSCARS/DCN)
- The DICE group has developed a standardized InterDomain Control Protocol (IDCP) for specifying the set up of segments of end-to-end VCs
 - While there are several very different InterDomain Controller implementations, they all speak IDCP and support compatible data plane connections
- The following organizations have implemented/deployed systems which are compatible with the DICE IDCP:
 - Internet2 Dynamic Circuit Network (OSCARS/DCN)
 - ESnet Science Data Network (OSCARS/SDN)
 - GÉANT2 AutoBahn System
 - Nortel (via a wrapper on top of their commercial DRAC System)
 - Surfnet (via use of above Nortel solution)
 - University of Amsterdam (OSCARS/DCN)
 - LHCNet (OSCARS/DCN)
 - LEARN (Texas RON) (OSCARS/DCN)
 - LONI (OSCARS/DCN)
 - Northrop Grumman (OSCARS/DCN)
 - Nysernet (New York RON) (OSCARS/DCN)
 - DRAGON (U. Maryland/MAX) Network
- The following "higher level service applications" have adapted their existing systems to communicate via the user request side of the IDCP:

- LambdaStation (FermiLab)
- TeraPaths (Brookhaven)
- Phoebus (UMd)





Extra Slide - Production OSCARS

Modifications required by FNAL and BNL

- Changed the reservation workflow, added a notification callback system, and added some parameters to the OSCARS API to improve interoperability with automated provisioning agents such as LambdaStation, Terapaths and Phoebus.

Operational VC support

- As of 12/2/08, there were 16 long-term production VCs instantiated, all of which support HEP
 - 4 VCs terminate at BNL
 - 2 VCs support LHC T0-T1 (primary and backup)
 - 12 VCs terminate at FNAL
 - 2 VCs support LHC T0-T1 (primary and backup)
 - For BNL and FNAL LHC T0-T1 VCs, except for the ESnet PE router at BNL (bnl-mr1.es.net) and FNAL (fnal-mr1-es.net), there are no other common nodes (router), ports (interfaces), or links between the primary and backup VC.
- Short-term dynamic VCs
 - Between 1/1/08 and 12/2/08, there were roughly 3650 successful HEP centric VCs reservations
 - 1950 reservations initiated by BNL using Terapaths
 - 1700 reservations initiated by FNAL using LambdaStation

