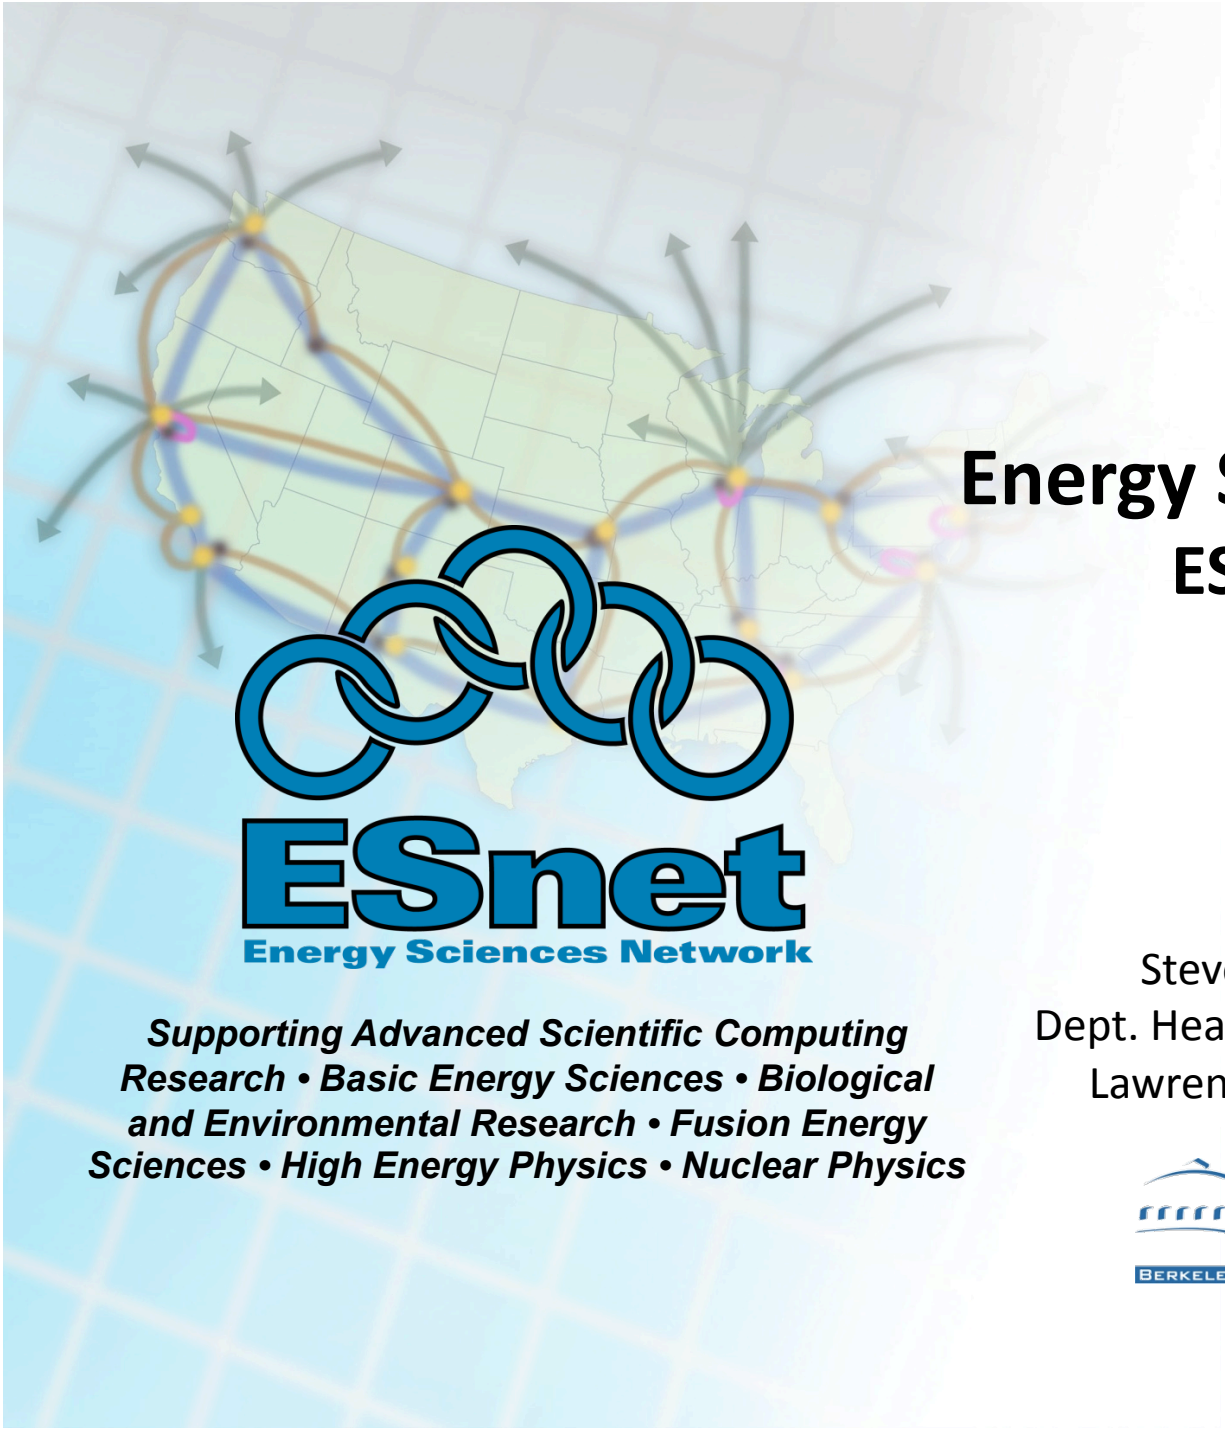

Energy Sciences Network
ESnet4 Update

July 21, 2009

ESnet
Energy Sciences Network

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

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

BERKELEY LAB

U.S. DEPARTMENT OF ENERGY
Office of Science

Agenda

- Network Update
  - Site Installs
  - Circuit Installs
  - Futures
- Services Update
- Futures

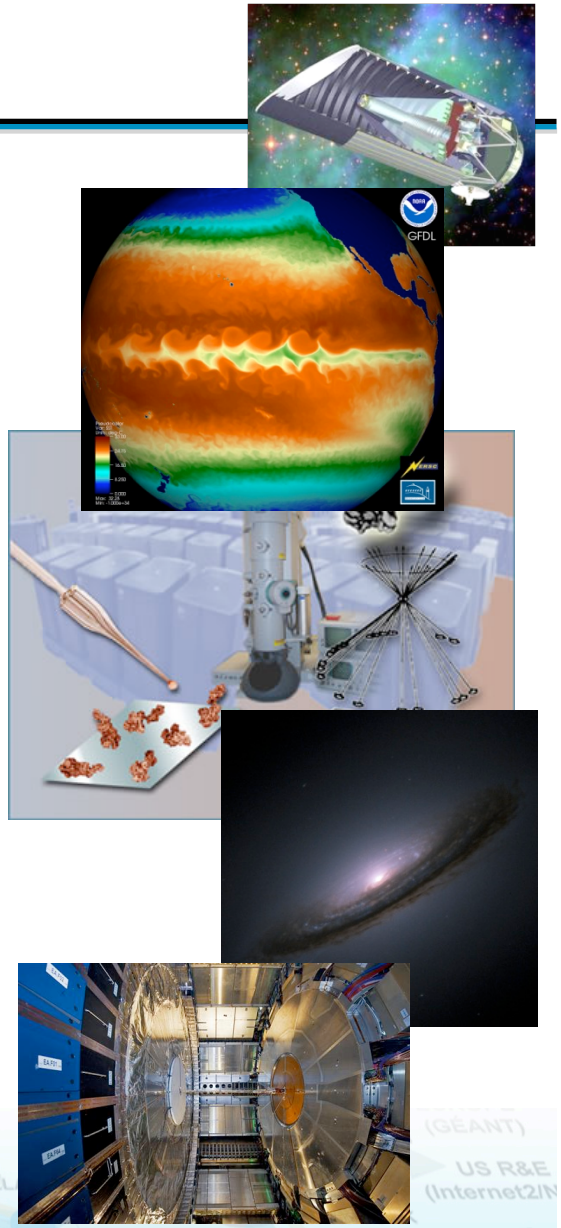

2009 Site Installs

Feb	Mar	Apr	May	June	July
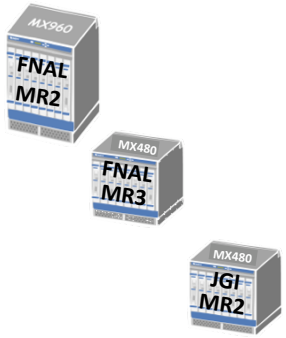					
Qwest AOA-HUB decommissioned	**Qwest DC-HUB decommissioned**				
FNAL MX40 & MX960 (replaced 6509) JGI MX480 (replaced 6509)	JGI MX480 (replaced 6509) BNL MX40 & MX960 (replaced 6509)	PAIX-PA M10i (replaced M10)	PAIX-PA M10i (replaced M10)	ANL MX960 (replaced 6509)	PAIX-PA M10i (replaced M10)
1 10GE LIMAN#3 AofA-BNL IP up Feb 2nd 1 10GE DF circuit between AofA-NEWY up on Feb 2nd 1 1GE PPPL-HEP-PU Feb 4th 1 10GE LIMAN#4 NEWY-BNL up on Feb 9th 1 10GE STAR-HUB USLHCnet E300 Feb 10th	1 DS3 OSTI-WASH-HUB (backup) up on Mar 3rd 1 10GE BOST-MIT IP peering enabled on Mar 19th		1 10GE peering BOST-HUB – NoX for BNL on May 27th	1 1GE SDN between PPPL (GFDL) – WASH on June 1st 1 10GE SDN PNWG-HUB – PNNL on June 6th 1 10GE NASH-ORNL-#2 SDN Wave on June 14th 1 1GE Peering PNWG-HUB – CSTNET (China) on June 17h 1 10M PNNL Sire office at OSTI June 19th *1 10GE KANS-Great Plains Net (GPN) for the DUSAL & OneNet Peering June 27th*	1 1GE Level3 private peering @ PAIX-PA up on July 2nd 1 OC12 GA-LASV-HUB up on July 9th 1 10GE SDN peering with ANL Site July 13th

Backbone Hub & Wave Count

Current Hub Count: 21

- 32 AofA, NEWY, WASH, ATLA, NASH, CLEV, BOST, CHIC, STAR, KANS, HOUS, ELPA, DENV, ALBU, BOIS, PNWG, SUNN, SNV(Qwest), LOSA, SDSC, LASV (SwitchNap)

Current Backbone Wave Count:

- Internet2/Level3 10G waves:
  - IP: 17 new/split for a total of 25
  - SDN: 25 new/split for a total of 30
- NLR 10G waves:
  - 5 10G waves

Circuit & Site Installs

Replace site 6509s (FNAL, ANL & BNL) with MX's

- FNAL: installed MX960 on Feb 9th and MX480 on Feb 20th
- BNL: installed MX480 on Mar 16th and MX960 on Mar 31st
- ANL: installed MX960 on June 30th

All BAMAN sites now have MX's

- JGI-MR2: installed on Feb 26th
- NERSC-MR2: installed on Mar 5th

Site and Hub router upgrades

- PAIX-PA M10: replaced by M10i on May 7th
- SNLA M10: will be replaced by M10i on July 30th

Circuit installs

- 3 10GE Lightower circuits LIMAN#3, LIMAN#4 & AofA-NEWY dark fiber (Feb 2009)
- 1 1GE PPPL to HEP (Princeton Univ.) Feb 4th
- 1 10GE STAR-HUB to USLHCnet E300 peering Feb 10th
- 1 DS3 back-up for ORAU to WASH-HUB Mar 3rd

Circuit & Site Installs (cont.)

Circuit installs (cont.)

- 1 10GE in BOST to MIT enabled on Mar 19th
- 1 10GE peering BOST-HUB to NoX for BNL May 27th
- 1 1GE SDN between PPPL (GFDL) – WASH June 1st
- 1 10GE SDN PNWG-HUB – PNNL June 6th
- 1 10GE NASH-ORNL-#2 SDN Wave June 14th
- 1 1GE Peering PNWG-HUB – CSTnet (China) June 17th
- 1 10M PNNL Site Office @ OSTI June 19th
- 1 10GE KANS HUB-GPN for the 1GE DUSAL & 1GE OneNet Peering June 27th
- 1 OC12 GA-LASV-HUB July 9th
- 1 10GE ANL SDN site peering July 13th

Future Circuit Installs

Future Circuit Installs:

- OC12 between DENV-HUB and Pantex (TBD)
- 1GE wave in BOIS to INL via IRON (TBD)
- 1GE links in D.C. Area for Germantown, IN to WASH-HUB (TBD)
- OC3 for NSTech-NV to LASV-HUB (on order)
- 10G peering at PNWG-HUB with Korea (KSTAR & KISTI) (TBD)

Agenda

- Network Update
- Services Update
  - IP & SDN Network
  - OSCARS
- Futures

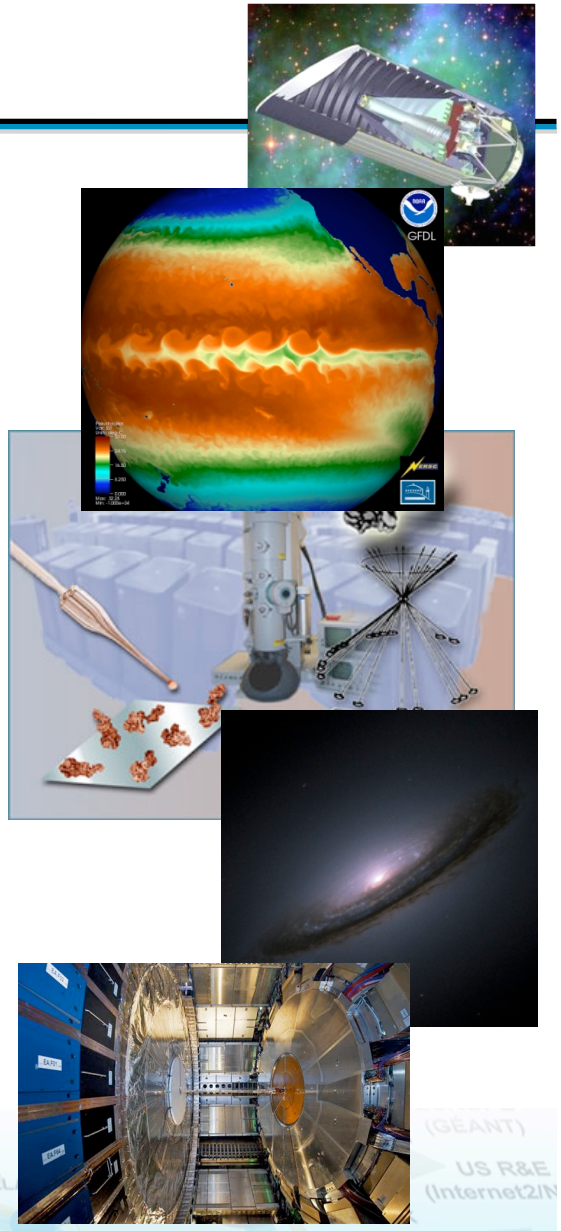

ESnet4 – July 09

Network Weather Map

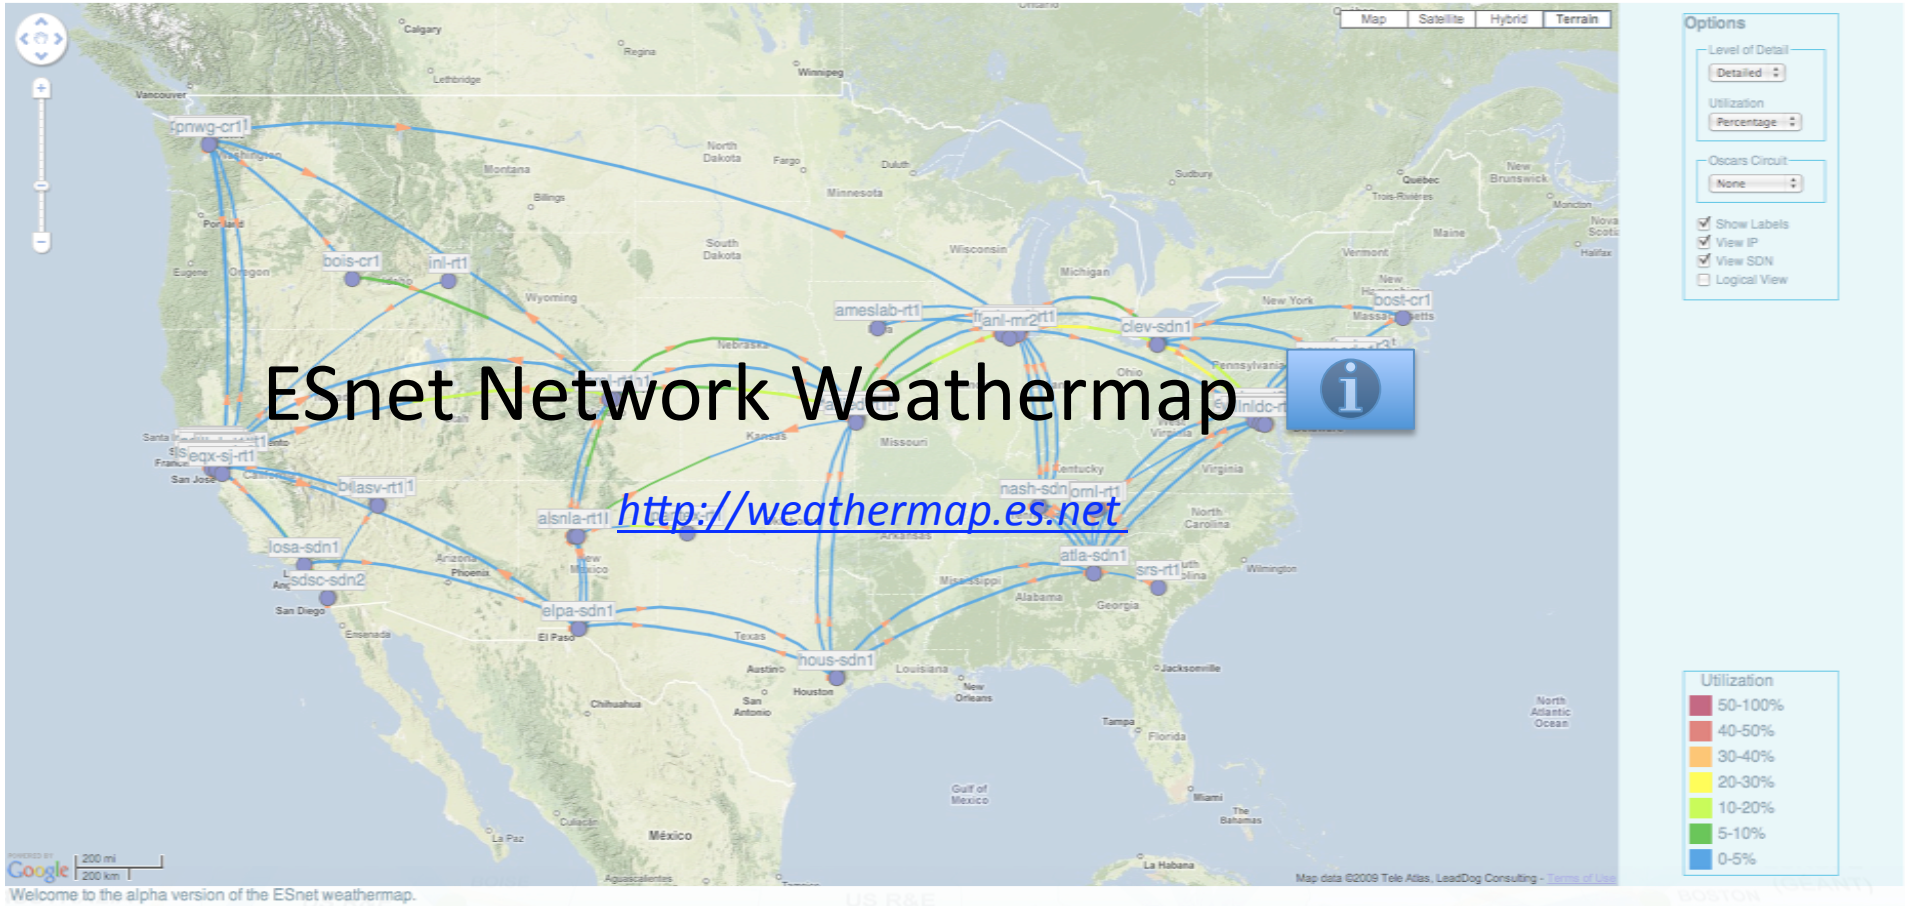

Exponential Traffic Growth

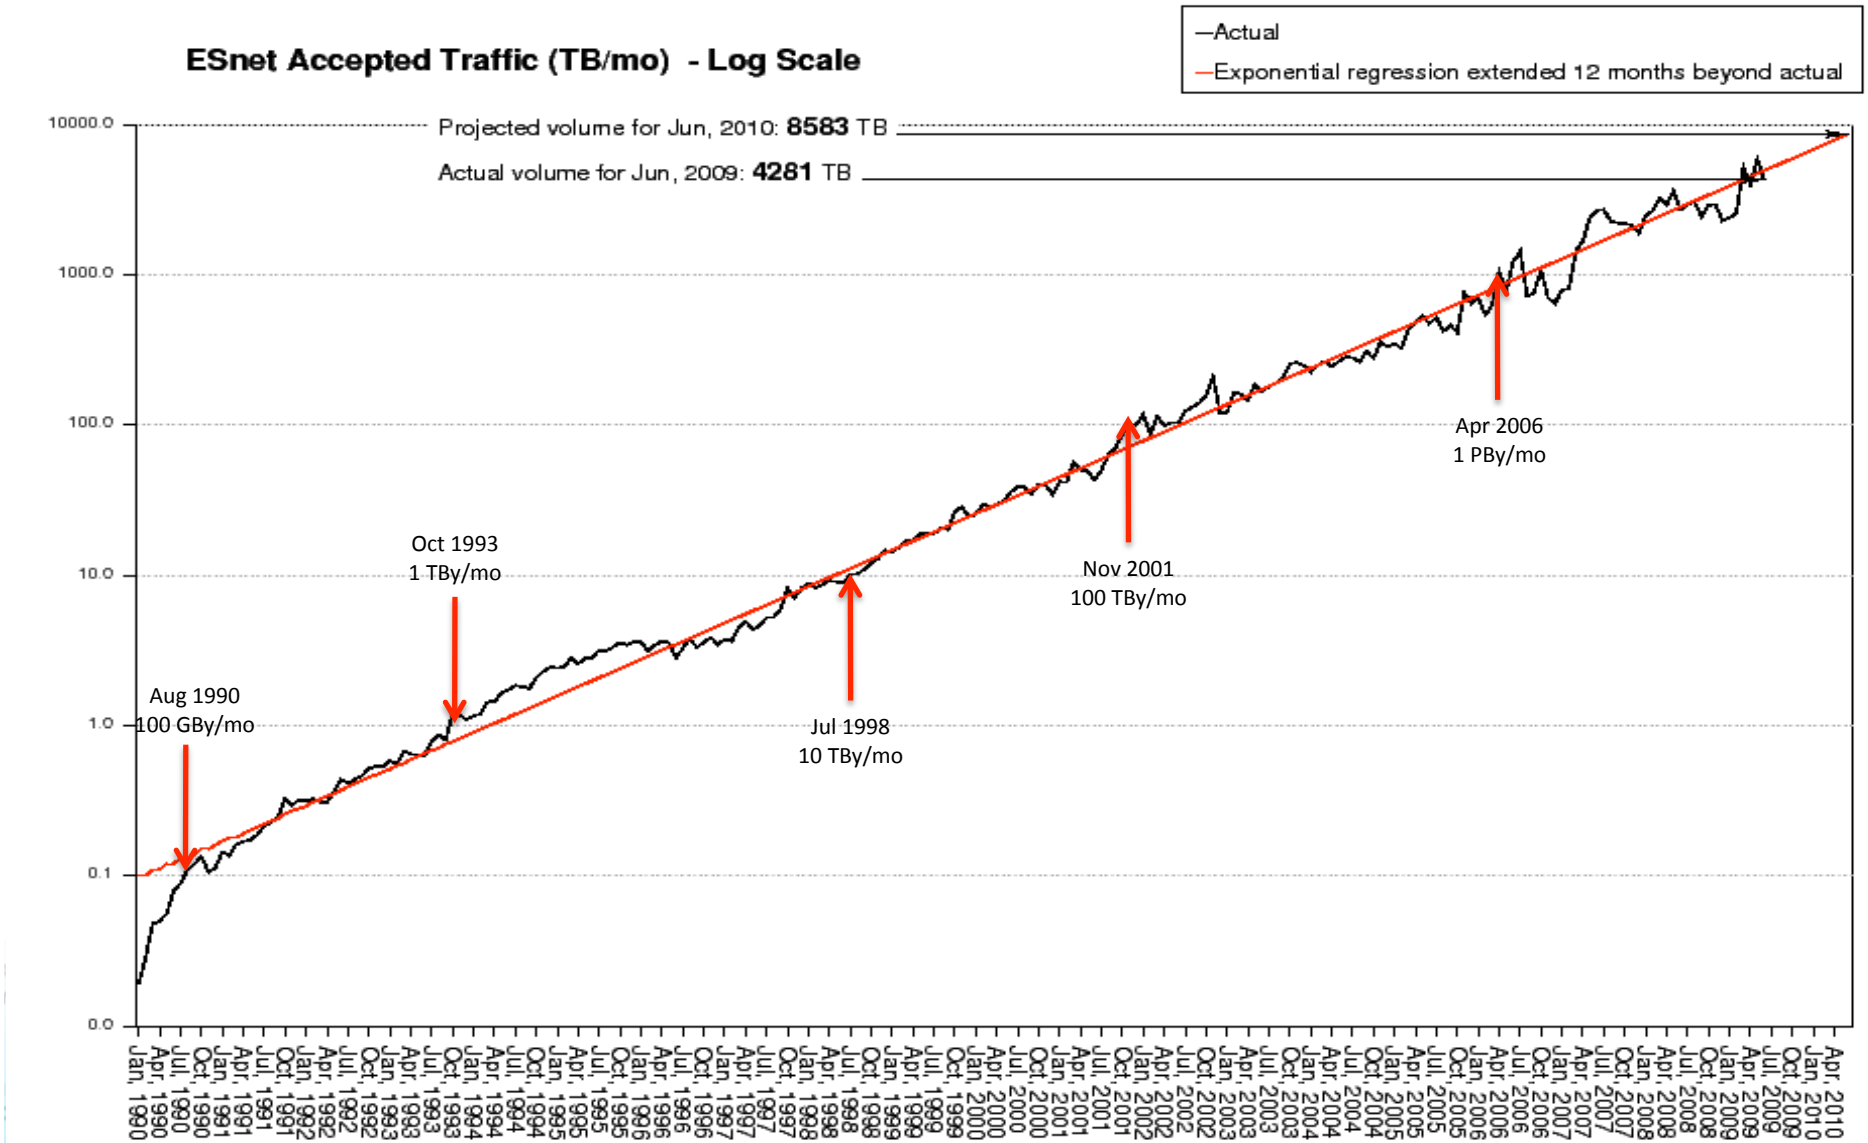

OSCARS: Multi-Domain VC Service

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)

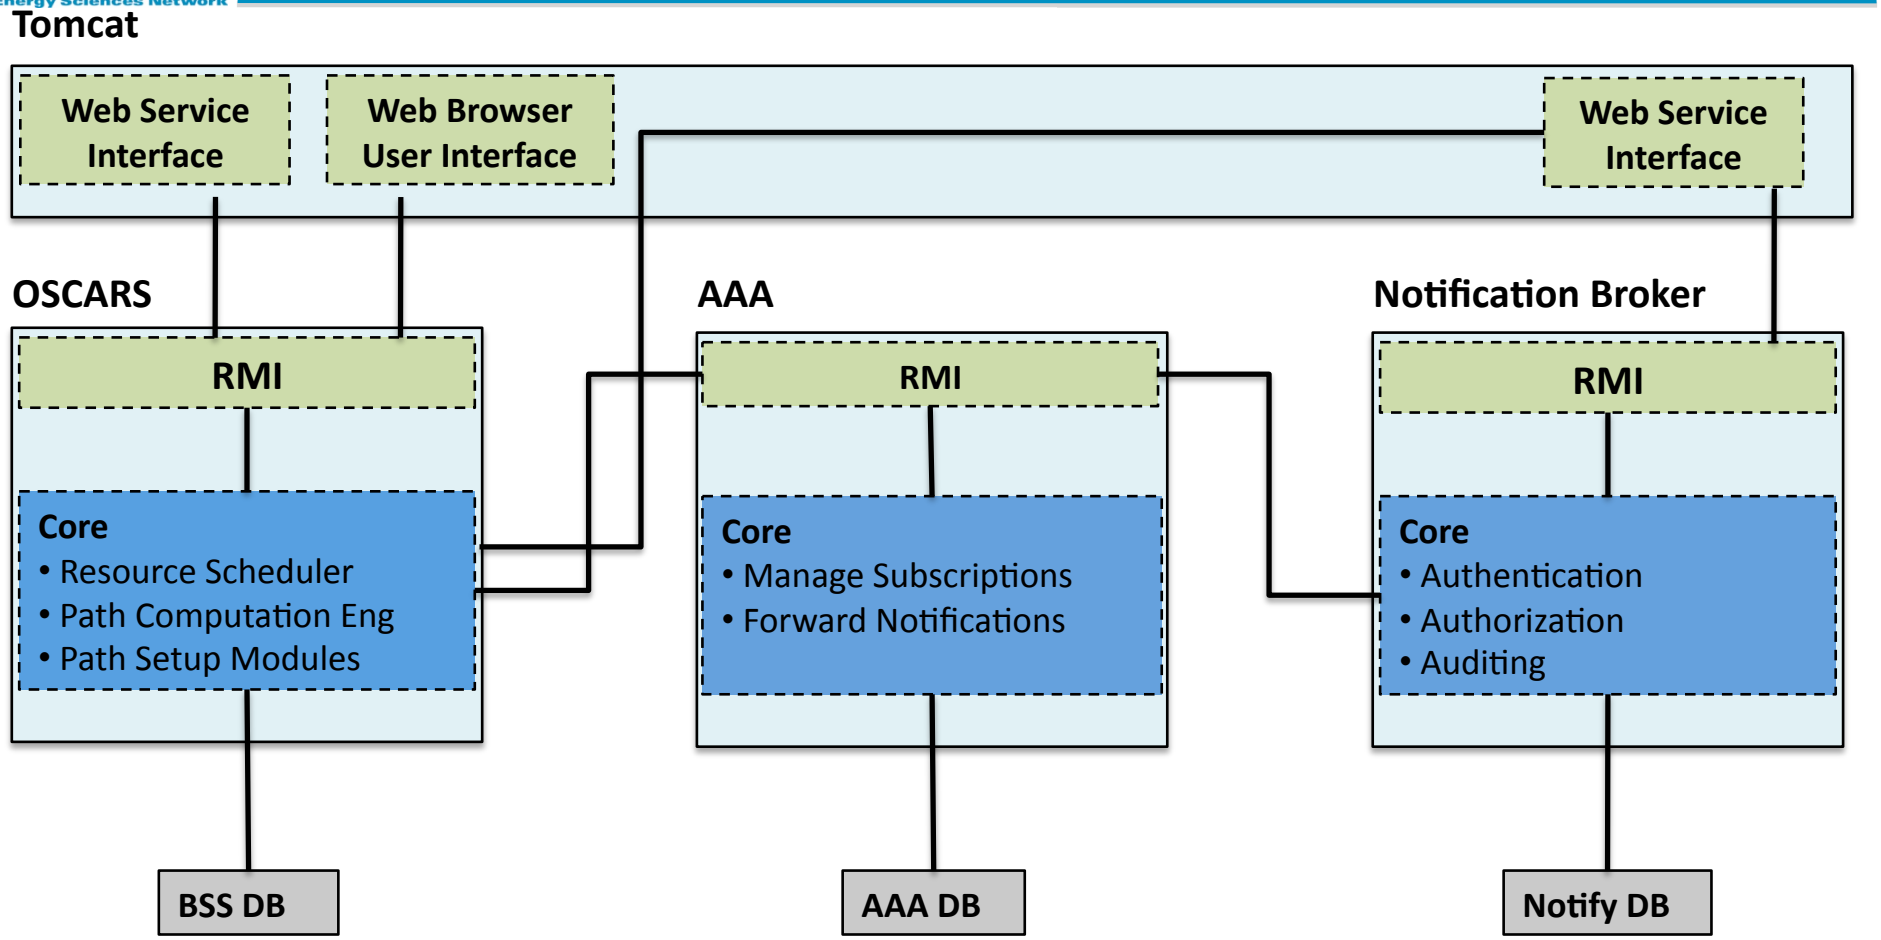

OSCARS 0.6 Design / Implementation Goals

- Support production deployment of service and facilitate research collaborations
  - Distinct functions in stand-alone modules
    - Supports distributed model
    - Facilitates module redundancy
  - Formalize (internal) interface between modules
    - Facilitates module plug-ins from collaborative work (e.g. PCE)
    - Customization of modules based on deployment needs (e.g. AuthN, AuthZ, PSS)
  - Standardize external API messages and control access
    - Facilitates inter-operability with other dynamic VC services (e.g. Nortel DRAC, GÉANT AuthBAHN)
    - Supports backward compatibility of IDC protocol

OSCARS 0.6 Architecture (Target 12/09)

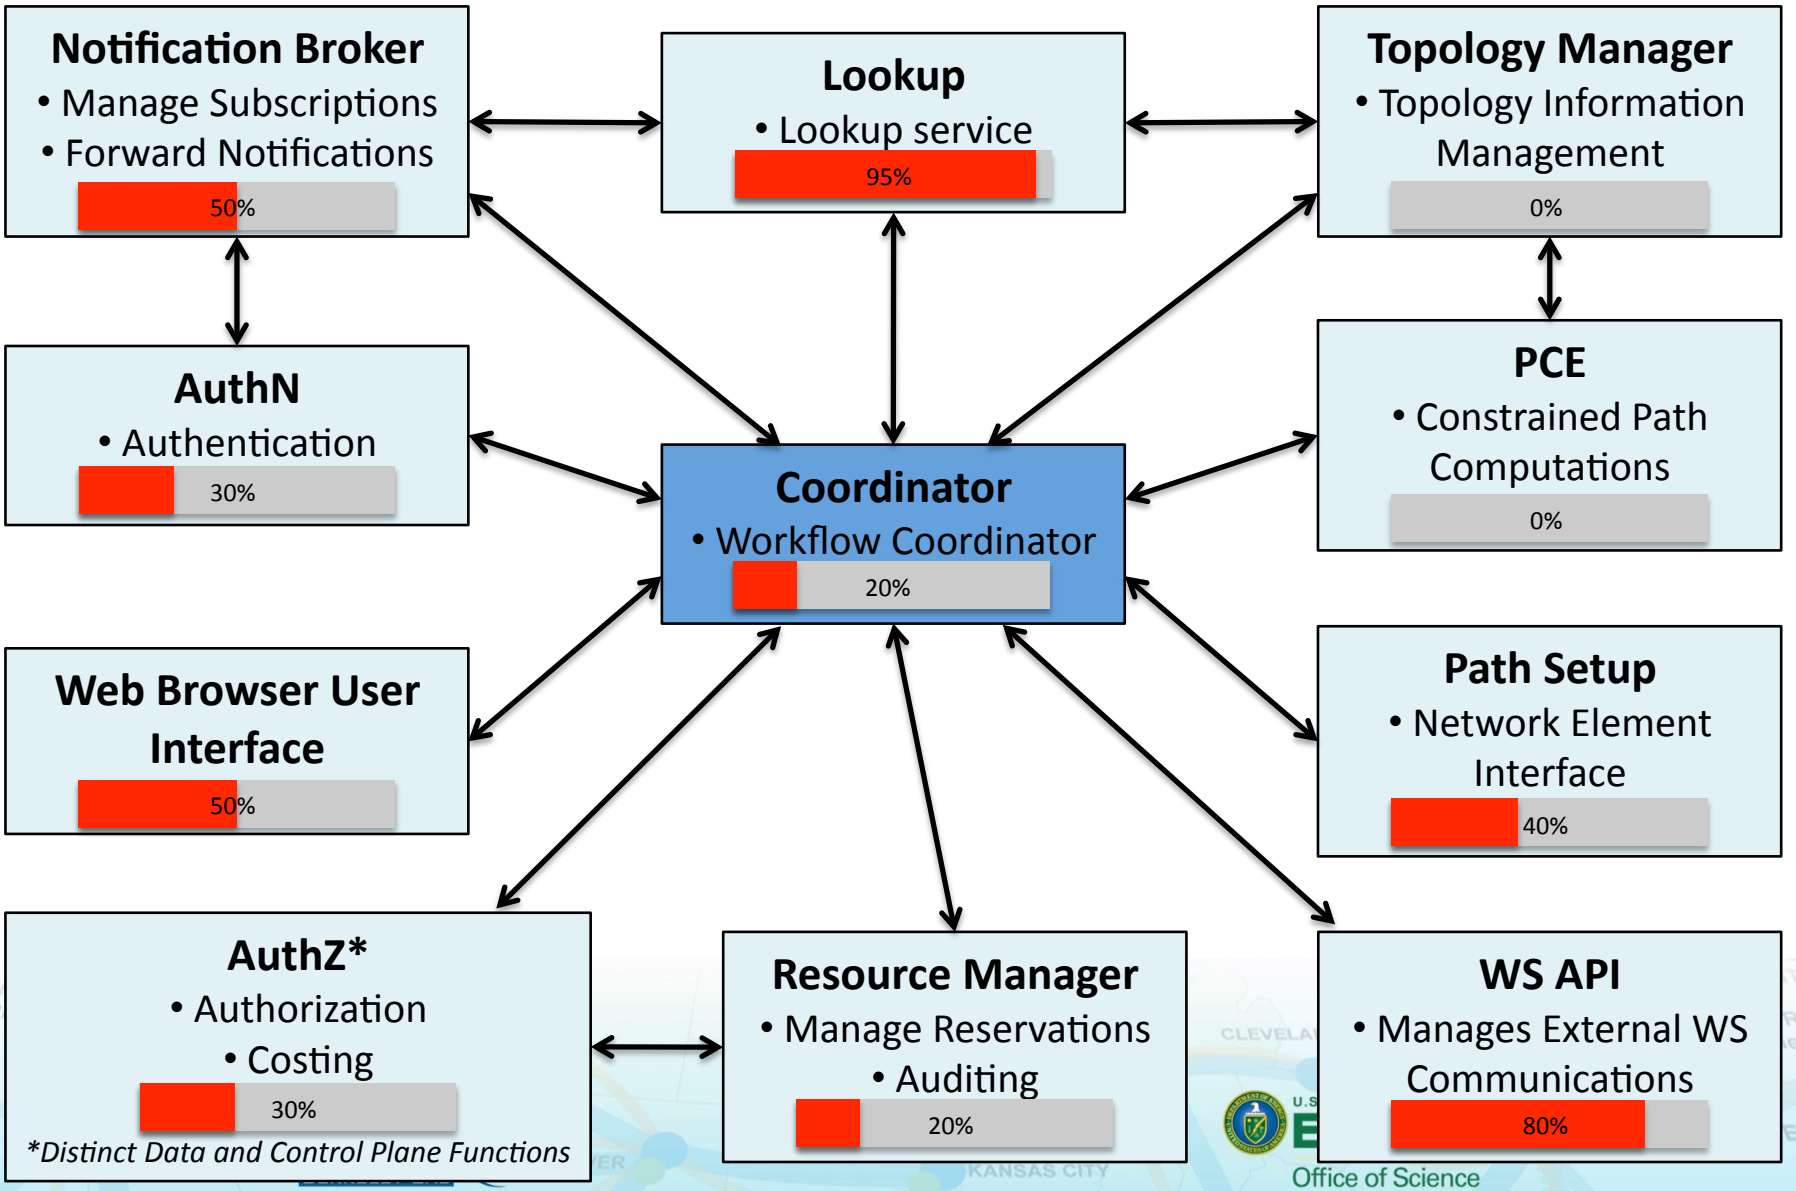

Office of Science

SDN Statistics https://oscars.es.net/stats/

Settings

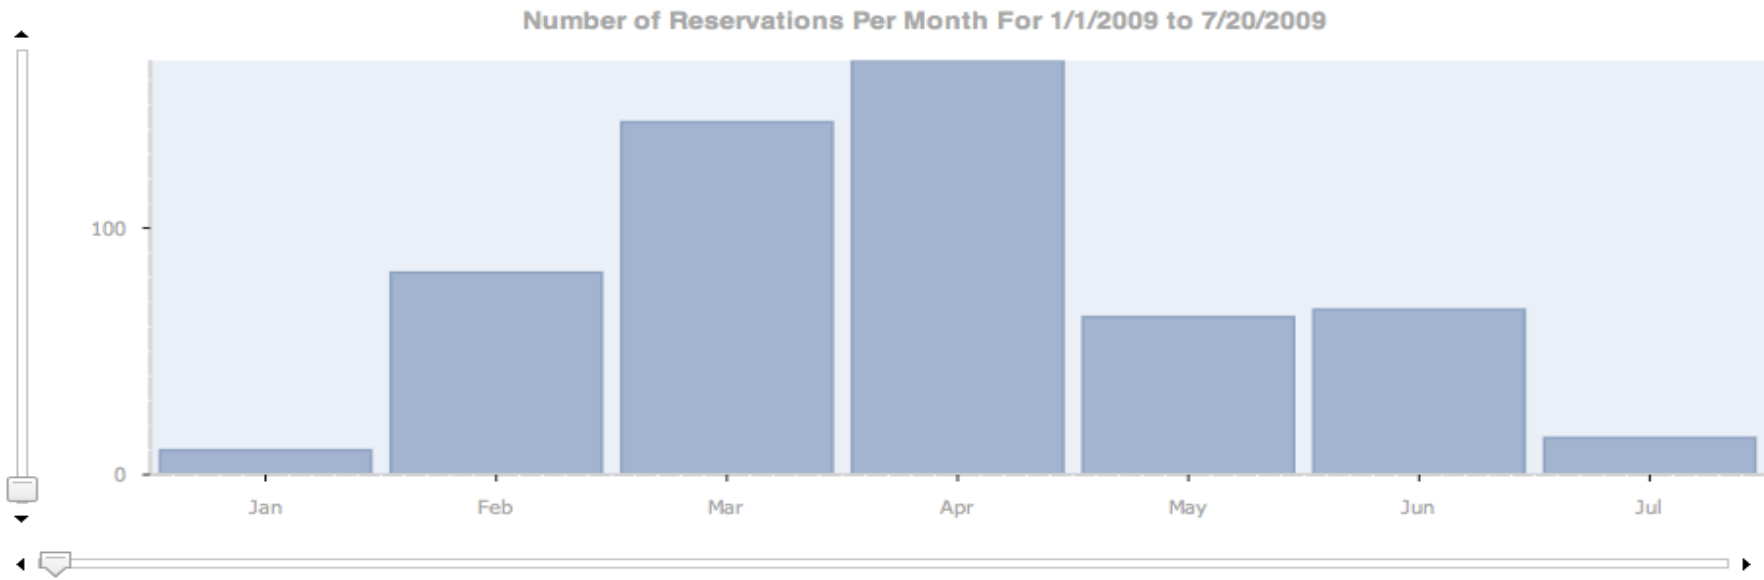

BW=Bandwidth (Mbps), Dur=Duration (Minutes)

X	# of Resv	BW Sum	BW Avg.	BW Max	BW Min	BW Std Dev	Dur. Sum	Dur. Avg.	Dur. Max	Dur. Min	Dur. Std Dev	Avg Hops	Max Hops	Min Hops	Hop Std Dev
Jan	10	10600	1060	5000	100	1575.56	2880029	288002.9	2619364	4	780993.92	4.8	9	2	2.48
Feb	82	137850	1681.1	10000	50	2272.97	142018292.1	1731930.39	5258884	3	2471154.49	3.62	8	2	1.8
Mar	143	193400	1352.45	7000	50	1337.74	9462905.87	66174.17	1051204	3.6	255280.85	3.64	7	2	1.23
Apr	168	160050	952.68	5000	100	763.18	31475933.75	187356.75	3146404	4	724499.84	5.29	9	2	1.43
May	64	64050	1000.78	5000	50	932.59	29107163.33	454799.43	3182404	3.6	1077214.03	5.27	9	2	1.42
Jun	67	53250	794.78	9000	50	1467.28	6333843.75	94534.98	3166564	4	538869.54	4.93	6	2	0.94
Jul	15	14864	990.93	3000	1	1128.22	1058494	70566.27	525604	4	178481.67	4.53	9	2	1.78

*Developed by Andy Lake (Internet2)

SDN - IP Traffic Breakdown

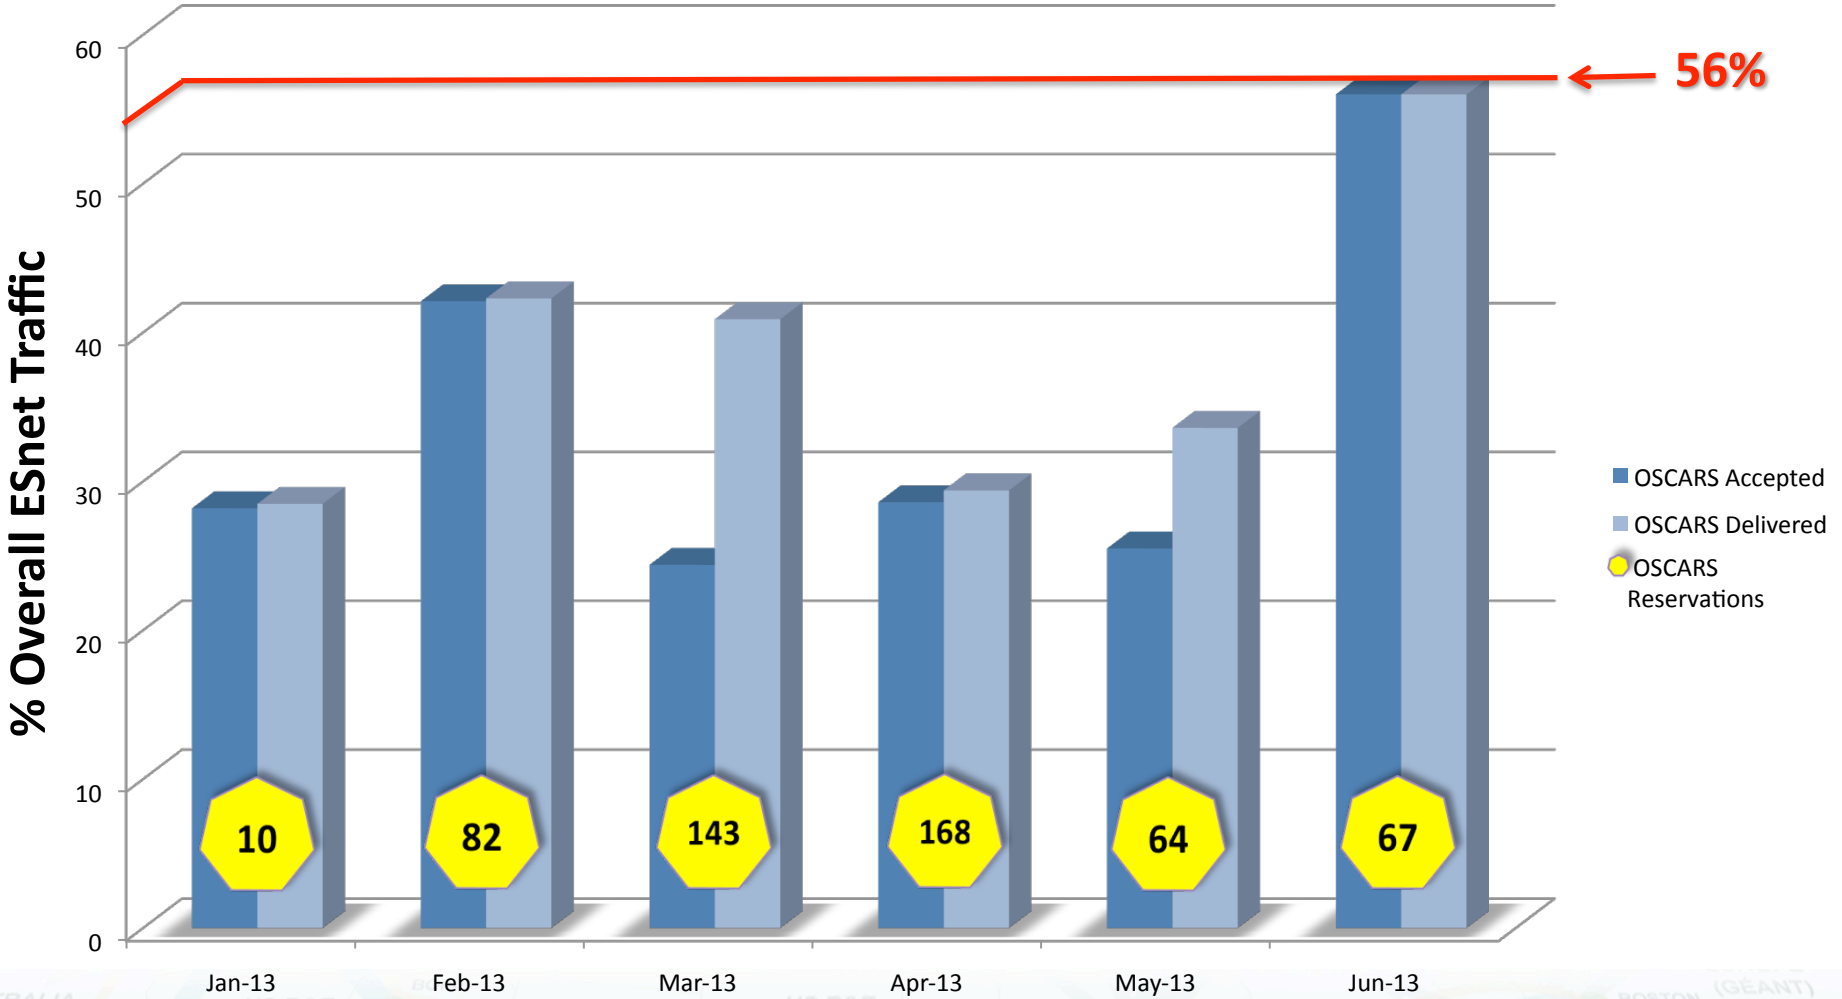

OSCARS Collaborative Research Efforts

- LBNL LDRD "On-demand overlays for scientific applications"
  - To create proof-of-concept on-demand overlays for scientific applications that make efficient and effective use of the available network resources
- GLIF GNI-API "Fenius" to translate between the GLIF common API to:
  - DICE IDCP: OSCARS IDC (ESnet, I2)
  - GNS-WSI3: G-lambda (KDDI, AIST, NICT, NTT)
  - Phosphorus: Harmony (PSNC, ADVA, CESNET, NXW, FHG, I2CAT, FZJ, HEL IBBT, CTI, AIT, SARA, SURFnet, UNIBONN, UVA, UESSEX, ULEEDS, Nortel, MCNC, CRC)
- DOE Projects:
  - "Virtualized Network Control" to develop multi-dimensional PCE (multi-layer, multi-level, multi-technology, multi-layer, multi-domain, multi-provider, multi-vendor, multi-policy)
  - "Integrating Storage Management with Dynamic Network Provisioning for Automated Data Transfers" to develop algorithms for co-scheduling compute and network resources
  - "Hybrid Multi-Layer Network Control" to develop end-to-end provisioning architectures and solutions for multi-layer networks

Agenda

- Network Update
- Services Update
- ARRA
  - Advanced Networking Initiative
  - Experimental Optical Testbed

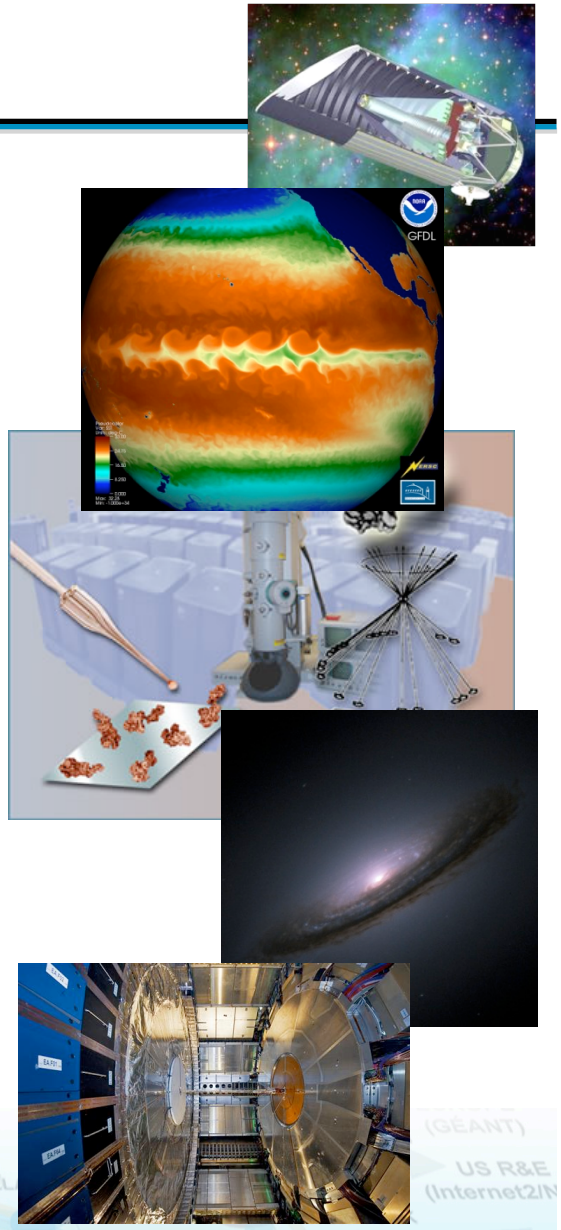

ESnet's ARRA Plans

ESnet designated to received ~$67M 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

Advanced Networking Initiative

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

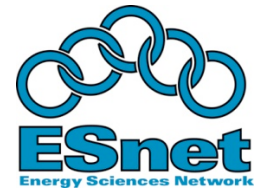
ESnet
Energy Sciences Network

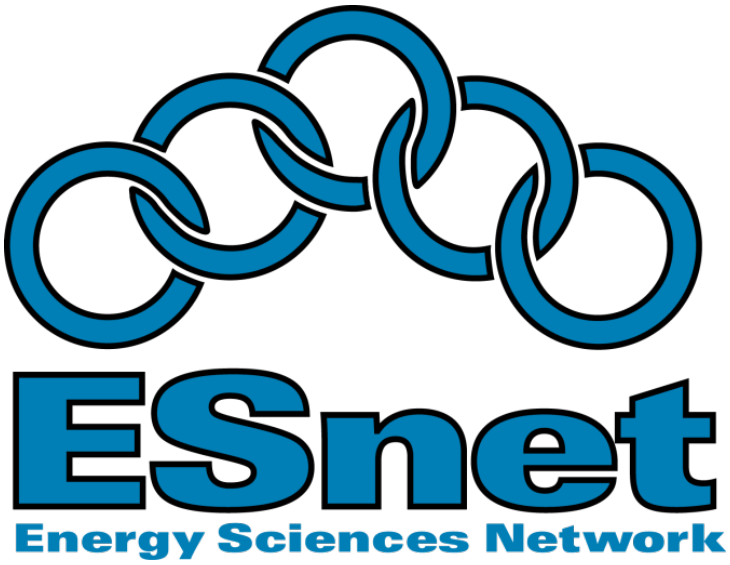
ESnet
Energy Sciences Network


U.S. DEPARTMENT OF
ENERGY
Office of
Science

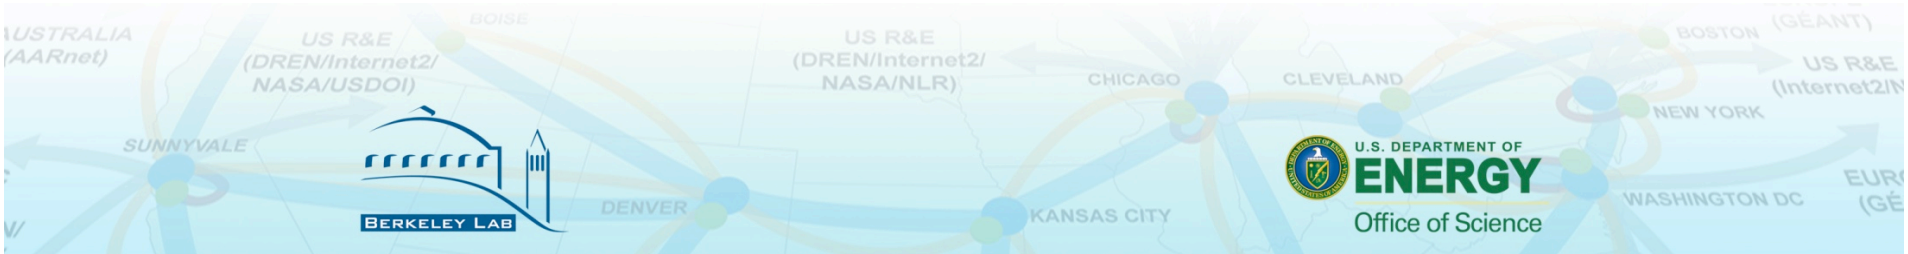
USTRALIA
(AARnet)
US R&E
(DREN/Internet2/
NASA/USDOI)
BOISE
SUNNYVALE
DENVER
US R&E
(DREN/Internet2/
NASA/NLR)
CHICAGO
CLEVELAND
KANSAS CITY
BOSTON
(GEANT)
US R&E
(Internet2/N
NEW YORK
WASHINGTON DC
BERKELEY LAB
U.S. DEPARTMENT OF
ENERGY
Office of Science