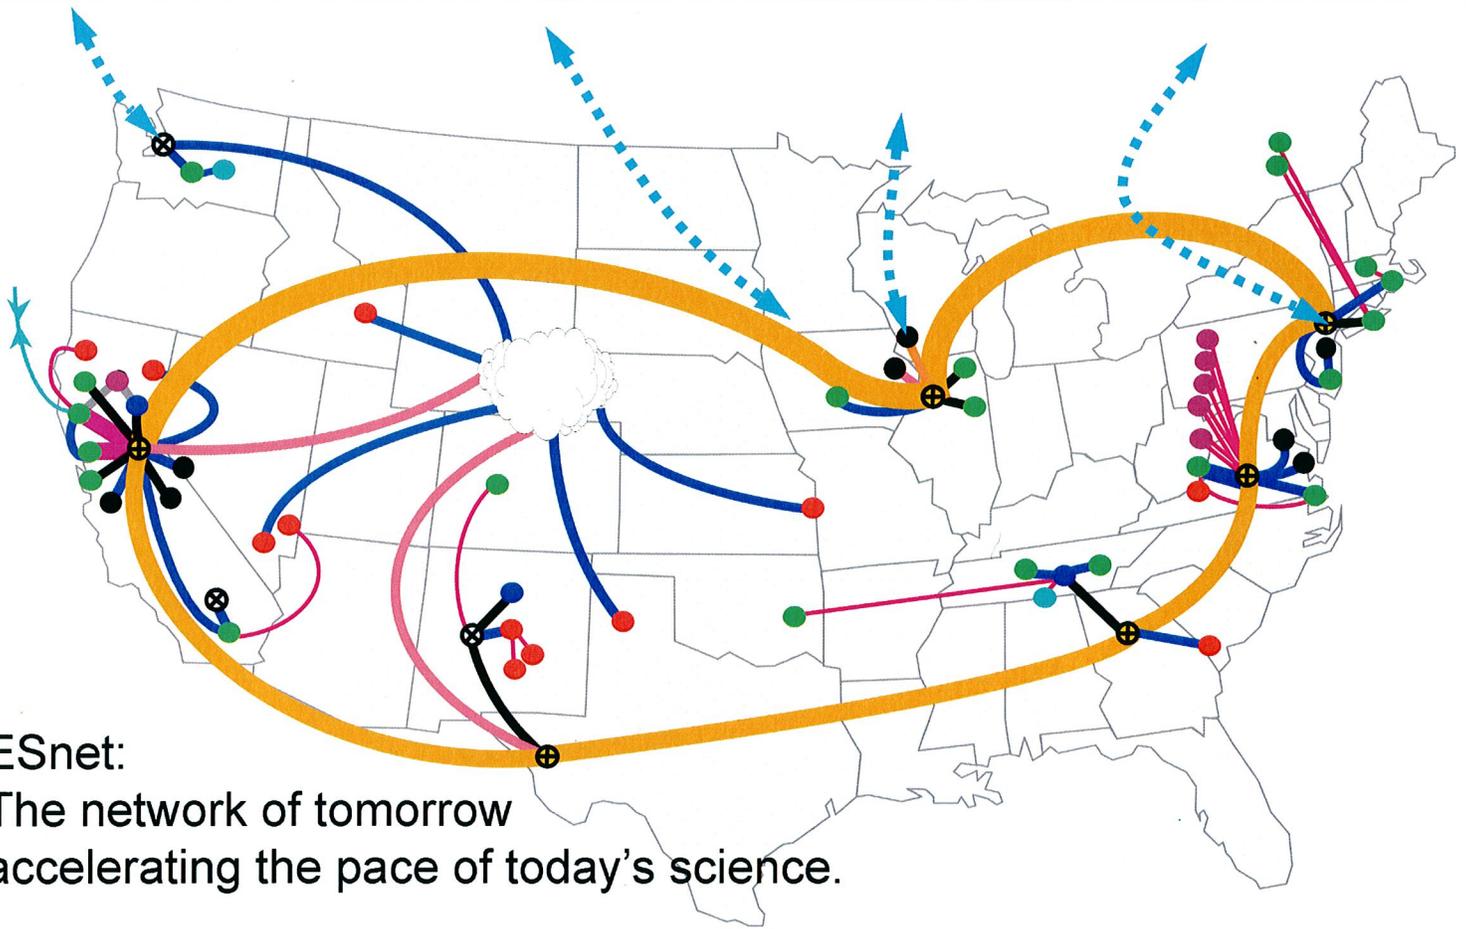


# Annual Program Review of ESnet

August 3, 2004



ESnet:  
The network of tomorrow  
accelerating the pace of today's science.



# Annual ESnet Financial and Program Review

August 3, 2004

## AGENDA

### Tuesday, August 3

8:30a – 9:00a      Executive Session      Abbott/Banda/Hitchcock/Johnston/  
Merola/Scott      50B-4231

### REVIEW

50B-4205

9:00a – 9:15a      Welcome – LBNL (Merola)

9:15a – 9:30a      Introduction – DOE HQ (Scott)

9:30a – 11:00a      ESnet – Part I (Johnston)

11:00a – 11:15a      BREAK

11:15a – 11:45a      ESnet – Finances (Fortney)

11:45a – 1:00p      Lunch – Executive Session      50B-4231

1:00p – 1:20p      ESnet – Part II (Johnston)

1:20p – 1:45p      ESnet – Network Infrastructure Engineering  
• MAN Milestones (Johnston)  
• Backbone Upgrades (Gagliardi)

1:45p – 2:45p      ESnet – Routing and Network Applications  
• Routing and ISP Services (Collins)  
• Usage Monitoring Systems (Chin)  
• NETINFO (Metzger)  
• Monitoring DOE Lab – University Connectivity (Metzger)  
• Performance Centers (Metzger)

2:45p – 3:00p      BREAK - refreshments

3:00p – 3:40p      ESnet – Routing and Network Applications (cont.)  
• Availability Measurement (O'Connor)  
• On-Demand Secure Circuits and Advance Reservation System - OSCARS (Chin)  
• Increasing the Robustness of ESnet to External Threats (Collins)

3:40p – 4:00p      ESnet – Operational Infrastructure  
• Trouble Ticket Process (Collins)  
• Disaster Recovery (Collins)  
• Infrastructure Support Services (Johnston)

4:00p – 4:30p      ESnet – Science Services (Johnston)

4:30p – 5:00p      Funding Profiles: FY05 and Beyond (Scott)

5:00p      Review Adjourn

5:00p – 5:15p      Closed - DOE Discussion      Abbott/Hitchcock/Scott/Sibal

# Annual ESnet Financial and Program Review

August 3, 2004

## AGENDA — Attendees List

### Attendees:

Ed Oliver	ASCR (via H.323)
Mike Strayer	MICS (via H.323)
Dan Hitchcock	MICS ✕
Mary Anne Scott	MICS ✕
Michael Banda	LBNL ✕
Joe Burescia	LBNL
Michael Collins	LBNL
Bill Fortney	LBNL
Jim Gagliardi	LBNL
Bill Johnston	LBNL ✕
Gizella Kapus	LBNL
Stan Kluz	LBNL
Sandy Merola	LBNL ✕
Lynne Rippe	LBNL
Horst Simon	LBNL ✕
Rajinder Singh	LBNL
Kim Abbott	BSO ✕

### PM Session to include:

Chin Guok	LBNL
Joe Metzger	LBNL (via H.323)
Mike O'Connor	LBNL (via H.323)

● ✕ - Lunch participant



INFORMATION TECHNOLOGIES  
& SERVICES DIVISION

**Sandy Merola**  
Division Director, ITSD

## Annual Program Review of ESnet

LBLN, Berkeley, CA – August 3, 2004



## Welcome and Purpose



### 1. Provide Program Manager

- o Status
- o Accomplishments
- o Issues
- o Plans
- o Finances
- o Business Practices

### 2. Reinvented and Re-engineered

- o Approach
- o Architecture
- o Technology
- o Business Practices

### 3. Incorporates the Request from the Lehman Review

INFORMATION TECHNOLOGIES AND SERVICES DIVISION



## Annual Program Review of ESnet August 2004

William E. Johnston, ESnet Dept. Head and Senior Scientist

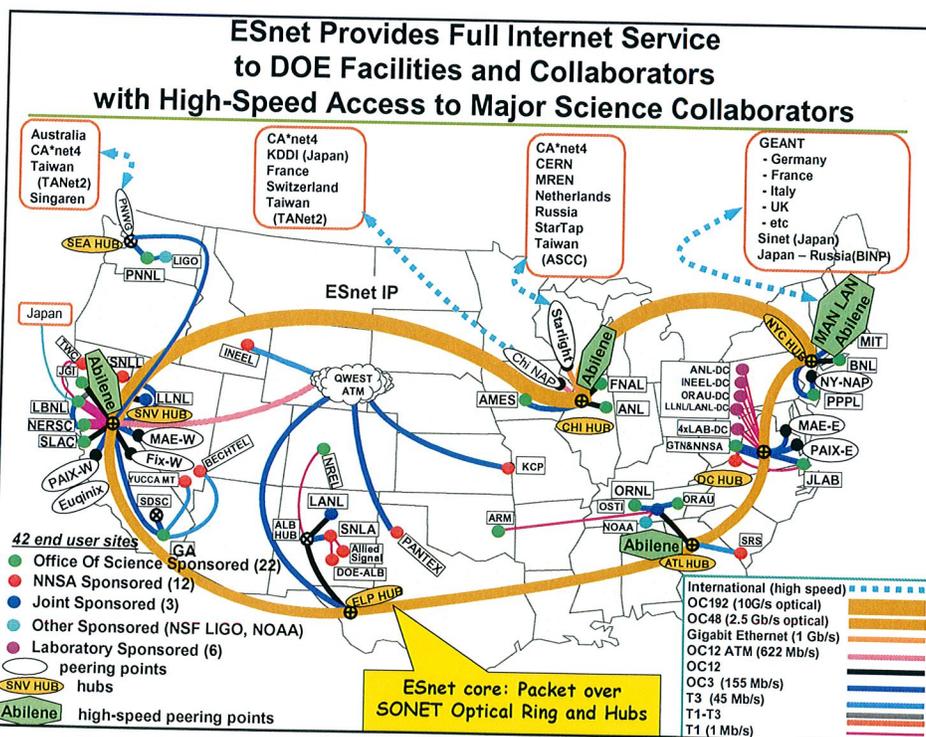
R. P. Singh, Federal Project Manager

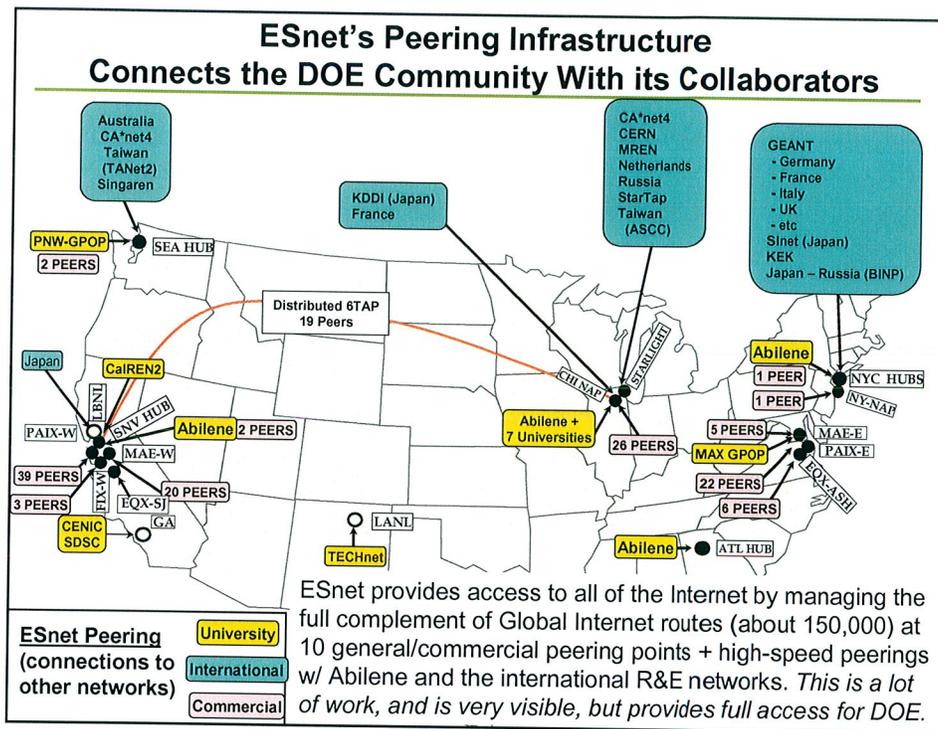
Michael S. Collins, Stan Kluz,  
Joseph Burreasca, and James V. Gagliardi, ESnet Leads

Gizella Kapus, Resource Manager

and the ESnet Team

Lawrence Berkeley National Laboratory





- ### Major ESnet Changes in FY04
- Dramatic increase in International traffic as major large-scale science experiments start to ramp up
  - Abilene high-speed crossconnects and CERNlink at 10 Gb/s
  - A new architectural approach to meet the Office of Science program needs
    - Second backbone
    - Metropolitan Area Networks (MANs)
  - Complete restructuring of the business practices of ESnet
  - Significant increases in the use of ESnet Science Services: PKI/Federated Trust and tele-collaboration

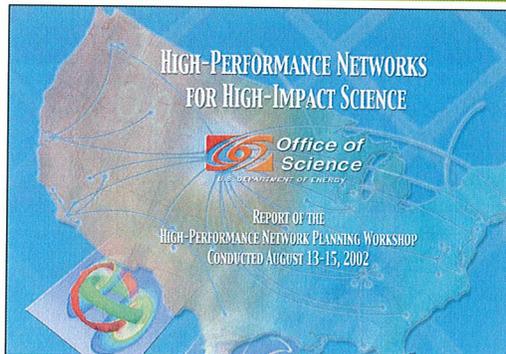
## The Message

The point of the first 2/3 of this talk is that

- 1) The predicted data intensive science environment is now real and can now be observed
- 2) ESnet has a new architecture and a plan for its implementation over the next 2-3 years to address the data increase
- 3) In the next major phase of ESnet (Fy07 and beyond) increased base funding will be needed to meet OSC requirements, though not huge amounts (probably 15-20% increase in base funding)

5

## Predictive Drivers for Change



August 13-15, 2002

Organized by Office of Science

Mary Anne Scott, Chair  
Dave Bader  
Steve Eckstrand  
Marvin Frazier  
Dale Koelling  
Vicky White

Workshop Panel Chairs

Ray Bair and Deb Agarwal  
Bill Johnston and Mike Wilde  
Rick Stevens  
Ian Foster and Dennis Gannon  
Linda Winkler and Brian Tierney  
Sandy Merola and Charlie Catlett

• Focused on science requirements that drive

- Advanced Network Infrastructure
- Middleware Research
- Network Research
- Network Governance Model

• The requirements for DOE science were developed by the OSC science community representing major DOE science disciplines

- Climate
- Spallation Neutron Source
- Macromolecular Crystallography
- High Energy Physics
- Magnetic Fusion Energy Sciences
- Chemical Sciences
- Bioinformatics

Available at [www.es.net/#research](http://www.es.net/#research)

## The Analysis was Driven by the Evolving Process of Science

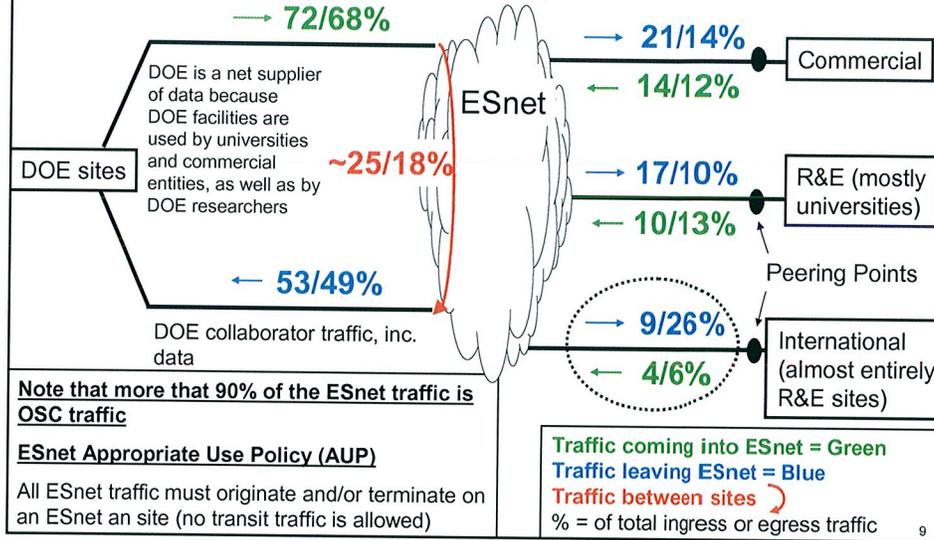
Discipline	Feature	analysis was driven by		Requirements	
		Vision for the Future Process of Science	Characteristics that Motivate High Speed Nets	Networking	Middleware
Climate (near term)	Analysis of model data by selected communities that have high speed networking (e.g. NCAR and NERSC)	<ul style="list-style-type: none"> <li>A few data repositories, many distributed computing sites</li> <li>NCAR - 20 TBy</li> <li>NERSC - 40 TBy</li> <li>ORNL - 40 TBy</li> </ul>	<ul style="list-style-type: none"> <li>Authenticated data streams for easier site access through firewalls</li> </ul>	<ul style="list-style-type: none"> <li>Server side data processing (computing and cache embedded in the net)</li> <li>Information servers for global data catalogues</li> </ul>	
Climate (5 yr)	Enable the analysis of model data by all of the collaborating community	<ul style="list-style-type: none"> <li>Add many simulation elements/components as understanding increases</li> <li>100 TBy / 100 yr generated simulation data, 1-5 PBy / yr (just at NCAR)                             <ul style="list-style-type: none"> <li>Distribute large chunks of data to major users for post-simulation analysis</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>Robust access to large quantities of data</li> </ul>	<ul style="list-style-type: none"> <li>Reliable data/file transfer (across system / network failures)</li> </ul>	
Climate (5-10 yr)	Integrated climate simulation that includes all high-impact factors	<ul style="list-style-type: none"> <li>5-10 PBy/yr (at NCAR)</li> <li>Add many diverse simulation elements/components, including from other disciplines - this must be done with distributed, multidisciplinary simulation</li> <li>Virtualized data to reduce storage load</li> </ul>	<ul style="list-style-type: none"> <li>Robust networks supporting distributed simulation - adequate bandwidth and latency for remote analysis and visualization of massive datasets</li> </ul>	<ul style="list-style-type: none"> <li>Quality of service guarantees for distributed, simulations</li> <li>Virtual data catalogues and work planners for reconstituting the data on demand</li> </ul>	

## Evolving Requirements for DOE Science Network Infrastructure

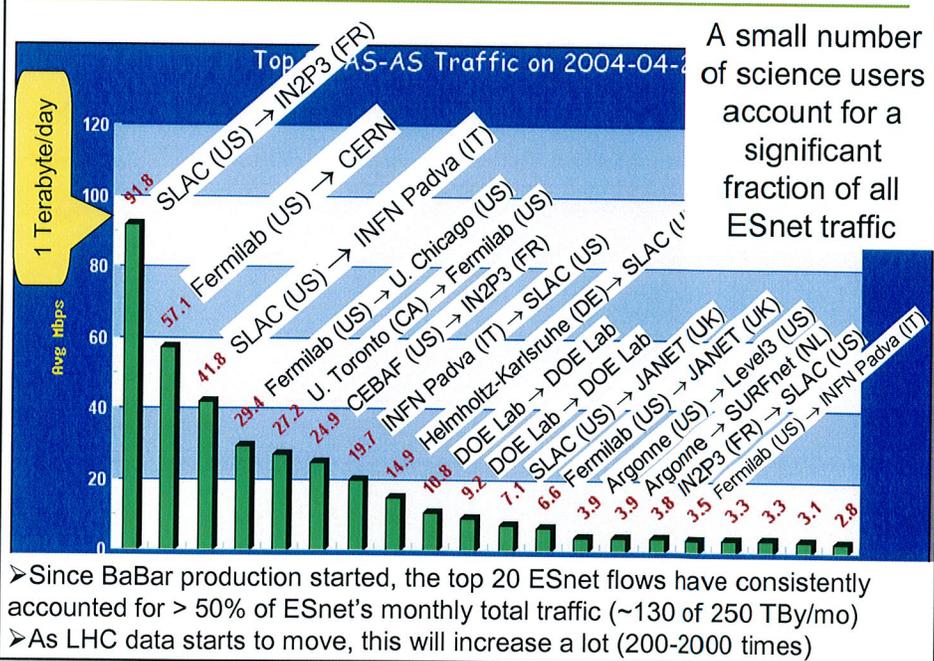
<b>1-3 yr Requirements</b>	<p style="text-align: center;">1-40 Gb/s, end-to-end</p> <p style="text-align: right;">             S storage              C compute              I instrument              C&amp;C cache &amp; compute           </p>	<b>2-4 yr Requirements</b>	<p style="text-align: center;">guaranteed bandwidth paths</p>
<b>3-5 yr Requirements</b>	<p style="text-align: center;">100-200 Gb/s, end-to-end</p>	<b>4-7 yr Requirements</b>	
<ul style="list-style-type: none"> <li>In the near term applications need higher bandwidth</li> </ul>	<ul style="list-style-type: none"> <li>high bandwidth and QoS</li> <li>network resident cache and compute elements</li> </ul>	<ul style="list-style-type: none"> <li>high bandwidth</li> <li>QoS</li> </ul>	<ul style="list-style-type: none"> <li>high bandwidth and QoS</li> <li>network resident cache and compute elements</li> <li>robust bandwidth (multiple paths)</li> </ul>

## Observed Drivers for Change

**ESnet Inter-Sector Traffic Summary,**  
**Jan 2003 / Feb 2004: 1.7X overall traffic increase, 1.9X OSC increase**  
 (the international traffic is increasing due to BABAR at SLAC and the LHC tier 1 centers at FNAL and BNL)

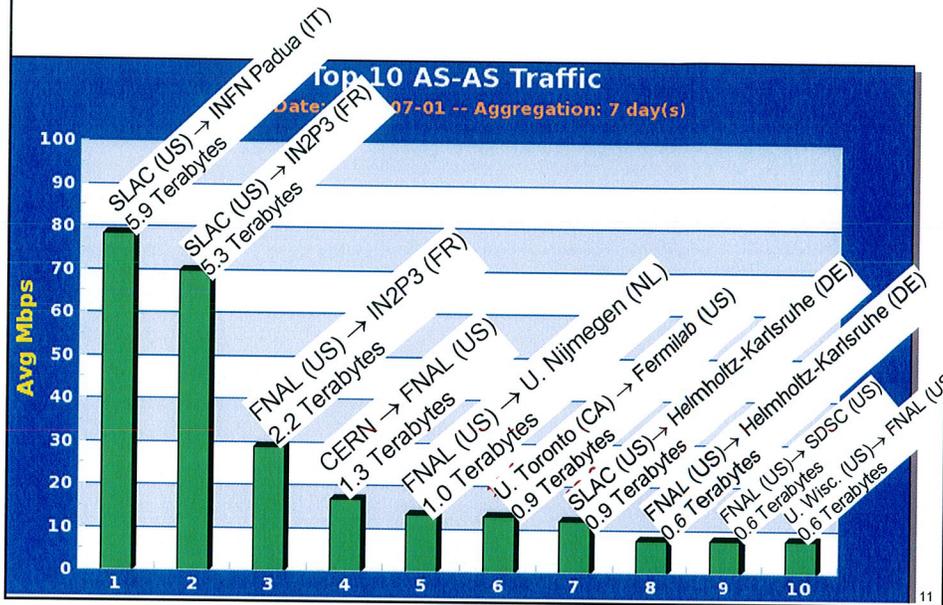


## ESnet Top 20 Data Flows, 24 hr. avg., 2004-04-20



## ESnet Top 10 Data Flows, 1 week avg., 2004-07-01

➤ The traffic is not transient: Daily and weekly averages are about the same.

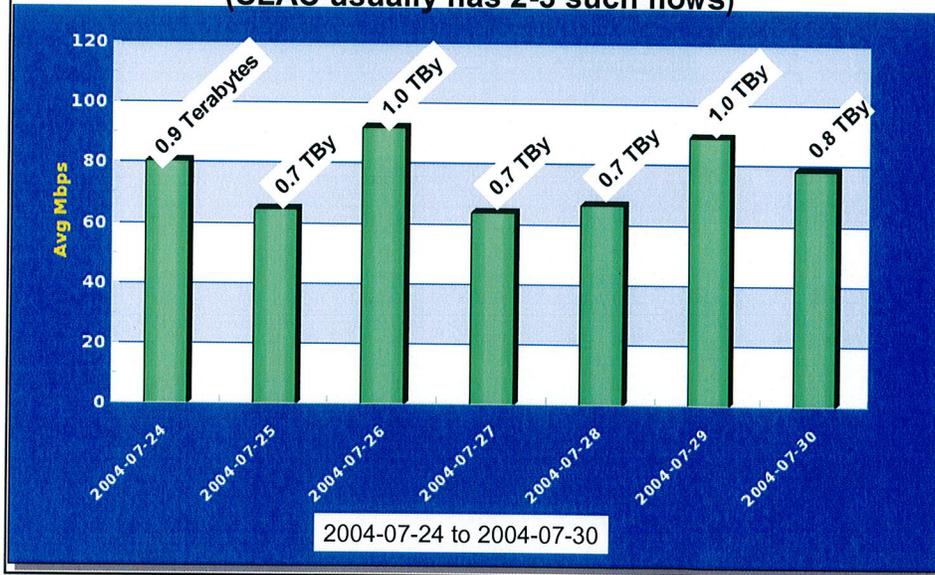


### A Typical Week at SLAC by Day for one flow:

SLAC to INFN, Padua

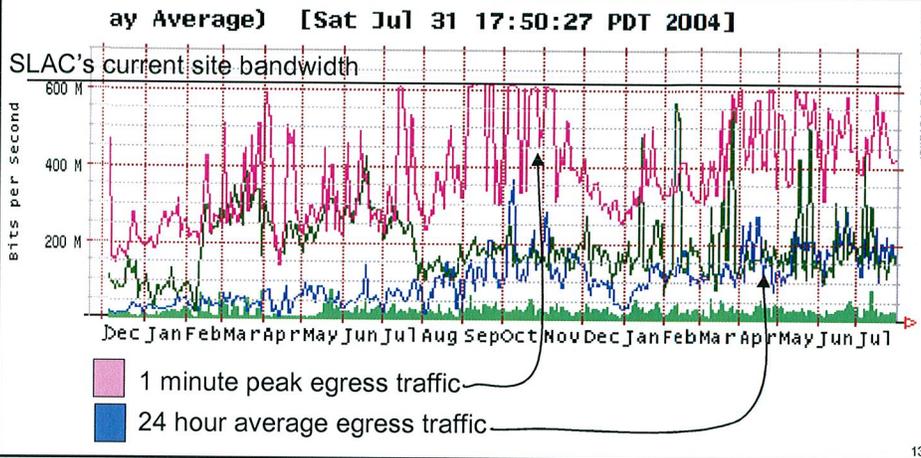
0.825 Terabytes/day, average

(SLAC usually has 2-3 such flows)



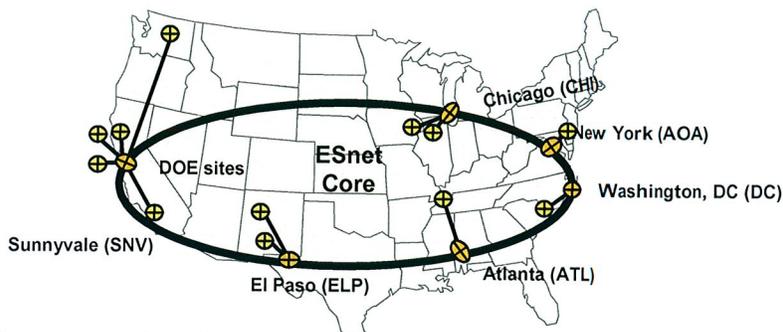
## SLAC – 24 hr average

- SLAC is a small prototype for what will happen when the big LHC experiments start moving data (200-2000 x BaBar)
- The flat tops on the 1 minute peaks means a lot of dropped packets, and programs that will not work well



## New ESnet Architecture Needed to Accommodate OSC

- The essential requirements cannot be met with the current, telecom provided, hub and spoke architecture of ESnet



- The core ring has good capacity and resiliency against single point failures, but the point-to-point tail circuits are neither reliable nor scalable to the required bandwidth

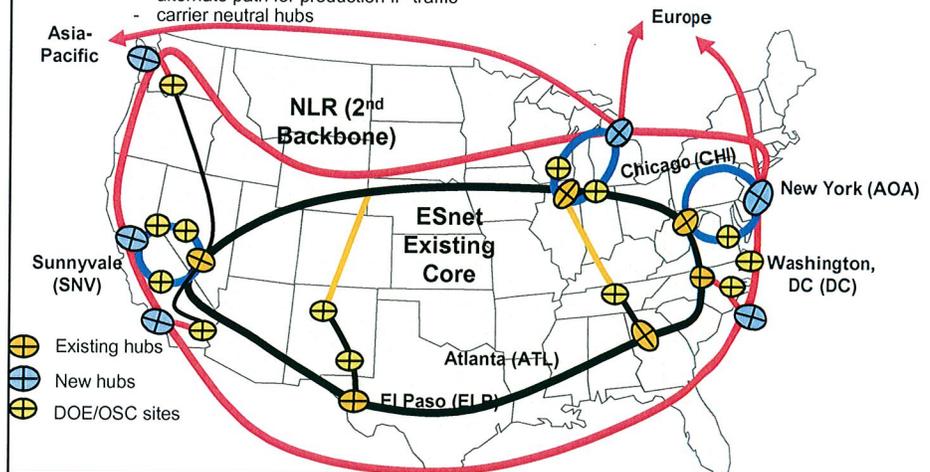
14

## A New Architecture

- With the current architecture ESnet cannot address
  - the increasing reliability requirements
  - the long-term bandwidth needs  
(incrementally increasing tail circuit bandwidth is too expensive – it will not scale to what OSC needs)
    - LHC will need dedicated 10 Gb/s into and out of FNAL and BNL
- ESnet can benefit from
  - Engaging the research and education networking community for advanced technology
  - Leveraging the R&E community investment in fiber and networks

15

- ESnet new architecture goals: full redundant connectivity for every site and high-speed access for every site (at least 10 Gb/s)
- Two part strategy
  - 1) MAN rings provide dual site connectivity and much higher site bandwidth
  - 2) A second backbone will provide
    - multiply connected MAN rings for protection against hub failure
    - extra backbone capacity
    - a platform for provisioned, guaranteed bandwidth circuits
    - alternate path for production IP traffic
    - carrier neutral hubs



## ESnet New Architecture: Why Start Now?

---

- Significant changes are needed to respond to the rapidly increasing demands of DOE science
  - It takes 2-3 years to make significant changes to the network with ESnet's small staff – LHC needs very high bandwidth in 2-3 years
  - No significant upgrades have occurred for two years and ESnet's ability to respond to new situations is decreasing
- The DOE science community is rapidly ramping up its network data transfers
  - LHC data challenges will start next year for both Atlas and CMS
  - Several OSC sites have inadequate bandwidth to the ESnet core
- Backbone changes are needed to accommodate the much higher speed site connectivity provided by the MANs
- High impact science capacity is needed for data intensive sites (site-to-site circuits)
- UltraScienceNet needs access to the Labs
- The Labs need better reliability than the current network can provide
- Future backbone transitions must be made easier and less costly (e.g. in FY07)

17

## ESnet New Architecture – Part 1

---

- MAN rings
  - SF Bay Area
  - Chicago MANs

(later)

18

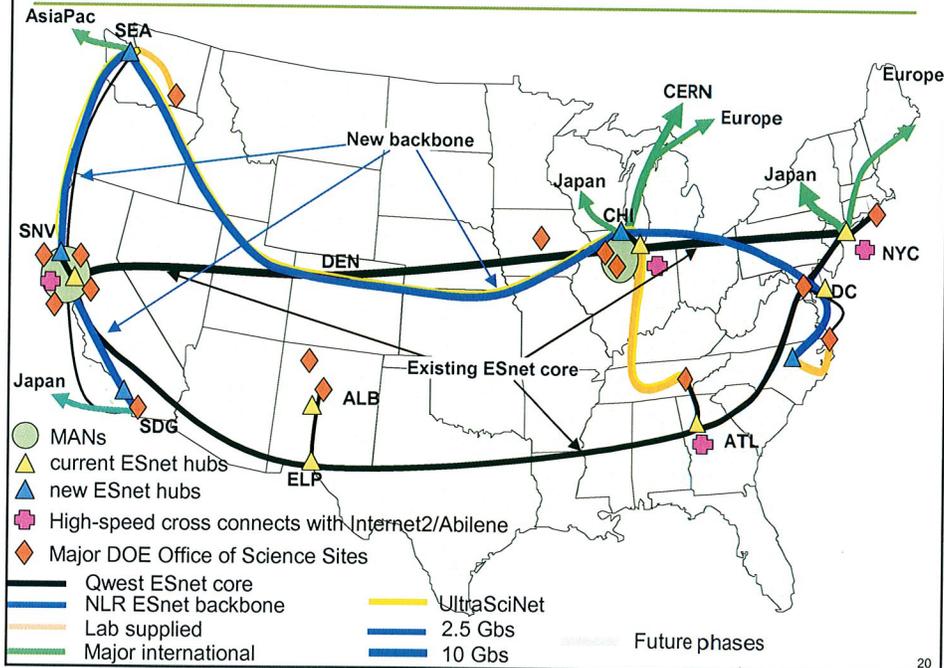
## ESnet New Architecture – Part 2: Second Backbone

Second Backbone: Phase 1 (FY05) and Phase 2 (FY06) Rationale

- Add major points of presence in carrier neutral facilities at Sunnyvale, Seattle, San Diego, and Chicago
  - Allow for more competition in acquiring circuits
  - Enable UltraSciNet cross-connect with ESnet
  - Provide access to NLR
- Initial steps toward second backbone (NLR)
  - Provide a second, independent path between major northern route hubs
    - Alternate route for ESnet core IP traffic
  - Provide for high-speed paths on the West Coast to reach PNNL, GA, and AsiaPac peering
  - Increase ESnet connectivity to other R&E networks
- Some interim fixes for outstanding problems in low access bandwidth to several OSC sites
  - Tail circuit upgrades
    - 10 Gb/s connection to CERN link
    - BNL to OC48
    - ORNL to OC48 and OC192 (Lab funded)
    - PPPL (maybe)
    - JLab (maybe – depends on SURAnet plans)

19

## ESnet New Architecture Goals – FY06

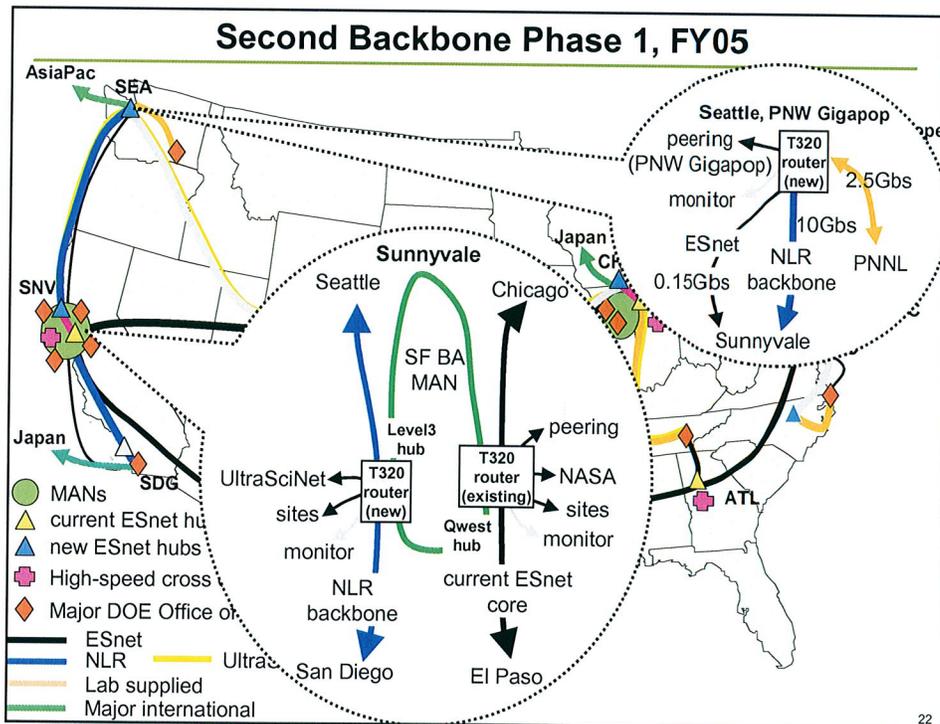


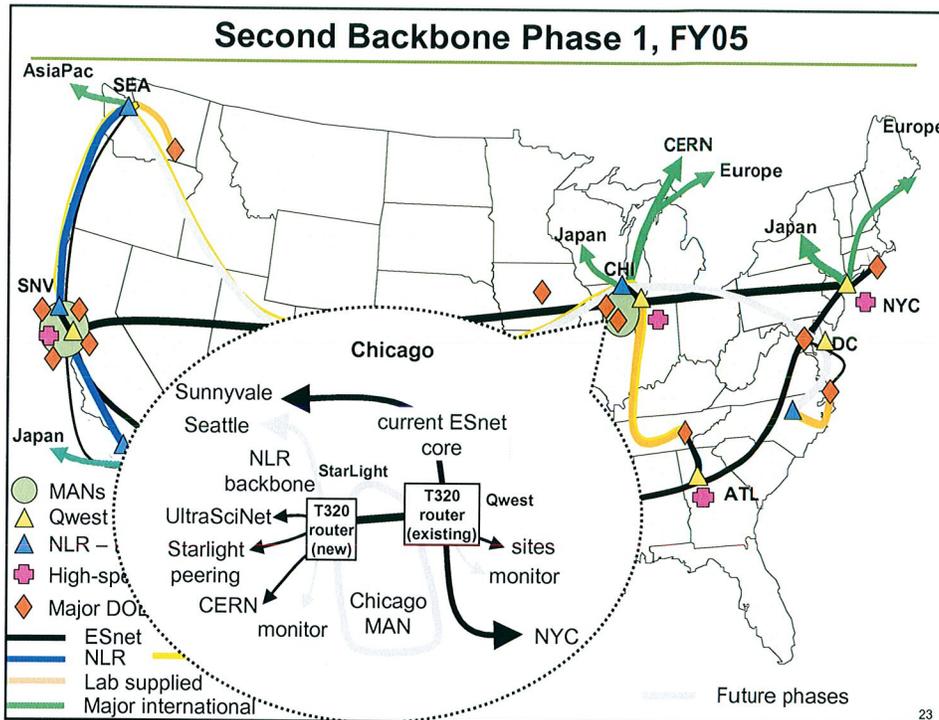
20

## How We Get to the New Architecture is a Compromise

- Have to phase in the new architecture
- Have to do some interim/temporary upgrades to site connections
- Two phases have been designed to get to fit within the available funds
- Chicago MAN is \$1.5M beyond current funds

21





### Second Backbone Phase 1 – FY05: Existing Funding

Action	Rationale
<ul style="list-style-type: none"> <li>• Add router to ESnet, Sunnyvale Level3 hub               <ul style="list-style-type: none"> <li>◦ \$ 1.1M cap. eq.</li> </ul> </li> <li>• Add router to ESnet, Chicago Level3 hub               <ul style="list-style-type: none"> <li>◦ \$1.1M cap. eq.</li> </ul> </li> <li>• Second backbone Sunnyvale to Seattle               <ul style="list-style-type: none"> <li>◦ NLR: \$ 105K install, \$14K/yr</li> </ul> </li> <li>• New router for Seattle hub               <ul style="list-style-type: none"> <li>◦ \$450K cap. eq.</li> </ul> </li> <li>• Second backbone Sunnyvale to San Diego               <ul style="list-style-type: none"> <li>◦ NLR: \$ 200K install, \$22K/yr</li> </ul> </li> <li>• New San Diego hub (10GE switch)               <ul style="list-style-type: none"> <li>◦ \$150K cap. eq.</li> </ul> </li> <li>• Upgrade PNNL to OC48 (2.5 Gbs)               <ul style="list-style-type: none"> <li>◦ \$ 100K</li> </ul> </li> <li>• Upgrade BNL to OC48 (2.5 Gbs)               <ul style="list-style-type: none"> <li>◦ \$49K cap. eq., +12K/yr</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• Add major points of presence in carrier neutral facilities at Sunnyvale, Seattle, San Diego, and Chicago               <ul style="list-style-type: none"> <li>◦ Allow for more competition in acquiring circuits</li> <li>◦ Enable UltraSciNet cross-connect with ESnet</li> <li>◦ Provide access to NLR</li> </ul> </li> <li>• Initial steps toward second backbone (NLR)               <ul style="list-style-type: none"> <li>◦ Provide a second, independent path between major northern route hubs                   <ul style="list-style-type: none"> <li>- Alternate route for ESnet core IP traffic</li> </ul> </li> <li>➢ Provide for high-speed paths on the West Coast to reach PNNL, GA, and AsiaPac peering</li> <li>◦ Increase ESnet connectivity to other R&amp;E networks</li> </ul> </li> <li>• Some interim fixes for outstanding problems in low access bandwidth to several OSC sites               <ul style="list-style-type: none"> <li>◦ Tail circuit upgrades                   <ul style="list-style-type: none"> <li>- 10 Gb/s connection to CERN link</li> <li>- BNL to OC48</li> <li>- ORNL to OC48 and OC192 (Lab funded)</li> <li>- PPPL (maybe)</li> <li>- JLab (maybe – depends on SURAnet plans)</li> </ul> </li> </ul> </li> </ul>

24

## Backbone Phase 1 – FY05: New Funding Needed

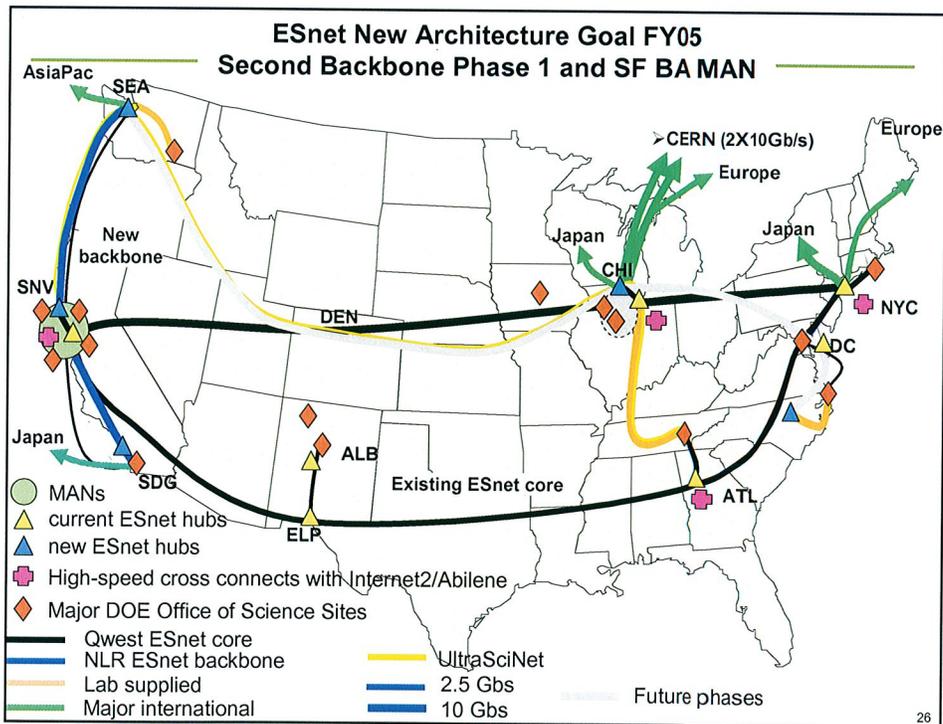
### Action

- Managing the added complexity of the network
  - 1 FTE (~220K)
  - + 200K/yr operational infrastructure (non cap. eq. equipment, software, licenses, etc.)

### Rationale

- ESnet is critically short on staff now

25



26

## Second Backbone Phase 2 – FY06: Existing Funding

Action	Rationale
<ul style="list-style-type: none"> <li>• NLR Seattle to Chicago               <ul style="list-style-type: none"> <li>◦ \$560K + \$51K/yr</li> </ul> </li> <li>• NLR Chicago to Raleigh               <ul style="list-style-type: none"> <li>◦ high speed access to JLab - (est.) \$300K + \$30K/yr</li> <li>◦ \$150 cap. eq.</li> <li>◦ first segment in southern ring</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• Add major points of presence in carrier neutral facilities (SNV, Seattle, San Diego, and Chicago)               <ul style="list-style-type: none"> <li>◦ Accommodate MANs</li> <li>◦ Allow for more competition in acquiring circuits</li> <li>◦ Enable UltraSciNet cross-connect with ESnet</li> </ul> </li> <li>Initial steps toward second backbone (NLR)               <ul style="list-style-type: none"> <li>◦ Increase routing diversity for reliability</li> <li>◦ Provide a second, independent path between major northern route hubs</li> <li>◦ Provide for high-speed paths on the West Coast: PNNL, GA, and AsiaPac peering</li> <li>◦ Increase ESnet connectivity to other R&amp;E networks</li> </ul> </li> <li>• Some interim fixes for outstanding problems in low access bandwidth to several OSC sites               <ul style="list-style-type: none"> <li>◦ Tail circuit upgrades                   <ul style="list-style-type: none"> <li>- BNL to OC48</li> <li>- ORNL to OC48 and OC192 (Lab funded)</li> <li>- PPPL (maybe)</li> <li>- JLab (maybe)</li> </ul> </li> </ul> </li> </ul>

27

## FY06 (New Funding Needed)

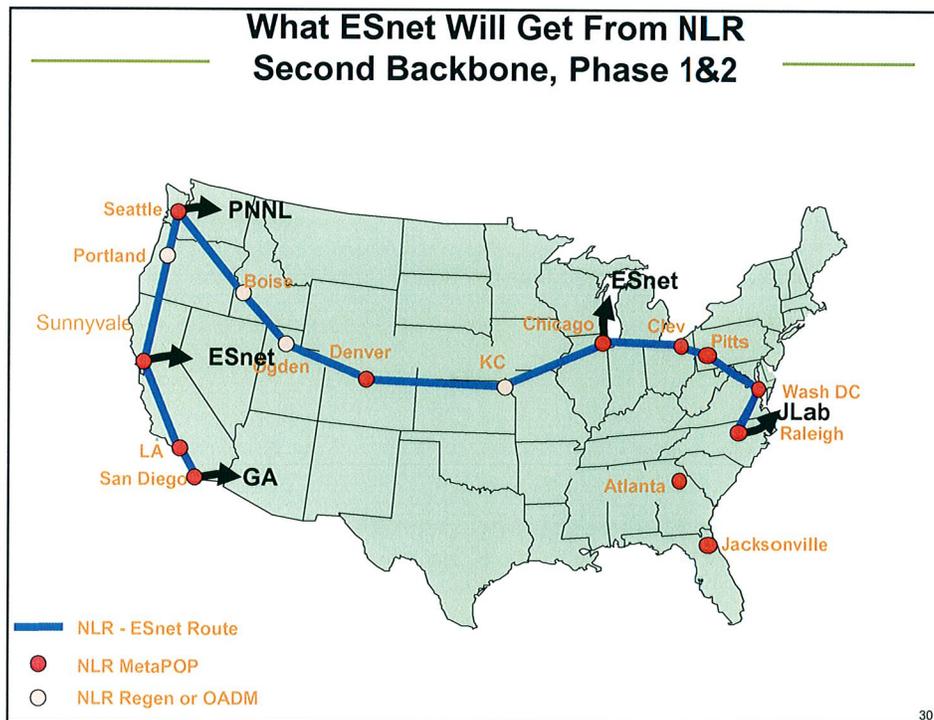
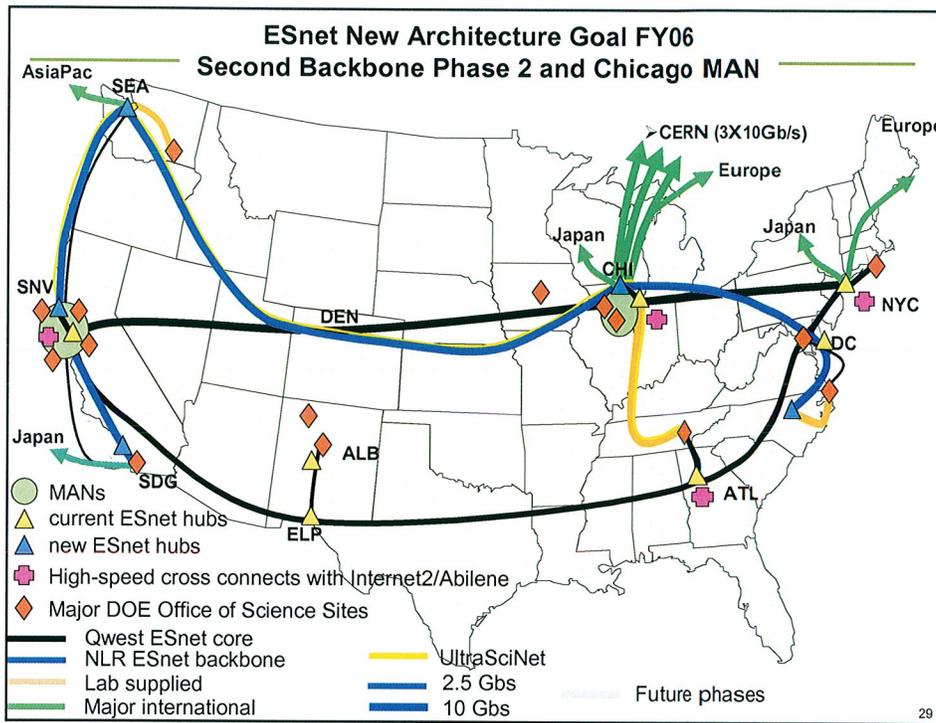
### Backbone Phase 2

Action	Rationale
<ul style="list-style-type: none"> <li>• Managing the added complexity of the network               <ul style="list-style-type: none"> <li>◦ 2 FTE (~220K/FTE)</li> <li>◦ + 400K/yr operational infrastructure (non cap. eq. equipment, software, licenses, etc.)</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• ESnet is critically short on staff now</li> </ul>

### MANs

- Chicago MAN
  - \$1.5M
  - \$500K of the \$1.5M total should be spent in FY05 for the fiber path FNAL to AN
    - Long lead time
    - The availability may not extend beyond FY05

28



## New Funding Needed – FY07

- **MANs**
  - LI MAN (FY06 or 07)
    - BNL at 10 Gb/s
    - Some ESnet participation will likely be required (guess \$1M)
    - Possible alternate access (alternate to MAN LAN)
  - VA MAN
    - Unknown circuit investment, if any
    - \$500K cap. eq.
  - NM MAN
    - Unknown circuit investment, if any
    - \$500K cap. eq.
- **Second Backbone Phase 3 (FY07)**
  - Complete the NLR ring (NYC loop and southern route)
    - ~\$1.3M for 5 years
      - \$900K of the \$1.3M paid in FY07 as an installation cost
      - \$86K/yr maint.
  - Note: the total cost of the NLR full ring, including FY05, 06, and 07, is ~\$2.8M for 5 years
    - This is about 1/10 what we currently pay for the ESnet core, and NLR provides more coverage
- **Monitoring for planning and cybersecurity**
  - ~\$150K/hub router (~\$1.5M total)

31

## ESnet FY07 and Beyond: One scenario may now be quantified

- Two primary assumptions
  - 1) Current Qwest 10 Gb/s Qwave @ \$4.2M/yr could continue to carry only production, 99.9% available, IP traffic
  - 2) 3 NLR lambdas would carry 30 Gb/s of high-impact science data
- NLR full ring (10 Gb/s) for high-impact science traffic
  - ~\$1M in FY07 for installation of the NLR southern route
  - ~\$450K/yr for subsequent 4 years for full ring
- NLR two additional 10 Gb/s circuits on northern route (San Diego to Jacksonville, FL)
  - ~\$1M paid in first year (FY08) for installation
  - ~\$450K/yr for subsequent 4 years
- If the full NLR ring is completed in FY07 and the second and third 10 Gb/s (northern) circuits are phased in, then the additional \$1M/yr will upgrade ESnet's overall capacity
  - from 10 Gb/s (northern) and 2.5 Gb/s (southern) today
  - to 40 Gb/s (northern) and 12.5 Gb/s (southern) by FY08/09

32

**ESnet FY07 and Beyond: One scenario may now be quantified**

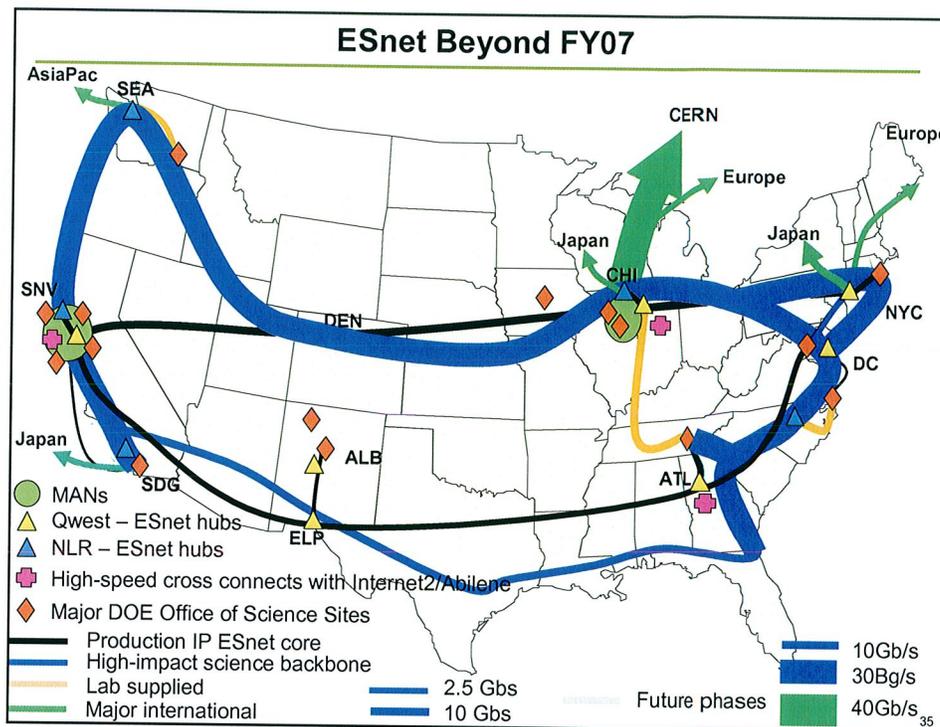
- The first assumption (Qwest 10 Gb/s production IP continuing to cost \$4.2M/yr) is pessimistic
  - Even if ESnet requires that the production IP is carried on full-up, commercial telecom circuits, that cost should drop significantly by FY07
  - A reduction of 50% in cost (not unreasonable) reduces communications cost by \$2.1M
    - This would mean that the whole 40 Gb/s northern +12.5 Gb/s southern scenario is cost neutral with respect to circuits

33

**Proposed NLR – ESnet Infrastructure – FY08**



34



### The Message

The point of the first 2/3 of this talk is that

- 1) The predicted data intensive science environment is now real and can now be observed
- 2) ESnet has a new architecture and plan for the next 2-3 years to address the data increase
- 3) In the next major phase of ESnet (Fy07 and beyond) increased funding will be needed to meet OSC requirements, though not huge amounts (probably 15-20%)

### Major ESnet Changes in FY04

- Dramatic increase in International traffic as major large-scale science experiments start to ramp up
- Abilene high-speed cross connects and CERNlink at 10 Gb/s
- A new architectural approach to meet the Office of Science program needs
  - Second backbone
    - Metropolitan Area Networks (MANs)
- Complete restructuring of the business practices of ESnet
- Significant increases in the use of ESnet Science Services: PKI/Federated Trust and tele-collaboration

37

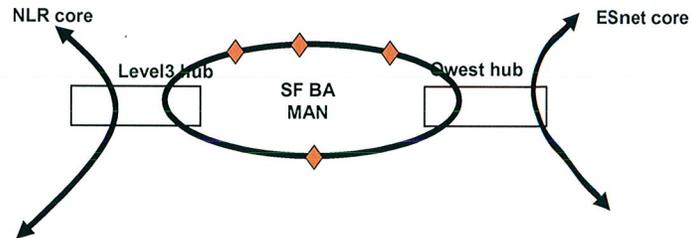
### Bay Area MAN

- We have in hand the funding and
  - an engineering study (as described in the proposal to MICS) from CENIC
    - CENIC is the Calif. higher-education-serving network that supplies high-speed networking to all UC campuses and all of the other major Calif. Universities (USC, Stanford, CalTech, etc.)
  - an as yet informal proposal from Qwest

38

## New Architecture: Integration of MANs and two Backbones

- To achieve the required full redundancy we need

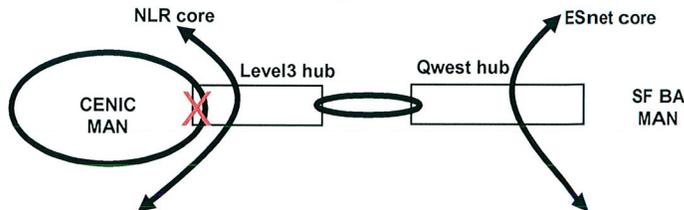


## Bay Area MAN

- CENIC
  - o ESnet partners with CENIC, and CENIC will
    - Provision another lambda on their existing infrastructure
    - Builds lateral out to the Labs where they do not now have CENIC access
      - They have build to most of the Labs – only LBNL and SLAC are close to CENIC to existing CENIC fiber
    - Purchase and operate the layer 2 network (10 Gbs Ethernet switches) as part of the proposal
    - Provide a topological ring that connects SLAC, LBNL, JGI, NERSC, and Level3 SNV
      - LLNL adds \$900K to the \$1.6M proposed
  - o Advantages
    - CENIC is a well established, important R&E network
    - Cost is good
  - o Disadvantages
    - They can only get into one of the two hubs in Sunnyvale, so ESnet will not have the resulting ring independently connected to both Qwest and Level3
    - They (so far) insist on operating the layer 2 network (Ethernet switches)
      - This potentially does not give ESnet sufficient flexibility to manage the network and to provide provisioned circuits

## The Issue of How the Ring Connects to the Backbones

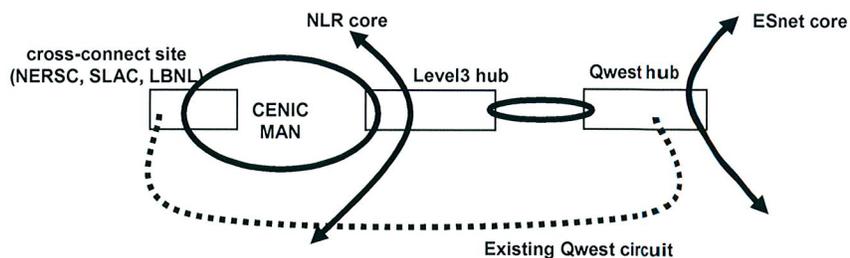
- CENIC disadvantage
  - Has a potential single-point failure in the Level3 hub that could isolate the MAN ring from both backbones



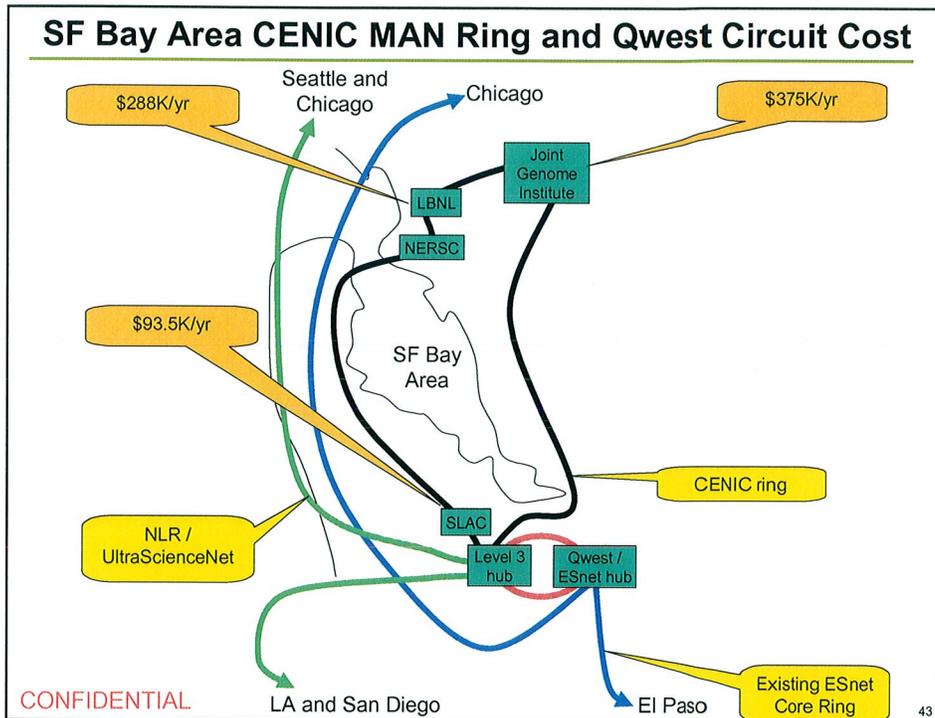
41

## The Issue of How the Ring Connects to the Backbones

- For full redundancy as in the new architecture, the CENIC proposal would require keeping some current circuits in place to connect both cores to the MAN ring, as below
  - Will need a router (\$400K existing funding) at a cross-connect site
  - Could not disconnect the cross-connect site Qwest circuit



42



### Bay Area MAN

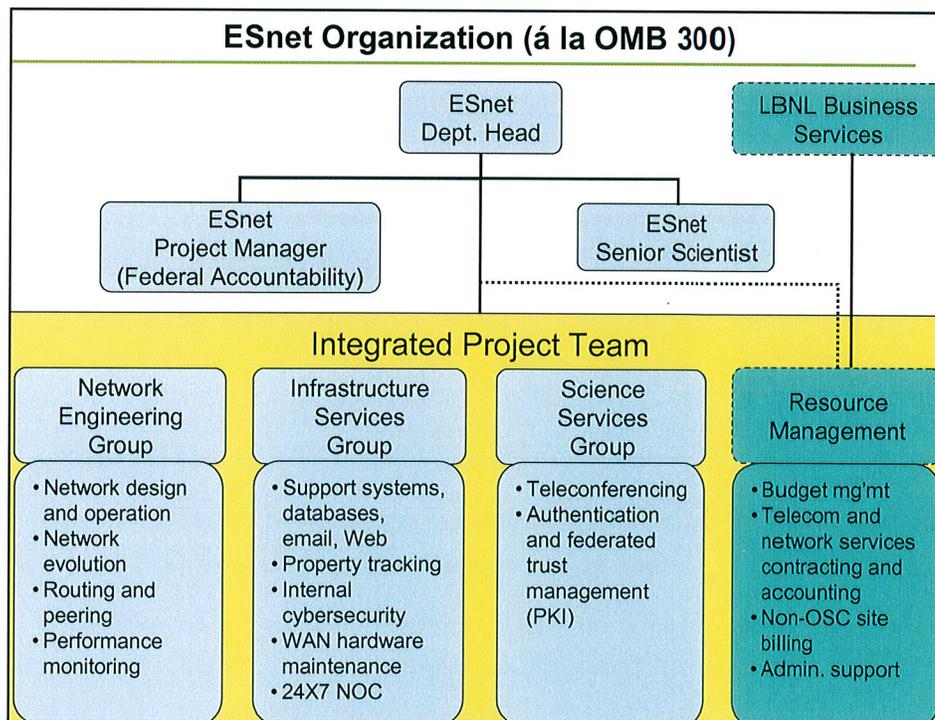
- Qwest informal proposal
  - Qwest has informally proposed 4 X 10 Gb/s rings
  - Qwest will purchase Ethernet switches as part of the deal (they probably get deeper discounts than ESnet does)
  - Advantages
    - Four times the bandwidth of the CENIC proposal
    - includes LLNL
    - touches both the Qwest hub and the Level3 hub
    - A simple ring configuration that involves no internal switches (CENIC ring collapses in the North Bay where several Ethernet switches implement the ring topology)
    - ESnet will own and operate the layer 2 network (Ethernet switches)
    - Cost is (maybe) the same as CENIC
  - Disadvantages
    - ESnet misses the opportunity to partner with CENIC
    - The actual cost may be somewhat more than CENIC
      - This is one reason for the \$500K contingency fund in FY05

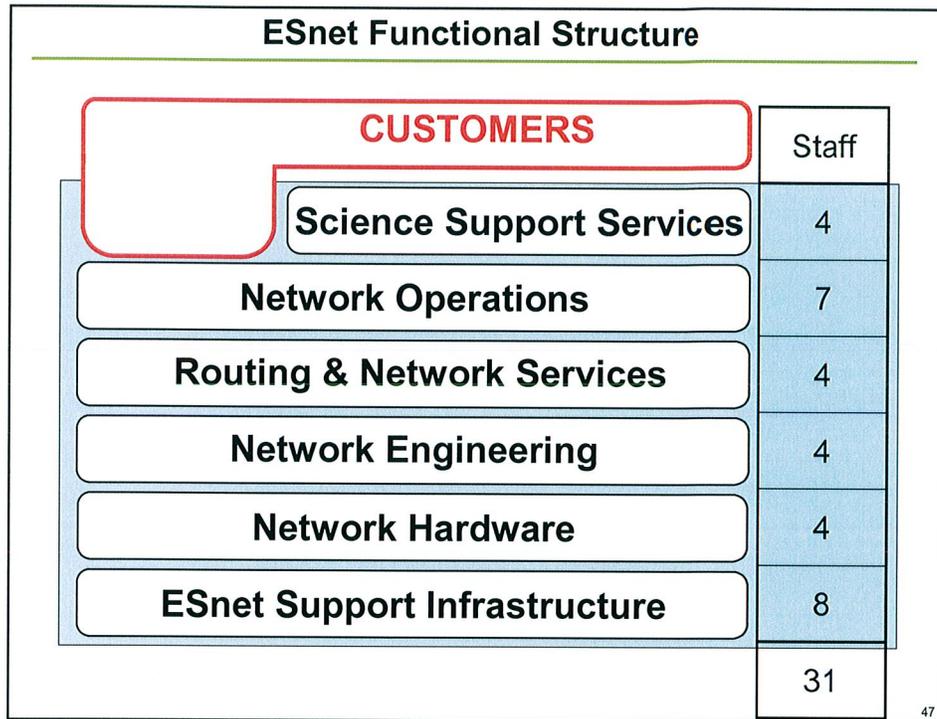
44

## Major ESnet Changes in FY04

- Dramatic increase in International traffic as major large-scale science experiments start to ramp up
- Abilene high-speed cross connects and CERNlink at 10 Gb/s
- A new architectural approach to meet the Office of Science program needs
  - Second backbone
    - Metropolitan Area Networks (MANs)
- Complete restructuring of the business practices of ESnet
- Significant increases in the use of ESnet Science Services: PKI/Federated Trust and tele-collaboration

45





- ### Work Breakdown Structure
- ESnet is not like a software or hardware construction project
  - ESnet has three components
    - 1) a steady state operation
    - 2) planned upgrades
    - 3) responses to external drivers, many of which are not predictable
      - major network problems like the Qwest ATM breakdown
      - routing failures, multicast failures, site outage, etc.
      - these can represent significant engineering challenges that require unplanned, on-the-fly work, sometimes on a large scale
      - such work will preempt, sometimes for long periods, progress on planned activities
  - Therefore a Work Breakdown Structure (WBS) analysis approach is used
- 48

WBS	Description	Technical	sub-tots	Admin
		FTE	5.25	FTE
1.0	Science Support Services			
1.1	Video, audio, data conferencing	1.50		
1.2	PKI	2.50		
1.3	E-mail lists	0.05		
1.4	Net News	0.05		
1.5	Web services	0.10		
1.6	Directory services (future)	1.00		
1.7	Conferences (ESSC)	0.05		
2.0	Routing and Network Services		4.90	
2.1	Network Support Applications and Services (DNS, etc.)	0.65		
2.2	Network measurements and monitoring	2.25		
2.3	Network testing lab	0.10		
2.4	New Network Capabilities (Engineering and deploying new capabilities in ESnet)	2.00		
3.0	Network Operations		4.20	
3.1	WAN Network Security	0.10		
3.2	On Call Support (OCS & POW)	2.00		
3.3	24x7 Help Desk - NERSC OSF	2.00		
3.4	Remote Centers			
3.4.1	NERSC OSF - Oakland, CA	0.05		
3.4.2	TalisWork Center - Livermore, CA	0.05		
3.4.3	Ames, IA	0.00		
3.4.4	BNL, NY	0.00		
4.0	Network Hardware		3.75	
4.1	Network Engineering	3.50		
4.2	Investigating Advanced Tech	0.25		
5.0	ESnet Support Infrastructure		7.90	
5.1	ESnet LAN Support	0.70		
	Network Conference Support (SC, I2, ESSC)	0.15		
	Unix servers at LBNL	1.75		
	Windows servers at LBNL	0.50		
	Staff desktops	0.75		
	Oracle, Remedy, Trouble Ticket System	0.75		
	Internal security	1.00		
5.9	Asset management	0.50		
5.9	Internal PKI	0.50		
5.10	Internal email	0.20		
5.11	Internal Web	0.10		
5.12	Data center (performance databases and archives)	1.00		
6.0	Management, Administration, and Accounting		2.00	
6.1	Manager	1.00		
6.2	Group Leads	1.00		
6.3	Project Administration			1.00
6.4	Staff Administration			1.00
6.5	Accounting			1.00
TOTAL FTE			20.00	3.00

### WBS Based on ESnet Functionality

We are tracking 30, or so, cost centers, and report at the level of the 6 top-level rollups.

Update

### Cost Centers

- The cost centers do not reflect the fact that the organization is so small that almost everyone is cross-trained in network operation
- As noted, in an emergency or configuration transition almost everyone will have to do network engineering and operation for extended periods of time even though that is not where their cost is centered, and regardless what other projects they might be working on

## Major ESnet Changes in FY04

- Dramatic increase in International traffic as major large-scale science experiments start to ramp up
  - A new architectural approach to meet the Office of Science program needs
    - Bay Area MAN
  - Complete restructuring of the business practices of ESnet
- Other progress – Next talks

51

## Conclusions

- ESnet is an infrastructure that is critical to DOE's science mission and that serves all of DOE
- ESnet is working on providing the DOE mission science networking requirements with several new initiatives and a new architecture
- ESnet is very different today in both planning and business approach and in goals than in the past

52

## Funding Elements

ESnet funding elements independent of FY

- Base (communications+operations+infrastructure+production Science Services)
- Research
  - PKI – SciDAC funding
    - has evolved to Federated Trust, with RADIUS One-Time-Password distributed, federated service as the current focus
  - OSCARS dynamic provisioned circuits – Networking R&D
- Capital equipment
  - Annual upgrades to all ESnet equipment to maintain viability
- Special projects
  - MANs

53

# ESnet Financial Review

**William Fortney**  
**Computing Sciences**  
**Business Manager**

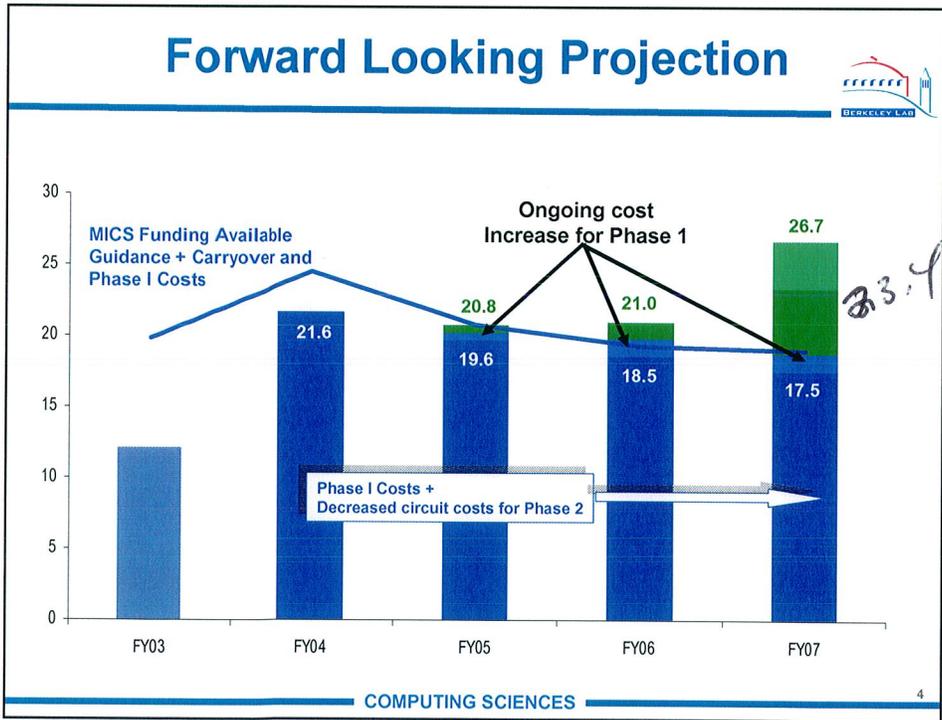
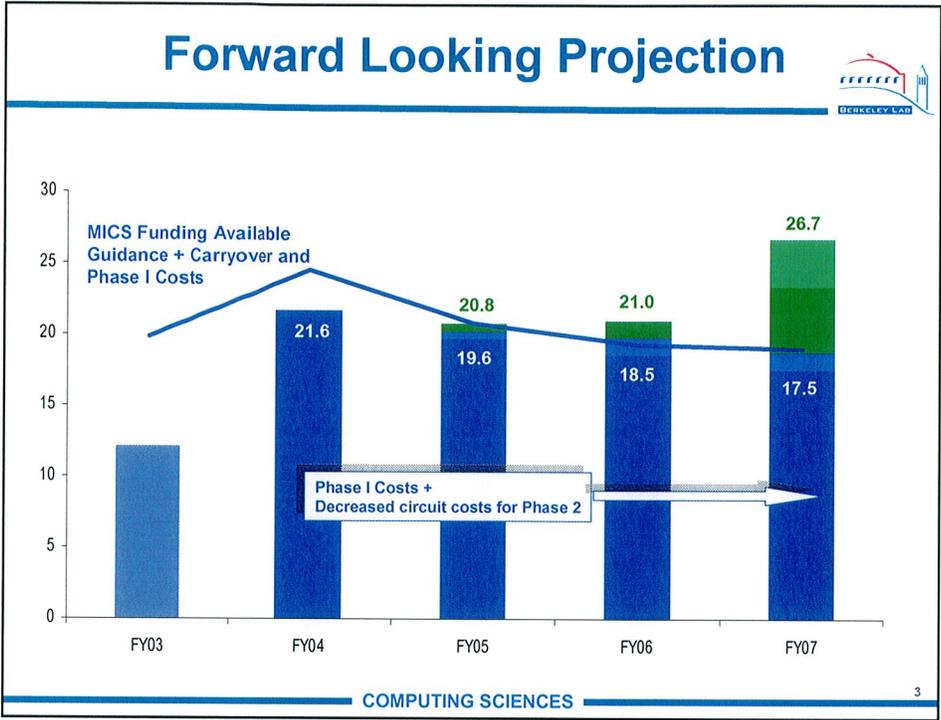
August 3, 2004

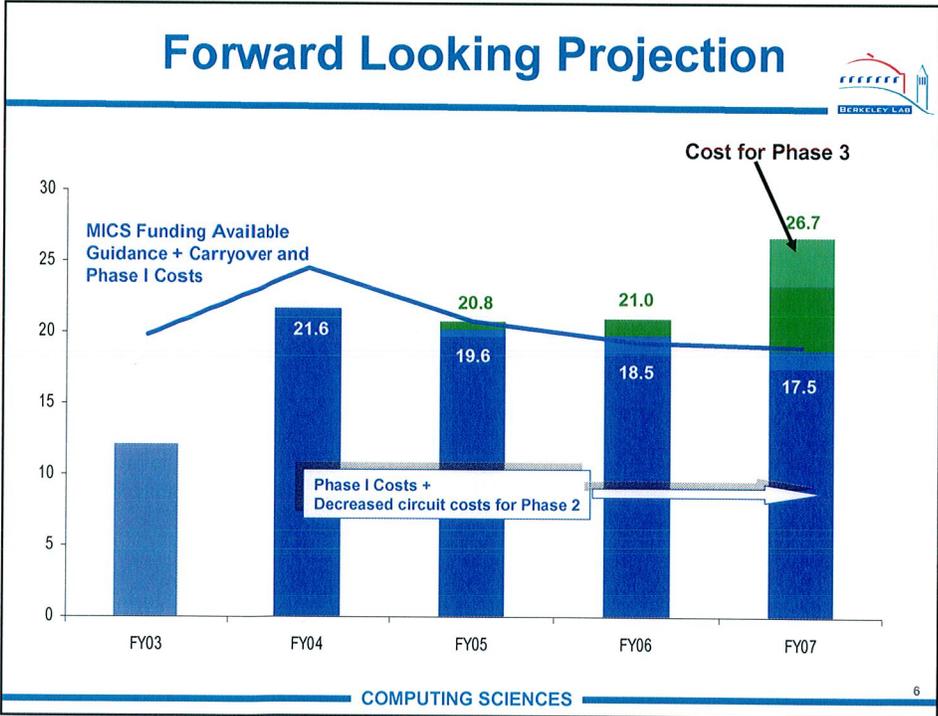
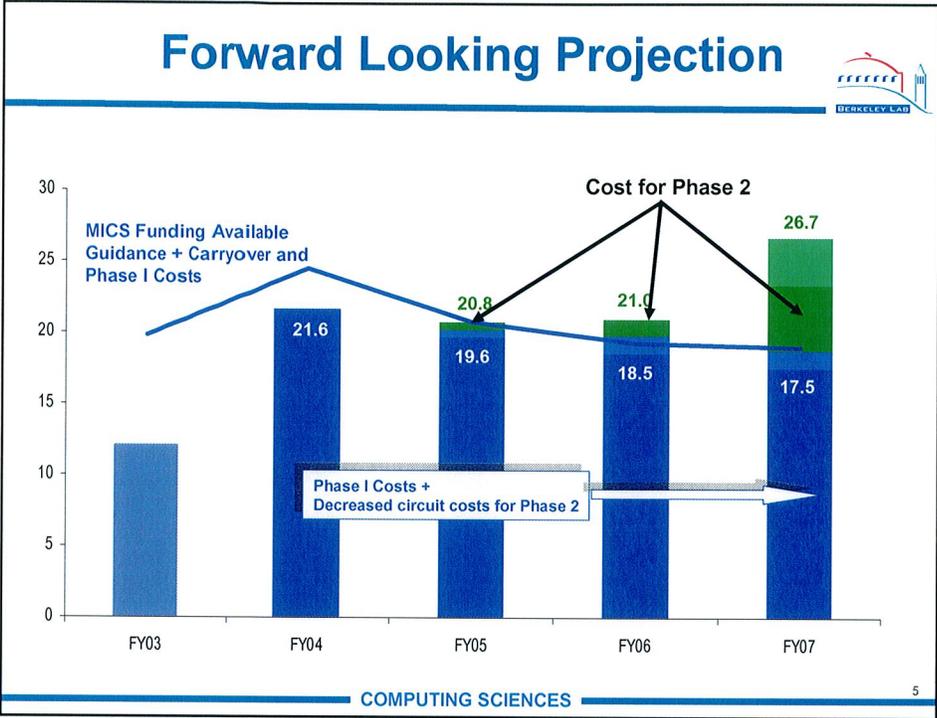


## Overview



- **Forward looking projection to FY07**
  - Bar graph forecast of total program costs
  - Detailed cost table
  - Table identifying additional funding
- **Business model discussion**
  - Financial reporting transparency
  - The organization and reporting relationships
  - Significant business processes and practices





## Projected Cost Details



	FY04	FY05	FY06	FY07
NLR Phase 3+ (Atlanta - San Diego)				3.302
CHI & LI MAN (Qwest disconnect)		.500	1.170 (.174)	1.270 (.504)
LI, VA, NM MAN Capital Investment				3.321
FTE Increase for Phase 1		.220 1 FTE	.680 3 FTE	.680 3 FTE
SBA & NLR Phase 1 Capital Investment	3.330			
NLR Phase 1		.335	1.010	.110
SBA MAN (Qwest disconnect)		1.867 (.331)	.267 (.662)	.267 (.662)
Ongoing Capital	.900	.900	.900	.900
Qwest Backbone Ring	4.322	4.322	4.322	4.322
Tail Circuits (net Qwest disconnects for SBA, CHI & LI MANs)	5.854	4.484	3.989	3.659
All other base & research costs	7.227	7.728	8.036	8.236
<b>TOTAL REQUIRED ESnet BUDGET</b>	<b>21.6</b>	<b>20.8</b>	<b>21.0</b>	<b>26.7</b>

*rephased*

*is c  
inf*

## Required Carry-over and additional Funding from DOE



	FY04 Actual funding	FY05 Projected Funding	FY06	FY07
Base Funding (Operating & Capital)	14.33	17.45	17.45	17.45
Research (OSCARS, Networking R&D/SciDAC)	.85	.35	.75	.75
SBA MAN (Operating)	1.60			
Funds Required to meet ESnet plans				
Carry-over(Operating & Capital)	7.73	2.89	1.05	.74
Add'l Funds Required	0	.07	1.73	8.17
<b>Total Funding</b>	<b>24.5</b>	<b>20.8</b>	<b>21.0</b>	<b>27.1</b>

Additional funds required to meet upgrade plans highlighted

## Required Carry-over and additional Funding from DOE



DOE Funding Table				
	FY04 Actual funding	FY05 Projected Funding	FY06	FY07
Base Funding	16.55	16.55	16.55	16.55
Research (OSCARS, Networking R&D/SciDAC)	.85	.35	.75	.75
SBA MAN	1.60			
Funds Required to meet ESnet plans				
Carry-over	7.73	2.89	1.05	.74
Add'l Funds Required	0	.07	1.73	8.17
<b>Total Funding</b>	<b>24.5</b>	<b>20.8</b>	<b>21.0</b>	<b>27.1</b>

Additional funds required to meet upgrade plans highlighted

**Projected costs are 21.6 resulting in a carry-over from FY04 to FY05**

## Required Carry-over and additional Funding from DOE



DOE Funding Table					
	FY04 Actual funding	FY05 Projected Funding	FY06	FY07	FY08
Base Funding	16.55	16.55	16.55	16.55	
Research (OSCARS, Networking R&D/SciDAC)	.85	.35	.75	.75	
SBA MAN	1.60				
Funds Required to meet ESnet plans					
Carry-over	7.73	2.89	1.05	.74	.46
Add'l Funds Required	0	.07	1.73	8.17	
<b>Total Funding</b>	<b>24.5</b>	<b>20.8</b>	<b>21.0</b>	<b>27.1</b>	

Additional funds required to meet upgrade plans highlighted

**Projected costs are 26.7, results in a .46 carry-over in FY08**

*cross*

# Overview

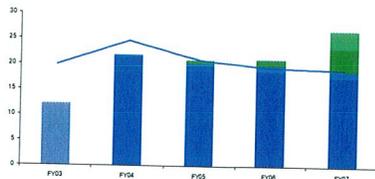


- Forward looking projection to FY07
  - Bar graph forecast of total program costs
  - Detailed cost table
  - Table identifying additional funding
- Business model discussion
  - Financial Reporting Transparency
  - The organization and reporting relationships
  - Significant business processes and practices

# Financial Transparency

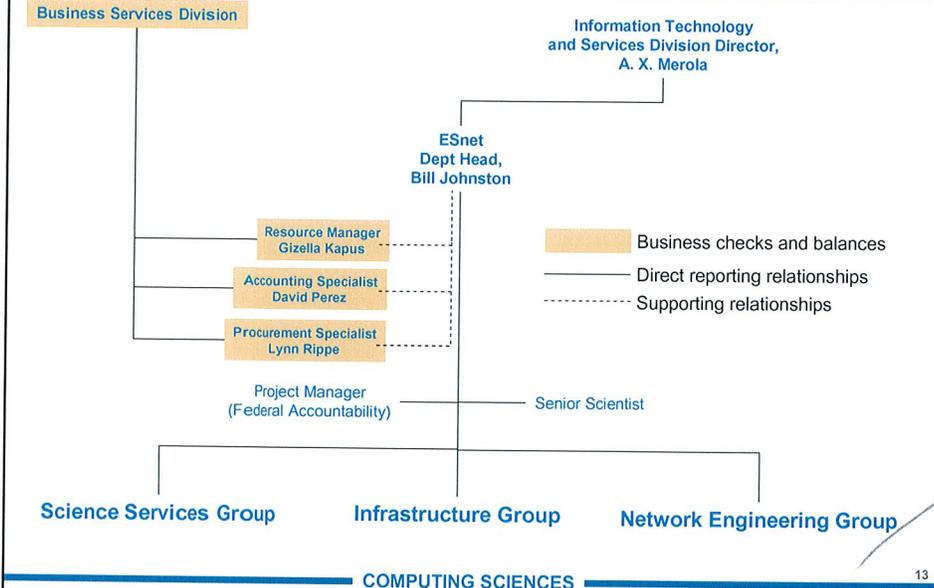


- Full disclosure through:
  - Financial Statements
  - Funding/cost breakdowns
- Describes impact of past, current and future plans



	FY99	FY00	FY01	FY02	FY03	FY04	FY05	FY06	FY07
<b>OPERATING FUNDING</b>									
Base						10.85	16.65	19.35	19.75
Research (OSCARs & PKI)						0.85	0.35	0.75	0.75
Special Projects						1.60			
<b>NEW GUIDANCE sub-total</b>	<b>14.35</b>	<b>13.90</b>	<b>16.90</b>	<b>18.15</b>	<b>16.47</b>	<b>13.30</b>	<b>17.00</b>	<b>20.10</b>	<b>20.50</b>
Carry-over	0.10	2.16	1.07	1.15	2.44	6.99	2.89	0.01	0.02
<b>TOTAL FUNDING</b>	<b>14.45</b>	<b>16.06</b>	<b>17.96</b>	<b>19.29</b>	<b>18.91</b>	<b>20.29</b>	<b>19.89</b>	<b>20.11</b>	<b>20.52</b>
<b>COSTS</b>									
Staff	2.92	3.21	3.42	3.78	4.07	4.04	4.55	5.12	5.22
Communications	5.37	6.90	7.88	7.25	2.96	10.18	11.51	10.75	10.57
OSCARs PKI									
Backbone						4.32	4.32	4.32	4.32
Tail Circuits						5.85	4.48	3.98	3.65
SBA MAN							1.87	0.27	0.27
NLR							0.34	1.01	1.06
LI MAN									1.10
CHI MAN							0.50	1.17	0.17
Network Infrastructure Maintenance	0.68	0.55	0.80	0.44	0.84	0.88	1.02	1.22	1.22
Operating Materials and Services	1.70	1.50	1.86	2.36	0.89	0.63	0.65	0.66	0.67
Travel & Training	0.35	0.23	0.29	0.21	0.23	0.15	0.21	0.21	0.21
Recharges	0.57	0.58	0.59	0.60	0.57	0.57	0.58	0.59	0.60
Misc	(1.15)	0.05	0.03	0.03	0.01	0.01	0.01	0.01	0.01
Operations Cost Recovery from ICOS						(1.45)	(1.46)	(1.46)	(1.46)
Burdens	1.86	1.97	1.95	2.19	2.33	2.49	2.82	3.00	3.05
<b>TOTAL OPERATING EXPENSES</b>	<b>32.29</b>	<b>34.99</b>	<b>36.82</b>	<b>36.85</b>	<b>31.91</b>	<b>47.10</b>	<b>49.88</b>	<b>50.09</b>	<b>50.09</b>
net balance	2.16	1.07	1.15	2.44	6.99	2.89	0.01	0.02	0.43
<b>CAPITAL FUNDING</b>									
Capital	0.68	1.50	0.89	0.66	0.90	3.60	0.90	0.90	6.80
less 3%	-0.03	-0.05	-0.03	-0.02	-0.03	-0.11	-0.08	-0.08	-0.20
<b>NEW GUIDANCE sub-total</b>	<b>0.65</b>	<b>1.46</b>	<b>0.86</b>	<b>0.64</b>	<b>0.87</b>	<b>3.49</b>	<b>0.82</b>	<b>0.82</b>	<b>6.60</b>
Carry-over	0.13	-0.11	0.07	0.15	0.04	0.73	0.00	0.00	0.00
<b>TOTAL FUNDING</b>	<b>0.98</b>	<b>1.34</b>	<b>0.93</b>	<b>0.79</b>	<b>0.91</b>	<b>4.23</b>	<b>0.82</b>	<b>0.82</b>	<b>6.60</b>
<b>COSTS</b>									
<b>TOTAL CAPITAL EXPENSES</b>	<b>1.09</b>	<b>1.27</b>	<b>0.79</b>	<b>0.75</b>	<b>0.18</b>	<b>4.23</b>	<b>0.87</b>	<b>0.87</b>	<b>6.57</b>
net balance	-0.11	0.07	0.15	0.04	0.73	0.00	0.00	0.00	0.03
<b>NET OPERATING &amp; CAPITAL</b>									
AVAILABLE FUNDS	15.4	17.4	18.9	20.1	19.8	24.52	20.76	20.98	27.11
GUIDANCE	15.2	15.4	17.8	18.8	17.3	16.79	17.87	20.97	27.10
COSTS	13.4	16.3	17.6	17.6	12.1	21.63	20.76	20.96	26.66

# Organization and Reporting Relationships



*ADD staffing to this*

# Service Orders



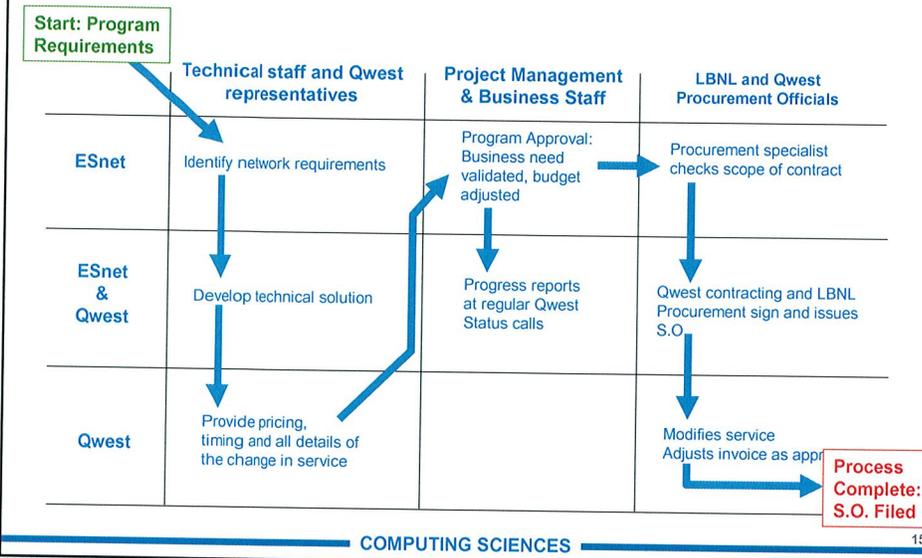
- Normal mode of doing business with Qwest
- Modifies the network connections, bandwidth and pricing so ESnet can meet OSC science requirements
- Involves the following:
  - Defining technical requirements
  - Establishing the business need
  - Review of the primary contract
  - Formal agreement between Qwest and an LBNL procurement specialist
- There have been 180 Service orders processed since the beginning of the contract, 20 of them in the last 12 months

*codify this so we need more support*

*copy files  
primary only*

*Best Some Select a w NLA*

## Processing Service Orders



## Criteria for Payment of Invoices

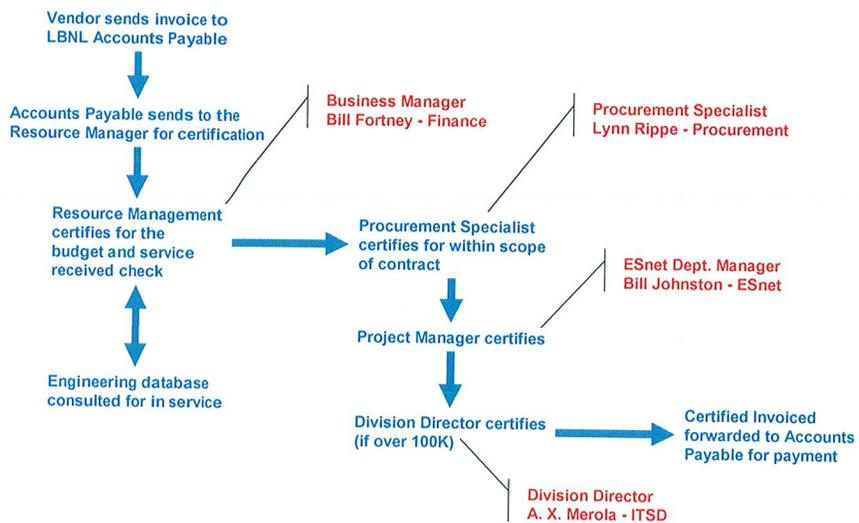


- Vendor is required to submit invoices directly to LBNL Accounts Payable before payment
- All communication invoices require certification from the ESnet Program
- Invoices over 100K require certification from the following area
  - Financial/Budget
  - Technical/Network Management
  - Procurement Specialist
  - An official with signature authority
- Invoice payments from Qwest are current and balanced as of 15 June 2004

# Invoice Certification Process



# Invoice Certification in Practice



## Conclusion



- **Financial transparency and control is improved since beginning of the year**
- **ESnet has a forecast for the forward looking vision presented today**
- **ESnet has a solid Business model that is extensible into the future**

## Overview



- **Forward looking projection to FY07**
  - Bar graph forecast of total program costs
  - Detailed cost table
  - Table identifying additional funding
- **Business model discussion**
  - Financial reporting transparency
  - The organization and reporting relationships
  - Significant business processes and practices

## Questions

## Annual Program Review of ESnet

August 3, 2004

all \$ in millions	FY99	FY00	FY01	FY02	FY03	FY04	FY05	FY06	FY07
<b>OPERATING FUNDING</b>									
Base						10.85	16.65	19.35	19.75
Research (OSCARs & PKI)						0.85	0.35	0.75	0.75
Special Projects						1.60	-	-	-
<b>NEW GUIDANCE sub-total</b>	<b>14.35</b>	<b>13.90</b>	<b>16.90</b>	<b>18.15</b>	<b>16.47</b>	<b>13.30</b>	<b>17.00</b>	<b>20.10</b>	<b>20.50</b>
Carry-over	0.10	2.16	1.07	1.15	2.44	6.99	2.89	0.01	0.02
<b>TOTAL FUNDING</b>	<b>14.45</b>	<b>16.06</b>	<b>17.96</b>	<b>19.29</b>	<b>18.91</b>	<b>20.29</b>	<b>19.89</b>	<b>20.11</b>	<b>20.52</b>
<b>COSTS</b>									
Staff	2.92	3.21	3.42	3.78	4.07	4.04	4.55	5.12	5.22
<b>Communications</b> <small>(OSC costs only)</small>	5.37	6.90	7.88	7.25	2.96	10.18	11.51	10.75	10.57
Backbone						4.32	4.32	4.32	4.32
Tail Circuits						5.85	4.48	3.98	3.65
SBA MAN							1.87	0.27	0.27
NLR							0.34	1.01	1.06
LI MAN							0.50	1.17	1.10
CHI MAN									0.17
<b>Network Infrastructure Maintenance</b>	0.68	0.55	0.80	0.44	0.84	0.88	1.02	1.22	1.22
<b>Operating Materials and Services</b>	1.70	1.50	1.86	2.36	0.89	0.63	0.65	0.66	0.67
Travel & Training	0.35	0.23	0.29	0.21	0.23	0.15	0.21	0.21	0.21
Recharges	0.57	0.58	0.59	0.60	0.57	0.57	0.58	0.59	0.60
Misc	(1.15)	0.05	0.03	0.03	0.01	0.01	0.01	0.01	0.01
<b>Operations Cost Recovery from ICOs</b>						(1.45)	(1.46)	(1.46)	(1.46)
Burdens	1.86	1.97	1.95	2.19	2.33	2.40	2.82	3.00	3.05
<b>TOTAL OPERATING EXPENSES</b>	<b>12.29</b>	<b>14.99</b>	<b>16.82</b>	<b>16.85</b>	<b>11.91</b>	<b>17.40</b>	<b>19.88</b>	<b>20.09</b>	<b>20.09</b>
<b>net balance</b>	<b>2.16</b>	<b>1.07</b>	<b>1.15</b>	<b>2.44</b>	<b>6.99</b>	<b>2.89</b>	<b>0.01</b>	<b>0.02</b>	<b>0.43</b>
<b>CAPITAL FUNDING</b>									
Capital	0.88	1.50	0.89	0.66	0.90	3.60	0.90	0.90	6.80
less 3%	-0.03	-0.05	-0.03	-0.02	-0.03	-0.11	-0.03	-0.03	-0.20
<b>NEW GUIDANCE sub-total</b>	<b>0.85</b>	<b>1.46</b>	<b>0.86</b>	<b>0.64</b>	<b>0.87</b>	<b>3.49</b>	<b>0.87</b>	<b>0.87</b>	<b>6.60</b>
Carry-over	0.13	-0.11	0.07	0.15	0.04	0.73	0.00	0.00	0.00
<b>TOTAL FUNDING</b>	<b>0.98</b>	<b>1.34</b>	<b>0.93</b>	<b>0.79</b>	<b>0.91</b>	<b>4.23</b>	<b>0.87</b>	<b>0.87</b>	<b>6.60</b>
<b>COSTS</b>									
Capital expenses	1.08	1.21	0.81	0.75	0.18				
<b>TOTAL CAPITAL EXPENSES</b>	<b>1.09</b>	<b>1.27</b>	<b>0.79</b>	<b>0.75</b>	<b>0.18</b>	<b>4.23</b>	<b>0.87</b>	<b>0.87</b>	<b>6.57</b>
<b>net balance</b>	<b>-0.11</b>	<b>0.07</b>	<b>0.15</b>	<b>0.04</b>	<b>0.73</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.03</b>
<b>NET OPERATING &amp; CAPITAL AVAILABLE FUNDS</b>	15.4	17.4	18.9	20.1	19.8	24.52	20.76	20.98	27.11
<b>GUIDANCE</b>	15.2	15.4	17.8	18.8	17.3	16.79	17.87	20.97	27.10
<b>COSTS</b>	13.4	16.3	17.6	17.6	12.1	21.63	20.76	20.96	26.66



# ESnet Annual Program Review August 3, 2004

William E. Johnston, ESnet Dept. Head and Senior Scientist

R. P. Singh, Federal Project Manager

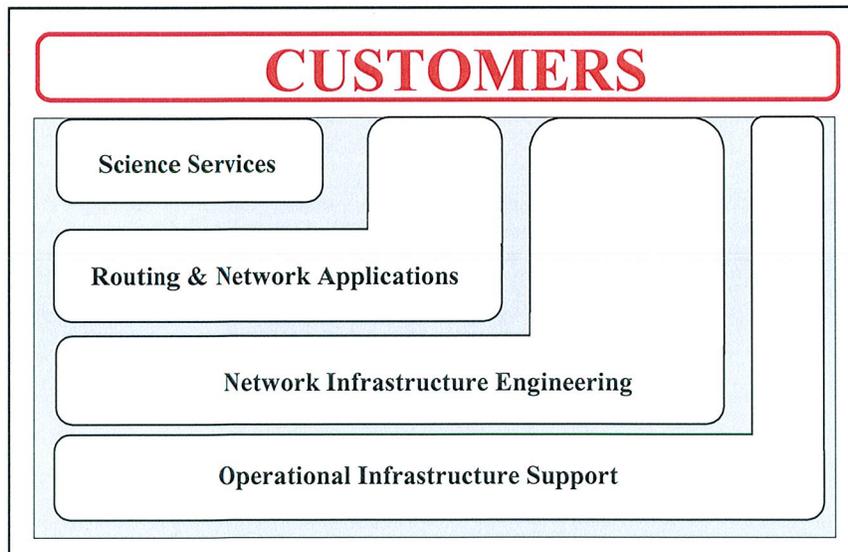
Michael S. Collins, Stan Kluz,

Joseph Burrestia, and James V. Gagliardi, ESnet Leads  
and the ESnet Team

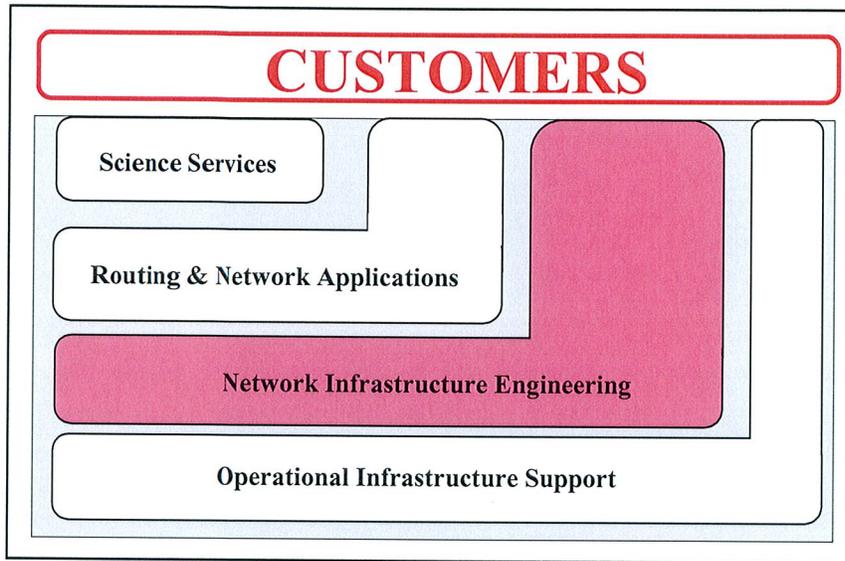
Lawrence Berkeley National Laboratory



## ESnet's Role in Support of Science



# ESnet's Role in Support of Science



MANs - Bill



# ESnet Backbone Upgrades ESnet Annual Program Review

August 3, 2004

James Gagliardi  
ESnet Network Technical Services Group  
Lawrence Berkeley National Lab  
jvg@es.net

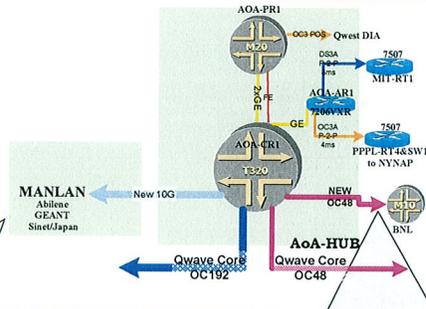


## Backbone Upgrades

We will be upgrading network services at these ESnet facilities:

- **AoA (NYC) Hub, MANLAN, and BNL**
  - This will increase ESnet University and International connectivity to 10Gb/s and increase BNL bandwidth from OC12 to OC48.
- **Chicago Hub, CERN, ORNL & Ames Lab**
  - This is the first step of getting higher speed access to CERN and Starlight and also moves ORNL to OC192.
- **Sunnyvale Hub**
  - This is in preparation for Phase 1 of new Architecture
- **Albuquerque Hub, LANL and SNLA**
  - This will allow the retirement of regional ATM equipment and prepare for a future NM GIGAPOP connection.
- **Atlanta Hub, ORNL**
  - This upgrades the current ORNL OC12 to OC48 with a reduction in cost.
- **ESnet site equipment**
  - This is a constant process of retiring, upgrading and enhancing ESnet equipment located at various sites.

# AoA (NYC) Hub Upgrades



## MAN LAN (The Manhattan Landing)

Purpose: To increase our University and International connectivity to 10Gb/s

- Moving from private peering's with Abilene @ GE, GEANT @ OC48 & SINET@ GE to 10G VLANS
- The Contract and 10G interface order is in progress
- Hardware cost: ~\$102,000
- MRC ~ \$1800
- Production 4Q04 – 1Q05

## BNL Upgrade

Purpose: To increase the access from OC12 to OC48 (for a cost increase of ~\$2,200.)

- Hardware cost: ~\$50,000
- MRC: Qwest circuit: ~\$26,600.00
- Qwest circuit & interface order to be placed
- Production 1 - 2Q05

# Chicago Upgrades

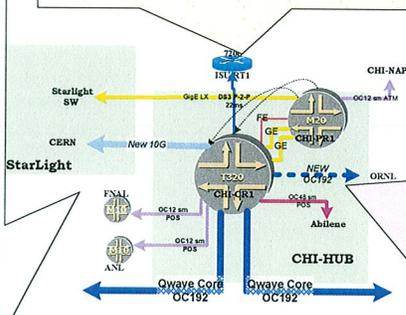
## AMES LAB Circuit Move

- Purpose: It will finish the security architecture for this HUB
- Move Ames Lab DS3 circuit from Peer M20 to Core T320
  - Hardware cost: 4port DS2 interface ~\$15,000
  - Production 4Q04

## CHI-HUB & STARLIGHT

Purpose: The 1st step of getting high speed access to CERN's router at Starlight. It moves CERN from the 1Gb/s vlan to a private 10Gb/s link

- Acquired fiber pair from (IWIRE) for an interim NBC to CERN circuit
- Qwest cross-connect and 10G interface orders are now in progress
- Hardware cost: 10G ~\$102,000
- MRC: \$500
- Production 4Q04



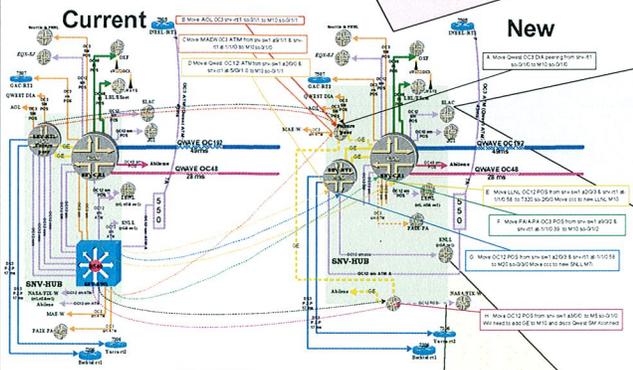
## ORNL's OC192

Purpose: ORNL funded OC192 link that moves it's access to the ESnet's 10G Northern Backbone

- ORNL is will supply OC192 card, ESnet will provide a port on the Chicago T320
- Routing changes needed to move ORNL primary access.
- Production 1 - 2Q05 ?

# Sunnyvale Upgrades

Very busy drawing intended to show the extent of the changes needed.



## SNV-HUB (1400 Kifer)

*Purpose: Clean-up of the hub's equipment & circuits for the future Bay Area MAN & Phase 1 of the New Backbone Architecture.*

- Move several (6) circuit from the ATM sw to the routers.
- Remove the switch
- Split the function of the Core Router between T320 & M20 and use the M20 for low speed circuits
- Add additional routers to serve as a Peer and shared Abilene router for access to Nasa Ames (FIX, MIX & NGIX peering)
- Hardware cost : \$2,800
- MRC - \$500
- Production 4Q04

## SNLL Upgrade

*Purpose: To combine the SNLL SW and router into one unit. This will allow the SecureNet hand-off to move from Sunnyvale to SNLL.*

- Juniper M7i router now on order
- Hardware cost : \$31,000
- Production 4Q04

# Albuquerque Upgrades

## LANL Changes

*Purpose: LANL will be moving its site peering to the Qwest facility in Alb.*

- Install GE-LX interfaces in ALB M20.
- Test and roll the LANL GE into production
- Ship M10 back to LBL
- Hardware cost: \$0 we have interfaces
- Production 4Q04

## ALB-HUB Changes

*Purpose: To move the LANL and SNLA circuits from the ATM switch to the Router.*

- Move the LANL (depending on there upgrade) & SNLA OC12 from the ATM SW to the M20.
- Ship ATM Switch back to LBL
- Production 4Q04

## SNLA Upgrade

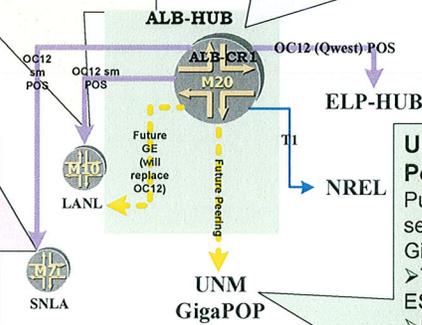
*Purpose: To combine the SNLA SW and router into one unit. This will allow the SecureNet hand-off to move from ALB to SNLA.*

- Juniper M7i router now on order
- Hardware cost : \$31,000
- Production 4Q04

## Univ. of NM Future Peering

*Purpose: UNM will be setting up a new GigaPOP.*

- This will be a future ESnet peering point.
- Peering interface and possible new router will be needed
- Production 1 - 2Q05

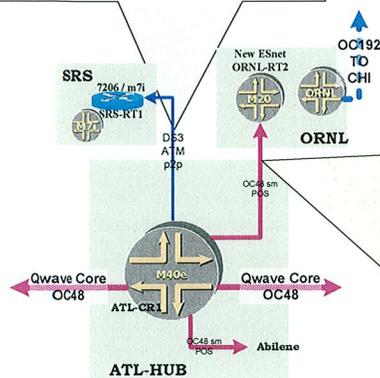


# Atlanta Hub & ORNL Upgrades

## Atlanta Hub Changes

Purpose: SecureNet moving to the IP encryptors the ATM switch serving SRS and ORNL's Y12 MPLS ATM paths can be removed.

- Move SRS DS3 Circuit from SW to Router port
- Production 4Q04 – 1Q05



## ORNL OC48 upgrade

Purpose: The upgrade from OC12 to OC48 comes with a reduction in cost over the existing OC12. It allows ORNL and their connected site high-speed access to the Southern 2.5G Backbone.

- OC48 service order sent to Qwest
- Upgrade the ORNL M10 to a M20 (now being configured at LBL)
- OC48 interfaces for ATL and ORNL on order
- Hardware cost : ~\$98,000
- MRC - \$19,400
- Production 4Q05

# General Sites Upgrades

## • Site Upgrades

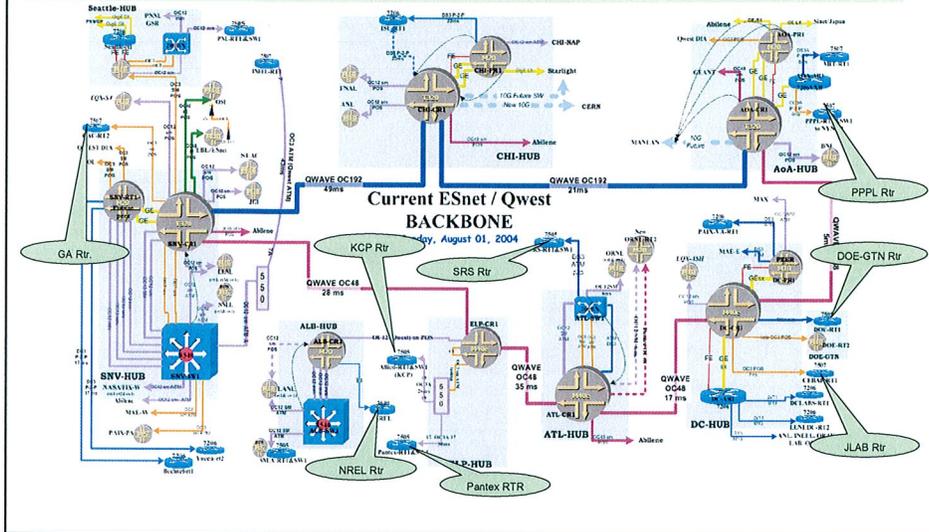
Purpose: With the constant change and upgrades going on in the network there are many hours spent "behind the scenes". This is to testing hardware and new vendor code prior to shipping equipment and reloading routers at the remote sites.

Some of the upcoming changes are:

- Retiring older hardware
  - ATM Switches that have been used for the SecureNet ATM Fastlane support (soon to be replaced by the IP Taclane)
  - "End of Live" routers and interfaces that will no longer be supported
- Enhancing the reliability of existing routers
  - In the process of replacing all the single power supply routers (Cisco 7505, 3640) with dual supply system like Cisco 7507, 7206 or Juniper M71
  - Will be evaluating the new Cisco RSP16 routing engine to upgrade and in some cases allow dual processors on the 7507 platform
  - Will be evaluating upgrades for the Cisco 7206 to expand their port density and performance.
- And always the New Cisco (IOS) and Juniper (Junos) router code upgrades and bug fixes
  - Normally needed 3 – 6 months UNLESS there are:
    - ✓ Security fixes that need immediate upgrades (happened with both vendors in the last 6 months)
    - ✓ New bug fixes like the multicast problem back in April with the Juniper code
    - ✓ New enhancement that are will enhance the network

# General Sites Upgrades

The routes highlighted below, are at sites pending upgrades that were not included in the previous slides. The changes will be for new line upgrades and router enhancements (dual power supplies, new processors, etc.)



## Future ESnet Backbone Architecture Phase 1

We will be upgrading network services at the current and new ESnet facilities to allow for the new backbone enhancements:

- **Chicago Hubs: Qwest CHI-HUB and Starlight**
  - This is phase 1 of getting higher speed access to the Chicago sites, a back-up 10G Backbone via NLR along with the moving to the proposed Chicago Area MAN
- **Sunnyvale Hubs: Qwest SNV-HUB and Level3**
  - This is a vital component of the MAN for high speed and redundant access to the Bay Area sites as well as the Phase 1 of the new backbone architecture
- **Seattle HUB: PNWGigaPop**
  - The Phase 1 step of the new NRL Backbone to provide high-speed access to PNNL, PNWGigaPop, Pacific Wave and International Peers on the West Coast.

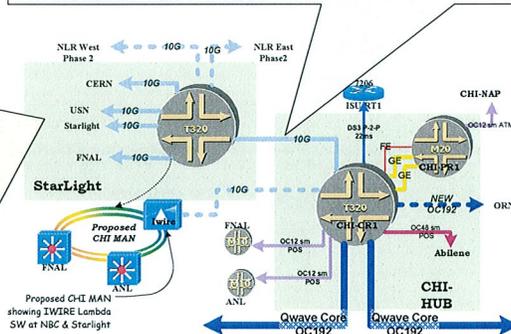
# Chicago Future Backbone Architecture Phase1

## Starlight

- Purpose: The installation of a ESnet T320 router will allow direct 10G connections to CERN, FNAL, the Starlight Peers, UltraScienceNet, NRL and the ESnet Qwest Backbone.
- Contract for collocation & 10G port at Starlight now in process. (.5 rack now full rack Dec. 2004)
  - New T320 order to be placed
  - Additional 10G port's for the Starlight SW being ordered by IWIRE
  - Hardware cost :
    - Starlight T320 ~\$870,000
    - Starlight Port - ~\$54,000
  - MRC - Starlight - TBD
  - Production 1-2Q05

## Qwest Hub

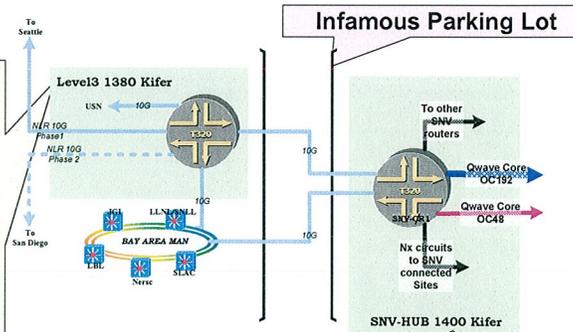
- Purpose: 10G interconnect between the ESnet T320 at Qwest Hub and Starlight
- Move the fiber from the 1G to the 10G interfaces
  - Move CERN peering from CHI-CR1 to Starlight T320
  - Future 10G interface needed for redundant routes (once CHI MAN is funded)
- Hardware cost : CHI 10G ~\$102,000
- MRC - Qwest - \$1000.
  - Production 4Q04 - 1Q05



# SNV Future Backbone Architecture Phase1

## Level3 Hub (1380 Kifer)

- Purpose: The installation of the ESnet T320 router will allow direct 10G connections to the Bay Area Man, UltraScienceNet, and the ESnet Qwest Backbone. As well as the NRL 10G Path to Seattle.
- Collocation & cross-connect contract now in process.
  - New T320 Router order to be placed
  - Hardware cost : T320 ~\$702,000
  - MRC -TBD - (Level3)
  - Production 1-2Q05



## Qwest Hub (1400 Kifer)

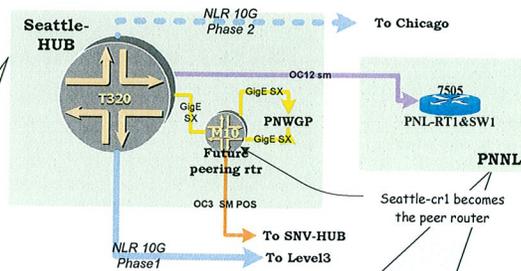
- Purpose: 10G interconnect between the ESnet T320's at the Qwest Hub and Level3
- Additional 10G and support interface to be ordered
  - Install and test the interfaces in the SNV-CR1 T320
  - Hardware cost : SNV 10G & 10x1G ~\$268,000
  - MRC \$2,500-(Qwest)
  - Production 1-2Q05

## Seattle Future Backbone Architecture Phase 1

### Seattle (PNWGP)

Purpose: The installation of the ESnet T320 router will allow 10G connections to the NLR and high-speed access to PNNL, PNWGigaPop, Pacific Wave and International Peers.

- Contract needed with NLR
- New T320 order to be placed
- Modification of the our existing collocation contract with Univ. of Washington for power & space upgrades
- Hardware cost :  
T320 ~\$448,000
- MRC -TBD - (NLR)  
TBD - (Univ. Of Wash.)
- Production 2Q05



### PNNL Upgrade

Purpose: Once the upgrades via NLR are in place the ESnet access to PNNL can accommodate the planned site access 2.5G upgrade.

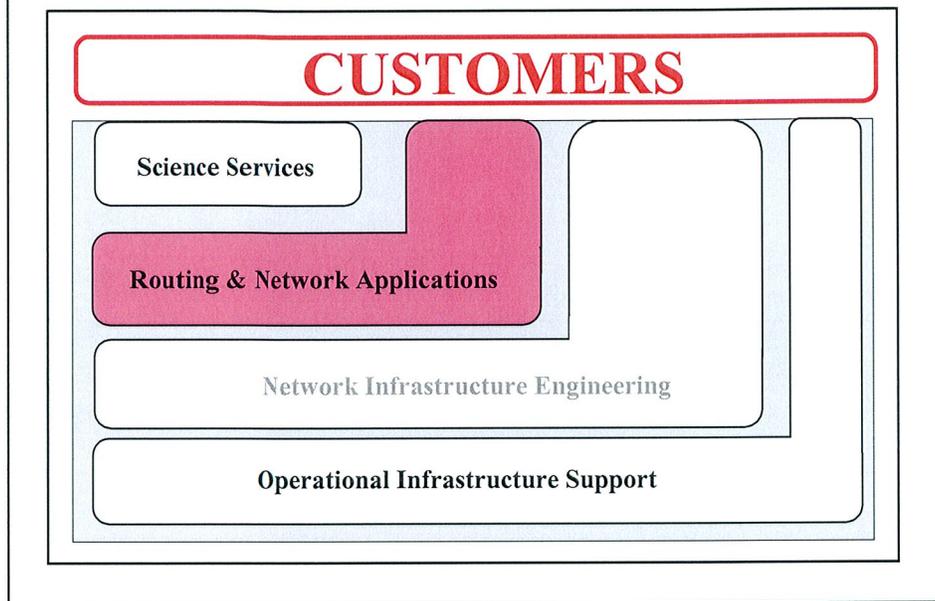
- Upgrade the ESnet PNNL router
- Hardware cost : TBD (~\$75,000)
- Production 2 -3Q05 ?

## Summary

In conclusion, these changes will prepare ESnet for:

- High speed connections to sites and international peers.
- Redundant high-speed site connections via the MAN architecture.
- A high-speed secondary Backbone using NLR.

# ESnet's Role in Support of Science



## ESnet Routing and ISP Services ESnet Annual Program Review

August 3, 2004

Michael Collins  
ESnet Network Engineering Group  
Lawrence Berkeley National Lab  
collins@es.net



## How Do Networks Work?

---

- When one types “google.com” into a Web browser to use the search engine, the following takes place
  - The name “google.com” is resolved to an Internet address by the Domain Name System (DNS) – a hierarchical directory service
  - The address is attached to a network packet (which carries the data – a google search request in this case) which is then sent out of the computer into the network
  - The first place that the packet reaches is a router that must decide how to get that packet to its destination (google.com)

## How Do Networks Work?

---

- Routers are in your site LANs and at your ISP, and each router typically communicates directly with several other routers
- The first router to receive your packet takes a quick look at the address and says, if I send this packet to router B that will probably take it closer to its destination. So it sends it to B without further adieu.
- Router B does the same thing, and so forth, until the packet reaches google.com
- What makes this work is routing protocols that exchange reachability information between all directly connected routers – “BGP” is the most common such protocol in WANs



## What happens when ISPs Peer?

- ISPs trade routing information at exchange points. The information exchanged is the networks that they can reach. The routers sort out the information to determine the best route to a particular prefix (LAN).
- The Border Gateway Protocol (BGP) is the protocol used to exchange routes between peers.
- ESnet daily peering report (top 20 of about 100). Daily churn.
- The Top 20 represent 360,000 prefixes. So, there are multiple announcements of each prefix that must be sorted out.
- ESnet has 167 Peering Sessions with 91 peers.
- ESnet routing management can be labor intensive.
- **If BGP breaks Science Stops**

Current Date: Sun, 1 Aug 2004 06:10:44 -0700 (PDT)				
Previous Date: Sat, 31 Jul 2004 06:05:44 -0700 (PDT)				
AS	Prefixes	Change	%	Name
1239	63747	34	0.1	SPRINTLINK
701	51764	31	0.1	ALTERNET-AS
3356	48087	47	0.1	LEVEL3
209	38517	-518	-1.3	QWEST
3561	38170	-2130	-5.3	CWUSA
7018	28392	11	0	ATT-INTERNET4
3549	19319	-8	0	GBLX
2914	15553	78	0.5	VERIO
1299	10316	83	0.8	TELIANET
5511	6350	69	1.1	OPENTRANSIT
6461	5896	76	1.3	ABOVENET
174	5572	11	0.2	COGENT-PSI-1
7473	5022	17	0.3	SINGTEL
3491	4425	13	0.3	CAIS-ASN
11537	3512	-2	-0.1	ABILENE
5400	3460	5	0.1	BT
2828	2978	-3	-0.1	XO-AS15
4323	2878	-2	-0.1	TWTC
6395	2776	-14	-0.5	BROADWING
4200	2520	4	0.2	ALERON-4200

## ESnet Routing Management (1)

- ESnet applies prefix filters to all BGP sessions
  - If an ISP announced the entire set of Internet routes, instead of just the routes it was supposed to announce, the Internet can break.
    - This has happened.
    - ESnet was protected by its filters.
  - ESnet constructs its filters using in-house scripts by using-
    - The Internet Routing Registries
      - All ISPs \*should\*, in our opinion, register the routes they are responsible for with an Internet Routing Registry. Many do not. ESnet does.
    - What ESnet learned yesterday with a check on the number of new prefixes.

If BGP Breaks – Science Stops

## ESnet Routing Management (2)

- Some ISPs have too many prefixes to filter.
  - ESnet applies a well-known bad route filter on those sessions.
  - They are usually well run networks.
- The filters are updated 2-3 times a week.
- What happens if Abilene and a commercial ISP announce the same prefix?
  - We manually configure the peering session so that the Abilene announcement is preferred over any others.
- What happens if Abilene announces CERN to ESnet?
  - We have to make sure that the direct CERN announcement is preferred over the Abilene announcement
- ESnet's routing is manually configured with preferences set for each peering session.
  - All of this requires planning and great attention to detail!

If BGP Breaks – Science Stops

## ESnet External Peering Policy

- If, because of a trouble call or direct observation, we think a better path to an external site exists, we try to setup the appropriate BGP session.
- If another entity requests to peer with ESnet, we generally agree to set up a session.
  - We are the Good Guys.

## Accepting A New Prefix from a Site

- The acceptance of a new prefix is a formal process requiring the involvement of the Site Coordinator.
- The protocol used between the Site and ESnet is generally BGP.
- ESnet filters the BGP session with the site.

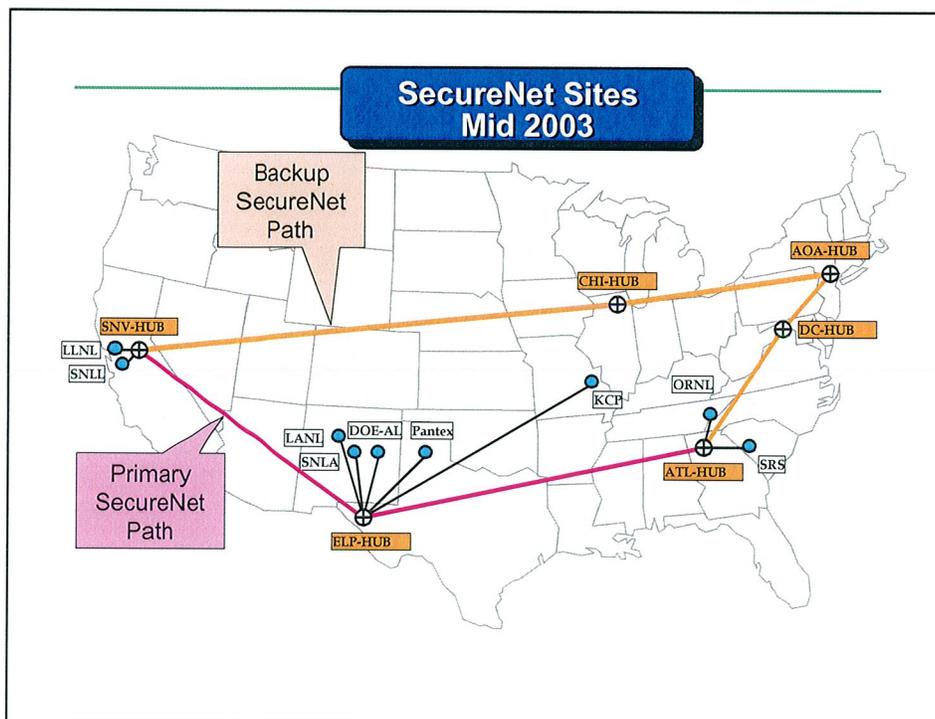
## ESnet Routes Native IPv6 Traffic

- IPv6 traffic uses the same physical interfaces and routers as IPV4 traffic.
  - A parallel infrastructure is not used.
- ESnet manages the distributed 6TAP, the first international IPv6 exchange point router.
  - Distributed between StarTap, StarLight and PAIX-PA.
  - ESnet Peers with the 19 6TAP members.
- ESnet also has direct IPV6 peering with Abilene, CERN and GEANT.
- Two sites are using IPv6. Two other sites are ready to come up.
- IPv6 usage is minimal.



## ESnet Uses MPLS Paths for SecureNet

- MPLS paths are used to route encapsulated and tagged SecureNet ATM cells through the IP backbone.
  - Packets are received over an ATM interface encapsulated in a Juniper proprietary protocol and then switched onto an MPLS path that terminates in a router close to the SecureNet site. The ATM cells are then stripped of the tag and encapsulation and switched to an ATM interface for ATM transport to the SecureNet destination host.
  - Emulates a fully meshed set of ATM PVPs between SecureNet sites.
- RSVP is used to dynamically set up the MPLS paths and to allow fail over.
- Most SecureNet sites are reached via the Southern route. If the Southern route goes down, the paths will be rerouted over the Northern route.
- All of the intercore router links are MPLS enabled.



## Who Owns The IP Address Space

---

- Technically the Internet Assigned Numbers Authority (IANA) “owns” the IP address space.
- As necessary, IANA “lends” blocks of address space to Regional Registries for further allocation to ISPs in their area.
- The American Registry of Internet Numbers (ARIN) “lends” portions of the addresses obtained from IANA to ISPs in North America.
- ESnet obtained its IPv4 CIDR and IPv6 address blocks from ARIN.
  - ESnet also has a /16 (Class B network) and several /24s (class C network).

## ESnet Assigns Address Blocks to ESnet Sites

---

- Upon request, ESnet can augment a site’s IPv4 address space by assigning blocks from the ESnet CIDR allocation.
- The process is formal. It requires the participation of the Site Coordinator and a justification for the size of the request.
- A note outlining the conditions of the assignment includes:

The following IPv4 CIDR block has been assigned to TheLab. By accepting and using this IPv4 CIDR block assignment TheLab agrees to the following:

TheLab will not announce the CIDR block to any other ISP.

The CIDR block assignment is "on loan" to TheLab.

TheLab will return the CIDR block assignment to ESnet if they decide to change their primary carrier from ESnet to another ISP.

Traffic from nodes using this prefix will meet ESnet's Acceptable Use Policy  
<http://www.es.net/hypertext/esnet-aup.html>

## Assigned Addresses Are SWIPed

- SWIPing is the process of registering contact information for the assignment of blocks from a larger CIDR block.
- We are in the process of completing the SWIP process for all of the CIDR assignments to sites.
  - Out of 79 blocks needing to be SWIPed, 50 are completed.
  - Some sites require a lot of assistance.
- If there are problems associated with the network, the complaint goes directly to the site.

## Sample Problem

Dear Mike Collins:

I am an authorized representative of the **Entertainment Software Association ("ESA")**, which represents the intellectual property interests of twenty-four (24) companies that publish interactive games for video game consoles, personal computers, handheld devices and the Internet. ESA is providing this letter of notification to make Energy Sciences Network aware of material available via its network or system that **infringes the exclusive copyright and trademark rights of one or more ESA members.**

This notice is addressed to you as an agent of Energy Sciences Network for purposes of receiving notifications of claimed infringement. We hereby affirm that the ESA is authorized to act on behalf of the ESA members whose exclusive copyright rights we believe to be infringed as described herein.

ESA has a good faith belief that the Internet site found at **198.125.1.11** continues to infringe the rights of one or more ESA members by offering for download one or more unauthorized copies of one or more game products protected by copyright, including, but not limited to:

Hoyle

- Robert L. Hunter, IV
- Entertainment Software Association
- 1211 Connecticut Avenue, N.W.
- Washington, DC 20036 USA
- Reference: Case ID 3293623 - Notice of Claimed Infringement
- Mr. Hunter:
- We are in receipt of your e-mail Notices of Claimed Infringement sent to ESnet on June 17, 2004 and July 6, 2004. ESnet is a private enterprise network supporting the scientific research mission of the US Dept of Energy. ESnet is managed by the E.O. Lawrence Berkeley National Laboratory (LBNL). LBNL adheres to the laws and regulations governing network use and copyright infringement. In your e-mail you identified the address 198.125.1.11 as the source of the infringement. We have investigated your claim. While **198.125.1.11 is in the network address space allocated to ESnet by ARIN (American Registry for Internet Numbers, which manages all Internet address space allocation in North America), it is from an unused region in the ESnet Classless Inter-Domain Routing (CIDR) block. There is no device associated with that address and ESnet does not route Internet traffic to or from that address. We can only conclude that the address has been stolen and used without our permission.** Further, because ESnet does not route traffic to or from unused portions of its address space, nothing can be done by ESnet to control traffic originating there. This situation is not unique to ESnet, but is the norm for Internet Service Providers.
- Regrettably, we cannot restrict or control the activity of those, unknown to us, who have stolen the address 198.125.1.11.
- Sincerely,
- A.X. Merola,
- Director, Information Technologies Division
- Chief Information Officer, LBNL
- CC: G. Woods, Laboratory Counsel
- W. Johnston, Head, ESnet

The Secret Society  
-OR-  
If I continue, you may never leave the room

**Message from the Society**

Message from external security group to ESnet member:

Hi, team.

Thanks to a combination of donated flow data and Matt's excellent narcbot, we've spotted bot. This botnet resides on 211.50.211.216 TCP 6667 AS3786 (ERX-DACOMNET DACOM Corporation)

Server IP: 211.50.211.216

Server AS: 3786 (ERX-DACOMNET DACOM Corporation)

Server Name: www4

We don't know the malware or all of the channels used for these botnets. The botnet server has 679 clients presently. The data was verified as of 2004-07-22 16:35:11 GMT.

The 211.50.211.216/32 prefix is being advertised by all of the DDoS route-servers.

**Unrelated message to a site**

To: security@\*\*\*\*.gov

Subject: Possible infected system at \*\*\*\*

Date: Mon, 26 Jul 2004 11:06:14 -0700

From: "Kevin Oberman" <oberman@ptavv>

Good day!

ESnet has received a report that dhcpvisitor217231.\*\*\*\*.gov has been scanning port 445 on system across the Internet.

The traffic was logged back on July 21, so it is likely that this has already been taken care of, but I wanted to pass it along, just to be sure.

Thanks

ESnet Manages it's IP address Space

- A commercial product, namesurfer, is used to maintain IPV6 and IPv4 address assignments.
  - Namesurfer replaced our ancient, home grown management scripts.
- We manage
  - Host address assignments for the ESnet LANS at LBNL, TWC, AmesLab and BNL.
  - Interface assignments for the ESnet backbone and associated equipment.
  - CIDR block assignments to the ESnet sites.

## Managing the Configuration Files of Routers and Switches

---

- The Really Awesome New Cisco Config Differ (RANCID - Opensource) is used to retrieve configuration files from routers and switches after a configuration change.
  - RANCID replaces our ancient and home grown scripts.
- An entry is appended to a log file stating the reason for the configuration.
- The configuration is stored in CVS
- The change from the previous configuration is e-mailed to the OCS Team and the NESG.
- 2000+ e-mail messages were sent between 7/1/2003 and 6/30/2004.
- There are other tools to upload configurations into the router or switch.

## Network Applications

---

- How do we track how much data is being transported and where is it coming from and going to?
  - Usage monitoring applications
- How do we insure ESnet is providing optimal network service?
  - Network performance monitoring
- How can we assist end users troubleshooting performance problems?
  - Performance Center testing
- How do we schedule, track and improve uptime due to outages?
  - Outage monitoring and reporting

**We are increasing the transparency of ESnet by making all of these applications and associated collected data accessible to our user community.**



# ESnet Usage Monitoring Systems ESnet Annual Program Review

August 3, 2004

Chin Guok

ESnet Network Engineering Group

Lawrence Berkeley National Lab

chin@es.net



## ESnet Usage Monitoring Systems

---

How does ESnet monitor traffic utilization?

Network Traffic Monitoring System (NTMS)

How does ESnet track where traffic is going to and coming from?

NetFlow Monitoring System (NFMS)

## Purpose Of The Network Traffic Monitoring System (NTMS)

---

### Motivation:

- To monitor traffic utilization on ESnet links.

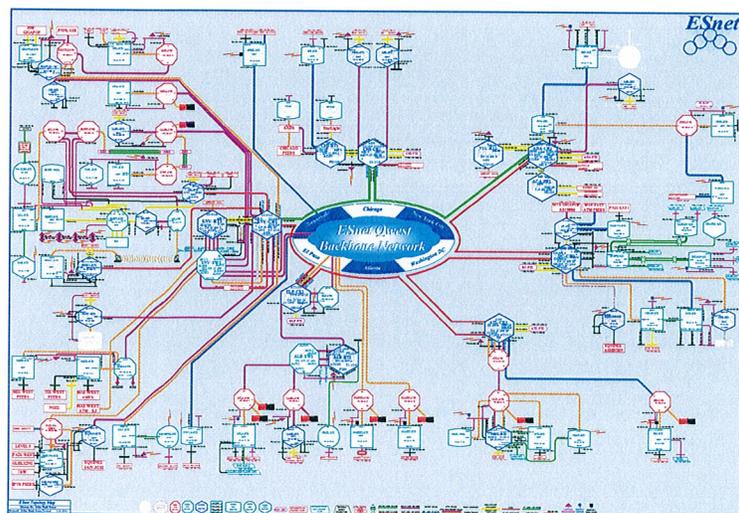
### Objective:

- Generate individual and aggregated graphs of traffic utilization.
- View historical data in detail.

## How Much Data Does NTMS Collect

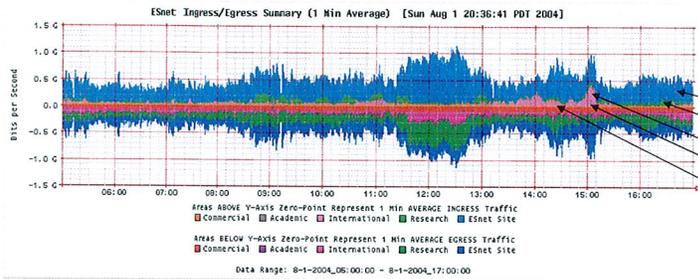
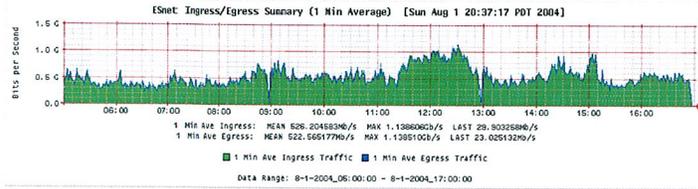
---

Traffic, error and discard counts on every router and switch interface are collected every 60 sec.



## How Does NTMS Display The Data (1/2)

Total volume of ingress and egress traffic flowing through ESnet

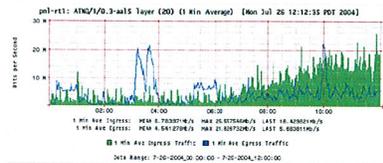


Traffic Utilization Broken Down By Categories:

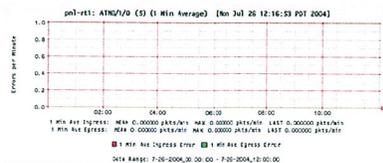
- ESnet Site
- Research
- International
- Academic
- Commercial

## How Does NTMS Display The Data (2/2)

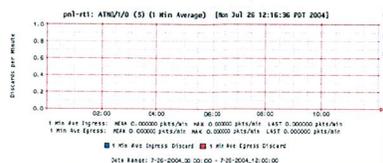
Traffic, error and discard graphs for a specific router interface.



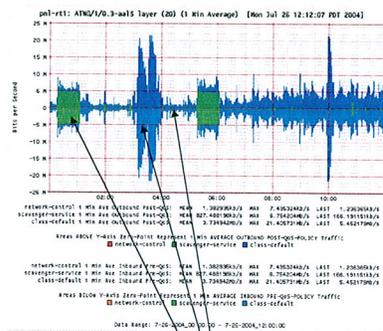
Traffic Utilization Graph



Error Count Graph



Discard Count Graph



Traffic Utilization Broken Down By Class-of-Service:

- Network-Control
- Best-Effort
- Scavenger-Service

## NTMS Milestones

---

Prototype in place.	3Q04
Increase current capability to view 16 hours of detailed data to 1 year.	1Q05
Import historical data into new system.	4Q05

## Purpose Of The NetFlow Monitoring System (NFMS)

---

### **Motivation:**

- To track down where traffic is coming from and going to.
  1. For planning purposes.
  2. For troubleshooting (e.g Denial of Service attacks).

### **Background:**

- NetFlow data is collected from about 50 ESnet border routers.
- About 15GB of NetFlow data is collected daily, which results in 3.5GB of compressed data being stored.

### **Objectives:**

- Identify top flows.
- Visualize top flows and their evolution.
- Query information for specific flows.

## How Does NFMS Display The Data

Report of the top 10 largest traffic exchange between a pair of Autonomous Systems (AS) (or network domains)

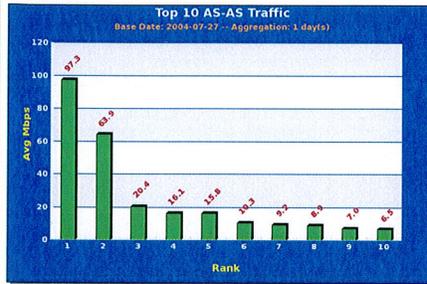
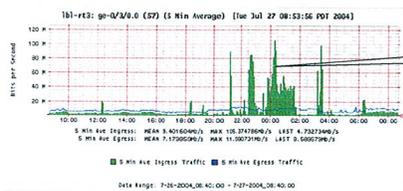


TABLE Base Date = 2004-07-27 -- Aggregation = Previous 1 day(s)

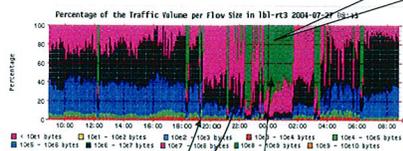
Rank	Router	Source	Destination	Avg Mbps	MBytes	Historic
1	aoa-cr1	SLAC (3671)	IN2P3 (789)	97.279	1050616	<a href="#">Week</a>
2	aoa-cr1	SLAC (3671)	ASGARR (137)	63.942	690577	<a href="#">Week</a>
3	aoa-cr1	FERMILAB (3152)	IN2P3 (789)	20.366	219957	<a href="#">Week</a>
4	chi-cr1	FERMILAB (3152)	BUFFALO-ASN (3685)	16.054	173380	<a href="#">Week</a>
5	aoa-cr1	SLAC (3671)	DFN-WIN-AS (680)	15.815	170805	<a href="#">Week</a>
6	aoa-cr1	FERMILAB (3152)	IANET (786)	10.269	110907	<a href="#">Week</a>
7	fnal-cr1	BNL-AS (43)	FERMILAB (3152)	9.152	98846	<a href="#">Week</a>
8	nl-cr1	BNL-AS (43)	MARLBOROUGH (18515)	8.864	97728	<a href="#">Week</a>
9	chi-cr1	BNL-AS (43)	U-CHICAGO-AS (160)	6.999	75584	<a href="#">Week</a>
10	slac-cr4	ASGARR (137)	SLAC (3671)	6.541	70644	<a href="#">Week</a>

## What Can NFMS Tell Us

Querying information for specific flows



Sharp increase in traffic volume between 12:00 am – 1:00 am.



Flow distribution shows that majority of traffic is comprised of a small number of large flows.

IP address indicates system in OSF used to store nightly backups.

KiloByte (10<sup>3</sup>) Flows  
 MegaByte (10<sup>6</sup>) Flows  
 Hundred MegaByte (10<sup>8</sup>) Flows

```

----- Report information -----
#
# Name:      Source/Destination IP
# Args:     flow-stat #10 #33
# Period:   2004-07-27 00:00 --> 2004-07-27 01:00
#
# src IPaddr  dst IPaddr  flows  octets  packets
#
198.128.3.23  128.55.128.84  136   15645837200  11044400
198.128.3.22  128.55.128.84   45    7056829500   4970600
198.128.1.136 128.55.128.84   27    1562002800   1041700
198.124.224.3 198.128.3.22   100   1172628000   781800
207.225.159.1 198.128.3.170  120    1151593200   1204700
...
    
```

## NFMS Milestones

---

Prototype in place.	3Q04
Analyze all data simultaneous by using a computing cluster platform	1Q05



# NETINFO

## ESnet Annual Program Review

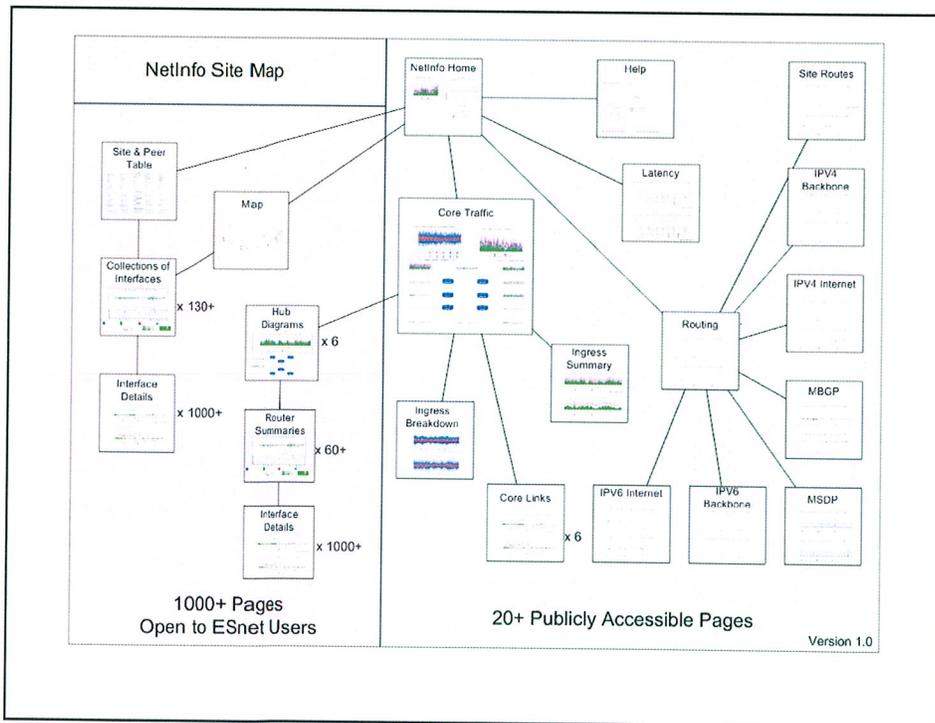
August 3, 2004

Joseph Metzger  
ESnet Network Engineering Group  
Lawrence Berkeley National Lab  
metzger@es.net

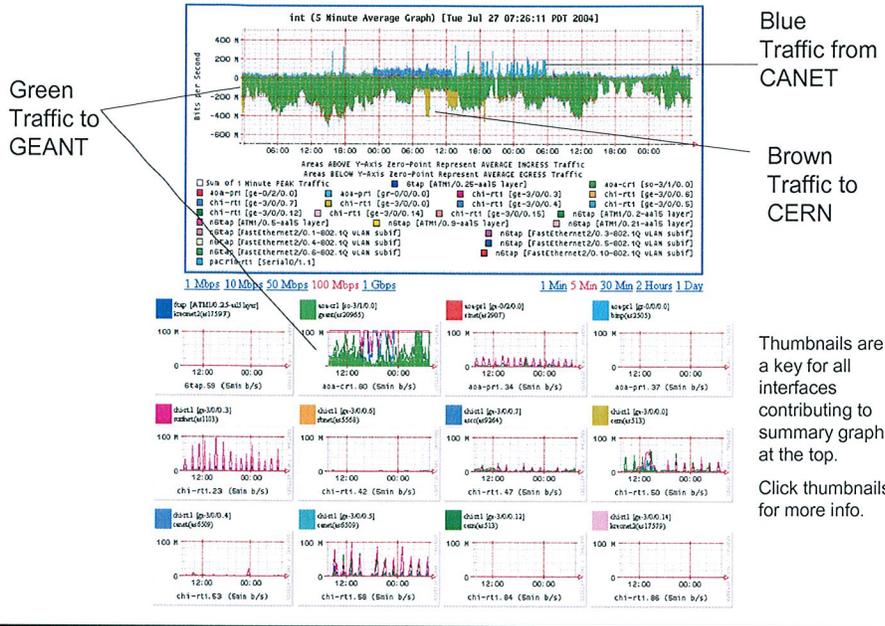


# Publishing Network Statistics

- We need to make it easy for ESnet users to find information about the portion of the network that affects them so they can debug and tune their applications.
- **NetInfo**: a web portal publishing a subset of existing ESnet statistics in an intuitive and useful manner.



## International Peers link from Connections Page



## Publishing Network Statistics Status

NetInfo Release to Beta Users	3Q04
NetInfo Release to all ESnet Users	4Q04
Add Owamp statistics to NetInfo	2Q05
Add MAN statistics to NetInfo	After Mans Deployed



## Monitoring DOE Lab - University Connectivity

ESnet Annual Program Review

August 3, 2004

Joseph Metzger

ESnet Network Engineering Group

Lawrence Berkeley National Lab

metzger@es.net



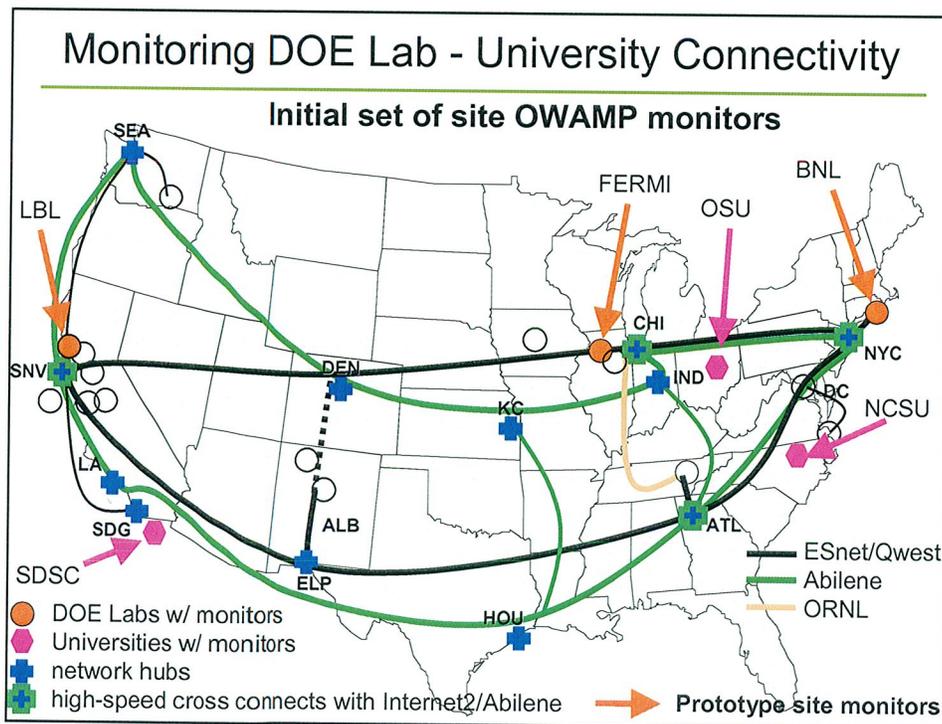
## Monitoring DOE Lab - University Connectivity

---

- We want to ensure that connectivity from any lab to any university is equivalent to lab to lab or university to university connectivity.
  - Installed High Speed cross connects in Sunnyvale, Chicago, New York and Atlanta
  - Setup Monitoring system to verify performance using latency testing.

## Latency Measurement with OWAMP

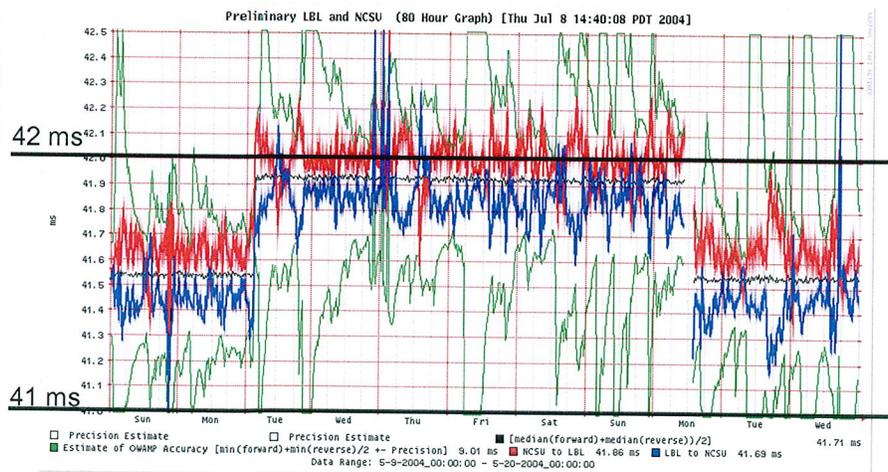
- We are using OWAMP to measure one way delays between test stations.
  - OWAMP One Way Measurement Protocol
  - OWAMP is a protocol documented in an IETF draft
    - <http://www.ietf.org/internet-drafts/draft-ietf-ippm-owdp-08.txt>
  - OWAMP is a set of tools developed by Internet2



## Lab University Connectivity Results

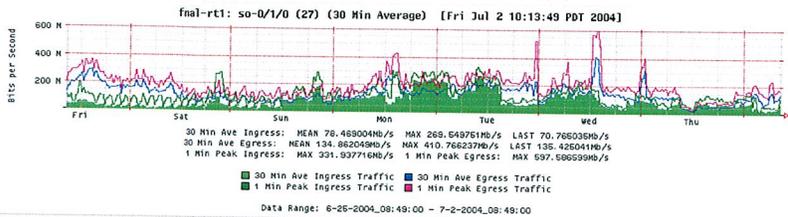
- We are not seeing any congestion or queuing on the high speed cross connects.
- Interesting issues with site access circuits are showing up.
  - North Carolina State University metro fiber reroute.
  - Congestion on FERMI access link.

## ESnet Visualization of OWAMP Data Showing NCSU's MAN Fiber Reroute

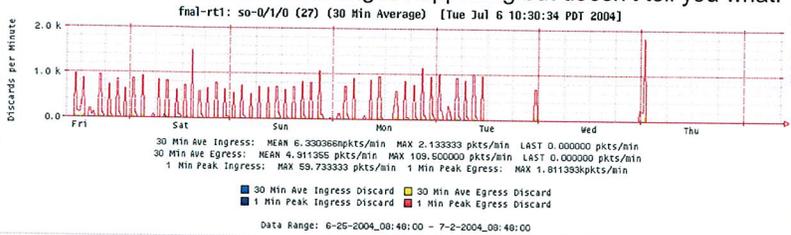


## Why multiple measurement systems? They show different aspects of performance.

Interface traffic counters do not show any problems.



Discard data shows something is happening but doesn't tell you what.

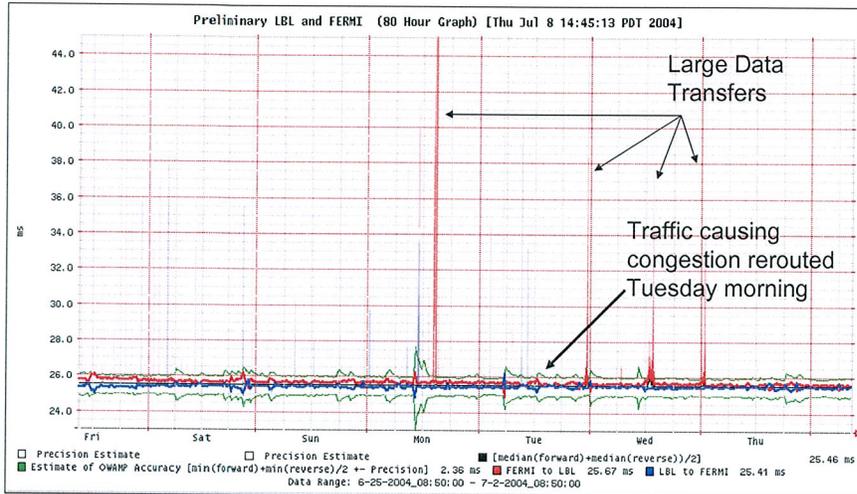


## NetFlow Analysis Identified Potential Source

```
# ----- Report Information -----
#
# Fields: Total
# Symbols: Disabled
# Sorting: Descending Field 3
# Name: Source/Destination IP
#
# Args: flow-stat -f10 -s3
#
#
# src IPaddr    dst IPaddr    flows    octets    packets    originating file
#
129.105.21.229 198.49.208.10 193      1140264700 1014000    fml-rt1.burstat.2004-06-23.2120-2004-06-23.2125
129.105.21.229 198.49.208.10 174      1138227500 1014600    fml-rt1.burstat.2004-06-24.0120-2004-06-24.0125
198.49.208.10 129.105.21.229 196      1106719500 1114000    fml-rt1.burstat.2004-06-24.0120-2004-06-24.0125
129.105.21.229 198.49.208.10 175      1086035800 980500     fml-rt1.burstat.2004-06-23.1920-2004-06-23.1925
198.49.208.10 128.100.190.11 182      1085264900 980500     fml-rt1.burstat.2004-06-23.1920-2004-06-23.1925
198.49.208.10 128.100.190.11 213      1062479100 960000     fml-rt1.burstat.2004-06-23.2120-2004-06-23.2125
198.49.208.10 129.105.21.229 180      1051220800 1093500    fml-rt1.burstat.2004-06-23.2120-2004-06-23.2125
128.100.190.11 198.49.208.10 242      1012027800 842100     fml-rt1.burstat.2004-06-23.0120-2004-06-23.0125
198.49.208.10 128.100.190.11 206      1007483100 916300     fml-rt1.burstat.2004-06-23.2120-2004-06-23.2125
128.100.190.11 198.49.208.10 200      1001671900 842300     fml-rt1.burstat.2004-06-24.0120-2004-06-24.0125
128.100.190.11 198.49.208.10 231      989225200  817700     fml-rt1.burstat.2004-06-23.1920-2004-06-23.1925
198.49.208.10 129.105.21.229 211      957567200  1050100    fml-rt1.burstat.2004-06-24.0120-2004-06-24.0125
131.215.144.227 198.49.208.10 198      946292400  876500     fml-rt1.burstat.2004-06-23.2120-2004-06-23.2125
131.215.144.227 198.49.208.10 209      936021800  882900     fml-rt1.burstat.2004-06-23.2050-2004-06-23.2055
131.215.144.227 198.49.208.10 196      932688300  857700     fml-rt1.burstat.2004-06-24.0850-2004-06-24.0855
131.215.144.227 198.49.208.10 206      904774900  848500     fml-rt1.burstat.2004-06-24.0650-2004-06-24.0655
```

Demzmon0.deemz.net at Fermi had large traffic flows only during the discard episodes.

## OWAMP Latency Testing Confirms Analysis



## DOE Lab - University Connectivity Monitoring Status

Start data collection at 2 sites.	3Q04
Start data collection between 3 labs and 3 universities.	4Q04
Deploy OWAMP measurement systems in ESnet hubs.	4Q04
Start bandwidth testing to at least 1 Abilene node.	1Q05
Evaluate generalized measurement frameworks (IEPM, PIPES, JRA1) to run tests and archive & publish results.	2Q05



# ESnet Performance Centers

## ESnet Annual Program Review

August 3, 2004

Joseph Metzger  
ESnet Network Engineering Group  
Lawrence Berkeley National Lab  
metzger@es.net



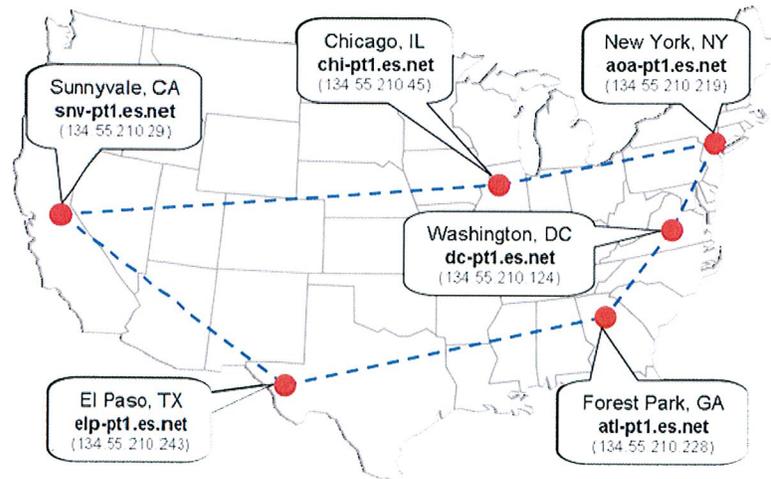
## ESnet Performance Center (EPC)

- Debugging network application performance problems is complex.
- The EPC simplify the task.
  - They are 'known good' test endpoints.
  - They are distributed across the ESnet Backbone to facilitate partial path analysis.
  - Users access the EPC via a secure web interface when and where they want.



## EPC Locations

---



## EPC Usage in the Last Year

---

- Designed as a diagnostic tool.
  - Facilitated debugging PSC to LBL performance problems.
  - Identified firewall performance problem at Ames Lab.
- Used by several sites as a maintenance verification tool.
  - BNL: Testing and Tune proxy servers
  - GA : Verifying performance during LAN redesign.



# Network Availability Measurement ESnet Annual Program Review

**August 3, 2004**

Mike O'Connor  
ESnet Network Engineering Group  
Lawrence Berkeley National Lab  
moc@es.net



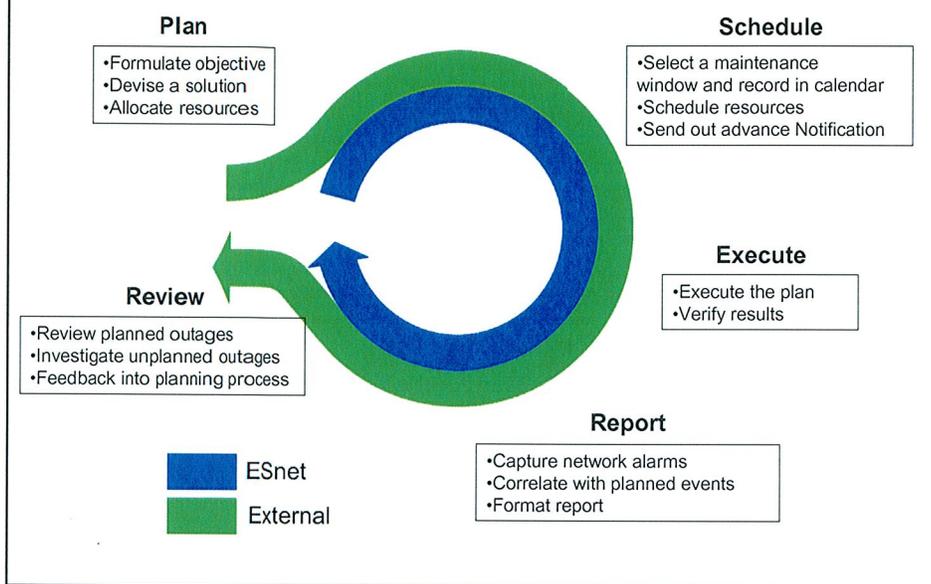
## Network Availability Measurement

---

### Project Objectives

- Develop and deploy a system which will track ESnet network availability, producing clear and concise reports.
- Accurately reflect both planned and unplanned network outages for ESnet, it's customer sites, contracted carriers and network peers.
- Provide a measure of "uptime" for any given site and for the ESnet Backbone.
- Improve maintenance planning and eliminate or improve response time to repetitive systematic network failures.
- Increase network availability for ESnet customers.

# Planned Maintenance Management Cycle



# ESnet's Network Management System (NMS)



The NMS tracks device availability in the ESnet, reporting where and when network components are unavailable.

- However:
- No alarm context
  - Redundant alarms
  - Irrelevant Alarms

**TOPOLOGY MAP**

**PERFORMANCE VIEW**

**ALARM LIST**

Alarm ID	Severity	Message	Time	Device
1000000001	Warning	APRISMA SPECTRUM: [Device Name] is down	2010-10-26 10:00:00	APRISMA SPECTRUM
1000000002	Warning	APRISMA SPECTRUM: [Device Name] is down	2010-10-26 10:00:00	APRISMA SPECTRUM
1000000003	Warning	APRISMA SPECTRUM: [Device Name] is down	2010-10-26 10:00:00	APRISMA SPECTRUM

# Tracking Network Changes

Daily network change reports are essential to maintaining an accurate NMS topology map.



[Main](#)

[Daily](#)

[Monthly](#)

[LinkDB](#)

[Deltas](#)



FILE:delta-06-7-04.html

**ADDED:** FNAL peering @ Starlight in Chicago

198.151.133.253		198.151.133.252			
chi-rt1	ge-3/0/0.13	198.151.133.253	198.151.133.253	255.255.255.252	85

**DELETED:**

**CHANGED:**

bechtel-ga.es.net		134.55.33.32			
-------------------	--	--------------	--	--	--

**CURRENT:**

bechtel-rt1	Serial3/0	134.55.33.34	bechtel-ga.es.net	255.255.255.224	3
gac-rt2	Serial5/1/0	134.55.33.33	ga-bechtel.es.net	255.255.255.224	9

**PREVIOUS:**

bechtel-rt1	Serial3/0	134.55.33.34	bechtel-ga.es.net	255.255.255.224	3
gac-rt2	Serial5/0/0	134.55.33.33	ga-bechtel.es.net	255.255.255.224	8

**CHANGED:**

192.73.7.33		192.73.7.32			
-------------	--	-------------	--	--	--

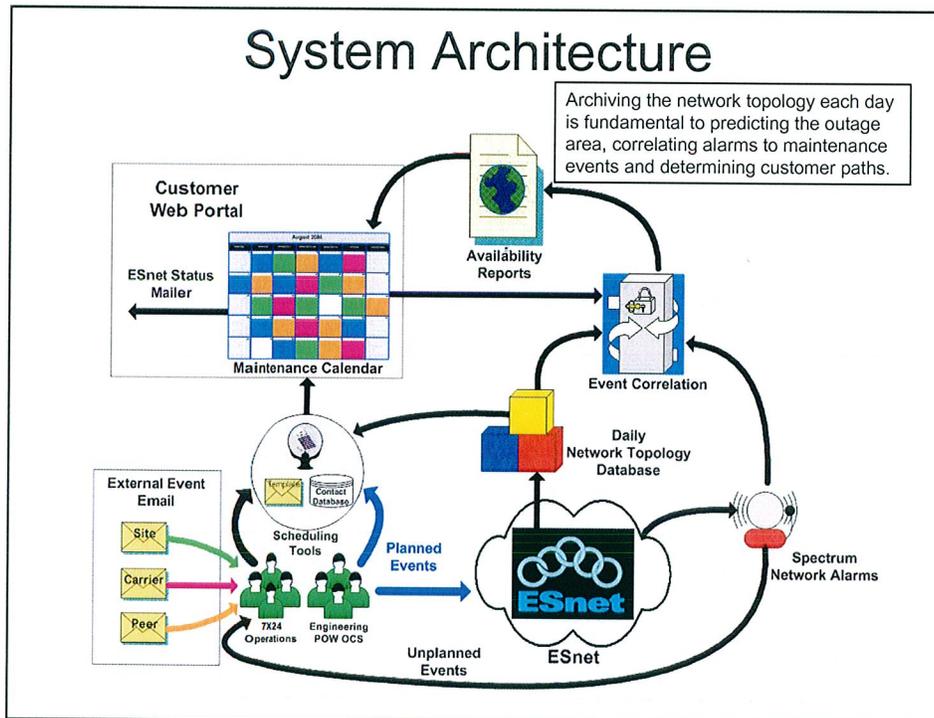
**CURRENT:**

gac-rt2	FastEthernet5/0/0	192.73.7.33	192.73.7.33	255.255.255.224	8
---------	-------------------	-------------	-------------	-----------------	---

**PREVIOUS:**

gac-rt2	FastEthernet6/0/0	192.73.7.33	192.73.7.33	255.255.255.224	13
---------	-------------------	-------------	-------------	-----------------	----

## System Architecture



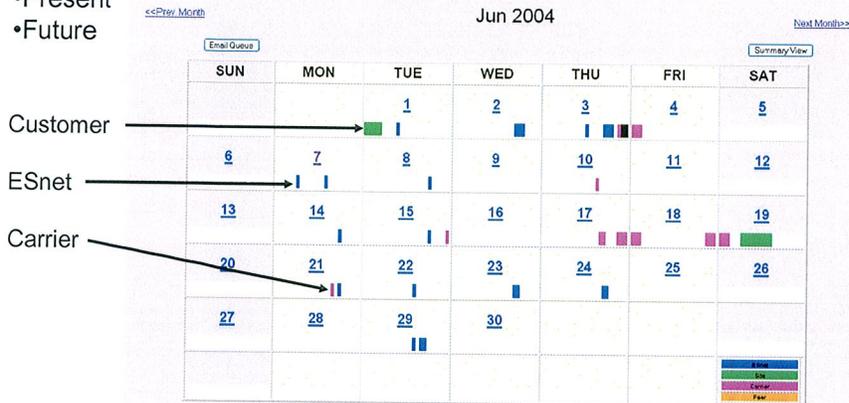
# Public Calendar Monthly View

ESnet network maintenance customer web portal

Network maintenance

- Past
- Present
- Future

ESnet Planned Maintenance Calendar



# Event Information View

- Bringing together:
- Scheduling
  - Customer impact
  - Email notifications
  - NMS Alarms
  - Organization/Culprit

**ESnet**

**Event View**

- Monthly View
- Event View
- Outage View
- Device View
- Add Email Msg
- Edit Event
- Delete Event

## [11606]-SLAC-Rt4 JunOS Upgrade

Start: Wed Jun 2 18:00:00 US/Pacific 2004 Thu Jun 3 1:00:00 GMT 2004  
End: Wed Jun 2 18:30:00 US/Pacific 2004 Thu Jun 3 1:30:00 GMT 2004

SLAC Planned

Message Sent: Tue Jun 1 6:55:03 US/Pacific 2004  
Subject: [11606]-SLAC-Rt4 JunOS Upgrade  
To: trouble@es.net esnet-status@es.net cxg@slac.stanford.edu hooker@slac.stanford.edu gtb@slac.stanford.edu

ESnet Outage Notification

[11606]-SLAC-Rt4 JunOS Upgrade

Begin: Wed Jun 2 18:00:00 US/Pacific 2004 Thu Jun 3 1:00:00 GMT 2004  
End: Wed Jun 2 18:30:00 US/Pacific 2004 Thu Jun 3 1:30:00 GMT 2004  
Location: SLAC

This outage has been estimated to be 00 (hrs) 10 (min) within the above maintenance window

Description:  
The SLAC-Rt4 router (R10) will be upgraded to JunOS 6.2.

Affected Devices:  
Stanford Linear Accelerator Center, router slac-rt4.es.net will be affected  
Link slac-p03-3mv.es.net  
Link 192.68.191.146  
Link 192.68.191.162

ENERGY SCIENCES NETWORK (ESnet)

24x7 NOC (510) 486-7607 Email: trouble@es.net http://www.es.net

Begin: Wed Jun 2 18:00:00 US/Pacific 2004  
End: Wed Jun 2 18:30:00 US/Pacific 2004

## Path Changes 6/03 to 6/04

---

ANL-DC	LBNL-DC
ANL-West	LLNL-DC
Allied	LLNL
Bechtel	MIT
DOE-ALB	NERSC
DOE-OAK	NREL
DOE	ORAU-DC
GAC	ORNL-DC
INEEL-DC	PNNL-DC
INEEL	PNNL
JLab	PNNL
LANL-DC	PPPL
LANL	Pantex
Yucca	SNLA

**Automated tools are essential for keeping up!**

**63 customer path changes**

**66% of ESnet customers experienced path changes last year**

**PNNL had the most with 5 due to the creation of the Seattle hub.**

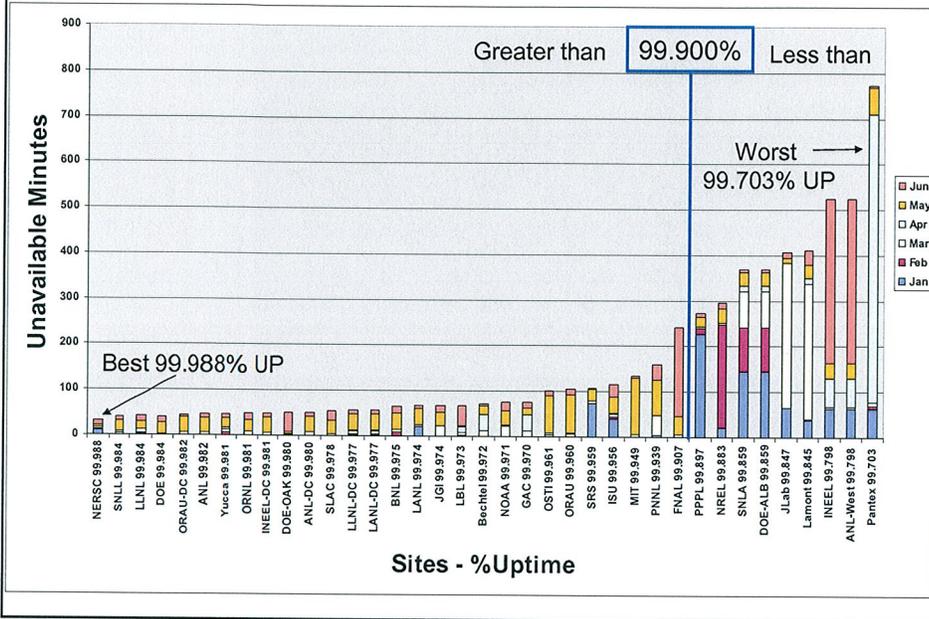
Customer availability is measured along the path to each individual customer site. The result is a "path based availability metric".

## Network Alarm Processing

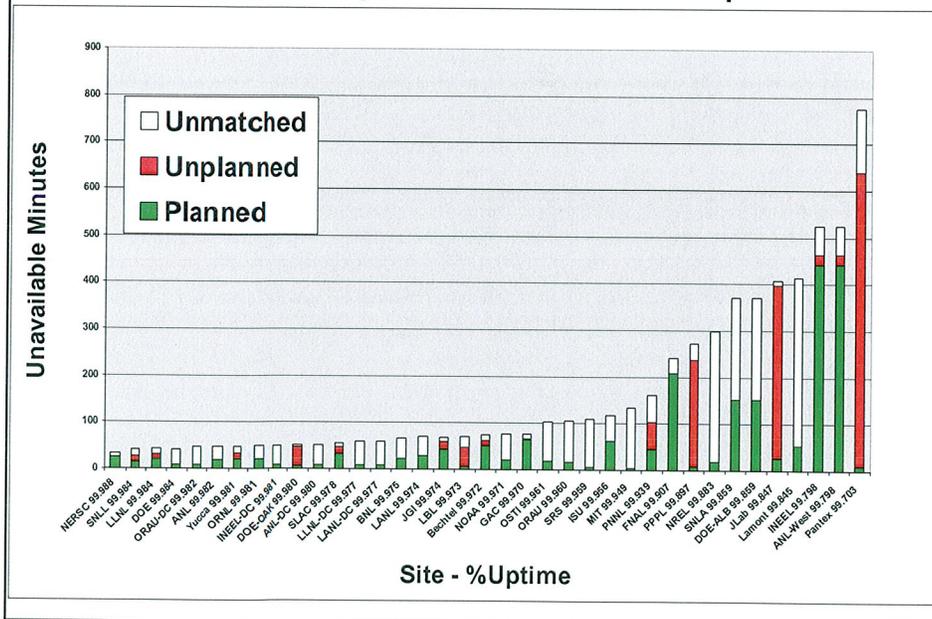
---

- A manageable alarm level is achieved through merging the redundant and suppressing the irrelevant alarms.
- Spectrum reported 16,342 device unavailable alarms between June 2003 and April 2004, 48.8 per day on average.
- Path based analysis isolated 1,448 customer relevant events, an average of 4.3 per day.
- 63% of customer outage minutes in 2004 were categorized using the new network maintenance calendar system.

## 2004 Availability by Month



## 2004 Availability Planned Vs. Unplanned







# On-Demand Secure Circuits and Advance Reservation System (OSCARS)

## ESnet Annual Program Review

August 3, 2004

Chin Guok

ESnet Network Engineering Group

Lawrence Berkeley National Lab

chin@es.net



## Purpose of The ESnet On-Demand Secure Circuits and Advance Reservation System (OSCARS)

---

### Motivation:

- Service sensitive applications (such as remote controlled experiments, time constrained massive data transfers, video-conferencing, etc.), require network guarantees.

### Objective:

- To develop and deploy a new service that can provide secure guaranteed bandwidth circuits within ESnet.

## Issues That OSCARS Must Address (1/2)

---

1. **Adopting The Appropriate Service Model** (1Q05)
  - Use Multi-Protocol Label Switching (MPLS) and Resource Reservation Protocol (RSVP) to create a Label Switched Path (LSP) with guaranteed Quality-of-Service (QoS).
2. **Configuring Acceptable Availability Levels** (1Q05)
  - Guaranteed bandwidth paths can consume up to a certain percentage of the line (e.g. 40%)with overrun traffic being dropped.
3. **Scheduling Bandwidth Reservations** (2Q05)
  - The network can report instantaneous usage, but a reservation scheduler is needed to provide a virtual future view of bandwidth availability.
  - Planned and unplanned network outages must also be taken into account. (3Q07)

## Issues That OSCARS Must Address (2/2)

---

4. **Creating A Guaranteed Bandwidth Path** (3Q05)
  - Prior to the start of a reservation, routers must be configured to setup the bandwidth guaranteed path. Tear down is done when the reservation expires.
5. **Having A Simple User Interface** (3Q05)
  - With an initial limited user base, a simple web-based user interface will be used.
6. **Securing The System** (2Q05)
  - Reserving bandwidth will initially be limited to a few users via web accounts.
  - DOEGrids certificate authentication will follow. (3Q06)
  - Routers access to setup/teardown bandwidth guaranteed paths will conform to ESnet's security model.

## Issues That OSCARS Must Address (3/3)

---

7. **Monitoring Usage** (2Q06)
  - To ensure that the service is not being abused, a monitoring system will be developed and deployed.
  
8. **Enforcing Usage Policies** (3Q05)
  - As the service evolves, a first-come-first-served method may become inadequate. *Design of the policy is beyond the scope of this project.*



# Increasing the Robustness of ESnet to External Threats

## ESnet Annual Program Review

August 3, 2004

Michael Collins

ESnet Network Engineering Group

Lawrence Berkeley National Lab

collins@es.net



## Increasing the robustness of ESnet to external threats

---

- What is ESnet's CyberSecurity defense against large scale CyberEvents?
  - Threat is DOS attacks, either on deliberate target(s) or on random targets generated by, for example, a worm.
  - The solution is to re-architect ESnet HUBs
    - Permitting a phased filtering response to external cyber threats.
    - Facilitating lifeline communications during a massive attack.
- How will ESnet recover from a failure of the primary NOC at LBNL? (discussed a little later).
  - Threat is LBNL NOC is down for extended time.
  - The solution is to replicate essential NOC services elsewhere.

## CyberSecurity Defense

---

### Project Goal

To insure that ESnet sites have the ability to communicate with each other and, to some extent, with the global Internet, even when confronted by major cyber attacks.

## CyberSecurity Defense Approach

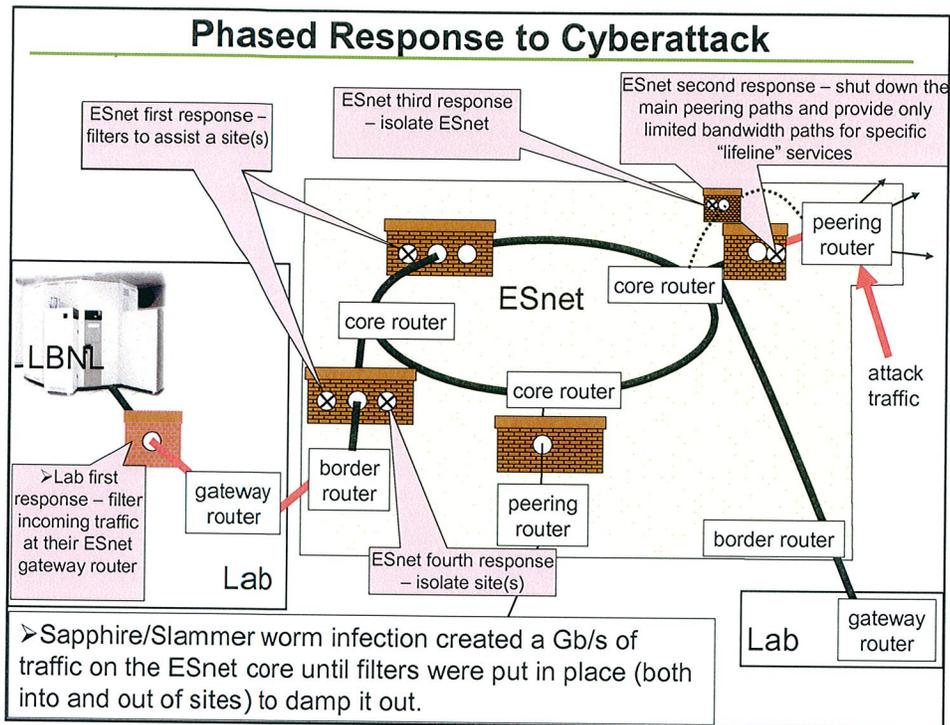
---

- Modify the ESnet HUB architecture to separate ESnet's core routing functionality from the external Internet connections by means of a hardware enhancement to our core architecture.
- Instantiate a rate limited path to the external Internet that will insure that Site-to-Site communication can not be subject to an external denial of service attack.

## CyberSecurity Architecture – Lifeline Filters

---

- Placing firewall filters on the rate limited path will allow minimal Internet access (dns, http, e-mail, ssh, etc.) during a massive attack.
  - Allow communication with other Agencies (CIAC, FBI).
  - Allow downloading of software patches.
  - Allow researchers to reschedule activities with Collaborators.
  - Allow communication with news servers



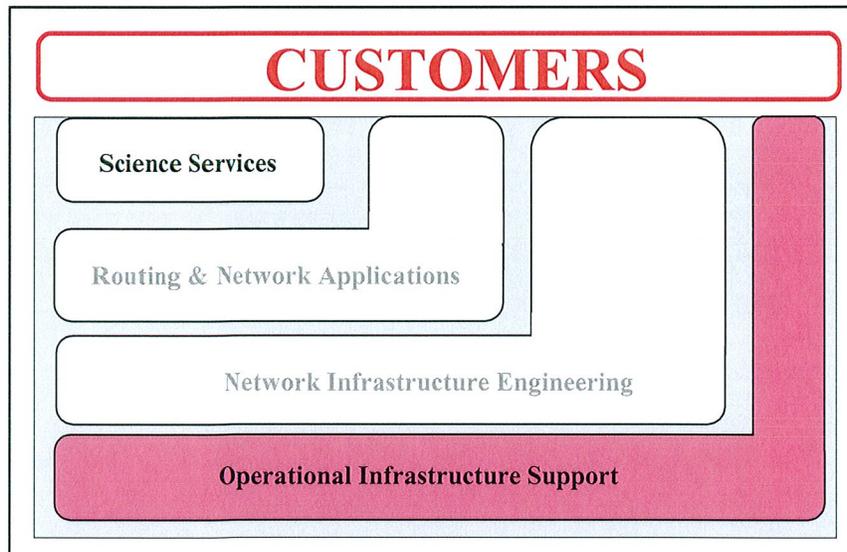
## CyberSecurity Architecture permits a phased response to external threats

- Phase I - Site self help
  - Site(s) install filters or other mitigating factors.
- Phase II - ESnet assistance
  - ESnet installs filters to protect site(s).
- Phase III - ESnet intervention
  - ESnet takes down high speed connects forcing traffic through choke point and lifeline filters.
- Phase IV - ESnet isolation
  - ESnet takes down the low speed choke points isolating sites from the Internet.
- Phase V - Site isolation
  - ESnet shuts down site's connection(s) to ESnet.

## CyberSecurity Defense Status Milestones

Complete new hardware architecture for all hub locations by installing peering routers and low speed interconnects	1Q05
Complete defining and testing lifeline filters.	1Q05
Configure filters and test failover	2Q05
In production	2Q05

## ESnet's Role in Support of Science





# ESnet Trouble Ticket Process

## ESnet Annual Program Review

August 3, 2004

Michael Collins

ESnet Network Engineering Group

Lawrence Berkeley National Lab

collins@es.net

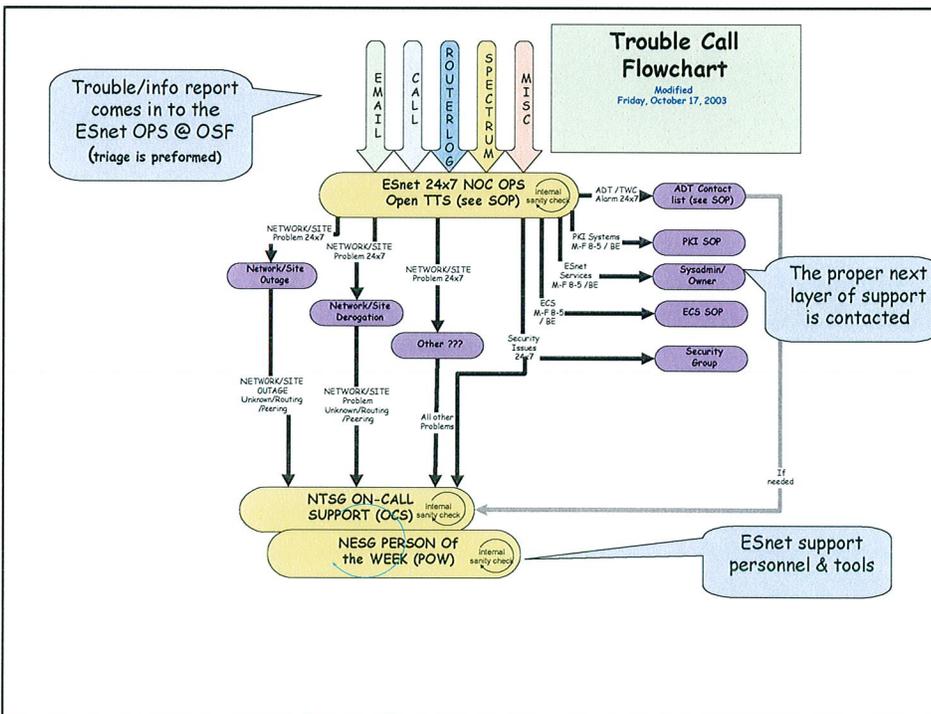
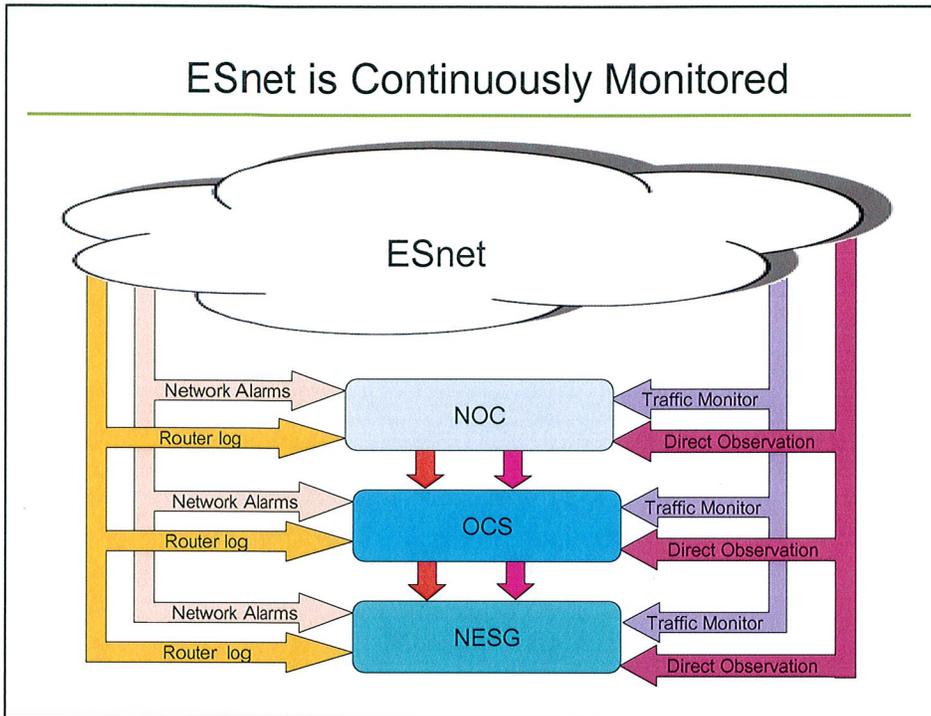


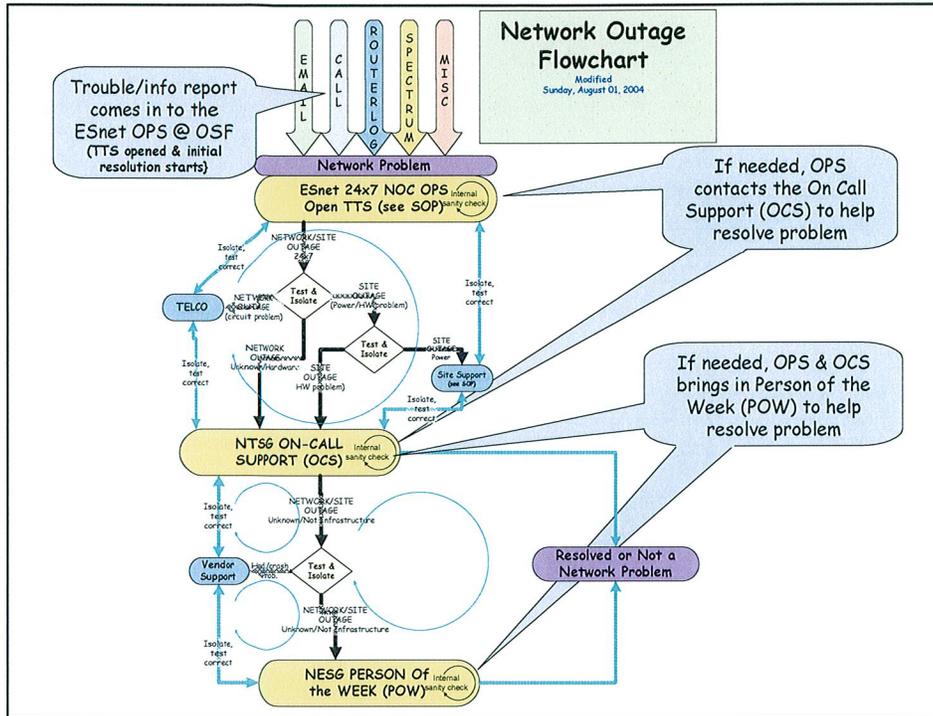
## Definitions

---

- **OPs** (Operations) a member of the shared OSF/ESnet operator staff at OSF
  - 24x7 coverage
- **OCS Team** comprised of the networking staff of the NTSG (Network Technical Support Group)
- **OCS** (On Call Support) a member of the OCS Team
  - Weekly rotation among 7 members
- **POW** (Person of the Week) a member of NESG (Network Engineering Services Group)
  - Weekly rotation among 7 members

# ESnet is Continuously Monitored





## Spectrum Alarm

Enterprise Alarm Manager: Main

File View Model Alarms Troubleshooter Options Help

Mail Filter Sort Order Select All Ack Silence Clear

System \* Probable Cause \* Events \* History \* Alarm Status Troub

**chi-r1**

**A BGP4 PEER BACKWARD STATE TRANSITION HAS OCCURRED**

SYMPTOMS:  
A BGP session with a peer has transitioned from a higher state to a lower state.

PROBABLE CAUSES:  
1) Communication with the peer has been interrupted.

Severity	Date/Time	Model Name	Network Address	Manufacturer
Major	Thu 29 Jul 2004 01:03:22	chi-r1	134.55.200.2	Juniper Networks
Major	Thu 29 Jul 2004 07:47:39	srv-r1	134.55.200.1	Juniper Networks
Minor	Thu 22 Jul 2004 01:40:41	dc-or1	134.55.200.28	Juniper Networks
Minor	Fri 30 Jul 2004 10:05:58	partex-srv1-e.es.net	134.55.30.130	Marconi

Search: [ ] Shown: [ ]

Filtered by Primary/Secondary State

Critical: 0 Major: 2 Minor: 2 Total: 4

Displays detailed information about the selected model

Servers

## Router Log: An example of the SNMP module in multiple routers detecting an unauthorized access attempt. Example of SNMP Scanning

---

Jul 28 10:56:54 **llnl-dc-rt2** 62: 2w6d: %SNMP-3-AUTHFAIL:  
Authentication failure for SNMP req from host 82.41.148.64

Jul 28 10:57:00 **twc-rt1** 96: 1w0d: %SNMP-3-AUTHFAIL:  
Authentication failure for SNMP req from host 82.41.148.64

Jul 28 10:57:01 **pnl-rt1** 173: 10w6d: %SNMP-3-AUTHFAIL:  
Authentication failure for SNMP req from host 82.41.148.64

Jul 28 10:57:15 **gac-rt2** 380: 7w2d: %SNMP-3-AUTHFAIL:  
Authentication failure for SNMP req from host 82.41.148.64

Jul 28 10:58:21 **llnl-dc-rt2** 63: 2w6d: %SNMP-3-AUTHFAIL:  
Authentication failure for SNMP req from host 82.41.148.64

The Miscreant is:

82-41-148-64.cable.ubr04.glen.blueyonder.co.uk



# Disaster Recovery ESnet Annual Program Review

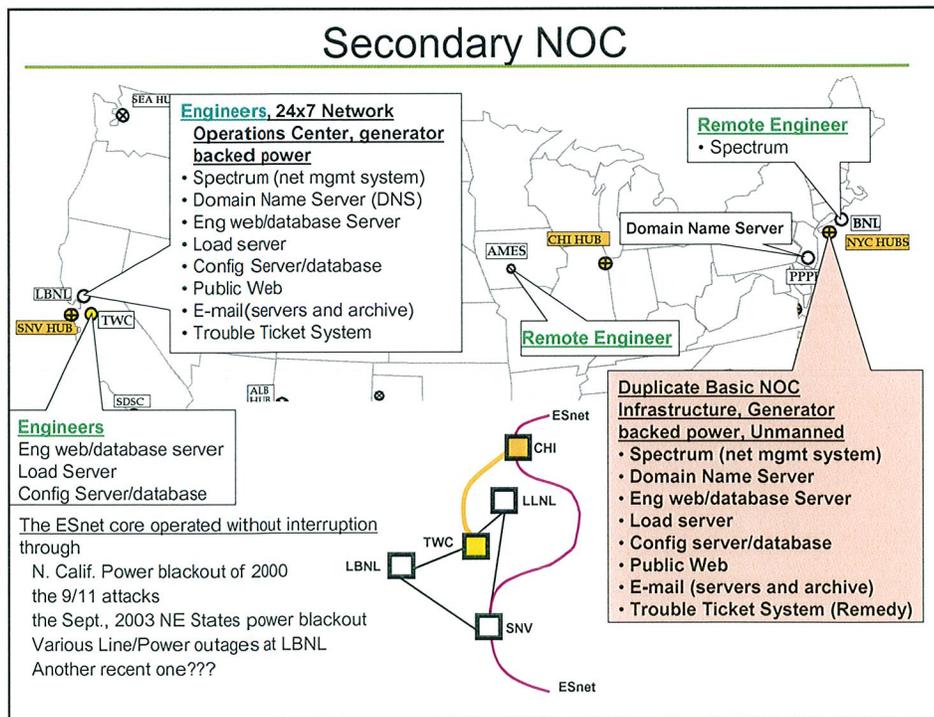
August 3, 2004

Michael Collins  
ESnet Network Engineering Group  
Lawrence Berkeley National Lab  
collins@es.net



# Disaster Recovery

- Goal: To Ensure Operation of ESnet in the Event of a prolonged outage of the primary LBNL NOC.
  - Power Grid Failure
  - Natural Disaster
- Approach: Build an unmanned Secondary NOC.
  - Replicate Basic NOC Services.
  - An unmanned Secondary NOC is possible because of distributed workforce.
  - Locate at AOA HUB because of space, generator power and it is part of the core ring.
  - Install a diverse path from TWC to CHI.



## Secondary NOC Milestones

---

The basic NOC functionality will be replicated at the AOA HUB in two phases. The DC powered systems for AOA are in the procurement process. The AC powered systems to support the NOC services at the other locations have arrived.

Phase 1 Configure Racks at AOA Install 6 systems at AOA (all but Trouble ticket system). Place Domain Name Server, Load Server, Configuration Server/database and Public web server into production	1Q05
Phase 2 Install trouble ticket system at AOA Place Spectrum, Eng Web/database Server, E-mail and Trouble Ticket System into production.	2Q05
Diverse routing to TWC	Open – under investigation



### The Internal ESnet Engineering Web Site

ESnet Annual Program Review

August 3, 2004

**Michael S. Collins**

**Joseph Burrescia**

**and the ESnet Team**

**Lawrence Berkeley National Laboratory**



ESnet Program Review 8/3/04



## The Source of Truth for Engineering Data

It is critical to track various engineering infrastructure databases, diagrams, documentation and tools on which ESnet is built.

- We do this through a variety of techniques linked together through the internal engineering web site.

## Internal Engineering Web Site

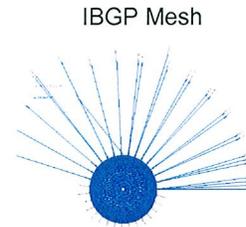
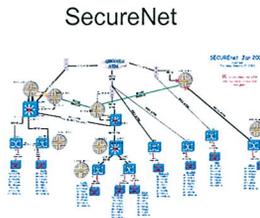
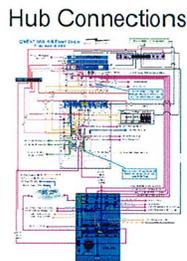
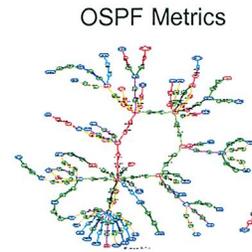
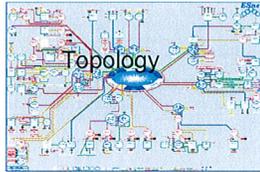
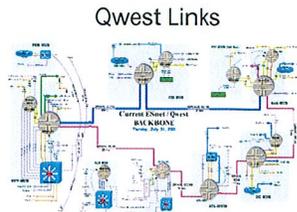
Top level hierarchical view.

The screenshot shows a web browser window displaying the ESnet Network Drawer. The interface includes a navigation menu on the left with items like Topo/Maps, Statistics, NetTools, Routing, ATM, IPv6, IPM/Access, Spectrum, Documents, and P.C's. The main content area is titled 'Network Drawer' and contains sections for 'ESnet Network Maps' (with links to Backbone Maps, WAN Topology Maps, Core Backbone Map, Services Map, LAN/WAN Maps, and MROSE Map) and 'Network Statistics' (with links to MRTG Graphs, Summary Graphs, Device Specific Graphs, Router Traffic, Router Errors, Switch Traffic, and Switch Errors). A table of routers is visible at the bottom, listing various router models and their polling intervals.

Callouts from pink speech bubbles point to specific parts of the interface:

- DB Portal**: Points to the 'Documents' link in the left navigation menu.
- Tools**: Points to the 'NetTools' link in the left navigation menu.
- Documentation**: Points to the 'Documents' link in the left navigation menu.
- Diagrams**: Points to the 'ESnet Network Maps' section.
- Network Statistics**: Points to the 'Network Statistics' section.

## Maps & Diagrams



## Database Applications

### Contact Information for Sites, Peers & Vendors

- Names, Phone, Addr & Notes
- Links to diagrams of site
- Links to miscellaneous docs provided by sites
- Links to peering agreements
- Circuit information for all wide area lines.
- Contact info & procedures for Circuit providers
- Phone numbers & port assignments for out-of-band access
- Support contract numbers

AS	Exchange	System	AS Name	Phone	Email	Notes
AS 100	STB-EX	stb-ex	100	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 101	STB-EX	stb-ex	101	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 102	STB-EX	stb-ex	102	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 103	STB-EX	stb-ex	103	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 104	STB-EX	stb-ex	104	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 105	STB-EX	stb-ex	105	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 106	STB-EX	stb-ex	106	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 107	STB-EX	stb-ex	107	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 108	STB-EX	stb-ex	108	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 109	STB-EX	stb-ex	109	714-244-1000	stb-ex@stb-ex.com	STB-EX
AS 110	STB-EX	stb-ex	110	714-244-1000	stb-ex@stb-ex.com	STB-EX

## Engineering Data

- **Address Allocations and Allocation Plans.**
  - IPv4, IPv6 & NSAP
- **Trouble Ticket Archive**
  - Plain text archive for searching and manipulation not supported by Remedy Trouble Ticket System.
- **Mailing List Archives**
  - Configuration changes back to 1995
  - Routing and Trouble email lists back to 1993

## Networking Tools



### Additional Networking Tools

- IP Address subnet & Mask calculators
- Interface Description generator
- Links to external "Looking Glasses"
- Internal DNS Request form
- Network File Management Drawer

## Documents

- **Routing Policy and Routing Notes**
- **BGP community usage information**
- **Troubleshooting and Configuration crib sheets**
- **Contact DB and Statistics maintenance documentation**
- **CIDR template**
- **RFC Sourcebook**
- **Lots more....**

## LANWAN

Another window into the Engineering Database

Important Links to Network information

Color button descriptions

4 Color function button for each site

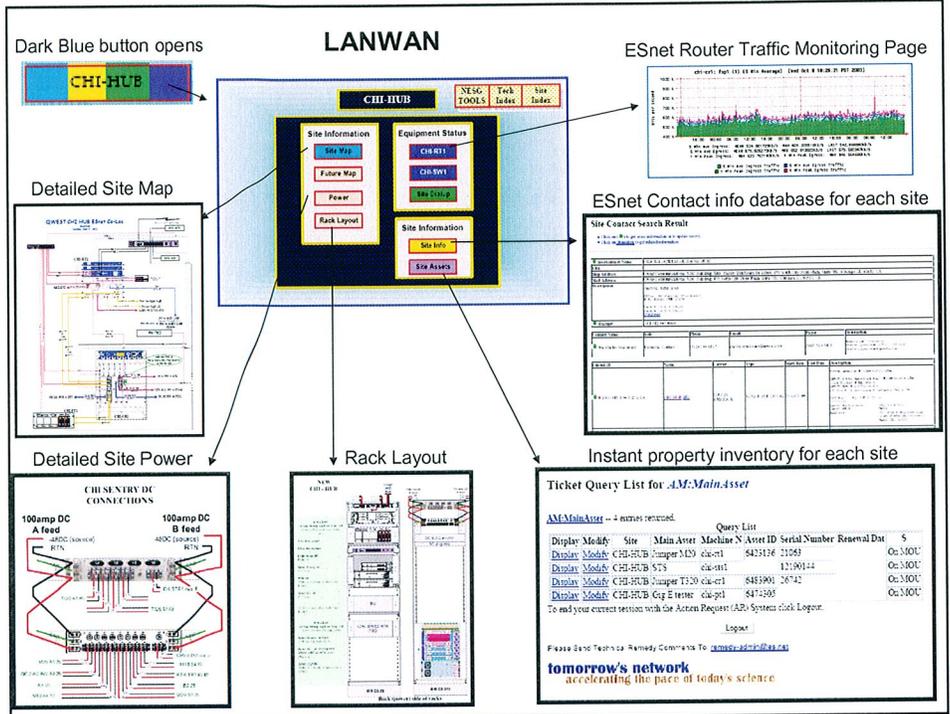
LANWAN Site Index Menu				
Map	Active On/Off	TOPO Index	Install Map	Calendar
ALB-HUB	CHI-HUB	JGI	NYNAP	SAN-DIA
ALLIED	DC-HUB	LANL	ORAU	Seattle
AMES-NAS	DC-LABS	LBL	ORNL	SLAC
ANL	DOE	LLNL	OSTI	SNLA
AOA-HUB	ELP-HUB	LLNLDC	PAIX-WEST	SNLL
ATL-HUB	FNAL	MIT	PAIX-EAST	SNV-HUB
BECHTEL	GAC	NERSC	PANTEX	SRS
BNL	INTEL	NEVIS	PNNL	TYWC
BRANDEIS	ISU	NREL	PPPL	YALE
CEBAT		NYC-HUB		YUCCA

ESnet Energy Enterprise Network

BERKELEY LAB

ESnet Program Review 8/3/04

Office of Science U.S. DEPARTMENT OF ENERGY





# ESnet Review



## Infrastructure and Science Services- Executive Summary

*William E. Johnston*

*August 3, 2004*

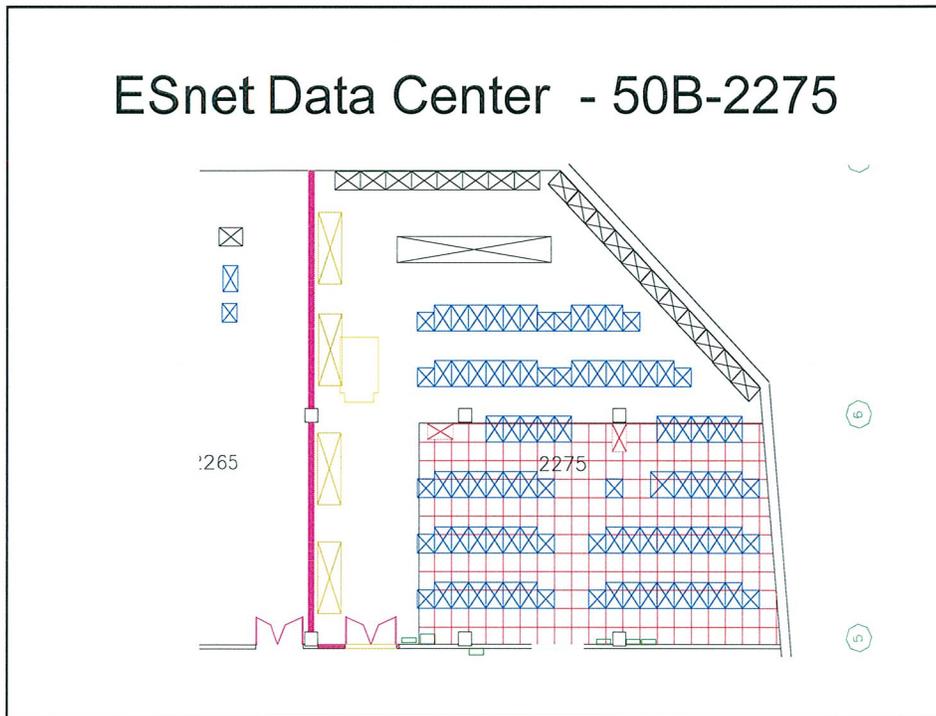
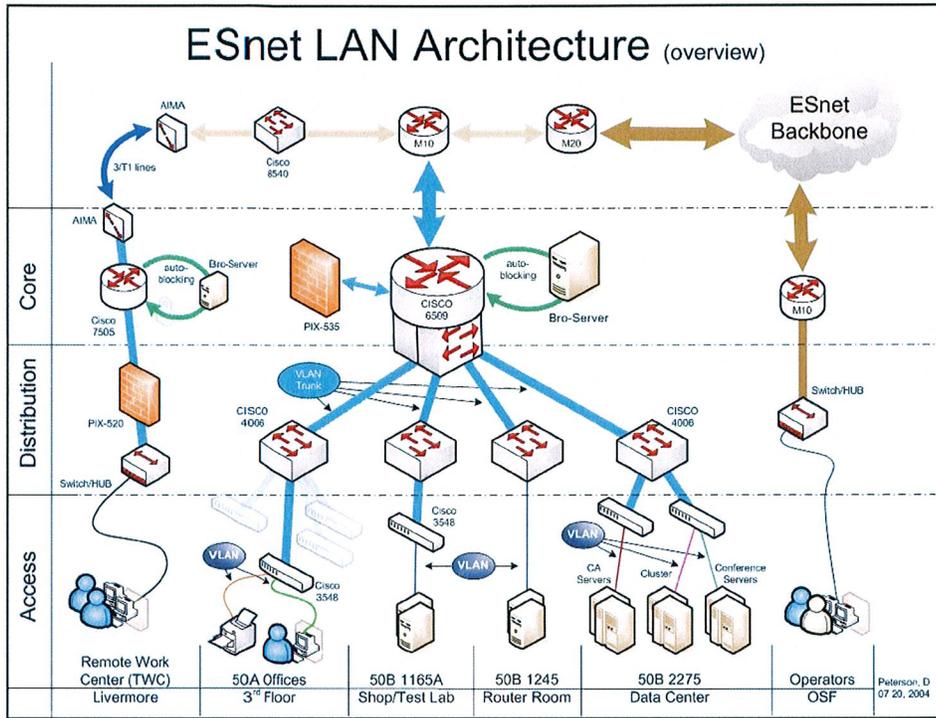
Lawrence Berkeley National Laboratory



## Infrastructure and Science Services

The Infrastructure and Science Services is broken down into the following groups:

- LAN
- Internal security
- Windows Support
- UNIX Support
- Web Services
- Trouble Ticket System
- E-mail and Lists
- Asset Management
- Data Center
- Collaboration Services
- PKI
- NetNews



## Internal Infrastructure Required Upgrades

The following is a list of the internal infrastructure required upgrades:

- LAN \$ 135K
- Internal security \$ 350K
- Windows Support \$ 320K
- UNIX Support \$ 450K
- Web Services
- Trouble Ticket System \$ 75K
- E-mail and Lists
- Asset Management \$ 4K
- Data Center \$2,160K
- NetNews

**TOTAL \$3,494K**



## ESnet Collaboration Services

- Current Production Conferencing
  - Scheduled and Ad-hoc Video (IP and ISDN) ~ 4500+ hours/mo.
  - Scheduled Audio conferencing ~2500 hours/mo.
  - Data conferencing ~150 hours/mo.
- Requirements to meet increasing demands in FY'05
  - Additional blade for IP video @ \$120k
  - Additional blade for Audio Conf @ \$230k
  - Scheduler Customization Eng. @ \$50k
- Projected Cost above current FY'05 Collaboration budget \$450K

# ESnet ATF Project

## **Authentication, Trust & Federation Services for Office of Science**

- Certification Authorities
  - 1800+ Valid Certificates for people, hosts, and Grid services
  - Advanced CA technology and deployment
    - Hardware Security Modules
    - Online public CA's and offline root CA
- Federated Trust
  - New initiative – generalize CA's and underlying approach
    - OTP – RADIUS Authentication Fabric
      - Support Labs' One Time Password token initiatives
      - Provide platform for further advancement
    - Advanced CA services
      - Generalize CA interfaces to support Grid services and other applications more directly

**ESnet needs \$350K per year to support these software and hardware commitments.**



## ESnet Review



# Infrastructure and Science Services

*August 3, 2004*

Lawrence Berkeley National Laboratory



## Infrastructure and Science Services

The Infrastructure and Science Services is broken down into the following groups:

- LAN
- Internal security
- Windows Support
- UNIX Support
- Web Services
- Trouble Ticket System
- E-mail and Lists
- Asset Management
- Data Center
- Collaboration Services
- PKI
- NetNews



# ESnet Review Local Area Network (LAN)

August 3, 2004

Dan Peterson  
John Paul Jones  
Chris Cavallo  
Mike O' Connor

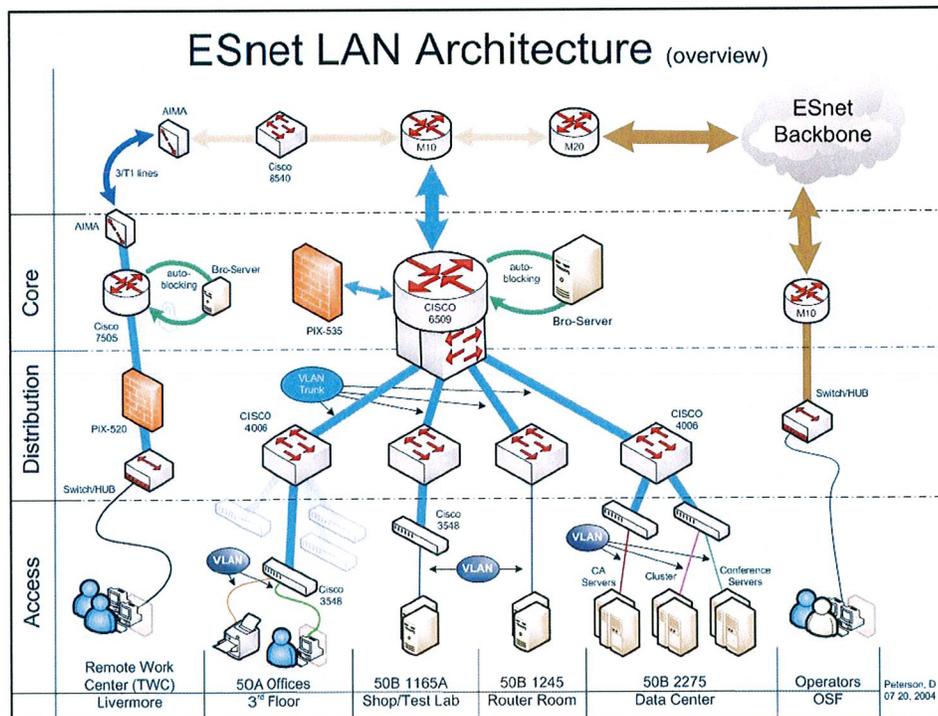


## ESnet LAN

- The ESnet Local Area Network (LAN) is where ESnet staff, *servers and services meet the Internet.*
- The LAN team's goal *is to provide the most flexible, easily managed LAN, with the highest possible levels of security and transparency for customers and staff.*

# ESnet LAN Architecture

- **Gigabit Core LAN Switch**
  - Cisco 6500 platform provides carrier class redundancy and throughput
  - Virtual LAN (VLAN) domain master
- **Gigabit Distribution Layer Switches**
  - Gigabit speeds in distributed wiring closets and clusters at reduced cost
- **VLAN capable access layer switches**
  - Fast flexible office connectivity



## LAN Milestones – FY05

- ESnet staff VPN project - \$30K
  - Office & 3rd Floor Upgrades - \$5K
  - 3rd Fl. Closet Cabling Completion - \$5K
  - CAT Switch Enhancements - \$12k
- Total - \$52K



## ESnet Review Internal Security

August 3, 2004

Ken Pon  
Don Varner  
Dan Peterson  
John Webster  
(With assistance from all ESnet staff)



## **ESnet Internal Security**

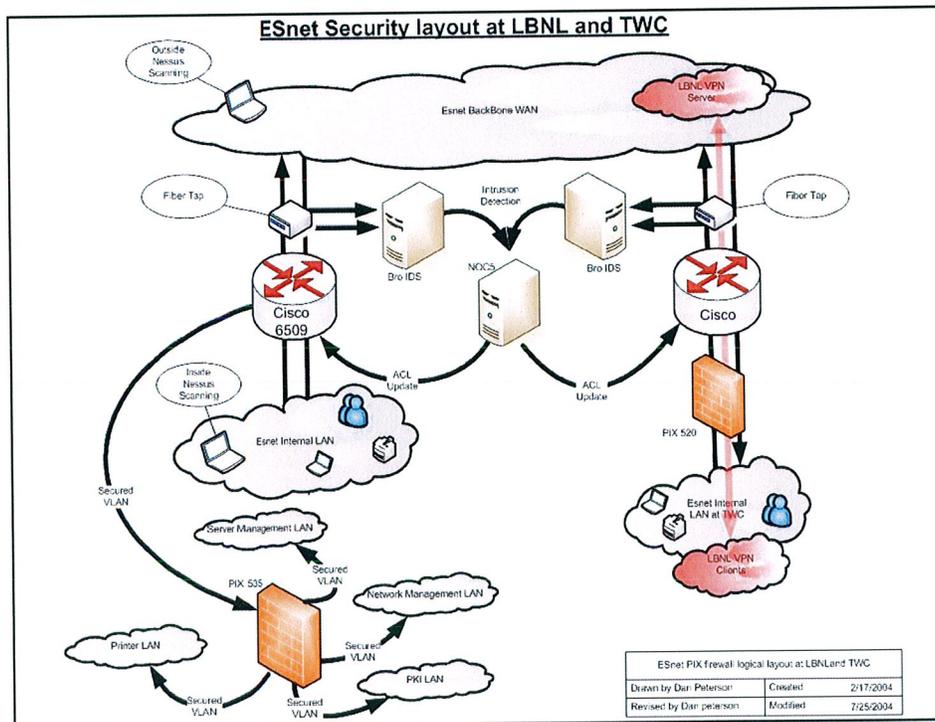
- ESnet's Internal security team mission:
  - *To secure and protect ESnet computing resources against external and internal threats to highest level possible, while continuing to providing an open computing environment to ESnet staff and collaborators.*
  - *To insure that the most dangerous and persistent attackers, worms, and viruses are blocked as quickly as possible.*

## **ESnet Internal Security**

- Vulnerability self assessment
  - Nessus scanning and reporting
  - Windows® auto hotfix scanning and reporting
  - Analyze BRO intrusion detection and other security related logs on a frequent basis.
  - Regular simulated attack scanning

# ESnet Internal Security

- Intrusion Detection System (IDS)
  - Two "Bro" IDS
    - Located at LBNL and TWC
    - Centralized auto-blocking to our Cisco 6509 (LBNL) and Cisco 7505 (TWC).
- Firewall
  - Cisco PIX 535 – LBNL
  - Cisco PIX 525 – LBNL
  - Cisco PIX 520 - TWC



## Internal Security Milestones FY05/06

### Security Upgrade (\$350K)

With the increase in malicious network activity it is necessary to upgrade the current security systems so they can keep up with the flood of activity. Upgrading will result in faster detection of intrusions and vulnerabilities. It will also increase the Security Team's ability to keep up with the rising activity.

- Phase 1 - Upgrade Vulnerability Assessment system to allow for passive and active determination of vulnerabilities. (\$250K)
- Phase 2 - Upgrade Intrusion Detection system to increase detection capabilities. (\$100K)



## ESnet Review Windows Active Directory

August 3, 2004

Dan Peterson  
Scott mason

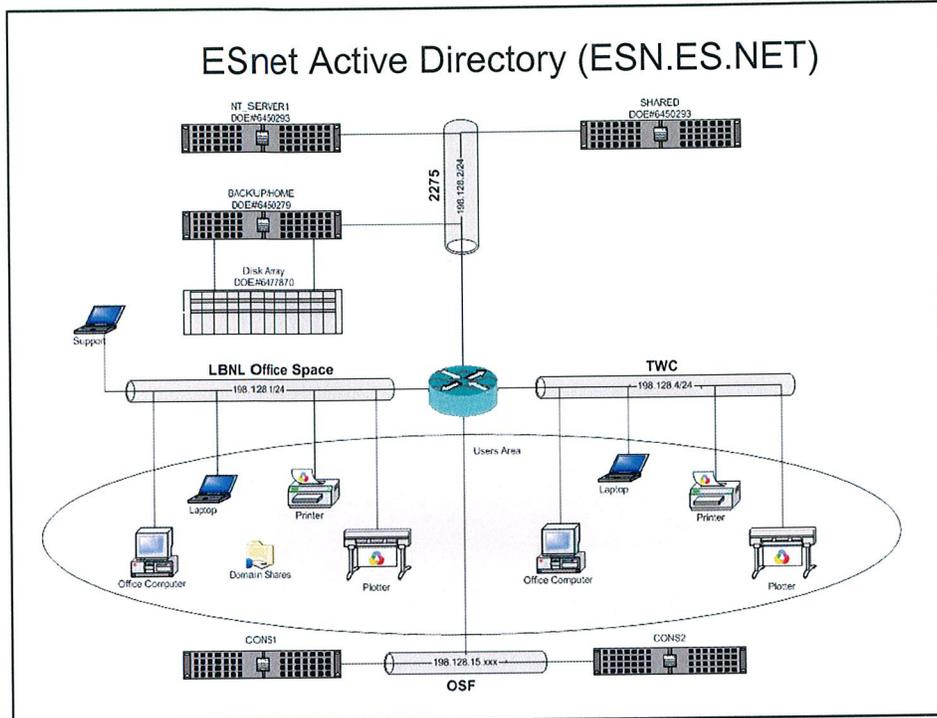


## **ESnet Windows Support**

- Active Directory (AD) provides:
  - Uniform user account database, with centralized user and host management for ~90 hosts and ~30 users.
  - Centralized, project and security specific group policy management including automatic hotfix scanning.
  - Centralized file, print and backup services to domain users.
  - Provide users with a secured remote access via VPN services.

## **ESnet Windows Support**

- Three DELL 2450/2550 servers.
  - Primary domain controller (DC), VPN server.
  - Applications, shared file area and print server.
  - Backups, user directories.
- DELL Power Vault Disk array.
  - Mass storage area for backups and user directories.  
(~750 Gigabytes available storage)



## Windows Milestones – FY05

- ESnet staff workstation application suit upgrades:

• Office 2003 pro	-	\$15k
• VISIO 2003	-	\$5k
• Hardware	-	\$300k
Total	-	\$320k



# ESnet Review

## Unix Support

August 3, 2004

John Webster



## ESnet Unix Support

- Support and maintain ESnet UNIX servers and workstations by performing backups, upgrades (hardware/software), and other preventive maintenance.
  - [Approximate count by Operating System]
  - 55 Solaris servers and workstations
  - 35 FreeBSD servers and workstations
  - 35 Linux servers
  - 1 Macintosh computer
- Monitor system log files for suspicious activity and/or indications of pending system hardware failures.
- Manage, configure and test new UNIX hardware and software installations.

## ESnet Unix Support

### Software/Services Supported

- Web servers - Apache and Netscape
- Oracle database servers
- PGP Key servers
- Network Statistics Collection
- Network Performance Testers
- Sun Ray System
- Unix Backup system
- Hosted-base firewalls
- AFS - Andrews File System
- Remedy - Trouble ticket system
- PKI - Public Key Infrastructure
- Spectrum - Network Monitor
- FTP
- DNS
- BRO - Intrusion Detection System
- Logging servers
- Modem pool

## ESnet Server Monitoring Tools

- Big Brother - monitors systems and services for availability. Status is displayed on a color-coded web page.
- Scripts - monitor system status and either page sysadmins if there are problems or restart services that are down.
- Watchdog functions - some services monitor each other and page sysadmins when services are not available.

## ESnet Unix Support Milestones FY05/06

### Secure Terminal Server Upgrade (\$200K)

The failure rate is increasing on the current Secure Terminal Servers . Also software updates are no longer available from the vendor. Upgrading with new hardware and software will result in increased security and less hardware failures.

- Phase 1 - Upgrade the Secure Terminal Servers on the Wide Area Network to ensure secure and reliable access to critical network switches and routers at sites and hub facilities. (\$100K)
- Phase 2 - Upgrade the Secure Terminal Servers on the Local Area Network to ensure secure and reliable access to critical systems on the local network. (\$100K)

## ESnet Unix Support Milestones FY05/06

### Backup System Upgrade (\$350K)

With the increasing amount of network statistical data being collected the current home-grown and labor-intensive backup system is overwhelmed. The upgrade will enable reliable backups and restores of critical data while decreasing the labor required. The upgrade will build upon the already owned tape library systems.

- Phase 1 - Upgrade software and servers to utilize tape library system. (\$250K)
- Phase 2 - Upgrade storage capacity of tape library systems with new tape drives. (\$100K)



# ESnet Review

## Web Services

August 3, 2004

**Marcy Kamps**

## ESnet Web Services

ESnet provides external web sites for our customers.

The virtual sites are running on an Apache web server.

In the process of upgrading our web server.

Main Web Site	<a href="http://www.es.net">http://www.es.net</a>
ECS	<a href="http://www.ecs.es.net">http://www.ecs.es.net</a>
ESCC	<a href="http://www1.es.net/escc/">http://www1.es.net/escc/</a>
ESSC	<a href="http://www1.es.net/essc/">http://www1.es.net/essc/</a>
SLCCC	<a href="http://www1.es.net/slccc/">http://www1.es.net/slccc/</a>
DOEGrids	<a href="http://www.doe grids.org">http://www.doe grids.org</a>
NetNews	<a href="http://www-news.es.net">http://www-news.es.net</a>

# ESnet Web Services

ESnet Collaboration Service

<http://www.ecs.es.net>

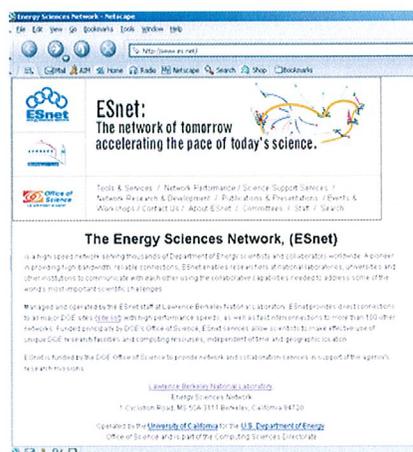
DOEGrids Certificate Service

<http://www.doe grids.org>



## ESnet Web Site Redesigned 4/04

- Tools & Services
- Network Performance
- Science Support Services
- Network Research & Development
- Publication & Presentations
- Events & Workshops
- About ESnet
- Committees



## Internal Staff Web Site

- Staff Calendar
- Internal Business Info
- Internal Services
- Network Information
- Network Maps
- Site Info
- Internal Documents
- NTSG & NESG web sites
- Photo & Video Gallery



## ESnet Review

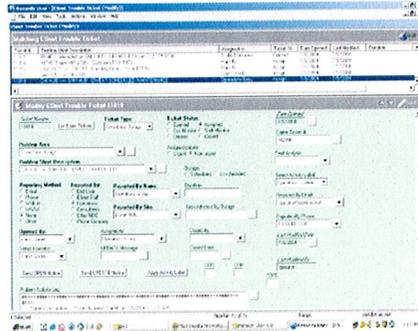
### Trouble Ticket System

August 3, 2004

Marcy Kamps

## ESnet Trouble Ticket System

- TTS used to track information about the Network, ECS, DOEGrids, Asset Management, NERSC, and other services.
- Running Remedy ARsystem server and Oracle database on a Sun Ultra workstation.
- Total external ticket = 11750 (1995-2004), approx. 1300/year
- Total internal tickets = 1300 (1999-2004), approx. 250/year



## Upgrade Trouble Ticket System

- In the process of upgrading our Remedy ARsystem server to version 6.x and our Oracle database to version 9i on a Sun Enterprise 420 machine.
- This project will replace an aging Trouble Ticket system with a new system based on high availability, security, growth, and ease of maintenance.
- Deploy our TTS applications on the web using a Remedy Mid Tier server with Tomcat Apache web server.
- Redundant system out of AOA is planned.
- Purchase Oracle Advanced Replication software and license additional database for mirrored environment approximately \$75k.



# ESnet Review



## Mail and Lists

*Roberto R. Morelli*

*August 3, 2004*

Lawrence Berkeley National Laboratory



## ESnet Mail

ESnet operates in a 24x7x365 role, mostly via email, and we have a need to have a email service that operates in the same fashion. Our internal mail service which is also used to support list and news services operates as follows:

- Cluster of 3 nodes for mail routing, list processing, Realtime Blackhole Lists (RBL), virus protection (bi-directional), content filtering, spam filtering, and access controls in both IPv4 and IPv6.
- Cluster of 2 nodes for IMAP and POP storage via secure channels and methods only. This system is used to store all email and provides secure access for users via a variety of mail clients or the Web.
- Disaster recovery mode and emergency outages. Redundant systems located at AOA, complete mirror of LBNL system.

## ESnet List Service

ISP-style list management and services that provides ESnet's communities, users, groups, and committees a neutral area to host mailing lists. List services can include everything from full filtering for spam, viruses, and/or content to archives accessible for private or public viewing.

Service operates as follows:

- IPv4 & IPv6 Service
- Highly available - 3 nodes
- Subscriptions via email commands
- Currently about 125 lists
  - 50 % Internal
  - 50 % External
- Currently an integrated function of the ESnet postal cluster and replicated at AOA.



## ESnet Review



### Asset Management

August 3, 2004

**Chris Cavallo**

**Scott Mason**



## Asset Management

The ESnet Asset Management System tracks all ESnet network and computing equipment throughout the country.

Over \$10 million in network assets and \$1.62 million in computer assets are tracked in the Remedy database at 49 locations in the US.

ESnet uses the Remedy Action Request System because it stores more information than the LBNL Sunflower Asset Management System.

### Asset Tracked in LBL's Sunflower

DOE# : 6302714  
Make : CISCO SYSTEM  
Model : 7505  
Serial# : 50011100  
Value : \$41,524.00  
Purchase Date : 7/01/1996  
Acct # : 738520  
PO# : 3643400  
Custodian : Cavallo Chris 813303  
Location : 50A-1165  
MOU# : MOU-049 SNLA  
Project ID : 738520

### Asset Tracked in ESnet's Remedy

DOE# : 6302714  
Make : CISCO SYSTEM  
Model : 7505  
Serial# : 50011100  
Value : \$41,524.00  
Purchase Date : 07/01/1996  
Acct # : 738520  
PO# : 3643400  
Custodian : Cavallo Chris 813303  
Location : PNL  
MOU# : MOU-049 SNLA  
Project ID : 738520  
Maint Contract : Yes  
Maint ID : 6515636  
Renewal Date : 1/31/2005  
Machine Name : pnl-rt1  
Site Contact : Cullen Tollbom  
Site Address : Richland, WA 99352  
Date Installed : 4/16/1999  
Date Inventoried : 10/1/2003  
Notes : Yes



## Deployed Assets

- Juniper Routers.....25
- Cisco Routers.....66
- Fore Switches.....17
- Terminal Servers.....14
- SSH Terminal Servers.....40
- Intel-Based Systems (PCs, Servers, Laptops)....167
- RISC-Based Workstations.....142
- Video Equipment.....39
- Printers.....30
- Other Equipment (Modems, UPS's, etc.).....86
  
- Total Assets.....626



## Wide Area Network

- Here is an example of the inventory found at a typical Hub site
- 6 Hubs across the country

Site: SNV-HUB  
 Control / MOU #: 9C000704  
 Site Contract: David Jones  
 Equipment Located: 1400 Kifer Rd  
 Sunnyvale CA 94086  
  
 LBL Custodian: Chris Cavallo  
 Employee ID: 813303  
 Office / Ext: 50A-3141 ext 8680



Main Asset	Machine Name	DOE#	S/N	Purchase Price*
Juniper	T320	snv-cr1	6483895 26891	425,000.00
Juniper	M20	snv-cr1	6418866 21048	134,263.00
Cisco	8540	snv-sw1	6418804 TBA04190530	100,800.86
Pro Tester		snv-pt1	6474275	4,000.00
Sency	4870-XL-4	4port-48VDC		3,900.00
Sency	4820-XL-8	8port-48V-DC		3,350.00
STS		snv-451		3,000.00
Site Value				5674,313.86



# Wide Area Network

- In addition to the 6 hubs, ESnet operates in 43 Key sites
- This report illustrates inventory at a typical Key site



Site: ANL  
 Control / MOU #: M0U-017  
 Site Contact: Scott Pinckton  
 Equipment Located: 9700 South Cass Av. Bldg 221  
 Argonne IL 60439  
 EBL Custodian: Chris Cavallo  
 Employee ID: 813303  
 Office / Ext: 50A-3141 ext:8680

Main Asset	Machine Name	DOE#	S/N	Purchase Price*
JWSR	M20	01-02	646864 6024	73,167.00
STS		01-01	0320059	3,000.00
Site Value				76,167.00



# Projected Costs

We want to purchase a software application and barcode scanners so we can scan all DOE tagged assets and transfer the barcode information directly into Remedy.

The cost of one PDA, two scanners, and the software would be approximately \$4,000.00.



## Asset Movement Tracking

- Shipping Documents created in Filemaker Pro
- Assets tracked through carrier's tracking system
- Set up and monitor RMAs with vendors



## Purchasing Equipment

- Verify equipment requirements
- Research equipment availability
- Obtain purchase authorization
- Order equipment and monitor order process
- Document equipment in Sunflower and Remedy
- Configure and install computer systems



# ESnet Review



## Data Center – Services & Support

*Roberto R. Morelli*

*August 3, 2004*

Lawrence Berkeley National Laboratory



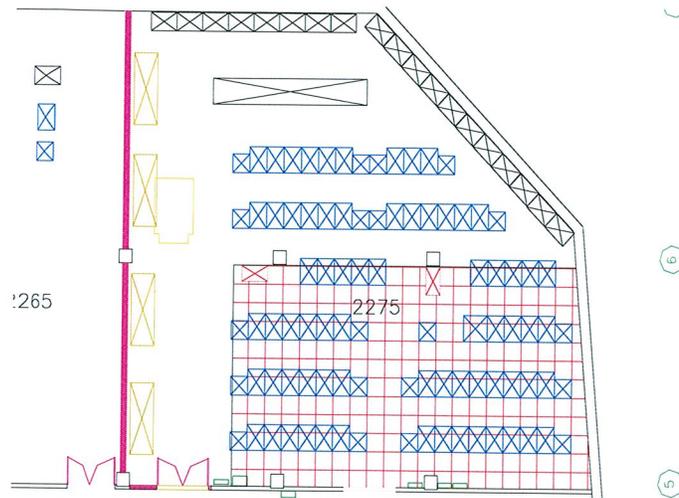
## ESnet Data Center

The ESnet Data Center allows for centralization of all current Esnet services, future planned services, and growth. The Data Center space is about 2550 sq. ft. and contains both production servers and production network equipment.

2275 is a production grade Data Center consisting the following:

- 120 tons of cooling (40 tons in service)
- 320KVA of power (128KVA in service, 48KVA on generator)
- 64 production racks

## ESnet Data Center - 50B-2275



## ESnet Data Center

The Data Center is broken down into the following groups and consists of the following services:

- Networking Services
- Collaboration Services
- PKI Services
- Web Services
- Unix Services
- Windows Services
- Mail & News Services
- Asset Management

## Milestones - FY05/06 (1)

2275:

- Infrastructure power control (\$80K)
  - This project allow us to complete the power control project for the data center. It will allow us to complete control all power from remote and in the advent of prolonged power outages, shut down unneeded service and equipment to keep critical service running. This also allows us to operate the facility in a lights out (unmanned) fashion.
- Infrastructure wiring - Part 2 (\$80K)
  - This project allow us to complete all copper and fiber wiring and bring into full production phase 2 of the data center. This allows us to populate 16 racks with production equipment.

## Milestones - FY05/06 (2)

2275:

- HVAC Upgrades (\$600K)
  - Phase 1 - Upgrade old high maintenance system to a high capacity systems. Current system is old, high maintenance, failing, leaking, and maxed out at about 40 tons. We need in excess of 80 tons to operate 2275 to full capacity. Hardware is currently owned, costs are to remove old system and duct, install new and re-duct and place into service at high capacity. (\$500K)
  - Phase 2 – Install condenser core (DX coils) and wire for possible generator service to allow reduced cooling capacity to all 2275 to operate at reduced capacity in the advent of prolonged PG&E power failure. (\$100K)
    - Alternate option, see slide next slide.

## Milestones - FY05/06 (3)

2275:

- Update ups power in 2275 - Part 1 (\$600K)
  - Replace current inefficient system consisting of 16 16KVA 208V UPS with a single highly efficient single 325KVA 480V ups.
  - Outcome of new system:
    - Generator power relocation on demand
    - Expandable without replacement (modular)
    - 3 phase system, generator friendly, power efficient
    - Full controllability from any ESnet NOC location
  - System can be installed in phased approach without disruption of services and existing equipment can be reassigned to other ESnet locations where it's capacity is more efficient.
  - System is expandable and flexible enough to even cover HVAC on ups power.

## Milestones - FY05/06 (4)

2275:

- Update ups power in 2275 - Part 2 (Generator)

Replace or add additional generator power, current system of 125KVA is full loaded and shared with other functions at LBNL.
- Phase 1 - Wiring and controls

Purchase and install all switching equipment, power panels, and control systems for new or additional generators. Install emergency generator plugs at parking lot to allow facility's emergency mobile generator a plug in point. (\$300K)
- Phase 2 - Generator purchase or lease & options

At this stage we have several options, lease to own or purchase a generator. A mobile 475KVA unit can be leased for about 3K per month including maintenance. Direct purchase price is about \$500K.



## ESnet Review



### ESnet Collaboration Services

*Mike Pihlman  
Clint Wadsworth*

*August 3, 2004*  
Lawrence Berkeley National Laboratory



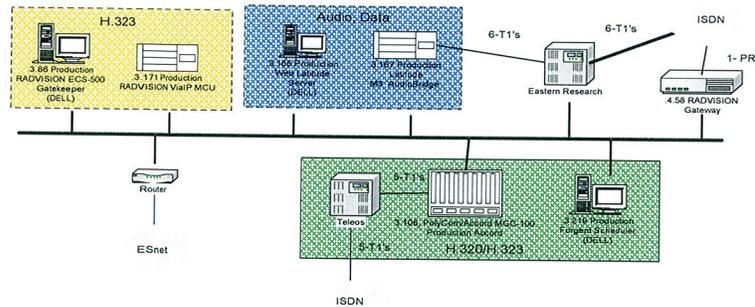
### ESnet Collaboration Goal

To provide seamless voice, video, and data collaboration services and applications for geographically dispersed collaborative groups.



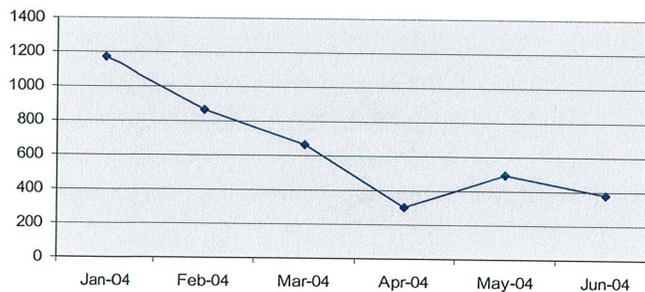
# ECS Production Services

- Web-based registration and scheduling
- Scheduled H.320/H.323 videoconferencing
- Scheduled audio and data conferencing
- Ad-Hoc H.323 videoconferencing



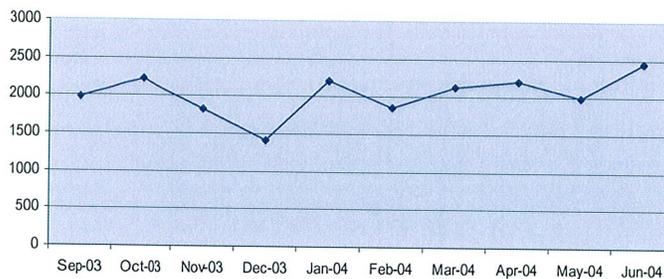
# H.320/H.323 Video Conferencing

- PolyCom/Accord
  - Usage declined rapidly after Jan 30, 2004
  - Will be taken off-line by 1Q05
  - Est. cost saving of ~\$70,000 per year.



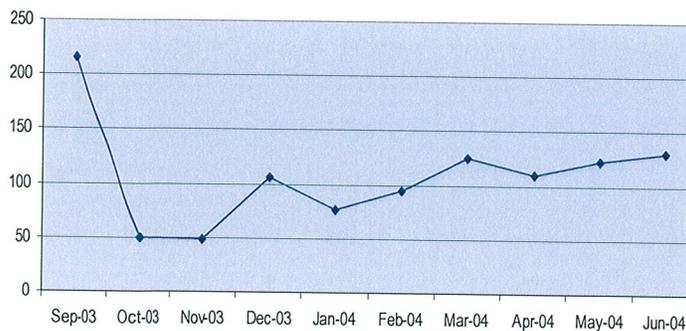
# Audio Conferencing

- Cisco/Latitude
  - 144 voice ports available
  - Usage continues to increase gradually
  - June 2004: 2455 port hours, 766 meetings



# Data Conferencing

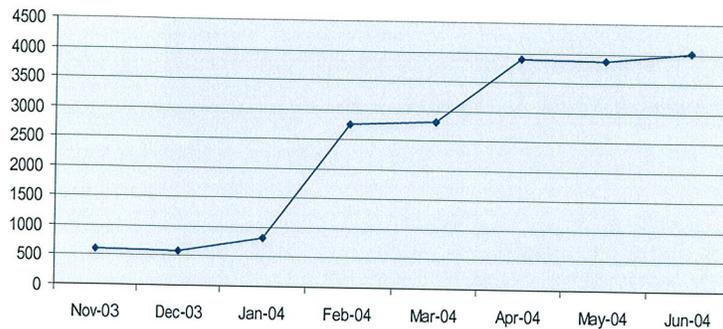
- Cisco/Latitude
  - >2000 data ports available
  - Usage starting to increase
  - June 2004: 130 port hours



# H.323 Video Conferencing

- Radvision

- 70 ports available at 384 kbps
- Usage increased dramatically, now leveling off
- June 2004: 4009 port hours



## FY05 Expectations

Increase H.323 Ad-Hoc capacity by another 70 ports	\$120,000
Increase audio bridge ports to meet expected demand	\$230,000
Streaming of H.323 meetings & additional H.323 capacity	\$50,000 (approved)
Add Network Management app for Radvision	\$8,000 (approved)
Add gateway to H.323 Ad-Hoc bridge	\$40,000 (approved)
Hold the ESnet Collaboration Workshop	\$30,000 (approved)
Total:	\$ 478,000

## Summary

ESnet collaboration is dedicated to providing the most advanced voice, video, and data collaboration technologies for the support of research in the Office of Science.

## ESnet Review

### ESnet ATF Project

03 Aug 2004

Michael Helm

## ESnet ATF Project Team

The Authentication, Trust, & Federation team consists of 3 core FTE's with additional support drawn from multi-disciplinary ESnet staff.

- ATF team:
  - Tony Genovese – software engineer
  - Michael Helm – project lead
  - Dhivakaran Muruganatham – CA Ops/development
- System\*
  - Roberto Morelli - special projects, 2275, RADIUS
  - Ken Pon - intrusion detection lead, security
  - \*John Webster - system lead
- Networking
  - Mike O'Connor - firewall lead
  - Kevin Oberman - networking, NESG security topics
- Web Servers
  - Marcy Kamps – Web developer and management
- ESnet staff

## ESnet ATF Project

Authentication, Trust & Federation Services for Office of Science

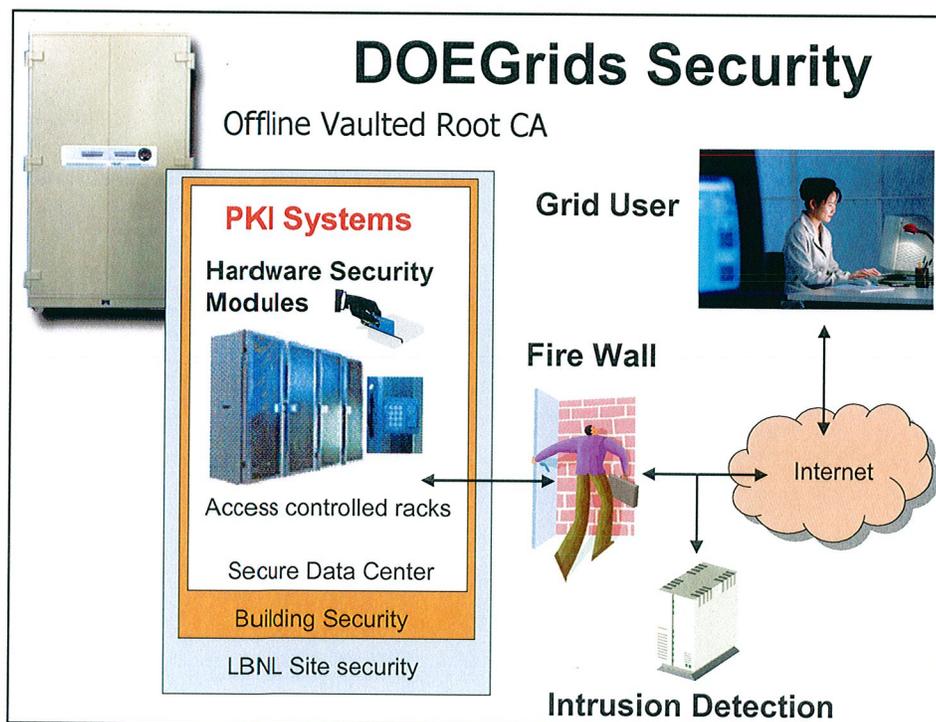
- Certification Authorities
  - ESnet Root CA
  - DOEGrids CA
  - NERSC CA – NERSC's "myProxy-NIM" integration
  - ESnet SSL Server CA – soon to expand
- Scope – X.509/PKIX certificates for Office of Science supported research and collaborations
  - Grids ; TLS ; Experimental uses
- Rigorous security
  - Industry best practices
  - Hardware Security Modules (HSM)
- Services
  - People, host, and service certificates
  - Key lifecycle management
  - User interface development and automation
  - Grid integration

## ESnet PKI Project (2)

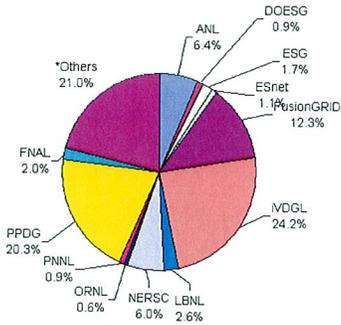
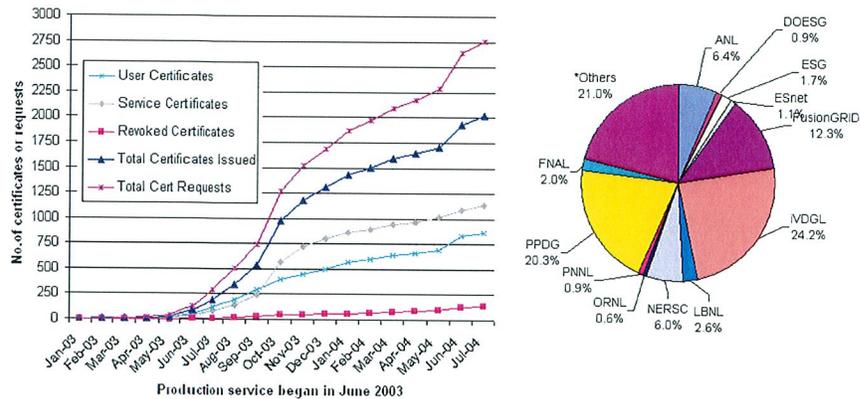
- Federation and Standards
  - DOEGrids supports 15 distinct “Registration Authorities”
    - Two are in progress for addition (LCG and EPA-NCC)
  - Regional peering – “Americas” PMA, TERENA, Asia-Pacific
  - Global Grid Forum
    - CAOPS (TG chair)
- PGP Key server

### New Initiatives:

- GIRAF – Grid Integrated RADIUS Authentication Fabric
- Fusion Grid PKI – support “myProxy” integration
- Remote Hardware Security Module operation
  - Response to ESnet’s challenge to provide redundant CA services
- Mozilla browser integration
- SIRS – Security Incident Response Services



# DOEGrids CA Usage Statistics



User Certificates	- 864	Total No. of Certificates - 2022
Service Certificates	- 1144	Total No. of Requests - 2759
Host/Other Certificates	- 14	
Internal certificates	- 31	

\* Report as of July 15, 2004

## Milestones

### Phase 1

- Advanced CA Services – see following
- Replace CA vendor -- see following -- \$50K
- Federation – Standards work (\$70K)
- Physical Security: (\$5K)
- FusionGrid CA - \$25K (servers)

### Phase 2

- GIRAF – Grid Integrated Radius Authentication Fabric -- phase 2 Pilot -- (\$205K)
- Add CA licenses - \$50K
- SIRS – Security Incident Response Services (\$50K)
- PGP Server replacement (\$20K)

Total: **\$475K + WebTrust**

## Advanced CA Services

### **WebTrust Certification of ESnet root CA - \$75K - \$250K**

Goal: The ESnet root CA in the default browser trusted CA list. WebTrust certification is a requirement by both major browser vendors (Microsoft Windows, Mozilla Foundation). This will greatly improve the usability and acceptance of our product in the DOE laboratory community.

Phase 1: Evaluation and compliance (\$20K-\$50K)

Phase 2: WebTrust Certification and vendor acceptance: \$55K-\$200K

Phase 3: Re-Certification as required: \$ yearly, unknown

### **AOL/Netscape CMS (CA software) - \$50K - \$100K**

Replace "End of Life" CA product with AOL/CMS.

Phase 1: Replace current licensed product: \$40K

Phase 2: Add additional licenses to support expansion: \$40K

Phase 3: Yearly maintenance: \$20K (estimate)

## Advanced Services

- **Developer – 2 contractors - \$500K**

Our community requires improved and new services, in both the PKI and the new OTP-Authentication Fabric initiatives.

**(1) OTP-Authentication Fabric**

Build out OTP support for Grid gateways

"Federation" middleware for Authentication Fabric

Integrate a SIPS (Site-Integrated Proxy CA) for ESnet deployment

Develop EAP middleware for other RADIUS clients, as needed

Develop API's for Authentication Fabric (RADIUS, OTP, and federation)

Support development of advanced Authentication Fabric protocols

**(2) Advanced CA and authorization-related services**

Extend the existing CA interfaces to real "Grid services"

Refine and extend the internal CA DBMS management to support VOMS projects

Develop a trustable "myProxy" credential store for wide area deployment

Integrate Shibboleth v 2.0

- **These two programs are independent but have many opportunities for synergy.**

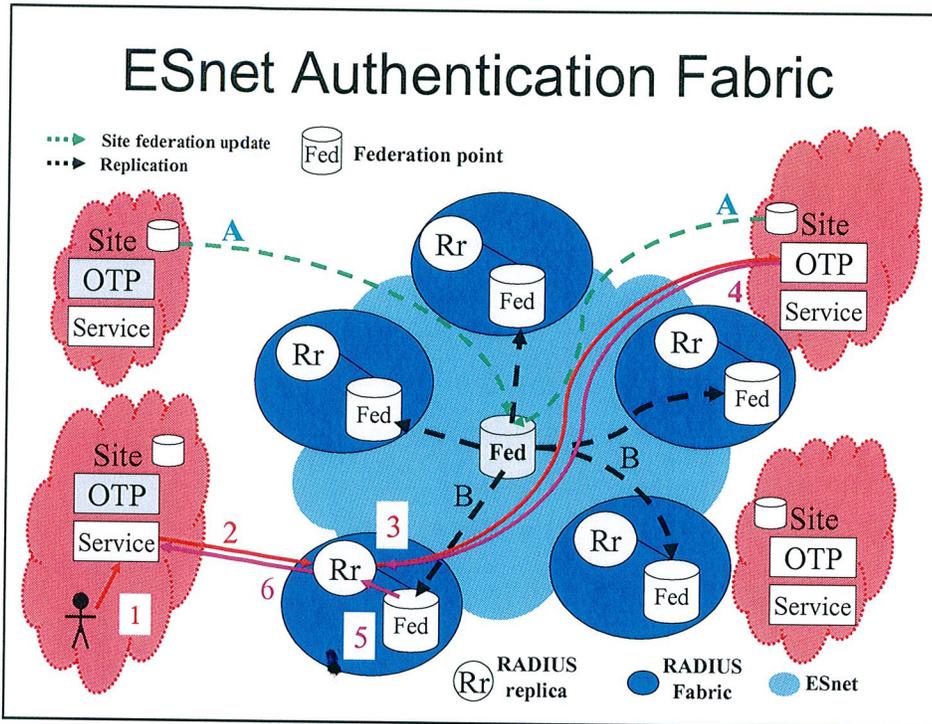
## The OTP - RADIUS Problem

Starting Point: One Time Password initiatives threaten to strangle Grids

- Globus response: Credential stores (“myProxy”) unlocked with OTP token, using RADIUS as *authentication service*
- Problems: multiple tokens, deployments, OTP interoperability limitations
- ESnet response: RADIUS Authentication Fabric\*

## Grid Integrated RADIUS Authentication Fabric (GIRAF)

- Redundant, highly available RADIUS replicas distributed around the core
- Login transactions routed to appropriate OTP backend no matter where you start
- Federation:
  - Consistent RADIUS “realm”\* across DOE Labs
  - Sites project appropriate data
  - Clients manage mapping(+)



## Authentication Fabric Issues

- RADIUS “back end” support
  - Insufficient proxy support in product
  - Mitigation: Schedule ; alternate product
- Rush to OTP solution by labs
  - Consequence: some duplication of tokens
  - Mitigation: live with it; add resources to speed the fabric development; custom interoperability agreements
- Federation\*
  - Need a working group / policy board
  - Mitigation: resource, direction
- Clients
  - Exploded beyond Grids, high expectations
  - Mitigation:



# ESnet Review



## NetNews

*Roberto R. Morelli*

*August 3, 2004*

Lawrence Berkeley National Laboratory



## ESnet NetNews

ISP-style news service that provides ESnet customers with direct access to individual news groups without the need and cost of sites running their own news service.

Service and server configured as:

- Dual CPU system running Dnews with 200+ GB mirrored new spool on a Gigabit Ethernet link.
- 66,000+ news groups via Qwest, processing 300,000 posts or 35GB per day, 24x7x365.
- News groups are made available after filtering and sanitizing. Currently 35,000 groups are available for direct access by customers.

## ESnet NetNews (2)

Sites using this service:

Production usage

- ORAU – Oak Ridge Associated University

Testing the service:

- LLNL – Lawrence Livermore National Lab.
- LBNL – Lawrence Berkeley National Lab.
- ORNL – Oak Ridge National Lab.

Interested in the service but not testing yet:

- LANL – Los Alamos National Lab.