



Designing a New Networking Environment for U.S. Research & Education

Steven Corbató corbato@internet2.edu

Director, Network Initiatives, Internet2 &
Center for High Performance Computing, Univ. of Utah

SURA Grid Application Workshop

Georgia State University

Atlanta

07 January 2005

One view of the U.S. landscape

- The packet network won, but now can it keep up?
 - Grid computing: a distinct view of the network as a *schedulable resource*
- The telecom/IP bust has created a ‘once-in-a-lifetime’ opportunity for previous non-players in facilities-based telecommunications
 - New R&E optical networking facilities are emerging on the regional and national scales
 - For the first time, higher ed networks are being built with significant up-front capital investment and thus lower ongoing costs including for incremental expansion
- Active investigation of new hybrid architectures is underway

Bandwidth-storage normalization

- A 10-Gbps wavelength is capable of transmitting roughly...
 - 100 TeraBytes per day
 - 35 PetaBytes per year

Grid-supporting network architectures

- **Dedicated capabilities ('lightpaths')**
 - Physical channels - λ 's
 - Subchannels – Gigabit Ethernet
 - Virtual channels – MPLS tunnels
- **High performance shared packet networks**
 - Excess bandwidth and absence of BW controls
 - QoS provisioning?
- **Both approaches require significant attention to end-to-end performance**

Topics for today

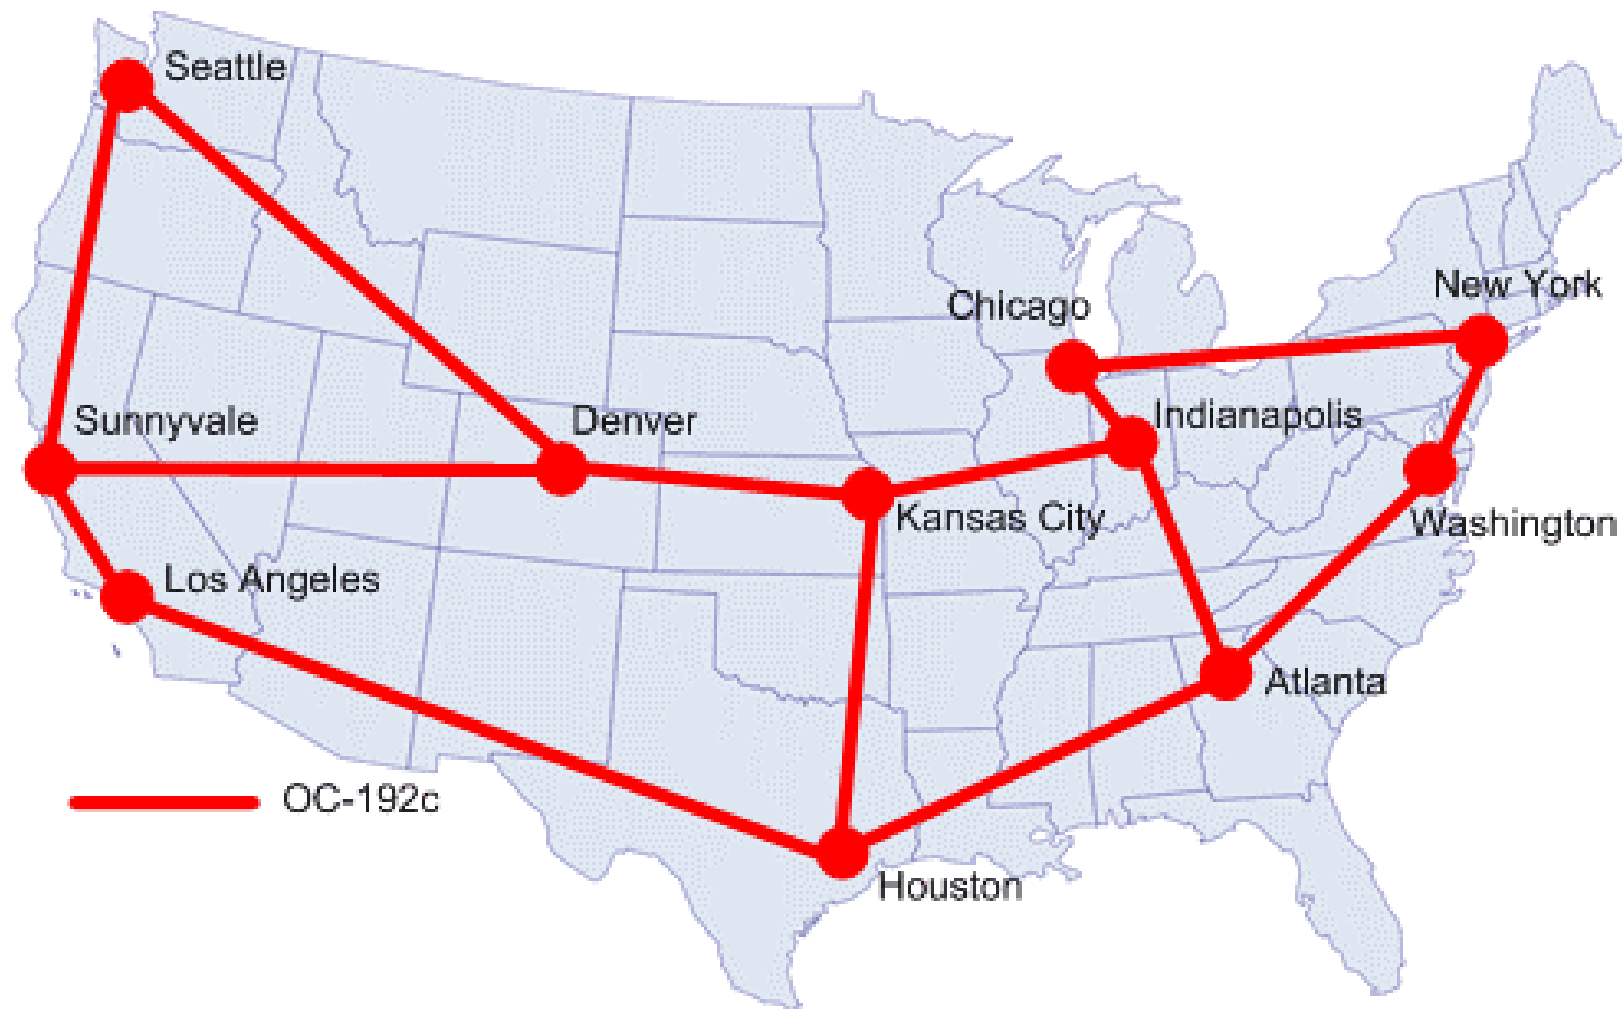
- High performance packet infrastructure
 - Abilene Network
- Regional Optical Networks (RONs)
 - FiberCo case study
- National optical facility
 - National LambdaRail (NLR)
- Future architectures
 - Hybrid Optical & Packet Infrastructure – HOPI
- Exchange point innovations
 - MAN LAN – New York City



Abilene Network

07 January 2005

Abilene Network – second generation



Abilene timeline

- **Apr 1998 Network announced**
 - Cisco Systems, Indiana Univ., Nortel Networks, and Qwest Communications initial partnership led by Internet2
 - 2.5-Gbps national backbone (OC-48c SONET)
- **Jan 1999 Network went into production**
- **Second generation network upgrade**
 - Oct 2001 Qwest MoU (DWDM+SONET) extension (5 years)
 - Apr 2002 Routers from Juniper Networks added
 - Dec 2003 10-Gbps upgrade complete
 - **Oct 2004 Transport agreement extended by at least one year**
- **Oct 2007 Transport MoU with Qwest ends**
 - The time frame for both next generation architecture finalization & decision on transport partner(s) is <15 months from now – **early spring 2005.**

Abilene scale

January 2005

- IPv4/v6-over-DWDM (OC-192c) backbone

- 44 direct connections (OC-3c → 10 GigE)
 - 2 (soon 3 – SoX!) 10-GigE connections (10 Gbps)
 - 7 OC-48c connections (2.5 Gbps)
 - 2 Gigabit Ethernet connections (1 Gbps)
 - 24 connections at OC-12c (622 Mbps) or higher

- 230+ participants – research universities & labs
 - All 50 states, District of Columbia & Puerto Rico

- Expanded access
 - 113 sponsored participants
 - 34 state education networks

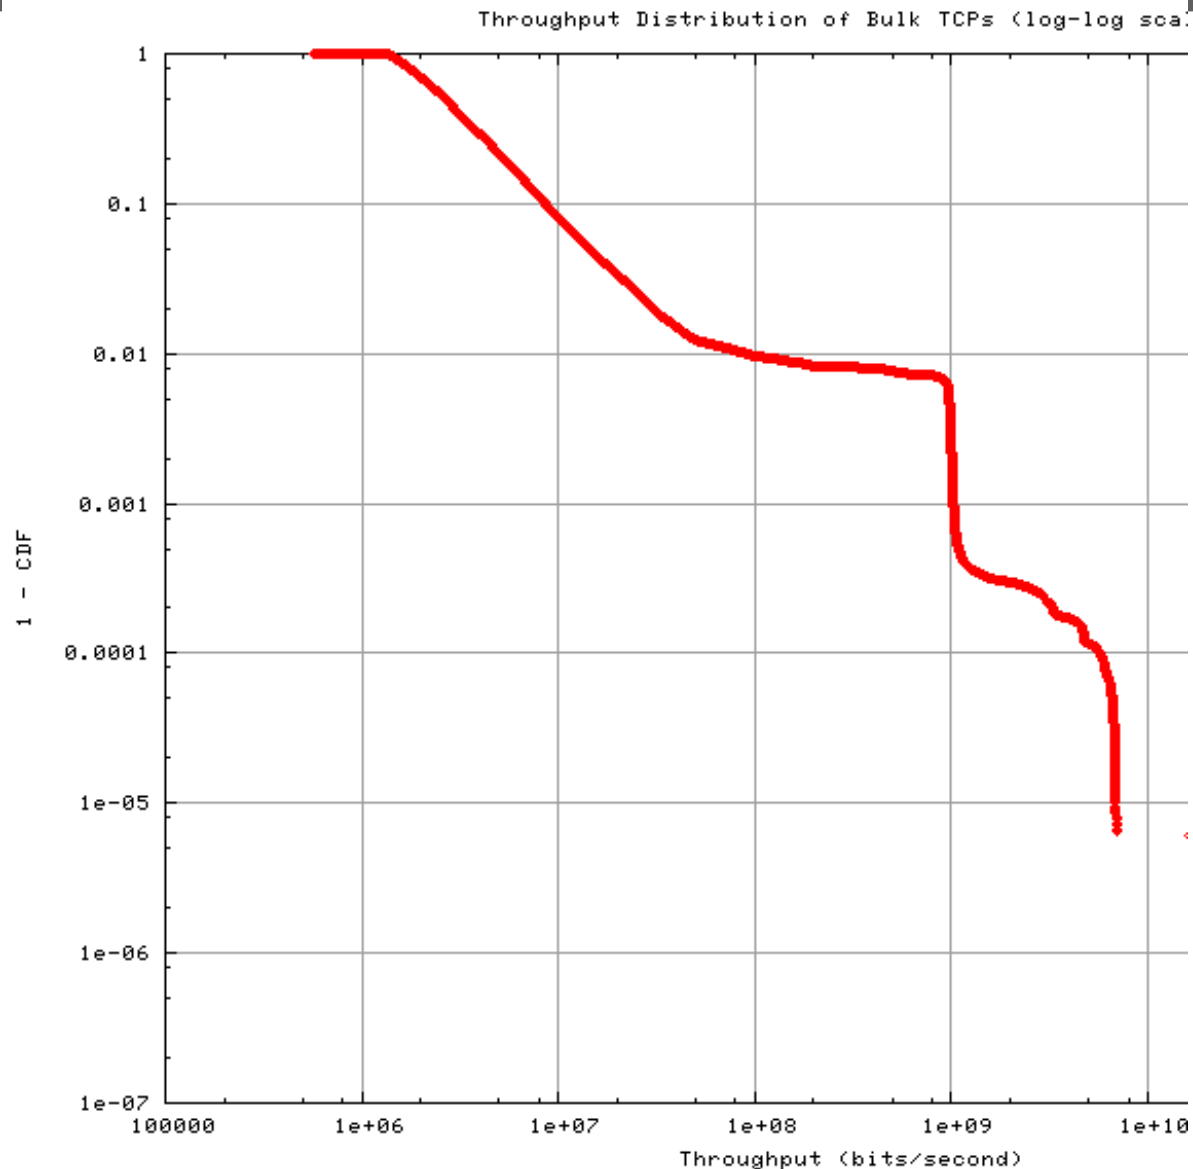
Abilene's distinguishing features

- Native advanced services – multicast & IPv6
- Ability to support large individual flows
 - Regular, routine testing: hourly 980+ Mbps TCP flows
 - Supporting multiple Internet2 Land Speed Records
 - Latest multi-stream TCP flow: 6.6 Gbps
- Home for community's advanced Internet initiatives
 - E.g., middleware, performance
- Cost recovery model
 - Pricing scales roughly logarithmically with bandwidth
 - Aim to is to encourage utilization and experimentation
- International R&E collaboration and peering
- Open measurement stance

Abilene Observatory

- A project designed to support the computer science network research and advanced engineering communities
- Two components
 - *In situ* experimentation
 - Access to comprehensive set of network performance data
- Hosted Projects
 - PlanetLab (Berkeley/Princeton/Intel Research/NSF)
 - AMP Project (SDSC/NSF)
- Access to Network Performance data
 - Objective is to maintain time-correlated data archive
 - Multiple time-corrected data views – traffic flows, passive measurements, routing data, SNMP and syslog data
- <http://abilene.internet2.edu/observatory/>

End-to-end performance: a persistent challenge



Bulk flow distribution
(aggregate payload >
10 MBytes)

Median: 2.7 Mbps

90%: 7.9 Mbps

99%: 38 Mbps


**Source: Stas Shalunov
(Internet2)**

Enabling high performance

- A number of problems continue to inhibit optimum e2e throughput
 - Application tuning
 - Host IP stack tuning
 - Host Ethernet interface duplex configuration
 - LAN capacity and operational issues
- In fact, optical networking is a ‘reversion’ in some respects to a more analog approach
- Routine passive measurement (e.g, PiPES) and active testing are essential components of the solution
- <http://e2epi.internet2.edu>

Very Long Baseline Interferometry (VLBI)

- Multiple antennae located at continental distance
- Each antennae collects data from the sky at speeds of 1-10 Gbps
- Transmit all data dynamically over Abilene (previously recorded data to tape) to correlation facility
- Correlation facility must process information from all antennas in real time (several computational challenges involved here)



MIT Haystack Observatory

Sources: Alan Whitney & David Lapsley (MIT/Lincoln Labs); Charles Yun (Internet2)



Regional Optical Networks (RONs)

07 January 2005

Underlying hypothesis

- The fundamental nature of regional networking is changing
 - The *GigaPoP* model based on *provisioned, high-capacity services* steadily is being replaced – on the *metro and regional scales*
- A model of *facility-based networking built with owned assets* – Regional Optical Networks (RONs) – has emerged
 - Notably, this change *increases* the importance of regional networks in the traditional *three-level hierarchy* of U.S. R&E advanced networking

Distance scales for U.S. optical networking

	<i>Distance scale (km)</i>	<i>Examples</i>	<i>Equipment</i>
Metro	< 60	Univ. Wash (Sea), USC/ISI(LA), MAX(DC/MD/VA)	Dark fiber & end terminals
State/ Regional	< 500	I-WIRE (IL), I-LIGHT (IN), CENIC ONI (CA), LONI (LA)	Add OO Amplifiers (or optical TDM)
Extended Regional/ National	> 500	TeraGrid 2 nd Gen Abilene, NLR	Add OEO regenerators & O&M \$'s

- Arkansas
- California (CALREN)
- Colorado (FRGP/BRAN)
- Connecticut (Conn. Education Network)
- **Florida** (Florida LambdaRail)
- Georgia (Southern Light Rail)
- **Indiana** (I-LIGHT)
- Illinois (I-WIRE)
- Louisiana (LONI)
- Maryland, D.C. & northern Virginia (MAX)
- **Michigan** (MiLR)
- Minnesota
- Minnesota
- New England region (NEREN)
- New York (NYSERNet, Cornell)
- **North Carolina** (NC LambdaRail)
- Ohio (Third Frontier Network)
- Oklahoma (OneNet)
- Oregon
- Pacific Northwest (Lariat – NIH BRIN, **PNNL**)
- Rhode Island (OSHEAN)
- SURA Crossroads (southeastern U.S.)
- Tennessee (OneTN)
- Texas (LEARN)
- Virginia (MATP)
- **Wyoming**

FiberCo

- Dark fiber holding company
 - Operates on behalf of U.S. higher education and affiliates – the Internet2 membership
 - Patterned on success of Quilt commodity Internet project
 - Assignment vehicle for the regionals and NLR
 - **Fundamentally, a dark fiber *market maker* for R&E**
- Project designed to *support* optical initiatives
 - Regional (RONs)
 - National (NLR)
- Not an operational entity
 - Does not light any of its fiber
- Concept was a spin-off from NLR governance discussions
 - Internet2 took responsibility for organizational formation
 - First acquisition of dark fiber through Level 3
 - 2,600 route miles (fiber bank) – 3/2003
 - Now has assigned over 5,600 route-miles to NLR and RONs
 - Subsequent strong working relationship with WilTel
- Complementary to SURA/AT&T dark fiber donation

Dark fiber: gauging community-wide progress

- Aggregate dark fiber assets acquired by U.S. R&E optical initiatives (segment-miles)
 - CENIC (for CalREN & NLR) 6,200
 - FiberCo (via Level 3 for NLR & RONS) 5,660
 - SURA (via AT&T) 6,000
 - Plus 2,000 route-miles for research
 - NLR Phase 2 (WilTel & Qwest) 4,000
 - OARnet 1,600
 - ORNL (via Qwest) 900
 - NEREN 670
 - Other projects (IN,IL,MI,OR, ...) 2,200+
- **Total (conservative estimate) 27,230+**
 - Over 55% of these assets are now outside NLR
 - NLR will hold ~11,250 route-miles

Starting a RON ... in stages

1. Convene enthusiastic/visionary regional partners
2. Identify science and other research drivers
3. Assemble a technical working group
4. Develop governance & capital approaches and preliminary business plan
5. Study dark fiber availability and procure
6. Select and procure optronics kit
7. Refine business plan (i.e., λ pricing/cost-recovery model)
8. Focus on means to extend new capabilities to the researchers on campuses
9. Learn how to operate and maintain the system
10. Install and commission plant
11. At last, provision λ 's and other services!

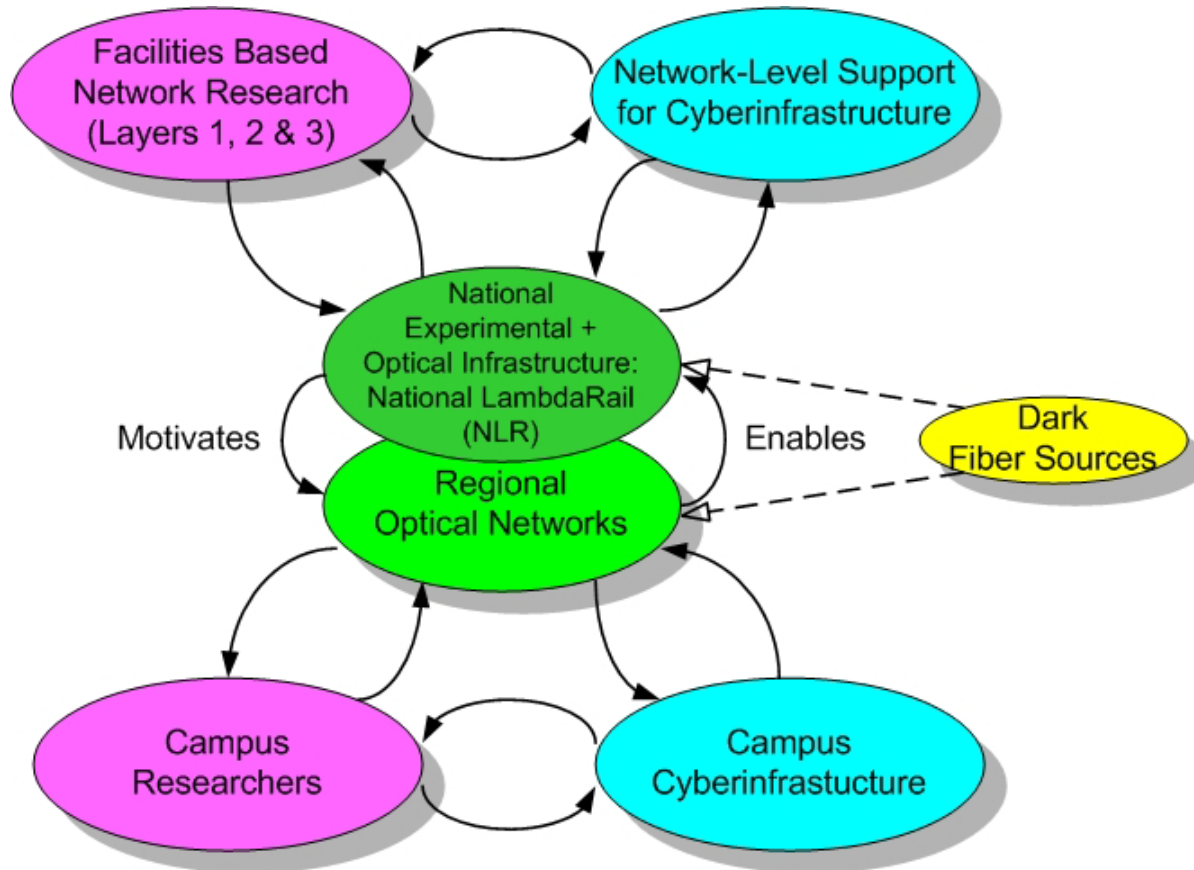


National optical networking and network research

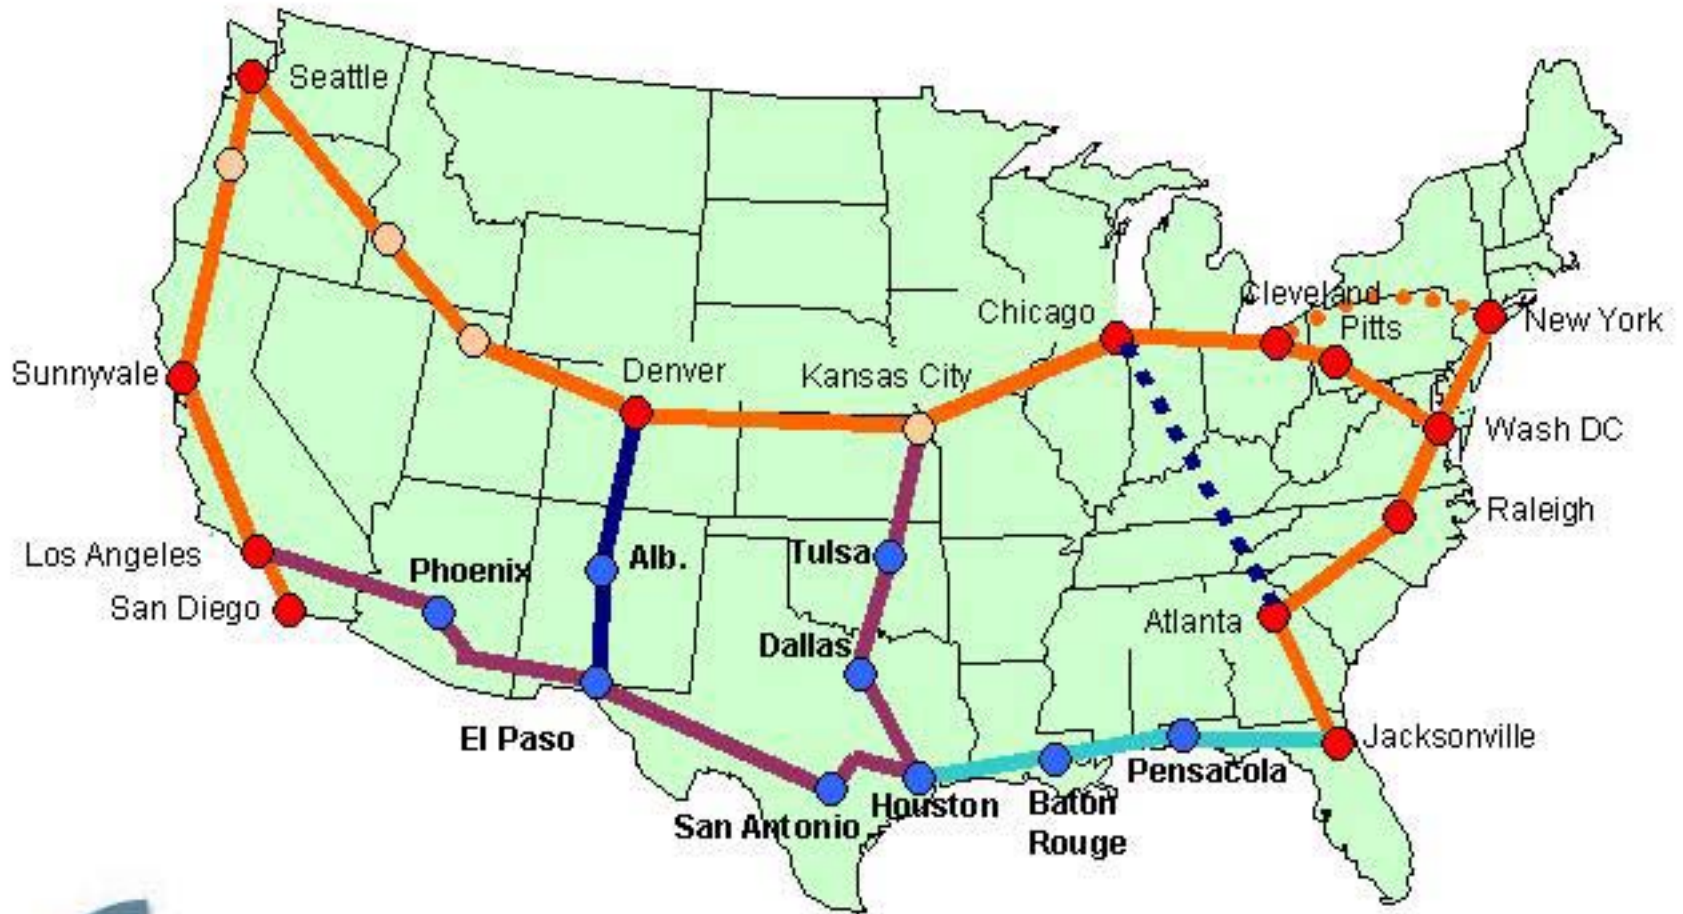
National LambdaRail (NLR)

07 January 2005

National LambdaRail - Motivations



National LambdaRail – Proposed Final Topology





NLR

distinguishing features

- Largest higher-education owned & managed optical networking & research facility
 - Over 10,000 route-miles of underlying dark fiber
 - Four 10-Gbps λ 's provisioned at outset
 - One allocated to Internet2
- First & foremost, an experimental facility for research
 - Optical, switching & experimental IP capabilities (layers 1, 2 & 3)
- Use of high speed Ethernet (10 Gbps) for wide area transport
- Sparse backbone topology
 - Each participant commits \$5M over 5 years and assumes responsibility for a regional node



Participating organizations

- CENIC (California)
- Pacific Northwest Gigapop
- Front Range and Intermountain Gigapops (CO, UT & WY)
- CIC (Midwest)
- Cornell (NY & New Eng)
- MATP (Virginia)
- Duke (North Carolina)
- Georgia Tech (Georgia)
- Florida LambdaRail
- Louisiana Board of Regents
- LEARN (Texas)
- Oklahoma State Board of Regents
- Pittsburgh Supercomputing Center and Univ. of Pittsburgh
- University of New Mexico
- Internet2
- Cisco Systems
- Affiliated organizations:
 - Case Western Reserve Univ.
 - SURA
 - Oak Ridge National Laboratory



Next steps for optical networking development

07 January 2005

Global Lambda Integrated Facility World Map – December 2004

Predicted international Research & Education Network bandwidth, to be made available for scheduled application and middleware research experiments by December 2004.



HOPi Project - Summary

- In the near future, we will see a richer set of capabilities available to network designers and end users
 - Core IP packet switched networks
 - A set of optically switched waves available for dynamic provisioning
- Fundamental Question: How will the core Internet architecture evolve?
- Examine a **hybrid** of shared IP packet switching and dynamically provisioned optical lambdas
- HOPi Project – Hybrid Optical and Packet Infrastructure
 - A white paper describing a testbed to model the above infrastructure is posted <http://hopi.internet2.edu>
 - Implement testbed over the next year
 - Coordinate and experiment with other similar projects
 - Design Team consisting of U.S. and international experiments

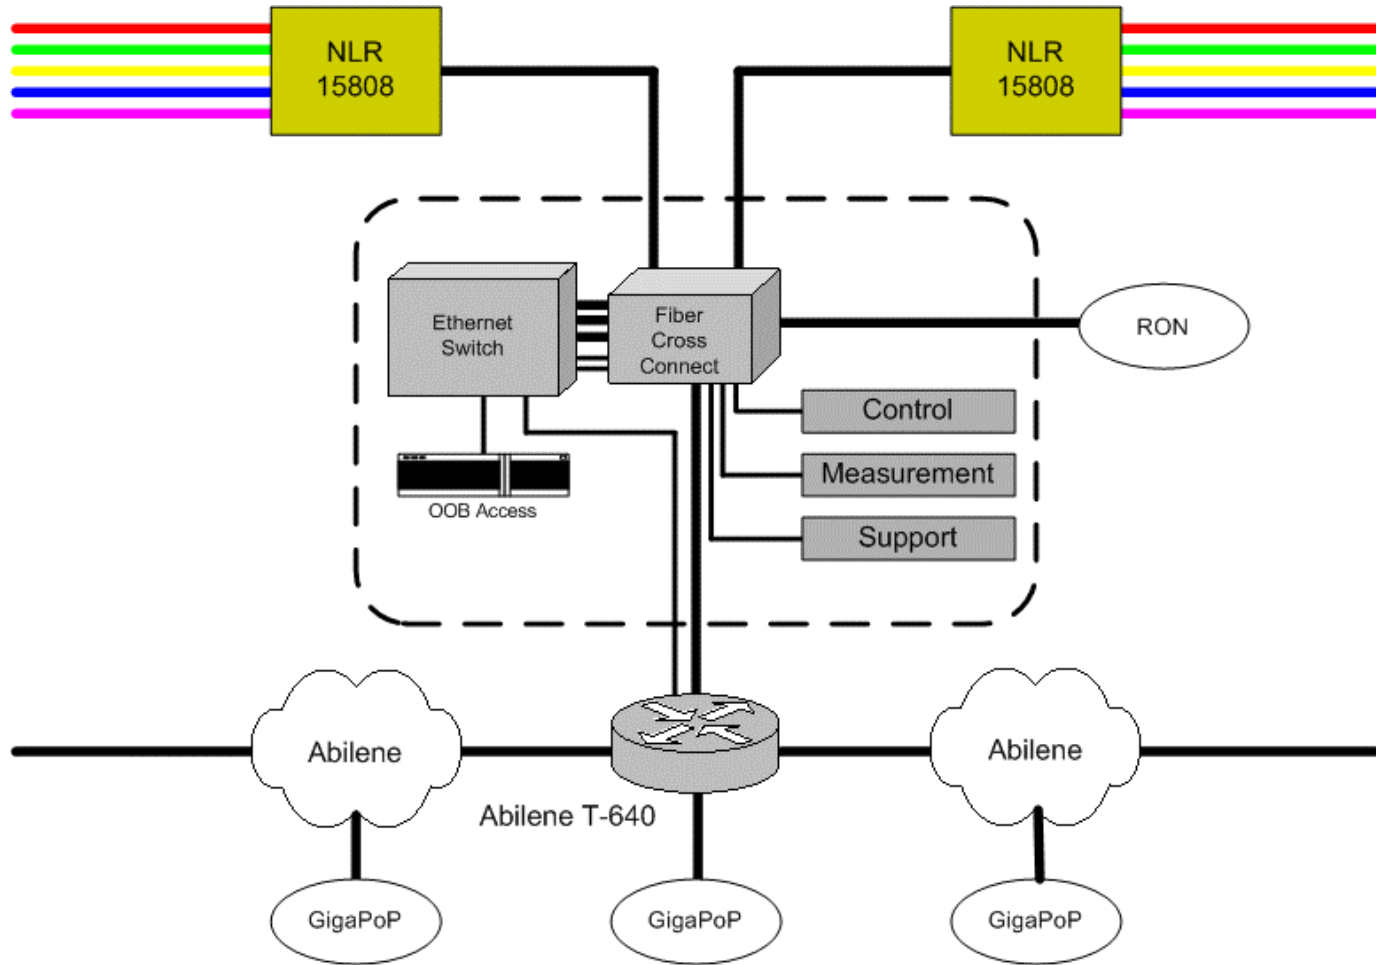
HOPI Resources

- The Abilene Network – MPLS tunnels and the 10-Gbps packet switched network
- Internet2's 10-Gbps λ on the NLR national footprint
- MAN LAN experimental facility in New York
 - IEEAF(Tyco Telecom) 10-Gbps lambda between NYC - Amsterdam
- Collaboration with the Regional Optical Networks (RONs) and other related advanced efforts (GLIF, DRAGON, SURFNet, GEANT-2)

HOPI Project

- Problems to understand
 - Goal is to look at architecture
 - Temporal degree of dynamic provisioning
 - Temporal duration of dynamic paths and requirement for scheduling
 - Topological extent of deterministic provisioning
 - Examine backbone, RON, campus hierarchy – how will a RON interface with the core network?
 - Understand connectivity to other infrastructures – for example, international or federal networks?
 - Network operations, management and measurement across administrative domains?

HOPi Node





Manhattan Landing (MAN LAN) Exchange Point

07 January 2005

U.S. R&E exchange points

- Star Light (Chicago)
- Pacific Wave (Seattle & LA)
- AMPATH (Miami)
- NGIX-East/MAX (DC/College Park MD)
- NGIX-West (SF Bay Area)
- **MAN LAN (New York City)**
- Current trend is to develop geographically distributed exchange points on both coasts
 - Pacific Wave (Seattle-soon Bay Area-LA)
 - Atlantic Wave (New York-Washington DC-Atlanta-Miami-Sao Paulo)

Manhattan Landing (MAN LAN)

- MAN LAN originally conceived as a high performance exchange point to facilitate peering between US and International Research and Education Networks
 - Facilitate peering between federal and international networks
 - Original design was layer 2, an Ethernet switch.

- MAN LAN was formed through a partnership with Indiana University, NYSERNet, Internet2, and the IEEAF
 - Indiana University provides NOC and Engineering services
 - NYSERNet provides co-location, hands and eyes, and interconnection support
 - IEEAF has provided donated 10-Gbps wavelength and OC-12c circuits to Amsterdam

- Located in 32 Avenue of the Americas in New York City
 - Collocated in the NYSERNet facility adjacent to the fiber meet me room – cross-connects simple to facilitate
 - Many other carriers maintain presences in 32 AoA
 - NYSERNet has co-location space available

MAN LAN service models

■ Production

- Layer-2 interconnection/peering for IPv4 and IPv6
- Layer-1 optical interconnection

■ Experimental facility

- Layer-1 optical interconnection
 - Partitioned from production service
 - Adjacent to one of first five HOPI nodes (linking Abilene IP and I2's λ over NLR)

Conclusions

- Abilene Network supports most of U.S. higher ed's collaboration needs
 - Observatory showing demonstrable impact in research facilitation
 - Network utilization growing; network capable of large flows
 - Next generation architecture and transport plan envisioned within 15 months
- NLR and the RONs are providing new options for U.S. advanced networking
 - RON development is a critical activity for research competitiveness
- HOPI and related projects are exploring a unified, hybrid architecture of packets and circuits for the near future

For more information...

- <http://abilene.internet2.edu>
- <http://abilene.internet2.edu/observatory>
- <http://ipv6.internet2.edu>
- [**http://e2epi.internet2.edu**](http://e2epi.internet2.edu)
- <http://www.fiberco.org>
- <http://networks.internet2.edu/manlan>
- <http://hopi.internet2.edu>
- <http://www.glif.is>

A large, stylized white number '2' is positioned vertically, overlapping the word 'INTERNET'. The '2' has a thick, rounded top and a long, sweeping tail that curves to the right.

INTERNET®

www.internet2.edu