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

Steve Corbató

Director, Network Infrastructure

Industry Strategy Council

Washington DC

16 April 2004

Outline

- National packet (IP) networking
  - Abilene Network
- New optical networking facilities
  - Regional projects
    - FiberCo supporting project
  - National initiative
    - National LambdaRail (NLR)
- Developing a hybrid architecture
  - Hybrid Optical and Packet Infrastructure (HOPI)

Abilene's 2nd generation, 10-Gbps backbone

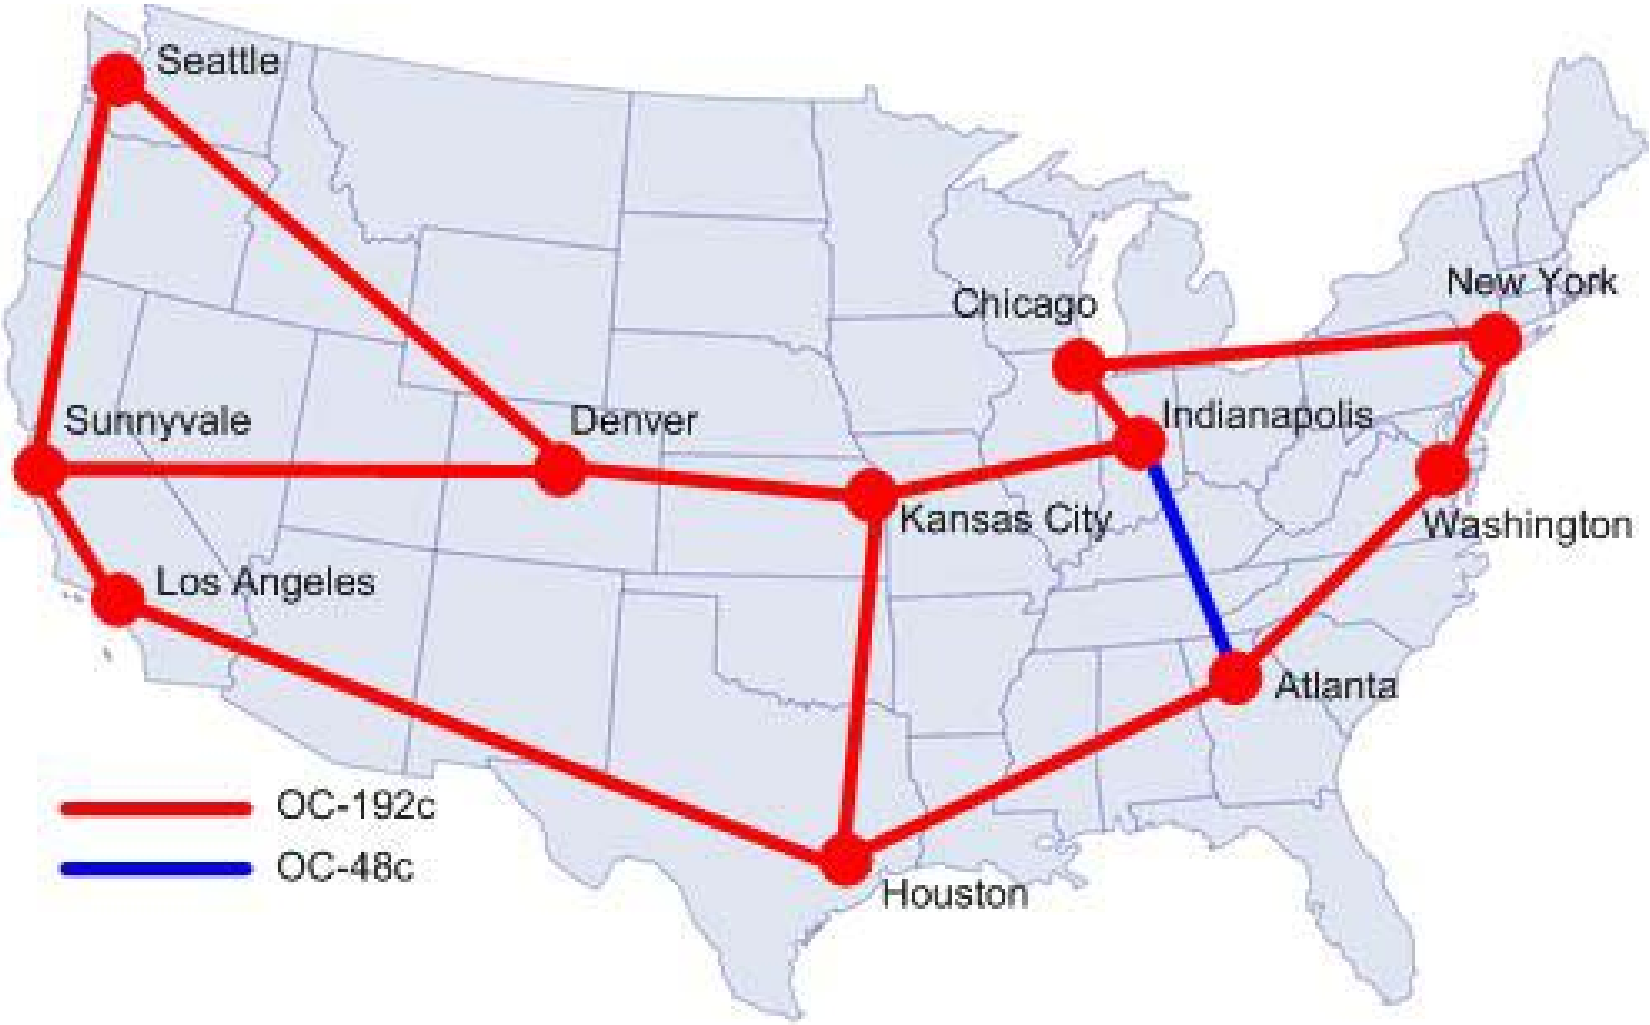

Abilene Participation
April 2004

- IP-over-DWDM (OC-192c) backbone
- 44 direct connections (OC-3c → 10 GigE)
  - 2 10-GigE connections (10 Gbps)
  - 6 OC-48c connections (2.5 Gbps)
  - 2 Gigabit Ethernet connections (1 Gbps)
  - 23 connections at OC-12c (622 Mbps) or higher
  - Connection fees lowered in 2004 to encourage upgrades
- 228 participants – research universities & laboratories
  - All 50 states, District of Columbia & Puerto Rico
  - Johns Hopkins as of today!
- Expanded access
  - 104 sponsored participants
  - 33 state education networks

Abilene International Peering

Abilene International Network Peers

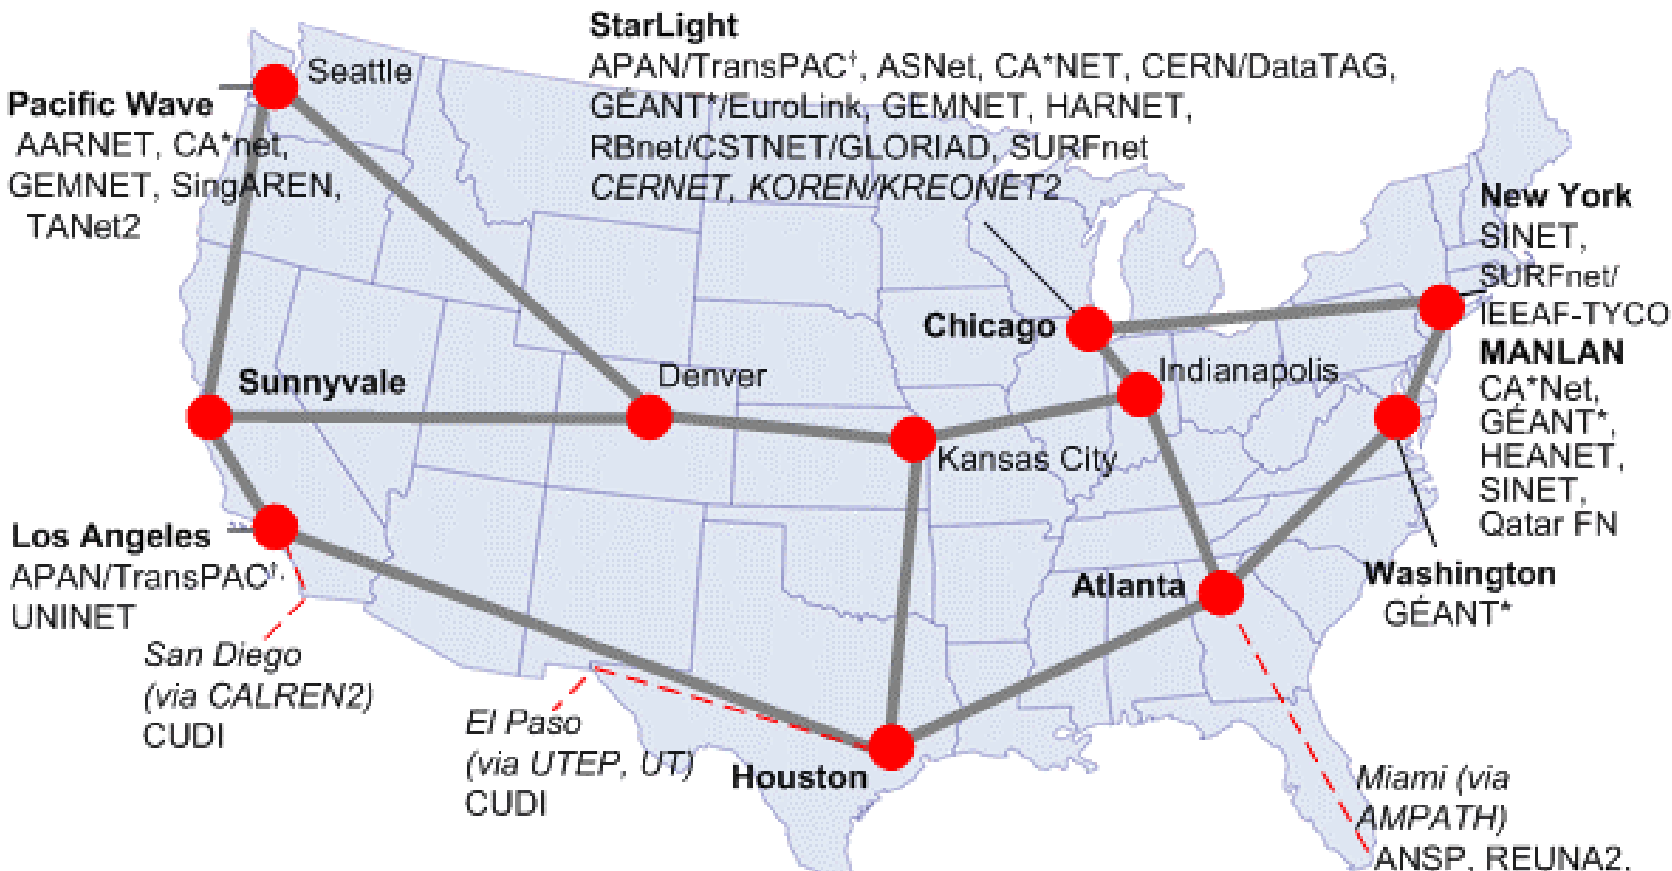

* *via GÉANT:* ACOnet, BELNET, CARNet, CESNET, CYNET, Forskningsnettet, EENet, Funet, Renater, G-WIN, GRNET, HUNGARNET, Rhnet, HEAnet, IUCC, GARR, LANET, LITNET, RESTENA, Univ. Malta, SURFnet, UNINETT, POL34, RCTS2, RoEduNet, RBnet, SANET, ARNES, RedIRIS, SUNET, SWITCH, JANET, ULAKBYM, CERN

† *via APAN/TransPAC:* WIDE/JGN, IMnet, CERNet/CSTnet/NSFCNET, KOREN/KREONET2, PREGINET, SingAREN, TANET2, ThaiSARN, WIDE (v6)

Abilene Focus Areas - 2004

- Support of high-throughput flows (multi-Gbps)
  - Collaboration with End-to-End Performance Initiative
  - Ensuring that large flows are the standard across the Internet2 Infrastructure
- Provisioning of dedicated capabilities (HOPI)
- IPv6
  - Furthering its end-to-end deployment in the I2 environment
  - Moon-v6 testbed (with DoD)
- Security
  - Collaboration with REN-ISAC and pending Cyber Trust effort
- Network resiliency
  - Fast network restoration in the face of outages
- Abilene Observatory

Abilene Observatory

- A project designed to support the computer science network research and advanced network engineering communities
  - 12 computer science research projects currently using Abilene
- Two components
  - *In situ* experimentation and demonstration
  - Access to comprehensive set of Abilene Network 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
- More information: abilene.internet2.edu/observatory/

Abilene's Annual Connection Fee Model

	Original	*2003*	*2004*
OC-3c (155 Mbps)	$110k (1998)	($110k)	($110k)
OC-12c (622 Mpbs)	$320k (1998)	$270k	$240k
Gig E (1 Gbps)	$325k (2001)	$325k	$280k
OC-48c (2.5 Gbps)	$495k (2000)	$430k	$360k
10 Gbps (SONET/ Ethernet)	$490k (2003)	$490k	$480k

Abilene timeline

- Apr 1998 Qwest & Cisco partnerships
- Sep 2001 Qwest MoU extension (5 years)
- Apr 2002 Juniper router selection
- Oct 2006 Transport MoU with Qwest ends

- Past experience suggests that the time frame for the next generation architecture definition & transport decision would be during the spring of 2005

Optical Network Development

16 April 2004

Underlying hypothesis

- The fundamental nature of *regional networking* is changing
  - The *GigaPoP* model based on *provisioned, high-capacity services* is being replaced steadily – on the *metro and regional scales*
- A model of *facility-based networking often 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

Leading & Emerging *Regional Optical Networks*

- 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
- Minnesota
- New York + New England region (NEREN)
- North Carolina (NC LambdaRail)
- Ohio (Third Frontier Network)
- Oregon
- Pacific Northwest (Lariat – supported by NIH)
- Rhode Island (OSHEAN)
- SURA Crossroads (southeastern U.S.)
- Texas (LEARN)
- Utah
- Virginia (MATP)
- Wisconsin

FiberCo model

- Dark fiber holding company
  - Operates on behalf of U.S. higher education and their R&E affiliates – the Internet2 membership
  - Assignment vehicle for 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
  - Will not light any of its fiber
- First acquisition of dark fiber through Level 3
  - 2,600 route miles – March, 2003

FiberCo assignment progress (route-mileage)

	Level 3	Genuity	Total
Completed assignments			
Indiana U	241	-	241
Assignments in progress			
NLR, Inc.	1,705	-	1,705
Florida/FLR	900	592	1,492
PNWGP (U Wyoming)	178	-	178
Michigan	963	-	963
Internet2	321	-	321
TOTAL	4,308	592	4,900

Gauging community-wide progress with dark fiber

- Aggregate dark fiber held for and assigned by U.S. R&E optical initiatives (segment-miles):
  - CENIC 6,200
  - **FiberCo** **4,900**
  - SURA 6,000
    - (plus 2,000 route-miles for research usage)
  - Oak Ridge National Lab. 900
  - Ohio 1,600
  - Other state projects (IN,IL,MI,OR, etc.) 1,500+
- **Total (conservative estimate) 21,000+**
- Pending procurements (TX, NY, NE) 1,700+

A new model for carrier support of advanced networking?

- Fundamental premise:
  - Total system ownership, control, and responsibility are neither the most cost-effective nor the optimum solution for higher education's optical networks
- Emerging optical networks offer newopportunities for carriers to provide services with potentially higher ROI
  - Higher education is willing to assume much of the initial and ongoing capital costs in this model

What higher ed really seeks at the minimum?

- Dark fiber ownership
  - Co-location capabilities
  - Open fiber interconnectivity
- Significant equity stake in the optronics
  - Ability to provision new waves at incremental cost
  - Full management visibility
  - Ability to switch waves
- Many opportunities for carrier services here
  - Fiber O&M
  - Colocation & power
  - Next gen optronics anchor tenant
  - Hands & eyes support
  - Higher level support services
  - Significant ancillary business (off-net waves, IP)

National optical networking and network research

National LambdaRail (NLR)

16 April 2004

NLR National Fiber Network

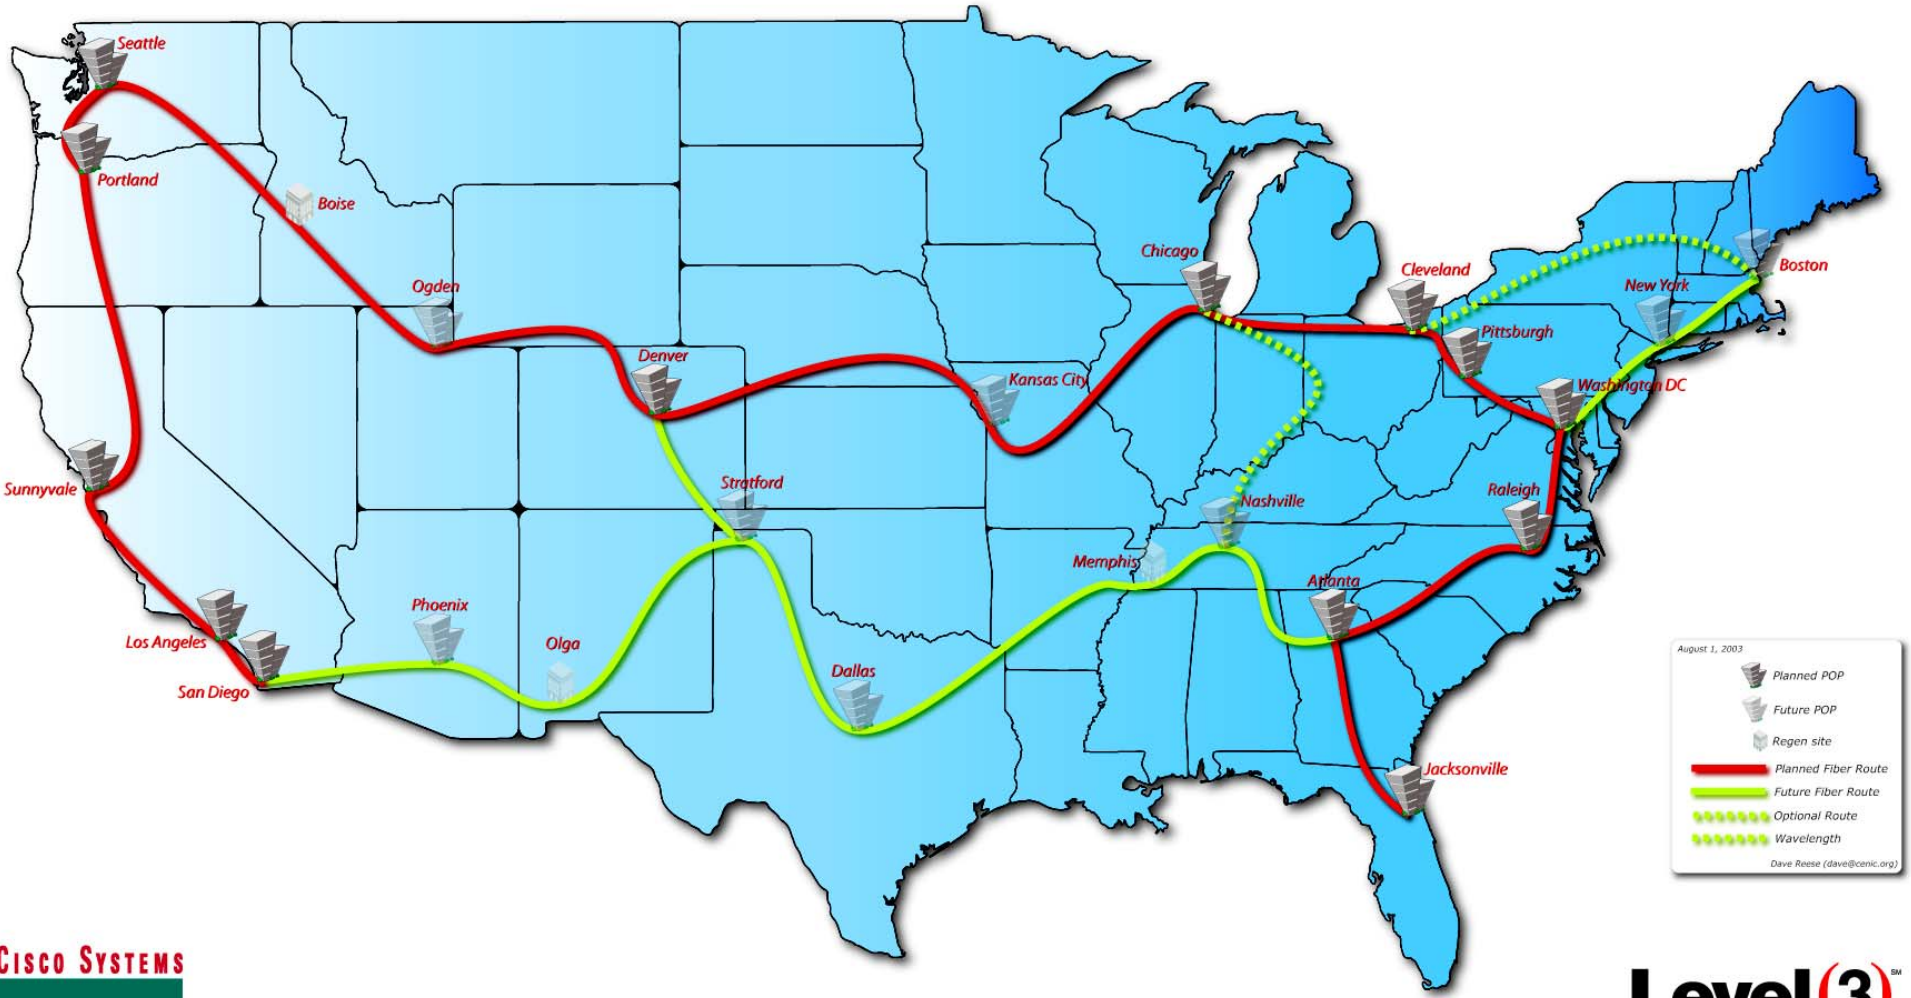

Cisco Systems

Level(3) COMMUNICATIONS

Initial Coordinating Participants

cenic Corporation for Education Network Initiatives in California

FLR

Georgia Institute of Technology

INTERNET2

IWIRE

NCAR

PACIFIC NORTHWEST GIGAPOP

PITTSBURGH SUPERCOMPUTING CENTER

Texas Consortium

Virginia Tech

NLR distinguishing features

- Largest higher-ed owned & managed optical networking & research facility in the world
  - Over 10,000 route-miles of dark fiber
  - Four 10-Gbps λ's provisioned at outset
    - **One λ on national footprint 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 node

NLR membership

- CENIC
- Pacific Northwest Gigapop
- Pittsburgh SC/CMU
- Duke Univ./NCLR
- Va. Tech/MATP
- Cisco Systems
- Internet2
- Florida LambdaRail
- Georgia Tech/SLR
- CIC

Pending:

- LEARN (Texas)
- Cornell (New York)
- NCAR (Colorado/Utah)
- LONI (Louisiana)
- SURA

- Collaboration
  - Oak Ridge Nat'l Lab

Planned NLR Capabilities

- Point-to-point wavelengths
  - 10GE LAN PHY, OC-192 using Cisco 15808 long haul and extended long haul and Cisco 15454 extended metro DWDM systems
- Switched Ethernet
  - Using Cisco 6500 switches
- Experimental routed IP
  - Using Cisco 12400 routers
- Dark fiber for optical layer research
- Traditional NOC services and Experiments Support Center”
  - Instrumentation, measurement, configuration management, tool development

Recent development

- NLR, Inc. granted federal 501(c)(3) tax exempt status – 2/2004
- Board at work on clear statement of NLR goals
- Dave Farber now serving as *pro bono* Chief Scientist
  - Research committee under development
- Debbie Montano now director of development and operations
  - Engineering group and NLR NOCs architecting Layer-2 and –3 capabilities

Next steps for optical networking development

16 April 2004

Considering Future Architectures

- Examine new infrastructures for the future
  - 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 wavelengths available for dynamic provisioning.
  - **Fundamental question: How will the core Internet infrastructure architecture evolve?**
    - **Campus, regional, national, and international components**
  - Thus, we are examining a *hybrid model* of shared IP packet switched networks and dynamically provisioned optical circuits
- HOPI effort – Hybrid Optical and Packet Infrastructure

HOPI Project Goals

- Examine new architectures for the future
  - Model a set of packet and optical capabilities
- Immediate goals
  - Develop a white paper by the April Member Meeting describing a testbed plan (DONE)
    - Request comments on the testbed plan
  - **Engage high-level, corporate technical participation**
  - Implement HOPI testbed over the next year
  - Coordinate and experiment with other similar projects
  - Coordinate current projects (e.g., GLIF)
- Have formed a design team to coordinate and recommend activities over the next year

HOPI Project Design Team

- Linda Winkler, Argonne (co-chair)
- Rick Summerhill, Internet2 (co-chair)
- Cees de Laat, U of Amsterdam
- Rene Hatem, CANARIE
- Mark Johnson, MCNC
- Tom Lehman, USC/ISI
- Peter O'Neil, NCAR
- Bill Owens, NYSERNet
- Philip Papadopoulos, UCSD
- Sylvain Ravot, Caltech/CERN
- David Richardson, U Washington
- Chris Robb, Indiana U
- Jerry Sobieski, U Maryland
- Steven Wallace, Indiana U
- Bill Wing, Oak Ridge
- Supporting Staff: Guy Almes, Heather Boyles, Steve Corbató, Chris Heermann, Christian Todorov, Matt Zekauskas (Internet2)

HOPI's raw materials

- The Abilene Network
  - 10-Gbps IPv4/IPv6 + MPLS tunnels
- Internet2's 10-Gbps λ over NLR
- MAN LAN exchange point – New York City
  - Ethernet Switch – layer 2 switching
  - TDM Switch – layer 1 switching
  - Venue for international collaboration
    - SURFnet, CANARIE, ESnet, SINET, Abilene
- IEEAF/Tyco 10-Gbps λ - NYC-Amsterdam
- Collaborations with Regional Optical Networks (RONs) and other related efforts (e.g., GLIF, DRAGON)

Abilene/NLR Topology

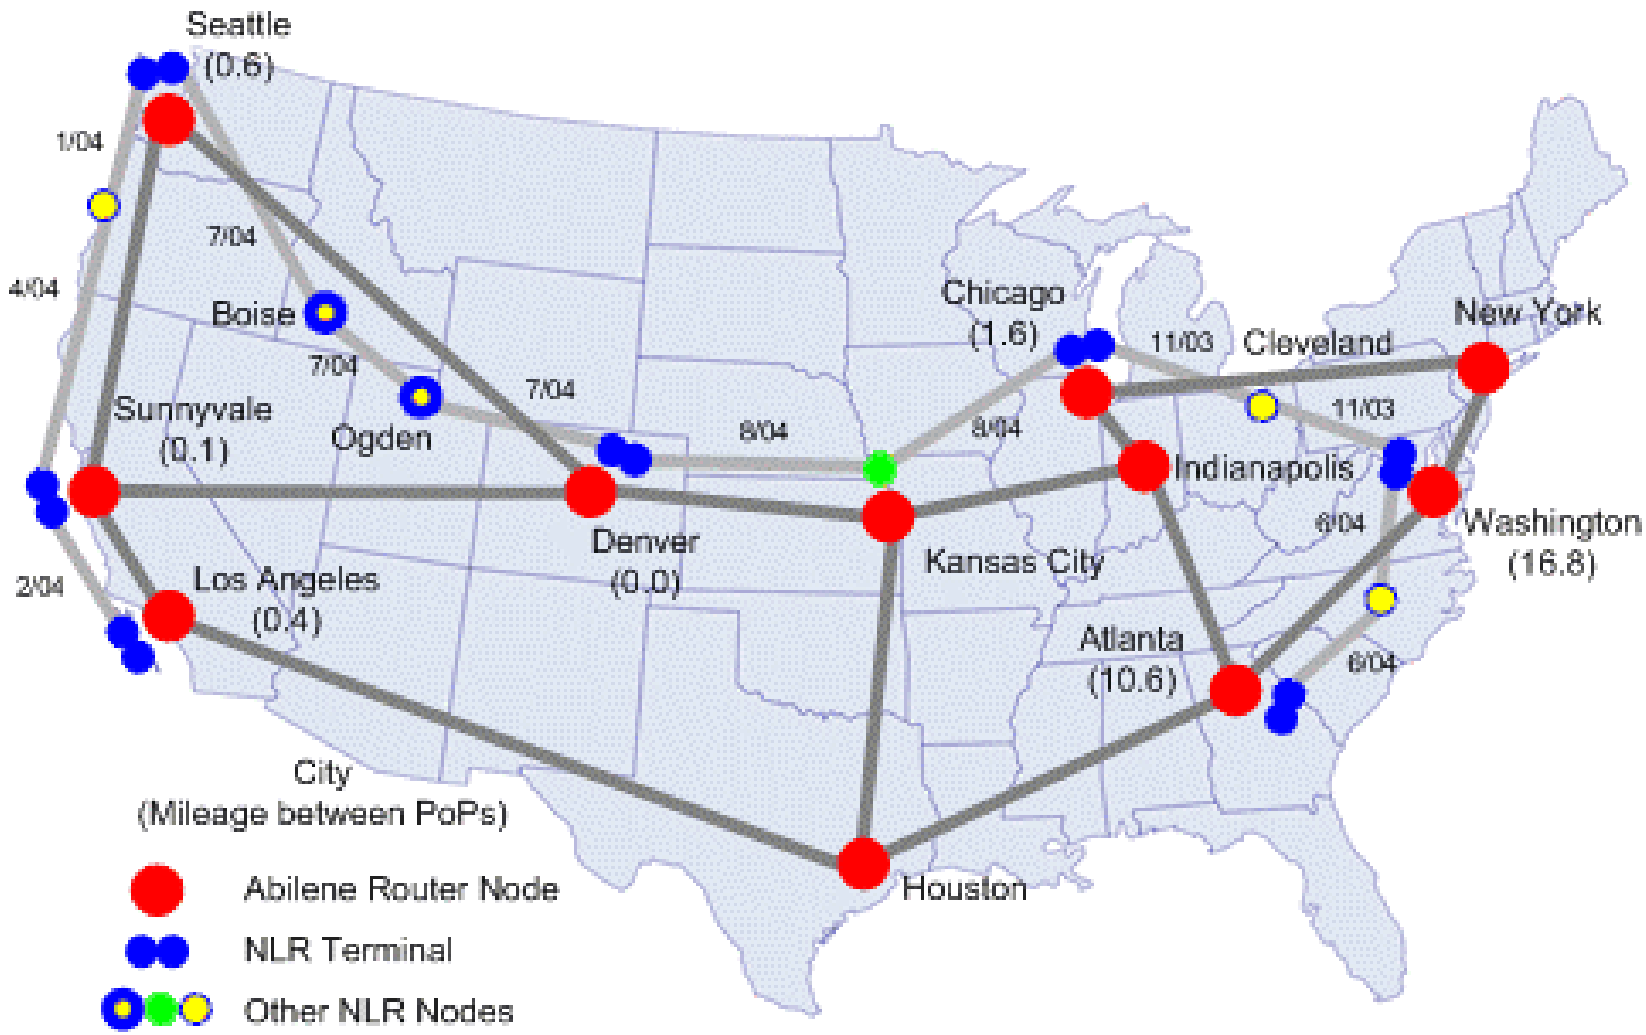

HOPI Basic Service

- Given resources availability, we cannot use individual wavelengths yet to model candidate architectures
- Instead, we will model them using lower-bandwidth, "deterministic" paths that resemble circuits – *"light paths"*
- Basic HOPI service: a 1/10 GigE point-to-point path with reasonable throughput, jitter, latency, and loss characteristics

HOPI Node

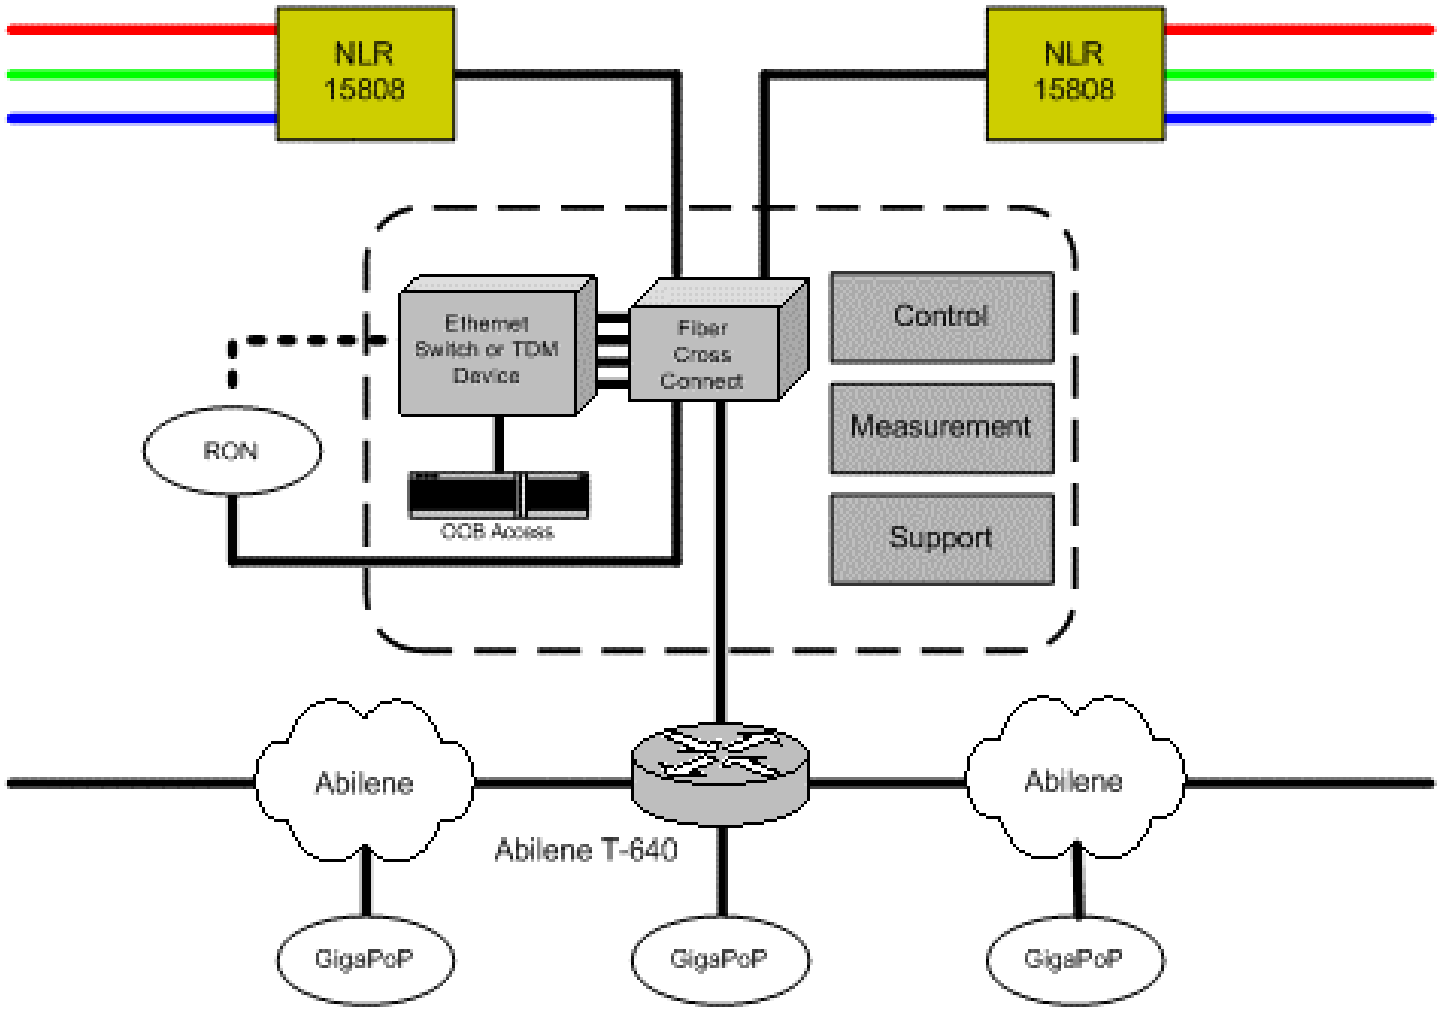
NLR
15808
NLR
15808
Ethernet
Switch or TDM
Device
Fiber
Cross
Connect
Control
Measurement
Support
RON
OOB Access
Abilene
Abilene
Abilene T-640
GigaPoP
GigaPoP
GigaPoP

Conclusions

- Abilene Network (packets over circuits)
  - Upgrade to 10-Gbps optical circuits complete
  - Focus on connector upgrades and advanced projects (e.g., Observatory)
  - Financial model on sound footing
- Evolution to a 3rd-generation network (packets + circuits over fiber)
  - Regional Optical Network (RON) development
    - FiberCo supporting project
  - National LambdaRail optical infrastructure
  - HOPI architecture and testbed underway

More information

- http://abilene.internet2.edu
- http://abilene.internet2.edu/observatory
- http://ipv6.internet2.edu
- http://www.fiberco.org
- http://hopi.internet2.edu
- http://www.nationalambdarail.org


INTERNET®
www.internet2.edu