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

Steve Corbató *corbato@internet2.edu*

Director, Network Initiatives, Internet2 &

Visiting Fellow, Center for High Performance Computing, University of Utah

CUDI

Reunión de Primavera 2005

Veracruz, Mexico

28 April 2005

Outline for this presentation

- Abilene Network
- U.S. R&E Exchange Points
- Critical role of dark fiber acquisition and State/Regional Optical Networks (S/RONs)
- National LambdaRail (NLR)
- Testbed for hybrid networking (HOPI)
- Planning for next generation network

Abilene Backbone

Abilene Focus Areas – 2005

- Support for high-throughput (multi-Gbps) flows
  - Supporting the End-to-End Performance Initiative
  - Objective: making large flows the norm across the Internet2 infrastructure
- Security
  - Expanded efforts in security motivated by bot-nets, DDoS, viruses, etc.
  - Enhancing proactive detection & response capabilities (with Abilene NOC, REN-ISAC & community)
- Provisioning dedicated capabilities (MPLS tunnles)
  - For the HOPI project
- Abilene Observatory
  - Supporting Network Research through an open measurement platform
- Providing experimental services on top of production network
  - Recent NSF awards supporting collaborating with network researchers
- IPv6
  - Roughly 2/3 of the connectors are IPv6 enabled
  - Roughly 1/2 of the peers are IPv6 enabled
  - We also do *experimental* IPv6 peerings with commercial ISPs

Abilene Observatory

- Provides co-location and measurement data for the network research community
- Co-location projects
  - PlanetLab – CS overlay network - recent NSF proposal submitted to change footprint
  - AMP (active measurement) project – SDSC/NLANR
  - PMA (passive measurement) project - the Indianapolis router clamp – SDSC/NLRANR
- Measurement data
  - Utilization, Netflow, Throughput, Latency, Routing, Router information, Syslog
  - Network access to all data - netflow requires an account

U.S. International Exchange Points

- Critical peering junctions and advanced service experimental platforms for R&E nets
  - Pacific Wave – West Coast
    - Distributed: Seattle, Los Angeles, and soon San Francisco Bay Area
  - Star Light - Chicago
  - MAN LAN - New York City
  - AMPATH - Miami
  - Atlantic Wave – East Coast
    - Distributed: NYC, D.C., Atlanta, and Miami – soon
- Increasingly co-located at major 'carrier hotels'
- Significant recognition of regional/international leadership through recent NSF IRNC awards

Abilene International Peering

Abilene International Network Peers

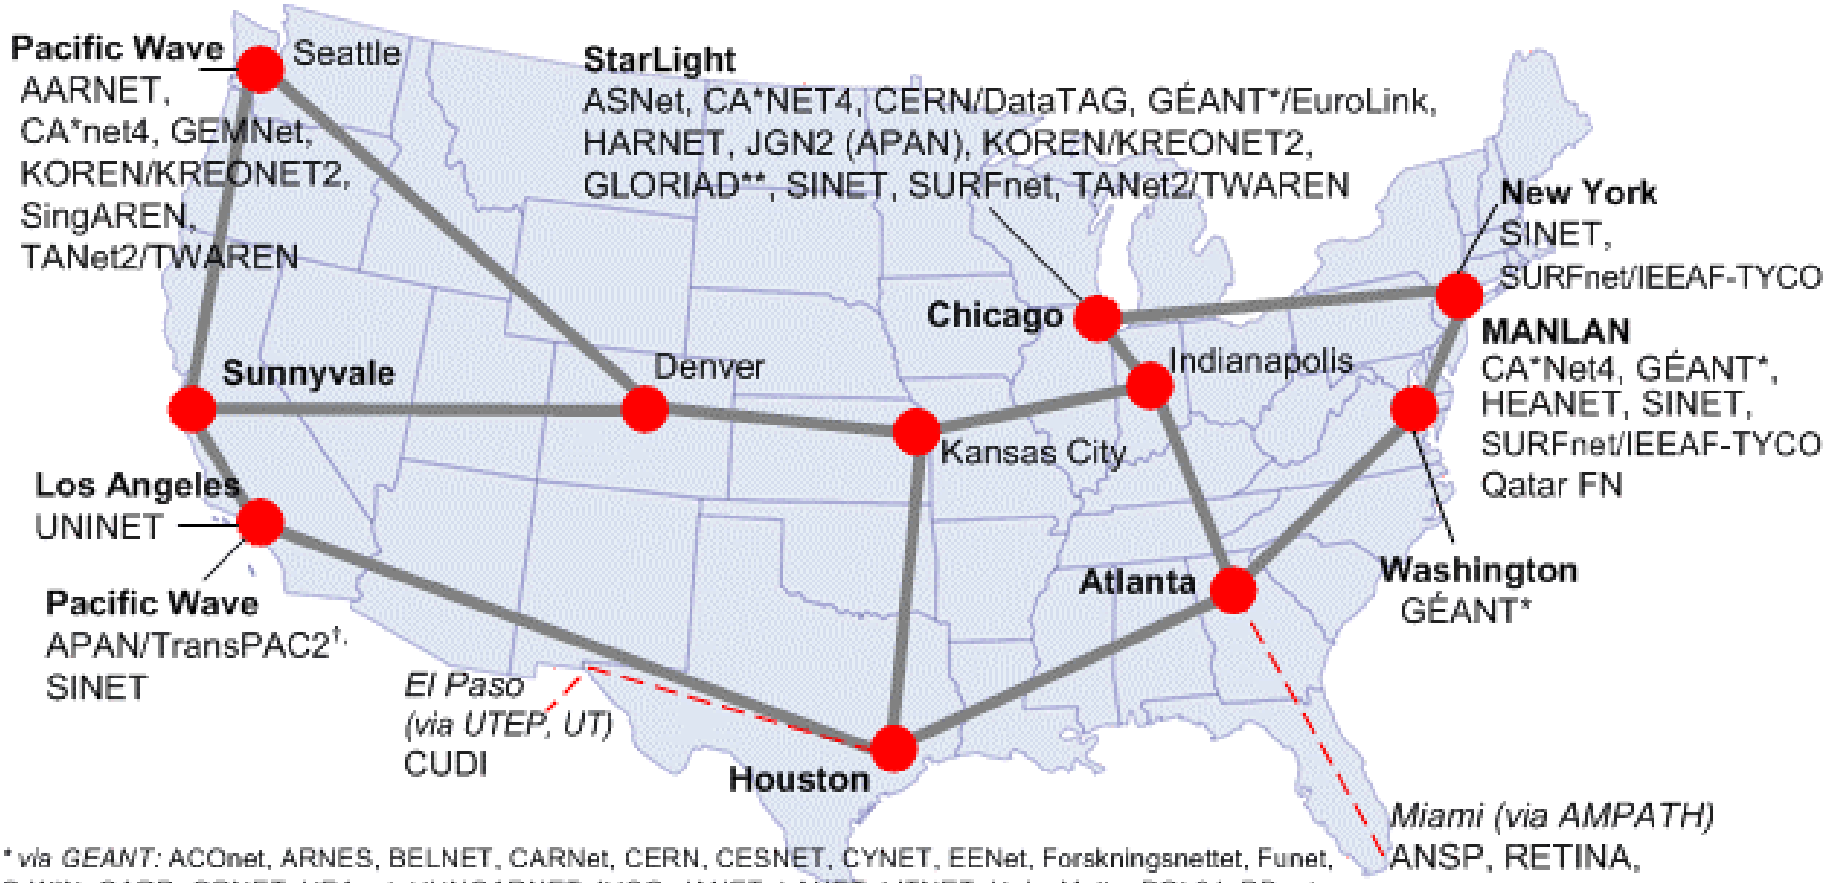

* *via GÉANT:* ACOnet, ARNES, BELNET, CARNet, CERN, CESNET, CYNET, EENet, Forskningsnettet, Funet, G-WIN, GARR, GRNET, HEAnet, HUNGARNET, IUCC, JANET, LANET, LITNET, Univ. Malta, POL34, RBnet, RCTS2, RedIRIS, Renater, RESTENA, REUNA2, Rhnet, RNP2, RoEduNet, SANET, SUNET, SURFnet, SWITCH, ULAKBYM, UNINETT

† *via APAN/TransPAC2:* WIDE/JGN, IMnet, CERNet/CSTnet/NSFCNET, KOREN/KREONET2, PREGINET, SingAREN, TANET2, ThaiSARN, WIDE (v6)

** *via GLORIAD:* CSTNET, RBnet

MAN LAN Exchange Point

- Manhattan Landing in New York City - partnership with NYSERNet, Indiana University, and the IEEAF
  - Provides a high performance, open exchange facility for all research and education networks
  - Located at 32 Avenue of the Americas (former AT&T Building) in lower Manhattan
- Easy interconnection to many national and international carriers and other research and education networks (Abilene, ESNet, NLR)
  - Peering model is open and bilateral
  - Cost recovery model - minimal connection charges for layer 2 facility, none for experimental layer 1 connections
- Focal point for Internet2's collaboration with CANARIE, GEANT-2, and SURFnet
- Participating in **Atlantic Wave** initiative
  - Emerging distributed exchange point along U.S. East Coast (NYC↔Miami)

MAN LAN services and capabilities

- Layer 2: Ethernet switch for IPv4/v6 peering with 1 GigE and 10 GigE interfaces
- Layer 1: TDM based optical equipment (SONET/Ethernet interfaces)
  - Cisco 15454
  - Nortel OME 6500
  - Nortel HDXc
- Layer 0: equipment to be installed soon
  - Optical cross connect to facilitate rapid circuit changes

Regional Optical Networks: Underlying hypothesis

- The fundamental nature of regional networking in the U.S. 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

	Distance scale (km)	*Examples*	*Equipment*
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 2nd Gen Abilene, NLR	Add OEO regenerators & O&M $'s

Leading & Emerging *Regional Optical Networks*

- Alabama
- 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
- 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

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 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

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

National LambdaRail Architecture

For more information regarding NLR see http://www.nlr.net or contact info@nlr.net

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 has commited $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

Network futures

What's different now?

- Within the last 5 years, the U.S. research universities have become **wholesale customers** of telecom assets & services
  - Over 27,000 route-miles of inter-city dark fiber are held by this community
  - ~25 Regional Optical Networks (RONs) have emerged - mostly state based and many with strong gubernatorial support (e.g., economic development)
- Single high-end PCs are capable of transmitting flows close to 10 Gbps over long distances
- Grid computing views the network as a *schedulable resource*
- Active examination of new service models (past best-effort IP as the common bearer service)

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 next generation architecture evolve?
- Examine a **hybrid** of shared IP packet switching and dynamically provisioned optical lambdas
- HOPI Project towards a Hybrid Optical and Packet Infrastructure
  - Immediate Goals
    - Implement testbed in 2005
    - Coordinate and conduct joint experiments with similar projects in U.S. and globally
  - Engaged Design and Corporate Advisory Teams

HOPI General Problem

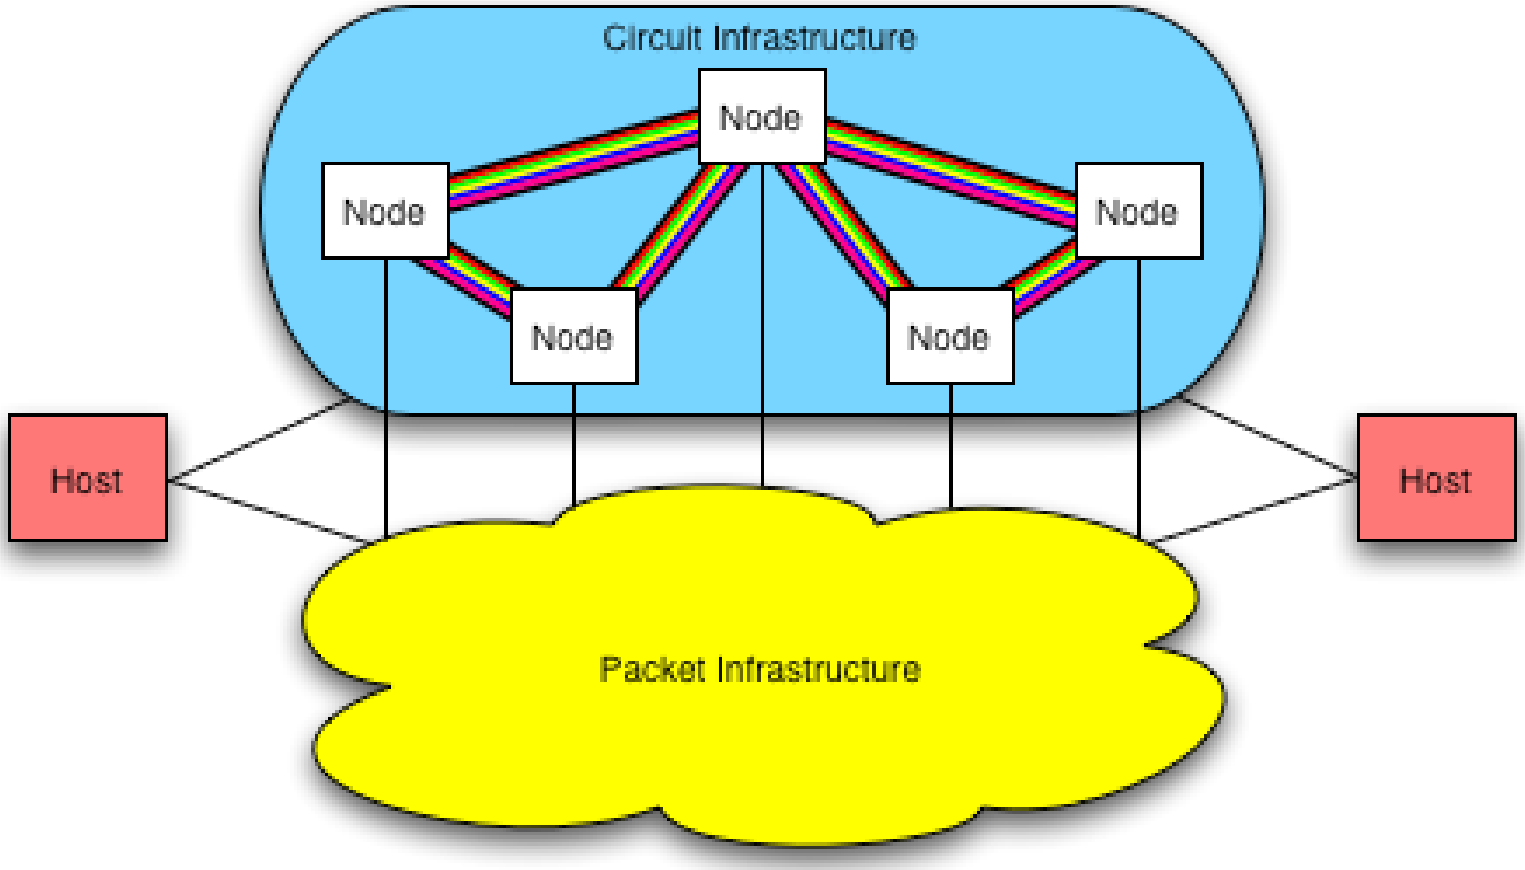

HOPI - General Problem

- How would one create a hybrid from these two infrastructures?
  - The nodes perform switching
  - The links are point-to-point circuit-like paths.
  - Each link may have attributes – for example, bandwidth.
    - Attributes may determine the ability to concatenate links.
  - Examples include:
    - Nodes can be λ switches with waves forming circuits – attributes include colors and bandwidth, etc.
    - Nodes can be SONET switches with paths being SONET links – attributes include channels, etc. For example, OC-3, OC-12, etc.
    - Nodes can be Ethernet switches with paths being point-to-point VLANS – attributes include bandwidth, etc.
      - HOPI will use this environment to examine different architectures
    - Nodes can be routers on a packet infrastructure and the point-to-point paths are MPLS L2VPNs

HOPI - Resources

- Abilene Network
  - MPLS tunnels and the packet switched network
- Internet2's 10-Gbps λ on the NLR national footprint
- MAN LAN exchange point in New York City
  - International 10-Gbps λ 's
    - TYCO/IEEAF - NYC – Amsterdam/SURFnet
    - NYC-London/GEANT2 (soon)
  - Layer 1 and 2 switching
- Collaborations with U.S. Regional Optical Networks (RONs) and other related efforts
  - E.g., GLIF, UltraLight, DRAGON

HOPI Node

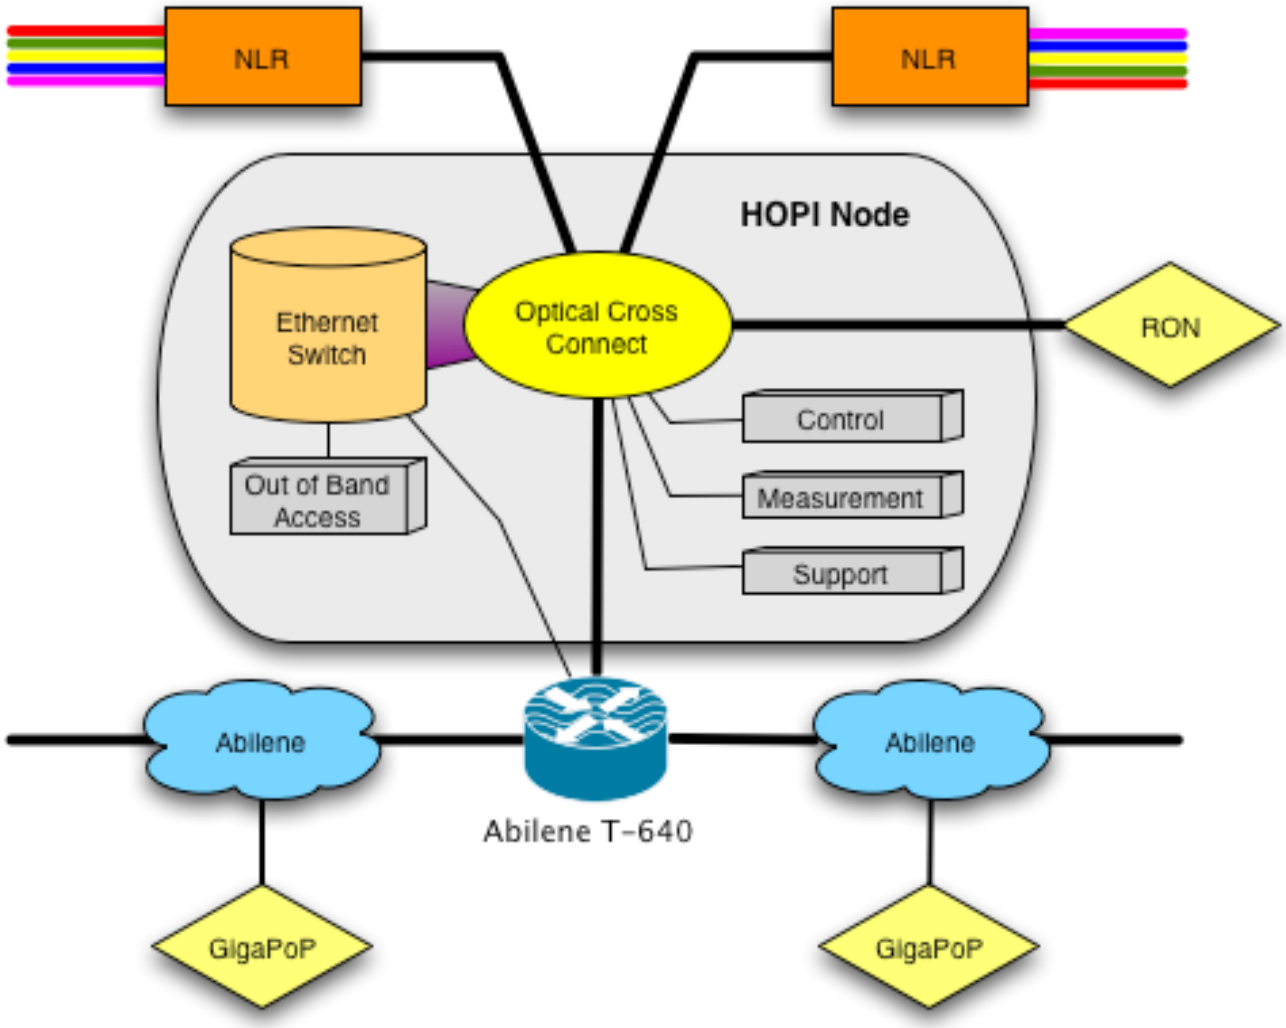
NLR
NLR
HOPI Node
Ethernet Switch
Optical Cross Connect
RON
Control
Measurement
Support
Out of Band Access
Abilene
Abilene
Abilene T-640
GigaPoP
GigaPoP

HOPI Topology

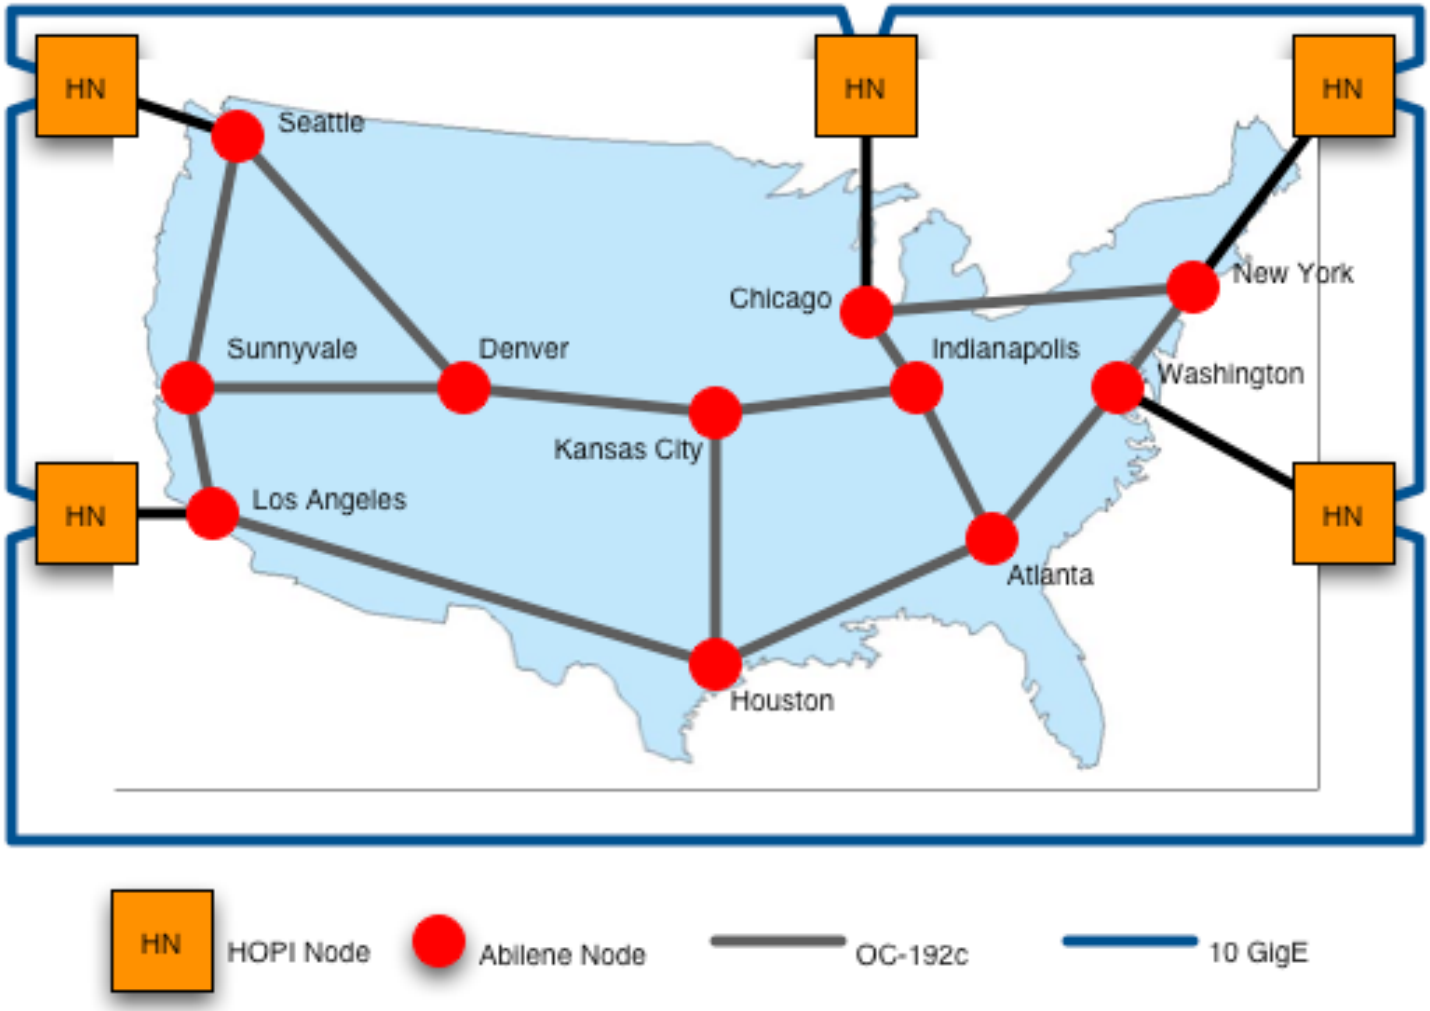

Abilene Network futures

- October 2007 – End of recent 1-year Abilene transport MoU extension
  - Sets 3rd-generation network planning timeline
    - Architecture definition: end 4Q05
    - Transport selection: end 1Q06
    - Router and other equipment selection: end 2Q06
    - Backbone deployed: end 4Q06
    - Connector transition (if necessary): starting 1Q07
  - Concurrently, review overall business plan and management model
  - Network design time frame: 2007-2012

Next Generation Network Design

- Critical factors
  - RON and International integration
  - Advanced service support
    - (Multicast, v6), High Performance Throughput, Measurement
  - Enhanced network research facilitation
  - Network and end-user security
  - Options for increased reliability
- Process
  - Hybrid architecture evaluation (HOPI)
    - IP core using 40?/10 Gbps transport
    - Dedicated capabilities (λ's, MPLS tunnels)
  - Evaluation of optical transport capabilities
    - NLR, commercial providers & RONs
  - Design & planning collaboration
    - U.S. & int'l partners (ESNet, TeraGrid, SURFnet, GEANT-2)

Evolving optical transport requirements

- Core wavelengths
  - 10 and hopefully 40 Gbps
  - Optimized availability
  - Applications: IP network backbone and backhaul, weather forecasting grids, radiological image transfer
  - To date, this service model has been the carrier model
    - Internet2 has significant interest in the potential hardening of NLR operational and maintenance models
- Flexible wavelengths
  - 10 and possibly 2.5 Gbps
  - Less stringent reliability requirements
  - Lower cost is critical objective
  - Potentially short duration (rapid provisioning and tear-down)
  - Applications: Grid clusters, IP network overflow, network research projects, highly redundant IP networks
  - By design, NLR is a natural source of this class of λ's

Potential Model for Services and Capabilities to be Offered in Next Generation Higher Education Network Infrastructure

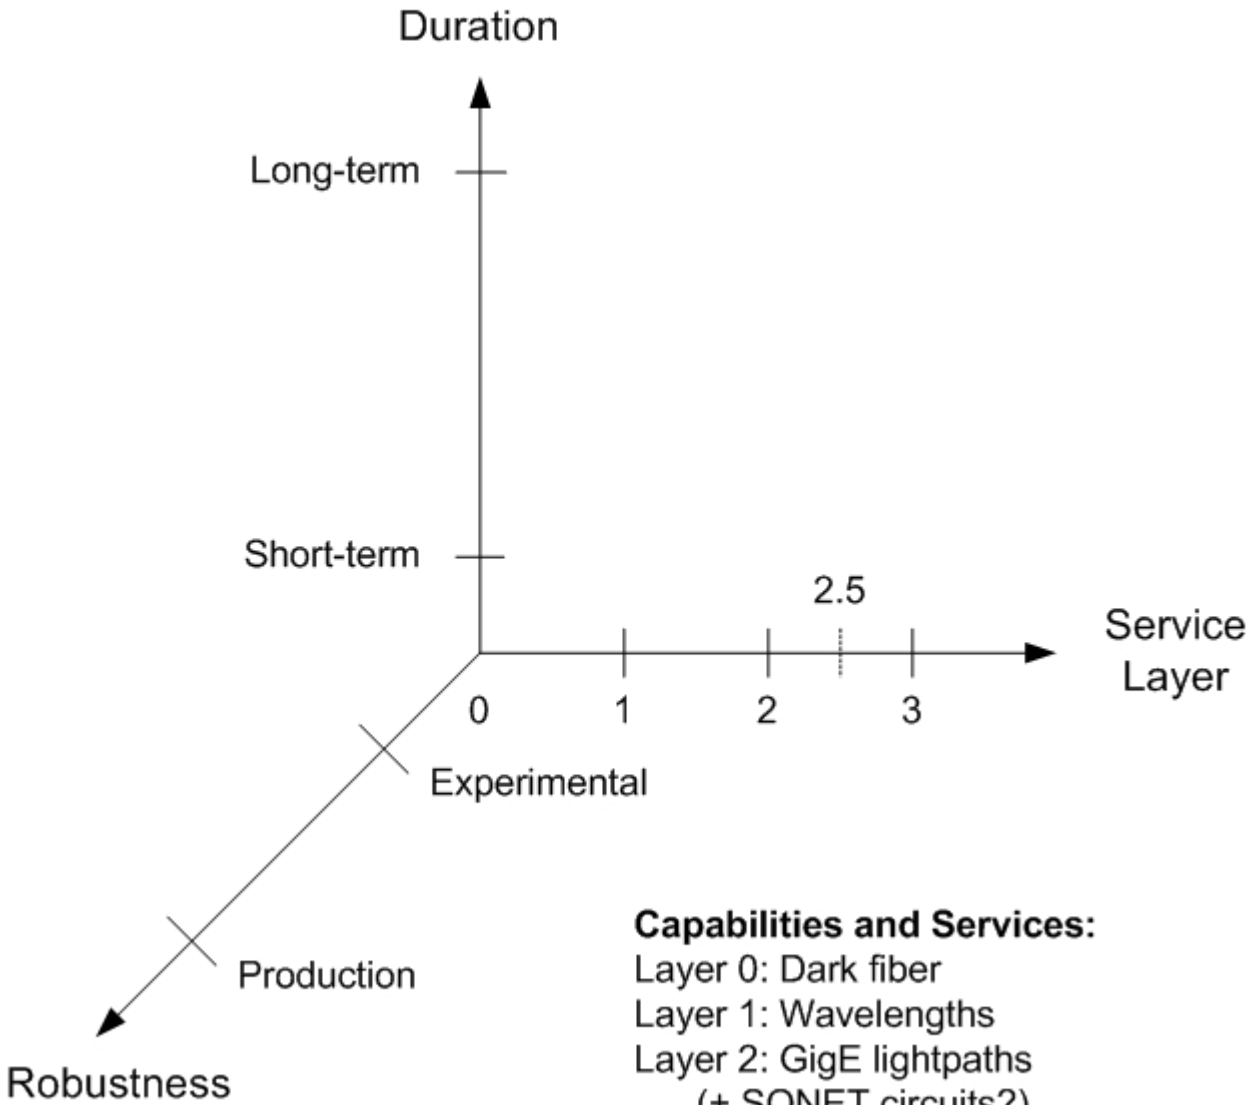

Capabilities and Services:
Layer 0: Dark fiber
Layer 1: Wavelengths
Layer 2: GigE lightpaths
(+ SONET circuits?)
Layer 2.5: MPLS tunnels
Layer 3: IPv4/v6

Service (re)differentiation

- Potential spectrum of services and capabilities
  - Dark fiber, wireless spectrum
  - Wavelengths
  - Subchannels
    - Gigabit Ethernet ‘circuits’
    - SONET circuits
    - MPLS tunnels
  - IPv4/v6
  - Overlay network support
- Need for new model of customer support and end-to-end connectivity delivery assurance
  - Working across campus, regional & national scales
  - Effective campus penetration of new services is a critical issue

Conclusions

- Abilene Network remains a viable packet infrastructure – supporting network research (Observatory) and architectural innovation (HOPI)
- Regional Optical Networks are transforming U.S. networking
- A national optical facility, NLR, is already partially operational – completion scheduled for later this year
- Projects such as HOPI are exploring the potential approaches to hybrid networking
- A third-generation U.S. higher education network architecture is being defined this year
- Additional details in tomorrow's talk (*Redes de la Nueva Generación*, 11 a.m.)


INTERNET®
www.internet2.edu