



Next Generation Abilene and U.S. Optical Networking Capabilities

Steve Corbató

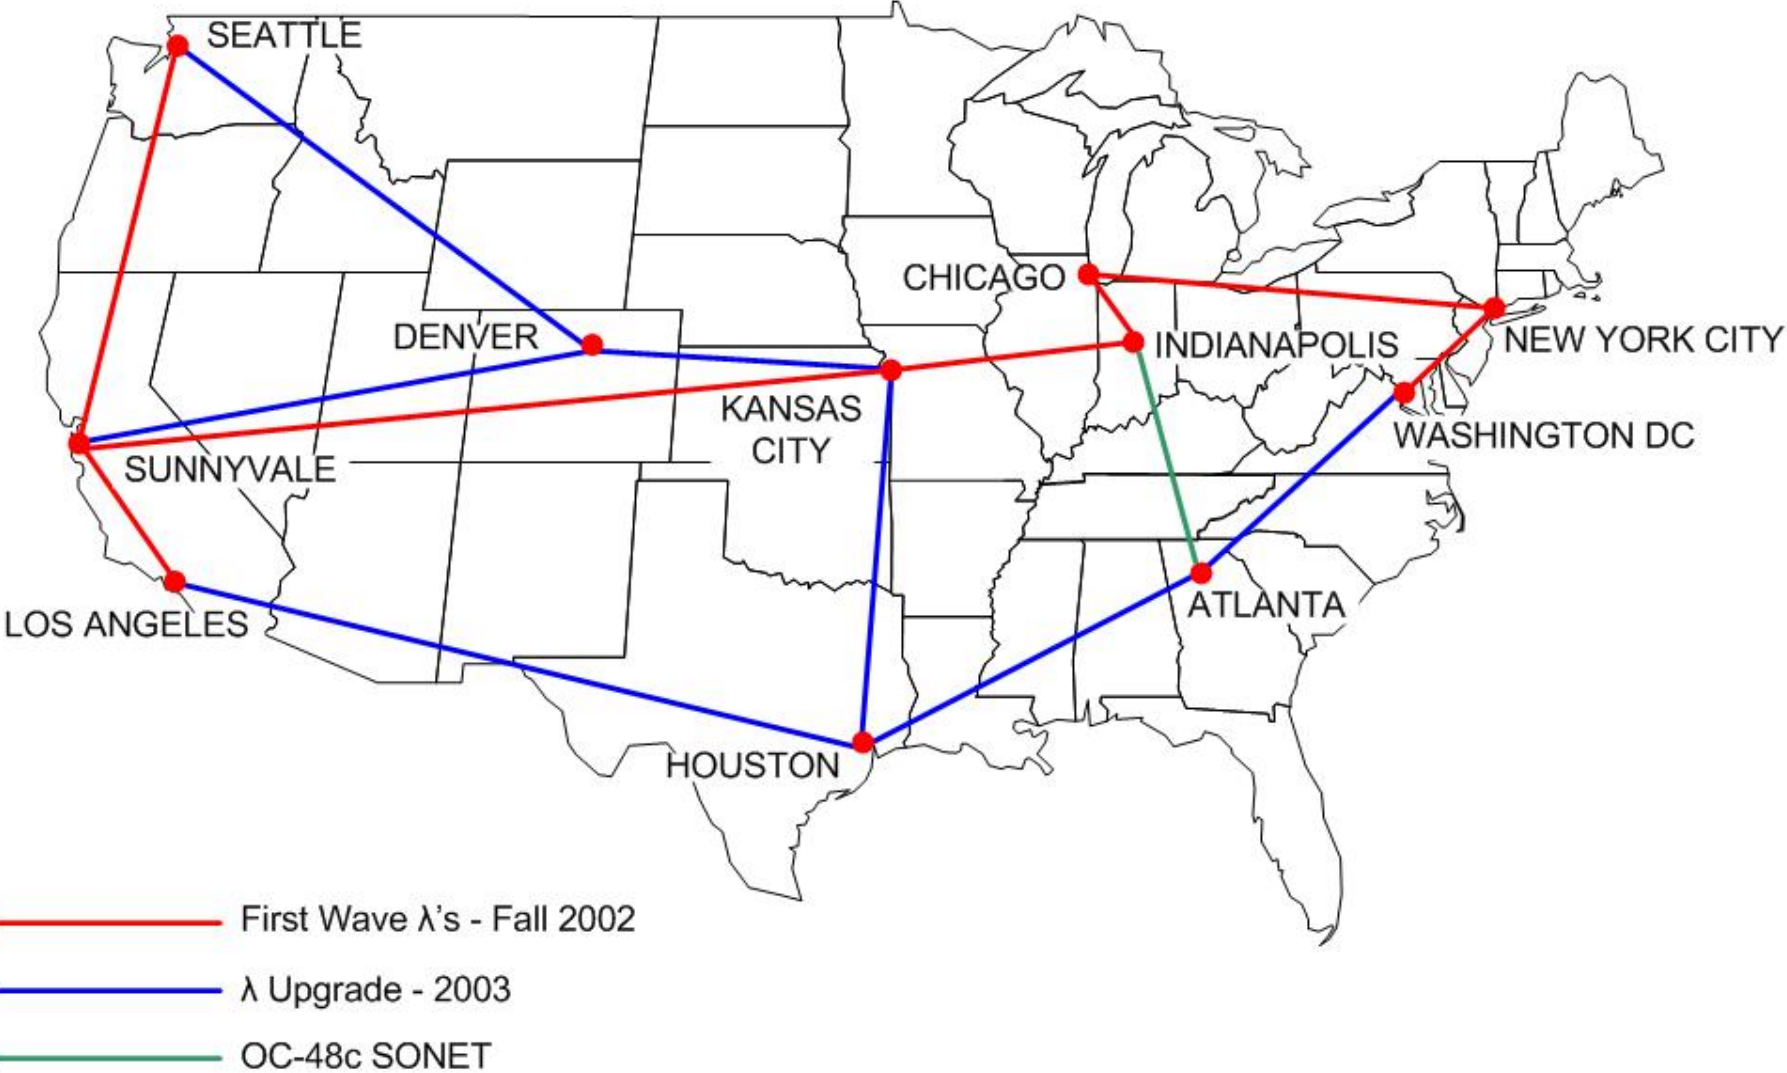
Director, Backbone Network Infrastructure

AMPATH Astronomy WG

Miami

31 January 2003

ABILENE NETWORK 10-Gbps OPTICAL UPGRADE - 2002-2003



50 direct connections (OC-3c → 10-Gbps)

- 7 OC-48c & 1 Gigabit Ethernet
- 2 10-Gbps connections pending (P/NW upgrade from OC-48)
- 24 connections at OC-12c or higher

224 participants – research univs. & labs

- All 50 states, District of Columbia, & Puerto Rico
- Oak Ridge Nat'l Lab – member participant as of today
 - Connected to SoX via 10-Gbps λ

Expanded access

- 64 sponsored participants
 - 5 free-standing medical schools
- 25 state education networks
 - Collaborative SEGP effort underway in Texas

Native IPv6 deployed

- v6 previously supported as overlay network
- Joins v4 multicast as natively supported advanced service
- UCSD demo (iGrid2002): 400 Mbps v6 - San Diego → Amsterdam

10 of 11 new Juniper T640 routers installed

- Very pleased with Juniper router performance and Cisco/Juniper interoperability

Transcontinental 10-Gbps path in place

- Six λ 's connected to network
 - First λ outage (3.5 hours): fiber cut in NYC in October
- ITEC network validation test
 - 8 Gbps of 2-way traffic (v4/v6 mix and 100% v6) transmitted over entire λ path without loss or reordering

Packetized, uncompressed High Definition Television (HDTV)

Raw HDTV/IP – single UDP flow of **1.5 Gbps**

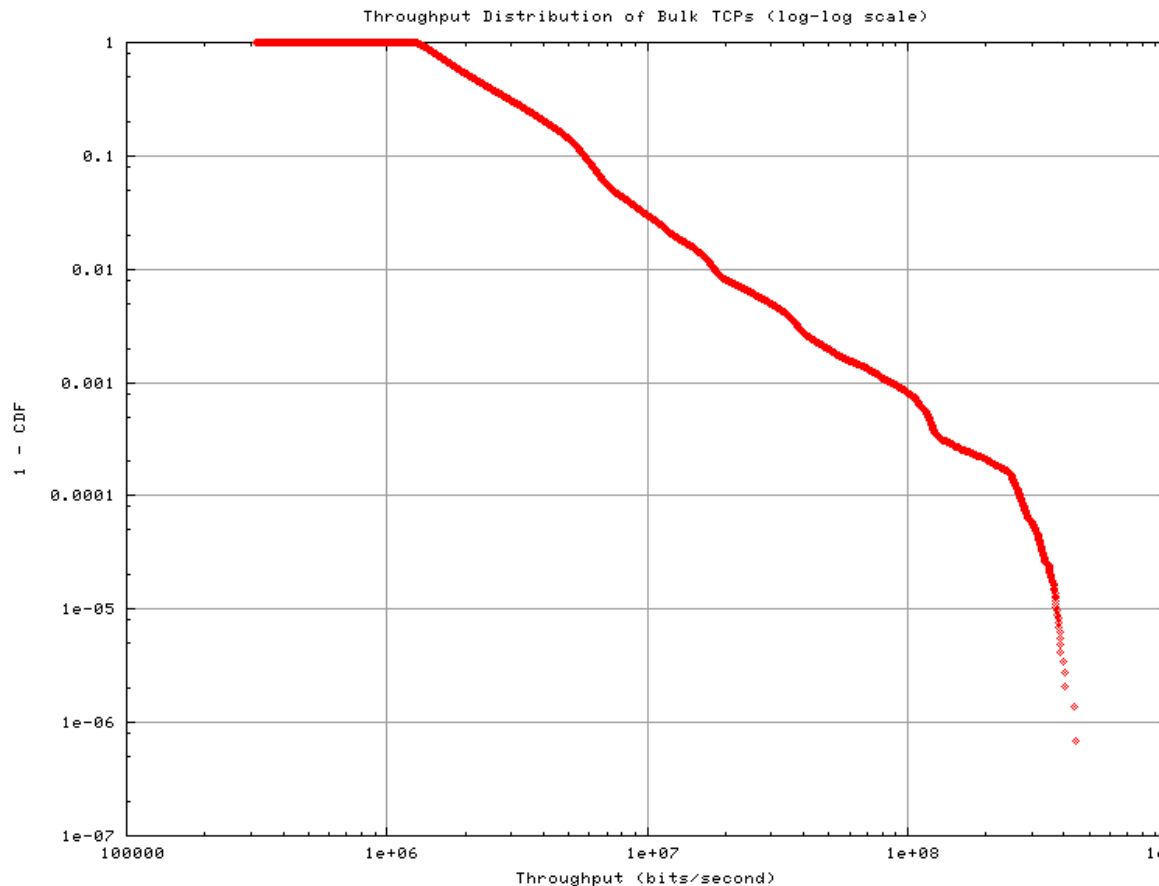
- Project of USC/ISIe, Tektronix, & U. of Wash (DARPA)
- 6 Jan 2002: Seattle to Washington DC via Abilene
 - **Single flow utilized 60% of backbone bandwidth**
- **18 hours: no packets lost**, 15 resequencing episodes
- End-to-end network performance (includes P/NW & MAX GigaPoPs)
 - Loss: <0.8 ppb (90% c.l.)
 - Reordering: 5 ppb
- *Transcontinental* 1-Gbps TCP requires loss of
 - <30 ppb (1.5 KB frames)
 - <1 ppm (9KB jumbo)
- Demo used 4.4 KB MTU



End-to-End Performance: 'High bandwidth is not enough'

Bulk TCP flows (payloads > 10 MBytes)

Median flow rate over Abilene: 2.1 Mbps



IEEAF transatlantic donations


- 10-Gbps λ (unprotected) and OC-12c SONET links
 - Now links Abilene in NYC and SURFnet in Amsterdam
 - λ 's from Silicon Valley to Amsterdam!
 - Joint effort in time for iGrid2002 event, Amsterdam
- Working collaboratively to extend reach in Europe

Importance of NYC router move

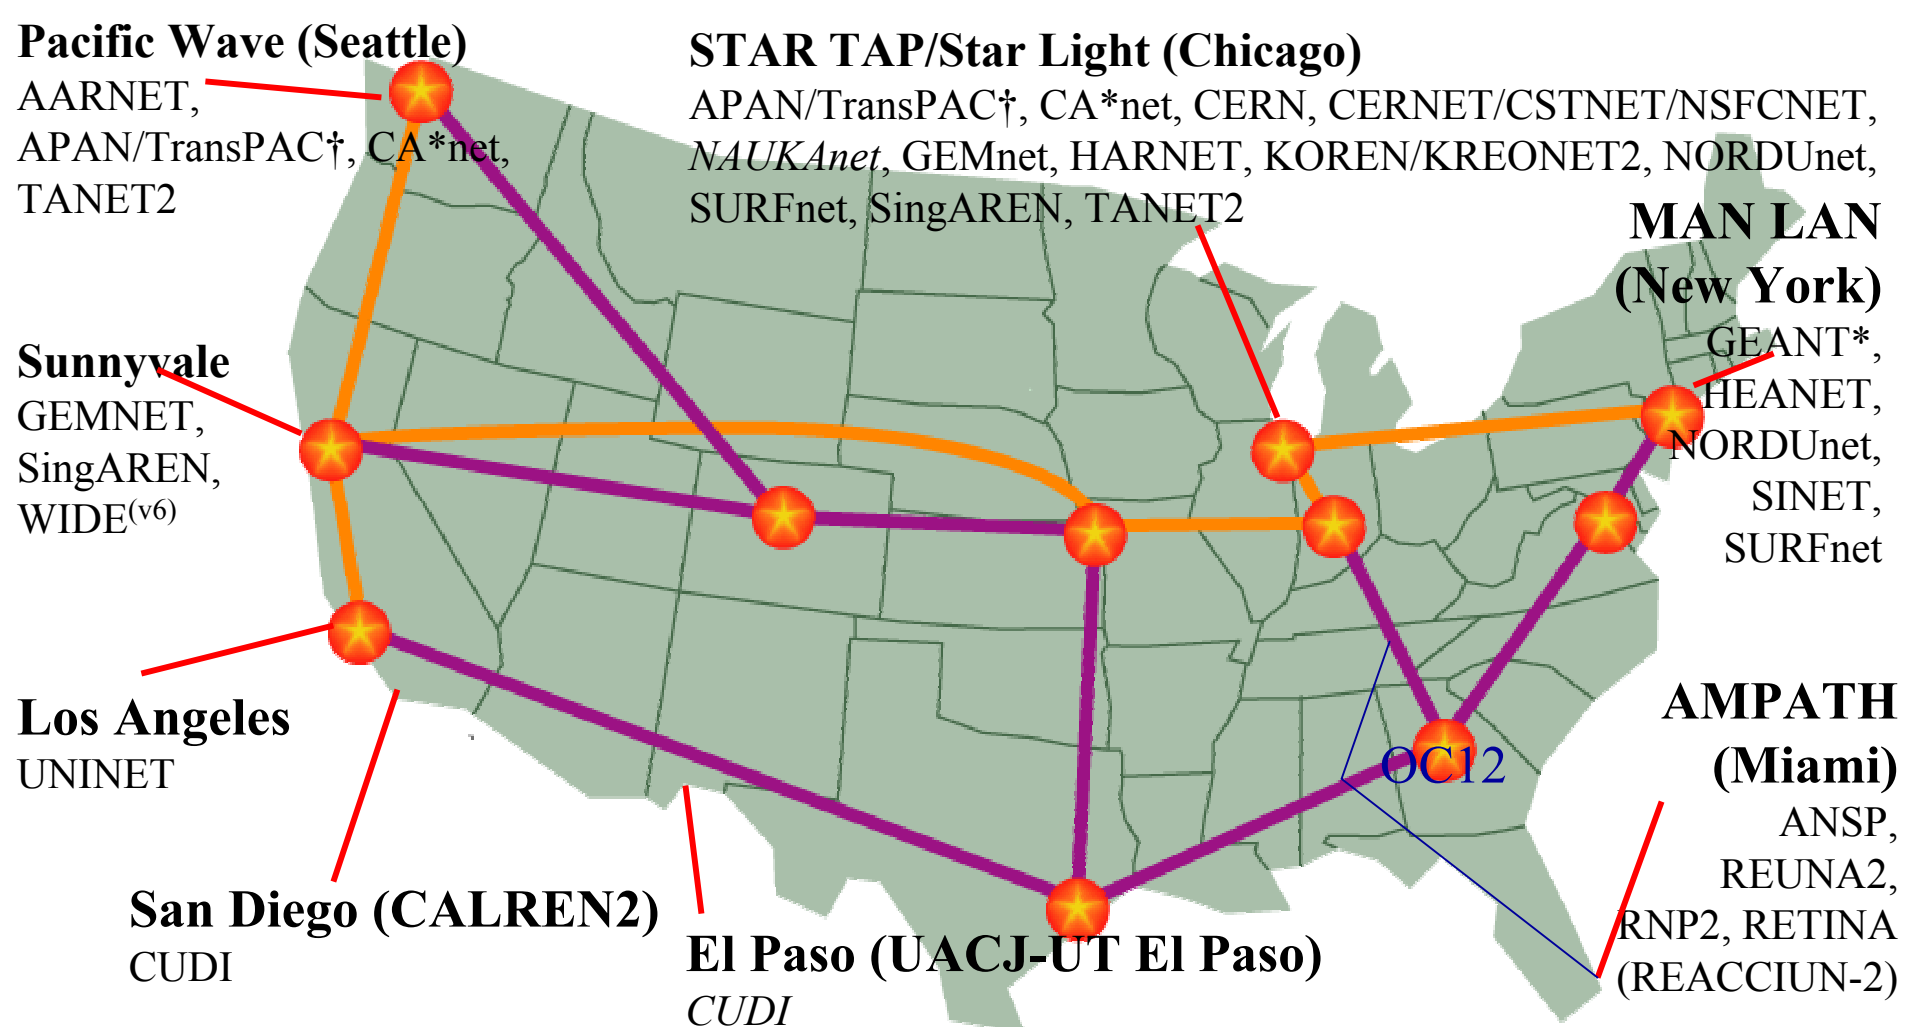
- 32 AoA as strategic carrier hotel (open FMMR)
- Manhattan Landing (MAN LAN) int'l R&E exchange point
- Collaborators: NYSERNET, Indiana U, IEEAF, DANTE

Other international exchange locations


- Chicago (StarLight), Seattle (P/WAVE), Miami (AMPATH)
- 10 Gig Ethernet to Star Light now and P/WAVE when ready



Abilene International Peering (January 2003)



•ARNES, ACONET, BELNET, CARNET, CERN, CESnet, CYNET, DFN, EENet, GARR, GRNET, HEANET, IUCC, JANET, LATNET, LITNET, NORDUNET, RENATER, RESTENA, SWITCH, HUNGARNET, GARR-B, POL-34, RCST, RedIRIS, SANET, SURFNET
 † WIDE/JGN, IMnet, CERNet/CSTnet/NSFCNET, KOREN/KREONET2, SingAREN, TANET2, ThaiSARN
 1/31/2003 9



Networks reachable via Abilene – by country

Europe-Middle East

Austria (ACOnet)
Belgium (BELnet)
Croatia (CARnet)
Czech Rep. (CESnet)
Cyprus (Cynet)
Denmark (UNI-C)
Estonia (ESnet)
Finland (FUnet)
France (RENATER)
Germany (G-Win)
Greece (GRnet)
Hungary (HUNGARnet)
Iceland (ISnet)
Ireland (HEANET)
Israel (IUCC)
Italy (GARR)
Latvia (LATNET)
Lithuania (LITNET)
Luxembourg (RESTENA)
Netherlands (SURFnet)
Norway (UNINETT)
Poland (PCSS)
Portugal (FCCN)
Romania (RNC)
Slovakia (SANET)
Slovenia (ARNES)
Spain (RedIris)
Sweden (SUNET)
Switzerland (SWITCH)
United Kingdom (JANET)
*CERN

Asia-Pacific

Australia (AARNET)
China (CERNET, CSTNET, NSFCNET)
Hong Kong (HARNET)
Japan (SINET, WIDE, IMNET, JGN)
Korea (KOREN, KREONET2)
Singapore (SingAREN)
Taiwan (TANET2)
Thailand (UNINET, ThaiSARN)

Americas

Argentina (RETINA)
Brazil (RNP2/**ANSP**)
Canada (CA*net4)
Chile (REUNA)
Mexico (CUDI)
United States
(Abilene, vBNS)

More information about reachable networks at
www.internet2.edu/abilene/peernetworks.html

Also, see www.startup.net

Advanced Services

- Multicast - high performance
- IPv6 - native, high performance
- Resiliency
- Security

Measurement

- Deploy/exploit new active & passive capabilities
 - Performance, traffic characterization, routing, SNMP
- e2e performance initiative support
- *Abilene Observatory*: correlated data archive for network research

Experimentation and collaboration

- *Abilene Observatory*: experiment/overlay co-location
- 'Lambda Grid' experimentation (MPLS lightpaths)
- International connectivity (Europe, Asia, CALA)
 - *MAN LAN* R&E exchange point in NYC
- TeraGrid interconnection (LA and Chicago)

Currently a program outline for better support of computer science research


1) Improved & accessible network data archive

- Need coherent database design
- Unify & correlate 4 separate data types
 - SNMP, active measurement data, routing, traffic characterization

2) Provision for direct network measurement and experimentation

- Resources reserved for two additional servers
 - Power (DC), rack space (2RU), router uplink ports (GigE)
- Initial candidates include large-scale overlay networks (e.g., PlanetLab)

- Abilene is now running native IPv6 over the entire Cisco 12008 and Juniper T640 backbone
 - Dual stack mode
 - IS-ISv6 used for internal routing
- Significant number of peers and connectors already have converted
- Tunnel support consolidated
 - IU-NOC provides support for existing tunnels
 - Not accepting any new tunnels
- Abilene provided addressing
 - 2001:468::/35 from ARIN for participants – 63% allocated
 - 3ffe:3700::/24 from 6bone for SEGP / sponsored users
- Native IPv6 (UCSD iGrid demo: 400 Mbps v6 SD-AMS)
- Kudos to Abilene NOC, IPv6 WG, Cisco, and Juniper



Abilene native IPv6 peerings – January 2003

Connectors (13)

- Great Plains Network
- Indiana Gigapop
- MAGPI
- MAX
- NYSERNet
- Oregon Gigapop
- Pittsburgh Gigapop
- SDSC
- WiscREN
- NoX
- South Florida Gigapop
- Front Range Gigapop
- ONEnet

Peers/Exchange Points (12)

- 6TAP
- APAN/TransPAC
- CUDI
- JGNv6/WIDE
- SingAREN
- SURFNET
- vBNS+
- AMPATH
- CA*NET(3)
- KREONet2
- HEAnet
- NORDUnet



U.S. Optical Networking Initiatives

Optical initiatives: Primary motivations and a caveat

Emerging requirements may not met by a high-performance, yet best-effort IP network


- DWDM: 10-Gbps channels now; 40-Gbps hard, but coming
- Computational science grids
 - Applications with deterministic network requirements
- Infrastructure for basic and applied network research

Period of unprecedented contrarian economic opportunity

- Distressed fiber assets available on national scale
- Optronics industry severely impacted by carrier woes


However, optical networking alone does not solve the end-to-end performance problem

- Host configuration (Web100)
- Local networking capability (DAST, Internet2 E2EPI)



Optical network project differentiation

	<i>Distance scale (km)</i>	<i>Examples</i>	<i>Equipment</i>
Metro	< 60	UW(SEA), USC/ISI(LA)	Dark fiber & end terminals
State/ Regional	< 500	I-WIRE (IL), I-LIGHT (IN), CENIC ONI	Add OO amplifiers
Extended Regional/ National	> 500	TeraGrid NG Abilene, Light Rail	Add OEO regenerators & O&M \$'s



Unique optical requirements in Higher Education Community (HEC)

10-Gbps: 10 Gigabit Ethernet preferred over OC-192c SONET

HPC could need 40-Gbps λ 's prior to the carriers

Integrated view of network management

- Transport & IP engineering/operational approaches are not intrinsically different
- SNMP preferable for network polling

HEC can provide experimental environment for development of 'rational', customer-focused optical switching

- Switching tightly integrated with optical transport
- Capacity for IP backbone expansion and p2p λ 's

Three current projects

- Fiberco
- USA Waves
- National Light Rail

Common factors

- National fiber footprint represents a strategic asset for HEC
- All leverage the much lower *incremental* cost of λ 's in an existing DWDM system (vs. the cost of the first λ)

Differentiating factors

- Scope
- Buy vs. build
- Production vs. research capabilities
- Participation cost

National facilities-based approach for optical networking and network research

- 15,000+ miles of fiber footprint
- HEC owned/managed fiber and optronics for p2p λ 's
- Shared experimental services: IP and GigE

Enabling innovative network research is key goal

Leadership: CENIC, Pacific Northwest Gigapop

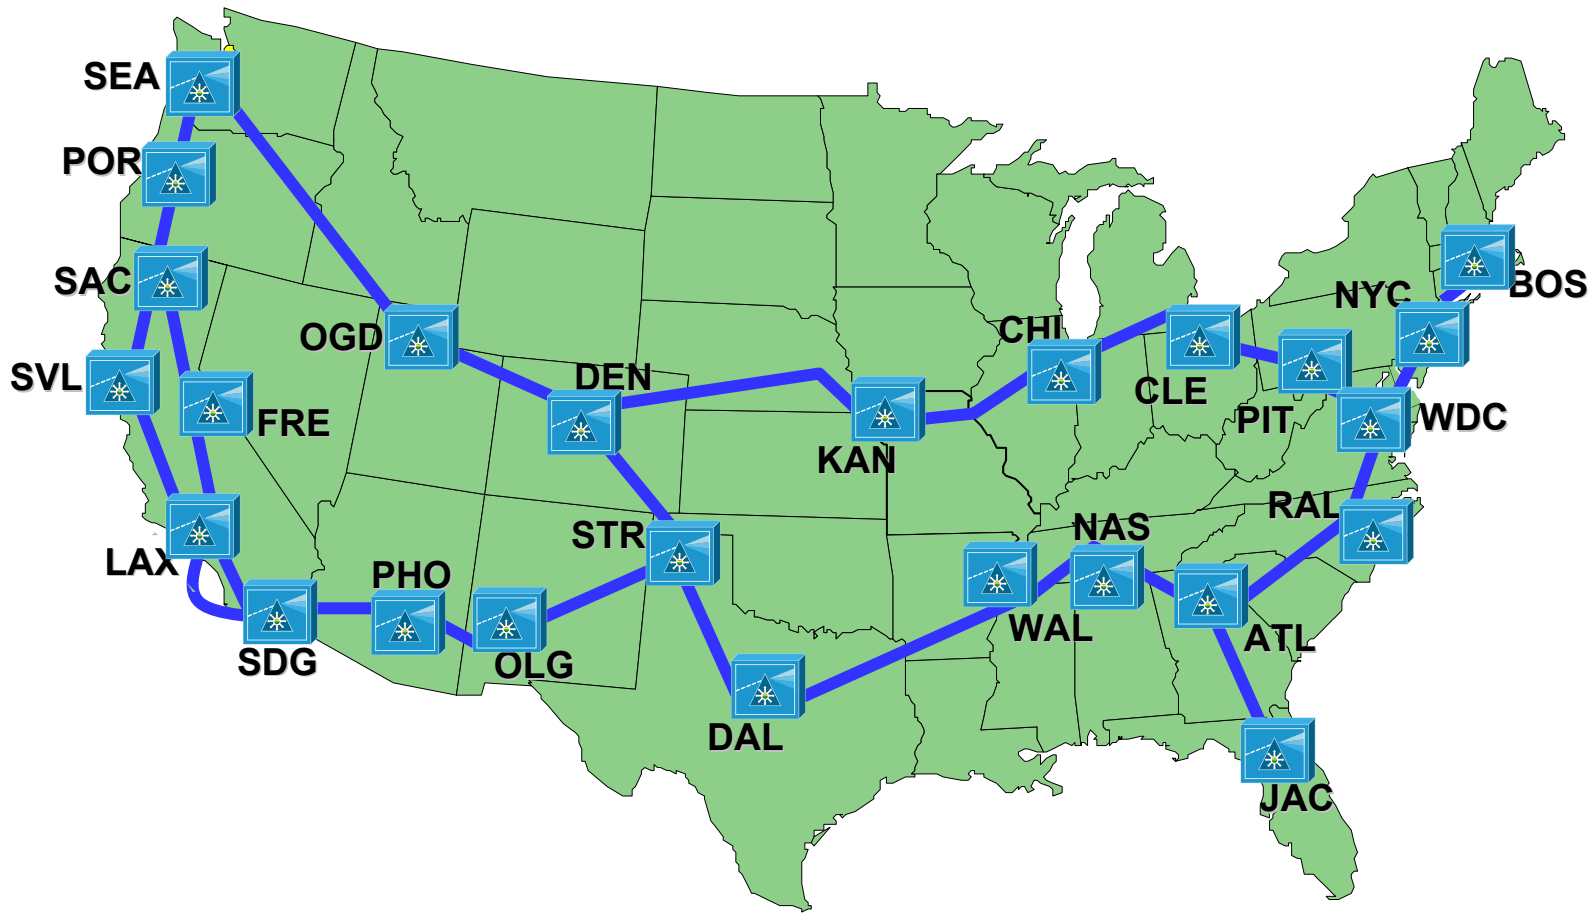
- Outgrowth of CENIC ONI regional project
- UCAID and multiple research universities collaborating

Corporate partners: Cisco, Level 3

Economics

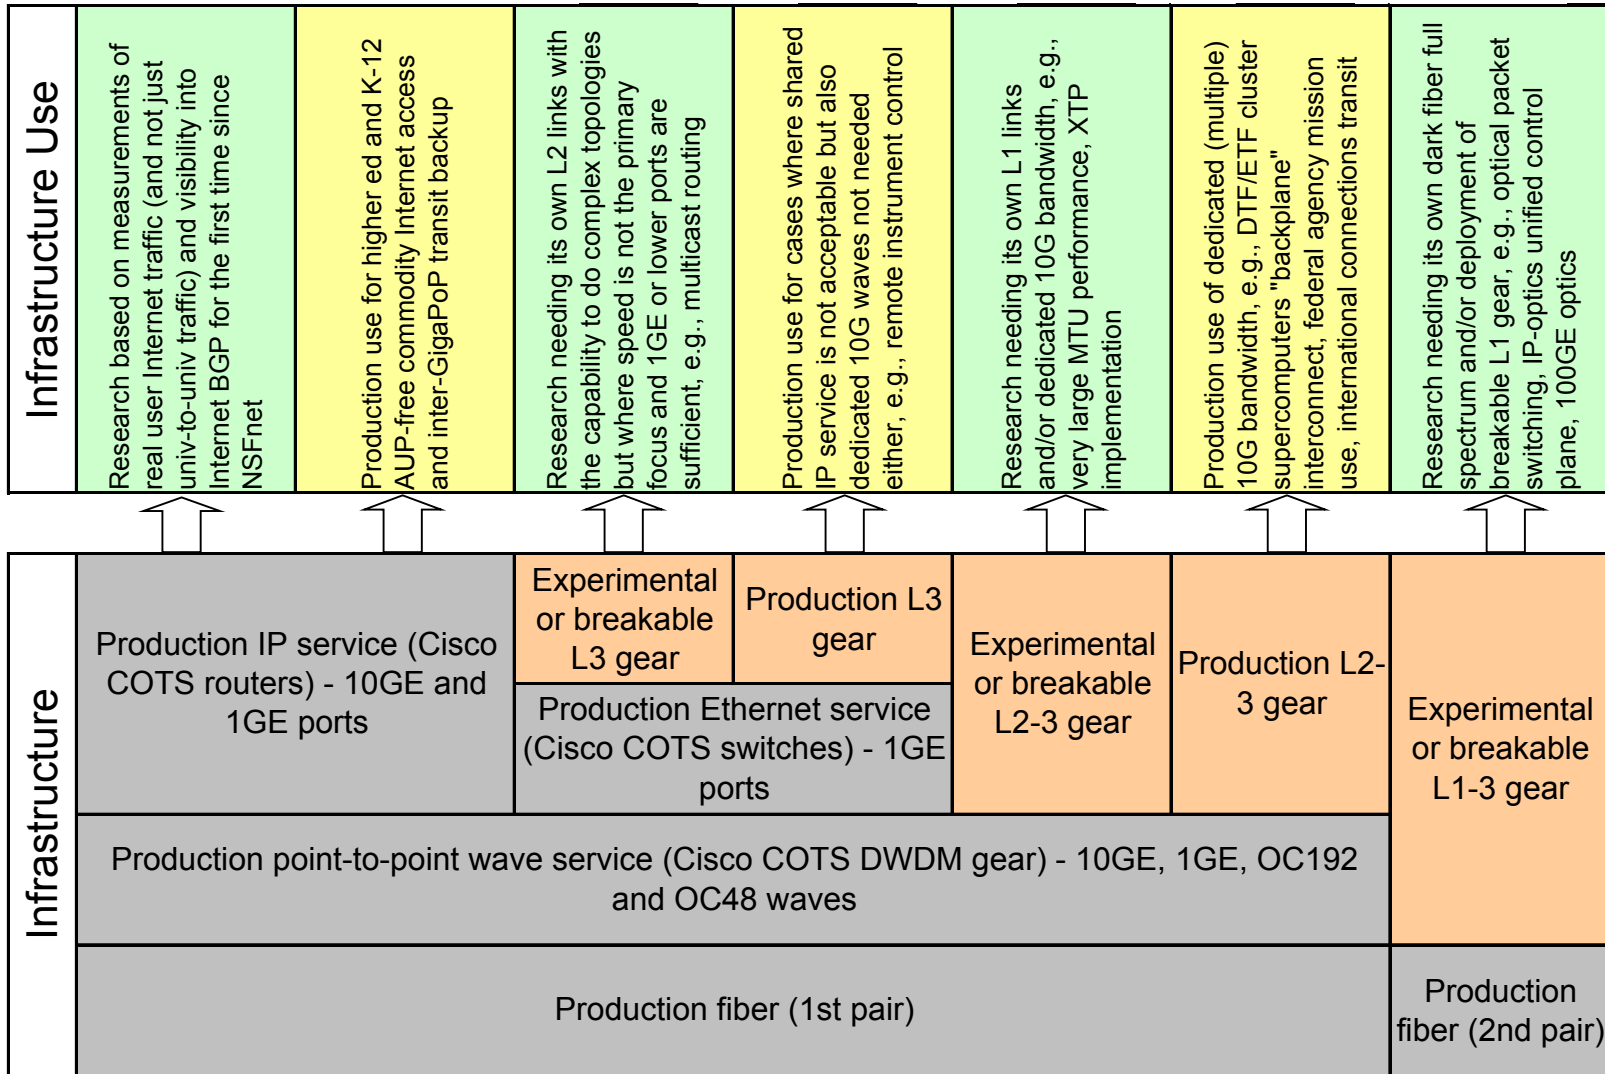
- 5-year cost for national footprint: \$83M (for 4 10-Gbps λ 's)
- Significant participation fee: \$5M over 5 years

NLR Footprint and Layer 1 Topology



 15808 Terminal, Regen or OADM site (OpAmp sites not shown)
 Fiber route

NLR networking research use vs. production (including science research) use



NLR operated
 NLR or its production customer or researcher operated
 Research use
 Production use

- 3 significant national optical networking initiatives underway in the U.S.
 - Fiberco, USA Waves, National Light Rail
- Higher education community will continue to acquire dark fiber assets on the national and regional scales in 2003
- Regional optical networks will be deployed
- Whether a national optical networking capability will be '*built or bought*' is an open issue
 - Possibility of hybrid approach
- In either case, expanding requirements of the computational science and network research communities must be addressed