Abilene Update Session

Steve Corbató and Rick Summerhill
Backbone Network Infrastructure

Internet2 Spring Member Meeting
Washington D.C.
10 March 2003

Abilene Update: Policy

Steve Corbató
Director, Backbone Network Infrastructure

Internet2 Spring Member Meeting
Washington D.C.
10 March 2003

Core features

- Higher ed's network
- IPv4+v6 common bearer service
- Bandwidth availability & utilization incentive
- Peering limited to national & int'l R&E nets
- Regional aggregation model
- "3-4 Nines" reliability
  - Advanced service deployment, 7x24 NOC
- Open measurement platform

Partnership approach

- The Abilene Network is an Internet2 high-performance backbone done in partnership by
  - Internet2
  - Cisco Systems (routers, switches, and access)
  - Juniper Networks (routers)
  - Nortel Networks (SONET kit)
  - Qwest Communications (SONET & DWDM circuits, co-location)
  - Indiana University (network operations center)

  - Internet2 Test & Evaluation Centers (ITECs)
    - North Carolina and Ohio

Abilene governance and advisory groups

- Internet2
  - Board of Trustees
  - Network Planning and Policy Advisory Council (NPPAC)
  - NRLC (research), ISC (industry), ASC (apps)
- Abilene
  - Executive Committee
  - Planning Committee
  - Technical Advisory Committee (TAC)
- OSAND
  - Board of Trustees
- The Quilt
  - Executive Committee
  - Steering Group

Abilene Timeline - I

- 1998
  - Qwest/Nortel & Cisco partnerships; Indiana Univ NOC
  - OC-3 & OC-12 connectors; native multicast
- 1999
  - 200-Mbps HDTV/IP; OC-48c BB complete
- 2000
  - network scale reached; 1st OC-48 connection
  - DARPA Supernet interconnect; Quilt SEGP effort
- 2001
  - First SEGPs; Qwest transport extension
  - 1st GigE connection; Quilt Project formalized

Abilene Timeline - II

- 2002
  - 1.5-Gbps HD/IP; native v6; Juniper T640 selection
  - DWDM BB transport; TeraGrid interconnect
  - IEEAF/Europe circuits; 25th SEGP
- 2003
  - 8-Gbps IPv4+v6 flows; 1st 10-Gbps (10 GigE) connectors
  - Abilene Observatory; MPLS experimentation
  - REN-ISAC; 10-Gbps λ BB completed
- 2004 - ???
- 2005
  - 3rd generation decisions & design
- 2006
  - Qwest extension expires (October)

Future of Abilene

- Abilene transport (DWDM & SONET) MoU with Qwest in place through October 2006
  - Final 10-Gbps λ's to be installed this summer
- New Juniper T640 routers deployed during summer 2002
- Current peak load ~10% of upgraded bandwidth
- Traffic doubling time ~ 1 year
- Engaged user community
- Ensemble of advanced networking projects

Limited peering model with experimental networks

- Abilene has a history of collaboration with experimental IP networks on national scale
- Interconnection and limited peering
- Objectives
  - Experiments (HDTV/IP, VLBI)
  - Demonstrations (SCxy)
  - Limited access (early TeraGrid access for PSC)
- Examples
  - DARPA Supernet (DC interconnect)
  - TeraGrid (Chicago and soon LA interconnect)

ABILENE NETWORK 10-Gbps OPTICAL UPGRADE - 2002-2003

Abilene scale – April 2003

- 50 direct connections (OC-3c → 10-Gbps)
  - 2 10-Gbps connections (both 10 Gigabit Ethernet)
  - 6 OC-48c & 1 Gigabit Ethernet
  - 24 connections at OC-12c (622 Mbps) or higher
- 221 participants – universities and labs
  - All 50 states, District of Columbia, & Puerto Rico
  - Recently: Oak Ridge National Laboratory
- Expanded access
  - 70 sponsored participants
  - 25 state education networks
    - Efforts underway in Texas and Massachusetts

Abilene Federal/Research Peering (Jan 2003)

Last updated: 17 January 2003

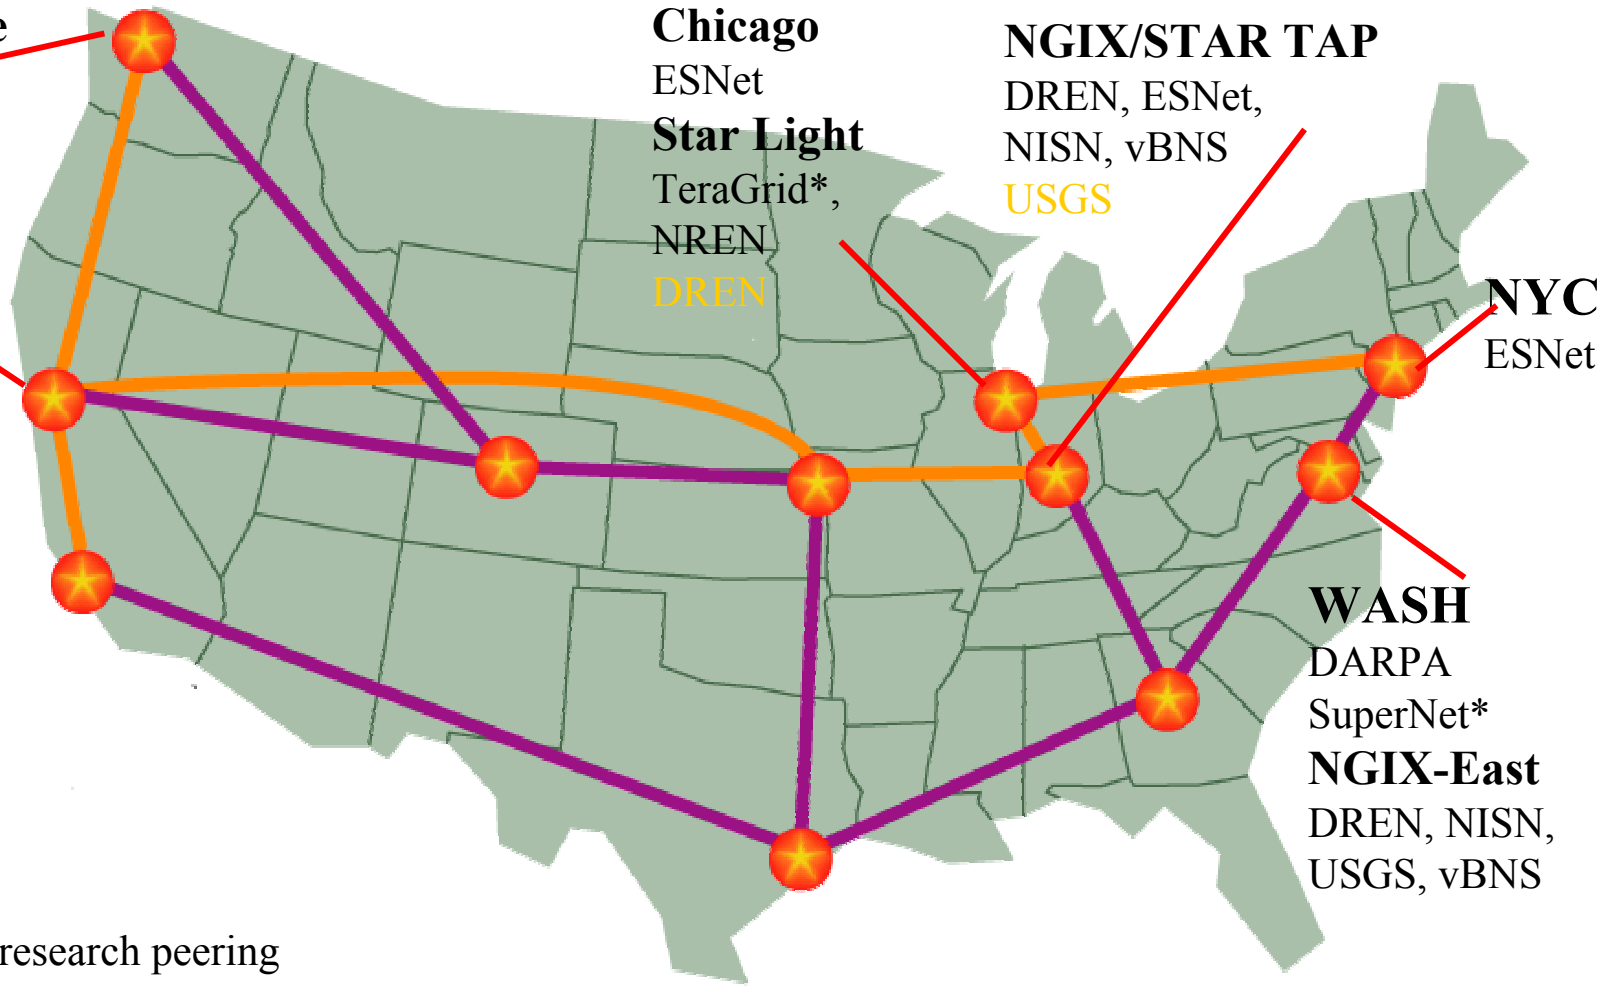

* Limited research peering

Abilene International Peering (January 2003)

Last updated: 17 January 2003

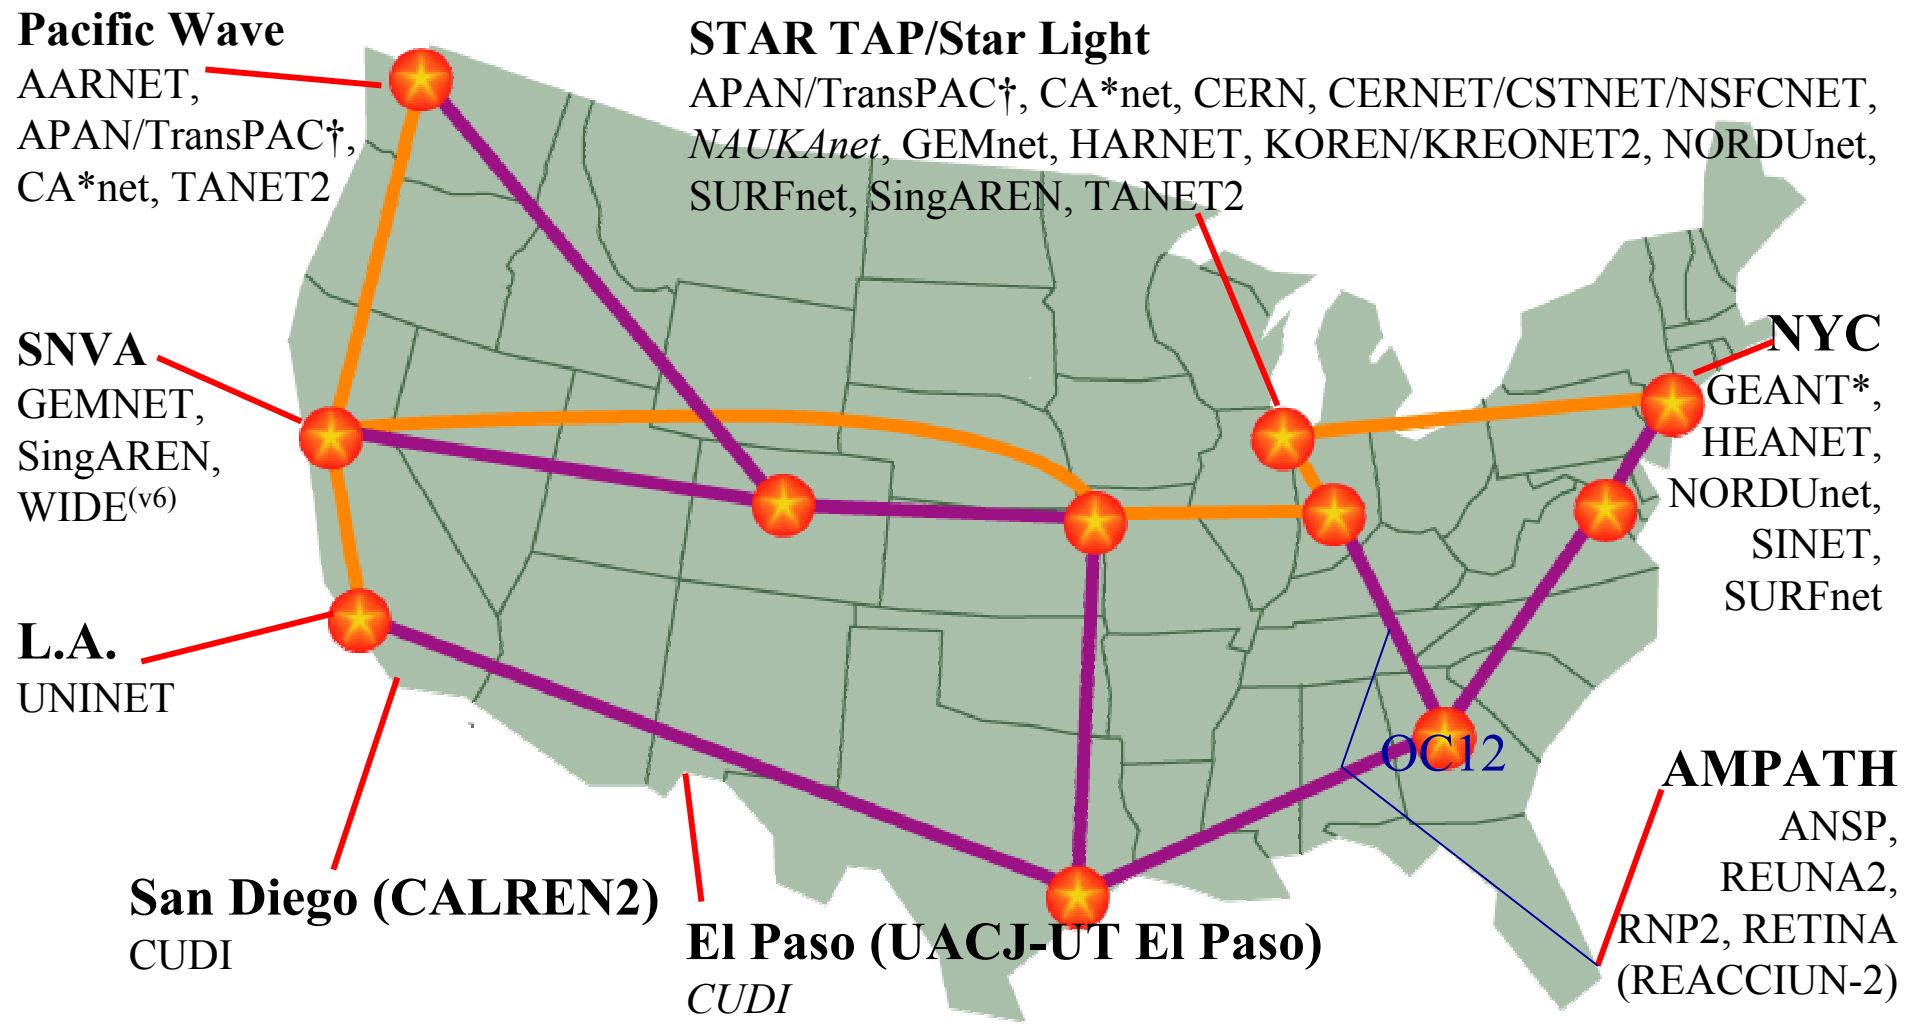

•ARNES, ACONET, BELNET, CARNET, CERN, CESnet, CYNET, DFN, EENet, GARR, GRNET, HEANET, IUCC, JANET, LATNET, LITNET, NORDUNET, RENATER, RESTENA, SWITCH, HUNGARNET, GARR-B, POL-34, RCST, RedIRIS, SANET, SURFNET

•† WIDE/JGN, IMnet, CERNet/CSTnet,/NSFCNET, KOREN/KREONET2, SingAREN, TANET2, ThaiSARN

U.S.-based International interconnection points

- STAR TAP/Star Light
  - Chicago
  - NSF-funded project
  - ATM-based STAR TAP
  - STAR LIGHT: GigE/10GE switch-based
- Pacific Wave
  - Seattle
  - Gigabit Ethernet-based
  - Pacific Northwest Gigapop
- MAN LAN
  - New York City
  - Internet2 project
  - Sited at NYSERNET colo space
  - GigE/10GE-switch based
- LALALAN
  - Los Angeles
  - Under development by CENIC
- AmPATH
  - Miami
  - Includes Global Crossing links to South America

Abilene cost recovery model - 2003

Connection (per connection)	**Annual fee**
OC-3 (155 Mbps)	$110,000
OC-12 POS (622 Mbps)	$270,000
Gigabit Ethernet (1 Gbps)	$325,000
OC-48 POS (2.5 Gbps)	$430,000
OC-192 POS/10GigE (10 G)	*$490,000*
Participation (per member)	$20,000

Higher bandwidth connections

- 10 Gigabit Ethernet (2)
  - CalREN-LA, Pacific Northwest
- OC-48c (6)
  - MAX, NCNI, NoX, OARnet, PSC, SoX
- Gigabit Ethernet (1)
  - MREN
- OC-12c (15)
  - CalREN-Sunnyvale, Front Range, Great Plains, Indiana
  - MAGPI, NYSERNet-Buffalo, OneNet, SDSC, SFGP
  - U Florida, NCSA, Northern Lights, Merit, WiscREN
  - NWVng

Backhaul

- SONET backhaul continues
  - Included as part of Abilene connection fee
  - Provided between nearest Abilene router and the closest Qwest PoP to the connector
  - OC-3c, -12c, and -48c SONET
- 10-Gbps DWDM backhaul available
  - Again part of Abilene connection fee
  - Unprotected service for now; OC-192c SONET framing
  - 3-year partial commitment required
- Ethernet connections
  - GigE and 10 GigE LAN PHY supported
  - No Ethernet backhaul capability yet
  - Dark fiber access to Abilene/Qwest PoP required

Abilene Update: Technical

Rick Summerhill
Associate Director, Backbone Network Infrastructure

Internet2 Spring Member Meeting
Washington D.C.
10 March 2003

Abilene Upgrade – April 2003 Future

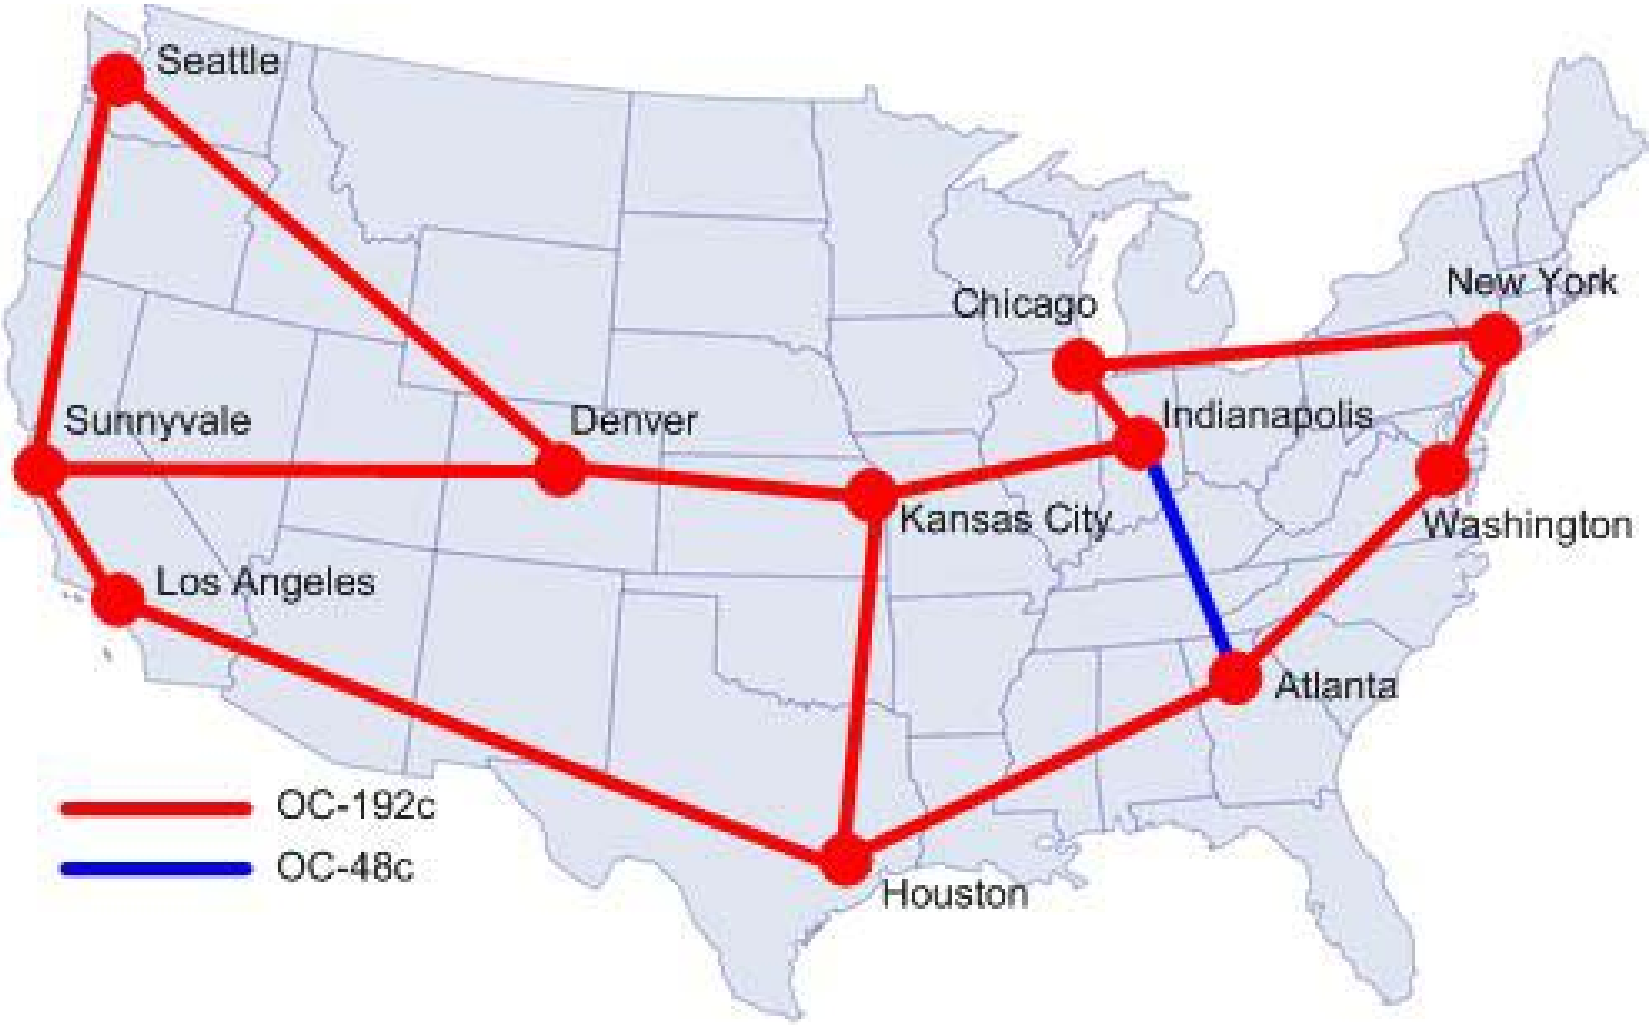

Abilene Upgrade – April 2003 Current

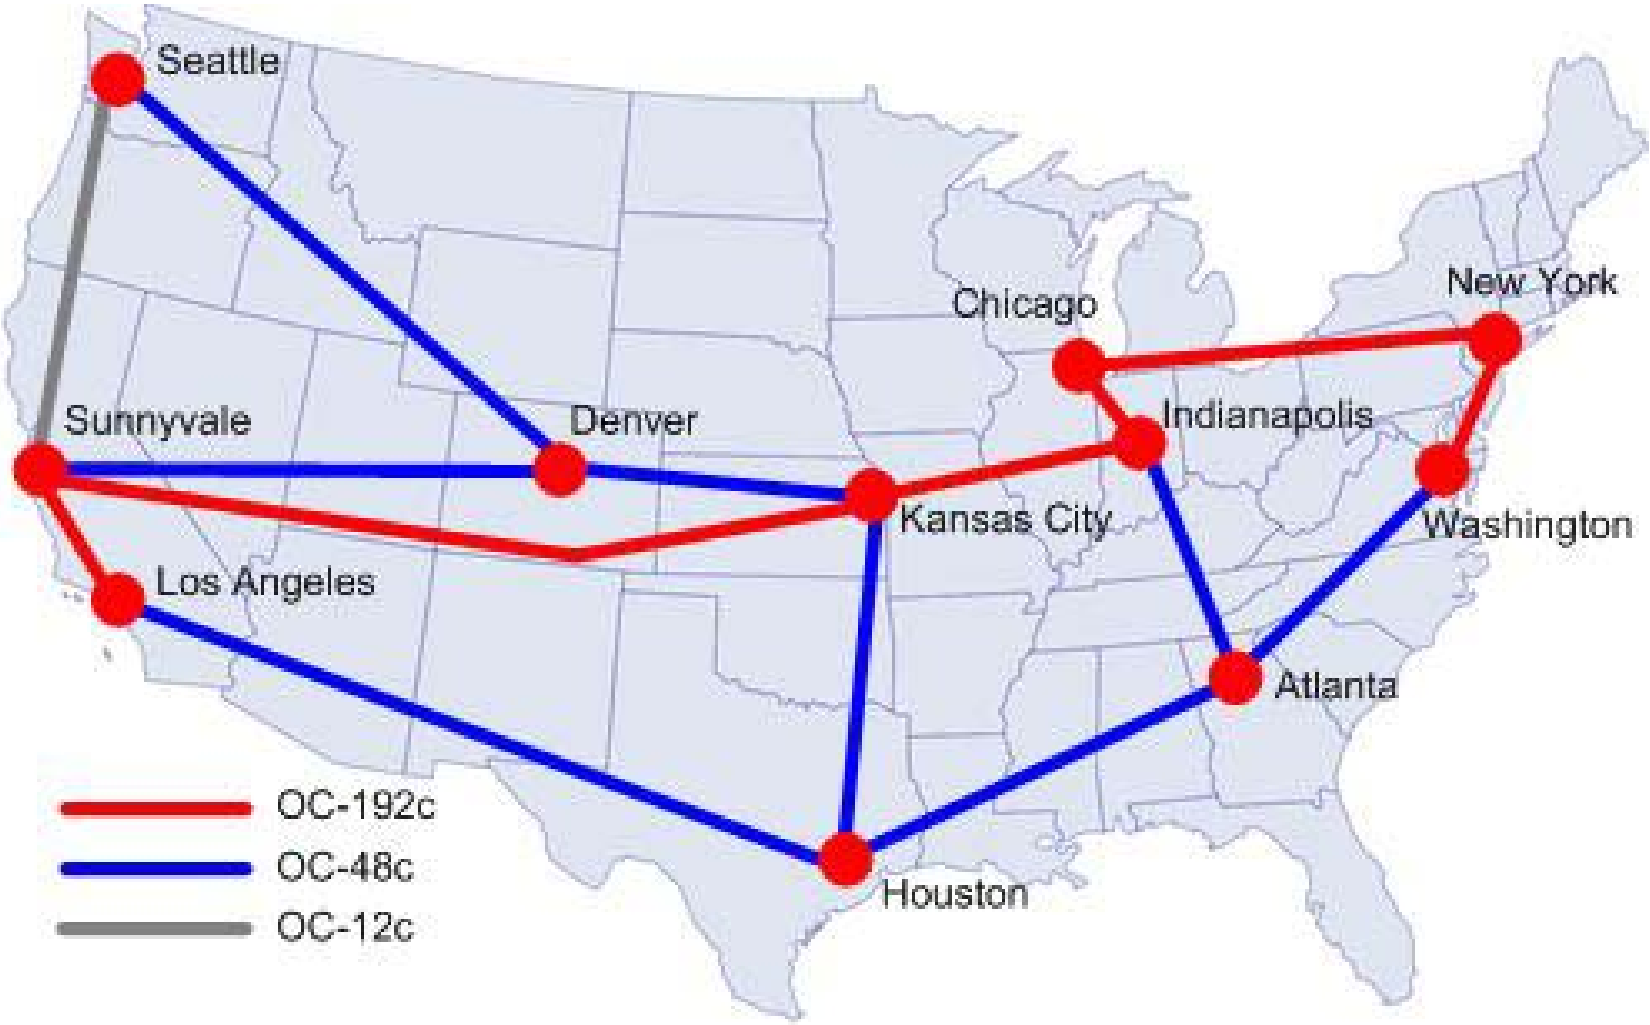

Abilene Upgrade – April 2003

- Upgrade has gone very smoothly – Kudos to the NOC, Qwest, and Juniper
- All routers installed except for Atlanta
  - Almost all connector/peer circuits have been groomed to the new routers
  - Currently removing old equipment
- Atlanta router node
  - Location now determined
  - Installation should occur within 4-5 weeks
  - Migrate the connectors to new router after installation

Abilene Upgrade – April 2003

- Backbone Circuits
  - In the coming two months, installation of the following circuits:
    - Seattle <-> Sunnyvale
    - Seattle <-> Denver
    - Kansas City <-> Houston
    - Washington <-> Atlanta
  - All nodes will have OC-192c connectivity when the above installs have been completed
  - Summer, installation of the following circuits:
    - Sunnyvale <-> Denver
    - Denver <-> Kansas City

Abilene Focus Areas - 2003

- High performance advanced services
  - Multicast
  - Native IPv6
- Facilitating end-to-end performance
- Supporting network research – Abilene Observatory
- Experimenting with MPLS/VPN on backbone
- Supporting large MTUs
- Security and the REN-ISAC

Abilene Focus Areas – 2003 Multicast

- Multicast
  - Continue to encourage high performance “flows”
  - Currently does not fall under the Abilene COU
  - Relatively simple to deploy on backbone type networks
  - Problems persist at the edges of the network
  - Provide better debugging tools
  - Abilene measurement Infrastructure to include multicast
- Internet2 workshops
  - Very successful hands-on sessions
  - Emphasis on debugging

Abilene Focus Areas – 2003 IPv6

- IPv6 on Abilene backbone
  - Fully deployed on all Abilene Routers
  - Native, dual stack service
  - Peering and connectivity does not fall under the Abilene CoU
  - Legacy tunnel connections still supported on single router at INDY
  - Tests – 8 Gbps across backbone, IPv6 only and mixed IPv6/IPv4
  - Abilene Measurement Infrastructure includes IPv6
- Internet2 workshops
  - Very successful, hands-on sessions

Abilene Focus Areas – 2003 IPv6

- Connectors/Peers
  - 16 native connections
  - 14 native peers
- IPv6 Multicast
  - Testing at North Carolina ITEC
  - SSM only, of course
- Applications
  - Almost common applications have been ported
  - Typical suspects fairly easy to deploy
    - Mail, Web, and News Servers
    - Oregon has organized a news service with v6 transport
  - Need to have large packages ported
  - Porting tools and combined v4/v6 tools in progress

Abilene Focus Areas – 2003 Facilitating E2E Performance

- Facilitating end-to-end performance
  - Bulk TCP flows
    - 2.3 Mbps (median)
    - 6.6 Mbps (90%)
    - 31 Mbps (99%)
  - Support of piPEs architecture implementation
  - Provide interactive access to measurement platform for network administrators as part of the Abilene Observatory

Abilene Focus Areas – 2003
Abilene Observatory

- A program to provide enhanced support of network research over Abilene
- Two Components of the Observatory
  - Measurements and data collected by the NOC and Engineering Team – The Abilene Measurement Infrastructure (AMI)
  - Access to Router Nodes for other projects through Collocation
- Initial web site
  - http://abilene.internet2.edu/observatory/

Abilene Focus Areas – 2003 Abilene Observatory

- Measurements taken by Abilene NOC and Engineering Teams
  - Four Dedicated Measurement PCs located in the Abilene Racks
  - Measurements are Local to the Router Nodes
  - Data eventually moved to a more central site
  - Database itself is virtual
  - Data to be correlated – Potential for the data mining projects to study basic network properties
  - CDMA Timing Devices Local to Nodes
  - Management Platform with extended access

Abilene Focus Areas – 2003
Abilene Observatory

- Types of Data
  - One-way latency, jitter, loss, reachability
  - Regular TCP/UDP throughput tests ~ 1 gbps
  - Netflow data, both in raw and summary formats (Ohio ITEC)
  - SNMP Data (NOC)
  - Routing Data, both IS-IS and BGP
  - E2E Performance Beacons, Interactive Potential
  - Optical Splitter Taps on Some Backbone Links
  - Open to recommendations for additional data types

Abilene Focus Areas – 2003
Abilene Observatory

- Access to network data archive
  - Almost all will be easily available, web services
  - Access to raw netflow data by special arrangement
    - Currently 4 research groups are accessing the raw netflow data
  - Elementary Graphical Representation of Data and Coherent views of the Abilene network
    - Summary View
    - Detailed View
  - Abilene Management Observation Platform
    - Publicly available software
    - Merit (Quilt) Participation in testing

Abilene Network Focus – 2003
Abilene Observatory

- Collocation Component
  - Space, power, access is limited, and decreases as more groups participate
  - Proposals reviewed by Abilene TAC and NRLC
  - Strict requirements concerning installation, access, and monitoring
  - Initial project is PlanetLab
    - http://www.planet-lab.org

Abilene Focus Areas – 2003 TeraGrid/MPLS Experiment

- TeraGrid
  - TeraGrid as a Computer Backplane
  - Network forms Backplane
  - Nodes consist of Computers
  - Limited peering at such nodes
  - Access to nodes is through networks such as Abilene
  - Access traffic never crosses backplane - Normal role of networks like Abilene is access to nodes
  - Backplane consists of lambdas

Abilene Focus Areas – 2003 TeraGRid/MPLS Experiment

- Abilene role
  - The problem is that a lambda is delayed from PSC node to Chicago Node
  - Can Abilene help support a temporary virtual circuit?
  - Goals
    - Use Abilene to emulate a circuit
    - PSC and Chicago use similar IP routing when lambda is in place
    - Opportunity to experiment and measure
  - Basically solving a routing issue
  - Experiment is limited in duration lasting only a few months

Abilene Focus Areas – 2003 TeraGRid/MPLS Experiment

- Solutions
  - Simple IP solution with change of routing
  - Simple GRE Tunnel would work
    - Can be setup without Abilene changes
  - Use MPLS Tunnel across Abilene, in potentially multiple ways
    - Several ways to do this, but preferable method is an MPLS tunnel across multiple administrative domains
  - Take this opportunity to do measurements

Abilene Focus Areas – 2003 TeraGRid/MPLS Experiment

- Testing
  - Lay-up configuration in lab – NC-ITEC
    - Test basic idea
  - Apply to Abilene backbone – non-intrusive
  - Opportunity for perform measurements under 3 scenarios:
    - Basic IP routing
    - MPLS solution
    - Comparison with lambda solution when installed
  - Will examine effects on normal traffic whenever possible

Abilene Focus Areas – 2003
Large MTUs on Abilene

- Theoretical Ideas
  - Matt Mathis formula: Throughput is directly proportional to MTU
  - To achieve large flows, support large MTUs
  - Relatively easy to support on backbones or regional networks
  - More difficult closer to the edges of the network
  - Fairly easy to support 9K MTUs, but larger MTUs are almost impossible at this time
    - Most new Ethernet NICs support 9K frames
    - Not all switches support large frames
    - Backbones and connector networks should support at least 9K frame sizes

Abilene Focus Areas – 2003
Large MTUs on Abilene

- Internet2 recommendations
  - Support 9K MTUs across the entire infrastructure
  - Larger MTUs when possible
- Very Large MTUs on Abilene
  - We may be able to support large MTUs on SONET backbone links
  - However, how do we get packets into the network?
  - Some test equipment available, but none with “stacks”. That is, only have packet blasters that use SONET interfaces.
  - Interest in gaining access to such devices
  - Investigating several candidates

Abilene Focus Areas – 2003 Security

- Security on backbone networks
  - Meaning and implications
  - Writing a security document
    - Preparations for security issues
    - Incidence response
    - Incidence reporting
  - Relationships to REN-ISAC
    - Information Sharing Analysis Center
      - facilitate communication, develop best practices, and disseminate security-related information
    - Define Abilene/NOC/REN-ISAC relationships
    - Policy and procedures concerning securities issues
  - Especially important to define this on an network providing advanced services

More information

- http://www.internet2.edu/abilene
- abilene@internet2.edu

Potential discussion topics

- Current Abilene focus areas
  - Performance (large MTUs)
  - Pushing Native IPv6 to the edge
  - Measurement: Abilene Observatory
  - MPLS VPN experiments
  - Security
- Conditions of Use (CoU)
  - Sponsored participants & SEGPs
  - Federal labs
- ‘Logarithmic’ pricing model
- Facilitating upgrades
- Backhaul
- Regional optical initiative interface
  - Dark fiber access into Qwest PoPs


INTERNET®
www.internet2.edu