



Internet2 Advanced Networking Initiatives

QUESTnet2000 - USQ, Toowoomba

Ben Teitelbaum
June 27th, 2000






Internet2 Overview



1/3 Internet2
Overview



Elevator Explanation

Mission


- *Develop and deploy advanced network applications and technologies, accelerating the creation of tomorrow's Internet*

Goals

- Enable new generation of advanced applications
- Re-create leading edge R&E network capability
- Transfer capability to global production internet

Means

- Advanced Applications
- Advanced Network Infrastructure
- New Network Capabilities
- Middleware
- Partnerships (Government, Industry, International)




Yesterday's Internet

Thousands of users

Remote login, file transfer

Applications capitalize on underlying technology




Today's Internet

Millions of users

Web, email, low-quality audio & video

Applications adapt to underlying technology




Tomorrow's Internet

Billions of users and devices

Convergence of today's applications and services

New technologies enable unanticipated applications (and create new challenges)




Virtual Laboratories

*Real-time access to
remote instruments*

*University of Pittsburgh,
Pittsburgh Supercomputing
Center*

3-D Brain Mapping






Virtual Laboratories

*Real-time access to
remote instruments*

*University of North
Carolina, Chapel Hill*

Distributed
nanoManipulator



Shared virtual reality

University of Illinois at Chicago

Virtual
Temporal
Bone



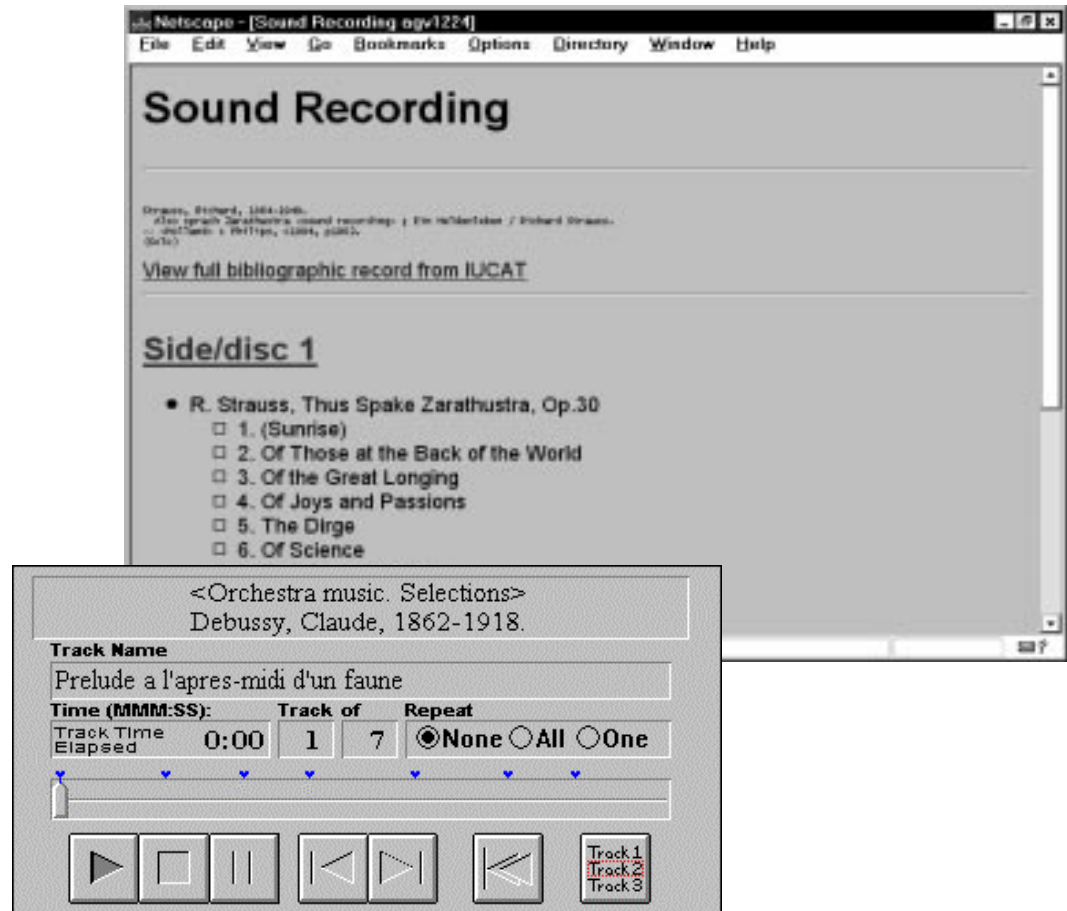



Source: University of Illinois-Chicago

Video and audio

Indiana University


Variations Project





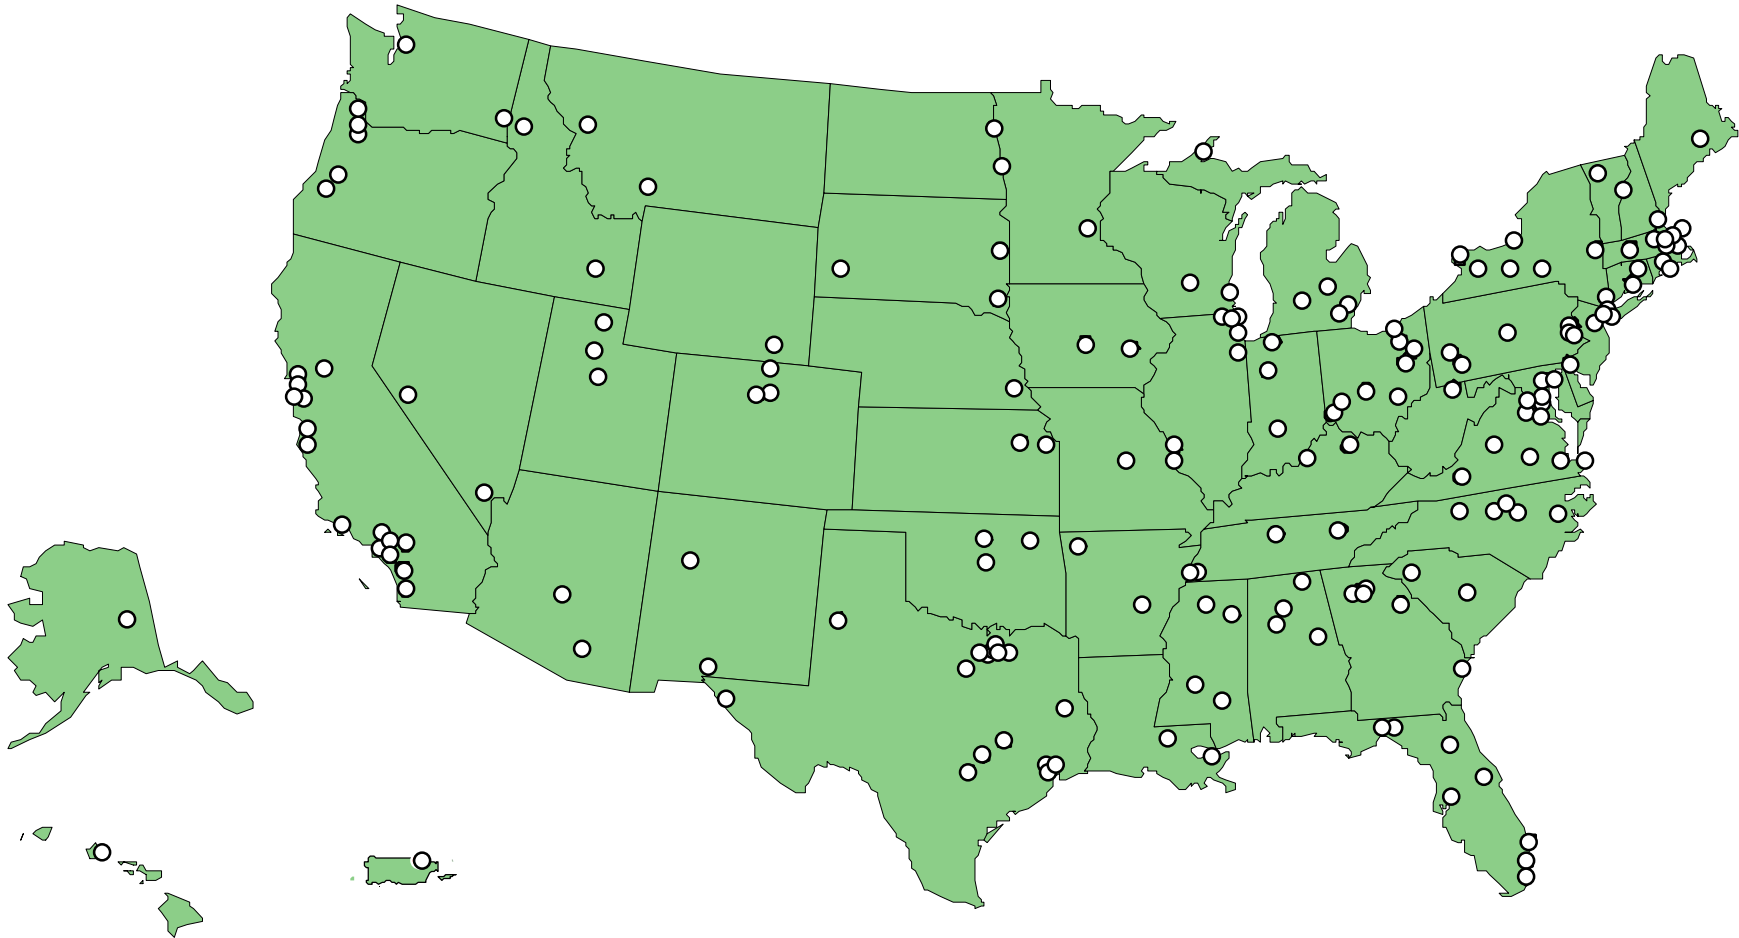
*Large-scale
computation*

*University
Corporation for
Atmospheric
Research*




Internet2 Universities

179 Universities as of June 2000



+ over 30 affiliate members



Internet2 Corporate Partners

3Com

*Advanced Network &
Services*

Alcatel

Ameritech

AT&T

Cabletron Systems

Cisco Systems

IBM

ITC^Deltacom

Lucent Technologies

Marconi

MCI Worldcom

Microsoft

Newbridge Networks

Netcom Systems

Nortel Networks

Qwest Communications

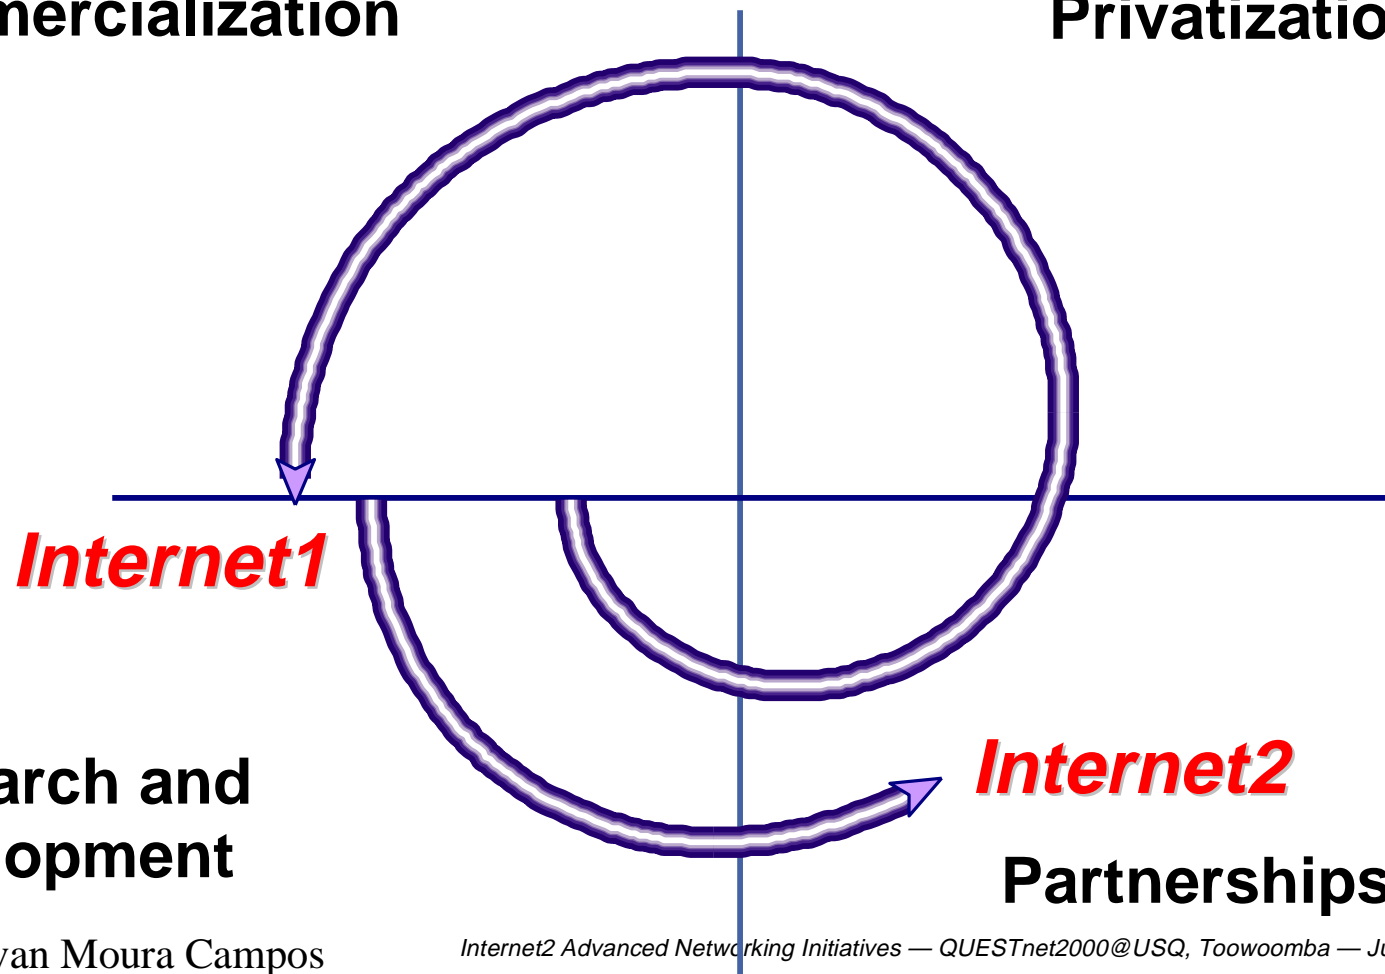
WCI Cable

+ over 70 corporate members

Internet Development Spiral

Commercialization


Privatization



**Research and
Development**

Internet2


Partnerships



Internet2 International Activities

Ensure global interoperability of advanced networking technologies and applications

Enable collaborations between US researchers at Internet2 institutions and their non-US counterparts




International Collaborations

Building peer to peer relationships


Looking for similar goals/objectives and similar constituencies

Mechanism: Memoranda of Understanding



International MoU Partners 32 as of June 2000


AAIREP (Australia)
APAN (Asia-Pacific)
APAN-KR (Korea)
ARNES (Slovenia)
BELNET (Belgium)
CANARIE (Canada)
CESnet (Czech Republic)
CERNET, CSTNET, NSFC (China)
CUDI (Mexico)
DANTE (Europe)
DFN-Verein (Germany)
Fundacion Internet 2 Argentina (Argentina)
GIP RENATER (France)
GRNET (Greece)
HEAnet (Ireland)
HUNGARNET (Hungary)
INFN-GARR (Italy)
Israel-IUCC (Israel)
JAIRC (Japan)
NORDUnet (Nordic countries)
POL-34 (Poland)
RCCN (Portugal)
RedIRIS (Spain)
RESTENA (Luxembourg)
REUNA (Chile)
RNP2 (Brazil)
SingAREN (Singapore)
Stichting SURF (Netherlands)
SWITCH (Switzerland)
TAnet2 (Taiwan)
TERENA (Europe)
JISC/UKERNA (UK)



Australian Advanced Internet Research and Education Program (AAIREP)

- *Australian Academic and Research Network (AARNet)*
- *Commonwealth Scientific and Research Organisation (CSIRO)*
- *Defence Science and Technology Organisation (DSTO)*
- *Australian Nuclear Scientific and Technology Organisation (ANSTO)*
- *Australian Institute of Marine Science (AIMS)*
- *Australian Partnership for Advanced Computing (APAC) initiative*
- *Advanced Computational Systems Cooperative Research Centre (ACSys)*

Motivate




Enables




Applications


Motivate

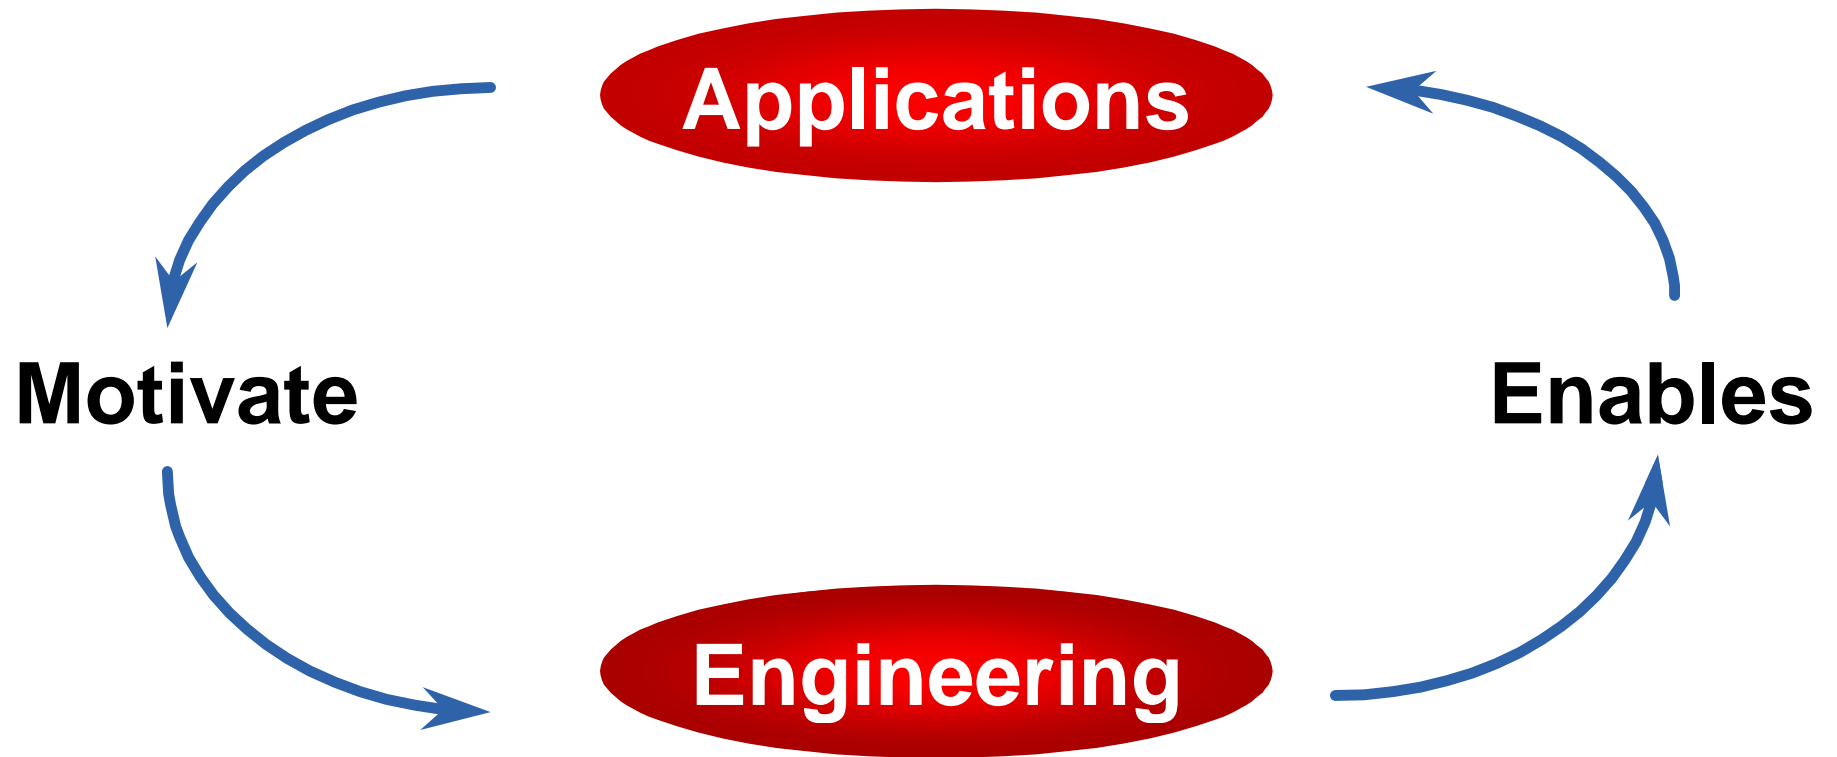



Enables



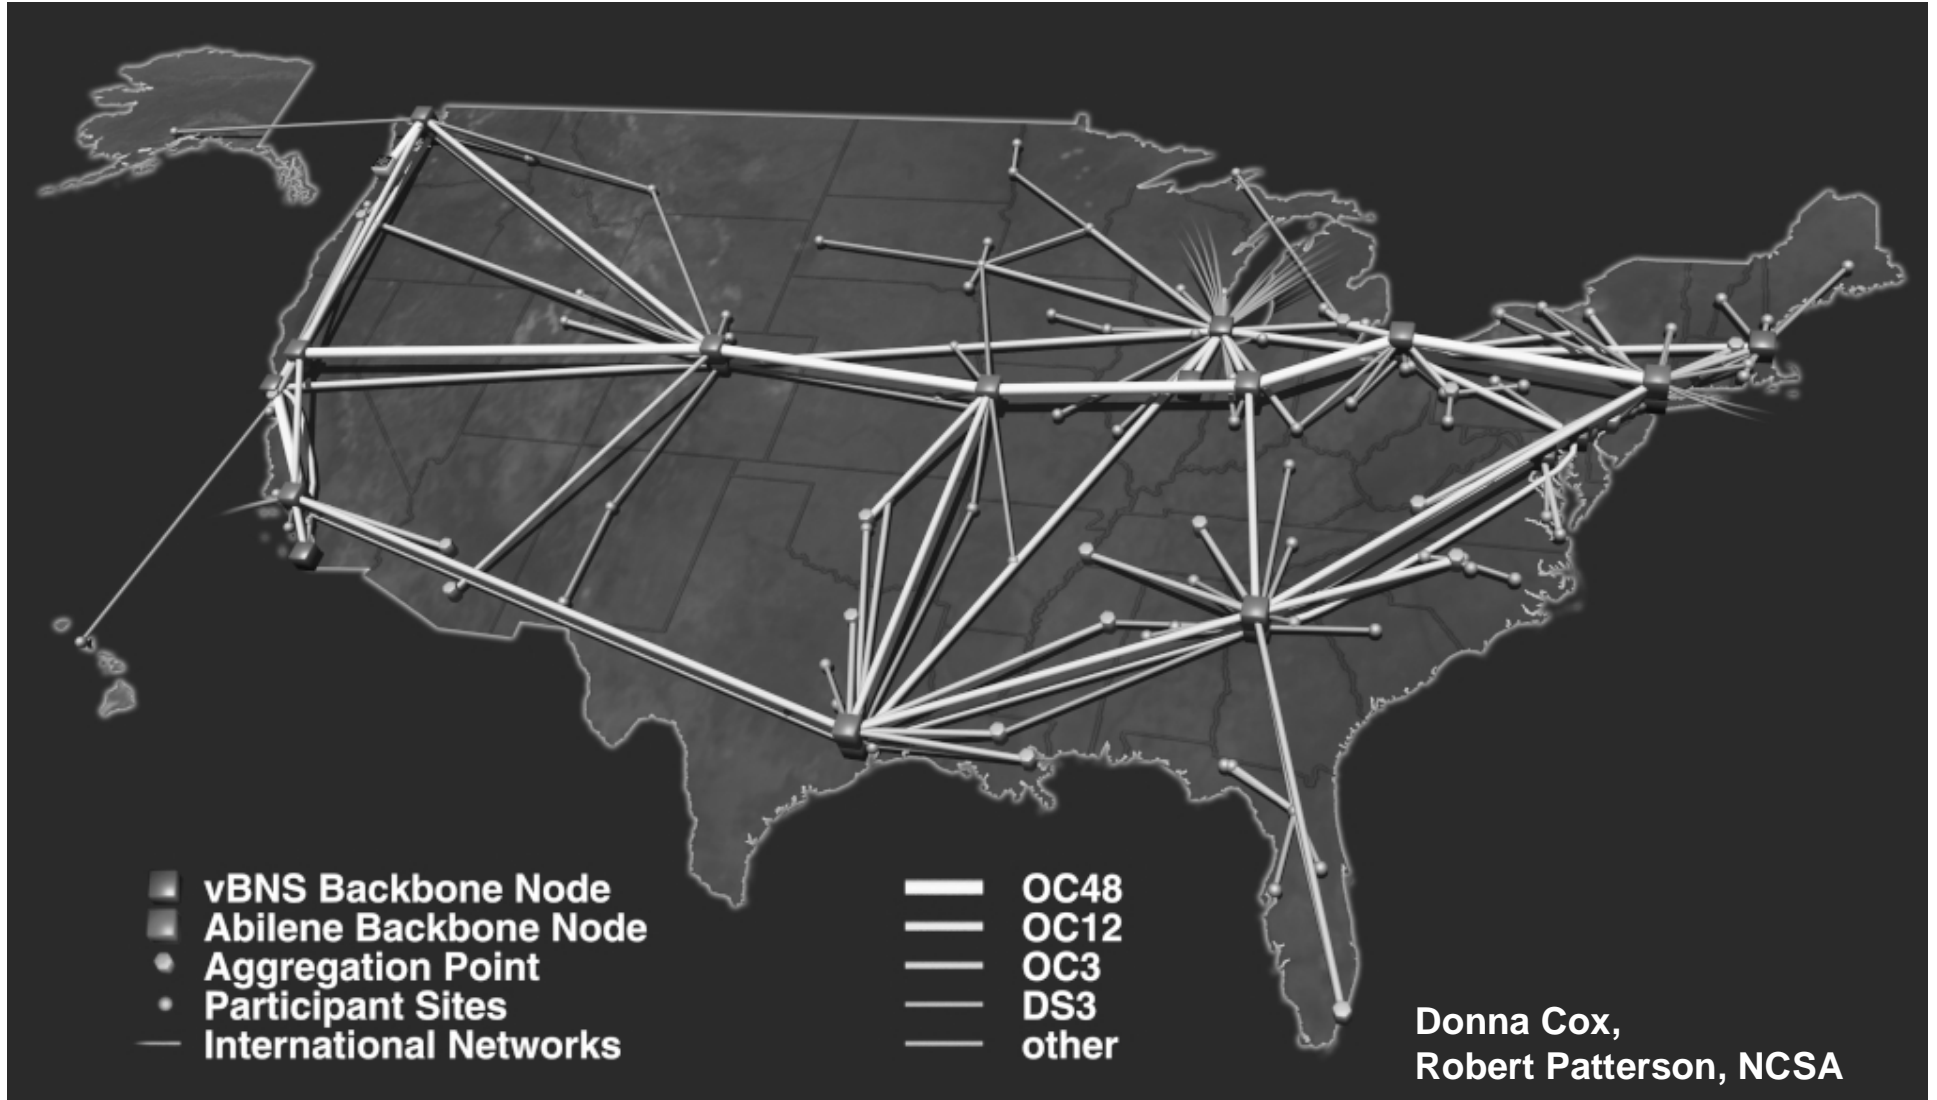






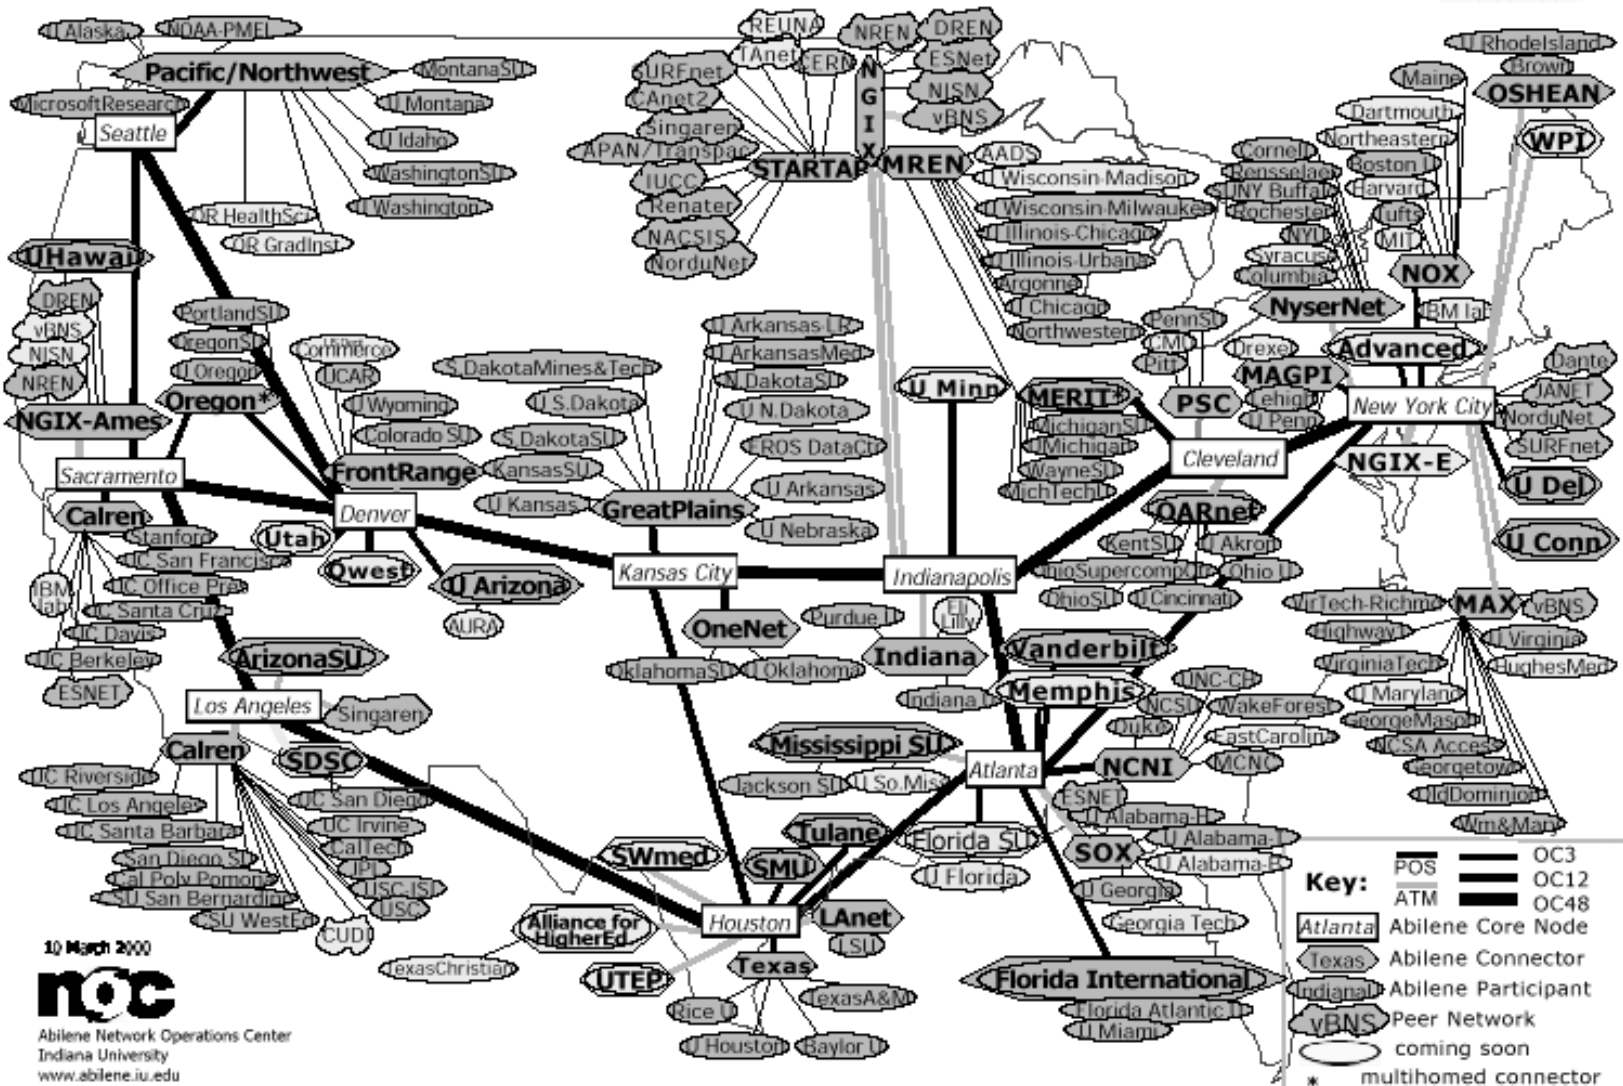



Internet2 Backbone Networks



Abilene Topology

March, 2000





STAR TAP International Transit Network (ITN)

STAR TAP has NSF funds to provide coordination for international connectivity

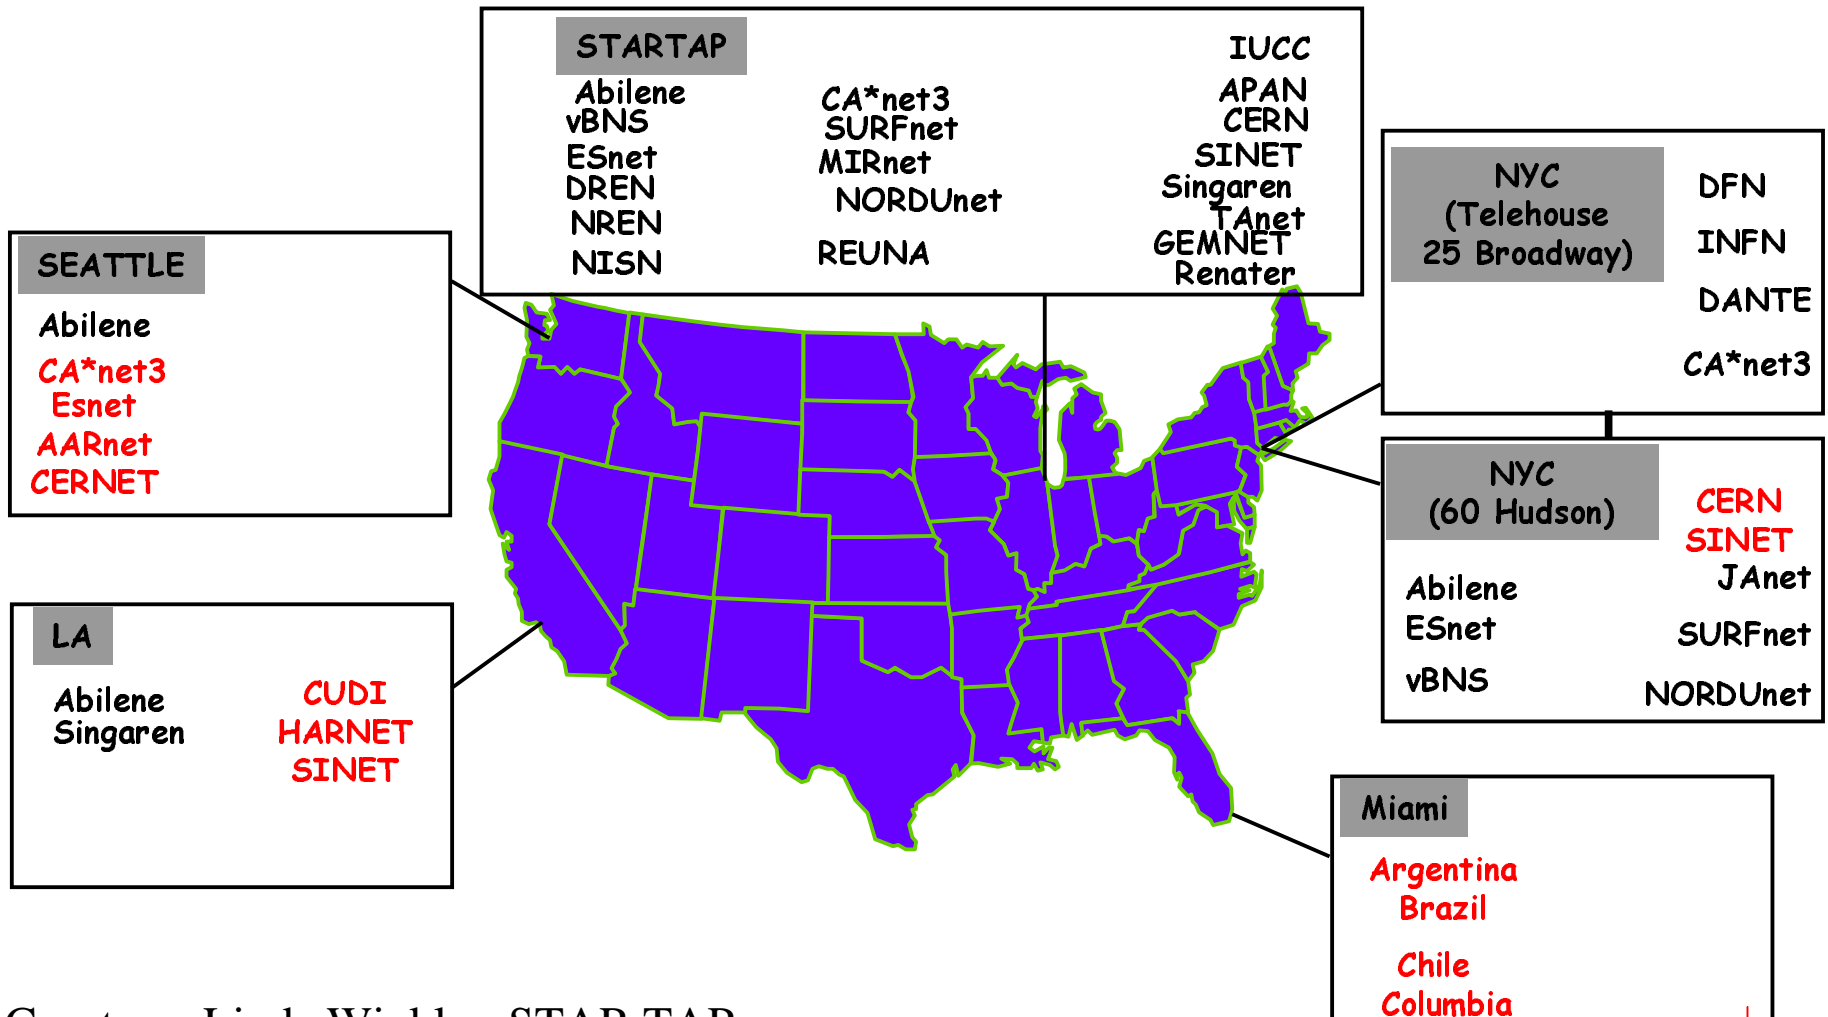
- UIC has cooperative agreement - Tom DeFanti

STAR TAP has also provided physical interconnection point - Chicago


How to get connectivity between these several peering places?

- Abilene and CANARIE (CANet) can contribute by providing transit
- STAR TAP continues coordination role
- Working on case by case basis to see how this rolls out

Developing International Peering





Courtesy: Linda Winkler, STAR TAP



New Network Capabilities

1/3 New
Network
Capabilities





New Network Capabilities

Quality of Service

- <http://www.internet2.edu/qos>

Scalable IP Multicast

- <http://www.internet2.edu/multicast>

IPv6

- <http://www.internet2.edu/ipv6>

Distributed Storage: I2-DSI

- <http://dsi.internet2.edu/>

Middleware

- www.internet2.edu/middleware/

Native multicast fully deployed / supported in the backbone

- Confidence in proper operation
- Good monitoring and reporting tools

Working to bring gigaPoPs / campuses online

- Trying to develop a formula for facilitating connections
- Developed new series of deployment tutorials

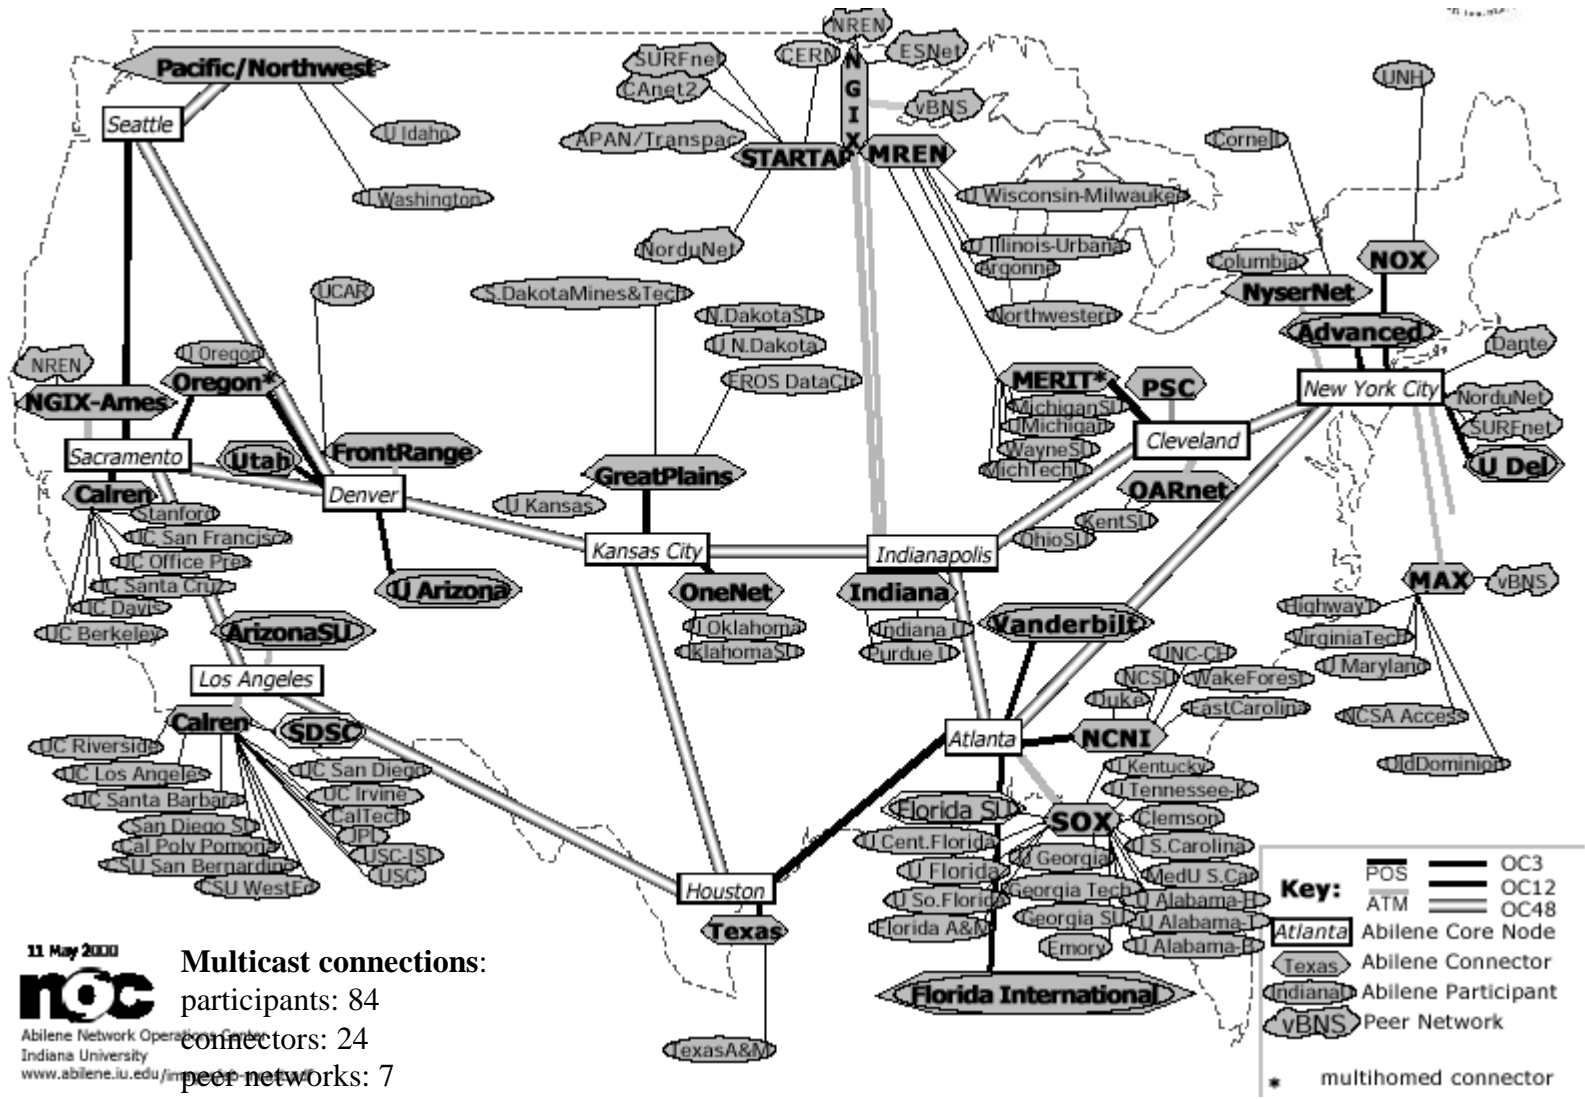
Peerings

- Including better peering with the commodity Internet

Support for network-wide events to encourage deployments (e.g. NetAid)

Abilene Multicast Deployment

(includes MBGP, MSDP, and PIM-sparse protocols)




11 May 2000



Abilene Network Operations Center
Indiana University
www.abilene.iu.edu/im...

Multicast connections:
participants: 84
connectors: 24
peer networks: 7



Future Focus of I2 Multicast WG

Key now is to push the next round of deployment

Deployment in gigaPoPs and campuses

Develop additional monitoring tools

Develop mechanisms for interacting with the commodity Internet

Track, encourage, and support high-speed applications

Future proofing your network

- You will run out of IPv4 addresses
- It is only a matter of time

New demands are going to be placed on your networks

- Mobility
- Security
- Multi-homing
- Integration of new devices


Set of core IPv6 routers managed and operated by the Abilene NOC

- Core consists of 4 Cisco 7200's donated by Cisco
 - Pittsburgh
 - Indianapolis
 - Atlanta
 - Denver
- Core routers configured in a full mesh of IPv4 tunnels
- Running 12.0 (x) IOS
- Native mode connection to the 6Tap located at MREN
- Currently peering with vBNS IPv6 network

3Com has also donated 6 routers for deployment within gigapops

These will serve as connection points for the gigapops and will connect to the core Ipv6 network

These can be either native or tunneled connections to the core



Internet2 Distributed Storage Infrastructure (I2-DSI)

Replicated hosting service for Internet content and applications

- Making storage an advance network resource
- Improved storage locality → increased performance

IS-DSI Channels

- CPAN Comprehensive Perl Archive Network
- Docsouth Documenting the American South (UNC-CH AAL)
- High MPEG High bandwidth MPEG-1 videos
- MetaLab Comprehensive Linux repository
- Mars Mars'98 Polar Lander mission. (NASA JPL)
- Netlib Mathematical software, papers, and databases. (UTK ICL)
- Open Video Video content for the research community (UNC-CH IDL)

Future: Low-Level Network Storage API

- Internet Backplane Protocol (IBP)

Replication

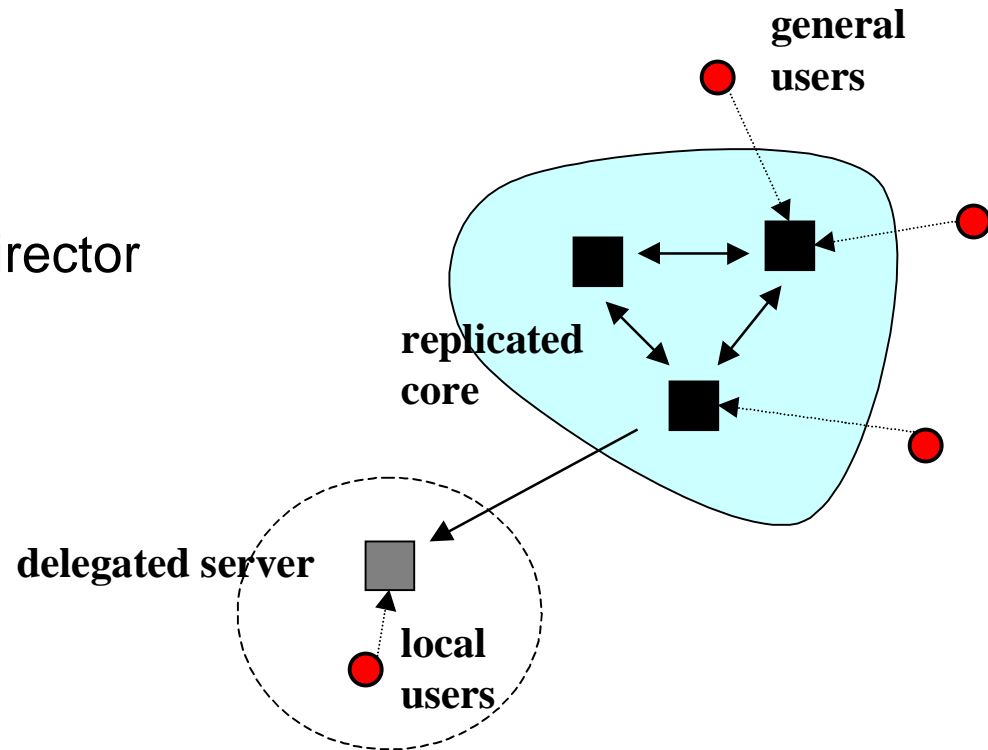
- Rsync

Resolution

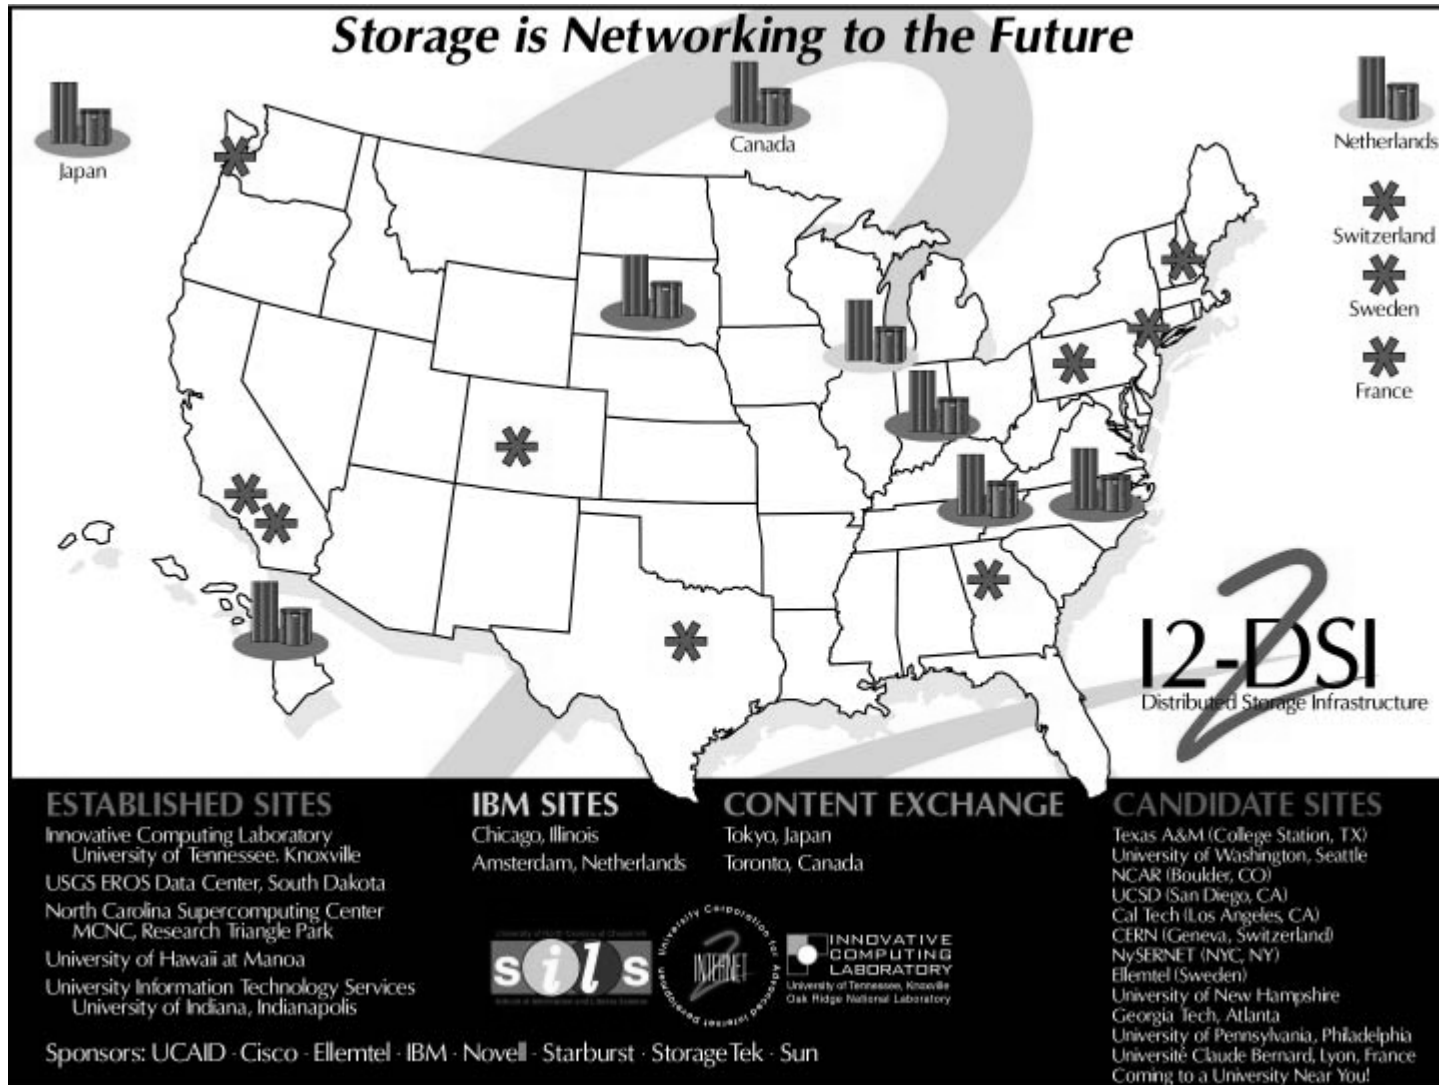
- Cisco Distributed Director

Delegation

- Streaming video only




I2-DSI: Current Deployment



A layer of software between the network and the applications


- Authentication
- Identification
- Authorization
- Directories
- Security




Internet2 Middleware Initiative

Internet2 community has unique needs and capabilities

- Middleware Architecture Committee for Education
- Early Harvest and Early Adopters
- PKI
- Shibboleth (authentication)
- Computational middleware (Beta Grid)
- Medical middleware
- Directories



Internet2 QoS




Status

- Continued focus on QBone Premium Service to enable loss/jitter intolerant applications

How We Got Here (short version)

- Began chanting:
“enable advanced applications,
enable advanced applications,…”
- Assessed requirements
- Recommended DiffServ
- Charted QBone initiative
- Specified QBone architecture
- Proceeded to implement it



Applications contract for specific QoS profiles

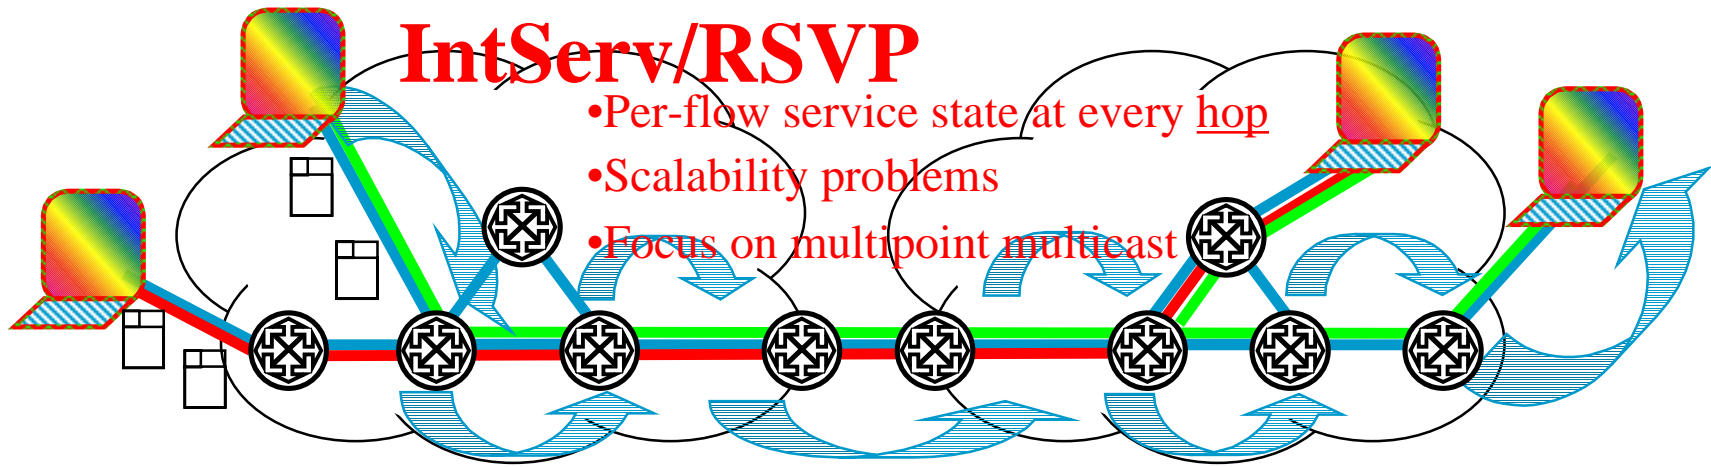
- Policing at network periphery
- “Color” packets with a few simple, differentiated per-hop forwarding behaviors (PHBs)
 - Indicated in packet header
 - Applied to PHB traffic aggregates
- **PHBs + policing rules = range of services**

DS domains contract with each other for aggregate QoS traffic profiles

- Policing at cloud-cloud boundary
- Supports simple, bilateral business agreements

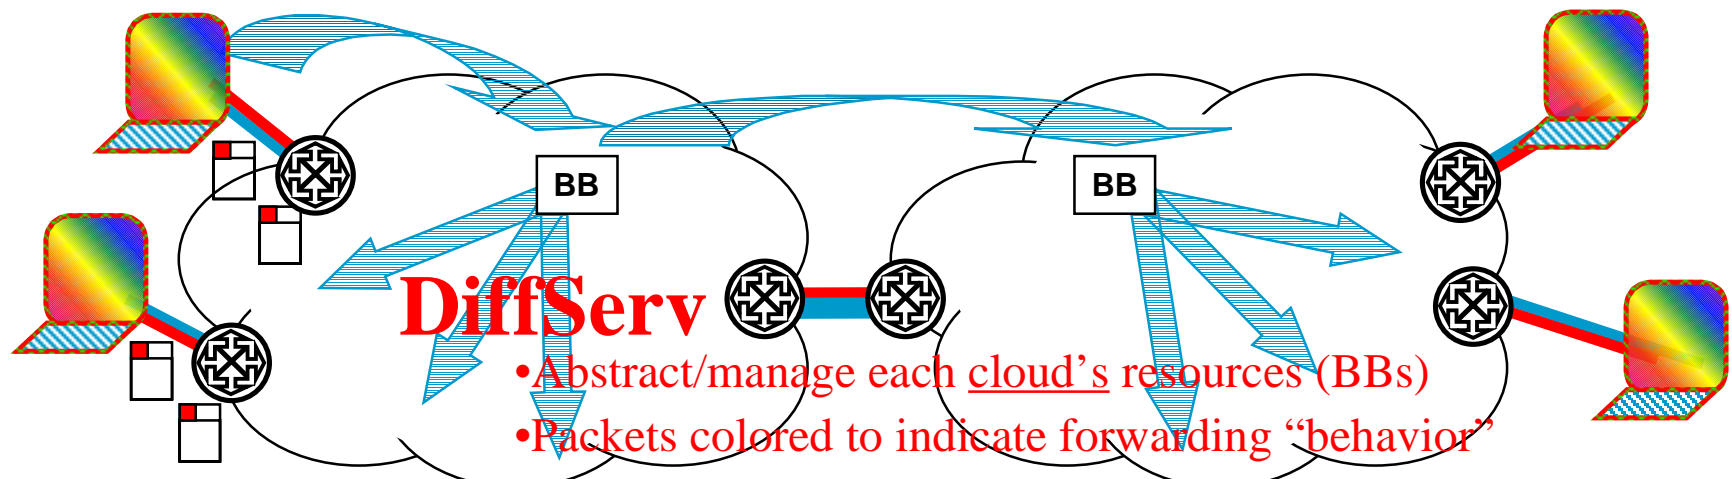
Exploits edge/core distinction for scalability

IntServ/RSVP vs DiffServ



IntServ/RSVP

- Per-flow service state at every hop
- Scalability problems
- Focus on multipoint multicast



DiffServ

- Abstract/manage each cloud's resources (BBs)
- Packets colored to indicate forwarding "behavior"
- Focus on aggregates not individual flows
- Policing at edge to get services

Assurance: like a leased line

PHB: Expedited Forwarding (RFC 2598)

- EF in separate queue configured with minimum departure rate
- Example mechanisms: strict priority, MDRR, WFQ

Policing: police to a specified peak rate and drop out-of-profile packets; effectively a leaky bucket with depth 1 MTU

Example Service #2: Controlled Load

Assurance: network looks “lightly-loaded”
for conforming traffic

PHB: Assured Forwarding (RFC 2597)

- 4 independent AF classes
- 3 drop preference levels within each class
- Example mechanisms: WRED, WFQ

Policing: police to specified rate and burst
profile, remarking out-of-profile packets to
have higher drop probability

Assurance: “better than Joe”

PHB: “drop the lower classes first ” (AF or class selector PHBs)

Policing: could be based on anything (e.g. higher priority for the CEO)

A.K.A. “Olympic” classes of BE service
(e.g. Gold, Silver, Bronze)

Specify an interdomain testbed architecture

- Leverage IETF DiffServ work; leading where appropriate


Build an interdomain testbed infrastructure

- Built on production R&E infrastructure
- Must balance networking **research** with providing a **service**

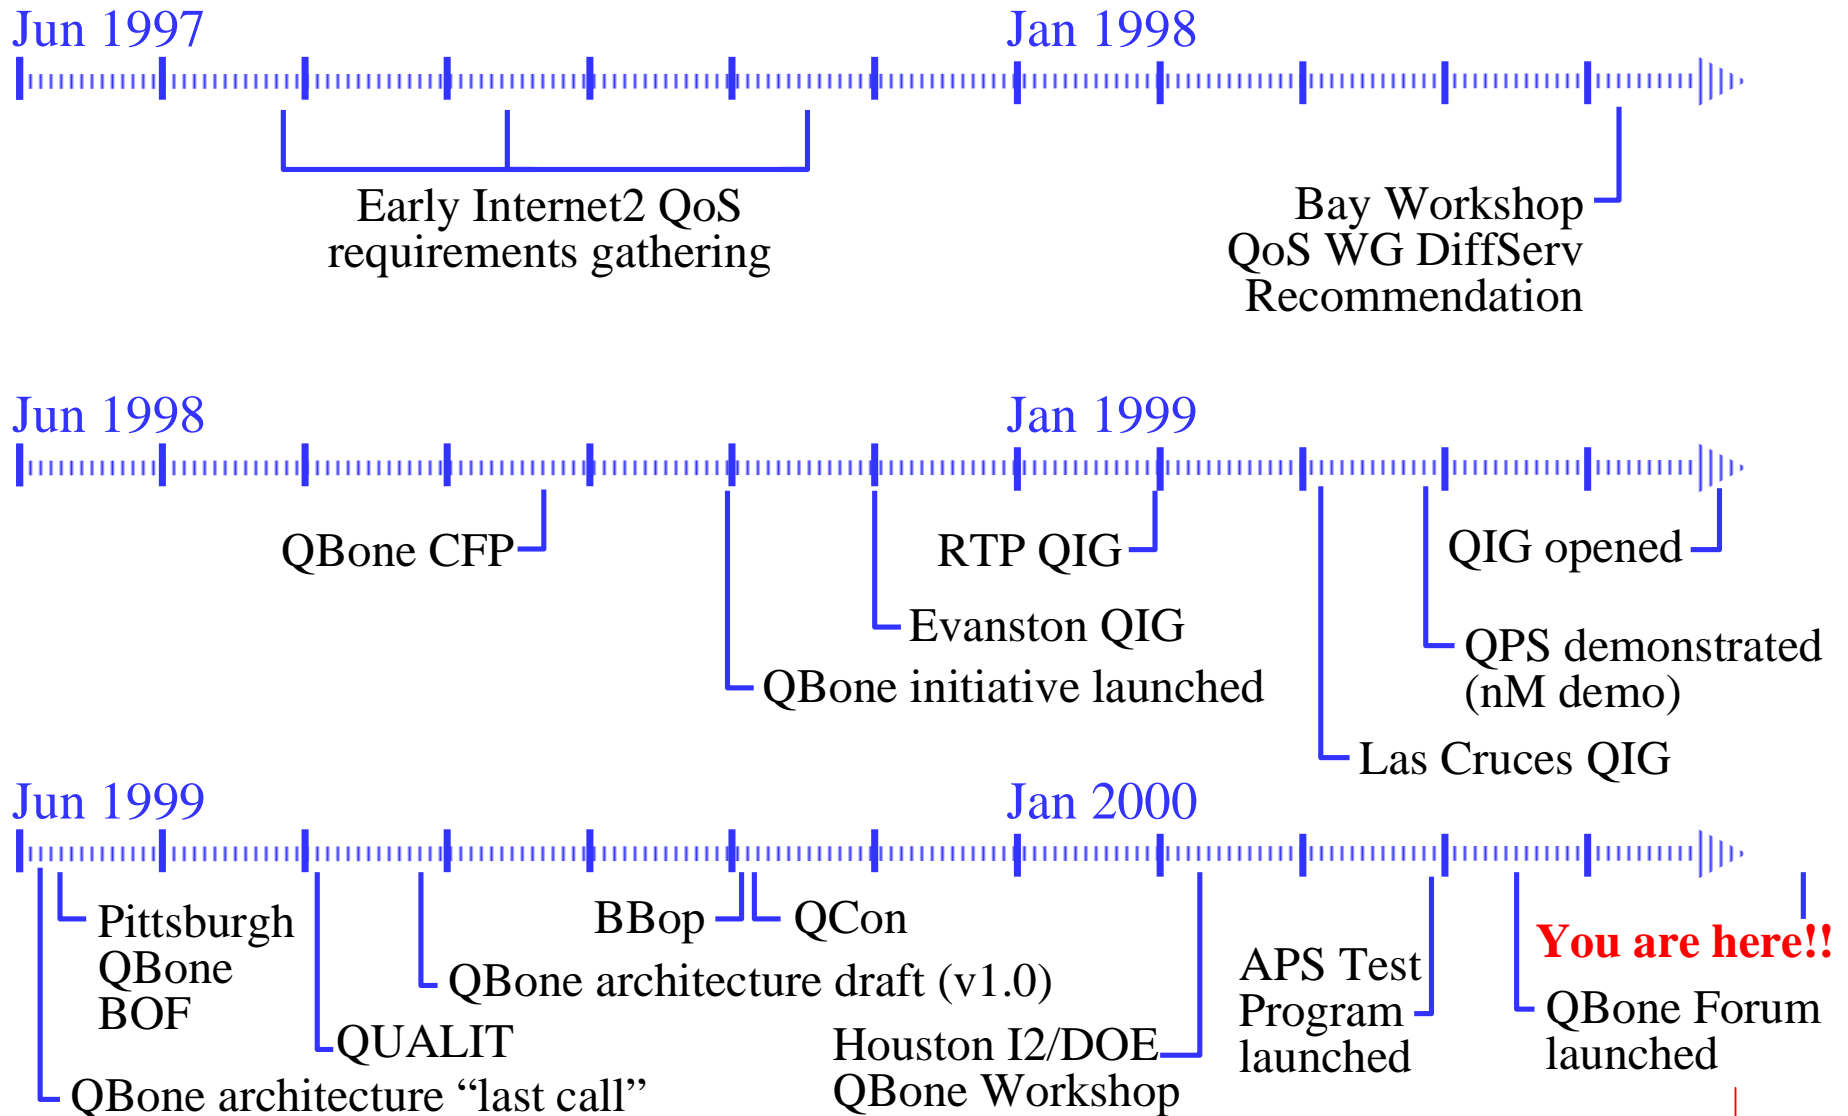
Communications and community building

- Facilitate exchange of information / experience
- Encourage inter-disciplinary / inter-institutional collaborations
- Examples: workshops, mailing lists, QBone Forum
- Cool logo:





How We Got Here (long version)



A Service: QBone Premium Service

- Built on Expedited Forwarding (EF) (RFC 2598)
- Assurance: near-zero loss & low, bounded jitter for marked traffic conforming to a specified peak rate
 - a.k.a. “virtual leased line”, “virtual wire”

Reservation Setup Protocol

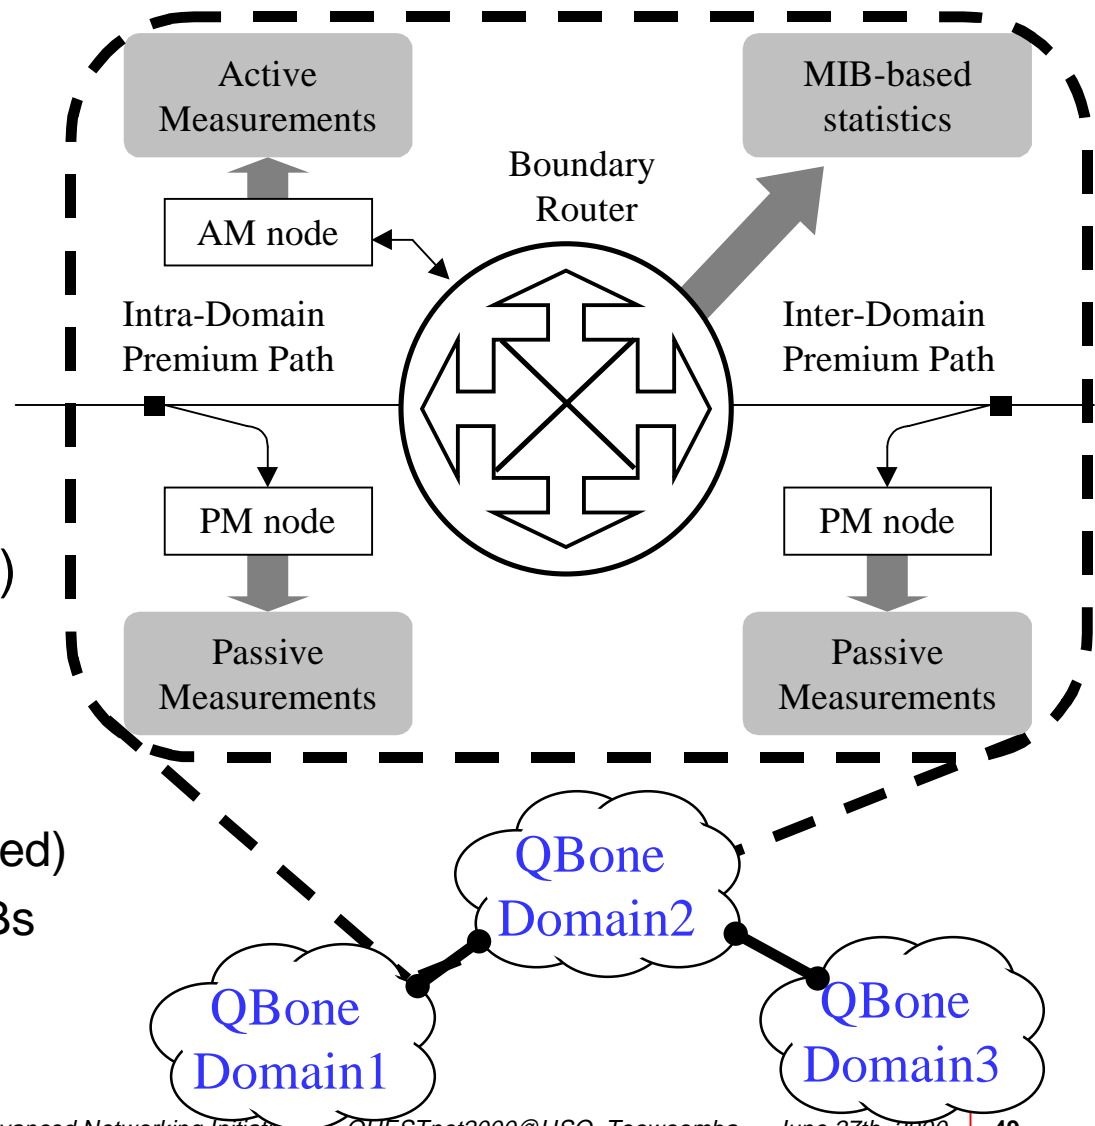
- **Now:** long-lived, manual setup
- **Proposed:** SIBBS protocol between QBone domains; RSVP end-to-end between hosts

QBone Measurement Architecture

- Uniform collection of QoS metrics
- Uniform dissemination interface

Collection

- \forall metrics, EF and BE...
- Active metrics (paths)
 - One-way delay-variation
 - One-way loss
 - Traceroutes
 - e.g. IPPM Surveyors
- Passive metrics (interfaces)
 - Load
 - EF reservation load
 - Discards (suggested)
 - Link bandwidths (suggested)
 - e.g. OCxMon, RTFM, MIBs



Dissemination

- real-time + archived measurements
- Canonical names for:
 - Metrics –Interfaces
 - Domains –Measurement nodes
- Standard metric aggregations:
 - Mostly 5-minute aggregations
- Standard URL name space for data:
 - `<root_URL>/<source_domain>/<dest_domain>/<first_hop>/<date>/<type>.<aggregation>.[suffix]`
- Future: mappings from IP addresses to QBone domains / measurement nodes via DNS or whois

SIBBS - Simple Interdomain Bandwidth Broker Signaling

Basic Idea

- Simple protocol for one QBone domain to request of another QBone domain an increase or decrease in an aggregate reservation of a globally well-known service


Design Goals

- Simple & extensible
- New protocol for now; looking at mapping to COPS


Status

- Protocol draft nearing completion
(see: <http://qbone.ctit.utwente.nl/BBroker/bboutline2.html>)

SIBBS: Basic Reservation Setup

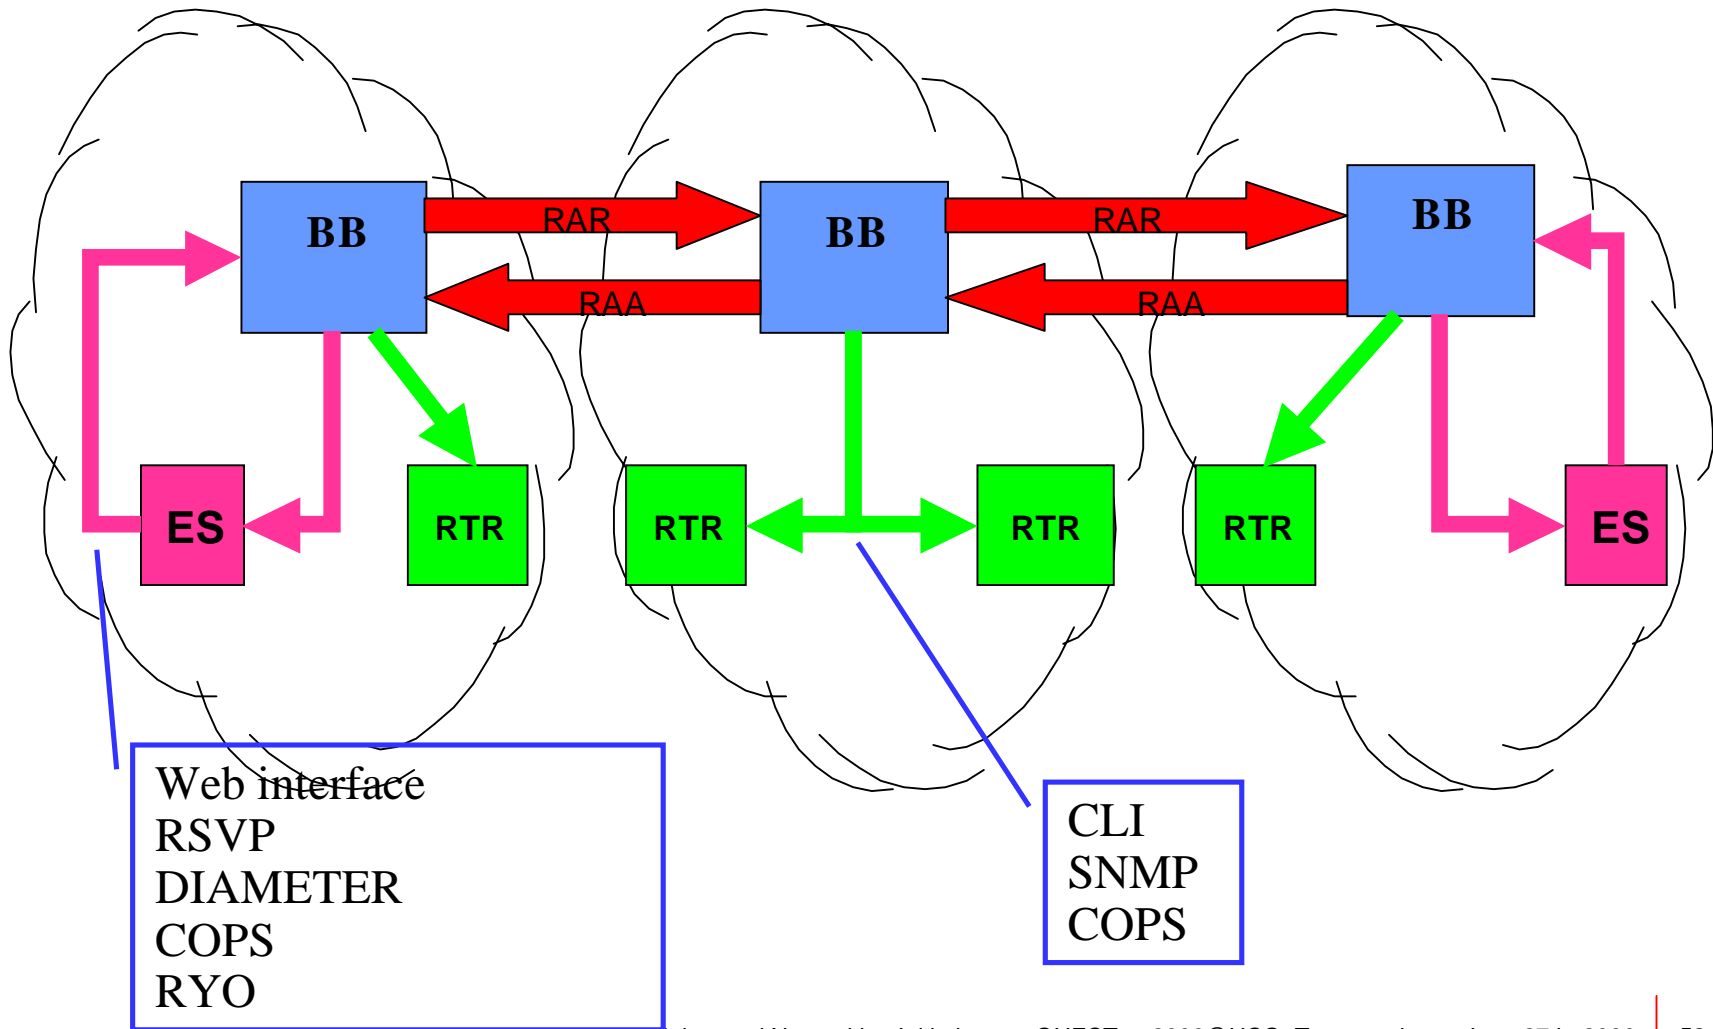


SIBBS: Basic Reservation Setup



- Web interface
- RSVP
- DIAMETER
- COPS
- RYO


SIBBS: Basic Reservation Setup



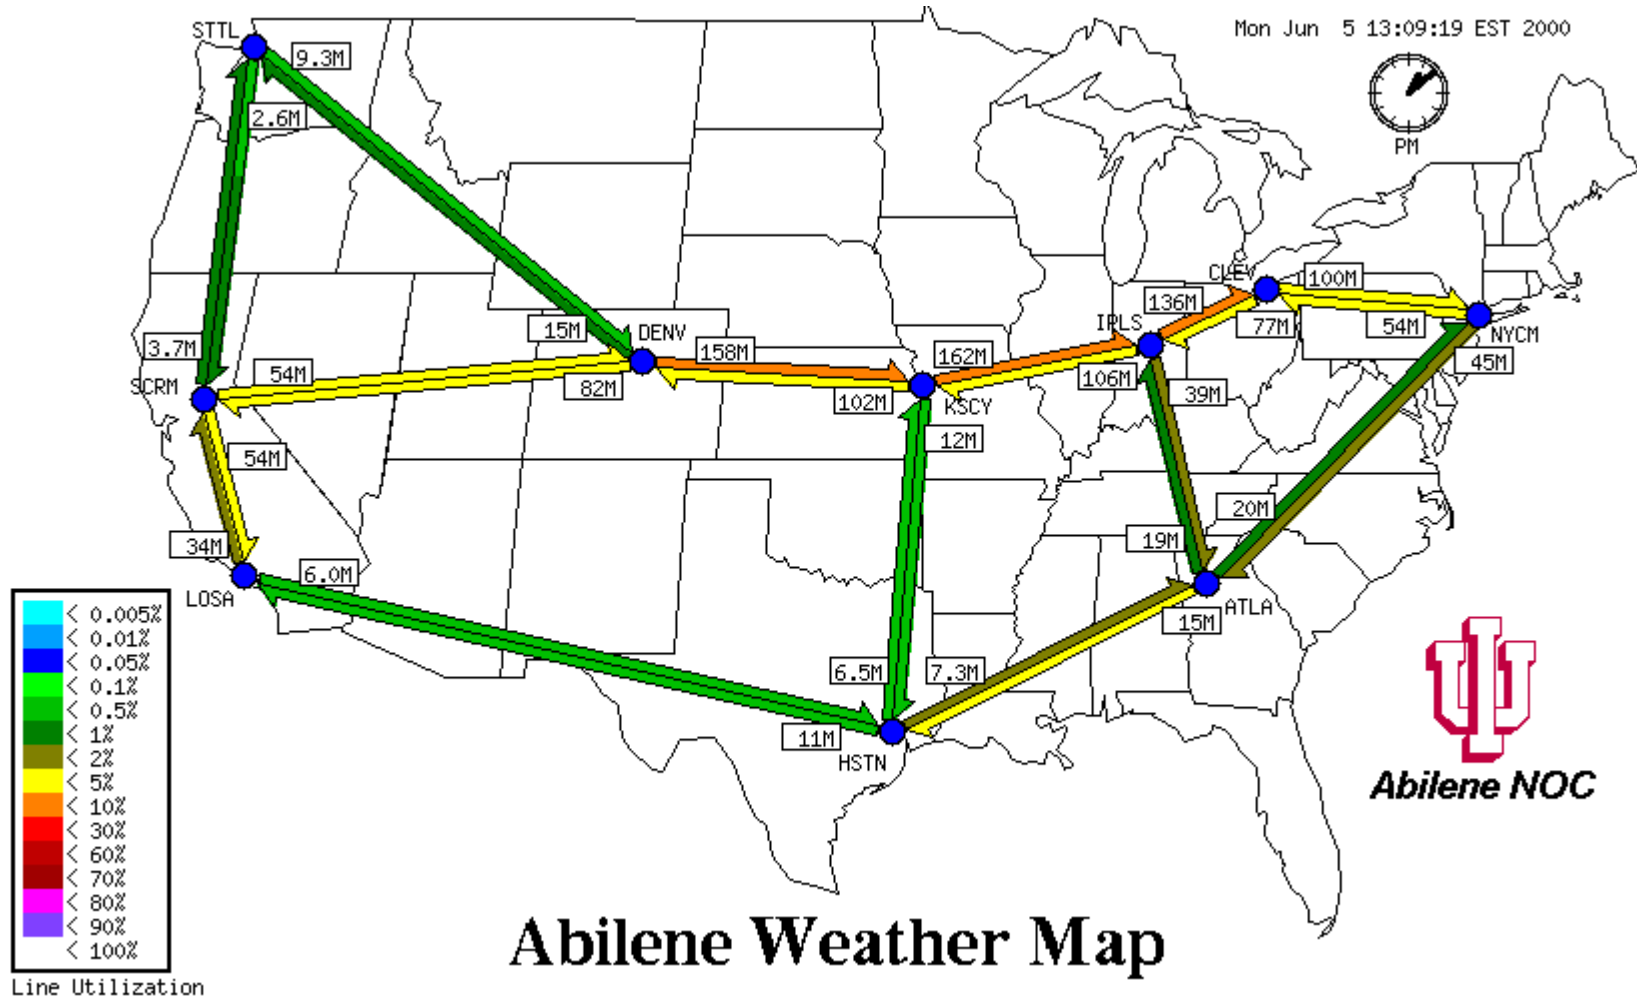
Typical ESNet/Abilene Performance

ORNL → Washington

- Minimum Delay
- 50th Percentile Delay
- 90th Percentile Delay



Abilene Weather Map



Abilene Weather Map



True...

- Edge-to-edge backbone measurements show that Internet2/NGI utilization is very low (~%5)
- Delay/jitter sensitive apps will work well today when end points are placed at or near the edge


But...

- Significant anecdotal evidence suggests that the e2e story is much less pretty ⇒ **it's the campus stupid!**
- Other congestion points exist ⇒ **e.g. submarine links**
- It's difficult to predict the future ⇒ **let's not be victimized by the first success catastrophe!**

Auckland → Washington

Delay statistics over 1-minute intervals starting 00:00 UTC, Tuesday, June 13, 2000


- Minimum Delay
- 50th Percentile Delay
- 90th Percentile Delay



Washington → Auckland

Delay statistics over 1-minute intervals starting 00:00 UTC, Tuesday, June 13, 2000

- Minimum Delay
- 50th Percentile Delay
- 90th Percentile Delay



Goal:


- Make APS a reference implementation of the QBone architecture

Why?

- **Catalyst:** Enable wide-area QoS experimentation and propagate Abilene's experiences to larger community
- **Safety Belt:** Congestion is only one success catastrophe away!

The Team:

- UCAID
- Indiana University
- I-TECs (NCNI, OARnet)
- Cisco
- Qwest
- Nortel
- NLANR/NCNE



Abilene Premium Service (APS) Test Program



Participation Criteria

- Abilene connector status necessary but not sufficient
- QBone participation
- Capability of access router / Abilene edge card

Current APS Participants

- MAGPI (U. Penn)
- iCAIR-NWU
- PSC (Penn State)
- OARNet (Ohio State)
- ANL
- DOE Science GRID (peering transit network)

Abilene QoS Service Phasing

<i>Sweetwater</i>	<i>Midland</i>	<i>Odessa</i>	<i>Pecos</i>
			
Measurement Infrastructure (Surveyor + SNMP + HTTP Dissemination)			
Edge Policing (CAR)			
Manual Setup (Whiteboard + CLI)			
	EF Core Forwarding (MDRR)		
		EF Edge Forwarding (MDRR)	
		Automated Setup (BB or ?)	
			Shaping (GTS)

First, crawl...then, walk...then, run...

- Ingress policing of a “firehose” EF aggregate
 - Can deepen token bucket policer to accommodate limitations of connector egress shaping
- Measurement feedback
- Limited technical support
- No priority queuing yet, but still get benefit of Abilene’s overprovisioning!

Goals


- Begin to establish operational practices for responding to admissions requests
- Provide participants with the experience of sitting behind a policer

QoS is a rich and varied space

- QoS addresses a **particular** class of applications (VoIP, video, other interactive/collaborative apps)
- There are many other things we **could** be doing

Anyone for...

- Managing BE performance by mapping low-priority traffic to an LBE/“scavenger” class?
- Two BE classes to trade off loss and delay?
- IP virtual trunking?
- Statistically assured services?
- Pricing feedback experiments?



For More Information...

Internet2

- <http://www.internet2.edu>

QBone Forum (Slash-based discussion forum)

- <http://qbone.internet2.edu>

Internet2 QoS Working Group

- <http://www.internet2.edu/qos/wg>

Abilene Premium Service

- <http://www.internet2.edu/abilene/qos>



INTERNETTM