VoIP & QoS: *You Can't Always Get What You Want*

VoIP Workshop
TAMU, College Station, Texas

Ben Teitelbaum <ben@internet2.edu>
April 4th, 2002

Outline

What does VoIP need?

Q: Will we ever get IP QoS?!

A: Yes, but probably not what you thought

What might VoIP get?

What Does VoIP Need?

...We interrupt this program to bring you the following flame...

- "This is the Internet, amigo. You should be grateful for what you can get and ask not what the network can do for you, but what you can do on the end−systems to make your application work."

...And now back to our regularly scheduled program...

Characterizing Voice Quality

Very subjective

Standard metric: mean opinion score (MOS)

Objective metrics do exist (PSQM, PESQ)

Quality dimensions

- **Clarity** − fidelity, clearness, and intelligibility of signal
- **Delay** − effect on interactivity (talker overlap minimized)
- **Echo** − distracting and confusing (caused by crosstalk between send and receive signals)

What Does VoIP Need?

Let's look at voice quality as a function of:

- Latency
- Jitter
- Loss

And, say a few words about:

- Bandwidth
- Reliability

Latency$_1$

Latency components

- Encoding
- Packetization
- Network delay
  - Queuing (QoS can help)
  - Propagation (QoS may help; TE will hurt)
  - Serialization and switching (QoS can't help)
- Receiver buffering
- Decoding

Latency$_2$

One−way delay budget

- Estimates vary from 100ms−300ms
- ITU−TG.114 recommends 150ms

Some rules of thumb

One-way Delay	Effect on Perceived Quality
<100-150ms	Delay not detectable
150-200ms	Acceptible quality; slight delay or hestitation noticeable
Over 200-300ms	Unacceptible delay; normal conversation impossible

Jitter

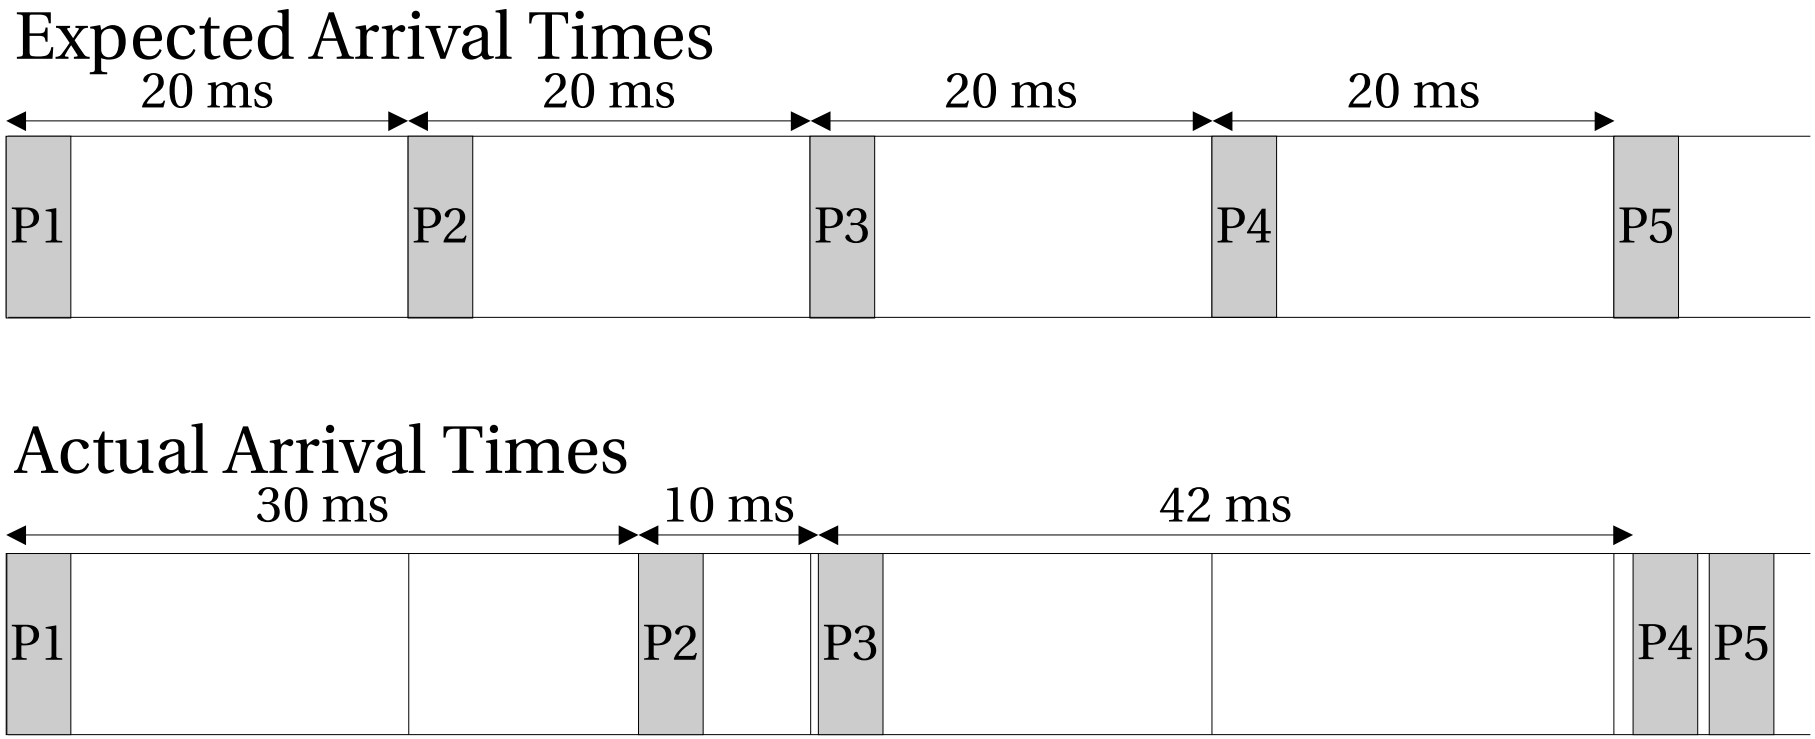

Smoothed by playback buffers (added delay)

*Receivers **adapt** the depth of these buffers*

⇒ sudden changes in jitter may cause loss

Loss

Relationship between packet loss and quality has many dependencies

- Codec used
- Packet size
- Existence of error protection / correction
- Loss pattern

Estimates of VoIP loss tolerance range from 1% to 5%

Bandwidth & Reliability

Bandwidth

- Generally modest (64 kbps or less)
- Depends on codec and use of silence suppression
- Examples:

Codec	Rate (kbps)
G.711	64
G.722	48-64
G.726	32
G.729 (A/B)	8
GSM FR	13

Reliability

- Does VoIP really need PSTN−level reliability?
- DOS attacks (QoS may help)
- Link failures (path redundancy, plus fast IGP convergence, plus fast EGP convergence)

Q: Will We Ever Get IP QoS?!

Picture is rosier in intradomain case...

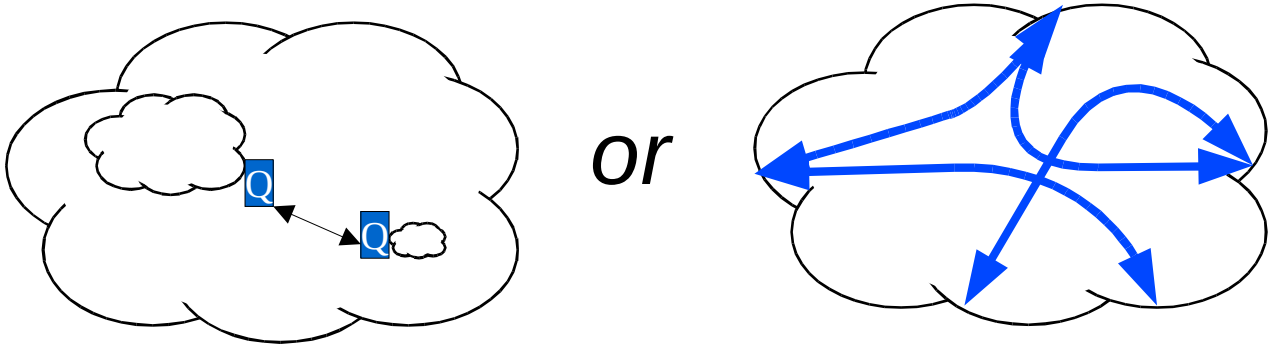

or if you assume that congestion is only at access circuits that you control...

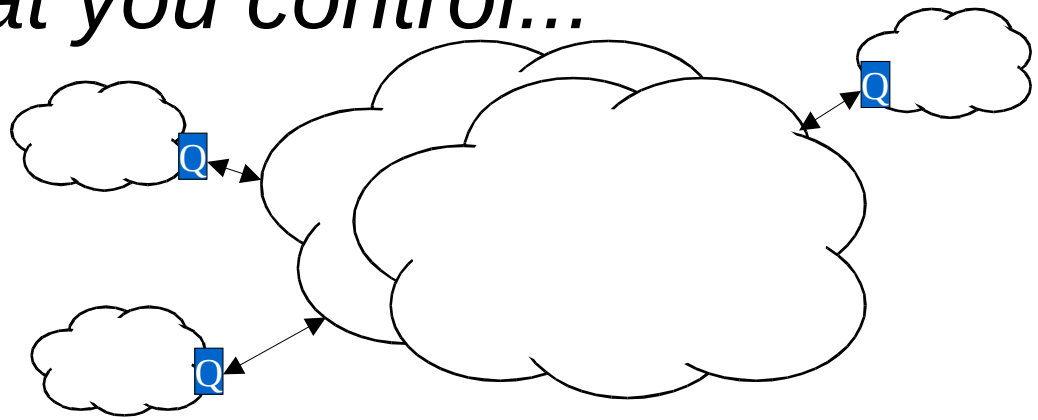

What About Interdomain Reservations?

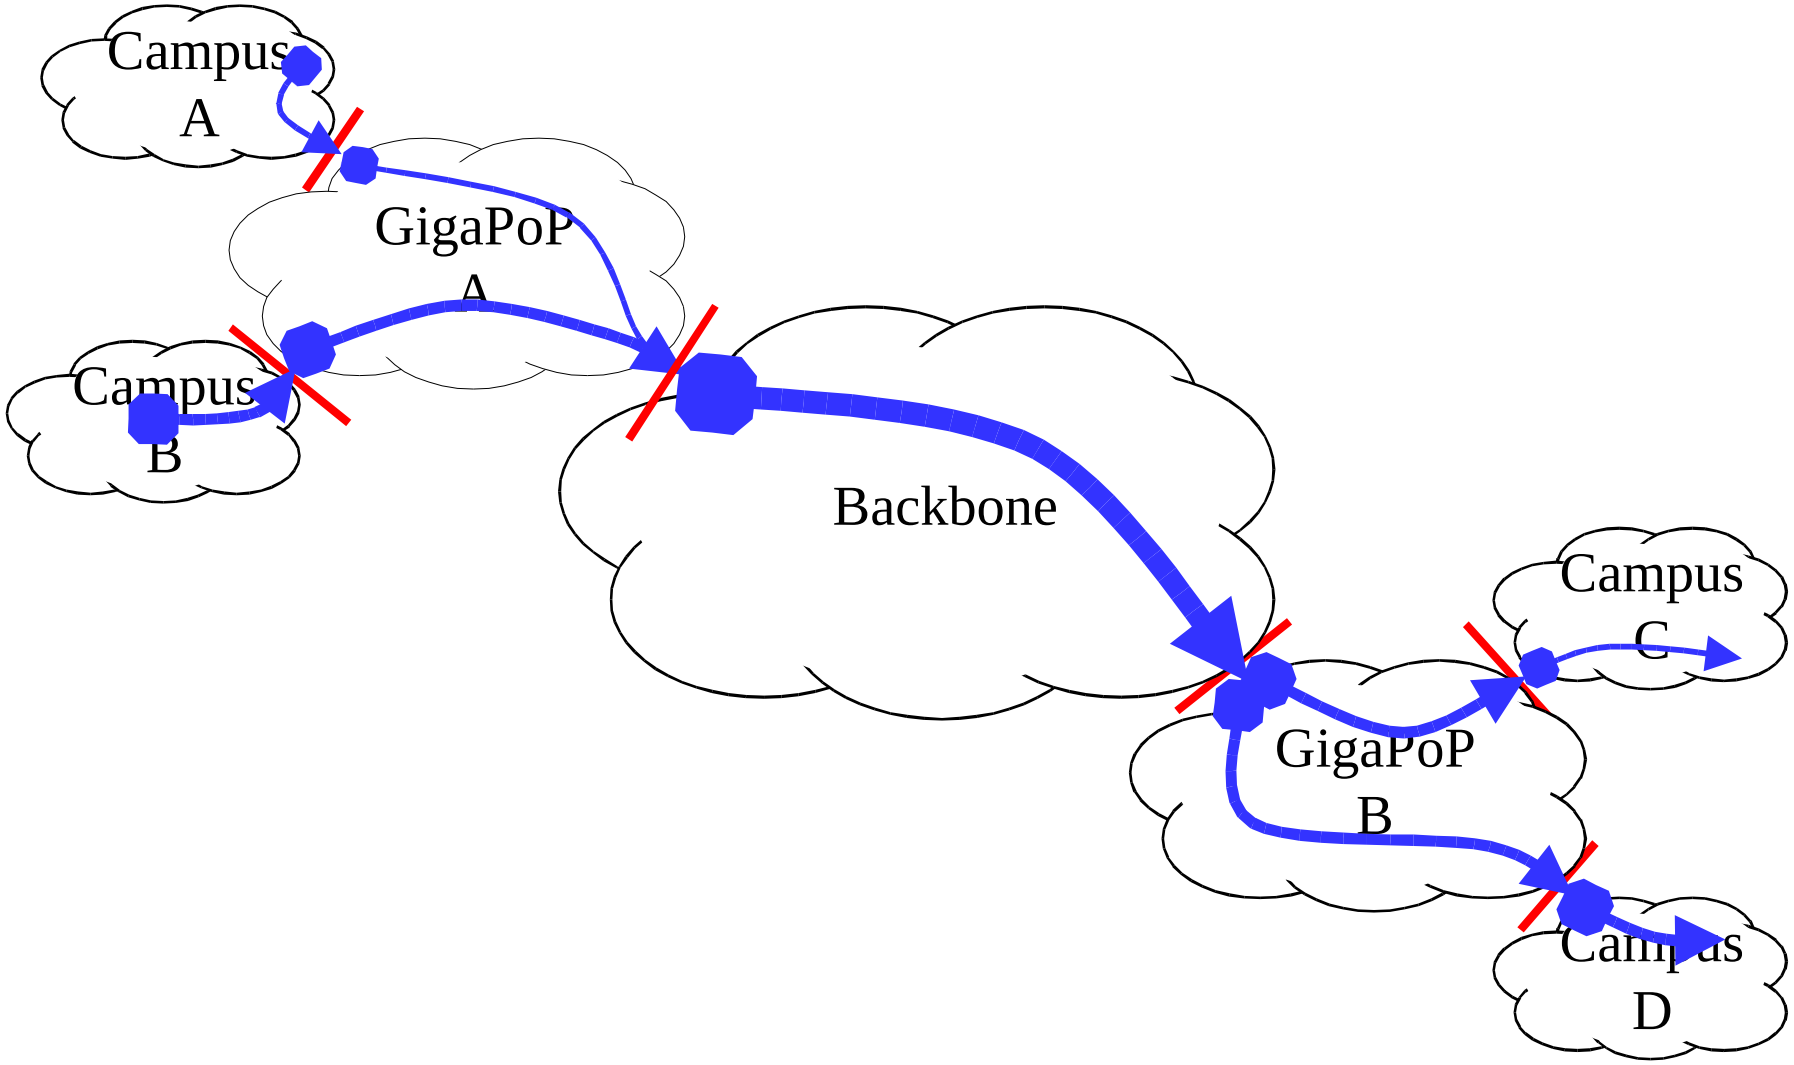

QBone Architecture (circa 1999)

A Service: QBone Premium Service

- IP circuit−emulation (a.k.a. "virtual leased line")
- Built on Expedited Forwarding (EF) (RFC 2598)

Reservation Setup Protocol

- **Initially**: long−lived, manual setup
- **Later**: SIBBS protocol between QBone domains; RSVP end−to−end between hosts

QBone Measurement Architecture

- Uniform collection of QoS metrics
- Uniform dissemination interface

QBone Architecture (30 kilofoot view)

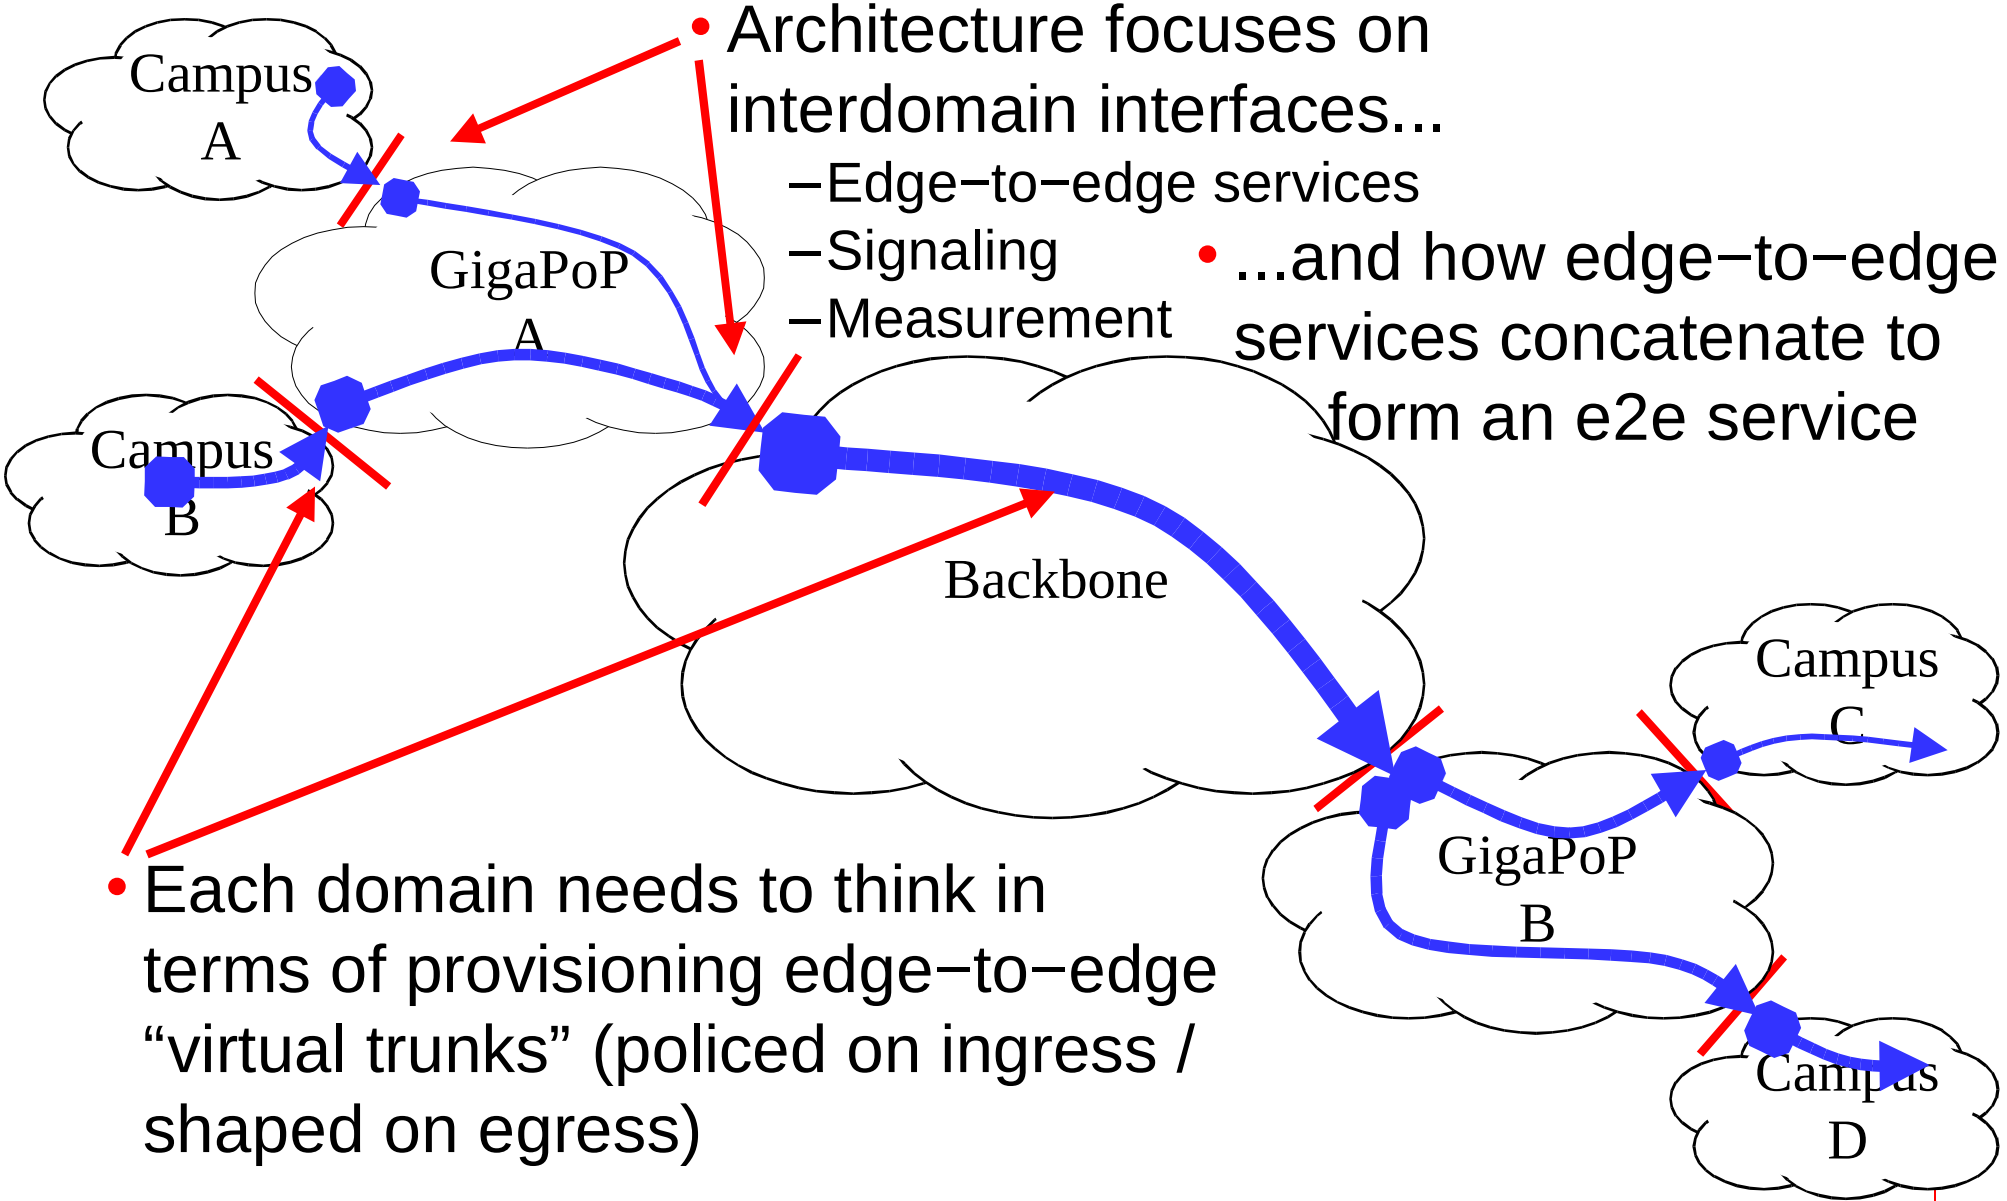

- Architecture focuses on interdomain interfaces...
  - –Edge–to–edge services
  - –Signaling
  - –Measurement
- ...and how edge–to–edge services concatenate to form an e2e service
- Each domain needs to think in terms of provisioning edge–to–edge "virtual trunks" (policed on ingress / shaped on egress)

QBone Architecture (30 kilofoot view)

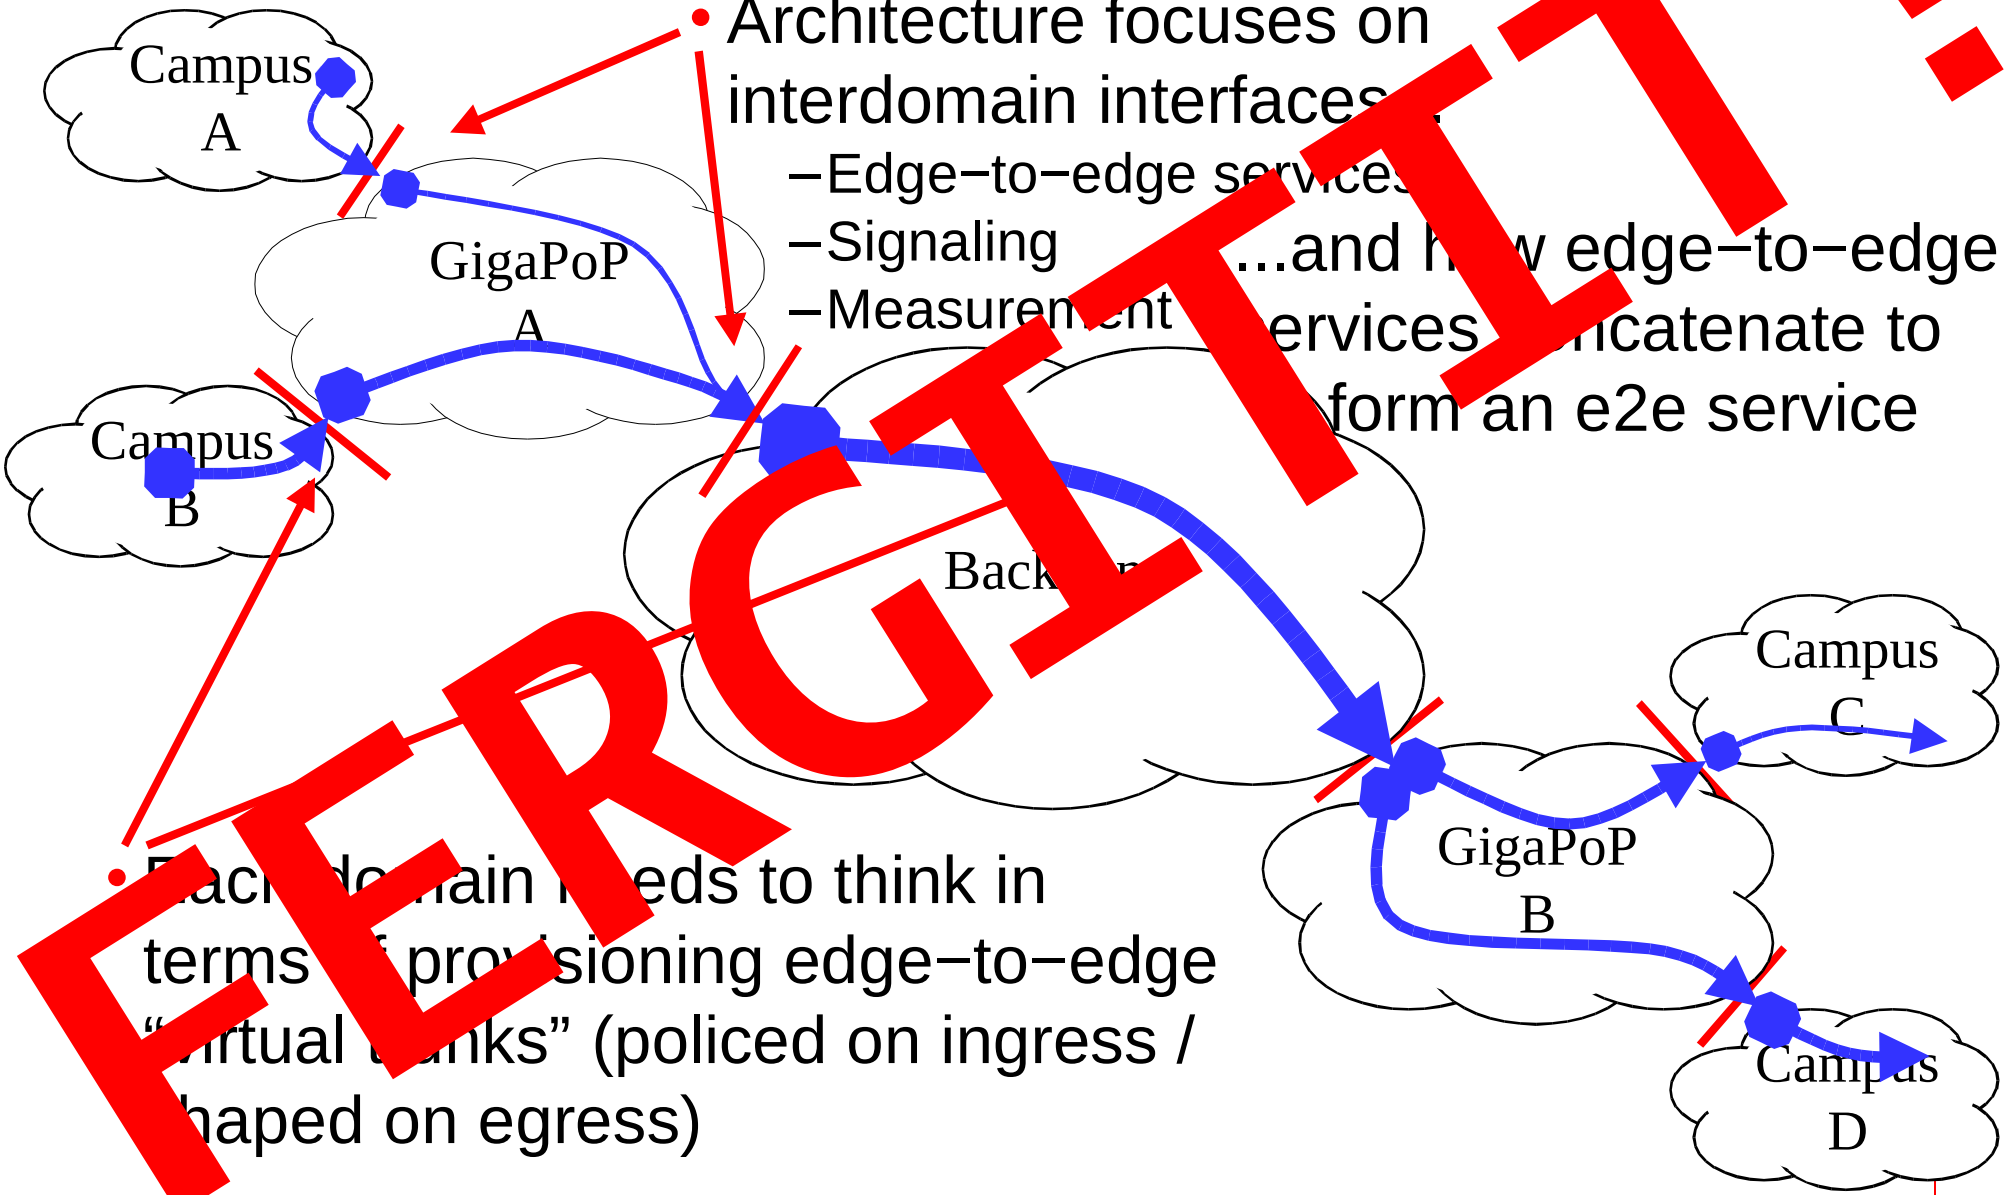

- Architecture focuses on interdomain interfaces
  - Edge-to-edge services
  - Signaling
  - Measurement
- ...and how edge-to-edge services concatenate to form an e2e service
- Each domain needs to think in terms of provisioning edge-to-edge "virtual links" (policed on ingress / shaped on egress)

FERGETIT!

Obstacles to Premium Deployment

Low demand

Current router support for DiffServ is spotty

Fundamental practical deployment difficulties

Fundamental theoretical problems

Obstacles to Premium Deployment

Low demand

- Classical "*chicken−and−egg" problem*
- *Artificially constrained BE load (more...)*

Current router support for DiffServ is spotty

Fundamental practical deployment difficulties

Fundamental theoretical problems

Utilization Paradox

Order ~10^4 hosts with nothing slower than switched 100Mbps Ethernet between them

Theoretically, ~25 of these could congest the 2.4 Gbps backbone

***Yet...** the backbone is lightly loaded!*

***Paradox**: Abilene is both under−provisioned and under−utilized*

Why is this?!

"Typical" E2E Internet2 Performance

- 2.6M bulk TCPs
- Week of 20020325
- Observed throughputs:
  - 50% 1.870M
  - 90% 4.898M
  - 99.9%106.3M

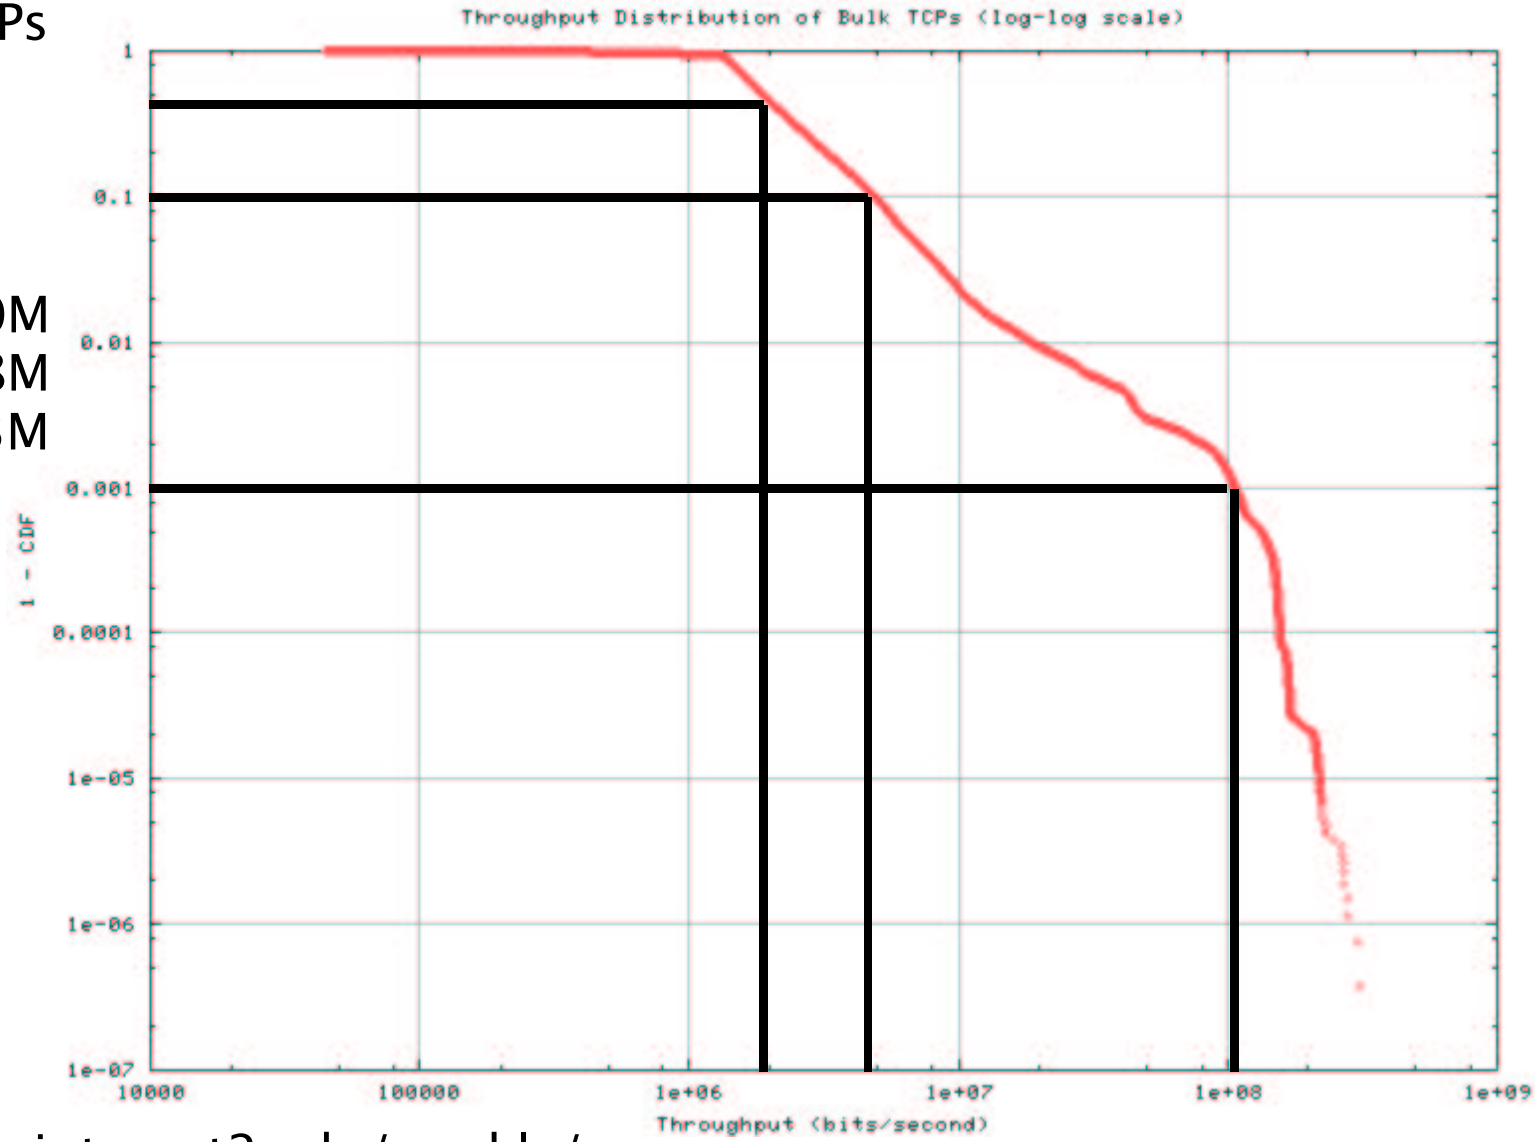

- http://netflow.internet2.edu/weekly/

Performance Faults Obviate QoS

Evidence suggests that most problems are in hosts and LANs

Common performance faults

- Broken TCP stacks (*e.g.* inadequate socket buffering, no window scaling)
- Ethernet duplex mismatch
- Crummy cabling (*e.g.* CAT3, shared, or damaged)

Internet2 End−to−End Performance Initiative

- Major initiative to work on this problem
- http://www.internet2.edu/e2epi/

Obstacles to Premium Deployment

Low demand

Current router support for DiffServ is spotty

Fundamental practical deployment difficulties

- Requires all−or−nothing network upgrades (*e.g.* all access interfaces must police)
- Service verification (by users or providers) difficult
- Dramatic changes to network operations, peering arrangements, and business models

Fundamental theoretical problems

“Non−Elevated Services

"Worse"

- QBone Scavenger Service (QBSS)
- Bulk Handling PDB (B. Carpenter, K. Nichols)

"Different−but−equal"

- Alternative Best Effort (ABE)
- Best−effort Differentiated Services (BEDS)

Why do we like these wacky services?!

- Require no policing, admissions, settlement, etc.
- Deploy incrementally at the granularity of single interfaces
- Consistent with end−to−end principle

QBone Scavenger Service

Basic idea

- **Voluntary** marking hints to network that degraded service is OK (like Un*x **`nice`** for the network)
- Scavenger traffic **may** be degraded at congestion points
- **Think**: thin, bottom−feeding best−effort network that can expand to full capacity in absence of congestion
- Formal service definition: **`http://qbone.internet2.edu/qbss/qbss-definition.txt`**

Goals

- A tool to preserve/extend uncongested BE experience for interactive applications

Motivations 1/2

All traffic is not equal w.r.t. loss and delay

- Mix of tolerant/intolerant traffic
- Since you may be competing with yourself for downstream resources, it's in your interest to identify tolerant traffic

Most routers support multiple queues

- Let's get some value and experience out of them!

Internet2 utilization very low

- **Pro**: interactive apps work fine; **Con**: what a waste!
- What new applications could be built if we weren't shy about filling the pipes?

Motivations 2/2

Fine−grained Netiquette

- Self−policing users exist
  - HEP community runs bulk−transfers “at night”
  - Network backups
  - CDN pre−fetching
- QBSS allows these apps to run continuously

Pricing

- Additional control over upstream commodity usage
- Potential point of negotiation for metered connectors

Policy

- Users/institutions could mark non−mission traffic

Current State of QBSS

Testing underway to support bulk transfer needs of HEP and astrophysics users

- SLAC, TransPAC (GRAPE), CERN, UKERNA

Gear tested and configs available for:

- Cisco 7200, 7500, GSR
- Juniper

Some operational traction

>1% QBSS on Abilene

QBSS Usage at Abilene CLEV

Biggest QBSS emitting ASes:

- RIT
- MORENET
- UMD
- MSUNET
- VA−TECH

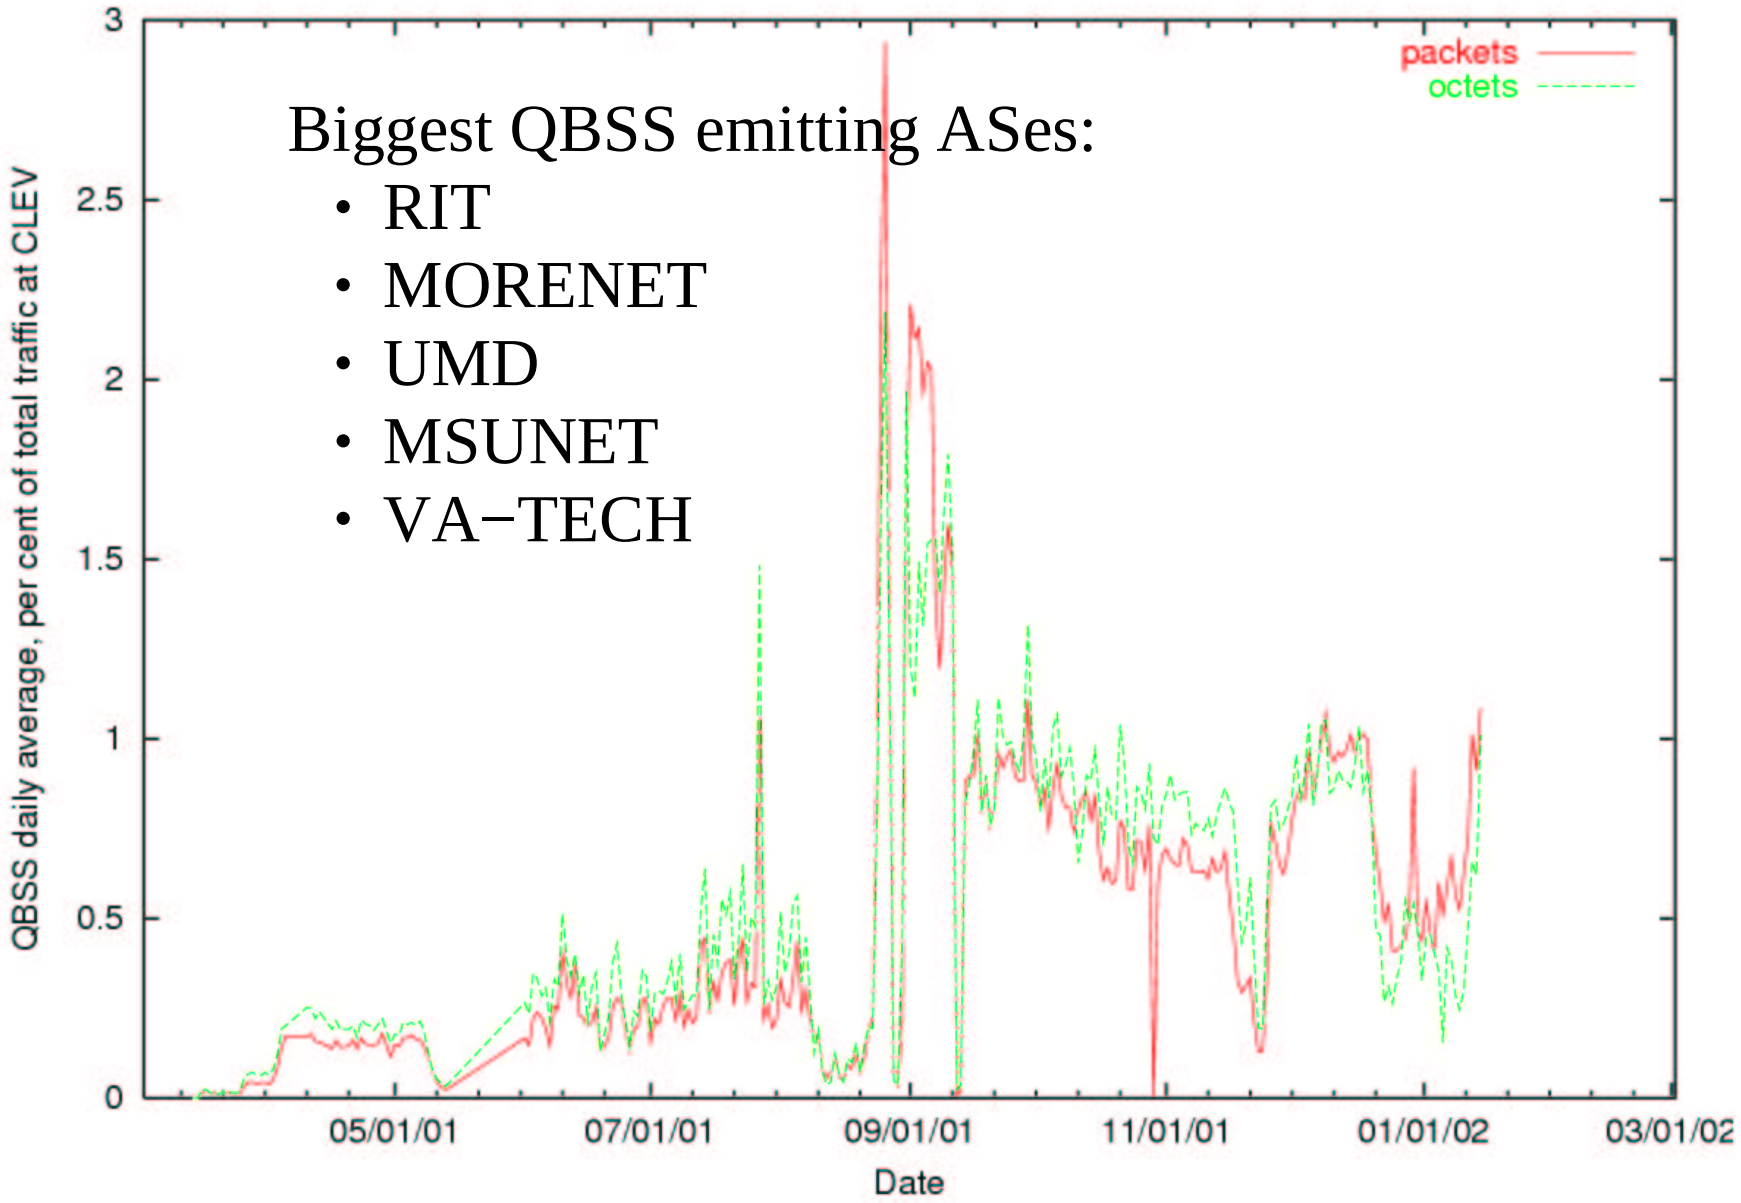

Alternative Best Effort (ABS)

Monolithic best−effort service class split into:

- **Blue** −lower loss / higher delay
- **Green** −higher loss / lower delay

Fairness relationship between classes

Each app knows its utility function and trades off loss for delay accordingly

See: http://www.abeservice.com/

Could We Do Something Like ABE Today?

Split BE into two classes, each with 50% link share (e.g. via WFQ)

Configure low−latency queue

- Steep RED curve to give high loss rate with short queue
- Queue−limit to tail drop before queue exceeds ~5ms

Map well−known DSCP to LLBE queue

TCPs could wind up in the LLBE queue, but they would be severely punished

Still vulnerable to DOS attacks

Application QoS Needs

Too much mythology and confusion about what apps really need

Goals:

- Build bridges between networkers and developers
- Promote best practices for developing and deploying adaptive multimedia applications

Activities in this area

- Detailed survey of application QoS needs and relationship between application utility and network performance
- Measurement and analysis to understand application performance and use of new services

For more information...

Internet2 QoS WG Home:

- http://www.internet2.edu/qos/wg/
- Links to all WG design teams may be found here

QBone Scavenger Service

- http://qbone.internet2.edu/qbss/

Application QoS Needs

- http://www.internet2.edu/qos/wg/apps/

QBone Home:

- http://qbone.internet2.edu/


INTERNET™