Internet2 QoS & Video: *Sharing Responsibility to Overcome Congestion Related Performance Problems*

SURA/ViDe 4th Annual Digital Video Workshop
U. Alabama, Birmingham

Ben Teitelbaum <ben@internet2.edu>
Amela Sadagic <amela@advanced.org>
April 25th, 2002

So, You Have a Problem...

Many video problems not related to network

- Poor mic placement, poor lighting, human factors, etc.
- But you all know more about these problems than I

OK, so you have a network problem...

- But, most performance problems are non−congestive
- Usually due to faults in or near hosts
  - Broken TCP stacks
  - Ethernet duplex mismatch
  - Crummy cabling

Typical E2E Internet2 Performance

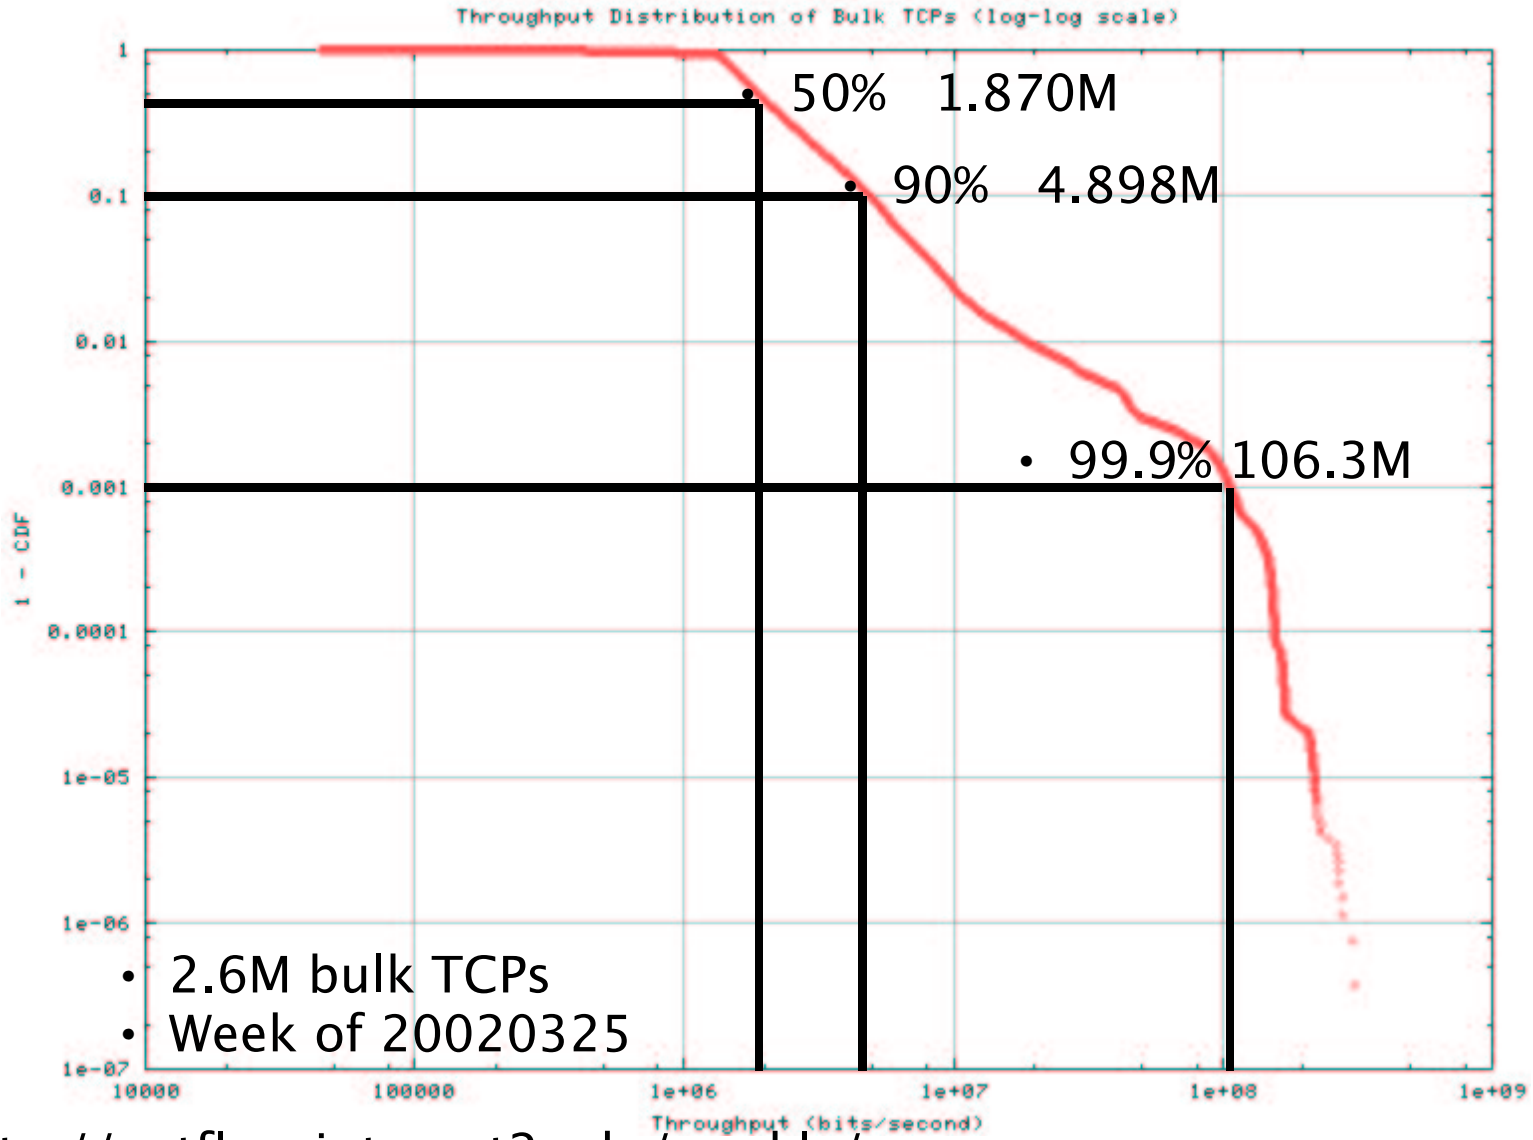

- http://netflow.internet2.edu/weekly/

OK, So You Actually Have Congestion...

This is largely an economic problem

- Economic solutions exist (pricing feedback mechanisms)
- Unfortunately Internet2 is a poor place to experiment with pricing

What are you going to do?

Options Do Exist When You Control the Congested Resources

Picture is rosier in intradomain case...

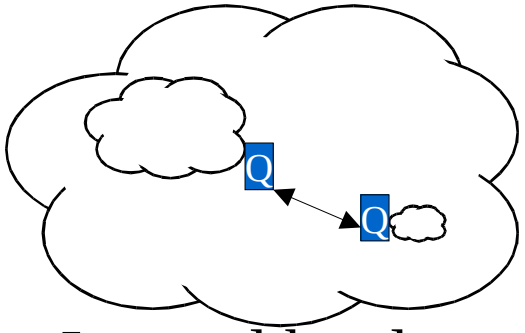

Internal bottleneck

or

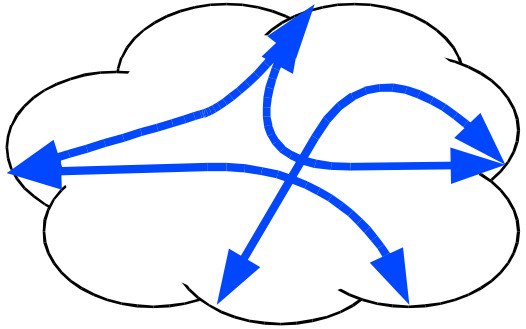

QoS−enabled VPN

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

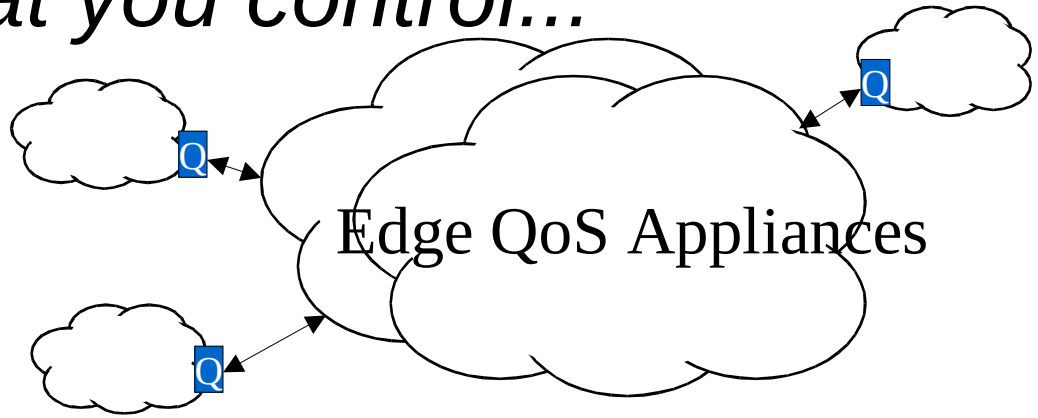

But, the Fantasy is Interdomain QoS!

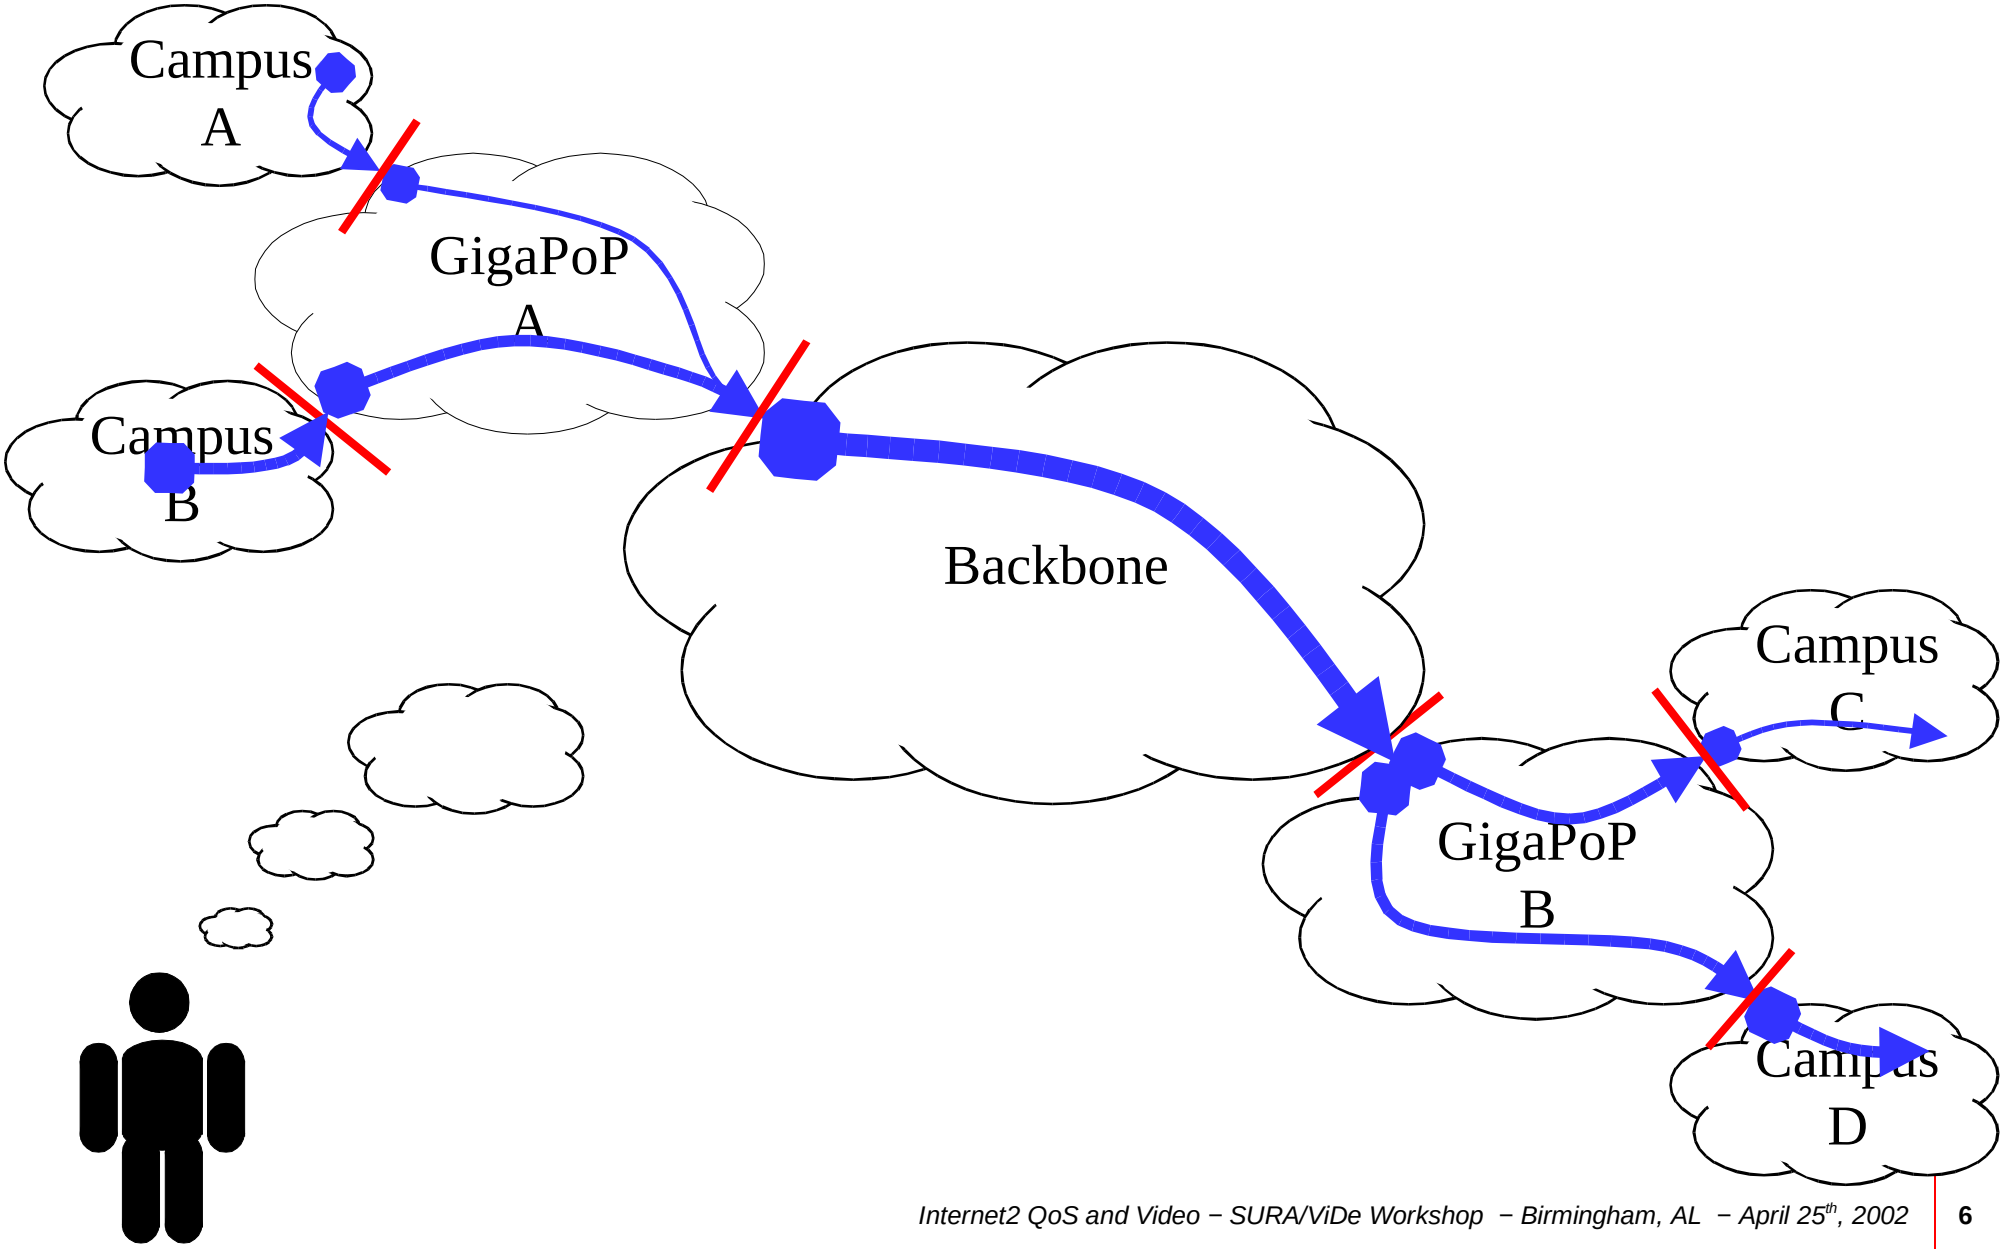

We Tried This (QBone Premium)

- Architecture focuses on interdomain interfaces...
  - –Edge–to–edge services
  - –Signaling
  - –Measurement
- ...and how edge–to–edge services concatenate to form an e2e service

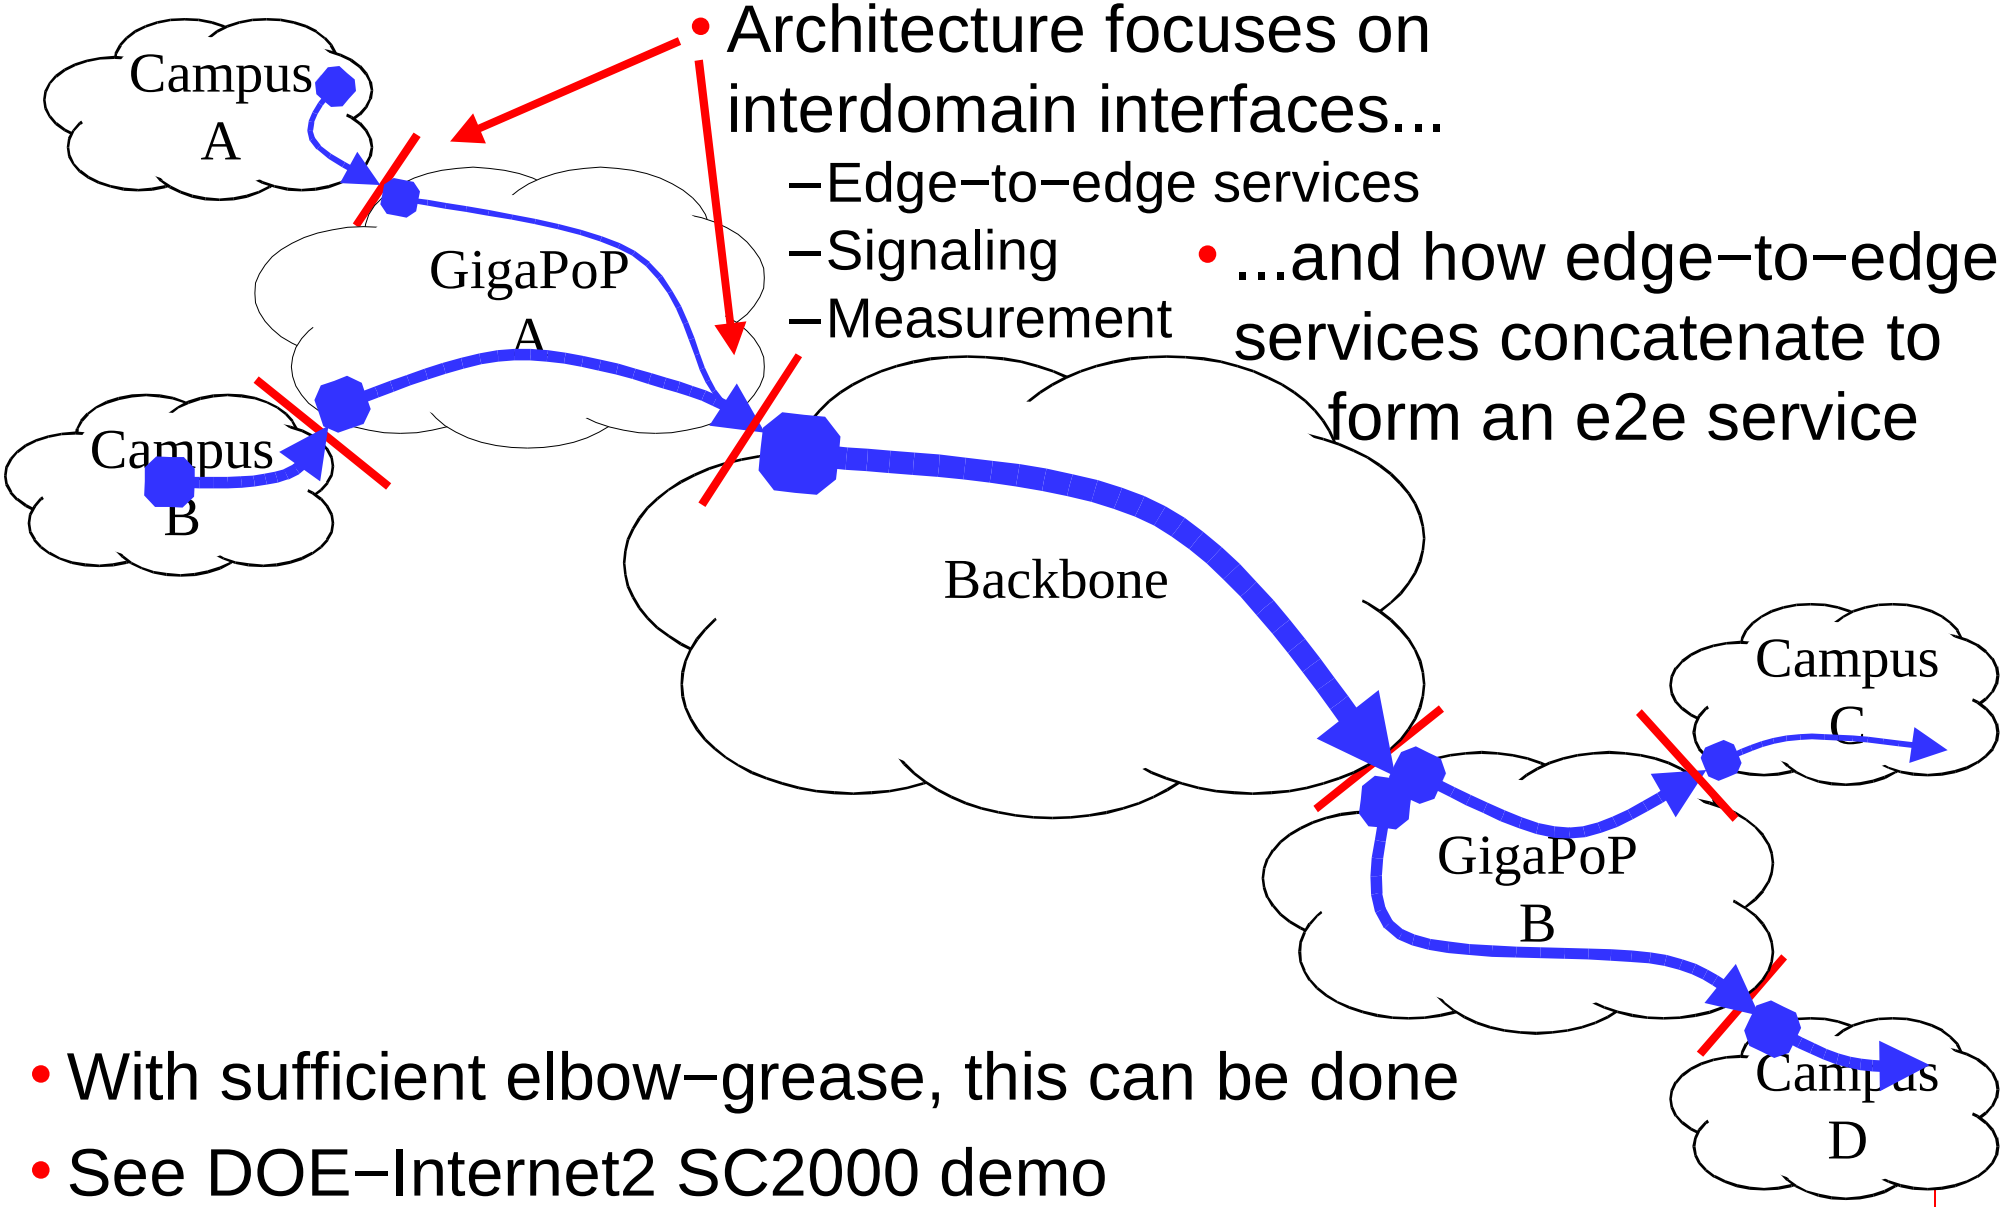

- With sufficient elbow–grease, this can be done
- See DOE–Internet2 SC2000 demo

We Tried This (QBone Premium)

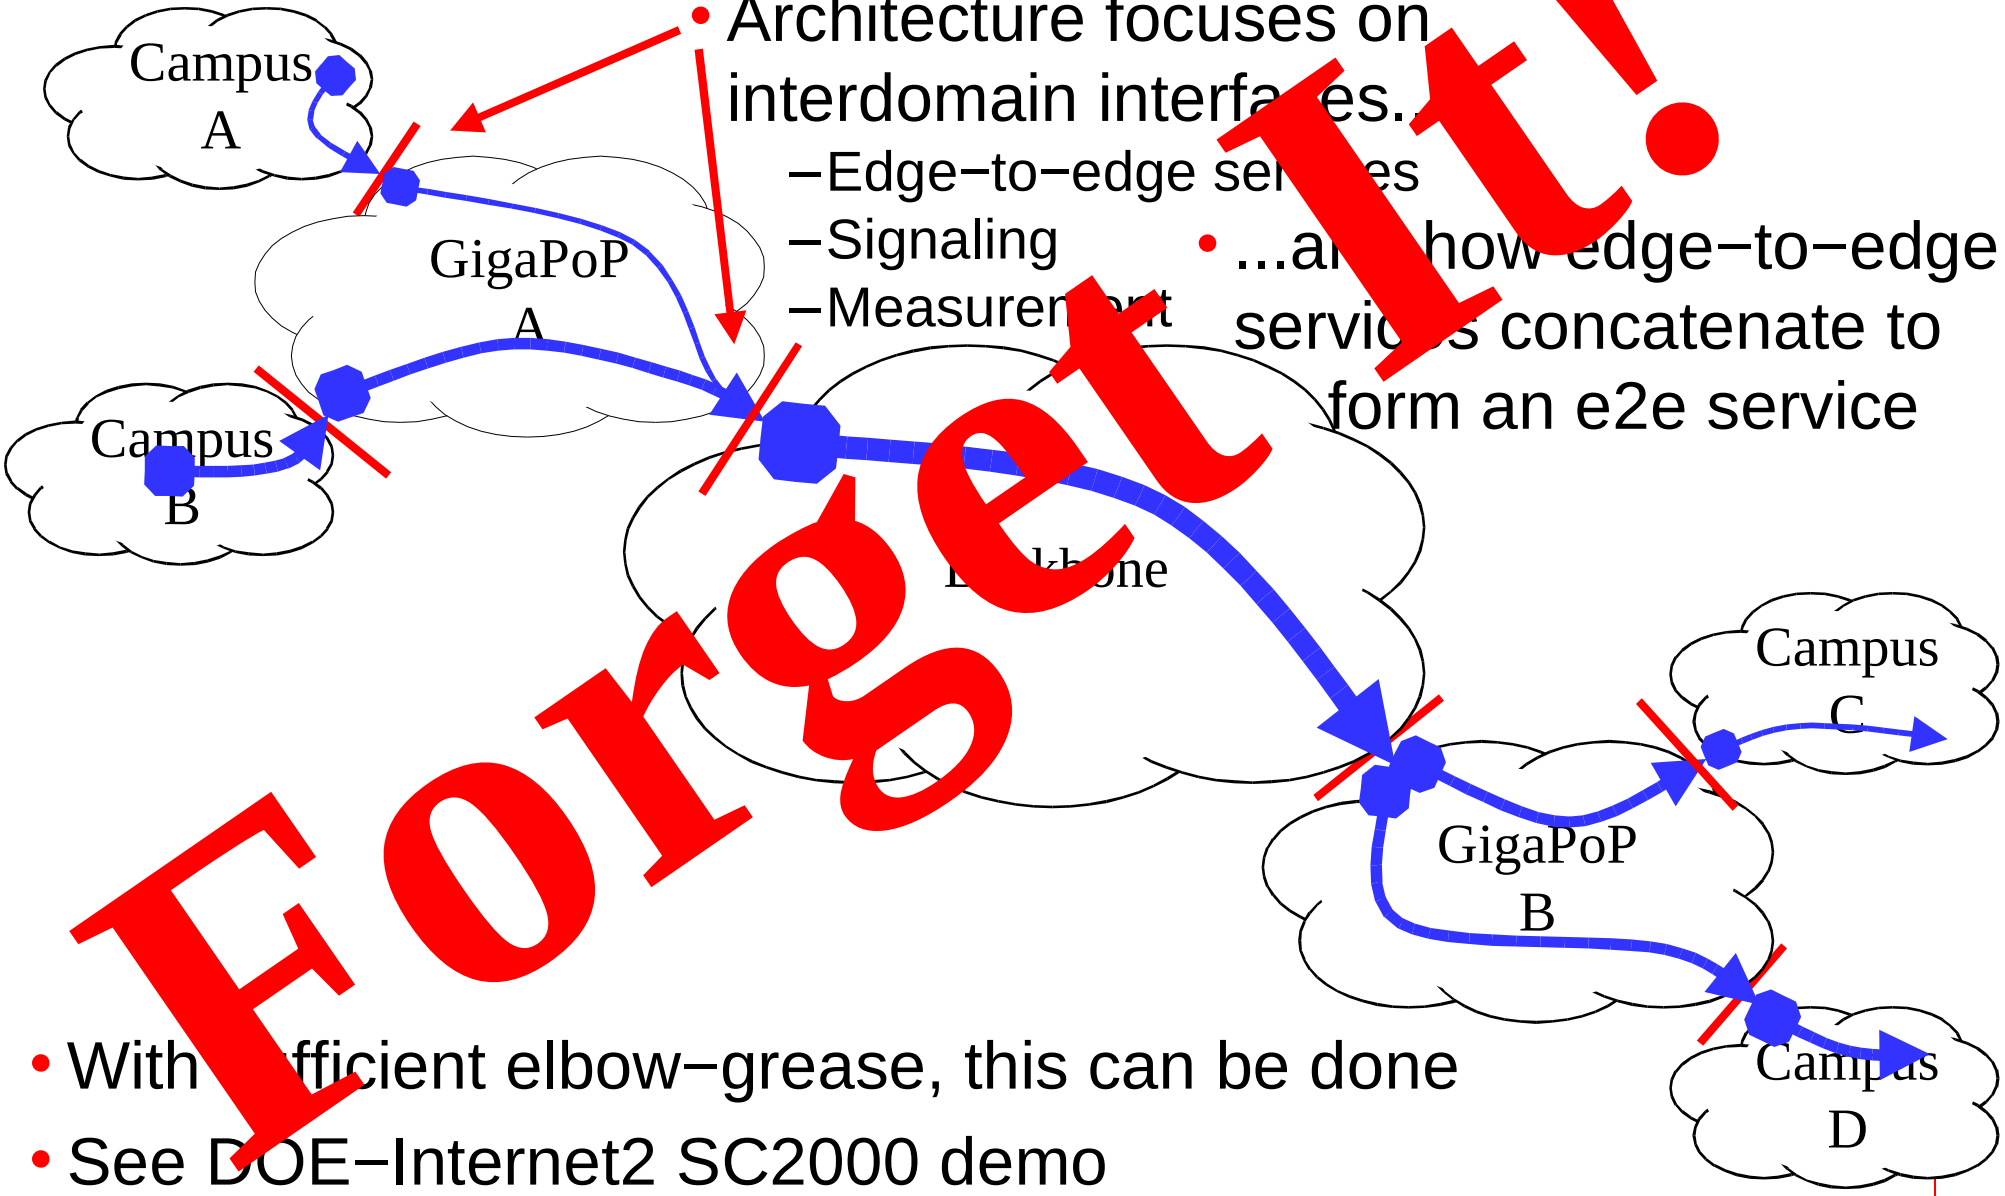

- Architecture focuses on interdomain interfaces...
  - –Edge–to–edge services
  - –Signaling
  - –Measurement
- ...and how edge–to–edge services concatenate to form an e2e service

Forget It!

- With sufficient elbow–grease, this can be done
- See DOE–Internet2 SC2000 demo

QBone Premium Problems[1]

Architecture not complete

- Exact shaper/policer provisioning never understood
- Signaling never gelled

DiffServ functionality still missing in modern routers or not available at line rate

- Route−based edge classification, anyone?
- How about multiple shaped aggregates within a PQ?

Per−net deployment granularity

- Must police EF traffic at every ingress interface

Service verification

- To "jiggle door" service provider or customer must launch EF DoS attack

QBone Premium Problems$_2$

Scary new business model

- Accounting, billing, etc.
- Complex new peering agreements with QoS SLAs

Scary new operational responsibilities

- Admissions control
- Increased vulnerability to DoS attacks

Finally, where's the demand that's going to make this all worth doing?!

Where Does This Leave Us?

We have "suspended" our Premium efforts

Working to fix common e2e performance faults and raise user expectations

Working to raise awareness of end−to−end principle and best practices of application adaptation

As for QoS, we are going with the theory that less is more[1]

1. More deployment anyway!

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

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

Could we do an ABE−like low−delay class today (e.g. with WFQ and RED)?

http://www.abeservice.com/

What Does Video Need?

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

- "This is the Internet, amigo. It's fast, cheap, and global, but there are no guarantees. 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 Audio 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 Audio 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

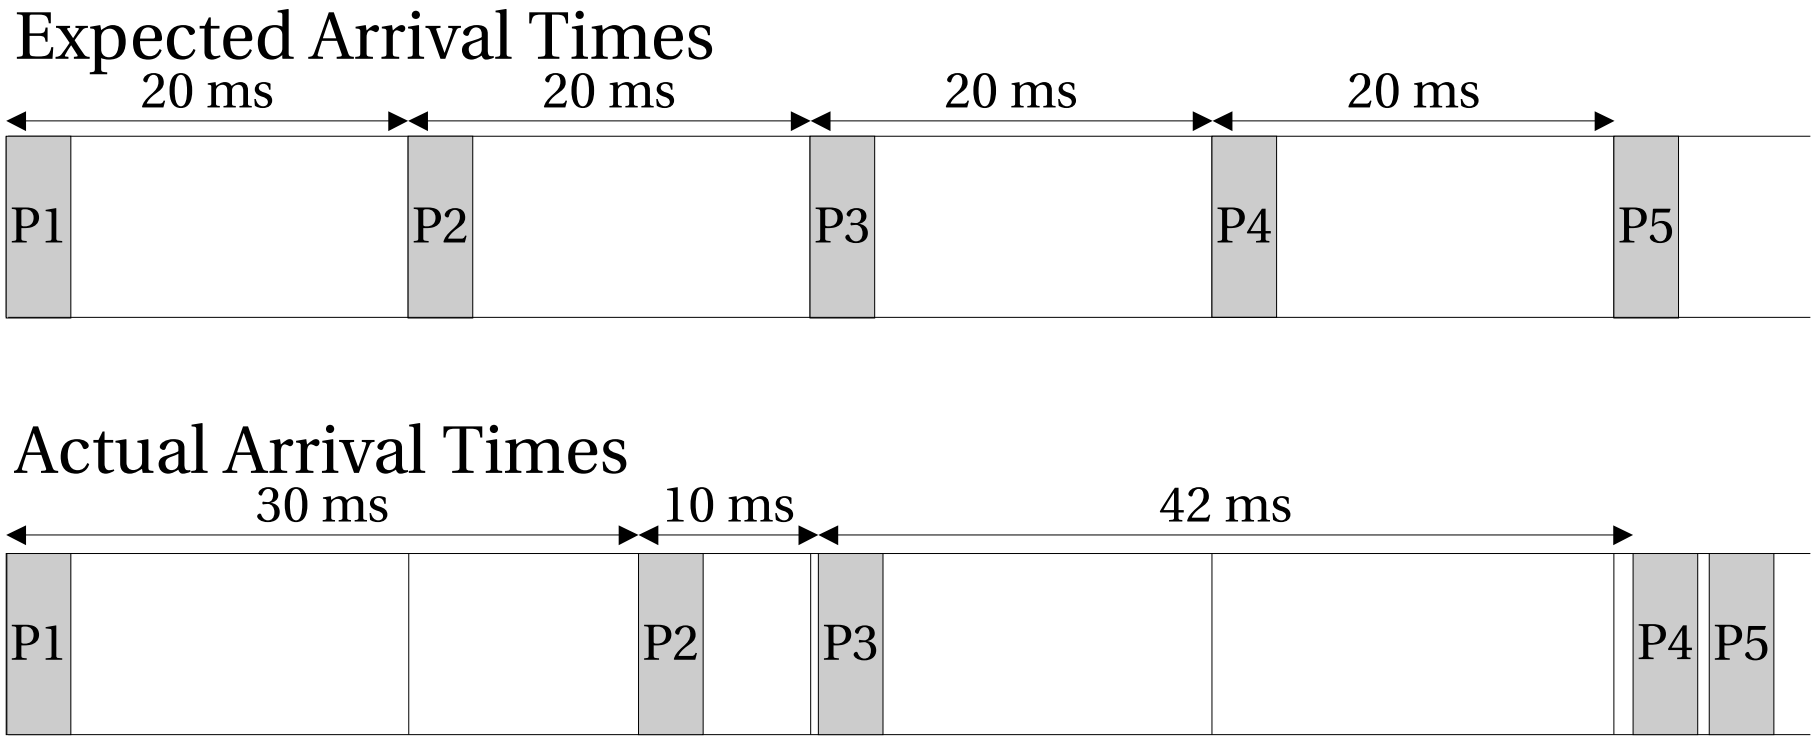

Smoothed by playback buffers (added delay)

*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)

What If We Add Video?

Object and temporal coherence → video less sensitive to data loss

Video less sensitive to latency

But video requires more bandwidth

Combining Video with Audio in One System

Bandwidth requirements increase greatly for high quality video

Mix of media (video, audio, data) and the context they are used in, changes the way we perceive them (different than each medium separately)

Synchronization with audio: a big issue

Do You Need QoS?

"Dumb network" and "network equal for all"
→ cornucopia of ideas & applications at the ends

Sure... as well as inevitable side effect of dumbing down the same applications (HCI issues).

...and so tele−medicine and tele−immersion are hard to imagine there...

How Well do I Have to Do?

Type	Latency	Bandwidth	Reliable	Multicast	Security	Streaming	DynQos
Control	< *30 ms*	64Kb/s	*Yes*	No	*High*	No	Low
Text	< *100 ms*	64Kb/s	*Yes*	No	Medium	No	Low
Audio	< *30 ms*	Nx128Kb/s	No	*Yes*	Medium	*Yes*	Medium
Video	< *100 ms*	Nx5Mb/s	No	*Yes*	Low	*Yes*	Medium
Tracking	< *10 ms*	Nx128Kb/s	No	*Yes*	Low	*Yes*	Medium
Database	< *100 ms*	> *1GB/s*	*Yes*	Maybe	Medium	No	*High*
Simulation	< *30 ms*	> *1GB/s*	Mixed	Maybe	Medium	Maybe	*High*
Haptic	< *10 ms*	> 1 Mb/s	Mixed	Maybe	Low	Maybe	*High*
Rendering	< *30 ms*	>*1GB/s*	No	Maybe	Low	Maybe	Medium

Source: Rick Stevens, Argonne National Lab

Where Can I Introduce QoS?

Network level (lowest level)

- ABE−like service, Scavenger, ...

H.323 suite level

- Codecs: great space for quality improvements
- Gatekeeper: bandwidth management

Application level

- Scene type: "talking heads" vs. "beach cam" scenes
- Importance of content: "talking heads" vs. entertainment video or laparoscopic surgery cam

Video Quality Assessment

Subjective

- Self−reported by subjects
- Many dimensions matter: subject's background, video content, correlation with audio, display size, resolution, viewing distance
- Subject is shown sequence pairs, reference and test (in-service i.e. as seen by end−user) sequence
- Exposure to short sequence (8−10 sec) vs. exposure to long sequence (20−30 min)
- Long clips: use slider scale from "bad" to "excellent" every 1−2 sec. But is it absolute or relative grading of successive video chunks?

Video Quality Assessment

Objective

- As calculated by an algorithm (computational models)
- Need to develop good quality metrics
- Some metrics rely on model of human vision system (eyes more sensitive to luminance than color), and Some on measuring features of perceptual distortions (compression artifacts and transmission errors)

Where More Work is Needed

Development of tools and experimental measurement procedures

- To quantify how different levels of service and resource guarantees translate into application level quality improvements
- They help assessing the benefits of service differentiation
- They will be application specific
- Critical to the successful deployment and usage of service differentiation in the Internet2

Application QoS Needs

Too much mythology and confusion about what apps really need

Goals:

- Build bridges between networkers and app 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/
- qos−appl−dt@internet2.edu
- Dimitrios Miras <d.miras@cs.ucl.ac.uk>

QBone Home:

- http://qbone.internet2.edu/


INTERNET™