



Internet2 QoS: Is Less More?

Columbia University


Ben Teitelbaum <ben@internet2.edu>
October 29th, 2001

Why QoS?

Best effort internet vulnerable to a “tragedy of the commons”

Internet2 doing everything it can to promote new, radically more demanding apps

QoS needed as “safety belt” to avert a success catastrophe



A Service: QBone Premium Service

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

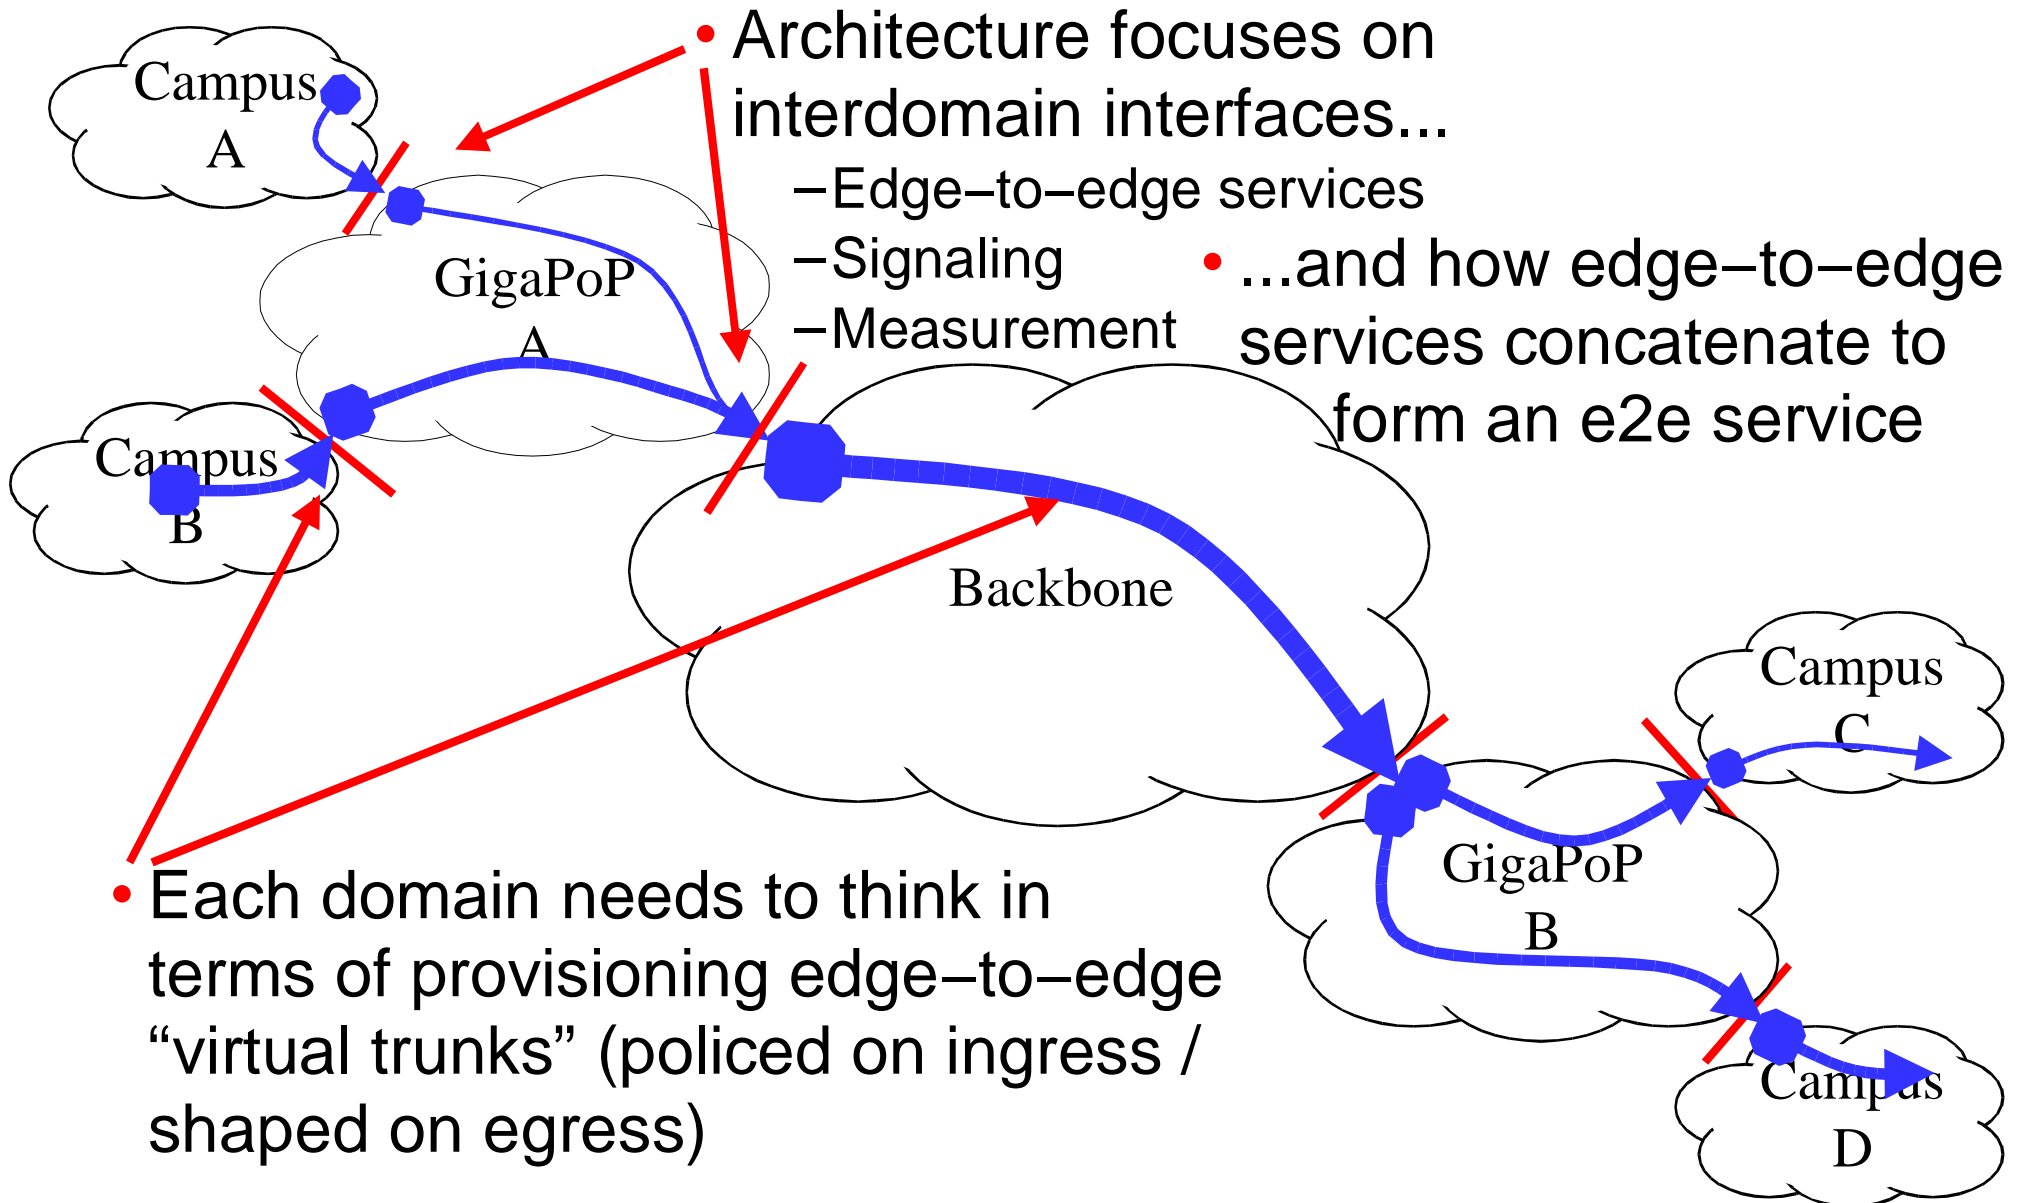
Reservation Setup Protocol

- **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

QBone Architecture (30 kilofeet view)



Low demand

Router support for DiffServ is spotty; e.g...

- No PQ
- DiffServ comes with a performance cost
- Limitations on token bucket depths
- Inflexible classification rules

Dramatic changes to network operations, peering arrangements, and business models

Requires all-or-nothing network upgrades (e.g. all access interfaces must police)

Utilization Paradox

Order $\sim 10^4$ 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?!

- 50,000 bulk TCPs observed on 6/19/01
- Sampled NetFlow at core router
- Observed throughputs:
 - Median: 880 Kbps
 - 10% \geq 3.9 Mbps
 - 1% \geq 23 Mbps



1. Draft paper at: <http://www.internet2.edu/abilene/tcp/>

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/>

Low demand

Router support for DiffServ is spotty; e.g...

- No PQ
- DiffServ comes with a performance cost
- Limitations on token bucket depths
- Inflexible classification rules

Dramatic changes to network operations, peering arrangements, and business models

Requires all-or-nothing network upgrades (e.g. all access interfaces must police)

“Worse”

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

“*Different–but–equal*”

- Alternative Best Effort (ABE)

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

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

All traffic is not equal

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

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

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



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/>

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

- Survey paper of application QoS needs
- Measurement and analysis to understand application performance and use of new services
- Investigating the value of a service-specific service initiative (e.g. a VoIP network service)

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/>

QBone Home:

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

“Future Priorities for Internet2 QoS” paper:

- <http://www.internet2.edu/qos/wg/papers/qosFuture01.pdf>



INTERNETTM