QoS in the Internet

Juha Oinonen
Development Manager, Funet services

Internet

Technology for
- providing global end-to-end connectivity
- via networks under no central coordination

QoS Bottleneck
- QoS, like connectivity, is end-to-end
- QoS of the network is the QoS of the worst part of the network

Problem of Traditional QoS Approach
- QoS requires coordination and joint effort
- There's no central control in the Internet (except for standards)

Global end-to-end requirement is a hard part when introducing new things

Aim of QoS (in the Internet)

To improve

- the observed/measured(?)
- “quality of service” f
- or some of the users/applications
- (by using standard IP techniques)

“Quality of Service”:

- IP service availability (connectivity, packet loss)
- IP service consistency (delay variation)
- upper level (perceivable) quality (sound, image, ...)

Aim of Commercial QoS Service:

- to make more income and profit

Reasons for Quality Problems in the Internet

Nature of packet networks and protocols currently used
- TCP flow control based on packet drop caused by congestion

Network design
- tree vs. ring topology; shared layer 2 links

Insufficient capacity (link, device)
- over-provisioning (on purpose or by mistake)

Poor network managing
- don’t know there’s a problem
- fixing takes a long time

Device problems (hw, sw)
- bugs; loss of power

Physical issues
- dirt, interference, ...

Internet QoS: Tools

DiffServ - generic set of tools

- allow different queue/drop behavior to packets marked differently
  - use two or more queues
  - dropping of packets possible even if no link congestion

Congestion control (for TCP):

- RED - handling mechanism on routers only
  - when congestion is near, router selectively drop (TCP) packets
  - can be implemented incrementally
- ECN - handling mechanism on both routers and end devices
  - when congestion is near, router marks packets so end devices can react
  - requires both router(s) and end devices to be ECN-capable

Traditional IP QoS: Premium service

ISP guarantees better-than-BE quality to selected packets

- assumption: BE is not good enough (for some applications/users)
- selected = more important = more expensive = more income
- lower queue delay and/or drop probability
- the ”guarantee” part problematic
  - must be implemented through the whole path
  - must define what the better quality means (SLA)
  - the amount of marked packets must be limited to a fraction of total bandwidth - limiting must be done close to the source, and on administrative borders
  - must measure in order to be able to say the quality is there (and money should be paid)
- in some cases, works on a “on-demand” basis, requiring manpower to activate and deactivate
- proposed bandwidth brokers add a new layer of control and complexity

Emerging IP QoS: Less-than-Best-Effort

User allows less-than-BE quality for selected packets

- assumpition: BE is good enough (for most cases everything)
- selected = less important/time-critical (= less expensive?)
- higher drop probability and/or queue delay
- many thing are easy:
  - *no SLA hassle*
  - works the better the more packets are selected (no risk of overloading or misuse as DoS)
  - can be implemented incrementally: transparency is enough to start with
- to work, requires a significant portion of packets selected
  - ok, we know 1% of flows takes 80% of bandwidth, so target those!
  - motivation must be found somehow

Based on available technology; no scaling issues :)

Current QoS Positioning

Commercial: Premium service

- business model (have SLA, pay more) and customers are there
- takes effort to put up, maintain, and verify SLAs
- does not scale to Internet: SLA between and over QoS domains...

Academic: turning to LBE

- everything's free (well, flat rate :), but it's polite not to waste
- easy to put up: no burden on defining or verifying SLAs
- scales to Internet (no end-to-end problem)
- commercial model missing at the moment

QoS in some Networks

Current status

- case Géant: Premium service, going towards LBE as well
- case Abilene: LBE only
  - Premium service considered and dropped
- case Funet: LBE pilot announced

- case FICIX: Premium service
  - some ISP:s offer Premium service
  - “FICIX SLA?” - we’ll see...

- many cases: plain Best-Effort
  - both ISP networks and LAN:s

QoS Bottleneck at the First Mile

ISP to customer, or inside a LAN

- shared media (Ethernet, ATM; cable, wireless)
- L2 devices must be L3-aware to allow DiffServ, RED and ECN
- DiffServ typically not implemented in low-end L3 devices

PCs and other IP devices

- not designed for QoS: linecard, driver, OS, application

QoS future (3 years perspective)

RED and ECN make the TCP ”slow start” problem (burst generation) easier

The QoS botteleneck in LANs and other last-mile solutions is still there

SLA based Premium service found impossible in Internet context

- still making money *inside operators' networks*

LBE will come to commercial networks and interconnects – part of current best effort traffic will be marked as LBE

- by ISP's own decision (straightforward to implement, but some risk with customer satisfaction as no choice given)
- by asking for volunteers and giving discount (socially better solution, but might require per-packet accounting per customer)
- ISP could offer both options per-customer basis