Adding Enhanced Services to the Internet: Lessons from History

k.c. claffy, David Clark
UCSD, MIT
September, 2015

Why enhanced services?

- Other terms
  - Quality of Service (QoS).
  - Differentiated services.
- It was understood, even in the 1980's, that certain applications, such as voice, video, and multi-player games*, benefitted from low-latency, low jitter delivery, and other applications benefitted from high-bandwidth service.
  - One service (a single class of "best effort"), may not provide the necessary support for this range of apps.

* SIMNET, a tactical tank-training simulator used (among other things) to rehearse for the first Gulf War.

This paper

- A look at 30+ years of failure to deploy differentiated services (QoS) in the public Internet.
  - A catalog of the many reasons why.
  - Some personal history.
- Some observations about what is different now.
  - Trends in the shape of the Internet ecosystem.

The paper—in decades

- The 1980's
  - Initial specification and operational experience.
- The 1990's
  - Standardization of mechanisms for differentiated services.
  - Non-technical barriers emerge.
- The 2000's
  - Frustration and surprise.
  - Regulatory concerns emerge.
- The 2010's and into the future.

The 1980's (early)

- The idea of differentiated services as been a part of the Internet from the beginning.
- The IP header has a field called Type of Service (ToS).
  - RFC 791 (1981)

```
Bits 0-2: Precedence.
Bit 3: 0 = Normal Delay, 1 = Low Delay.
Bits 4: 0 = Normal Throughput, 1 = High Throughput.
Bits 5: 0 = Normal Relibility, 1 = High Relibility.
Bit 6-7: Reserved for Future Use.
   0     1     2     3     4     5     6     7
+-----+-----+-----+-----+-----+-----+-----+-----+
|                 |     |     |     |     |     |
|   PRECEDENCE    |  D  |  T  |  R  |  0  |  0  |
|                 |     |     |     |     |     |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

Precedence

- 111 – Network Control
- 110 – Internetwork Control
- 101 – CRITIC/ECP
- 100 – Flash Override
- 011 – Flash
- 010 – Immediate
- 001 – Priority
- 000 – Routine

Design principles for ToS

- The end-node (application) sets the ToS field, the network honors it.
  - In contrast to a scheme where the network elements peek at what the application is and picks the ToS.
- The ToS specifies a service (e.g., low delay), not the network mechanism to achieve it.
  - RFC 795 (1981)

The 1980’s (mid)

- ARPAnet is replaced by the NSFnet.
- First generation backbone links are 56 kb/s.
  - Faster than the ARPAnet!!
- Congestion emerges as serious problem.
  - A particular problem for interactive apps like remote login.
- Solution: give telnet packets higher precedence.
  - But the end-nodes are not setting the ToS field, so the routers peek at the header.

The 1980’s (late)

- Van Jacobson proposes algorithm to limit congestion.
  - Called “slow start”, it has evolved, but is still in use today.
- Slow start does not fix all QoS issues.
  - Queues still form in routers.
    - So episodes of higher delay and jitter.
  - Interactive apps still suffer.

The 1990's (NSF)

- As NSF backbone gets faster (1.5 mb/s, then 45 mb/s) immediate pressure from congestion abates.
- But wise people see the future: the popularity of real-time applications "*bodes ominously for an infrastructure not able to: distinguish among certain traffic types; provide more than a best-effort guarantee to datagram traffic; or upgrade in a time efficient way towards an availability of higher bandwidth (if only due to lack of accounting and billing mechanisms to enable cost recovery*)."
  - (Bohn et al., 1994)

The 1990's (IETF)

- Idea emerges of an "integrated services network": a single network that can carry every sort of traffic.
  - Phone company liked the idea, but not IP.
    - IP "best effort" service judged too unpredictable.
  - Their answer: Asynchronous Transfer Mode (ATM).
    - ATM supports differentiated service for data and voice.
- What was the IETF to do?
  - Stick to best effort or accept differentiated services?
  - They decide to go for differentiated services.

Two services

- Guaranteed service
  - Mimics guarantee of phone system.
  - Computable bound on jitter.
  - Led to the *intserv* standards.
- Predictive service
  - Gives probabilistic bound.
  - Led to the *diffserv* standards.

intserv

- Gives per-flow guarantee.
- Very complex: per-flow setup and per-flow state in router.
  - Quite inconsistent with core philosophy of Internet.
- The bound on the jitter is computable, a tight bound, and for traffic with any degree of burstiness, not appealing.
  - Must provision at close to peak rate to get reasonable guaranteed bound.
    - PhD thesis by Abhay Parekh at MIT

diffserv

- Deals with aggregates, not individual flows.
  - All packets in aggregate are mixed and treated alike.
- Set of standards, each defining different treatment of the aggregate.
  - Most common is Expedited Forwarding (EF) treatment.
    - Provides better jitter characteristics than normal best effort.
  - Bound is not computable in advance.
  - Hence "predictive" service.
    - Think about HOV lanes.
  - Requires some sort of admission control.
- ToS field redefined to select different treatments.
  - Differentiated Service Code Points (DSCPs).

Service or local function?

- Original ToS field specified a *service.*
  - For example, *low jitter* service.
- *Intserv* specified both a service and what the router needs to do.
  - *"The end-to-end behavior provided to an application by a series of network elements providing controlled-load service tightly approximates the behavior visible to applications receiving best-effort service *under unloaded conditions* from the same series of network elements."*
- But the IETF had no tradition of standardizing services.

Diffserv—what to specify?

- IETF retreated from specification of service to specification of what a component does.
  - Specify *per-hop behavior* (PHB).
- Diffserv overview RFC 2475
  - "It is strongly recommended that an appendix be provided with each PHB specification that considers the implications of the proposed behavior on current and potential services."
- EF RFC 2598
  - "Note that the EF PHB only defines the behavior of a single node. The specification of behavior of a collection of nodes is outside the scope of this document."

Troubling implication

- What an application wants to invoke is a *service*.
  - For an application to be designed to use a service, that service needs to be standardized.
- What an ISP might want to offer to its customers is a *service*.
- There was (is) no venue in which to standardize the service.
  - A substantial barrier to deployment
  - Especially for multi-provider QoS, where ISPs have to agree on what the service is and how to realize it.

Economics

- If service discrimination is not a *service*, can an ISP make money from it?
  - A service is something you sell.
  - Technology is something you buy.
- My personal experience trying to “sell” QoS mechanisms.

The 2000's (early)

- SigComm holds a wake for QoS.
  - A workshop called RIPQOS.
- Big surprise: QoS mechanisms are alive and well.
  - Davie, B. (2003). Deployment experience with differentiated services
  - Technically, they work fine.
  - But not in the public Internet.
  - In private IP intranets.

The 2000's (mid)

- MIT tries to get ISPs to talk about interprovider QoS deployment.
  - We did produce a document, which does not seem to have had much impact.
- What we found were the barriers.
  - Cannot talk about pricing at all.
    - Fear of anti-trust.
  - Necessity of sharing operational information is severe problem.
    - To offer a characterized service, must share much more information among operators.

The 2000's (later)

- Fears of abuse shape the discourse.
  - If ISPs cannot make practical use of QoS to enhance certain flows, perhaps they can use these same tools to *disadvantage* certain flows.
  - Traffic shaping, for example.
    - P2P in UK.
- Was the development of QoS tools a necessary step to support apps like voice, video and games? Or was it a fundamental threat to the open Internet?

The perverse incentive

- QoS enhancement (using packet scheduling) only matters if there is congestion.
  - No congestion means no queues, means no jitter, means no queue of packets to schedule.
- In the early days, congestion was chronic.
  - But perhaps now we live in a time of plenty?
- Should an ISP, dealing with increasing traffic:
  a) invest to add more capacity?
  b) make money selling QoS over the congested network?

The 2010's (now)

- Two important trends:
  - Interconnection: from transit to peering to *direct interconnection.*
  - The rise of multi-firm IP networks distinct from "the Internet".
- Two observations:
  - Congestion (overloads) only occur in specific places today.
  - What matters to users is Quality of Experience (QoE).

Direct interconnection

- Connections from content providers (CDNs) are *not* the same thing as interconnections among peers.
- Majority of traffic flowing into broadband provider networks is over direct interconnection.
- Implications (for QoS):
  - Inside CDNs, there are no prohibitions against use of QoS.
  - The resulting traffic only crosses one "ISP" to reach destination.
- Is this a service opportunity with legitimate benefit?
  - FCC has forbidden "paid priority".

Multi-firm IP networks

- Today we see “carrier-grade VoIP” carried over IP, but over distinct networks.
- What distinguishes these networks?
  - Better control over QoS parameters.
    - Restricted range of services.
    - Careful provisioning.
    - Traffic discrimination at interconnection points.
- Once again (as in the 1990s) should “we” continue to argue that best effort is “good enough” for the Internet, or instead permit beneficial use of QoS to allow OTT apps to compete with these networks?
  - Third-party content may start to move onto these alternative networks.

Where is congestion today?

- In specific places:
  - Across the access link.
  - On cellular networks.
  - At some interconnection points.
  - In the developing world.
- Perhaps point solutions are the right idea.
  - But congestion may continue to move.
    - Gigabit access links may be shifting congestion back into the core of the net.

Beneficial uses

- FCC has argued that QoS must be application-agnostic or under control of user.
  - Former is not addressing this problem, latter requires standardization.
- How could "we" allow the use of QoS but only for beneficial uses?
  - Disclosure by ISPs of their practices and the justification for them.
    - Arguable this causes no harm to business.
  - A venue for discussion/standardization of service offerings.

Quality of Experience (QoE)

- What we typically measure about networks are technical parameters:
  - Throughput, loss rate, delay, jitter, etc.
- What the user cares about is the performance of applications:
  - Rebuffering in video streaming, lag in games and VoIP.
- Arguments about benefit should be cast in terms of QoE.
  - But we have no agreed methods to measure it.
  - Our research challenge for the future.

Final thought

- Differentiated service tools (e.g., diffserv) are not pixie dust.
- They cannot solve all QoE impairments.
  - For example, will not solve problems of overload on direct interconnection between Netflix and access provider.