

Adding Enhanced Services to the Internet: Lessons from History

kc claffy and David D. Clark
kc@caida.org and ddc@csail.mit.edu

September 7, 2015

Contents

1	Introduction	3
2	Related technical concepts	4
2.1	What does “enhanced services” mean?	4
2.2	Using enhanced services to mitigate congestion	5
2.3	Quality of Service (QoS) vs. Quality of Experience (QoE)	6
2.4	Limiting the use of enhanced services via regulation	7
3	Early history of enhanced services: technology and operations (1980s)	7
3.1	Early 1980s: Initial specification of <i>Type-of-Service</i> in the Internet Protocol suite	7
3.2	Mid 1980s: Reactive use of service differentiation to mitigate NSFNET congestion	10
3.3	Late 1980s: TCP protocol algorithmic support for dampening congestion	10
4	Formalizing support for enhanced services across ISPs (1990s)	11
4.1	Proposed short-term solution: formalize use of IP Precedence field	11
4.2	Proposed long-term solution: standardizing support for enhanced services	12
4.3	Standardization of enhanced service in the IETF	13
4.4	Revealing moments: the greater obstacle is economics not technology	15
5	Non-technical barriers to enhanced services on the Internet (2000s)	15
5.1	Early 2000s: a funeral wake for QoS	15
5.2	Mid 2000s: Working with industry to gain insight	16
5.3	Late 2000s: QoS becomes a public policy issue	17
6	Evolving interconnection structure and implications for enhanced services (2010s)	20
6.1	Expansion of network interconnection scale and scope	20
6.2	Emergence of private IP-based platforms to support enhanced services	22
6.3	Advancing our empirical understanding of performance impairments	23
7	Lessons learned	25

Abstract

We revisit the last 35 years of history related to the design and specification of Quality of Service (QoS) on the Internet, in hopes of offering some clarity to the current debates around service differentiation. We describe the continual failure to get QoS capabilities deployed on the public Internet, including the technical challenges of the 1980s and 1990s, the market-oriented (business) challenges of the 1990s and 2000s, and recent regulatory challenges. Our historical perspective draws on, among other things, our own work from the 1990s that offered proposals for supporting enhanced services using the Internet Protocol (IP) suite, and our attempts to engage both industry and policymakers in understanding the dynamics of the Internet ecosystem. In short, the engineering community successfully developed protocols and mechanisms to implement enhanced services (QoS), and a few individual service providers have deployed them internally or in trusted two-party scenarios. The long-standing failure has been to deploy this capability across the public Internet.

We reflect on lessons learned from the history of this failure, the resulting tensions and risks, and their implications for the future of Internet infrastructure regulation. First, the continued failure of QoS over the last three decades derives from political and economic (business) obstacles as well as technical obstacles. The competitive nature of the industry, and a long history of anti-trust regulation (at least in the U.S.) conflicts with the need for competing providers to agree on protocols that require sharing operational data with each other to parameterize and verify committed service qualities. Second, QoS technology can yield benefits as well as harms, so policymaking should focus on harms rather than mechanisms. To assure the benefit to consumers, regulators may need to require transparency about the state of congestion and provisioning on networks using such mechanisms. Third, using QoE as the basis for any regulation will require research, tools and capabilities to measure, quantify, and characterize QoE, and developing metrics of service quality that better reflect our understanding of QoS and QoE for a range of applications. Finally, profound shifts in interconnection arrangements suggest a reshaping of the debate over QoS on the public Internet. Some access networks are interconnecting their private IP-based network platforms to support enhanced services, and using this interconnected platform to vertically integrate infrastructure and applications. Access networks are also connecting directly to large content providers to minimize the risk of performance impairments. These changes trigger new regulatory concerns over the fate of the public Internet, including capital investment incentives and gaps across different bodies of law.

Barriers to the deployment of scalable interprovider QoS may be unsurmountable, but since any Internet of the future will face them, it is worth developing a systematic understanding to the challenge of enhanced services, and documenting successes and failures over the history of the Internet as carefully as possible.

1 Introduction

Recent rhetoric around network neutrality might suggest that the Internet has a history of neutrality, and deviations from neutrality are a recent consequence of profit-seeking operators adapting to rapidly growing traffic demands. In this paper, we draw on over 30 years of personal involvement in the design and specification of enhanced services on the Internet (also described in terms of Quality of Service or QoS), to put the current debates into context, and dispel some confusion that swirls around service differentiation. We describe the thirty-year failure to get QoS capabilities deployed on the public Internet, including technical, economic and regulatory barriers, and we consider implications of this failure for the future.

We first review terminology and technical issues that shape QoS implementations, and how they can affect applications running over the Internet (Section 2). We then offer a historical perspective. The first two decades of history (Sections 3 and 4) include dramatic episodes of dangerous levels of congestion on Internet backbones, proposals to augment the IP protocol suite to support enhanced services, and proposals of interim solutions to alleviate congestion while new protocols are being developed and deployed. Although protocols now exist to support enhanced services, nowhere on the global Internet today are these services supported across multiple ISPs.

Section 5 reviews the third decade of history, when it became clear that although technical mechanisms for QoS were (and are) deployed in private IP-based networks, global deployment of QoS mechanisms foundered due to coordination failures among ISPs. Although the Internet has a standards body (the IETF) to resolve technical issues, it lacks any similar forum to discuss business issues such as how to allocate revenues among competing ISPs offering enhanced services. In the U.S., ISPs feared such discussions would risk anti-trust scrutiny. Thus, lacking a way to negotiate the business implications of QoS, it was considered a cost rather than a potential source of revenue. Yet, the relentless growth of a diversity of applications with widely varying performance requirements continued on the public Internet, with ISPs using relatively primitive, and not always completely benign, mechanisms for handling them.

Section 6 describes a transformation of the interconnection ecosystem that has codified this decade, driven by performance and cost optimizations, and how this transformation is reshaping the landscape and the role of enhanced services in it. The role of Internet exchanges (IXes) as anchor points in the mesh of interconnection has enabled, as well as benefited from, the growing role of content providers and Content Delivery Networks (CDNs) as major sources of traffic flowing into the Internet. By some accounts over half the traffic volume in North America now comes from just two content distributors (Youtube and Netflix). This shift constitutes the rise of a new kind of hierarchy in the ecosystem, bringing new constraints on existing players who need to manage traffic on their networks to minimize congestion. Evidence of trouble has increased this decade, resulting in tussles among commercial players as well as between the private sector and regulatory bodies, at the expense of users suffering degraded quality of experience.

More recently, a few large ISPs have built private interconnected IP-based networks, using the IP technology but distinct from the public Internet, and able to support higher service qualities. The emergence of these private networks raises the question: are we better off driving capital investment into private interconnected (and unregulated) IP networks that can offer enhanced QoS, thus limiting the public Internet to best-effort and lower revenue services?

With an engineering perspective on this evolution, we reflect on lessons learned from three decades of experience trying to solve the "multi-provider QoS problem", the resulting tensions and risks, their implications for the future of Internet infrastructure regulation.

First, the continued failure of QoS over the last three decades derives from political and economic (business) obstacles as well as technical obstacles. The competitive nature of the industry,

and a long history of anti-trust regulation (at least in the U.S.) conflicts with the need for competing providers to agree on protocols that require sharing operational data with each other to parameterize and verify committed service qualities.

Second, QoS technology can yield benefits as well as harms. Thus in our view policymaking should try to focus on regulating harms rather than mechanisms. If deployed appropriately, QoS technology can benefit latency-sensitive traffic without impairing performance of latency-insensitive traffic. But to assure consumers and regulators that deployment of QoS technology (or other traffic management mechanisms) will fall into this “no collateral harm” category, i.e., not impair the Quality of Experience (QoE) for users, regulators may need to require transparency about the state of congestion and provisioning on networks using such mechanisms.

Third, we emphasize the daunting technical obstacles to the use of QoE impairment as a basis for regulation. It will require research, tools and capabilities to measure, quantify, and characterize QoE, and developing metrics of service quality that better reflect our understanding of QoS and QoE for a range of applications.

Finally, shifting patterns of interconnection warrant a closer examination of the challenges and risks associated with enhanced services on the public Internet. The risk is the potential to incentivize artificial scarcity as opposed to pro-growth strategies. Network operators can respond to congestion on their networks by either increasing capacity, or selling enhanced services to those willing to pay for their applications to work well during congestion periods. The tension is exacerbated by the lack of competition at the (wired and wireless) access level, and a lack of transparency into whether these perverse incentives are prevailing at any given time.

The co-evolution of regulatory, legal, business, and technological capabilities, all at different paces, is tightly coupled in the case of enhanced services – a quintessential interdisciplinary challenge. While, as we document, the barriers to the deployment of scalable interprovider QoS on today’s Internet may be insurmountable, the issues and challenges remain. If any Internet of the future is to be the platform for the full range of useful applications, either the basic service of that Internet must be so improved that no support for differentiated services is necessary, or it will be necessary to overcome these challenges in some way. For this reason, it is worth developing a systematic understanding to the challenge of enhanced services, and documenting successes and failures over the history of the Internet as carefully as possible.

2 Related technical concepts

In this section we review terminology and concepts related to enhanced services on the Internet.

2.1 What does “enhanced services” mean?

The traditional service model of the Internet is called *best effort*, which means that the networks that make up the Internet try their best to deliver packets in a timely and reliable manner, but there are no specifications that define a minimal acceptable service. IP networks may lose packets, deliver them out of order or more than once, or deliver them after inexplicable delays. Some IP networks, for example wireline ISPs in most parts of the developed world do a pretty good job of delivering packets. Other parts of the Internet, including some wireless networks or networks in developing regions, do less well.

One layer of the IP protocol suite (Postel, 1981) is called the Transmission Control Protocol (TCP). TCP runs in end nodes attached to the Internet and tries to mitigate these sources of unreliability from lower layers. Two end nodes participating in a TCP connection with each other will:

number packets they send, track which packets arrive, retransmit those that are lost, re-order them if necessary, and deliver the data to the receiver as it was initially sent by the sender. However, TCP cannot remedy all problems. Delays due to congestion or retransmission of lost packets cannot be reversed. Different sources of unreliability, combined with TCP's transmission algorithms, can lead to long and variable delays from the sending application to the receiving application. There is no simply no guarantee of a certain amount of bandwidth from sender to receiver.

For many applications, this variation in delay and bandwidth is tolerable. If a web page loads a few seconds more slowly, or an email is delayed a few seconds, the user is not really disadvantaged. But for some applications, especially *real-time* applications, this variation can be disruptive. An audio or video conversation that suddenly suffers seconds of delay becomes essentially unusable. Similarly, unpredictable variation in bandwidth can disrupt the viewing of video.

There are competing views on how to improve this situation. One approach is to increase our expectation of minimum "best effort" to a threshold that allows real-time application to work. Many ISPs in the developed world advertise that their basic services enable real-time multimedia applications to work quite well. This assertion may often be true. However, the two of us, as collaborators on opposite sides of the country, often must disable video in a Skype call to get the audio to work. Disruption of real-time applications is a reality today. Further, for the typical user, there are no tools to help understand why.

Whether parts of the Internet are working well enough today to support demanding applications, congestion and lack of capacity is not equally easy to remedy to this level of performance in all parts of the Internet, so another view is that ISPs should provide different treatment of packets as they cross the Internet, giving preferential treatment to packets from more demanding applications, and shifting less sensitive applications to lower priority forwarding. Support for this sort of differential traffic treatment is also called *enhanced services* (Clark, Shenker, & Zhang, 1992), and one goal of this paper is to see what lessons we can glean from the (failed) history of trying to get enhanced services deployed in the Internet.

2.2 Using enhanced services to mitigate congestion

The term *congestion* has multiple interpretations (Bauer, Clark, & Lehr, 2009), but most simply, congestion arises when the offered load exceeds the capacity at points within the network. Different networks deal with congestion differently. The traditional telephone network, a *virtual circuit* network, dealt with this possibility by reserving enough capacity for a call at every link along the path through the network when the call was being initiated. If there was not sufficient capacity, the call would not go through.¹ This *admission control* approach made sense when the network only carried one sort of traffic (phone calls). In contrast, the Internet was designed under the assumption that it would carry a wide range of traffic, from long-lived flows to short message exchanges. The designers rejected any sort of admission control: most applications (like file transfer) did not have a well-defined target service requirement, and the overhead of admission control seemed excessive for small interactions, which might be completed with few packets in the time it would take to process the admission control request. The design approach in the Internet is that applications simply sent traffic as needed. This approach, in contrast to virtual circuits, was called a *datagram* network. In this case, the consequence of excess offered load is that, in the short run, queues of packets form, and over the longer term, either some congestion control algorithm causes the senders to slow down or else overloaded routers run out of memory to store queued

¹The busiest day of the year for the phone system was reputed to be Mother's Day, and in past time fears of not getting through made dutiful children call home early in the day. The indication that a call was blocked due to lack of resources was a "busy signal" delivered twice as fast as normal.

packets and thus drop some packets. In contrast to the telephone network, this *best effort* service of the Internet made no performance commitments, which implies potential impairment of certain applications, in particular so-called *real-time* applications like voice.

The idea of enhanced services is to add mechanisms to the network that differentially treat traffic in order to mitigate these impairments. There are two general approaches: use different paths through the network for the different flows, or treat the packets from different flows differently as they pass through the same router. Early designers of the Internet contemplated but never implemented the first option. The other approach – distinct treatment of distinct flows through a single router – mostly applies when a router experiences congestion. When a link leaving a router becomes overloaded, a queue of packets forms; the router can implement different service qualities by scheduling packets in different ways, e.g., giving packets of real time flows priority by jumping their packets to the head of the queue. If there is no congestion, there is no queue, so there are no options for re-arranging the queue.²

As concerns over Internet congestion waxed and waned over the years, so has interest in enhanced services. Interestingly, the locus of congestion has shifted over time. In the early days of the Internet, where the trans-continental links were 50 kb/s and there were not yet effective algorithms to control overloads, congestion in the core of the Internet was commonplace. By the time dialup residential access became common, the capacity of the core had increased, and because dialup speeds were initially at such low rates (9600 b/s, for example), the access link was often the point of overload, and the core of the network was less often congested.

At the time, there was a lot of (investment-bubble funded) fiber available to build out the core, and even the advent of broadband residential access was not enough to shift congestion to the core. Today, in many parts of the Internet, especially wireline access providers, the facilities of individual ISPs are often provisioned with adequate capacity to carry the traffic from their customers, in which case the access link is the point of overload, unless the overload occurs at points of interconnection among ISPs (Bauer et al., 2009). On the other hand, some wireless access networks show evidence of congestion in parts of their infrastructure. This variation raises a design question: if enhanced services are added to the Internet, should the mechanisms be general enough to deal with congestion wherever it manifests, which calls for multi-provider coordination with all the complexity it implies, or customized to deal with a specific type of congestion prevalent at a particular region and time, which might be easier to implement but of diminishing (or at least inconsistent) value over time as the nature and location of congestion evolves.

2.3 Quality of Service (QoS) vs. Quality of Experience (QoE)

The term Quality of Service (QoS) is used to describe technical measures of network performance, such as throughput, latency, jitter and packet loss. However, users do not directly perceive these parameters. Users perceive how their applications work, and whether there is degradation in the way the application performs. The term used to describe user-perceived impairments is Quality of Experience (QoE). QoE is a subjective measure of quality, and assessment of QoE is application-specific. Was voice quality degraded? Is the video stream experiencing re-buffering delays? Impairments to QoE derive in part from underlying QoS parameters (e.g., jitter impairs voice QoE but not email) and can arise anywhere in the network.

Network engineers can measure parameters of QoS, but they cannot directly measure QoE. The application designer can attempt to estimate QoE, perhaps by a questionnaire at the end of the

²Here we ignore the particular case where packets from a specific flow are rate-limited, perhaps as a result of a service agreement, even though there is excess capacity.

application session, or by observing user behavior (did the user abandon the application session in the middle). Laboratory research with human subjects can try to correlate variation in QoS with the resulting variation in QoE by subjecting subjects to controlled variation in QoS and asking them to report on their assessment of QoE. But this information informs the network engineer only indirectly. As a practical matter, the linkage between QoS and QoE is sometimes obscure; many technical quirks in the system can produce observable impairment in QoE. As well, our understanding of QoE and how to measure it is not as well developed as our understanding of the technical dimensions of QoS. So the argument that deployment of enhanced services is justified depends on a line of reasoning that is subject to argument—is the proposed differentiation actually going to mitigate the QoE impairment?

2.4 Limiting the use of enhanced services via regulation

The notion of *best effort* suggests that all packets receive the same treatment, which superficially implies a sense of fairness even if it actually impairs certain applications without measurably benefiting any others. Much regulatory attention has focused over the last several years on the possibility that Internet Service Providers (ISPs) will discriminate among traffic flows for business rather than technical reasons. The rallying cry for intervention to prohibit this sort of behavior is called “*network neutrality*”, and the more nuanced but still ill-defined term used by regulators such as the FCC is *reasonable network management*, which requires a “*reasonable*” technical rather than business justification.

3 Early history of enhanced services: technology and operations (1980s)

In this section we review the early intentions of the Internet designers to support distinct types of service qualities in the architecture, and protocol features to implement them.

3.1 Early 1980s: Initial specification of *Type-of-Service* in the Internet Protocol suite

The idea of enhanced or differentiated services has been a part of the Internet’s design from the beginning. The Internet Protocol (first defined by RFC 760 in 1980) had as a part of its header an 8-bit field called Type of Service, or ToS. Box 1 describes how the designers conceived of this ToS field at the time. This field enabled specification of service quality using five parameters, some of which were supported by some networks at the time: precedence, stream vs. datagram, reliability, speed vs. reliability, and speed. The specific precedence terminology mirrored then-current controls in military command and control networks, and allowed the sender to specify the importance of the message. Military message networks of the time allowed different messages to be tagged with different precedence, which influenced the order of delivery. “Flash Override” was the most urgent. At the time, it was unclear how this concept mapped to the packet level of the Internet, but the designers were encouraged to add this feature to make the network more familiar to potential military users.

Box 1
Definition of Type of Service field (RFC 760)

The Type of Service provides an indication of the abstract parameters of the quality of service desired. These parameters are to be used to guide the selection of the actual service parameters when transmitting a datagram through a particular network. Several networks offer service precedence, which somehow treats high precedence traffic as more important than other traffic. A few networks offer a Stream service, whereby one can achieve a smoother service at some cost. Typically this involves the reservation of resources within the network. Another choice involves a low-delay vs. high-reliability trade off. Typically networks invoke more complex (and delay producing) mechanisms as the need for reliability increases. (RFC 760, page 11)

- Bits 0-2: Precedence.
- Bit 3: Stream or Datagram.
- Bits 4-5: Reliability.
- Bit 6: Speed over Reliability.
- Bits 7: Speed.

0	1	2	3	4	5	6	7
+-----+-----+-----+-----+-----+-----+-----+-----+							
PRECEDENCE			STRM RELIABILITY		S/R		SPEED
+-----+-----+-----+-----+-----+-----+-----+-----+							
111-Flash Override			1-STRM		11-highest		1-spd 1-hi
110-Flash			0-DGRM		10-higher		0-rlb 0-lo
11X-Immediate				01-lower			
01X-Priority				00-lowest			
00X-Routine							

This version of the IP specification was “pre-standard”, in that RFC 791 was issued a year later (1981), which is generally taken as the first official standard for the Internet Protocol. RFC791 presented an updated definition of the ToS field (Box 2). While the designers still had no experience with implementing any ToS function, they decided to separate the bits used to code delay, throughput and reliability requirements. They also removed the distinction between stream vs. datagram models of flows, as the design philosophy had evolved to a pure datagram model. The additional precedence levels reflect the realization that network control traffic was necessary to keep the network itself working, and thus of the highest importance.

At the time, there was no use of these bits to control how routers treated packets. The intention was to enable differential control of traffic on networks that made up the early Internet, e.g., the ARPAnet, the packet radio network PRNET, or the satellite network SATNET.³ The idea was that the Internet ToS field would be used to set the network-specific ToS features of the various network

³RFC760 (Postel, 1980) pg 27 states, “For example, the ARPANET has a priority bit, and a choice between “standard” messages (type 0) and “uncontrolled” messages (type 3), (the choice between single packet and multipacket messages can also be considered a service parameter). The uncontrolled messages tend to be less reliably delivered and suffer less delay.”

technologies being used in the Internet of the time. RFC 795 (Jon Postel, 1981) catalogued these service mappings from IP ToS to network-specific ToS mechanisms. Although based on the then best understanding of service variation, this specification was in some respects a place-holder, and these fields were substantially redefined in the 1990's. (Section 4).

Box 2
Definition of the IP TOS field as of RFC 791

The IP Type of Service has the following fields:

Bits 0-2: Precedence.
 Bit 3: 0 = Normal Delay, 1 = Low Delay.
 Bits 4: 0 = Normal Throughput, 1 = High Throughput.
 Bits 5: 0 = Normal Reliability, 1 = High Reliability.
 Bit 6-7: Reserved for Future Use.

0	1	2	3	4	5	6	7
PRECEDENCE	D	T	R	0	0		

Precedence values:

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

The important point of this early specification was not exactly how the fields were defined, but the intention to support endpoint requests for a specific quality of service via setting bits in the packet header. An alternative approach would have been for the router to deduce what application was sending the data (perhaps by looking at the port field in the TCP header) and map the application to a service class. Drawbacks to this approach were well-understood at the time. First, not all applications used port fields in the same way, even then. Second, the port fields are in a higher-layer (transport) protocol header, which might be encrypted. The design intent was that code running on the endpoint, and reflecting the desires of the end-node for the requested service, would set the TOS fields, and the network would honor them. For some sub-fields, such as the precedence field, there was no way the router could consistently deduce the correct value, because the precedence field was conceived as differentiating messages of the same sort (e.g., email) based on the sender's assessment of the importance of the content, and the necessity of rapid delivery. Only the sender could determine the importance of the message.

3.2 Mid 1980s: Reactive use of service differentiation to mitigate NSFNET congestion

As related in (Bohn, Braun, Claffy, & Wolff, 1994), in the mid-1980's, the U.S. National Science Foundation commissioned a national backbone network to replace the original ARPANET, which was planned for decommissioning by ARPA. The original NSFNET backbone links were 56kb/s, and gradually became congested. In response the NSFNET engineers deployed an emergency measure to provide certain interactive network applications, specifically remote login (telnet), preferential treatment. This is probably the first example of enhanced services (or service differentiation) being used in the Internet.

The routers that made up the early NSFNET backbone were programmed by David Mills at the University of Delaware. The hardware for each backbone router was an LSI-11, a modest computer, even for those times, and the device was called, perhaps lovingly, a "Fuzzball". The Fuzzball software supported queue management schemes (primarily a priority scheduling scheme) to support multiple service classes, which allowed for prioritizing the most delay-sensitive applications of the time, remote login (telnet). While the Fuzzball software was designed to use the TOS field to trigger service differentiation, the remote login software running on end-nodes was not setting these bits. So in perhaps the first example of routers peeking at higher-level protocol fields in the header, the Fuzzball looked at the port field to see if the protocol being used was remote login. In 1988, (Mills, 1988) wrote: *"However, many implementations lack the ability to provide meaningful values and insert them in this [TOS] field. Accordingly, the Fuzzball cheats gloriously by impugning a precedence value of one in case the field is zero and the datagram belongs to a TCP session involving the virtual-terminal TELNET protocol."*⁴

Priority treatment allowed interactive users to continue working under highly congested circumstances. But because the backbone administrators did not have any way to provide an incentive to not use the highest priority, they did not publicize the priority-based treatment of traffic, and end users did not know it was possible to give high precedence to other applications. It was a short-term operational tactic, not a long-term strategy.

3.3 Late 1980s: TCP protocol algorithmic support for dampening congestion

While the initial design of the Internet included the specification of a mechanism to signal congestion (the ICMP Source Quench, see RFC 792), there was no success in designing an effective overall mechanism to dampen congestion. The traffic prioritization described above only sorted through packets that were not dropped by routers, but did nothing to dampen congestion. In the late 1980's, Van Jacobson proposed a workable scheme for congestion control (Jacobson, 1988). With refinements, this scheme is still the basis of TCP congestion control in the Internet, and immediately mitigated congestion in the overloaded core of the NSFNET backbone.

However, it did not resolve all congestion-related issues. First, it did not prevent queues from forming. The signal of congestion from the network to the traffic-originating host was a dropped packet, which only occurred (in the original version) when the queue in the router was full. Delays and variance caused by congestion were still present. Hosts that encountered a lost packet would slow down, which prevented massive overload in the core. However, hosts might suffer extremely low rates (only a few packets per second in extreme cases), inducing high variation in achieved throughput. Taming excess flows into the network was critical in preventing massive packet drops in the core, but did not remove the justification for offering enhanced services to applications that were intolerant of excess delay, jitter, and highly variable throughput.

⁴Thus, routers have been looking at parts of the packet beyond the IP header well before 1990.

4 Formalizing support for enhanced services across ISPs (1990s)

In this section we relate proposals from the 1990s to actually use existing but underutilized protocol features that supported enhanced services, and attempts to augment the Internet protocol suite with more formal and fine-grained support for real-time media services.

4.1 Proposed short-term solution: formalize use of IP Precedence field

NSFNET leadership recognized that even twenty years ago software developers were building network applications that could consume as much bandwidth as network operators could provide. In particular, real-time video and voice applications did not exhibit the same stochastic burstiness characteristics and modest resource requirements of traditional applications such as file transfer and email. In their view, 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 authors predicted that in the long run the community would redesign networks and network protocols to support complex resource reservation and accounting. In the meantime, however, they proposed an interim strategy that would shield, in a limited way, the existing environment from traffic whose behavior conflicts with the nature of resource sharing, e.g., high volume continuous traffic. Specifically, they proposed a scheme for voluntarily setting and respecting priorities (precedence levels) for Internet traffic that included incentives to limit one’s use of high precedence levels (Bohn et al., 1994). Application developers could set default priority levels for the IP precedence field, which users could modify based on their own criteria, e.g., latency-sensitivity. By default, the value is typically 0, the lowest priority. If users encountered no problems with their traffic, they would see no reason to change. Administrators of participating networks would update router software to maintain multiple queues based on the IP precedence field (well within router capabilities at that time). ISPs could selectively enable IP precedence-based queuing as they saw a need, and disable it in case of misuse.

An obvious risk of this enhanced services scheme was that users would set all their traffic to the highest priority. The third part of the proposed scheme addressed this incentive problem. The scheme included support for ISPs to set soft quotas on the total volume of traffic sent at higher (although not at lower) precedence levels. As before, there would be no quota on traffic sent at the lowest (say, lowest two) precedence levels. Quotas on higher precedence levels would be only loose incentives, monitored after the fact. If a customer exceeded its quota, its provider could assess a penalty, such as a fee on the next bill. If a customer found that it was exceeding its quota regularly (similar to some of today’s residential broadband service traffic volume quotas), it would have three choices: negotiate a higher quota, put and enforce quotas on its own customers (or internal users), or pay the assessed penalties. The proposal also clarified that quotas would be measured and enforced only at gateways between networks, and optionally. The authors imagined a variety of co-existing quota systems. They also imagined (more naively) that peer pressure and informal mechanisms (such as publishing a list of offenders) could be effective methods of enforcement. The scheme was not tied directly to billing and pricing, but over time it could establish a knowledge base from which to tie multiple service qualities more directly to monetary incentives for compliance, such as paid prioritization.

The authors acknowledged and discussed approaches to potential issues with the scheme, including incentives to discriminate against competitors’ packets by placing them on lower priority queues. They recognized the social obstacles as at least as great as the technical ones, but they

also saw the potential for a cultural shift as people saw massive video streams slowing down traditional transactions. Specifically, they considered it possible that transparency, e.g., publishing statistics on usage and congestion, could raise awareness and promote this culture shift. Yet the proposal was acutely academic in nature, and published at the height of the privatization and commercialization of the previously largely academic and government-funded NSFNET community. Nothing like this scheme was ever deployed.

4.2 Proposed long-term solution: standardizing support for enhanced services

Assuming that in the long-term, architectural support for enhanced services would be required on the Internet, starting in the 1990's there was a major push by the research and standards community to produce a fully-defined approach to adding differentiated services to the Internet. In 1992, one of us (Clark), along with two co-authors, presented a paper (Clark et al., 1992) at SIGCOMM (the major networking conference at that time) that laid out a proposal for adding what was called *Integrated Services* to the Internet. They chose this term to align the language of the Internet community with the (then) language of the telephone industry, which was arguing that the failure of the Internet to provide such services would doom the Internet in the marketplace. The concept of *Integrated Services* suggested that the network should equally be able to support the traditional best-effort services of the Internet (of that era), as well as applications such as real-time voice, what we now call Voice over IP or VoIP. The telephone industry, which held itself to high standards with respect to voice quality (cellular had not yet entered the picture), was arguing that a single integrated network would be economically beneficial, but that the Internet could not be the basis for such a system because of the under-specified best effort delivery model.

In parallel, the telephone industry was developing its alternative to Internet packet switching, which was called *cell switching*; this work grew out of early work at ATT Labs starting in the early 1970's on a system called DataKit. This technology was *virtual circuit* rather than *datagram* and offered a more predictable and stable service than the best-effort service of the Internet, a service better matched to the telephone industry's conception of how to deliver digital real-time voice. In 1992, Fraser *et al.* (Fraser, 1972) documented this alternative view of networking in the Communications of the ACM, and the telephone industry set to work to standardize and commercialize it, which led to the technology called *Asynchronous Transfer Mode*, or ATM. ATM provided a number of service classes, and thus embedded the idea of differentiated or enhanced service in its architecture.⁵ The ATM forum was established in 1991⁶, and set about standardizing what the members saw as an alternative and replacement for the Internet Protocol.

The Internet community was faced with a dilemma: they could continue to argue that the single, best-effort service would prove good enough to support all the commercially relevant services that might emerge in the context of data networking, or they could take the defensive position that argued for adding new functionality to the packet carriage layer of the Internet. To some, this was a mandatory step to preserve the competitive advantage of the Internet protocols; to others it was heresy. In a 1992 talk to the Internet Engineering Task Force, Clark predicted that the emergence of real-time applications such as interactive voice, and the proposals for an alternative network architecture to meet these requirements, would be one of the major challenges the Internet community would face in the decade (Clark, 1992). At a March 1993 IETF meeting, he laid out the challenge to the Internet community (Clark, 1993), and in June 1994, RFC 1633 was published, drawing substantially on (Clark et al., 1992). The Internet community accepted the challenge of

⁵The different service classes were called ATM Adaptation Layers or (AALs). There were AALs designed to carry fixed rate traffic like voice, and variable rate traffic like data.

⁶See <https://www.broadband-forum.org/about/forumhistory.php>

defining differentiated services for the Internet, and set up a number of working groups that produced the standards that are in use today (if not in the public Internet), which we describe next.

4.3 Standardization of enhanced service in the IETF

(Clark et al., 1992) described the concept of an Integrated Services network in terms of *services*. The example used to justify this effort was a “low latency, low jitter” service that would meet the needs of real time voice. As the IETF progressed, they considered two variants of an enhanced forwarding service: a *guaranteed service* and a *predictive service*. They intended the guaranteed service to match the service that the phone system claimed to offer: it would give a strict guarantee as to the variation in delay (jitter) for arriving packets, and thus a computable bound on the maximum end-to-end delay. In contrast, the predictive service did not provide any guarantee, and no algorithm for computing the bound. Rather, the application would have to measure the service in operation, estimate the bound and adapt to it. The predictive service was conceived as giving a low bound (better than the default best-effort service), but the application did not know in advance what the bound would be. The service was called *predictive* because the application needed to include an estimation algorithm to predict the current delay bound. The outcome of the standards process resulted in two groups of standards, one called *intserv*, which realized the guaranteed service, and the other *diffserv*, which realized the predictive service.

When the IETF began to debate these services, it heard from a constituency not usually heard at IETF meetings: the designers and marketers of advanced network applications. They contended that the predictive service was too unpredictable and they would only be interested in networks that offered the guaranteed service. The guaranteed service is much harder to implement because it requires per-flow state to be set up in every router along the path, a mechanism totally at odds with the overall philosophy of the Internet. However, there was progress on this effort. The *intserv* standards, which realize the guaranteed service, are specified in RFC 1633 (“Integrated Services in the Internet Architecture: an Overview” (Mo et al., 1994)). The core standard (RFC 2211) (“Specification of the Controlled-Load Network Element Service” (John Wroclawski, 1997)) described a complex set of mechanisms, including a flow setup phase where the sending machine specifies key parameters of the anticipated flow, and the state setup protocol RSVP (RFC2205) (Braden, Zhang, Berson, Herzog, & Jamin, 1997).

There are two potential components to a standard that describes a technique for traffic differentiation: the operation of the various network components and the resulting end-to-end service. The distinction is subtle but important. A service is what would be described to a higher level in the service stack, for example an application. “Low latency/low jitter” is the description of a service. The original specification of the ToS field in the IP header describes an abstract service; the mapping of the service to the technology of different networks was relegated to a separate RFC (RFC 795). But what a router designer needs to know is what mechanisms and functions to program into that device. The end-to-end service must result from the suitable behavior of the components along the path, and the standard must specify what those components must do in sufficient precision that the end-to-end behavior is achieved. Most of the effort in defining the controlled load service was specifying what network elements should do. At the same time, the Controlled-Load RFC (John Wroclawski, 1997) described the resulting end-to-end service, as follows: “*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.*” The word “tightly”, while not quantitative, was meant to suggest that deviations from the behavior that would result in a unloaded network would be so minor as to not impair the operation of the application.

The IETF standardization effort was reshaped by a pivotal thesis done at MIT by Abhey Parekh, the results of which are described in (Parekh & Gallager, 1993). This thesis used reasonable assumptions to derive a lower bound on the jitter that any such guaranteed service could provide, and it was a *tight* bound; he displayed a worst-case arrival pattern of packets that would generate an arrival delay at this bound. The bound was a function of the ratio of the capacity of the links in the network to the average speed of the flows across the link. Capacity can be related to cost, so an informed network engineer could compute the cost to provision a network so as to achieve a reasonable bound. For traffic with any degree of variation in sending rate, the answers were not appealing. To build a network that gave an absolute worst-case guarantee would require substantial over-provisioning (cost). This realization shifted the attention of the IETF from the guaranteed service to the predictive service, which could not give a guaranteed bound, but gave a much lower bound most of the time, with much lower complexity and better bandwidth utilization.

During the course of these debates, the IETF participants became uncertain about the desirability of what they were doing. The IETF did not have a tradition of specifying or standardizing a service, although they had essentially done so (informally) for the Controlled Load standard. The IETF normally standardized mechanisms and protocols, but not services, and many felt strongly that specification of a service was out of scope for the IETF. The IETF, after a struggle to understand how such a specification might be more precise, backed off and instead decided that for the diffserv standards, what they would standardize was the functional behavior in the router. They decided to standardize *technology* rather than *service*. To make this point clear, the IETF working groups defined a new term, “per hop behavior” or PHB, to describe what the router would do, and they redefined the ToS field in the IP header so that instead of specifying a service, as it did in the original IP specification described above, it now specified a particular PHB that the router should apply to the packet. The concept of per-hop behavior, or PHB, is discussed in RFC 3086, section 2.4.

The diffserv standards were intended to provide the predictive service described above and initially specified in RFC 2475 (“An Architecture for Differentiated Services” (Blake et al., 1998)). The architecture provides a framework to specify different PHBs; each such specification would result in a standard. With respect to defining the resulting service, RFC 2475 gives this advice: “*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.*”

The diffserv architecture redefined the IP TOS field so that different values of the field (in the diffserv context, called “Differentiated Services Code Points” or DSCPs) select different PHBs.⁷ The most commonly used PHB today within the diffserv architecture is called “An Expedited Forwarding PHB” described in RFC 2598, which maps the PHB to a service as follows: “*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.*” The specification of the Expedited Forwarding PHB, informally, is simple: a router offering EF service must specify an aggregate traffic rate R such that so long as the aggregate arrival rate does not exceed R , the router will forward EF packets with minimal delay.⁸

One key difference between the intserv and diffserv approaches is that intserv deals with individual flows, while diffserv deals with aggregates of flows. Routers aggregate all packets with the same code point value and give them the same collective treatment. By working with aggregates,

⁷The TOS field is not used for the intserv scheme. Since intserv treats each flow separately, the packets must be classified using the IP addresses and port numbers in the packet, not the TOS field.

⁸RFC 4594 (Babiarz, Chan, & Baker, 2006) defines 12 DSCP values, including the default service of best-effort forwarding.

the design avoids any need to keep track of individual flows, which leads to great simplicity compared to the intserv approach. However, for the diffserv approach to result in an overall good service, some larger mechanism must ensure that the flows marked with the specific code point do not result, in aggregate, in too large a total demand. But this mechanism is outside the scope of the specification of the PHB.

This design shift from service to PHB made the participants at the IETF more comfortable, but created a serious dilemma for the overall concept of enhanced services. What applications care about is the service they are going to receive, not the technological approach used to achieve it. The decision of the IETF to standardize on PHBs and remain silent on how these PHBs were to be used to create a service meant that an ISP considering offering such a service had a lot of work to do beyond what the IETF had done. It would be necessary to define the service that the ISP is offering, and define how the end-node requests that service for a flow of packets. Further, application designers write code based on the assumption that it can be used in any part of the Internet, which meant applications needed a standard way to request a service, which each network would translate into the best technical approaches to achieve that service. However, the IETF retreated from writing any such standard, as the effort was hard, and inextricably related to business models, which many thought were out of scope for the self-defined mission of the IETF.

4.4 Revealing moments: the greater obstacle is economics not technology

A personal anecdote illustrates some of the problems facing technical innovation in QoS. One of us (Clark) was trying to push for the deployment of QoS technology in the mid-1990s. He spoke to the major vendor of routers (Cisco), and an engineer from Cisco responded that they added features when they had a customer. So he went to the CTO of the (then) largest wide-area ISP and suggested that deploying QoS would be a good idea. The CTO had a simple answer: “No”. His elaboration was revealing: “Why should I spend a lot of money deploying QoS so Bill Gates can make a lot of money selling Internet Telephony for Windows?” In one sentence, this comment captures one of the central economic dilemmas of the Internet. If money is spent at one layer (e.g., the facilities and packet carriage layer), and is being made at another layer (the application layer), and the open interface of the packet carriage service platform makes it hard for revenue to flow across this interface, there may be a mismatch of incentives across these layers: between the Internet Service Provider and the application provider. This same CTO went on to say, more ominously: “And if I did deploy QoS, why do you think I would turn it on for anyone besides me?” Again, in one sentence, this brought into focus all the concerns we now address under the banner of *network neutrality*.

5 Non-technical barriers to enhanced services on the Internet (2000s)

The 2000s was the decade that operational and research attention shifted away from purely technical issues inherent in QoS, toward the economic and business/market drivers of and obstacles to deployment of QoS. It was also the decade that the issue of QoS hit public policy debates, and by the end of this decade, it was unclear which challenge was more daunting: the business coordination and incentive issues, or regulatory resistance.

5.1 Early 2000s: a funeral wake for QoS

The work on standardization of differentiated services resulted in many research papers on the topic, but no successful deployment on the Internet. Frustration with this failure motivated ACM

SIGCOMM to host a workshop in 2003 called “Workshop on Revisiting IP QoS: Why do we care, what have we learned? (RIPQOS)”. The organizers framed this workshop as a place for examination of reasons behind the failure. From the workshop summary, *“Large segments of the operational community simply cannot see the point of adding QoS to networks that are humming along quite nicely as they are. A broad spectrum of people can’t entirely agree on what QoS actually is.”*

One paper presented at this workshop (Bell, 2003) proposed that the obstacles to wide-area deployment of QoS were the complexity of the mechanisms, the immaturity of the code, and the lack of clear motivation for deployment. A related issue is the varying time scales at which different stakeholders evaluate the costs and benefits of proposed QoS schemes, which leads to a structural rift between, for example, the designers of protocols and QoS architectures, and operators of QoS-enabled networks. Several participants felt that the appropriate response to contention for network bandwidth is usually not QoS support but rather creating more bandwidth, but acknowledged that it is not always possible to easily upgrade wide-area network links. Internet2 engineers (a U.S. research and education network) observed, *“Neither customers nor Internet service providers need or want hard performance guarantees. Instead, each wants tools to understand and manage risk.”* (Benjamin Teitelbaum and Stanislav Shalunov, 2003)

Bruce Davie from Cisco (Davie, 2003) surprised many at the workshop when he reported that the QoS technology developed by the IETF was in use in many IP-based enterprise networks, e.g., corporate intranets. Although intra-domain (as opposed to inter-domain) deployment of QoS is not necessarily visible to researchers, he assured us that the technology worked, and it was delivering solid benefits to the networks that deployed it. The workshop summary (Armitage, 2003) emphasized deep non-technical reasons that differentiated services had not been deployed in the Internet, including complexity of coordination among multiple ISPs, and lack of an economic framework to allocate revenues across providers and management complexity. The conclusion from the organizers was: *“QoS is not dead, but as an IP QoS R&D community we need to reach out and include business, systems control, and marketing expertise in our efforts to get IP QoS meaningfully deployed and used.”*

In the next few years this same debate continued in the research literature. AT&T published a study in 2007 that, while it provided no traffic or cost data about AT&T’s network, reported traffic and user growth estimates from elsewhere to portend an imminent bandwidth and investment problem that will require new QoS-based revenue models to solve (Houle, Ramakrishnan, Sadhvani, Yuksel, & Kalyanaraman, 2007). They used modeling and simulation techniques to compare the amount of extra capacity required of a best-effort network to support the same performance requirements of a simple two-class differentiated services network, for a range of traffic scenarios. The conclusion is that a relatively small fraction of premium traffic could require considerable extra capacity to make sure those traffic needs are satisfied (consistent with (Parekh & Gallager, 1993)), and the revenue for that capacity would simply not exist under best-effort revenue models. Two years later, Andrew Odlyzko offered a more academic view: *“To evaluate claims about need for additional revenues... one needs solid cost data and a dynamic model of the industry. At the moment we do not have either one.”* (Andrew Odlyzko, 2009).

5.2 Mid 2000s: Working with industry to gain insight

One benefit of the best effort service model of the Internet is that it allows ISPs to interconnect without having to negotiate any performance parameters in order to provide that service. ISPs already have to participate in a protocol that exchanges routing information, and in doing so they of necessity reveal something about their patterns of interconnection—this revelation is considered normal. But tight specification of performance characteristics of an end-to-end service

require more complex inter-ISP protocols. Such protocols have never been proposed, let alone standardized, partly due to the considerable resistance to the inherent need for ISPs to reveal significant internal operating information in order for those protocols to work. In addition, the 1990s discussions among ISPs to develop a general approach to defining and marketing cross-ISP differentiated services raised considerable anxiety related to anti-trust concerns, since the conversation inherently touched on pricing and revenue sharing.

In 2006, the MIT Communications Futures Program convened a series of meetings involving Internet Service Providers (engineers) to try to work through the issues that would have to be solved to provide true multi-provider QoS. The attendees, all network engineers, were not willing to discuss the problem of pricing and revenue sharing, and so an alternative business framework was proposed in order to sidestep this problem. This group chose the hypothetical use case of provisioning an intranet for a global corporation, where no one ISP has the required service footprint to cover all outposts of the corporation. The business assumption was that one ISP would contract to provide the service, and subcontract with other regional ISPs to build a global footprint. In this scenario, the question of pricing seemed less vexing and could thus be set aside.

The discussions, which initially focused on the technical aspects of delivering a differentiated service (a low-jitter service for VoIP) revealed a number of complex and thorny issues. The participants were skeptical that an unquantified service could successfully be brought to market. They felt that some *quantitative* service specification would be required, such as “jitter will be no more than 20 ms. end-to-end.” But then the challenge was to sort out how ISPs along a path would cooperate to achieve a quantitatively specified outcome. One answer is to allocate a static *jitter budget* to each ISP. If there are two ISPs in the path, each would be allowed a jitter budget of 10 ms. But one of these ISPs might be well-provisioned, and could easily avoid inducing jitter without special effort, while the other ISP might be overloaded and hard-pressed to hit the 10 ms target. If all of the 20 ms budget was allocated to the second ISP, it would be a better outcome from an optimization point of view, but in this case, the second ISP would have to disclose to the first its state of under-provision, which was unacceptable. Worse, what if the actual service provided did not meet its specification? What sort of information sharing and inter-ISP measurement would be required to assign the responsibility for the fault to one of the providers? The recurring problem in the design of cross-ISP service provision was that to ensure the service was meeting its commitment, ISPs would have to share information about their internal operation that they view as proprietary. ISPs are competitors, even as they cooperate to interconnect.

The report of this effort is available at (MIT Communications Futures Program, 2006). Some of the firms that participated in the discussions intended to bring the document forward into a suitable standards forum, but in retrospect it does not appear to have had much visible impact. This tension between providing a better specified service across a multi-ISP Internet and the resistance to the sharing of operational information is an issue for other public policy challenges as well, the most acute of which today is improving the Internet’s resistance to distributed denial of service (DDoS) attacks and other security vulnerabilities.

5.3 Late 2000s: QoS becomes a public policy issue

The exercises of the 2000’s convinced many that the coordination problems and business-related barriers to QoS were so severe as to preclude any possible deployment in the public Internet, even though they were seen as delivering benefits in private IP networks. At the same time, fear of a dark side of QoS began to emerge. As we have described the use of QoS, and as the IETF standardized the mechanisms, the goal was to benefit applications that required enhanced service. But an ISP could also use QoS techniques to *disadvantage* certain packet flows, e.g., a

third-party video service competing with the carrier's own. In response to this fear, the goal (and vocabulary) of *network neutrality* emerged, with the objective of preventing potential harms from traffic discrimination.

Furthermore, while using QoS as a tool for performance improvement seems to require cooperation of every ISP in the path, any single ISP could use QoS technology to induce impairment of a flow crossing its network. So, from a technical perspective, the concern is certainly legitimate.

The concern is even stronger from an economic perspective, because any QoS capability has the potential to incentivize anti-consumer business behavior based on technologies that induce artificial scarcity, as opposed to pro-growth and innovation-driven strategies. This problem has been called the "dirt road future for the information super highway." Queue management (e.g., scheduling packets in a queue to give one set of packets a different service) only makes sense if there is a queue. If there is enough capacity in the network to carry all of the offered traffic, then queues do not form, and there is no queue to manage. So QoS tools only become relevant when the network is congested, and the QoS technology creates a possible perverse incentive. As traffic builds up on a network and periods of congestion start to occur, the network operator can either spend money to increase the total capacity, or instead sell enhanced services to its users to allow critical applications to continue to function well in the face of congestion. The question a rational ISP faces could be cast as whether to spend money on improving infrastructure or make money off services. This tension is exacerbated by the lack of competition at the access level (in both wired and wireless infrastructures), and a lack of transparency into whether the perverse incentives prevail. For this reason, discussions about acceptable uses of tools for traffic differentiation trigger discussions about measurement of performance, and disclosure and transparency with respect to practices and network conditions.

In the U.S., we know of no instances of the use of QoS technology to slow down certain flows, but ISPs have used blocking and termination of flows – extreme forms of traffic discrimination – to disadvantage certain applications. In 2005 the FCC fined Madison River, a small ISP, for blocking VoIP traffic based on its port number. In 2006, a researcher discovered that Comcast was blocking BitTorrent uploads from their customers, by injecting a TCP *reset* packet into the data stream (Federal Communications Commission, 2008). Comcast argued that they were enhancing performance for non-heavy users, but shortly thereafter changed their network throttling to be agnostic to the application.

In contrast, British Telecom used traffic shaping tools to limit but not block BitTorrent during peak hours, which did not trigger U.K. regulatory reprimand. BT argued that this step was performance-enhancing for latency- and jitter-sensitive applications such as VoIP and games.⁹

By the late 2000's, the U.S. FCC realized it needed to ramp up its understanding of the how technical and business issues connected, in order inform any future regulatory framing around QoS.¹⁰ The TCP/IP architecture had become a remarkably stable platform for communications, one of the goals of its original designers, in order to facilitate innovation above and below this network layer. Indeed, many aspects of networking technology below and above the Internet layer – bandwidth provisioning, data processing, and storage efficiencies – made phenomenal (Moore's Law) advancements in the first twenty years of the Internet's largely unregulated global deployment. The political economy of just about every aspect of the ecosystem experienced upheaval as

⁹For a detailed and thoughtful analysis of the British Telecom rate limiting situation, see (Cooper, 2013). It is interesting to contemplate what would have happened in the U.S. if Comcast had chosen to slow rather than block BitTorrent uploads, disclosed in advance that they were going to do so, and maintained that it was performance-enhancing.

¹⁰FCC Technical Advisory Workshop on Broadband Network Management, <https://www.fcc.gov/events/technical-advisory-process-workshop-broadband-network-management>. Dec. 2009. Some material from this section was presented at this workshop (http://www.caida.org/publications/presentations/2009/traffic_historical.context/)

the U.S. government continued the policy it began in the mid-1990s to privatize the Internet infrastructure and governance. The result was a evolutionary shift in capital and industry structure, and associated business models, of not only core and access infrastructure providers but also the naming (DNS) provisioning and content (CDN) provisioning industry.

As it became clear that most telephone calls would eventually use IP infrastructure, the U.S. FCC began to realize that it would have to exercise stewardship over that infrastructure, and initiated attempts (not very successfully) to learn more about its dynamics and economics.¹¹ During the Obama Administration's Broadband America program in 2009, the NTIA constructed a set of requirements for data collection associated with \$7B of "stimulus" funds allocated to broadband. The nature of these requirements¹² would have provided the U.S. government with some (limited) empirical data on demand for some Internet infrastructure resources, and possible some insight into drivers for QoS. To our knowledge the requirements were never implemented.

The last years of that decade (2008-2009) also provided examples of how advanced countries (Japan and Canada) were handling possible harms that could arise from traffic management practices, in particular *packet shaping*, which is traffic discrimination against heavy users to improve overall service for most users. That is, the term packet shaping is associated with a subset of QoS behaviors that degrade QoS of specific users in order to improve overall QoS, which triggers the question of how to constrain packet shaping behavior to minimize harm to users. In May 2008, a group of Japanese ISP industry associations published an industry consensus (not legally binding) on guidelines¹³ that described circumstances in which they considered packet shaping acceptable. They cited specific examples that emphasized the relationship between "secrecy of communications" (which has strong weight in Japan) and fairness in use under Japanese business law. The guidelines affirmed that packet shaping should be implemented only in exceptional circumstances to "facilitate necessary network management" or to "protect users". Such packet shaping must: (1) be in response to congestion of specific heavy users that is degrading or will likely degrade service of general users; (2) must be substantiated by objective data. For example, content examination, e.g., looking for copyright infringement based on payload, was not deemed reasonable because one cannot do it accurately for a single user, nor reasonably for all users. They emphasized the need for further study of the impacts of different packet shaping approaches to adapt to the increase in streaming video traffic, as well as the need for more information-sharing among players.

An international case closer to home was Canada, whose regulator (CRTC) published Internet traffic management best practices in October 2009. In spirit similar to those expressed by the U.S. FCC's, and Japan's ISPs, the CRTC's guidelines required transparency about ISP traffic management practices (Canadian Radio-television and Telecommunications Commission, 2009). They emphasized that any packet shaping must be narrowly tailored for purpose (technically

¹¹Another anecdote: An FCC engineer called one of us around 2006 to ask whether the research community knew how to define "outage" of a voice call over the Internet, since the FCC is required to track telephony outages, and clear metrics for doing so.

¹²Awardees receiving Last Mile or Middle Mile Broadband Infrastructure grants must report, for each specific BTOP project, on the following:

- i. The terms of any interconnection agreements entered into during the reporting period;
- ii. Traffic exchange relationships (e.g., peering) and terms;
- iii. Broadband equipment purchases;
- iv. Total & peak utilization of access links;
- v. Total & peak utilization on interconnection links to other networks;
- vi. IP address utilization & IPv6 implementation;
- vii. Any changes or updates to network management practices.

¹³(Japan Internet Providers Association (JAIPA), Telecommunications Carriers Association (TCA), Telecom Services Association (TELESA), and Japan Cable and Telecommunications Association (JCTA), 2008)

“efficient”), and minimize harm to consumers. They recognized the challenge of defining “reasonable” network management, but affirmed that targeting specific content or applications was not allowed, and that any techniques used should be based on transparent quantifiable data.

By 2010, the FCC formalized their own need for transparency with respect to network traffic management, and implemented disclosure requirements in their 2010 Report and Order (Federal Communications Commission, 2010). As part of this order, they also established the Open Internet Advisory Committee to facilitate evaluation of the risks and benefits associated with QoS.¹⁴

6 Evolving interconnection structure and implications for enhanced services (2010s)

In pursuit of performance and cost optimizations in the face of demanding new services and applications, the interconnection ecosystem has been profoundly transformed in this decade. In this section we describe several dimensions of this transformation, and implications for the theory and practice of regulation of enhanced services on the Internet.

6.1 Expansion of network interconnection scale and scope

An evolutionary shift in the pattern of interconnection in the Internet has codified over the last 10 years: from one where ISPs typically exchanged traffic via intermediate larger *transit providers* (the 1990s Internet equivalent of a long-distance telephone companies) to a pattern where many ISPs exchange traffic directly, over *peering links*. ISPs can peer with each other bilaterally via a private interconnection, or peer *multilaterally*, i.e., with multiple ISPs, through an Internet Exchange (IX), which is a location with a central switch that can route traffic among many connected ISPs. The term *peering* arose in the early 1990s because interconnecting large ISPs viewed each other as equal players (peers) in the market. Traditionally, such peering was revenue-neutral (settlement-free), because large ISPs perceived the arrangement as of approximately equal value.

As smaller ISPs began to peer with each other, avoiding larger transit providers, (sometimes called *donut peering*), they sent less traffic – and thus less revenue – to the transit providers, diminishing their importance the marketplace. Many transit providers from the 1990s disappeared into mergers or bankruptcies by the 2010s. The result has been called the *flattening* of the Internet topology (Dhamdhere & Dovrolis, 2010). Another market sector of the interconnection ecosystem that enabled, as well as thrived on, this peering expansion was Internet exchanges (IXes), which served as anchor points in the mesh of interconnection. A third related trend was the growing role and importance of content providers and Content Delivery Networks (CDNs) as major sources of traffic flowing into the Internet. Proliferation of IXes around the world facilitate interconnection among networks within a region, allowing traffic to flow along cheaper and lower-latency routes, which CDNs leverage to place (cache) content as close to users as practical.

We now see a new shift in the structure of ISP interconnection. This shift constitutes the rise of a new kind of hierarchy in the ecosystem, bringing fundamentally new constraints on existing players who need to manage traffic on their networks to minimize congestion. One option for large CDNs who can afford it is to install content caches adjacent to (or inside) major access networks, and maintain their own global networks to serve content to these caches. A CDNs may use these caches to deliver its own content or content on behalf their customers. When ISPs connect

¹⁴<https://www.fcc.gov/encyclopedia/open-internet-advisory-committee>. This committee met several times in 2012-2013 and published a report (available at that URL), after which the committee chair stepped down and they have not met since.

directly with content providers, we use the term *direct interconnection* to describe this mode. The economies of scale of technology of the high-bandwidth content delivery market, with its concomitant requirements for storage, bandwidth, computation and IT expertise, and legal licensing support, has led to concentration of content ownership and delivery capability among a few large content providers. By some accounts over half the traffic volume in North America now comes from just two content distributors (Youtube and Netflix). Most of this traffic flows over direct interconnection paths.

The emergence of direct interconnection has been accompanied by disputes between ISPs and content providers over the terms of interconnection. Evidence of trouble has increased dramatically in the last five years, resulting in tussles among commercial players as well as between the private sector and regulatory bodies, at the expense of users suffering degraded quality of experience. These disputes have to some extent masked the emergence of direct interconnection as a distinct phenomenon, in part because some content providers have referred to this mode of interconnection as another form of *peering*, which suggests that it should be revenue-neutral, as peering between approximately similarly-sized players had traditionally been (i.e., in the 1990s). Through heated debates in the press and to the regulator (at least in the United States), access ISPs have argued that interconnection between largely content providing networks and access networks did not fit the traditional paradigm of settlement-free peering, among other reasons because traffic exchange was heavily asymmetric. The apparent resolution of these disputes¹⁵ has been that content providers and CDNs are paying for direct interconnection, and ISPs view them as customers, to whom they make service commitments.

In the context of a debate over enhanced services on the public Internet, the important aspect of this shift toward direct interconnection is that content delivery networks are not offering mass-market services to the public. While they are a part of the global Internet, they either serve a single firm to deliver its content (e.g., Netflix or Youtube) or they serve as a third-party delivery service to firms that market content (e.g., Akamai). For such networks, there are currently no regulatory restrictions on the use of traffic differentiation; they can operate their networks as best supports the services running over them. But equally important, with direct interconnection, there is only one ISP in the path between the source of the content and the destination—the broadband access network itself. This pattern implies that for an access ISP to offer enhanced QoS to a flow from a content provider to a user via direct interconnection, no negotiation with any other ISP is necessary. The technical issues in the MIT study (Section 5.2), such as how to allocate a delay budget or isolate service impairments across several ISPs, do not arise. It is likely not a coincidence that business relationships (and associated interconnection) have evolved to enable technology to mitigate the inability to support truly interprovider QoS.

However, at least this year in the U.S., the obstacle to enhanced QoS in this specific circumstance would be regulatory, because the FCC has ruled out the possibility for the content provider, as a customer of the access ISP, to pay for any service enhancement to the end user (Federal Communications Commission, 2015). The FCC termed this *paid priority*, and concluded that it violates the neutrality that the ISPs are expected to provide, both with respect to end users and content providers (what the FCC calls *edge providers*). The FCC has stated (Federal Communications Commission, 2015, paragraph 221) that service differentiation is more likely to be acceptable if it is application-agnostic, or if it is under the control of the user, thus presumptively not harming the user or his data flow. Unfortunately, as described in Section 4.3, there is currently no framework within which a user (or application code acting on the user's behalf) could take steps to invoke or

¹⁵The details of these negotiations and resulting agreements are under non-disclosure agreements determined by the interconnecting companies.

control the use of service discrimination.¹⁶

Even without QoS technology, the use of direct interconnection itself implies a form of traffic differentiation to enhance the performance of the applications being served over that interconnection. This differentiation is not related to how packets are scheduled at a point of congestion, but the number and location of interconnection points, and the algorithm used to select the source for a particular piece of content. By adding more points of interconnection, and by picking a source close to the destination, the content provider can reduce latency, both improving end-to-end performance and bandwidth efficiency of the access ISP. Cooperation between the ISP and the content provider can potentially lead to more efficient and higher quality delivery of the service. The fact that one provider of content may negotiate more points of direct interconnection compared to their competitor has raised the issue of whether, through pricing, an access ISP can cause harms to bear on one CDN relative to another. However, this sort of discrimination has thus far been outside the scrutiny of regulators, because it seems a pure business practice (price discrimination), with no technical component such as packet scheduling.

6.2 Emergence of private IP-based platforms to support enhanced services

The traditional discourse on network neutrality (and specifically the FCC’s Open Internet principles) approaches the question of regulating QoS (or not) primarily by considering constraints on the Internet offering itself. But as the world converges on the Internet Protocols as a universal method to implement networks, these protocols are being used to deliver far more than just access to the public Internet. Recently, a few large ISPs have built private interconnected IP-based networks, using the IP technology but distinct from the public Internet. We recently used multi-sided platforms (MSP) theory to explore the range of options ISPs have to offer and interconnect enhanced services in today’s IP-based ecosystem (claffy & Clark, 2014). In particular, we identified the single and multi-firm IP platforms as alternative delivery options for consumer-facing services.

These private IP networks will allow ISPs operating them to develop new business models, support enhanced services, and vertically integrate infrastructure and applications. Such networks are a natural industry response to the need for stable network infrastructure, but they present a serious challenge to the theory and practice of regulation, particularly vivid in the now quite muddled debate over enhanced services on the Internet. First, this shift leaves the public Internet as a place for activities of lower importance and lower profitability, and perhaps starved for capital investment. Second, the new networks constitute a *shadow* activity, serving a role previously served by a regulated sector that is no longer regulated. Without regulation, these activities may carry substantial systemic risk, amplified in this case by gaps across different bodies of law, which hinder policymakers’ ability to respond to problems.

For example, a platform operator (broadband provider) may allocate a share of the IP platform as an alternate, logically separate, multi-firm platform, which is IP-based but not interconnected with the public Internet. On that share, the platform operator may allow third-party complementors to offer consumer-facing services, perhaps with desirable qualities compared to the public Internet: it may have better performance; a curated library of available applications like the Apple app store; more security; or be billed in ways that are not available over the Internet. For the context of this paper, we consider a platform with superior QoE, either through the use of

¹⁶This issue sheds light on another reason why QoS tools have been successfully deployed within private IP networks and private regions of the Internet: A private network operator is free to use classification rules to assign QoS service classes to specific services and applications using inspection of the packets, thus imposing its own policy constraints on its traffic.

QoS technology, or capacity engineering combined with limits on traffic allowed onto the platform. (If the ISP limits the platform to specific applications, it is easier to provide congestion-free provisioning than it is for the Internet, where the user can send any sort of traffic desired.)

The emergence of these alternative platforms signals an important point: ISPs would not position services on these alternative platforms unless the result was new revenue opportunities. These platforms, in one way or another, must reflect an opportunity for a facilities-owning service provider to offer services using their own facilities in ways that give them a competitive advantage over competing over-the-top services on the Internet. Under these circumstances, policymaking must consider the side-effects of constraining the Internet service relative to what is allowed on a different platform serving the same participants.

Similarly, we must consider the implications of QoE impairments on the public Internet for edge-provider innovation: do they inhibit development of classes of applications? Real harms might arise from an access provider's ability to negotiate direct interconnection on its private IP platform with arbitrary levels of QoS. One harmful trajectory is that innovative edge providers who need special QoE, and can afford it, will migrate away from the global Internet onto alternative specialized service platforms, eroding the centrality of the global Internet as the universal platform for innovation. The future might be an all-IP world, but with parallel IP platforms competing both for end user and edge-provider access. Inherent harms of such a fragmented world include inconsistent access to end users, and constant concerns about discriminatory treatment.

Any theory of regulation that tries to limit activities of the platform owner must consider the degrees of freedom that the platform owner has, i.e., unintended side-effects. These alternative IP-based platforms greatly increase those degrees of freedom. Banning discriminatory traffic treatment on the Internet may prevent certain harms, but may also prevent the public Internet from competing with alternative private IP platforms with superior QoE, an undesirable outcome. A structured focus on potential harms and benefits of specific discrimination and pricing behavior of ISPs can help frame a debate about how to maintain the Internet as a vigorous platform that can compete with alternative private platforms. Specifically, allowing ISPs to sell QoS enhancements that lead to improved QoE on the public Internet may reduce the drive to offer specialized services, preserving the Internet as the unified service platform. It is worth exploring possibilities to accommodate such enhancements using monitoring to detect potential collateral impairments.

6.3 Advancing our empirical understanding of performance impairments

There has been research on the technical aspects of QoE: how it relates to aspects of underlying service quality. More central to this paper is the policy question—what is driving observed QoE impairments today: is the underlying cause technical issues that are hard to mitigate or decisions about level of investment and exploiting deliberate scarcity? Unless tools to enhance QoS (and thus QoE) are a response to actual technical issues, they may be nothing but a way for infrastructure owners to maximize return on their capital investment.

Today, traffic from users is capped by an access link with a fixed peak speed, which is often the limiting factor for throughput, i.e., by design this segment path will congest first. For some access technologies, e.g., DSL, these limits on throughput are fundamental. When the capacity of the access link to the Internet is the limiting factor in performance, interaction among different traffic flows for a single user across that link can lead to impairments of QoE. Today, a phenomenon called *buffer bloat* (Jean-Philippe Lang, 2014) can induce jitter on access links. Some ISPs are moving to a new generation of access technology that uses service differentiation techniques to mitigate these problem. For example, the cable industry has proposed a new generation of their DOCSIS standard (DOCSIS 3.1), which will provide isolation among classes of flows for a single

Internet user so congestion from one class of flow would not trigger impairment in another user flow such as over-the-top VoIP, e.g., Skype. This transition will require upgrading cable modems and other equipment of the cable provider. Access providers use a crude form of this isolation today to separate Internet traffic from specialized services.

Thus far, using QoS tools in this context has not triggered regulatory concerns. There is a strong argument that traffic differentiation in this context will improve QoE. But there is a lurking regulatory dilemma. On the one hand, allowing QoS enhancement for the Internet service to the user can make the QoE of Internet services more competitive relative to other services offered over that link (those “private IP platform” services discussed in Section 6.2). On the other hand, perhaps the right goal for regulation should be to encourage investment in higher-capacity Internet access, which would remove some of the impairments without the need for QoS, and improve the users’ ability to run a wider range of services with good QoE. Allowing the use of traffic differentiation for Internet services over the access link could incentivize ISPs to build business models based on (and extract revenues from) scarcity rather than invest to reduce the need for QoS.

One approach to regulation of differentiated services is to link its use to a strong requirement for disclosure and justification. Since any traffic management mechanism, including QoS, can be used in ways that either benefit the user or cause harms, we have advocated explicit attention to possible harms from traffic management approaches, their causes, and means to prevent them (Clark, Bauer, & Claffy, 2014). One approach is to require ISPs to publish an analysis of harms and benefits of their proposed traffic management techniques, which explains why the benefits outweigh the harms. This approach has many challenges: ensuring such analyses are sufficiently detailed to allow independent, third-party evaluation; obtaining such evaluations; adapting to changing reasonable expectations about QoE over time; and capturing specific as well as more general societal harms, such as the effects of under-investment in capacity.

A more fundamental problem is that while impaired QoE is a meaningful indicator of harm, tools to measure and analyze QoE today are primitive. Using QoE as a basis for regulation will require research, tools and capabilities to measure, quantify, and characterize QoE, and developing metrics of service quality that better reflect our understanding of QoS and QoE for a range of applications. Indeed, there are motivations for measurement, disclosure and transparency that arise for reasons unrelated to differentiated services. Persistently congested links are an indicator that users with traffic flowing over those links may not achieve their desired service quality, and this circumstance should attract regulatory attention. Although it does not exist today, we can imagine a scenario where ISPs provide enough transparency about the state of congestion and provisioning on their networks to assure both consumers and regulators that deployment of some proposed traffic management mechanism falls into this “no collateral harm” category. One option is a policy that triggers when persistent congestion is detected on major paths in the Internet, to compel access to information from ISPs relevant to the cause of the congestion. In this context, how might a regulator deal with issues around QoS? Which actor would be responsible for detecting such congestion? In a framework where impaired QoE is considered a harm, any cause of persistent impairment would be a basis for a regulator to require disclosure of relevant information.

An interesting question is whether, by provision of enough capacity in the Internet, we could eliminate congestion (and thus the impairments caused by congestion) everywhere, or whether congestion would shift to the least recently upgraded part of the infrastructure (see Section 2.2). Recently, firms in Japan have invested to support gigabit access services, and are subsequently seeing congestion move from access networks to the backbone networks (Kamei, T, T.Inoue, & Nishatani, 2013).¹⁷ As long as users are mixing a broad range of services, it is hard to imagine

¹⁷“Since fiber to the home (FTTH) has rapidly spread all over Japan, bottlenecks in IP networks have been shifting from access

eliminating congestion altogether. Today's computers can easily exchange data among themselves at over a Gb/s. If users were offered truly unlimited service, so that they could consider streaming constant data at these speeds, Internet traffic loads might rise by orders of magnitude.

7 Lessons learned

"In such situations of moderate scarcity, however, not all people can have whatever means of communications they want. The means are rationed. The system of rationing may or may not be equitable or just. There are an infinity of ways to partition a scarce resource – egalitarian.. meritocratic.. [recognizing] privilege.. cultural values .. [rewarding] skill and motivation, as that which allows communications institutions to earn profits that depend on their efficiency."

Ithiel de Sola Pool, Technologies of Freedom, 1983, p.240.

Many parts of the Internet are well-provisioned, and usually provide quite adequate service, even for applications like VoIP, without any use of enhanced services. In other places, this presumption does not hold, including some mobile networks, residential access links, and interconnection points between ISPs. Some level of congestion is inevitable over time, and ISPs must find effective and efficient ways to handle it. The public policy dilemma is how to balance respect for the judgment of network operators in managing service qualities on their own networks with respect for consumers who do not have the capability to measure performance, nor in some cases the ability to switch to a different provider if their provider underperforms.

With an engineering perspective on this evolution, we reflect on lessons learned over the last three decades, and their implications for the future of policy. First, as with cybersecurity, obstacles to enhanced services on the Internet today are multi-dimensional with inter-related root causes that span politics, economics (business), and technology. We learned from the 1990s that ISPs lack a neutral venue in which to develop clear rules on how to define and negotiate cross-ISP services. The competitive nature of the industry conflicts with the need for providers to agree on protocols that fundamentally require sharing operational data with each other to parameterize and verify committed service qualities. Furthermore, some ISPs in the 2006 MIT inter-provider QoS discussions (Section 5.2) felt that forcing conversations about value flow and settlements would threaten the entrenched norm of revenue-neutral peering for best-effort service. The relative prevalence of paid-peering arrangements today (although all bilateral and under NDAs) may mitigate this resistance, i.e., revenue-neutral peering can no longer be assumed.

Thus, there might be several services that could be characterized using tools standardized by the IETF, but there has been no effective forum to support discussion of the necessary business coordination practices to bring enhanced services to market. Similarly, the reaction of the IETF to the standardization of differentiated service, and in particular their resistance to discussing a service specification, has created a hole in the path from research to deployment. What an ISP wants to offer, and what an application designer needs to understand, is a service specification, not the behavior of a router (the PHBs). The IETF limits itself to standardizing equipment behavior, not ISP behavior, perhaps because IETF participation is dominated by equipment suppliers to ISPs, not ISPs themselves.

Second, QoS technology can yield benefits as well as harms, thus in our view policymaking should focus on regulating harms rather than mechanisms. If deployed appropriately, QoS technology can benefit latency-sensitive traffic without impairing performance of latency-insensitive

networks to backbone networks and equipment, such as bandwidth between ISPs and capacity in IXs" (Kamei et al., 2013).

traffic. But to assure consumers and regulators that deployment of QoS technology (or other traffic management mechanisms) will fall into this “no collateral harm” category, i.e., not impair the Quality of Experience (QoE) for users, regulators may need to require transparency about the state of congestion and provisioning on networks using such mechanisms.

Third, in the meantime, we have little empirical understanding of QoE because business obstacles to enhancing end-to-end QoE have in turn constrained strategic investment in measurement capability. Using the concept of QoE impairment as a basis for regulation will require (at least) research, tools and capabilities to measure, quantify, and characterize QoE, and developing metrics of service quality that better reflect our understanding of QoS and QoE for a range of applications.

Finally, if a generalized solution is too hard to define and bring to market, more specific approaches to mitigating impairments are likely to prevail (Section 6.3). The unregulated use of IP technology to build and interconnect private network platforms as alternatives to the public Internet, and direct interconnection between content providers and access ISPs to avoid congestion on the public Internet, are both natural reactions to the demand for high service qualities and demanding workloads, but may have the unintended consequence of drawing capital investment away from the public Internet. When considering policies that try to limit the use of QoS to prevent harm to consumers, we must also take into consideration the possibility of collateral harm from the policy itself: incenting innovation to move to alternate platforms.

The co-evolution of regulatory, legal, business, and technological capabilities, all at different paces, is tightly coupled in the case of enhanced services – a quintessential interdisciplinary challenge. While barriers to the deployment of scalable interprovider QoS on today’s Internet may be insurmountable, the issues and challenges remain. If any Internet of the future is to be the platform for a full range of useful applications, either the basic transport service must be so improved that no support for differentiated services is necessary, or it will be necessary to overcome these challenges. For this reason, it is worth developing a systematic understanding to the challenge of enhanced services, starting with cataloging their successes and failures over the history of the Internet as carefully as possible.

References

- Andrew Odlyzko. (2009, March). Network Neutrality, Search Neutrality, and the Never-ending Conflict between Efficiency and Fairness in Markets. *Review of Network Economics*, 8(1). (<http://www.dtc.umn.edu/~odlyzko/doc/rne81.pdf>)
- Armitage, G. J. (2003). Revisiting IP QoS: Why Do We Care, What Have We Learned? ACM SIGCOMM 2003 RIPQOS Workshop Report. *SIGCOMM Comput. Commun. Rev.*, 33(5). doi: 10.1145/963985.963995
- Babiarz, J., Chan, K., & Baker, F. (2006). *Configuration guidelines for diffserv service classes*. Retrieved from `\url{http://www.ietf.org/rfc/rfc4594.txt}`
- Bauer, S., Clark, D., & Lehr, W. (2009). The evolution of internet congestion. In *Proceedings 2009 the 37th research conference on communication, information and internet policy (tprc)*.
- Bell, G. (2003). Failure to thrive: Qos and the culture of operational networking. In *Proceedings of the ACM SIGCOMM Workshop on Revisiting IP QoS: What Have We Learned, Why Do We Care?* doi: 10.1145/944592.944595
- Benjamin Teitelbaum and Stanislav Shalunov. (2003). What QoS Research Hasn’t Understood About Risk. *ACM SIGCOMM 2003 Workshop: Rest in Peace QoS*.
- Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., & Weiss, W. (1998). *An architecture for differentiated services*. Retrieved from <http://www.ietf.org/rfc/rfc2475.txt>

- Bohn, R., Braun, H., claffy, k., & Wolff, S. (1994, Nov). Mitigating the coming Internet crunch: multiple service levels via Precedence. In (Vol. 3, p. 335-349). (Available at <http://www.caida.org/publications/papers/1994/mcic/>)
- Braden, R., Zhang, L., Berson, S., Herzog, S., & Jamin, S. (1997). *Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification*. (<http://www.ietf.org/rfc/rfc2205.txt>)
- Canadian Radio-television and Telecommunications Commission. (2009). *Review of the Internet traffic management practices of Internet service providers*. (<http://www.crtc.gc.ca/eng/archive/2009/2009-657.htm>)
- claffy, k., & Clark, D. (2014, Sep). Platform Models for Sustainable Internet Regulation. *Journal of Information Policy*, 4, 463-488.
- Clark, D. (1992). *A cloudy crystal ball–visions of the future*. Internet Engineering Task Force.
- Clark, D. (1993). An architecture for resource management in networks. In (p. 619-623). (Available at <https://www.ietf.org/proceedings/26.pdf>)
- Clark, D., Bauer, S., & claffy, k. (2014, Sep). Approaches to transparency aimed at minimizing harm and maximizing investment. *Federal Communications Commission (FCC) Commission Documents*. (http://www.caida.org/publications/papers/2014/approaches_to_transparency_aimed/)
- Clark, D., Shenker, S., & Zhang, L. (1992). Supporting real-time applications in an integrated services packet network: Architecture and mechanism. *Proceedings of SIGCOMM 1992*.
- Cooper, A. (2013). *How regulation and competition influence discrimination in broadband traffic management: A comparative study of net neutrality in the united states and the united kingdom*. Unpublished doctoral dissertation, Oxford University. (Available at <https://www.alissacooper.com/files/Thesis.pdf>)
- Davie, B. (2003). Deployment experience with differentiated services. In *Proceedings of the acm sigcomm workshop on revisiting ip qos: What have we learned, why do we care?* doi: 10.1145/944592.944598
- Dhamdhere, A., & Dovrolis, C. (2010). The internet is flat: Modeling the transition from a transit hierarchy to a peering mesh. In *in proceedings of acm conext*.
- Federal Communications Commission. (2008). *In the Matters of Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications; Broadband Industry Practices Petition of Free Press et al. for Declaratory Ruling that Degrading an Internet Application Violates the FCCs Internet Policy Statement and Does Not Meet an Exception for Reasonable Network Management*. (FCC 08-183, https://apps.fcc.gov/edocs_public/attachmatch/FCC-08-183A1.pdf)
- Federal Communications Commission. (2010). *Report and order, preserving the open internet, fcc10-201*. (http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.Rcd.pdf)
- Federal Communications Commission. (2015). *In the matter of protecting and promoting the open internet: Report and order on remand, declaratory ruling, and order, gn docket no. 14-28*. (https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-24A1.pdf)
- Fraser, A. G. (1972). On the interface between computers and data communications systems. *Commun. ACM*, 15(7). doi: 10.1145/361454.361471
- Houle, J. D., Ramakrishnan, K. K., Sadhvani, R., Yuksel, M., & Kalyanaraman, S. (2007). The evolving internet- traffic, engineering, and roles. In *Proceedings of the Telecommunications Policy Research Conference*.
- Jacobson, V. (1988, August). Congestion avoidance and control. *SIGCOMM Comput. Commun. Rev.*, 18(4). doi: 10.1145/52325.52356

- Japan Internet Providers Association (JAIPA), Telecommunications Carriers Association (TCA), Telecom Services Association (TELESA), and Japan Cable and Telecommunications Association (JCTA). (2008). *Guideline for Packet Shaping*. (http://www.jaipa.or.jp/other/bandwidth/guidelines_e.pdf)
- Jean-Philippe Lang. (2014). *Bufferbloat*. (<http://www.bufferbloat.net>)
- John Wroclawski. (1997, September). *Specification of the Controlled-Load Network Element Service*. (<http://www.ietf.org/rfc/rfc2211.txt>)
- Jon Postel. (1981). *Service Mappings*. (Available at <http://www.rfc-archive.org/getrfc.php?rfc=795>)
- Kamei, S., T. M., T. Inoue, & Nishatani, T. (2013). *The p2p network experiment council's activities and experiments with application-layer traffic optimization (alto) in japan*. (<http://www.ietf.org/rfc/rfc6875.txt>)
- Mills, D. L. (1988). The fuzzball. In *Proc. acm sigcomm 88 symposium* (p. 115-122).
- MIT Communications Futures Program. (2006). *Inter-provider Quality of Service*. (Available at <http://cfp.mit.edu/docs/interprovider-qos-nov2006.pdf>)
- Mo, I. M., Shenker, S., Zhang, L., Estrin, D., Jamin, S., Wroclawski, J., & Herzog, S. (1994). *Integrated services in the internet architecture: an overview status of this memo*. (<http://www.ietf.org/rfc/rfc1633.txt>)
- Parekh, A. K., & Gallager, R. G. (1993, June). A generalized processor sharing approach to flow control in integrated services networks: The single-node case. *IEEE/ACM Trans. Netw.*, 1(3). doi: 10.1109/90.234856
- Postel, J. (1980, January). *Internet protocol, DARPA Internet Program Protocol Specification, RFC760*. (<http://www.ietf.org/rfc/rfc791.txt>)
- Postel, J. (1981). *Internet protocol, network working group request for comments 791*. (<http://www.ietf.org/rfc/rfc791.txt>)