BITAG

Broadband Internet Technical Advisory Group

Differentiated Treatment of Internet Traffic

A BROADBAND INTERNET TECHNICAL ADVISORY GROUP
TECHNICAL WORKING GROUP REPORT

A Uniform Agreement Report

Issued:

October 2015

Copyright / Legal Notice

This document and the information contained herein is provided on an "AS IS" basis and BITAG AND THE CONTRIBUTORS TO THIS REPORT MAKE NO (AND HEREBY EXPRESSLY DISCLAIM ANY) WARRANTIES (EXPRESS, IMPLIED OR OTHERWISE), INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR A PARTICULAR PURPOSE, OR TITLE, RELATED TO THIS REPORT, AND THE ENTIRE RISK OF RELYING UPON THIS REPORT OR IMPLEMENTING OR USING THE TECHNOLOGY DESCRIBED IN THIS REPORT IS ASSUMED BY THE USER OR IMPLEMENTER.

The information contained in this Report was made available from contributions from various sources, including members of Broadband Internet Technical Advisory Group, Inc.'s Technical Working Group and others. Broadband Internet Technical Advisory Group, Inc. takes no position regarding the validity or scope of any intellectual property rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this Report or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights.

About the BITAG

The Broadband Internet Technical Advisory Group (BITAG) is a non-profit, multi-stakeholder organization focused on bringing together engineers and technologists in a Technical Working Group (TWG) to develop consensus on broadband network management practices and other related technical issues that can affect users' Internet experience, including the impact to and from applications, content and devices that utilize the Internet.

The BITAG's mission includes: (a) educating policymakers on such technical issues; (b) addressing specific technical matters in an effort to minimize related policy disputes; and (c) serving as a sounding board for new ideas and network management practices. Specific TWG functions also may include: (i) identifying "best practices" by broadband providers and other entities; (ii) interpreting and applying "safe harbor" practices; (iii) otherwise providing technical guidance to industry and to the public; and/or (iv) issuing advisory opinions on the technical issues germane to the TWG's mission that may underlie disputes concerning broadband network management practices.

The BITAG Technical Working Group and its individual Committees make decisions through a consensus process, with the corresponding levels of agreement represented on the cover of each report. Each TWG Representative works towards achieving consensus around recommendations their respective organizations support, although even at the highest level of agreement, BITAG consensus does not require that all TWG member organizations agree with each and every sentence of a document. The Chair of each TWG Committee determines if consensus has been reached. In the case there is disagreement within a Committee as to whether there is consensus, BITAG has a voting process with which various levels of agreement may be more formally achieved and indicated. For more information please see the BITAG Technical Working Group Manual, available on the BITAG website at www.bitag.org.

BITAG TWG reports focus primarily on technical issues, especially those with the potential to be construed as anti-competitive, discriminatory, or otherwise motivated by non-technical factors. While the reports may touch on a broad range of questions associated with a particular network management practice, the reports are not intended to address or analyze in a comprehensive fashion the economic, legal, regulatory or public policy issues that the practice may raise. BITAG welcomes public comment. Please feel free to submit comments in writing via email at comments@bitag.org.

Executive Summary

The Internet is composed of interconnected networks, each having its own architecture and technical characteristics. The data transmitted across these networks is formatted as packets containing information payloads encapsulated within one or more headers, which in turn provide the information needed by networks to deliver the packets to their destinations. As these packets travel across networks, they contend with other packets for network resources. Contention can occur at any point where two or more packets can compete for a resource at the same time. The simplest way to handle such requests would be on a first come, first served basis (also known as First In First Out, or FIFO). In practice, however, network operators make many exceptions to FIFO, using the packet header information to classify packets into flows and treating those flows differently, for example rearranging the order or the timing with which packets are sent, or sending them along different network paths.

Differentiated treatment of Internet Access Service traffic has been a subject of debate and regulatory scrutiny. In February 2015, the Federal Communications Commission (FCC) adopted Open Internet rules that address paid prioritization as well as other topics [1]. This report touches on a broad range of questions associated with differentiation, but is not intended to address or analyze the economic, legal, regulatory, or public policy issues that the differentiated treatment of Internet access service traffic may raise, focusing instead on the technical issues.

The ability to treat traffic differentially has been built into Internet protocols from the beginning. The specifications for both IPv4 and IPv6 have included fields to support traffic differentiation since their inception (initially IPv4's Type of Service or ToS field) to indicate to routers the quality of service desired, in terms of queuing precedence and routing parameters around delay, rate, and reliability. This was changed to more generic service descriptions with the definition of the Differentiated Services Field, and implemented in IPv4 and IPv6. Notably, traffic differentiation in this sense has not been implemented in multi-provider environments, although it is extensively used within specific networks. End to end deployment would require the harmonization and cooperation of a large number, if not all, of the relevant network operators.

In its broadest sense, traffic differentiation includes any technique that classifies and applies potentially different treatment to two or more traffic flows contending for resources on a network (a flow being a group of packets that share a common set of properties). Differentiated treatment of network traffic is a two-part process: (1) traffic is classified into traffic streams, and (2) a prescribed set of actions is applied to each stream. This treatment may determine the order in which routers and switches send packets from different flows across the link, the rate of transmission of a given flow, or even whether certain packets are sent at all.

While the techniques used for traffic differentiation overlap with those used to manage congestion, differentiation has a broader purpose that includes meeting service level agreement (SLA) guarantees and selecting paths for traffic from different applications, among other things. Differentiated treatment of traffic can also contribute both to the efficiency of a network and to the predictability of the manner in which network resources are shared.

Differentiation can be complex, and a common vocabulary is key. This report uses the terms “differentiated treatment” or “differentiation,” as opposed to “prioritization” when referring to the full range of treatments that may be applied to traffic flows. The technical definition of “prioritization” is narrow and generally applies only to certain scheduling, dropping, and marking techniques. This report uses “differentiation” in a much broader sense, including most of the ways in which packets may be treated differently from each other while en route to their respective destinations across one or more networks. The scope of differentiation in this report encompasses the classic techniques of scheduling, shaping and queue management by which packets are processed at a network node, and also includes the techniques by which traffic flows are segregated or forwarded onto different physical or logical network paths where they may encounter greater or lesser propagation delays or contention for resources.

This report addresses differentiation applied to traffic on Internet access services, as well as the impacts to Internet access services when differentiation is applied to other traffic carried over the same network. Traffic for mass-market Internet access services is often carried over a common infrastructure with traffic associated with other IP services, as well as the network management traffic used to control devices and report status from them. Since differential treatment of other network traffic has the potential to affect the performance of Internet access services, it is considered here.

The subjective experience perceived by the user of a networked application is known as Quality of Experience, or QoE, and the factors that contribute to QoE vary significantly from one application to the next. In contrast, Quality of Service, or QoS, describes the performance of a network service using objective metrics such as throughput, delay, delay variation, and loss. The relationship between QoS and QoE is highly dependent on the type of application, but variations in QoS have been mapped to corresponding variations in QoE for a number of applications. It is possible to use knowledge about the relationships between network performance parameters and their effects on QoE to attempt to optimize the performance of network flows for their intended applications. Differentiation is often also used to address impairments to QoS.

Broadband networks use different network architectures and access technologies. Several of these network architectures have developed to take advantage of existing access infrastructure that was originally deployed for other services – for example, telephone service over twisted copper pairs or video over coaxial cable. Other networks were developed to meet specific needs, such as for mobility or for access in remote rural areas. In many cases, differences in network design can be traced to the different characteristics of the access technology used. Access technologies can require different approaches to differentiation of traffic.

Observations. From the analysis made in this report and the combined experience of its members when it comes to the differentiated treatment of Internet traffic, the BITAG Technical Working Group makes the following *observations*:

- **TCP causes recurring momentary congestion.**

  When TCP transfers a large file, such as video content or a large web page, it practically guarantees that it will create recurring momentary congestion at some point in its network path. This effect exists by design, and it cannot necessarily be eliminated by increasing capacity. Given the same traffic load, however, the severity of the momentary congestion should decrease with increased capacity.

- **A nominal level of packet discard is normal.**

  Packet discard occurs by design in the Internet. Protocols such as TCP use packet discard as a means of detecting congestion, responding by reducing the amount of data outstanding and with it self-induced congestion on the transmission path. Rather than being an impairment, packet discard serves as an important signaling mechanism that keeps congestion in check.

- **The absence of differentiation does not imply comparable behavior among applications.**

  In the absence of differentiation, the underlying protocols used on the Internet do not necessarily give each application comparable bandwidth. For example:

  - TCP tends to share available capacity (although not necessarily equally) between competing connections. However, some applications use many connections at once while other applications only use one connection.
  - Some applications using RTP/UDP or other transport protocols balance transmission rate against experienced loss and latency, reducing the capacity available to competing applications.

- **Differentiated treatment can produce a net improvement in Quality of Experience (QoE).**

  When differentiated treatment is applied with an awareness of the requirements for different types of traffic, it becomes possible to create a benefit without an offsetting loss. For example, some differentiation techniques improve the performance or quality of experience (QoE) for particular applications or classes of applications without negatively impacting the QoE for other applications or classes of applications. The use and development of these techniques has value.

- **Access technologies differ in their capabilities and characteristics.**
  Specific architectures and access technologies have unique characteristics which are addressed using different techniques for differentiated treatment.

- **Security of traffic has at times been downgraded to facilitate differentiation techniques.**
  Encrypted traffic is on the rise and it has implications for current differentiation techniques. In response to this increase, some satellite and in-flight network operators have deployed differentiation mechanisms that downgrade security properties of some connections to accomplish differentiation. The resulting risks to the security and privacy of end users can be significant, and differentiation via observable information such as ports and traffic heuristics is more compatible with security.

<u>Recommendations.</u> The BITAG Technical Working Group also has the following *recommendations*:

- **Network operators should disclose information on differential treatment of traffic.**
  In previous reports, BITAG has recommended transparency with respect to a number of aspects of network management. BITAG continues to recommend transparency when it comes to the practices used to implement the differential treatment of Internet traffic.

  Specifically with respect to consumer-facing services such as mass-market Internet access, network operators should disclose the use of traffic differentiation practices that impact an end user's Internet access service. The disclosure should be readily accessible to the public (e.g. via a webpage) and describe the practice with its impact to end users and expected benefits in terms meaningful to end users. The disclosure should include any differentiation amongst Internet traffic and should disclose the extent and manner in which other services offered over the same end user access facilities (for example video services) may affect the performance of the Internet access service.

- **Network operators and ASPs should be encouraged to implement efficient and adaptive network resource management practices.**
  In a previous report BITAG recommended that ASPs and CDNs implement efficient and adaptive network resource management practices; we reiterate that recommendation here, extending it to network operators. Examples of such practices might target the minimization of latency and variation in latency induced in network equipment, ensuring sufficient bandwidth for expected

traffic loads, and the use of queue management techniques to manage resource contention issues.

- **Quality of Service metrics should be interpreted in the context of Quality of Experience.**

  Common Quality of Service metrics, often included in commercial service level agreements, include capacity, delay, delay variation, and loss rate, among other things. From the viewpoint of the end user application, these metrics trade off against each other and must be considered in the context of Quality of Experience. For example, since TCP Congestion Control and adaptive codecs depend on loss to infer network behavior, actively trying to reduce loss to zero leads to unintended consequences. On the other hand, non-negligible loss rates often directly reduce the user's Quality of Experience. Hence, such metrics should be interpreted in the context of improving user experience.

- **Network operators should not downgrade, interfere with, or block user-selected security in order to apply differentiated treatment.**

  Network operators should refrain from preventing users from applying over-the-top encryption or other security mechanisms without user knowledge and consent. Networks should not interfere with, modify, or drop security parameters requested by an endpoint to apply differentiated treatment. Given the potential for possible exposure of sensitive, confidential, and proprietary information, prior notice should be given to end users of traffic differentiation features that affect security properties transmitted by endpoints.

Table of Contents

1 Introduction

The Internet is composed of interconnected networks, each having its own architecture and technical characteristics. The data transmitted across these networks is formatted into packets, which are composed of information payloads encapsulated within one or more headers, which in turn provide the information needed by networks to deliver the packets to their destinations. As these packets travel across networks, they contend with other packets for network resources. Contention can occur at any point where two or more packets can compete for a resource at the same time – for example, at a network switch where traffic from multiple input ports is forwarded to a common output port. The simplest way to handle such requests would be on a first come, first served basis (also known as First In First Out, or FIFO). In practice, however, network operators make many exceptions to FIFO, using the packet header information to classify packets into flows and treating those flows differently, for example rearranging the order and/or the timing with which packets are sent, or sending them along different network paths. Such "differentiated treatment" of network traffic is the subject of this report.

Differentiated treatment of Internet Access Service traffic has been a subject of debate and regulatory scrutiny. In February 2015, the Federal Communications Commission (FCC) adopted Open Internet rules that address paid prioritization as well as other topics [1]. This report touches on a broad range of questions associated with differentiation, but is not intended to address or analyze the economic, legal, regulatory, or public policy issues that the differentiated treatment of Internet access service traffic may raise, focusing instead on the technical issues.

Differentiation can be a complex topic, and a common vocabulary is important. This report uses the terms "differentiated treatment" or "differentiation" as opposed to "prioritization" when referring to the full range of treatments that may be applied to traffic flows. "Prioritization" has a narrower technical definition that applies only to certain scheduling, dropping, and marking techniques. This report uses "differentiation" in a broader sense, including most of the ways in which packets may be treated differently from each other while en route to their respective destinations across one or more networks. The scope of differentiation in this report encompasses the classic techniques of scheduling, shaping, and queue management by which packets are processed at a network node, and also includes the techniques by which traffic flows are segregated and/or forwarded onto different physical or logical network paths where they may encounter greater or lesser propagation delays or contention for resources.

This report addresses differentiation applied to traffic on Internet access services,[1] as well as the impacts to Internet access services when differentiation is applied to other traffic carried over the same network. Traffic for mass-market Internet access services is often

[1] These services are largely analogous to Broadband Internet Access Services (BIAS) in the recent Open Internet Report and Order published by the FCC [1]. The FCC Order uses the term "non-BIAS data services" to refer to services that share "last mile" connections with BIAS yet are not BIAS. Note that although the FCC emphasizes last mile connections at times in its Report and Order, this report addresses differentiated treatment at any point in the network.

carried over a common infrastructure with traffic associated with other IP services, as well as the network management traffic used to control devices and report status from them. Since differential treatment of other network traffic has the potential to affect the performance of Internet access services, it is considered here.

The report is organized as follows: Section 2 gives an overview of how and why differentiated treatment of traffic exists in current networks, reviews the history of differentiation, and discusses the potential impacts of traffic differentiation in terms of both Quality of Service (QoS) and Quality of Experience (QoE). Section 3 addresses the techniques used to differentiate traffic, and Section 4 shows how these techniques are applied in different access network architectures. Section 5 illustrates the impact of these techniques with a number of examples of network practices associated with traffic differentiation. Section 6 provides a number of observations, and Section 7 provides recommendations. In addition, the report includes references and a glossary, as well as an appendix listing relevant standards.

2 Differentiation in IP networks

In its broadest sense, traffic differentiation includes any technique that classifies and applies potentially different treatment to two or more traffic flows (groups of packets that share common properties [2]) contending for resources on a network. Differentiated treatment of network traffic is a two-part process: (1) traffic is classified into traffic streams, and (2) a prescribed set of actions is applied to each stream. This treatment may determine the order in which routers and switches send packets from different flows across the link, the rate of transmission of a given flow, or even whether certain packets are sent at all.

While the techniques used for traffic differentiation overlap with those used to manage congestion [3], differentiation has a broader purpose than just congestion management. Differentiation is used to deal with impairments due to congestion. It is also used to ensure that service level agreement (SLA) guarantees are met. Differentiation can be used to schedule packets or to select a path that minimizes delay for delay-sensitive applications, select a path that experiences low corruption of bits for loss-sensitive applications, or even select a path that keeps the traffic on the network of the provider offering a guaranteed SLA.

Differentiated treatment of traffic is practiced in nearly every provider network. Some of the many reasons for traffic differentiation are:

- Network operators routinely use shaping to limit each customer’s traffic to their purchased rate, and use scheduling to manage traffic from different customers at times of congestion. Since Internet access services are typically offered at a variety of rates, both shaping and scheduling may use different parameters for different customers.

- Many networks carry a mix of traffic, including customer traffic and traffic whose purpose is solely network control or management, such as routing protocol messages or device configuration updates. Network operators typically prioritize control and management traffic above other traffic to ensure timely delivery, which in some cases can be necessary for the stability of the network.

- Many network operators offer multiple IP-based services to consumers over a common access link. A typical combination of services includes Internet access, IPTV, and carrier grade voice, frequently referred to as "triple play." Traffic for the voice and video services may be differentiated to ensure that each service is delivered to the customer with its required Quality of Experience (QoE, see Section 2.3).
- Many networks carry traffic for a variety of business services in addition to the consumer services noted above. Business connectivity services, such as Carrier Ethernet, are typically sold with an associated service level agreement (SLA) that specifies the service requirements for some or all of the traffic carried by the service. Traffic for these services is differentiated to enable its delivery within the QoS parameters set by the SLA.
- Mobile access networks have to deal with constantly changing capacity and congestion conditions based on the mobility of their customers. These networks differentiate services to ensure proper balance of signaling, voice and data to ensure the proper experience for each aspect of the network.

2.1 History and evolution of differentiation methods

The ability to treat traffic differentially has been built into Internet protocols from the beginning. The IPv4 protocol, first specified for the Defense Advanced Research Projects Agency (DARPA) in 1981 [4], is still the dominant protocol in the Internet today – although there is an increasing movement to IPv6 as the number of assignable IPv4 addresses dwindles [5,6,7]. Every packet sent across the Internet uses either IPv4 or IPv6 to provide end-to-end addressing and other information. The specifications for both IPv4 and IPv6 have included fields to support traffic differentiation since their inception, providing a set of control bits in the Internet Protocol header (initially the Type of Service or ToS field) to indicate to systems en route, including routers, middleboxes, and the destination host, the quality of service desired [8]. Originally, the ToS field was described in terms of precedence, latency, throughput, and reliability requirements [8]. With the definition of the differentiated services architecture [9], the ToS field was redefined to include a differentiated services codepoint (DSCP) whose values were defined in terms of the localized differentiated treatment (or "per-hop behavior") requested of routers in the network path.

One example of traffic differentiation from the early days of the Internet (the 1980s) concerned interactive traffic from remote login sessions. This traffic was given priority over other traffic to improve the perceived performance of the interactive session [10]. It is worth noting that the original backbone of the Internet (the ARPAnet) had long distance links that ran at 56 kb/s, and that at the time, persistent congestion was widespread [11,12].

Early networking standards recommended that applications on host computers sending traffic should specify the correct setting of the ToS field, and that routers should respect this setting (either by processing it or by passing it to the data link control layer) [13, 14, 8]. This design put the sending host, and not the router (or its operator) in control of selecting what sort of traffic treatment the packets would receive. The Differentiated Services Architecture makes the same assumption, although it allows the network to override the setting [9].

The inclusion of the ToS control field in the IP header allowed that field to be acted on by contemporary routers, which (at least in theory) only examined fields within that header. Newer routers, however, regularly look at port numbers in the Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) header to classify, and sometimes differentiate between traffic from different applications (frequently due to security policies).[2] However, there is no guarantee that applications will use port numbers in the expected way, or indeed that port fields will always be present in the header. In the increasingly encrypted Internet, even the basic assumption of the visibility of those fields may be suspect [19]. As a result, routers that look into higher layer headers are taking advantage of a common convention, but not a feature assured by the architecture.

While they have been used in specific networks, such as US Navy SPAWAR and in individual public and private networks, the IPv4 ToS field and the IPv4/IPv6 DSCP, have not been deployed or used across network interconnects for both engineering and economic reasons [20], and would require the harmonization and cooperation of the relevant network operators. Proposals to that end are being discussed in the IETF, however [21].

2.2 Differentiated treatment and allocation of resources

Differentiated treatment of traffic can affect the manner in which network resources are shared. Different methods of sharing resources might affect:

- The amount of time that each sender is sending,
- The amount of data that each sender sends (in terms of packets or bits), or
- The average rate of each session.

It is also possible to share resources at different levels of aggregation, including for example:

- Individual flows, for example as defined by their 5-tuple (see Section 3.2),
- All flows associated with the same service and user, or
- All flows associated with the same user regardless of service.

[2] The Internet Corporation for Assigned Names and Numbers (ICANN), under contract by the National Telecommunications & Information Administration (NTIA) to perform the Internet Assigned Numbers Authority (IANA) functions, maintains a registry of “well known” port numbers associated with different applications [15]. Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) are two of the core protocols in the Internet Protocol suite, with TCP being the protocol that many major Internet applications rely on [16,17,18], see also Sec. 3.1 and 5.2.

The literature includes significant discussion of network resource allocation among applications and transport protocols [22,23]. TCP and similar transport protocols tend to share available capacity (although not necessarily equally) between competing connections. In the absence of scheduling or other differentiation techniques, however, this sharing of capacity may be skewed by the applications sending traffic over TCP. For example, even if we assume that each TCP connection receives a roughly equal share of capacity, an application that opens many connections will receive much more capacity than an application opening a single connection. In addition, not all transports or applications share capacity in the same way; for example, voice and video applications using RTP/UDP transport will often balance transmission rate against experienced loss and latency, reducing the capacity available to competing applications.

2.3 Quality of Experience (QoE) and Quality of Service (QoS)

Customers of Internet access services use those services for a broad range of applications. However, customers rarely notice the underlying transfer of data across the network that enables these activities, except when a performance issue causes a perceptible reduction in quality in the application they are using. The subjective experience perceived by someone using an application is known as Quality of Experience, or QoE [24]. QoE has a number of contributing elements, including network performance, the platform used by the customer, and the application itself.

The subjective factors that contribute to QoE vary significantly from one application to the next. For voice communications, contributing factors include whether the received voice is garbled or missing in places, whether the speaker's echo is audible, and how much delay is introduced by the combination of communications channel, application and equipment [25]. For streaming video, factors include whether blocking or other artifacts corrupt the received video, whether it freezes or stutters, the amount of time before playback begins and the quality of the source content including its encoding algorithm and encoded bit rate [26]. The QoE associated with a web browsing application may be affected by the speed at which pages load and whether all of the content is received correctly. Some applications may exhibit multiple behaviors simultaneously (e.g. video inside a browsing session) and this can make QoE difficult to assess. QoE can be measured and quantified for a given application and set of conditions (for instance, by use of Mean Opinion Scores [27]), but the specific measurement methods vary from one application to another [28, 29].

A related concept that is frequently confused with QoE is Quality of Service, or QoS.[3] While QoE describes subjective user experience, QoS describes the performance of a network service using objective metrics such as throughput, delay, delay variation, and loss [33]. Variations in QoS have been mapped to corresponding variations in QoE for a number of applications [34, 35]. The contributions made to QoE by specific network performance parameters are highly dependent on the type of application. For example, two-way

[3] Contributing to the confusion is the fact that before the term QoE came into general use in a networking context, QoS (and variations such as PQoS) was commonly used to describe both network performance and subjective user experience [29,30,31,32].

interactive voice communication is sensitive to round trip delay and delay variation, but the throughput required is orders of magnitude lower than typical broadband service rates. Conversely, video streaming QoE is sensitive to packet loss and variations in throughput below a threshold rate, but less sensitive to delay variation or round trip delay. Video streaming also accounts for the majority of current throughput requirements on a typical broadband service.

Figure 1 shows a generalized shape for the relationship between QoS impairment and the QoE for an application [36]. The curve shows three regions of varying degrees of QoS impairment. The first region (QoS impairment less than x1) denotes a range in which the impairment has no discernible effect on QoE. The size of region 1, which may be zero in some cases and significant in others, depends on the impairment and the application. For example, in the absence of other impairments, a one-way delay of up to 150 msec from the speaker's mouth to the listener's ear has no discernible effect on interactive voice QoE [37]. In the second region (QoS impairment between x1 and x2), increasing QoS impairment corresponds to decreasing QoE. In the third region (QoS impairment greater than x2) the QoE is so poor that most users consider it unacceptable and may stop using the application. The value of x2 is dependent on the impairment, the application, and user tolerance.

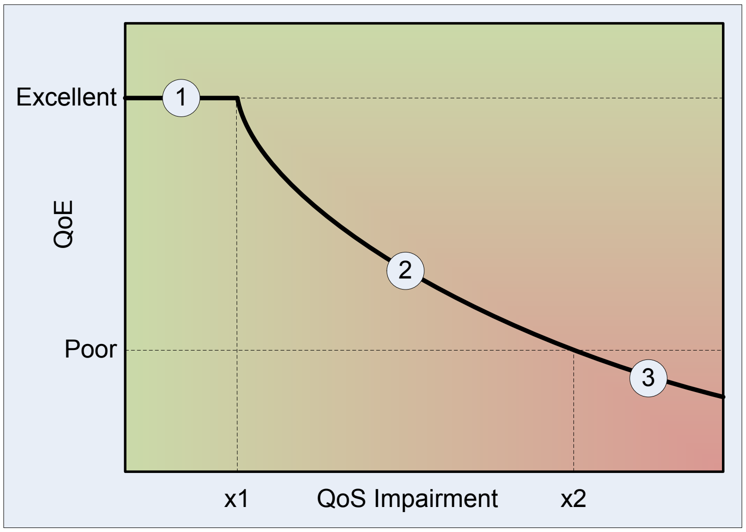

Figure 1: General shape of the mapping curve between QoS and QoE [36]

It is possible to use knowledge about the relationships between network performance parameters and their effects on QoE to optimize the performance of network flows for their intended applications. In many cases, differentiating between flows can improve the QoE for some applications without materially degrading the QoE for other applications. Section 5 describes a number of examples of this type of optimization.

2.4 Contributors to QoS

The network contributors to QoS discussed below have a direct bearing on differentiated treatment of traffic. Other contributors, such as packet corruption or reordering, affect network performance but only have an indirect bearing on differentiation (e.g., reordered packets increasing delay).

2.4.1 Delay

Delays across the network derive from four basic components [38]:

- *Serialization delay*: the amount of time it takes to send a packet on a communications link. Serialization delay for a packet is calculated as the length of the packet divided by the rate of the link.
- *Processing delay:* the amount of time required to calculate how to forward the packet within a router or switch.
- *Propagation delay*: the amount of time it takes packets to travel the physical network path from source to destination. This is calculated as the distance traveled divided by the speed of propagation in the transmission medium.
- *Queuing delay:* occurs when packets must wait in a buffer before being transmitted.

Serialization and processing delays do not usually change significantly due to differentiation, although they are affected by the number of hops in a network path. Propagation delay can be affected by the choice of network path, and queuing delay can be affected significantly by differentiated treatment.

Delay is characterized by median or average latency and by variation in latency (also known as "jitter"). Queuing delay is usually the largest contributor to jitter. Jitter may be increased by techniques in which a technology queues several packets (introducing momentary delay) and then sends them in a burst. Examples exist in IEEE 802.11 WLAN and DOCSIS technologies.

2.4.2 Packet discard and Transmission Control Protocol (TCP) Congestion Control

When a packet arrives at a buffer, it may be processed immediately (if the resource fed by the buffer is available), buffered for later processing, or discarded. In the simplest case, packets are discarded when the buffer has no room for new traffic ("tail drop"). More sophisticated algorithms discard packets before the buffer is full to signal congestion to TCP [39]. Some algorithms also differentiate based on the drop precedence marked in packet headers (Section 3.2), discarding some packets more aggressively than others [40].

Packet discard is fundamental to the design of the protocols supporting the Internet today. TCP probes for available capacity by continuously increasing the amount of data placed on the network until it detects lost packets, which it interprets as congestion.[4] Upon detecting congestion the protocol decreases the amount of outstanding data before it once again

[4] Nodes in the network path can also signal congestion to TCP explicitly using bits in the TCP header [41] – however, packet discard is used to implicitly signal congestion in most cases.

starts probing for capacity. Rather than necessarily being an impairment, packet discard can serve as an important signaling mechanism that keeps congestion in check.

2.4.3 Fragmentation

Packet fragmentation in IPv4 occurs when a host attempts to send packets that are larger than the maximum packet size (known as Maximum Transmission Unit, or MTU) that can be processed by a network segment [42]. When a packet larger than the segment's MTU arrives, it may be subdivided into two or more packets before being forwarded, or dropped in cases when fragmentation is not desired.

Fragmentation can cause additional latency (as a receiver must buffer and reassemble packets), additional CPU utilization, and additional memory usage or packet loss when buffers are exhausted. When packets are fragmented, only the initial fragment has the TCP/UDP port number, and differentiated treatment that depends on the port number for classification (Section 3.2) may not work [42].

3 Differentiation techniques

Differentiated treatment of network traffic is a two-part process: (1) traffic is classified into traffic streams, and (2) a prescribed set of actions is applied to each stream. The classification rules and the action rules are combined to form service policy rules [43].

3.1 Layered network model

To understand how differentiation is performed, it helps to first have a background understanding of the layered network model used to describe how networks operate, as some differentiation techniques are implemented at different layers. While a number of different models exist, each uses the concept of "layers" to abstract away the internal structure and technology of a network, as well as group common functions together. This report uses the 5-layer TCP/IP model [44] as shown in Figure 2 below.

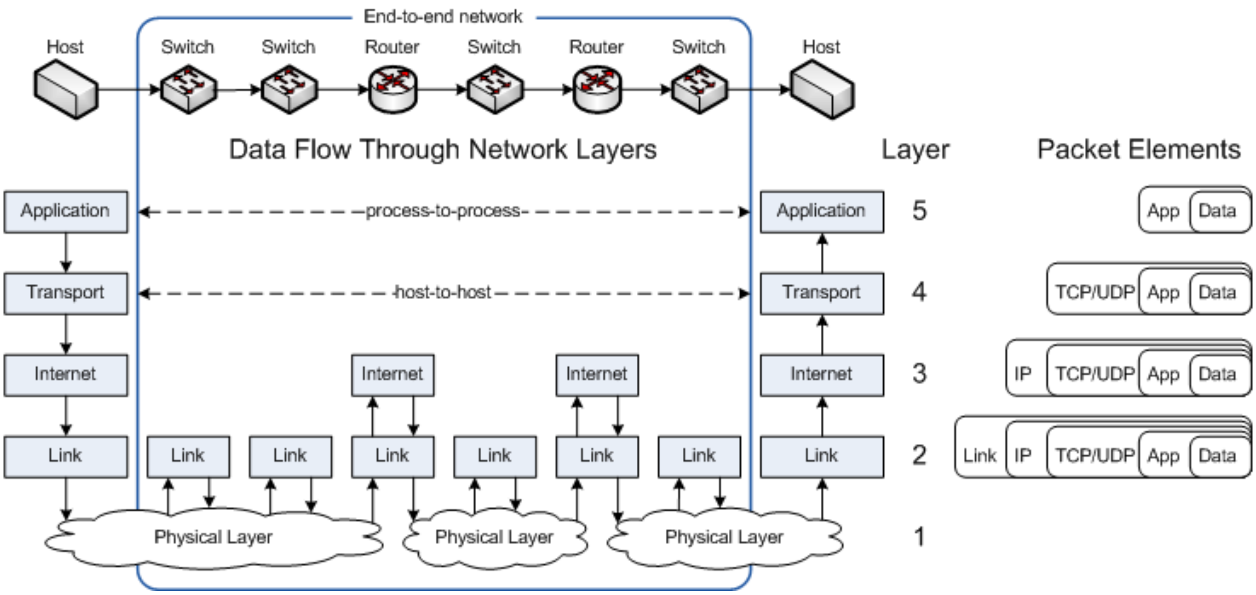

Figure 2: 5-layer model of the Internet and IP Packet Elements [44]

Layer 1 – Physical Layer. The physical layer conveys the bit stream on the transmission media (fiber, copper, radio wave) at the electrical and mechanical level – converted to electrical impulses, light waves, or radio signals.

Layer 2 – Data Link Layer. The data link layer encompasses the technologies and protocols used to send traffic across a sub-network, or "link."[5] At the lowest level, link layer protocols manage access to the physical media and encode traffic into frames such as Ethernet frames, Frame Relay frames, or ATM cells [44]. These protocols are sometimes designed in conjunction with a specific physical layer, such as IEEE 802.11 or DOCSIS. Link layer protocols also support classification and marking to facilitate scheduling, shaping, and other differentiation functions that may occur in the nodes that perform switching within a link.

In some network architectures, multiple protocols may operate between the physical layer and the Internet layer. Examples of this are MPLS and the use of Ethernet MAC over ATM and PPPoE over Ethernet MAC [46,47,48].

Layer 3 – Internet Layer. The Internet layer delivers packets across the end-to-end network from source endpoint to destination endpoint. The Internet Protocol (IPv4 or IPv6) at this layer supports end-to-end addressing, as well as classification and marking. Routers perform scheduling, shaping, and policing as well as routing at this layer.

Layer 4 – Transport Layer. At the transport layer, the Internet transport protocol (typically TCP or UDP) delivers a flow of packets across the network with characteristics determined by the protocol used. TCP provides end-to-end flow identification, packet

[5] The IETF uses the term "link" for a sub-network in which traffic flows between two or more IP interfaces [45]. In this context, a "link" can include multiple physical segments as well as switches where traffic may be differentiated.

sequencing, error recovery, and flow control for reliable data transfer [17]. UDP provides flow identification and error correction [18].

Layer 5 – Application Layer. The application layer represents all the functions that are performed by the application endpoints (e.g. client and server) to manage application-to-application level communication, such as controlling the transfer of a large file. One example is the hypertext transfer protocol (HTTP), the protocol for the transmission of web pages [49]. Other examples include the file transfer protocol (FTP), the Dynamic Host Configuration Protocol (DHCP), and the email protocols POP, IMAP and SMTP [50,51,52,53,54].

3.2 Classification

Traffic classification can be performed in most layers of the network model, though the available classification elements differ at each layer. Classification below the Application layer (Layer 5) uses pattern analysis on elements within packet headers. Classification at the Application layer may use pattern analysis or other, more complex techniques.

- **Layer 2.** Traffic is often classified at Layer 2 in converged networks that deliver multiple services such as high-speed Internet access service and carrier grade voice. Classification can be performed using any element in the Layer 2 frame headers such as MAC address, virtual LAN (VLAN) tags, and multiprotocol label switching (MPLS) labels.

- **Layer 3 and 4** classification is performed on the elements in the IP (Internet layer) and TCP/UDP (Transport layer) packet headers. The IP packet header includes IP addresses (source and destination), type of service (TOS), and protocol (TCP or UDP). The TCP and UDP headers both contain source and destination port numbers that can be used to identify certain applications. Five elements in the Internet layer and Transport layer headers (IP source and destination addresses, protocol, and TCP/UDP source and destination port numbers) are referred to as the "5-tuple," and they uniquely identify a connection or flow between two application layer entities [55]. The term “IP flow” is often used to refer to all the packets that have the same 5-tuple.

- **Layer 5.** Application layer classification is performed on elements above the Transport layer, including the higher layer headers and the data payload. Unencrypted traffic can be classified through pattern matching and/or more advanced techniques. Many applications use standard protocols such as hypertext transfer protocol (HTTP), session initiation protocol (SIP), and file transfer protocol (FTP) as part of their communications and expose elements that can be used for classification [49,56,50]. Classification at Layer 5 is sometimes referred to as "Deep Packet Inspection." Encryption generally interferes with attempts to perform pattern analysis or deep packet inspection at this layer [57].

  Traffic that is encrypted or that does not use standard protocols may still be classifiable using signature- or heuristic-based techniques. Heuristic analysis involves inspecting a large set of traffic for behavior patterns [58,59]. It is often possible to infer the type of traffic by examining how many endpoints are talking to each other. For example, encrypted VoIP traffic can sometimes be classified by looking for IP flows both to and from many end-points communicating to a soft-switch, combined with numerical analysis that examines the flow rates and packet payload sizes [60]. Malicious traffic may also be inferred by looking for a many-to-one IP-flow relationship [61].

3.3 Application of service policies

After traffic has been classified, certain service policies can be applied. In addition, packets may be marked so that other processes or network nodes can apply the assigned service policies more readily.

3.3.1 Traffic Markings

Traffic can be marked at Layer 2 and Layer 3 of the network model by setting or changing some element in one of the headers.

- ***Layer 2 Marking (Data Link Layer).*** The commonly used Layer 2 technologies such as ATM, Frame Relay, and Ethernet all include options for marking the Layer 2 frame or packet [62,63]. ATM and Frame Relay both include a field in their headers that can be used to indicate whether the cell or frame can be dropped during periods of congestion. The IEEE 802.1Q standard defines priority and VLAN fields, both of which can be used to mark Ethernet frames [64].

- ***Layer 3 Marking (Internet Layer).*** At Layer 3 the IP header has fields that can be used to mark traffic. The IP header has a field that can be used to indicate either Type-of-Service (TOS) or used to specify a diffserv code point (DSCP) intended to indicate a desired per-hop behavior [9].

3.3.2 Service Policies

Once a packet has been classified, it can be treated according to the assigned service policy. Service policies include scheduling policies (e.g., queuing, shaping, dropping) as well as routing decisions, such as what egress port to use on the network element or whether the packet is eligible to be cached.

3.3.2.1 Scheduling Policies

When there is contention for a network resource, for example the egress port on a network element such as a router or switch or for access to a shared medium, the network element may use a scheduling algorithm to determine the order in which packets are transmitted. Many scheduling algorithms fall into one of three categories: 1) priority scheduling [65], which schedules higher priority traffic before lower priority traffic; 2) rate-based scheduling (such as round-robin scheduling [66] or weighted fair queuing [67]), which allocates resources to isolate the effects of different flows on each other; or 3) deadline scheduling [68], which limits the maximum time allowed before a packet is either transmitted or discarded.

The above algorithms are not mutually exclusive and can be used in combination with each other. For example, a network provider might use a rate-based scheme to separate users from each other in a shared access regime, and then use priority within each share to allow each user to favor their latency-sensitive traffic relative to other traffic. In addition, some network elements support more specialized scheduling algorithms than those described here.

3.3.2.2 Traffic shaping

Traffic shaping is the process that a network might use to limit the rate that a sender (i.e., a device, application, user) can send traffic on a particular link [69]. An operator may use a traffic shaper for example, to implement bandwidth limitations; to limit the rate at which traffic bursts (averting unnecessary delay and loss), or to control the effects of buffer bloat [70]. Shaping of different flows to different parameters implements differentiated treatment.

There are several mechanisms that providers can use to shape traffic; each of these mechanisms has different characteristics. Some of these mechanisms limit a flow to a certain average rate. Other mechanisms allow a sender to periodically "burst" (i.e., send traffic at a higher rate for a period of time before they are shaped to a lower sustained rate). These mechanisms can include leaky bucket, token bucket, and composite shaping (which combines leaky and token bucket shaping) [71].

3.3.2.3 Resource reservation

Resource reservation is a technique appropriate for applications that require a minimum level of network resources in order to function adequately. Resources can be statically configured to reserve them for certain users or application traffic, or resource reservation and associated admission control (denying or granting application requests for special treatment of certain flows based on availability of network resources) can be done dynamically [72,73].

3.3.2.4 Routing policies

In addition to packet scheduling policies, classification can be used by a network node to assist routing decisions. The node, which may be connected to multiple networks, can attempt to optimize the QoE for certain traffic by forwarding it to a path with QoS characteristics that may be aligned with the traffic's requirements such as lower latency, lower round-trip time, less congestion, etc.

4 Differentiation in access network architectures

Differentiation techniques are most often deployed in the access and aggregation networks that operate close to end users. Differentiation can make more of a difference in the lower speed network segments near the network edge that aggregate smaller numbers of flows, because in these segments the relative effect of each flow on other flows is magnified compared to the highly aggregated segments in the network core. When differentiation is used in the core it frequently takes the form of routing along engineered paths. This section focuses on the architectures and access technologies deployed near the network edge.

Broadband networks have been deployed with a number of different network architectures and access types. Several of these network types have developed to take advantage of existing access infrastructure that was originally deployed for other services – for example, telephone service over twisted copper pairs or video over coaxial cable. Other networks were developed to meet specific needs, such as for mobility or for access in remote rural areas. Most broadband networks are engineered to support multiple services sharing common infrastructure. While the designs differ, they are conceptually similar to the degree that they are designed to meet similar requirements. Each of the underlying broadband access networks provides a means to isolate services from one another at the link layer (Layer 2) by creating logical channels. In DOCSIS cable networks the logical channels are called service flows, in telco networks they are called VLANs, and in 3GPP mobile networks they are called bearer channels. Each broadband access technology includes the capabilities to:

- Classify and map traffic to the assigned logical channel.
- Limit the rate at which traffic is delivered over the logical channel.
- Control how traffic in each logical channel is delivered relative to other channels when contention occurs.

In many cases, network design can be traced to the characteristics of the access technology used. Specific access technologies can present unique challenges that require different approaches to differentiation of traffic sent over the access link, as documented in the following sections.

4.1 Telco fixed broadband network architectures

Telecom broadband networks trace their heritage to the telephone networks that delivered analog voice service over twisted copper wire pairs (or “loops”) to households across the developed world for much of the twentieth century. Digital Subscriber Loop (DSL) technology was developed to support broadband data services over these same copper loops. While optical fiber has been increasingly deployed in the access network and maximum DSL speeds have increased by orders of magnitude, in many cases copper loops still deliver data over the last link to the consumer. The evolution from voice to data networks and the physical challenges imposed by the copper loop environment have each influenced how telco broadband network architectures have evolved.

The Broadband Forum (BBF) has specified architectures for broadband access through a series of technical reports [74,75]. A representative architecture for the Multi-Service Access Network is shown in Figure 3 below. In the figure, traffic is transported over a regional access network between the network’s interconnection interface to other networks (interface A10) and the interfaces to customer equipment (interface T). The points at which traffic may be differentiated can be examined by tracing the path of Internet access traffic sent from other networks to the subscriber. This traffic, sourced by a server in a remote network, is transported across the remote network and possibly one or more intermediate transport networks (labeled NSP$_2$) before arriving at the interconnection point A10 for the destination network. The traffic enters and is transported across the regional IP network at the IP layer (Layer 3). At nodes in the regional network, the traffic is aggregated with Internet access traffic destined for other subscribers and may be scheduled along with traffic associated with other IP services, using DSCP markings or other means of classification to determine the differentiated treatment provided during scheduling.

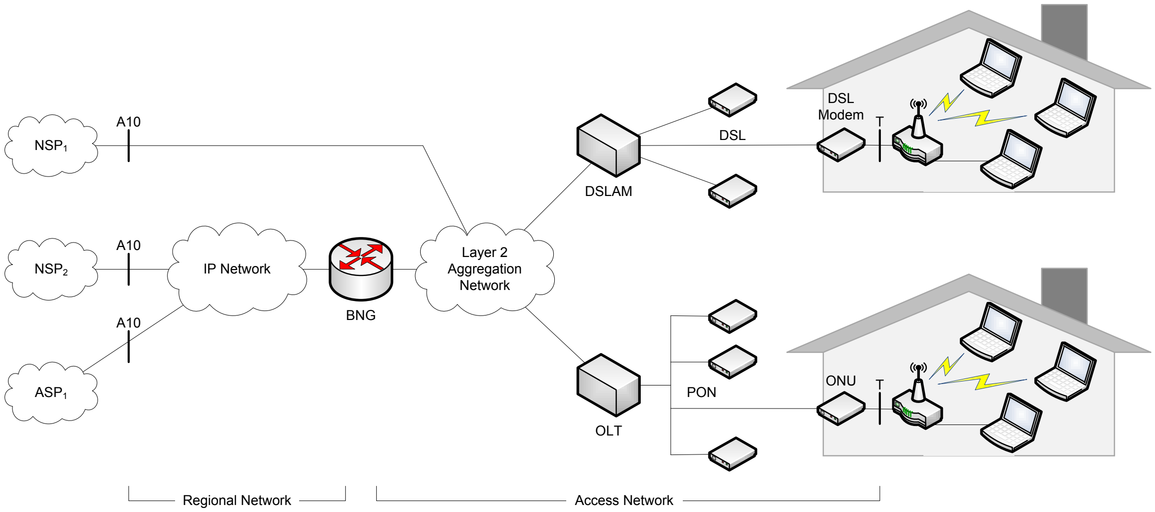

Figure 3: Broadband Multi-Service Reference Model [74]

A node called the Broadband Network Gateway (BNG) provides the primary interface between the IP-based networks on the left and the Layer 2 network – so called because with few exceptions, the nodes within this network ignore any packet information above the Layer 2 (e.g., Ethernet) header – extending to the subscriber on the right. The BNG may also provide per-subscriber shaping to enforce each subscriber's service rate, and scheduling to enforce policy between subscribers. Finally, the BNG may isolate each subscriber's Internet access traffic into a separate VLAN for transport across the access network (note that there are variations in how isolation is implemented by different network operators, including isolation by subscriber and isolation by type of service).

The Multi-Service Access Network supports a variety of IP services in addition to Internet access, including residential services such as IPTV and voice. Traffic for these services may come from network providers (NSP$_2$) or application providers (ASP$_1$) across A10 as IP traffic, or (for services such as Layer 2 business connectivity), from another network provider (labeled NSP$_1$) as Ethernet or other Layer 2 traffic. This traffic may be multiplexed with Internet access traffic in the regional or access network as shown, and may be scheduled alongside Internet access traffic to generate the desired QoS for each service.

The last links used in these networks to reach the customer typically run over either twisted copper pairs known as Digital Subscriber Loops (DSL) or optical fibers, which may be either shared between multiple customers as Passive Optical Networks (PON) or dedicated to a single customer. Each access medium has different characteristics, some of which lend themselves to specific differentiation techniques:

- **Digital Subscriber Loop (DSL)**. DSL technology is used to provide broadband access across twisted pair copper loops. Depending on the network design, DSL technology and loop length, the capacity available across the DSL link may range from one to hundreds of megabits per second (Mbps). Of particular interest are links that provide capacity on the order of 20 to 40 Mbps, which is high enough to support both IPTV and Internet access services, but which may not support both services concurrently at their maximum expected rates. Under these conditions, IPTV traffic may be prioritized to prevent interruption of video programming when the DSL link becomes congested [76].
- **Optical Fiber.** The optical fibers over which Passive Optical Networks (PONs) transmit can carry gigabits per second. However, PONs can be shared by 32 or more customers, so shaping and scheduling (frequently performed at the BNG) enforce policy in the downstream direction. In the upstream direction, only one customer on the PON can transmit data at any given time. The specification for granting upstream allocations includes provisions for differential treatment including weighting, prioritization, and guaranteed bandwidth [77].

4.2 Cable operator network architectures

Cable networks were originally deployed to deliver television services to subscribers over coaxial cables. Starting in the 1990s, these networks evolved to support two-way data communication and have since seen several generations of the Data Over Cable Service Interface Specifications (DOCSIS) standards that specify how broadband access services are provided to residential and small-to-medium business customers [78]. As with telecom networks, the architectures of modern Hybrid Fiber/Coaxial (HFC) cable networks have been influenced both by their history and by the mix of services they offer.

A single cable system typically serves a metropolitan area, including outlying communities. An example of a cable system is shown in Figure 4 below. In a typical cable system, a hub site might provide service to an area consisting of 10,000 to 20,000 households. The hub site connects a hybrid fiber-coax portion of the cable network (the “Access Network”) to the regional data network via the Cable Modem Termination System (CMTS), which connects IP services to customers' cable modems [79].[6]

[6] In some cable systems, some or all of the access network is deployed using Ethernet Passive Optical Network (EPON) technology instead of HFC. CableLabs DOCSIS Provisioning of EPON (DPoE) specifications allow EPON devices to mimic the functionality of a DOCSIS CMTS and cable modem.

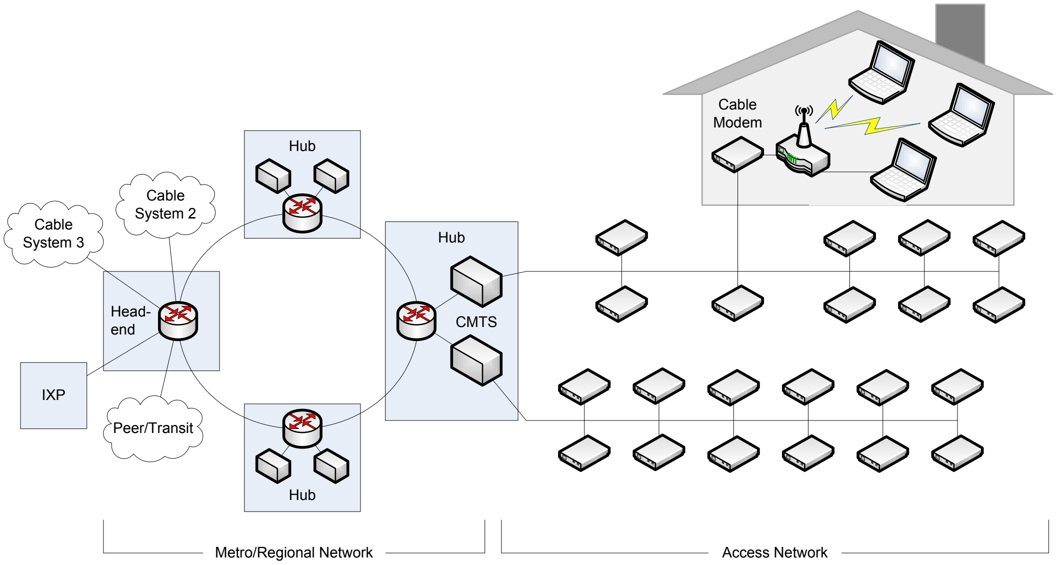

Figure 4: Example Cable System

In the access network, cable spectrum is divided into channels as specified by DOCSIS, which in turn are grouped into Service Groups [79]. Each Service Group consists of a set of upstream and downstream DOCSIS channels whose total capacity is shared by a number of cable modems. Each cable modem typically serves one customer. A single CMTS supports dozens of Service Groups, with each Service Group providing service to dozens of customers. The total capacity of each Service Group can be managed to provide the desired performance to the customers served by that group. Services are shaped and scheduled in the CMTS, and bandwidth for specific services may be reserved to enforce policy, manage contention, and provide each service's QoS. In DOCSIS, services are configured using Service Flows, which are uni-directional (upstream or downstream) logical channels between a CMTS and a cable modem [79]. Each cable modem can support a dozen or more Service Flows, though for basic residential broadband service only two are configured (one upstream and one downstream). Each Service Flow is configured with a QoS using controls that include rate shaper parameters, reserved rate and traffic priority, allowing services to be optimized for applications such as digital phone service or business service level agreements (SLAs). IP packets are classified into Service Flows by classifiers in the cable modem (for upstream flows) or the CMTS (for downstream flows) [79].

One example of service differentiation in cable systems is the hosting of public WLAN hotspots in residential gateways, as described in Section 4.6.

4.3 Satellite Internet

Satellite is used by a variety of Service Providers to deliver broadband services that include Internet access service, voice, video, and enterprise business applications. While often used as a secondary option to terrestrial broadband when addressing challenged service areas, it is also used by a number of providers as the primary option in delivery of broadband services [80,81,82]. The use of satellite faces a number of unique challenges when compared to terrestrial alternatives. This stems from the significant propagation delays of

Geosynchronous Earth Orbit (GEO) satellites and the variable propagation delays and loss inherent to Low-Earth Orbit (LEO) satellites [83]. Given the implications of high and variable latency, link impairments, asymmetry and packet loss, traffic differentiation plays a key role in maximizing the usability of satellite in delivering broadband services.

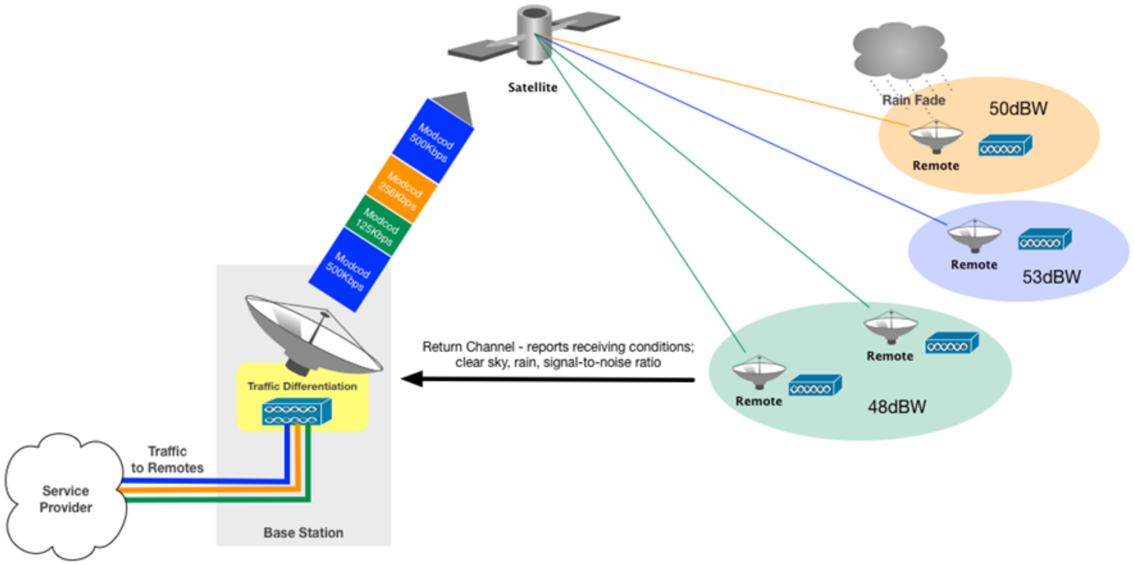

Figure 5: Satellite access architecture [84]

While traffic differentiation may be performed at both base station and remote sites, the focal point is at the base station where traffic flows are multiplexed to the satellite uplink. Both base station and remote sites possess the QoS capabilities of a traditional IP router (classification, marking, scheduling, shaping), with added requirements for cross-layer communication with data-link (layer 2) and physical (layer 1) layers. Due to the unique receive conditions of each remote terminal (clear sky, rain fade, interference), bandwidth resources will vary per site; thereby compromising efficient utilization of the return link and any meaningful enforcement of quality of service (QoS) service level agreements (SLAs). To address this issue, dynamic bandwidth allocation schemes use a feedback loop from remote sites to base stations to communicate the status of bandwidth demand and the conditions of the satellite interface. With such information available, the base station is able to assign bandwidth to its multiplexed traffic flows in a more deterministic manner.

An additional method in which satellite architectures facilitate traffic differentiation is through the use of Performance Enhancing Proxies (PEPs) with split connection capabilities. PEPs have been used to improve the performance of protocols across networks with suboptimal link or subnetwork characteristics [85]. PEPs have no inherent layer restrictions and may be implemented in isolation or through cross-layer coordination in order to achieve holistic performance improvements. Satellite networks have used a variety of PEPs such as TCP compression/acceleration, HTTP compression/acceleration, Link-layer compression, HTTP caching, and TCP spoofing [85,86]. In the context of traffic differentiation, TCP PEPs may be used to give priority to urgent, interactive connections while lower priority connections would yield bandwidth resources by slowing or being suspended when bandwidth congestion arises. This process may be performed by steering

TCP connections, based upon DSCP classifier, to the TCP PEP where they may be compressed and/or accelerated. All other TCP connections follow the normal data path and are not affected by the TCP PEP.

4.4 Mobile (3GPP) architecture

Like their fixed telco counterparts, mobile networks have evolved from voice-only to a rich mix of data and voice services on the network. Mobile networks, however, face a number of unique challenges to delivery across the radio access network, including but not limited to: rapidly changing performance as a user moves within a cell; the need to hand users over from one cell to another; finite spectrum and interference from other cells; and aggregate capacity limitations that constantly change as users move, among other things.

The Third Generation Partnership Project (3GPP) and other organizations have specified sophisticated QoS mechanisms to cope with the dynamic mobile environment [87]. This report focuses on the LTE (and LTE-A) specifications developed by the 3GPP, as these specifications enabled the first end-to-end all-IP mobile networks, are supported by most mobile network operators, and provide the architecture for most current and near future mobile networks [88,89].

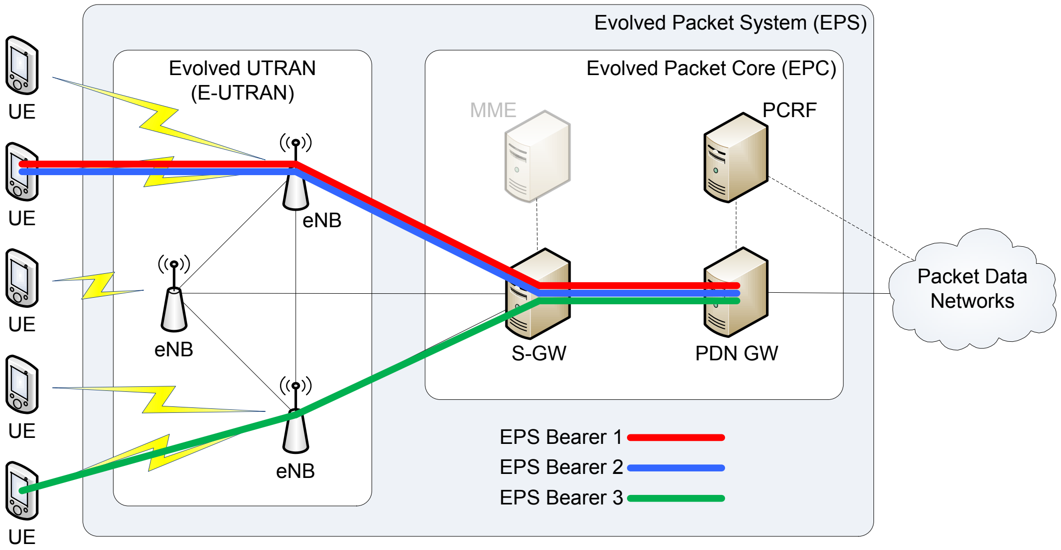

Figure 6: LTE Architecture [90,91]

Figure 6 shows a simplified version of the 3GPP LTE architecture. In this architecture, traffic is transported between a subscriber's user equipment (UE) and remote endpoints across the Evolved Packet System (EPS), which includes the Evolved Node-B (eNB) base station, the Serving Gateway (S-GW), and the Packet Data Network Gateway (PDN GW). LTE networks typically have many eNBs and S-GWs to serve mobile users, and multiple PDN GW interfaces to other networks through which traffic may be routed, depending on whether the traffic is associated with Internet access, voice service, or a different service such as a business VPN.

LTE uses EPS bearers (three of which are shown in Figure 6) to isolate and differentiate between services across the EPS [92]. An EPS bearer is a connection set up between a user equipment and an interface linking a PDN GW to an external network, and is identified by three characteristics: the user equipment at one end; the PDN GW interface at the other end; and a QoS profile that identifies the performance objectives associated with the traffic. The QoS profile contains several parameters including the QoS Class Indicator (QCI), which defines a number of QoS-related characteristics for the bearer as shown in Table 1.

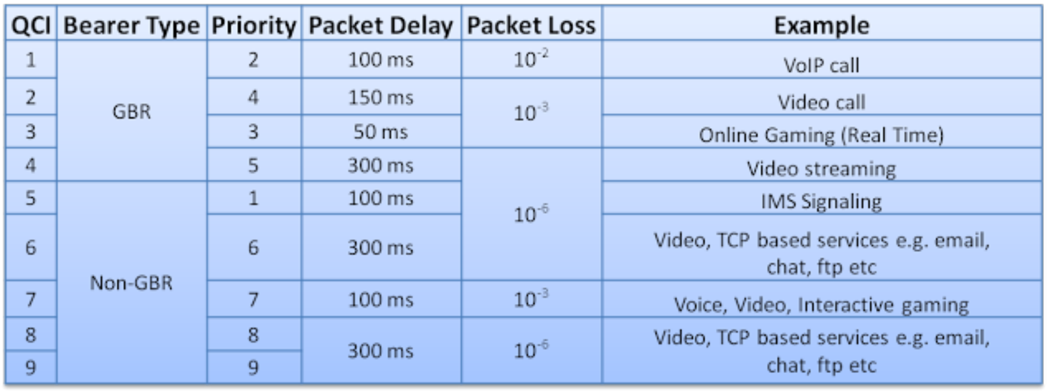

QCI	Bearer Type	Priority	Packet Delay	Packet Loss	Example
1	GBR	2	100 ms	10^{-2}	VoIP call
2		4	150 ms	10^{-3}	Video call
3		3	50 ms		Online Gaming (Real Time)
4		5	300 ms	10^{-6}	Video streaming
5	Non-GBR	1	100 ms		IMS Signaling
6		6	300 ms		Video, TCP based services e.g. email, chat, ftp etc
7		7	100 ms	10^{-3}	Voice, Video, Interactive gaming
8		8	300 ms	10^{-6}	Video, TCP based services e.g. email, chat, ftp etc
9		9			

Table 1: Standardized QCI types [92]

Multiple nodes in the Evolved Packet System play a role in implementing QoS and policy management:

- The Policy and Charging Rules Function (PCRF) uses available network information and operator-configured policies to create service session-level policy decisions [93].
- The Policy and Charging Enforcement Function (PCEF) located in the PDN GW enforces the policy decisions forwarded from the PCRF by establishing bearers, mapping service data flows to bearers, and performing traffic policing and shaping [93].
- Both the eNB and the user equipment may allocate bandwidth and schedule traffic using the parameters associated with each bearer.

Mobile networks have always relied on tight performance requirements in the backhaul networks that transport traffic between cellular base stations and the nodes that make up the mobile core. As those backhaul networks have migrated from time division multiplexing (TDM)/synchronous optical networking (SONET) to IP packet networks, differentiation techniques have been key to ensuring that the most time sensitive traffic is transported with minimum delay.

Industry standards provide guidance for network operators implementing mobile backhaul [94,95]. One standard defines up to four classes of service with increasingly stringent performance requirements to enable bearer traffic with time-sensitive QCIs to be scheduled first over Ethernet-based services [94]. Another provides support for multiple classes of service, including the classes identified by the first standard, over MPLS networks [95]. When these standards are used, bearers are mapped based on QCI value into the appropriate class of service and then scheduled within the backhaul network to achieve the necessary performance.

4.5 Fixed wireless network architecture

Fixed wireless access networks frequently serve rural areas that are characterized by long distances between subscribers and lack of high-speed wired communications infrastructure. As a result, these networks frequently make use of multiple wireless links, both in the "last mile" used to reach the subscriber and in the "middle mile" backhaul and aggregation links that send traffic between towers. These links are subject to capacity limits and performance variation based on the spectrum used and atmospheric conditions.

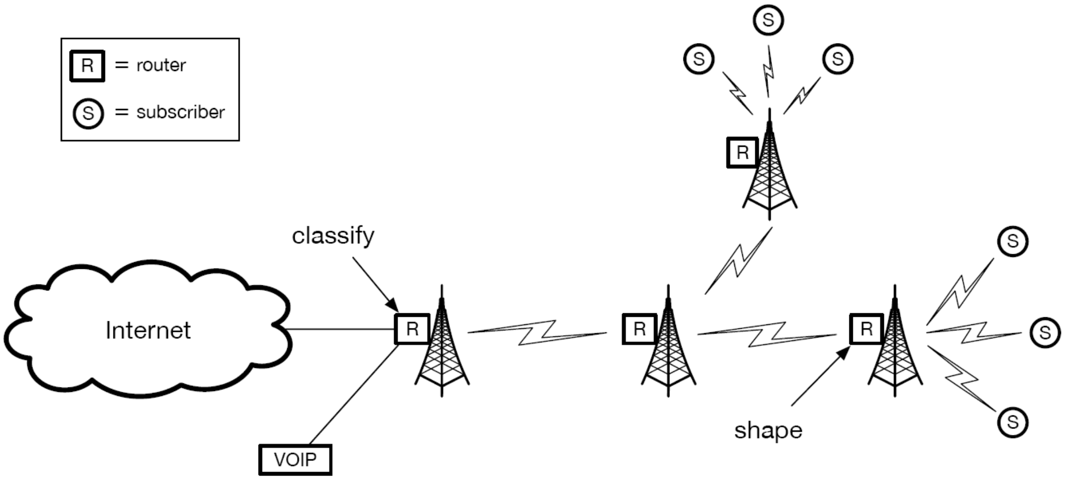

Figure 7: Example fixed wireless access network architecture

Figure 7 shows an example fixed wireless access network. In the figure, downstream traffic is classified and any differentiated treatment is applied at the edge router on the left where the wireless Internet service provider (WISP) network interconnects to other networks. Routers at each relay tower may schedule traffic as necessary, and per-service shaping to subscriber service rates is typically performed at access node towers such as the ones shown on the right. Most WISPs provide business-grade Internet access services in addition to consumer-grade services, and some offer business connectivity services as well.

4.5.1 Middle Mile

Most middle-mile links in WISP networks use 5 GHz unlicensed spectrum, which with modern equipment can deliver 100 to 200 Mbps over distances of 20 miles or more [96] and which make it feasible to deploy service to remote and sparsely populated areas. As demand for capacity has grown, WISPs have begun to deploy fiber middle-mile connections where they are available, supplemented with microwave backhauls in licensed spectrum in core areas. Short range, gigabit capacity radios using unlicensed 24 GHz spectrum have also become popular for short backhaul links [97].

Despite the increased capacity of the middle-mile options available to WISPs, exponential increases in demand continue to strain capacity limits. This has led to a rise in the use of packet processing systems by WISPs, to shape network flows and more efficiently use capacity in overloaded backhauls [98]. End-user data can traverse a daisy-chain of multiple middle-mile microwave connections before reaching a fiber backhaul. In this environment, differentiated treatment can enable delivery of services with predictable jitter and latency.

4.5.2 Last Mile

The "last mile" links in a WISP are typically point-to-multipoint links in which one tower serves multiple subscribers. A typical unlicensed access point has a variable capacity that depends on spectrum interference levels, channel sizes, distance of customers from the access point and propagation characteristics of the spectrum utilized. Under real-world conditions, a modern fixed wireless access point with 20 Mhz of clear spectrum may have a capacity of up to 50 Mbps [99]. Without differentiated treatment, a small number of high bandwidth video streams can consume most of the capacity of the access point, leading to degraded service for the rest of the end users. Configuring rate limits on flows such as video streams is one tool used by WISPs to ensure the access point approaches but never reaches the point of overload without blocking services. This allows for a significant increase in oversubscription levels without a noticeable decrease in performance to the end user [100]. For WISP customers in remote areas with limited alternatives, this is one of the few ways to provide a quality user experience without upgrading every segment between the end-user and the network core.

4.6 Wireless LAN Public Hotspot Networks

Wireless Local Area Networks (WLANs) based on IEEE 802.11 radio technology are frequently used to provide wireless broadband access to users at homes, enterprises, and venues such as restaurants, stores, and airports. These WLANs can generally be considered as an extension of any of the network architectures discussed above. From a differentiation and subscriber viewpoint, they can be categorized in several ways:

- Home WLANs are generally considered private (although many have little or no security, they are still intended only for use by members and guests of the household). With the exception of the shared WLAN technologies discussed below, they are not considered further.

- Venue-based WLANs and public hotspots may be designed for larger populations and may implement policies to prevent misuse, including: isolation of devices so that they cannot communicate directly with each other on the WLAN to prevent LAN-based attacks (although devices that need to communicate can still do so through the IP layer); and limiting individual users' rates to prevent any one user from monopolizing the available WLAN bandwidth. Some of these venues also offer multiple tiers of service shaped to different rates.

Of particular interest are public hotspot networks with multiple geographic locations and centralized management, such as those operated by fixed and mobile network providers and over-the-top (OTT) WLAN network operators [101,102].   In a number of these networks, subscribers use in-home wireless routers that support separate logical WLANs – one WLAN for the subscriber's private use and a second WLAN for use as a public hotspot. A broadband network provider can either manage the hotspot as an extension of their network or contract with an OTT WLAN provider to manage the service. Conversely, an OTT WLAN provider can manage a hotspot with or without coordinating with the subscriber's broadband provider. When the broadband provider is involved in managing the hotspot, the management domain covers both the WLAN network and its upstream connectivity – otherwise, the management domain only covers the WLAN network. Depending on the scope of the hotspot operator's management domain, the following options may be available to differentiate the treatment of public WLAN traffic compared to the subscriber's private traffic:

- Within the WLAN, public traffic can be scheduled and shaped so that it does not impact the subscriber's traffic.
- In the uplink, public traffic can be isolated from the subscriber's traffic in a separate logical broadband connection, with a firewall between the WLANs.
- In the broadband network, public traffic can be excluded so that it does not count against subscriber usage limits.

In a fully managed solution the operator may choose to apply differentiation techniques applicable to the access network technology used along with those available in the WLAN specifications. In a solution where the operator only manages the WLAN network device, only WLAN specific differentiation techniques are available to the operator. The differentiation techniques may be used to isolate public traffic from subscriber traffic that is using the WLAN, as well as aiding in the seamless handover between WLANs and cellular networks.

4.7 Network Function Virtualization (NFV)

Network operators are beginning to deploy Network Function Virtualization (NFV) in the above architectures to make them more flexible and responsive to changing business and technical requirements. NFV allows network operators to implement functions as software in virtual machines using commoditized hardware, instead of in dedicated hardware appliances. NFV hardware is often deployed in data centers, central offices or other locations where pooled resources can be managed and orchestrated; however, some network operators are also experimenting with hosting network functions in lightweight "containers" (a lower overhead alternative to virtual machines) and in smaller hardware resources distributed at the edge of the network.

NFV has several implications for the differentiated treatment of Internet traffic:

- Virtualization of network functions makes it possible to deploy and modify these functions much more flexibly in response to new service offerings and requirements. For example, virtualized traffic shapers or firewalls can be deployed in the logical paths where they are needed without physically deploying new hardware across the network perimeter.
- NFV can facilitate new services that may be impractical with conventional approaches. For example, by virtualizing the routing and network address translation (NAT) functions in residential gateways, network operators have better visibility of flows to optimize features such as parental control of children's online activities.
- Conversely, NFV has the potential to degrade performance by introducing additional latency or jitter in the physical path. For example, traffic flows may need to be "steered" through a sequence of network functions; if those functions are hosted at distant network locations, the latency of those flows may suffer. Standards bodies and the broader networking research community are actively exploring the performance and security of NFV, as well as designing architectures and algorithms for improved management and orchestration [103].
- Since virtual machines can be isolated, NFV allows network operators to host functions that are deployed and managed by third parties. For example, a content provider might perform CDN functions such as caching or server selection on the network provider's NFV infrastructure, improving performance for that content.

5 Examples

Network operators often deliver multiple services such as Internet access plus other IP-based, non-Internet services over a common infrastructure. The examples in this section illustrate multiple services being delivered over common infrastructure each with their own delivery requirements using the differentiation techniques described earlier in this report.

5.1 Interactive service differentiation

Interactive applications and services often have stringent delivery requirements to meet the interactive nature of the application or service. Example applications or services include voice and video.

5.1.1 Effects of carrier grade interactive voice on Internet access services

Interactive voice is a real-time application that requires low network delay and jitter in order to provide a good QoE to the user [104]. While jitter buffers can compensate for a limited amount of jitter, they add to the average end-to-end delay, which in turn places more stringent requirements on the network's contribution to end-to-end delay. Hence the engineering of carrier grade voice services seeks to minimize both delay and jitter in the network.

Carrier grade voice services commonly rely on a combination of “connection admission control” (allowing new calls only if there are sufficient network resources) and bandwidth reservation. Carrier grade voice traffic may also be prioritized and scheduled ahead of traffic for Internet access services on networks carrying both types of traffic.

3GPP mobile networks and some fixed networks use the IP Multimedia Subsystem (IMS) for call admission control and bandwidth reservation [105]. The IMS creates new voice bearers, differentiated by the 5-tuple (Section 3.2), and scheduled via a combination of resource admission control and priority. IMS allows 3rd-party voice providers to request differentiation and QoS for mobile roaming users. More commonly, 3rd-party voice providers mark packets with a differentiated services code point requesting appropriate treatment [9].

Cable networks (Section 4.2) use PacketCable™ [106] for the call admission control and bandwidth reservation. Similar to how IMS works, PacketCable dynamically creates service flows for the admitted calls. The dynamic service flows are differentiated by the 5-tuple, assigned bandwidth and priority, and scheduled.

Telco networks (Section 4.1) use Ethernet VLANs, and IP DSCP markings, and priority code points to isolate, schedule, and route carrier grade voice traffic.

Carrier grade voice traffic, which is transported over UDP, is differentiated primarily to mitigate degradation in its QoE in converged networks, caused by TCP traffic competing for resources. Because TCP tends to fill up the queue feeding a congested link, it increases the delay and jitter experienced by all traffic in the same queue – which includes voice traffic if it is not differentiated. There are two fundamental approaches to improving voice QoE, both of which remove the voice traffic from the queue used by TCP traffic. One approach is to divert voice traffic to another network path engineered to the purpose, and the other is to queue voice traffic separately and schedule it before TCP traffic. Neither of these approaches degrades QoE for the TCP traffic:

- If voice traffic is diverted to another path, the TCP traffic benefits in that it no longer competes with the voice traffic.
- If voice is prioritized over TCP traffic, the TCP traffic maintains the same QoE as it would without prioritization. This is true because both the capacity used by voice traffic (sent over UDP, which does not change rate in response to congestion) and the overall capacity of the network links are independent of whether or not the voice traffic is prioritized. Simple subtraction shows that the remaining capacity available for TCP traffic – that is, the difference between the overall capacity and the capacity used by voice – is also independent of whether or not the voice traffic is prioritized. The only observable difference in the QoE of the TCP traffic is an increase in delay variation relative to the optimal delay. If the increase is relatively small TCP performance will not be adversely affected. Since TCP is insensitive to

moderate amounts of relative delay variation, its QoE is not affected by the prioritization of the voice traffic.

In either case, differentiation can provide a significant benefit to the QoE for voice traffic without degrading the QoE for the TCP traffic.

5.1.2 Managing the impact of streaming video on other traffic

A typical video stream, as sent by a server, consists of a series of large bursts of traffic, or "chunks," where each chunk consists of multiple packets transmitted as quickly as possible. Sequential chunks are separated by time periods that can span seconds [107]. The transmission rate for each chunk is much higher than the average rate of the encoded stream, which is a function of the average chunk size and the time between chunks.[7] The video client buffers the chunks and then plays them out at the encoded rate.

When a chunk from a video stream arrives at a bottleneck link, it can cause significant delay and jitter for other traffic sharing the same link, causing severe degradation in the QoE of time-sensitive applications such as interactive voice. This problem can be mitigated via a technique known as pacing [110], in which the video stream is differentiated and traffic shaped to a rate equal to or greater than the stream's average rate, but still lower than the bottleneck link's rate. Pacing spaces out the video packets in time, allowing other traffic in between the chunks and in doing so may reduce the latency and jitter experienced by other traffic. Since the first packet in each chunk is not delayed, the net effect of pacing on streaming video is to deliver video packets to the receiver at a more consistent rate without creating any additional delay in video playback. In effect, network pacing performs the same "smoothing" function in the received video content that the receive buffer in the video client would have performed had the chunks been received in discrete high speed bursts, so the QoE for the streaming video may be maintained because the content in each chunk is still received before the decoder needs it.

Pacing is an example of differentiated treatment that is implemented in mobile networks and that acts on the traffic within Internet access services. It may also be implemented by the sending service or application, reducing the need for differentiation in the network. As noted above, this technique can improve the QoE for other traffic without degrading the QoE for OTT video streams.

5.2 Transmission Control Protocol (TCP) performance optimizations

Transmission Control Protocol (TCP) is the dominant transport protocol used in the Internet. The protocol is designed to deliver traffic reliably from one endpoint to another, as quickly as the network will allow [17]. Under certain network and traffic conditions however, the protocol can perform poorly and can even contribute unnecessarily to network congestion. Two key elements of the protocol's design are:

[7] Most modern over-the-top video streams use adaptive rate technology: the server offers each video at multiple bitrates, and receivers request the best rate supported by the network, switching between rates mid-flow if necessary to avoid stalling [108,109].

- TCP provides reliable delivery by having the receiver acknowledge receipt of traffic from the sender. If no acknowledgement is received within some time frame, the sender will assume that the unacknowledged packet is lost and will retransmit it. The sender may also reduce the number of unacknowledged packets it sends into the network based on the assumption that the loss was due to congestion.

- TCP adapts to available capacity, increasing the amount of traffic it sends until it detects congestion. One consequence of this is that when there is enough traffic to send, TCP virtually guarantees that it will create recurring momentary congestion at some point in its network path [111]. This effect exists by design, and it cannot necessarily be eliminated by increasing capacity.[8]

When there is a large amount of upstream traffic – especially in networks with asymmetric speeds and capabilities (e.g., mobile, satellite) – upstream congestion from one or more TCP flows can cause acknowledgements associated with downstream flows to be dropped. This degrades performance and may add to congestion in the downstream direction, because downstream traffic that was sent and received must now be sent again. It also causes a multiplicative effect in that one upstream flow can degrade the performance of many downstream flows [112].

Satellite and cellular networks commonly prioritize empty acknowledgements (i.e., packets that include the acknowledgement flag but no data) to mitigate the above issues [113]. Empty acknowledgements are short packets that create only minimal delay for other traffic if they are prioritized, so this use of prioritization either does no harm (when there is no congestion) or results in a net performance gain (when congestion is present). Some home routers also allow users to prioritize outgoing acknowledgements [114].

Another TCP performance optimization commonly used is in the transition between dissimilar networks. For example, a mobile network includes a radio access link with high latency and moderate packet loss (some of which is due to transmission errors rather than congestion), which interconnects with wired networks having low latency and low loss. In this environment, a client (in the high latency/moderate loss environment) can struggle to get a TCP connection up to speed, which affects application performance. Network based techniques are often also used to mitigate this. For example, a proxy may convert from a TCP model suitable for low-loss/low-latency networks to one suitable for moderate-loss/high latency networks [115,116]. [9] Using this model, the mobile network may be more efficient as it spends less time transitioning in and out of idle states. The end-user also achieves a faster page-load time in web applications.

[8] Increasing capacity will give more bandwidth to competing flows, but TCP flows – by design – will send as fast as they are able until the new increased capacity path is congested. Given the same traffic load, however, the severity of the momentary congestion should decrease with increased capacity.

[9] This technique of differentiation is commonly used in mobile networks. The optimization device recognizes which flows will benefit from a higher initial burst (e.g. web) versus ones that will not (e.g. long-form video). The device then proxies the flows to be optimized, using a TCP congestion model such as TCP 'cubic' or 'TCP reno' on the Internet side, and a congestion model such as TCP 'Westwood+' (which is more optimized for long-latency, moderate loss environments) on the access side. For more information see [116].

5.3 User-defined differentiated treatment

Recent experimentation suggests possible trends towards network differentiation that is configurable by users. For example, participatory networking (PANE) allows the network to expose a configuration API to users, applications and end hosts, potentially allowing these systems to specify scheduling or other differentiation parameters [117]. PANE's user-facing API exposes abstractions that allow users and applications to make requests for guaranteed minimum bandwidth, or to prefer (or avoid) network paths that have specified performance properties. For example, the user might route bulk traffic through a packet shaper during busy hours, or avoid certain parts of the network. There are also network controllers that configure QoS using application-described requirements and a database of network state [118], as well as controllers that can perform application classification and automatically assign application traffic flows to the appropriate traffic classes [119].

5.4 Differentiation in the presence of secure traffic

Secured traffic that uses encrypted transport protocols (e.g., TLS) to protect the confidentiality and integrity of the communication session necessarily obfuscates the payload of packets. Because of this, classification is limited to either an explicit control mechanism (e.g., IMS) or header fields that fall outside the encrypted payload (e.g., DSCP, SNI [120], IP, protocol number). Encrypted traffic has become increasingly prevalent on the web and the larger Internet in recent years [121].

Some satellite and in-flight network operators have deployed proxy systems that allow differentiation of encrypted traffic. They do this by breaking the end-to-end encryption principle in favor of two encrypted segments, or in some cases with one segment being unencrypted entirely. Examples of these systems include:

- Satellite operator ViaSat has developed a modified version of the Chrome browser that decrypts traffic inside their network, in order to optimize performance [122].
- To improve performance on retail in-flight WLAN network access, the networks provided by Gogo Inflight Internet, for a period of time, dynamically forged TLS certificates in order to shape traffic or block high-bandwidth uses such as video streaming (due to popular outcry, this practice was ceased shortly after discovery) [123].

In each of these cases, the network provider has made a decision to trade security for performance in order to differentiate between different data flows, and serve as a man-in-the-middle for an otherwise secure communication. As an additional byproduct they generally weaken the security of the connection against other more malicious attackers. For example, in the Gogo case, the provider mandated that the user disregard certificate authentication warnings. This would have applied indistinguishably both to those certificates generated by Gogo as well as those made by an attacker elsewhere on the network. The risks to the security and privacy of end users can be significant.

6 Observations

From the analysis made in this report and the combined experience of its members when it comes to the differentiated treatment of Internet traffic, the BITAG Technical Working Group makes the following observations.

6.1 TCP causes recurring momentary congestion

As mentioned in Sections 2.4.2 and 5.2, when TCP transfers a large file, such as video content or a large web page, it practically guarantees that it will create recurring momentary congestion at some point in its network path [111]. This effect exists by design, and it cannot necessarily be eliminated by increasing capacity. Given the same traffic load, however, the severity of the momentary congestion should decrease with increased capacity.

6.2 A nominal level of packet discard is normal

As mentioned in Sections 2.4.2 and 5.2, packet discard occurs by design in the Internet. Protocols such as TCP use packet discard as a means of detecting congestion, responding by reducing the amount of data outstanding and with it self-induced congestion on the transmission path. Rather than being an impairment, packet discard serves as an important signaling mechanism that keeps congestion in check.

6.3 The absence of differentiation does not imply comparable behavior among applications

As discussed in Sections 5.1 and 5.2, in the absence of differentiation, the underlying protocols used on the Internet do not necessarily give each application comparable bandwidth. For example:

- TCP tends to share available capacity (although not necessarily equally) between competing connections. However, some applications use many connections at once while other applications only use one connection.
- Some applications using RTP/UDP or other transport protocols balance transmission rate against experienced loss and latency, reducing the capacity available to competing applications.

6.4 Differentiated treatment can produce a net gain in Quality of Experience (QoE)

As introduced in the Section 2 discussion on the relationship between QoS and QoE and later in Section 5.1, when differentiated treatment is applied with an awareness of the requirements for different types of traffic, it becomes possible to create a benefit without an offsetting loss. For example, some differentiation techniques improve the Quality of Service (QoS) or Quality of Experience (QoE) for particular applications or classes of applications without negatively impacting the QoE for other applications or classes of applications. The use and development of these techniques has value.

6.5 Access technologies differ in their capabilities and characteristics

Specific architectures and access technologies have unique characteristics that are addressed using different techniques for differentiated treatment. Section 4 describes how differentiation is accomplished in various access network architectures.

6.6 Security of traffic has at times been downgraded to facilitate differentiation techniques

As discussed in Section 5.4, encrypted traffic is on the rise and it has implications for current differentiation techniques. In response to this increase, some satellite and in-flight network operators have deployed differentiation mechanisms that downgrade security properties of some connections to accomplish differentiation. The resulting risks to the security and privacy of end users can be significant, and differentiation via observable information such as ports and traffic heuristics is more compatible with security.

7 Recommendations

This section of the report presents recommendations of the BITAG Technical Working Group (TWG).

7.1 Network operators should disclose information on differential treatment of traffic.

In previous reports, BITAG has recommended transparency with respect to a number of aspects of network management [124,3,125,126,127,128]. BITAG continues to recommend transparency when it comes to the practices used to implement the differential treatment of Internet traffic.

Specifically with respect to consumer-facing services such as mass-market Internet access, network operators should disclose the use of traffic differentiation practices that impact an end user's Internet access service. The disclosure should be readily accessible to the public (e.g. via a webpage) and describe the practice with its impact to end users and expected benefits in terms meaningful to end users. The disclosure should include any differentiation amongst Internet traffic and should disclose the extent and manner in which other services offered over the same end user access facilities (for example video services) may affect the performance of the Internet access service.

7.2 Network operators and ASPs should be encouraged to implement efficient and adaptive network resource management practices

In a previous report BITAG recommended that ASPs and CDNs implement efficient and adaptive network resource management practices [3]; we reiterate that recommendation here, to include network operators. Examples of such practices might target the minimization of latency and variation in latency induced in network equipment, ensuring sufficient bandwidth for expected traffic loads, and the use of queue management techniques to manage resource contention issues.

7.3 Quality of Service metrics should be interpreted in the context of Quality of Experience

Common Quality of Service metrics, often included in commercial service level agreements, include throughput, delay, delay variation, and loss, among other things. As noted in Section 2.3 and 6.4, from the viewpoint of the end user application, these metrics trade off against each other and must be considered in the context of Quality of Experience. For example, since TCP Congestion Control and adaptive applications depend on loss to infer network behavior, actively trying to reduce loss to zero leads to unintended consequences. On the other hand, non-negligible loss rates often directly reduce the user's Quality of Experience. Hence, such metrics should be interpreted in the context of improving user experience.

7.4 Network operators should not downgrade, interfere with, or block user-selected security in order to apply differentiated treatment.

Network operators should refrain from preventing users from applying over-the-top encryption or other security mechanisms without user knowledge and consent. Networks should not interfere with, modify, or drop security parameters requested by an endpoint to apply differentiated treatment. Given the potential for possible exposure of sensitive, confidential, and proprietary information, prior notice should be given to end users of traffic differentiation features that affect security properties transmitted by endpoints. (See Section 5.4)

8 References

[1] Federal Communications Commission, “In the Matter of Protecting and Promoting the Open Internet,” FCC 15-24A1, February 26, 2015. https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-24A1.pdf.

[2] J. Lennox, “Connection-Oriented Media Transport over the Transport Layer Security (TLS) Protocol in the Session Description Protocol (SDP),” RFC5472, July 2006, https://tools.ietf.org/html/rfc4572.

[3] Broadband Internet Technical Advisory Group (BITAG), Real-time Network Management of Internet Congestion, October 2013, http://www.bitag.org/documents/BITAG_-_Congestion_Management_Report.pdf.

[4] Information Sciences Institute, University of Southern California, “Internet Protocol, DARPA Internet Program Protocol Specification,” RFC791, September 1981, https://tools.ietf.org/html/rfc791.

[5] G. Huston, IPv4 Address Report, August 2, 2015, http://www.potaroo.net/tools/ipv4/.

[6] The Hosting News, ColoGuard Offers IPv6 Addresses as IPv4 Addresses are Exhausted, July 31, 2015, http://www.thehostingnews.com/cologuard-offers-ipv6-addresses-as-ipv4-addresses-are-exhausted.html (last visited Sept. 3, 2015).

[7] Cisco, 6Lab, IPv6 Adoption Maps, http://6lab.cisco.com/ciscolive (last visited Sept. 3, 2015).

[8] P. Almquist, “Type of Service in the Internet Protocol Suite,” RFC1349, July 1992, https://tools.ietf.org/html/rfc1349.

[9] K. Nichols, S. Blake, F. Baker, D. Black, “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers,” RFC 2474, December 1998, https://tools.ietf.org/html/rfc2474.

[10] Mills, D.L. The Fuzzball. Proc. ACM SIGCOMM 88 Symposium (Palo Alto CA, August 1988), 115-122.

[11] Computer History Museum, Internet History, 1980s, 1986, http://www.computerhistory.org/internet_history/internet_history_80s.html (last visited Sept. 3, 2015).

[12] J. Nagle, Congestion Control in IP/TCP Internetworks, RFC 896, January 1986, https://tools.ietf.org/html/rfc896.

[13] R. Braden, “Requirements for Internet Hosts – Communication Layers,” RFC1122, October 1989, https://tools.ietf.org/html/rfc1122.

[14] R. Braden, “Requirements for Internet Hosts – Application and Support,” RFC1123, October 1989, https://tools.ietf.org/html/rfc1123.

[15] Internet Assigned Numbers Authority (IANA), Service Name and Transport Protocol Port Number Registry, July 2015, http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml.

[16] A. L. Russell, OSI: The Internet that wasn’t, How TCP/IP eclipsed the Open Systems Interconnection standards to become the global protocol for computer networking, IEEE Spectrum, July 30, 2013, http://spectrum.ieee.org/computing/networks/osi-the-internet-that-wasnt (last visited Sept. 3, 2015).

[17] Information Sciences Institute, University of Southern California, “Transmission Control Protocol, DARPA Internet Program Protocol Specification,” RFC 793, Sept. 1981, https://tools.ietf.org/html/rfc793.

[18] J. Postel, User Datagram Protocol, RFC 768, Aug. 1980, https://tools.ietf.org/html/rfc768.

[19] S. Kent, “IP Encapsulating Security Payload (ESP),” RFC4303, December 2005, https://www.ietf.org/rfc/rfc4303.txt.

[20] Huston, Geoff, "Quality of Service - Fact or Fiction?" The Internet Protocol Journal 3, no. 1 (2000): 30. Accessed July 29, 2015. http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_3-1/ipj_3-1.pdf.

[21] R. Geib and D. Black , "Diffserv interconnection classes and practice", Internet Draft, July 1, 2015, https://tools.ietf.org/html/draft-ietf-tsvwg-diffserv-intercon.

[22] N. Lynch, N. Griffeth, M. Fischer, and L. Guibas, “Probabilistic Analysis of a Network Resource Allocation Algorithm,” Journal of Information and Control, vol. 68, Issue 1-3, Jan/Feb/March 1986, pp. 47-85, *a*vailable at http://www.sciencedirect.com/science/article/pii/S0019995886800286.

[23] K. Nahrstedt and R. Steinmetz, “Resource Management in Multimedia Networked Systems,” 1994, Technical Reports (CIS), Paper 331, *available at* http://repository.upenn.edu/cis_reports/331.

[24] Kilkki, K., “Quality of Experience in Communications Ecosystem ,” Journal of Universal Computer Science, vol.

14, no. 5, pp.615-624, March 2008.

[25] A. Charny, J. Bennet, K. Benson, J. Boudec, A. Chiu, W. Courtney, S. Davari, V. Firoiu, C. Kalmanek, K. Ramakrishnan, “Supplemental Information for the New Definition of the EF PHB (Expedited Forwarding Per-Hop Behavior),” RFC 3247, March 2002, https://tools.ietf.org/html/rfc3247.

[26] R.K.P. Mok, E.W.W. Chan, and R.K.C Chang, ”Measuring the quality of experience of HTTP video streaming,” Conference: Proceedings of the 12th IFIP/IEEE International Symposium on Integrated Network Management, IM 2011, Dublin, Ireland, 23-27 May 2011.

[27] International Telecommunication Union, Standardization Sector (ITU-T), “Mean Opinion Score (MOS) terminology,” Recommendation P.800.1, July 2006, https://www.itu.int/rec/T-REC-P.800.1-200303-S/en.

[28] J. Babiarz, K. Chan, and F. Baker, “Configuration Guidelines for DiffServ Service Classes,” RFC4594, August 2006, https://tools.ietf.org/html/rfc4594.

[29] International Telecommunication Union, Standardization Sector (ITU-T), “End-user multimedia QoS categories,” Recommendation G.1010, November 2001.

[30] A. Bouch, A. Kuchinsky, N. Bhatti, “Quality is in the Eye of the Beholder: Meeting Users' Requirements for Internet Quality of Service,” Proceedings of the SIGCHI conference on Human Factors in Computing Systems (CHI '00), pp.297-304, April 2000.

[31] International Telecommunication Union, Standardization Sector (ITU-T), “Definitions of terms related to quality of service,” Recommendation E.800, September 2008, http://www.itu.int/rec/T-REC-G.1010-200111-I.

[32] H. Koumaras, A. Kourtis, D. Martakos, and J. Lauterjung, “Quantified PQoS assessment based on fast estimation of the spatial and temporal activity level,” Multimedia Tools and Applications, vol. 34, issue 3, pp. 355-374, September 2007.

[33] R. Beuran, M. Ivanovici, B. Dobinson, N. Davies, and P. Thompson, “Network Quality of Service Measurement System for Application Requirements Evaluation,” International Symposium on Performance Evaluation of Computer and Telecommunication Systems, SPECTS’03, Montreal, Canada, July 20-24, 2003, pp. 380-387, https://ivanovic.web.cern.ch/ivanovic/articles/qos_measurement.pdf.

[34] W. Cherif, A. Kentini, D. Negru, and M. Sidibe, "A_PSQA: PESQ-like non-intrusive tool for QoE prediction in VoIP services," 2012 IEEE International Conference on Communications (ICC), vol., no., pp.2124,2128, 10-15 June 2012.

[35] D. Hernando, J.E.L. de Vergara, D. Madrigal, and F. Mata, "Evaluating quality of experience in IPTV services using MPEG frame loss rate," 2013 International Conference on Smart Communications in Network Technologies (SaCoNeT), vol.03, no., pp.1,5, 17-19 June 2013.

[36] M. Fiedler, T. Hossfeld, and P. Tran-Gia, "A generic quantitative relationship between quality of experience and quality of service," IEEE Network, vol.24, no.2, pp.36,41, March-April 2010.

[37] International Telecommunication Union (ITU), Telecommunication Standardization Sector, Series: Transmission Systems and Media, Digital Systems and Networks, One Way Transmission Time, May 2003, ITU G.114, https://www.itu.int/rec/dologin_pub.asp?lang=e&id=T-REC-G.114-200305-I!!PDF-E&type=items.

[38] Cisco, Understanding Delay in Packet Voice Networks, February 2, 2006, http://www.cisco.com/c/en/us/support/docs/voice/voice-quality/5125-delay-details.html (last visited Sept. 3, 2015).

[39] F. Baker and G. Fairhurst, IETF Recommendations Regarding Active Queue Management, RFC 7567, July 2015, https://tools.ietf.org/html/rfc7567.

[40] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, Assured Forwarding PHB Group, RFC 2597, June 1999, https://tools.ietf.org/html/rfc2597.

[41] K. Ramakrishnan, S. Floyd, and D. Black, “The Addition of Explicit Congestion Notification (ECN) to IP,” RFC 3168, Sept. 2001, https://tools.ietf.org/html/rfc3168.

[42] Information Sciences Institute – University of Southern California, “Internet Protocol: DARPA Internet Program Protocol Specification,” RFC 791, Sept. 1981, https://tools.ietf.org/html/rfc791.

[43] K. Chan, R. Sahita, S. Hahn, and K. McCloghrie, “Differentiated Services Quality of Service Policy Information Base,” RFC 3317, March 2003, https://tools.ietf.org/html/rfc3317.

[44] R. Braden, “Requirements for Internet Hosts – Communications Layers,” RFC 1122, Oct. 1989, https://tools.ietf.org/html/rfc1122.

[45] E. Meyer, “Conversations with Stever Crocker (UCLA),” RFC 49, April 25, 1970, https://tools.ietf.org/html/rfc49.

[46] E. Rosen, A. Viswanathan, and R. Callon, “Multiprotocol Label Switching Architecture,” RFC 3031, Jan. 2001, https://tools.ietf.org/html/rfc3031.

[47] D. Grossman and J. Heinanen, “Multiprotocol Encapsulation over ATM Adaptation Layer 5,” RFC 2684, Sept. 1999, https://tools.ietf.org/html/rfc2684.

[48] L. Mamakos, K. Lidl, J. Evarts, D. Carrel, D. Simone, and R. Wheeler, “A method for transmitting PPP over Ethernet (PPPoE),” RFC 2516, Feb. 1999, https://tools.ietf.org/html/rfc2516.

[49] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, “Hypertext Transfer Protocol -- HTTP/1.1,” RFC 2616, June 1999, https://tools.ietf.org/html/rfc2616.

[50] J. Postel and J. Reynolds, “File Transfer Protocol (FTP),” RFC 959, Oct. 1985, https://tools.ietf.org/html/rfc959.

[51] R. Droms, “Dynamic Host Configuration,” RFC 2131, March 1997, https://tools.ietf.org/html/rfc2131.

[52] J. Myers and M. Rose, “Post Office Protocol – Version 3,” RFC 1939, May 1996, https://tools.ietf.org/html/rfc1939.

[53] M. Crispin, “Internet Message Access Protocol – Version 4rev1,” RFC 3501, March 2003, https://tools.ietf.org/html/rfc3501.

[54] J. Klensin, “Simple Mail Transfer Protocol,” RFC 5321, Oct. 2008, https://tools.ietf.org/html/rfc5321.

[55] M. Bagnulo, P. Mathews, and I. van Biejnum, “Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers,” RFC 6146, April 2011, https://tools.ietf.org/html/rfc6146

[56] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, “SIP: Session Initiation Protocol,” RFC 3261, June 2002, https://tools.ietf.org/html/rfc3261.

[57] Michael Hibberd, “Encryption: Will it be the death of DPI?”, Feb. 13, 2012, Telecoms.com, http://telecoms.com/39718/encryption-will-it-be-the-death-of-dpi/ (last visited Sep. 3, 2015).

[58] A. W. Moore and K. Papagiannaki, “Toward the Accurate Identification of Network Applications,” Passive and Active Measurement Conference (PAM), 2001, https://www.cl.cam.ac.uk/~awm22/publications/moore2005toward.pdf

[59] W. John and S. Tafvelin, “Heuristics to Classify Internet Backbone Traffic based on Connection Patterns,” http://www.sjalander.com/wolfgang/publications/John_ICOIN08_CR.pdf.

[60] C. Wright, F. Monrose, and G. Masson, “On Inferring Application Protocol Behaviors in Encrypted Network Traffic,” Journal of Machine Learning Research (JMLR): Special issue on Machine Learning for Computer Security, volume 7, pp. 2745-2769, 2006, available at http://www.jmlr.org/papers/volume7/wright06a/wright06a.pdf.

[61] A. Seewald and W. Gansterer, “On the detection and identification of botnets,” Feb. 2010, http://www.researchgate.net/profile/Wilfried_Gansterer/publication/220614106_On_the_detection_and_identification_of_botnets/links/0deec518cacaeb6574000000.pdf.

[62] C. Brown and A Malis, “Multiprotocol Interconnect over Frame Relay,” RFC 2427, September 1998, https://tools.ietf.org/html/rfc2427.

[63] C. Hornig, “A Standard for the Transmission of IP Datagrams over Ethernet Networks,” RFC 894, April 1984, https://tools.ietf.org/html/rfc894.

[64] IEEE Standards Association, IEEE Standard for Local and metropolitan area networks – Bridges and Bridged Networks, IEEE Std 802.1 Q-2014, http://standards.ieee.org/getieee802/download/802-1Q-2014.pdf.

[65] R.G. Miller, Priority Queues, The Annals of Mathematical Statistics, v31, number 1 (1960), pp. 86-103, *available at* http://projecteuclid.org/download/pdf_1/euclid.aoms/1177705990.

[66] P. J. Rasch, “A Queueing Theory Study of Round-Robin Scheduling of Time-Shared Computer Systems, Journal of the ACM, v17 Issue 1, Jan. 1970, pp. 131-145, *available at* http://dl.acm.org/citation.cfm?id=321569?.

[67] A.K., Parekh and R.G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: The single-node case," *IEEE/ACM Transactions on Networking* 1993, v1, issue 3, pp. 344-357,

available at http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=234856.

[68] W.H. Hesselink and R. M. Tol, "Formal Feasibility Conditions for Earliest Deadline First Scheduling, 1994, http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.55.8609.

[69] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services," RFC 2475, Dec. 1998, https://tools.ietf.org/html/rfc2475.

[70] CeroWRT, CeroWRT Overview, http://www.bufferbloat.net/projects/cerowrt (last visited on Sept. 3, 2015).

[71] G.R. Ash, Traffic Engineering and QoS Optimization of Integrated Voice & Data Networks, Appendix A.4.1.1 and A4.1.2 Leaky Bucket Algorithm and Token Bucket Algorithm, p. 417-418, Nov. 3, 2006, Morgan Kaufman

[72] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSerVation Protocol (RSVP)," RFC2205, Sept. 1997, https://tools.ietf.org/html/rfc2205.

[73] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," RFC 1633, June 1994, https://tools.ietf.org/html/rfc1633.

[74] DSL Forum, TR-144: Broadband Multi-Service Architecture & Framework Requirements, Technical Report, August 2007, https://www.broadband-forum.org/technical/download/TR-144.pdf.

[75] Broadband Forum, TR-145: Multi-service Broadband Network Functional Modules and Architecture, Technical Report, November 2012, https://www.broadband-forum.org/technical/download/TR-145.pdf.

[76] AT&T, Broadband Information: Information about the Network Practices, Performance Characteristics & Commercial Terms of AT&T's Mass Market Broadband Internet Access Services, http://www.att.com/gen/public-affairs?pid=20879 (last visited Sept. 3, 2015).

[77] International Telecommunications Union, Gigabit-capable passive optical networks (G-PON): Transmission convergence layer specification, Recommendation G.984.3, January 2014, https://www.itu.int/rec/T-REC-G.984.3-201401-I/en.

[78] CableLabs, Specifications Library, http://www.cablelabs.com/specs/ (last visited Sept. 3, 2015).

[79] CableLabs, Data-over-Cable Service Interface Specifications DOCSIS 3.0, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv30.0-I28-150827, *available at* http://www.cablelabs.com/wp-content/uploads/specdocs/CM-SP-MULPIv3.0-I28-150827.pdf.

[80] HughesNet, About HughesNet, http://www.hughesnet.com//company (last visited Sept. 3, 2015).

[81] DISH, About Dish: Company Info, http://about.dish.com/company-info (last visited Sept. 3, 2015).

[82] Exede, What is Exede, http://www.exede.com/what-is-exede/, (last visited Sept. 3, 2015).

[83] Bruce R. Elbert, The Satellite Communications Applications Handbook, Artech House, pg. 24, Jan. 1, 2004.

[84] European Telecommunications Standards Institute, Digital Video Broadcasting (DVB); Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Services, News Gathering and other broadband satellite applications (DVB-S2), ETSI EN 302 307 V1.2.1 (2009-08), http://www.etsi.org/deliver/etsi_en/302300_302399/302307/01.02.01_60/en_302307v010201p.pdf.

[85] J. Border, M. Kojo, J. Griner, G. Montenegro, Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations," RFC 3135, June 2001, https://tools.ietf.org/html/rfc3135.

[86] H. Balakrishan, V.N. Padmanabhan, G. Fairhurst, and M. Sooriyabandara, "TCP Performance Implications of Network Path Asymmetry," RFC3449, Dec. 2002, https://tools.ietf.org/html/rfc3449.

[87] The 3rd Generation Partnership Project (3GPP), About 3GPP, 2015, http://www.3gpp.org/about-3gpp/about-3gpp (last visited Sept. 3, 2015).

[88] Next Generation Mobile Network (NMGN) Alliance, Next Generation Mobile Networks: Beyond HSPA & EVDO, White Paper, December 2006, http://ngmn.org/uploads/media/Next_Generation_Mobile_Networks_Beyond_HSPA_EVDO_web.pdf.

[89] F. Firmin, The Evolved Packet Core, 3GPP, http://www.3gpp.org/technologies/keywords-acronyms/100-the-evolved-packet-core (last visited Sept. 3, 2015)

[90] M. Nohrborg, LTE Overview, 3GPP, http://www.3gpp.org/technologies/keywords-acronyms/98-lte (last visited Sept. 3, 2015).

[91] Christopher Cox, An Introduction to LTE: LTE, LTE-Advanced, SAE and 4G Mobile Communications, pp. 36-

37, John Wiley& Sons 2012.

[92] 3rd Generation Partnership Project (3GPP), Policy and Charging Control Architecture, 3GPP TS 23.203, http://www.3gpp.org/DynaReport/23203.htm.

[93] 3GPP, ETSI policy and charging control architecture — 3GPP TS23.203 v.12.9.0 release 12, http://www.etsi.org/deliver/etsi_ts/123200_123299/123203/12.09.00_60/ts_123203v120900p.pdf.

[94] Metro Ethernet Forum, MEF 22.1: Mobile Backhaul Phase 2, January 2012, http://www.mef.net/PDF_Documents/technical-specifications/MEF_22.1.pdf.

[95] Broadband Forum, TR-221: Technical Specifications for MPLS in Mobile Backhaul Networks, October 2011, https://www.broadband-forum.org/technical/download/TR-221.pdf.

[96] Mimosa, Whisper Internet, http://mimosa.co/home/Products/docs/case-studies/wisper-isp.html (last visited Sept. 3, 2015).

[97] Ubiquiti Networks, AirFiber 24: 24 GHz Point-to-Point Gigabit Radio, https://www.ubnt.com/airfiber/airfiber24/ (last visited Sept. 3, 2015).

[98] Reuters, Procera Deployed by Multiple Rural Wireless Internet Service Providers, Press Release, Oct. 22, 2012, *available at* http://www.reuters.com/article/2012/10/22/idUS75943+22-Oct-2012+MW20121022.

[99] Ubiquiti Networks, Rocket AC Base Station, Datasheet, *available at* http://dl.ubnt.com/datasheets/RocketAC/Rocket5ac_DS.pdf.

[100] Juniper.net, Hierarchical Rate Limits Overview, https://www.juniper.net/techpubs/en_US/junose10.3/information-products/topic-collections/policy-management/policy-mgm-rate-limit-hierarchical.html (last visited Sept. 3, 2015).

[101] AT&T, AT&T Wifi and Hotspots, https://www.wireless.att.com/businesscenter/solutions/wireless-laptop/wifi-hotspots.jsp (last visited Sept. 3, 2015).

[102] Cable Wifi (Bright House, Cox, Optimum, Time Warner Cable, Comcast), http://www.cablewifi.com/ (last visited Sept. 3, 2015).

[103] European Telecommunications Standards Institute (ETSI), Network Functions Virtualisation: NFV in ETSI, http://www.etsi.org/technologies-clusters/technologies/nfv (last visited on Sept. 3, 2015).

[104] International Telecommunications Union, Standardization Sector (ITU-T), "One-way transmission time," Recommendation G.114, May 2003, https://www.itu.int/rec/T-REC-G.114/en.

[105] 3GPP, IP Multimedia subsystem (stage 2), Specification detail, http://www.3gpp.org/DynaReport/23228.htm (last visited Sept. 3, 2015).

[106] CableLabs, PacketCable Technical Report, Multimedia Architecture Framework, Oct. 29, 2009, http://www.cablelabs.com/wp-content/uploads/specdocs/PKT-TR-MM-ARCH-V03-091029.pdf.

[107] R. Pantos, "HTTP Live Streaming," April 15, 2015, draft-pantos-http-live-streaming-16, https://datatracker.ietf.org/doc/draft-pantos-http-live-streaming/.

[108] International Organization for Standardization (ISO), Adaptive Streaming over HTTP (DASH), ISO/IEC 23009-1:2014, https://www.iso.org/obp/ui/#iso:std:iso-iec:23009:-1:ed-2:v1:en.

[109] Apple, HTTP Livestreaming Overview, https://developer.apple.com/library/ios/documentation/NetworkingInternet/Conceptual/StreamingMediaGuide/Introduction/Introduction.html (last visited Sept. 3, 2015).

[110] Tellabs, Mobile Video Optimization Concept and Benefits, White Paper, 2011, http://s3.amazonaws.com/zanran_storage/www.tellabs.com/ContentPages/2438991029.pdf.

[111] C. Bastien, T. Klieber, J. Livingood, J. Mills, and R. Woundy, "Comcast's Protocol-Agnostic Congestion Management System," RFC 6057, December 2010, https://tools.ietf.org/html/rfc6057.

[112] D.J. Leith, P. Clifford, "TCP Fairness in 802.11e WLANs," http://www.hamilton.ie/peterc/downloads/wicom05tcpfairness.pdf (last visited Sept. 3, 2015).

[113] Z. Sun, Satellite Networking: Principles and Protocols, Dec. 13, 2005, Wiley and Sons.

[114] Trendnet.com, AC3200 Tri Band Wireless Router, http://www.trendnet.com/products/proddetail.asp?prod=100_TEW-828DRU (last visited Sept. 3, 2015).

[115] S. Mascolo, C. Casetti, M. Gerla, M.Y. Sanadidi, and R. Wang, TCP Westwood: Bandwidth Estimation for

Enhanced Transport over Wireless Links, ACM Sigmobile July 2001, pp. 287-297, http://cpham.perso.univ-pau.fr/TCP/2001-mobicom-0.pdf.

[116] C. Casetti, M. Gerla, S. Mascolo, M.Y. Sanadidi, and R. Wang, TCP Westwood: End-to-End Congestion Control for Wired/Wireless Networks, Wireless Networks v. 8, 2002, pp. 467-479, http://netlab.cs.ucla.edu/publication/download/41/Wn_02.pdf.

[117] Brown University, Participatory Networking: A user-level API for SDNs, http://pane.cs.brown.edu/ (last visited Sept. 3, 2015).

[118] W. Kim, P. Sharma, J. Lee, S. Banerjee, J. Tourrilhes, S. Lee, and P. Yalagandula, "Automated and Scalable QoS Control for Network Convergence," 2010, *available at* https://www.usenix.org/event/inmwren10/tech/full_papers/Kim.pdf.

[119] Georgia Tech, FlowQoS: Providing Per-Flow Quality of Service for Broadband Access Networks, http://flowqos.noise.gatech.edu/ (last visited Sept. 3, 2015).

[120] D. Eastlake, "Transport Layer Security (TLS) Extensions: Extension Definitions," RFC6066, January 2011, *available at* https://tools.ietf.org/html/rfc6066.

[121] Finley, Klint, "Encrypted Web Traffic More than Doubles After NSA Revelations," Wired (May 16, 2014), *available at* http://www.wired.com/2014/05/sandvine-report/.

[122] P. Lepeska, "Trusted Proxy and the Cost of Bits," IETF 90 HTTP WG presentation, *available at* https://www.ietf.org/proceedings/90/slides/slides-90-httpbis-6.pdf (last visited Sept. 3, 2015).

[123] A. Kingsley-Hughes, "Gogo in-flight Wi-Fi serving spoofed SSL certificates," ZDNet January 5, 2015, http://www.zdnet.com/article/gogo-in-flight-wi-fi-serving-spoofed-ssl-certificates/ (last visited Sept. 3, 2015).

[124] Broadband Internet Technical Advisory Group (BITAG), VoIP Impairment, Failure, and Restrictions, May 2014, *available at* http://www.bitag.org/documents/BITAG_-_VoIP_Impairment,_Failure,_and_Restrictions_Report.pdf.

[125] Broadband Internet Technical Advisory Group (BITAG), Port Blocking, Aug. 2013, *available at* http://www.bitag.org/documents/Port-Blocking.pdf.

[126] Broadband Internet Technical Advisory Group (BITAG), SNMP Reflected Amplification DDoS Attack Mitigation, Aug. 2012, *available at* http://www.bitag.org/documents/SNMP-Reflected-Amplification-DDoS-Attack-Mitigation.pdf.

[127] Broadband Internet Technical Advisory Group (BITAG), Implications of Large Scale Network Address Translation (NAT), Mar. 2012, *available at* http://www.bitag.org/documents/BITAG_TWG_Report-Large_Scale_NAT.pdf.

[128] Broadband Internet Technical Advisory Group (BITAG), IPv6 AAAA DNS Whitelisting, Sept. 2011, *available at* http://www.bitag.org/documents/BITAG_TWG_Report-DNS_Whitelisting.pdf.

9 Glossary of terms

All definitions of terms are solely for the purposes of this report, and many are adapted from publications of the Internet Engineering Task Force (www.ietf.org). Readers should be aware that a number of terms have alternate definitions, particularly when used in different or non-networking contexts.

Admission control: Decisions that determine which flows are allowed to begin, based on available resources.

Application: A program that originates or receives data.

Bottleneck: The link or node in a network path where demand is highest relative to capacity.

Bursty: A traffic flow is bursty if its volume changes rapidly over time.

Capacity: The capacity of a link is the number of bits per second that it can transmit. The capacity of a router is the number of packets or bytes per second that it can transmit.

Classification: The categorization of traffic based on identifying characteristics.

Congestion: The effect upon network performance during time periods in which instantaneous demand exceeds capacity

Deadline scheduling: A scheduling algorithm in which each packet is marked on ingress with a maximum period of time it may be delayed (a "deadline"), and the queuing system either ensures the packet's transmission within that interval or drops the packet.

Demand: The volume of traffic that is presented to a link at a given point in time, typically measured in bits per second.

Differentiated treatment, Differentiation: The application of a traffic policy based on classification of the traffic to a class of traffic such as a session, aggregate of sessions, or other traffic flow. Examples of such policies include resource allocation, scheduling algorithms, drop precedence, and routing.

Differentiated Services Architecture (also Diffserv): An architecture for the differentiated handling of IP traffic. Primarily described in RFC 2474 and RFC 2475.

Drop precedence: Examples of this include Frame Relay Discard Eligibility, ATM Cell Loss Priority, the IP "Assured Forwarding" service, or a "less than best effort" service. In such cases, included in the service level agreement is an understanding that a specified subset of traffic (that exceeding some rate, or perhaps all of it) is more likely to be dropped than other traffic.

Flow: A group of packets that share a common set of properties.

Internet Access Service: In this document we use the term 'internet access service' to refer to a mass market service that provides the ability to transmit and receive data to the global address space associated with the Internet.

Mark: To set specific bits of a packet header (see marking) to specific values.

Marking: Specific bits of a packet header that indicate the classification of a packet or that indicate congestion.

Network Operator: A business that operates one or more communications networks. The networks may include access networks that interconnect directly with retail customers, other network types such as transport networks or content delivery networks, or all of the above. In the case of access networks, the network operator may also act as an ISP, offering Internet services to customers. The network operator may also act as an NSP, offering transport or other network services to other network operators.

Over-the-top (OTT): An application or traffic flow that is carried over an Internet access service.

Policing: The dropping or reclassification of packets that exceed the maximum capacity allocated to a flow.

Prioritization: The application of a traffic policy that prefers one or more classes of traffic over one or more other classes of traffic. This may be in actual traffic sequence, in drop probability, or other mechanisms.

Quality of Experience (QoE): The subjective quality of a networked application as perceived by the user.

Quality of Service (QoS): the amount of impairment experienced by traffic during its transmission across one or more networks. QoS is expressed in terms of delay, delay variation, packet loss, and other objective metrics.

Scheduling: The reordering or other treatment of packets according to an algorithm.

Service Level Agreement (SLA): A contractual agreement between network operators or between users and network operators that delineates aspects of the service, often including the upstream and downstream bit rates at the boundary between the operators' networks, the maximum delay across an operator's network, sometimes the maximum proportion of packets to be dropped or other QoS characteristics, and sometimes specifications of payments.

Traffic shaping: Rate limiting of flows in a network.

Packet: A formatted unit of data carried by a packet-switched network. A packet consists of one or more headers containing control information and a payload containing user data.

10 Document Contributors and Reviewers

- Fred Baker, *Cisco*
- Steven Bauer, *Massachusetts Institute of Technology (MIT)*
- Richard Bennett, *American Enterprise Institute*
- Don Bowman, *Sandvine*
- Lily Chen, *Verizon*
- kc claffy, *University of California, San Diego; Center for Applied Internet Data Analysis (CAIDA)*
- David Clark, *Massachusetts Institute of Technology (MIT)*
- David Cooper, *Level 3*
- Amogh Dhamdhere, *University of California, San Diego; Center for Applied Internet Data Analysis (CAIDA)*
- Amie Elcan, *CenturyLink*
- Michael Fargano, *CenturyLink*
- Nick Feamster, *Princeton University*
- Joseph Lorenzo Hall, *Center for Democracy & Technology (CDT)*
- Kevin Kleinsmith, *Union Wireless*
- Ken Ko, *ADTRAN*
- Gary Langille, *EchoStar*
- Matt Larsen, *Vistabeam*
- Jason Livingood, *Comcast*
- Patrick McManus, *Mozilla*
- Chris Morrow, *Google*
- Barbara Stark, *AT&T*
- Matthew Tooley, *National Cable and Telecommunications Association (NCTA)*
- Jason Weil, *Time Warner Cable*
- Greg White, *CableLabs*
- David Winner, *Charter Communications*

*BITAG would like to thank Professor Marvin Sirbu of Carnegie Mellon University for his presentation regarding the structure and composition of different types of access networks.

11 Appendix: Standards, Standards Organizations, and Industry References

A variety of standards organizations and industry alliances have published standards, technical references, informational documents, and best practices that are relevant to the topic of differentiated treatment. Some of the most prominent and influential of these organizations and their publications are listed in this section. In addition, some corporations have created technologies that are widely used and are sometimes considered "de facto" standards. A list of some of these is also included.

- **Internet Engineering Task Force (IETF)**: The IETF produces technical documents that influence the way people design, use, and manage the Internet (www.ietf.org). Relevant publications include:
    - *Internet Protocol* [RFC 791]: Defines IPv4 header formats. The various elements of the IPv4 header (e.g., source IPv4 address, destination IPv4 address, type of service) can be used for classification of traffic (xref to 5.1.1).
    - *Internet Protocol, Version 6* (IPv6) Specification [RFC 2460]: Defines IPv6 header formats. The various elements of the IPv6 header (e.g., source IPv6 address, destination IPv6 address, traffic class, flow label) can be used for classification of traffic (xref to 5.1.1).
    - *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers* [RFC 2474]: Defines the layout of the IPv4 type of service header field and IPv6 traffic class header field, and a base set of packet forwarding treatments, or per-hop behaviors, that can be used for scheduling (xref to 5.1.2).
    - Additional per-hop behaviors, including *Assured Forwarding* [RFC 2597] and *Expedited Forwarding* [RFC 3246/3247], that can be used for scheduling (xref to 5.1.2).
    - C*omcast's Protocol-Agnostic Congestion Management System* [RFC 6057]: Describes Comcast's congestion management system that was deployed December 31, 2008. This is used for scheduling traffic (xref to 5.1.2), but note that it this is not a standard.
    - *Multiprotocol Label Switching Architecture* [RFC 3031]: Describes the MPLS architecture commonly used in current DSL and PON access networks, beyond the first mile, and is used for service level agreement (SLA) assurance in service provider networks.
    - *Resource ReSerVation Protocol* (RSVP, RFC 2205): a protocol designed for soft reservation of bandwidth within a network. Reservation is "soft" in the sense that while it ensures bandwidth is available when the subject traffic flow is present, it is available for other traffic when the subject traffic flow is not or does not use it all. RSVP is also used for the management of SLA-sensitive MPLS LSPs.
    - *Integrated Services in the Internet Architecture* [(RFC 1633]): the architecture that defines Real Time vs. Elastic applications, and their requirements.

- **Institute of Electrical and Electronics Engineers - Standards Association** (IEEE-SA): The IEEE-SA drives the functionality, capabilities and interoperability of a wide range of products and services (standards.ieee.org). Some of the most relevant (to differentiation) working groups inside IEEE-SA include 802.1 (Higher Layer LAN Protocols Working Group), Ethernet Working Group (802.3), 802.11 (Wireless LAN Working Group), and Broadband Wireless Access Working Group (802.16). There are many other working groups that include efforts on such topics as resilient packet rings, wireless coexistence, mobile broadband, and powerline networking (not an 802 working group).
    - *IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture* [IEEE 802]: Contains descriptions of the IEEE 802® standards published by the IEEE for frame-based data networks as well as a reference model (RM) for protocol standards.
    - *IEEE Standard for Ethernet* [IEEE 802.3]: Relevant parts define Ethernet, including physical layer protocols (e.g., 10BASE-, 100BASE- [Fast Ethernet], 1000BASE- [Gigabit Ethernet], Ethernet in the First Mile [EFM]), and the format of the Media Access Control (MAC) frame, commonly called the Ethernet frame. The various elements of the MAC frame (includes

source and destination MAC addresses, Ethertype, and extensions defined in other IEEE standards) can be used for classification of traffic (xref to 5.1.1).
    - *IEEE Standard for Local and Metropolitan Area Networks – Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks* [IEEE 802.1Q]: Relevant aspects of this standard are that it defines VLAN Tag extensions, which include the VLAN ID and Priority Code Point (PCP). Both of these fields are useful for classification (xref to 5.1.1), with the Ethernet Priority used primarily to determine scheduling behavior (xref to 5.1.2) and the VLAN ID used primarily to determine routing (xref to 5.1.3).
    - *IEEE Standard for Local and Metropolitan Area Networks –v Media Access Control (MAC) Bridges* [IEEE 802.1D]: Relevant aspects of this standard describe bridging behavior based on the Priority Code Point in the VLAN Tag.
    - *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications* [IEEE 802.11]: Relevant aspects of this standard are that it defines several physical layer protocols commonly referred to as "802.11a," "802.11b," "802.11g," and "802.11n," and defines the MAC frames and QoS mechanisms for these physical layers.
    - *IEEE Standard for Service Interoperability in Ethernet Passive Optical Networks (SIEPON)* [P1904.1]: This standard describes the system-level requirements needed to ensure service-level, multi-vendor interoperability of Ethernet Passive Optical Network (EPON) equipment. The specifications complement the existing IEEE 802.3 and IEEE 802.1 standards, which ensure the interoperability at the Physical Layer (PHY) and Data Link Layer. Included in this specification are:
        - EPON system-level interoperability specifications covering equipment functionality, traffic engineering, and service-level quality of service/class of service (QoS/CoS) mechanisms;
        - Management specifications covering equipment management, service management, and EPON power-saving mechanism.

- **Broadband Forum (BBF)**: The BBF develops multi-service broadband networking specifications addressing interoperability, architecture and management (www.broadband-forum.org). This organization has a strong focus on DSL and PON ISP access networks and providers. The former ATM Forum, IP/MPLS Forum, MFA Forum and MPLS & Frame Relay Alliance organizations have all merged with BBF, and all of these organizations' publications are now available from the BBF website.
    - *Migration to Ethernet-Based Broadband Aggregation* [TR-101 Issue 2]: Describes an Ethernet (at the MAC layer) architecture for DSL and PON access networks. This describes an access and aggregation network architecture (xref 5.1.3) including basic methods to support classification, scheduling and other differentiation techniques.
    - *Multi-service Broadband Network Functional Modules and Architecture* [TR-145]: Extends the TR-101 architecture to support multiple services and network interfaces.
    - *Using GPON Access in the context of TR-101* [TR-156]: Extends TR-101 to GPON.
    - *Using EPON in the Context of TR-101* [TR-200]: Extends TR-101 to EPON.
    - *Multiprotocol Label Switching (MPLS) standards*: MPLS was originally defined by IETF, but much subsequent work was done in the IP/MPLS Forum, which has since merged with BBF.

- **International Telecommunications Union - Telecommunication Standardization Sector (ITU-T)**: ITU-T is the telecommunications standards sector of the ITU, which is the United Nations specialized agency for information and communication technologies.
    - *Asymmetric digital subscriber line (ADSL) transceivers* ADSL: [ITU-T G.992.1], *Asymmetric digital subscriber line transceivers 2 (ADSL2)* [ITU-T G.992.3], *Very high speed digital subscriber line transceivers (VDSL)* [ITU-T G.993.1], *Very high speed digital subscriber line transceivers 2 (VDSL2)* [ITU-T G.993.2]: Define the various DSL physical layer technologies over copper twisted pair (phone lines).
    - *Broadband Passive Optical Network (BPON) for telecommunications Access networks* [ITU-T G.983],

    - *Gigabit-capable Passive Optical Networks (GPON)* [ITU-T G.984]: Define the BPON and GPON physical layer technologies
    - J Series: ITU-T officially named the CableLabs' DOCSIS specifications as standards through several of its J Series publications. DOCSIS 1.0 was J.112 Annex B (1998), DOCSIS 1.1 is J.112 Annex B (2001), DOCSIS 2.0 is J.122, and DOCSIS 3.0 is J.222.

- **CableLabs:**
    - *DOCSIS 1.0* - A set of eight specifications that define the first version of the Data Over Cable Service Interface Specification (DOCSIS). These specifications describe requirements for Cable Modems (CMs) and Cable Modem Termination Systems (CMTSs) to provide basic broadband IP connectivity over the hybrid fiber coax cable network, including per modem rate shaping, link encryption, and network management functions.
    - *DOCSIS 1.1* - A set of three specifications that extend DOCSIS 1.0 in order to provide Quality of Service controls, a Public Key Infrastructure based CM authentication and CM firmware validation mechanism, and enhanced network management tools.
    - *DOCSIS 2.0* - A set of three specifications that extend DOCSIS 1.1 with a new upstream physical layer technology that provides a 3x increase in upstream channel capacity, and optional support of IPv6 management.
    - *DOCSIS 3.0* - A set of five specifications that extend DOCSIS 2.0 with support for channel bonding to increase capacity to over 200 Mbps in the upstream and over 1 Gbps in the downstream, and full IPv6 support.
    - *DOCSIS 3.1* - A set of five specifications that extend DOCSIS 3.0 with a new upstream and downstream physical layer enabling multi-Gbps service in both directions, as well as Active Queue Management and hierarchical Quality of Service.
    - *DOCSIS Provisioning of EPON, version 1.0* (DPoEv1.0) – A set of nine specifications describing the translation of DOCSIS provisioning procedures to provision and manage Ethernet Passive Optical Networks via IPv4. A comprehensive set of extended OAM messages is defined to allow the Optical Line Terminal (OLT) to provision and manage Optical Network Units (ONUs). The specifications also describe the device requirements for supporting IP high speed data and Ethernet Private Line (EPL) Metro Ethernet service models in an MSO environment.
    - *DOCSIS Provisioning of EPON, version 2.0* (DPoEv2.0) – A collection of nine specifications that extend DPoEv1.0 specifications, includes management using IPv6, multicast services, and network synchronization. The support of commercial services is expanded to include Virtual Private LAN and Tree services, as defined by the Metro Ethernet Forum.

- **Wi-Fi Alliance (WFA):** WFA is an industry alliance that has created IEEE 802.11 implementation profiles and developed certification programs for the same. "Wi-Fi" is its registered trademark brand.
    - Wi-Fi CERTIFIED™ n in both 2.4 and 5 GHz, and Wi-Fi CERTIFIED™ ac in 5 GHz: Certification programs and implementation profiles offered by WFA for physical layer wireless technologies based on the IEEE 802.11 standard. The former Wi-Fi CERTIFIED a in 5 GHz and Wi-Fi CERTIFIED b/g in 2.4 GHz are no longer available.
    - Passpoint™: Certification program and specification (developed with GSMA and the Wireless Broadband Alliance) that enables SIM and non-SIM mobile devices to discover, select and connect to Wi-Fi networks without user intervention. Also known as Hotspot 2.0.
    - •WMM® (Wi-Fi Multimedia™): Certification program and implementation profile for IEEE 802.11 Quality of Service (QoS) mechanisms.

- **Widely-used Proprietary Technologies** (primarily Cisco, Juniper, Motorola, and Alcatel-Lucent)
    - Cisco's Netflow, Juniper Sflow
    - Motorola Canopy