Broadband Internet Technical Advisory Group

1550 Larimer Street, Suite 168
Denver, CO. 80202

BITAG Publishes Report: Differentiated Treatment of Internet Traffic

Denver, CO (October 8, 2015): Today, the Broadband Internet Technical Advisory Group ("BITAG") announced the publication of its technical report on the subject of Differentiated Treatment of Internet Traffic. The executive summary of the report is attached here – including the full set of observations and recommendations – and the report itself can be found at: http://www.bitag.org/documents/BITAG_-_Differentiated_Treatment_of_Internet_Traffic.pdf.

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. 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 data transmitted across the networks that make up the Internet is formatted as packets, which contain information payloads encapsulated within one or more headers. These headers provide the information needed to deliver the packets to their destinations. As these packets travel across networks, they contend with other packets for network resources. The simplest way to handle this contention 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. This is done for various reasons, including 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. The ability to treat traffic differentially has been built into Internet protocols from the beginning but has not been deployed end-to-end due to a number of issues.

Observations. From the analysis made in the 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* (See attached Executive Summary or full report for the complete explanation of each):

- TCP causes recurring momentary congestion.
- A nominal level of packet discard is normal.

- The absence of differentiation does not imply comparable behavior among applications.
- Differentiated treatment can produce a net improvement in Quality of Experience (QoE).
- Access technologies differ in their capabilities and characteristics.
- Security of traffic has at times been downgraded to facilitate differentiation techniques.

Recommendations. The BITAG Technical Working Group also has the following *recommendations* (See attached Executive Summary or full report for the complete explanation of each):

- Network operators should disclose information on differential treatment of traffic.
- Network operators and ASPs should be encouraged to implement efficient and adaptive network resource management practices.
- Quality of Service metrics should be interpreted in the context of Quality of Experience.
- Network operators should not downgrade, interfere with, or block user-selected security in order to apply differentiated treatment.

Ken Ko, Senior Staff Scientist at ADTRAN, and Fred Baker, a Fellow at Cisco, served as the lead editors of the report. Douglas Sicker, Executive Director of BITAG, Chair of BITAG's Technical Working Group, Department Head of Engineering and Public Policy and a professor of Computer Science at Carnegie Mellon University, chaired the review itself.

About BITAG. 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.

This is BITAG's eighth report. BITAG's previous reports have focused on: Internet interconnection; VoIP impairment, failure, and restrictions; Real-time network management of Internet congestion; Port blocking; SNMP DDoS attack mitigation; Large scale network address translation; and IPv6 whitelisting. Copies of these technical reports can be found on the BITAG website at www.bitag.org.

Questions or Comments? BITAG welcomes any questions, comments or suggestions. Please contact our Executive Director, Douglas Sicker, at dsicker@bitag.org or our Deputy Director, Kaleb Sieh, at ksieh@bitag.org.

– ATTACHMENT –

Executive Summary of BITAG Report on Differentiated Treatment of Internet Traffic

Full Report: http://www.bitag.org/documents/BITAG_-_Differentiated_Treatment_of_Internet_Traffic.pdf

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.

Recommendations. 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.