

MR-404

Traffic Management in Multi-Service Access Networks

Issue: 01
Issue Date: September 2017

Issue History

Issue Number	Approval date	Publication Date	Issue Editor	Changes
01	4 September 2017	13 October 2017	Christele Bouchat, Nokia	Original

Comments or questions about this Broadband Forum Marketing Draft should be directed to help@broadband-forum.org

Editors	Francois Fredricx	Nokia	Francois.fredricx@nokia.com
	Florian Damas	Nokia	Florian.damas@nokia.com
	Ing-Jyh Tsang	Nokia	Ing-jyh.tsang@nokia.com
Innovation Leadership	Christele Bouchat	Nokia	Christele.bouchat@nokia.com
	Mauro Tilocca	Telecom Italia	mauro.tilocca@telecomitalia.it

Executive Summary

Traffic Management is a widespread industry practice for ensuring that networks operate efficiently, including mechanisms such as queueing, routing, restricting or rationing certain traffic on a network, and/or giving priority to some types of traffic under certain network conditions, or at all times. The goal is to minimize the impact of congestion in networks on the traffic's Quality of Service. It can be used to achieve certain performance goals, and its careful application can ultimately improve the quality of an end user's experience in a technically and economically sound way, without detracting from the experience of others.

Several Traffic Management mechanisms are vital for a functioning Internet carrying all sorts of Over-the-Top (OTT) applications in "Best Effort" mode. Another set of Traffic Management mechanisms is also used in networks involved in a multi-service context, to provide differentiated treatment of various services (e.g. Internet access service carrying any OTT application, business VPNs, specialized VoIP, or video services), where these services share a common infrastructure. This white paper introduces Traffic Management, and specifically describes the rationale and mechanisms for applying some Traffic Management techniques to access networks in the context of multi-service.

There is still on-going research on new Traffic Management mechanisms and techniques. It is an evolving field of study.

The necessity of deploying some set of network management practices is widely recognized by government agencies around the world. For example, the United States FCC order on Protecting and Promoting the Open Internet states "*Reasonable network management shall not be considered a violation of [the rule that end users' and edge providers' abilities to use the Internet shall not be unreasonably interfered with or disadvantaged]*" [1]. This same order also specifically allows the offering of specialized services, which are IP-services that do not travel over broadband Internet access service, such as "*facilities-based VoIP offerings, heart monitors, or energy consumption sensors*" [1]. In Europe, the BEREC considers in the Net Neutrality Guidelines that as long as Traffic Management is done independently of applications and end-users, the traffic is normally considered to be treated equally [2]. Plus, the Open Internet Regulation by the European Parliament and Council allows "*reasonable Traffic Management*" which may be used to differentiate between "*categories of traffic*" (e.g. be defined by reference to application layer protocol or generic application type). [3]

The Broadband Forum is actively working on several topics that linked together would allow Broadband Assured Services for IP services. A description of the Broadband Forum vision can be found at <https://www.broadband-forum.org/about-the-broadband-forum/about-the-forum/broadband-20-20>.

Introduction to Traffic Management

All network resources are limited physically (e.g. by the number and speed of links) and by economic considerations. There has always been the need to optimize the use of equipment and other network resources and this can lead to congestion, especially in large networks. Even in the circuit-switched POTS world, limited switching and multiplexing capacity meant call blocking was possible when there was high call volume. In the IP, packet-switched Internet world, limited link, routing and aggregation capacity can cause throughput limitations, increased latency (packet delay), latency variation (jitter), and packet loss.

Network Operators have the responsibility to keep the network operational. This requires the use of Traffic Management functions. Traffic Management mechanisms are a set of tools that allow a network operator to ensure the continued function of the network during times of congestion at congested nodes. These mechanisms are also useful in supporting Service Level Agreements (SLAs), and the delivery of different types of services. One of the main purposes is to avoid, reduce and/or delay the adverse effects of congestion on the different types of traffic that use the network. Effective Traffic Management is essential to minimize the impact of congestion on real-time video, VoIP, streamed video and even web browsing, which in turn influences the experience of users. Besides managing congestion, it is also useful to identify and monitor congestion, as described in [4].

Traffic Management influences both Quality of Service (QoS) and Quality of Experience (QoE) metrics. As described in [5], QoS is “*a measure of performance at the packet level from the network perspective*”, whereas QoE is “*the overall performance of a system from the point of view of the users. QoE is a measure of end-to-end performance at the services level from the user perspective and an indication of how well the system meets the user’s needs*”. QoS is a measure of throughput, latency, packet delay variation, and packet loss, while QoE is a subjective measure of a user's perception of the performance of a particular service. The relationship between QoS and QoE depends on the type of service. For example, users of streamed video can tolerate some latency, but not packet loss or high variation in latency. Users of real-time voice do not tolerate significant latency or variation in latency, but can tolerate some packet loss. Both types of application have minimum throughput requirements. Users of email have a very high tolerance of latency, latency variation, throughput variation, and packet loss (because the protocols used by email can easily recover lost packets).

There are many Traffic Management mechanisms, such as

- traffic classification
- traffic metering and shaping,
- packet marking and/or dropping
- packet scheduling
- admission control and resource reservation
- routing decisions
- caching

These tools can be combined in various ways to achieve the network provider’s Traffic Management policy. But note that it is not necessary or even beneficial to apply all these tools

together in any portion of the network.

The next sections focus on the use of some differentiated Traffic Management mechanisms in the context of multi-service access networks.

Congestion in access networks

Access networks (the first-mile lines and the nodes terminating these lines) have historically been a bottleneck. However successive waves of technological innovation have allowed a tremendous evolution, from dial-up modems through to the introduction of DSL and Cable modems, and now the ever-deeper FTTx deployments bringing Gigabit access connectivity. A similar increase is happening in the mobile networks with the advent of each new generation. While greatly increased access rates are needed to significantly improve the broadband experience, they do not eliminate all congestion, which can still occur for the following reasons;

- Access operators are not only responsible for providing the mobile or fixed access, but also for multiplexing and aggregating traffic from all user connections higher up in the network. A trade-off needs to be made between aggregation capacity and investment, resulting in an acceptable statistical multiplexing factor, considering expected concurrency, peak and average rates¹. Complete avoidance of congestion through pure capacity growth (dimensioning the whole network for concurrent, continuous peak rate capacity (line rate) for all users) is not economically feasible.
- From a dimensioning point of view, the attempt to completely avoid congestion would raise the question of how much bandwidth allowance is needed per user. This is difficult to answer as new undefined bandwidth-hungry applications can appear at any time. Further, self-similarity of internet traffic has been demonstrated, meaning it is bursty on all timescales; hence avoiding congestion would require network dimensioning for peak rates for all users.
- The dimensioning of provider-to-provider links is important. As an access network is connected to multiple provider networks, it is not realistic to size all these links to support 100% of all user traffic in the access network. If such a provider network sends more traffic than the link is dimensioned for, the link will become congested. Note that providers are aware of the dimensioning of provider-to-provider links and choose on which of those links to send their traffic. If such a provider network sends more traffic than the link is dimensioned for, the link will become congested.
- From a technical perspective, the nature of IP traffic (UDP and TCP) and current control mechanisms mean some congestion will always happen. Every TCP connection, by design, will try to fill the available link capacity, sending more and more packets until congestion is detected, triggering its congestion control mechanisms and leading to a back-off reaction. There are different TCP flavours, and all try to optimize throughput, control congestion, and be fair to other flows. These three goals are directly impacted by the queue length and packet drop/marketing mechanism used in the queue. No matter the queue depth, current

¹ The overprovisioning is only for Non-guaranteed traffic. Guaranteed traffic on the other hand must always be able to count on available end-end capacity.

TCP flavours will always lead to queue filling and congestion.

When TCP is mixed with UDP traffic without any distinction, unbounded UDP traffic (applications without flow control) leads to congestion and packet drops for both UDP and TCP data, which could starve the TCP flows.

- In an access network, there are typically several high-capacity links (e.g. 1, 10 Gbit/s) on the network-side of an access node, which forward the traffic to and from a multitude of lower-capacity user links (e.g. ~ 100 Mbit/s for VDSL). In the downstream direction, the incoming bursty high-speed traffic needs to be buffered onto the lower-speed user lines. The ingress/egress speed mismatch can cause queue filling, even for flows below the individual user line rate.

Applications being carried over access networks have broadened in scope and greatly increased the traffic volume, going from basic Internet access service to triple play to fixed-mobile convergence to cloud applications. Such applications also bring new requirements on QoS. For instance, mobile front-haul can have very stringent latency and jitter requirements on the transport between the distributed and centralized nodes constituting the Baseband Unit. Also, future Next-Generation applications such as Ultra Reliable Low Latency Communication will impose very low end-to-end latency at application level (order of 1ms).

The persistence of congestion and the broadening of the amount of services and their traffic load and QoS expectations illustrate the continuing importance of applying proper Traffic Management techniques in access networks.

Current use of Traffic Management in multi-service access networks

As described in section 4.1 of [6], a 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) or application providers (ASP) across [the provider to provider interface] as IP traffic, or (for services such as Layer 2 business connectivity), from another network provider 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*”

In such networks, it is common practice to apply differentiated treatment to the traffic, by classifying it into different packet streams on a per-user and/or per-application basis, and by applying a prescribed set of actions to the different streams. For an in-depth background on Traffic Management and traffic differentiation, please see [6] and [7].

An example of the structure and functional blocks of a typical FTTx access network is given in Figure 1. It shows the connectivity for a mix of internet and operator-managed services with a typical example of Traffic Management measures used at the various points in the network. This section focuses on the access part of the network (from Access Node to the end-user Residential Gateway). Similar mechanisms can be used in other parts of the network (Border Network Gateway, intermediate aggregation switches). For more information on the architectures of fixed Broadband networks, please see [8], [9], [10], [11], [12], and for more information on the involvement of the Residential Gateway see Appendix A in [13].

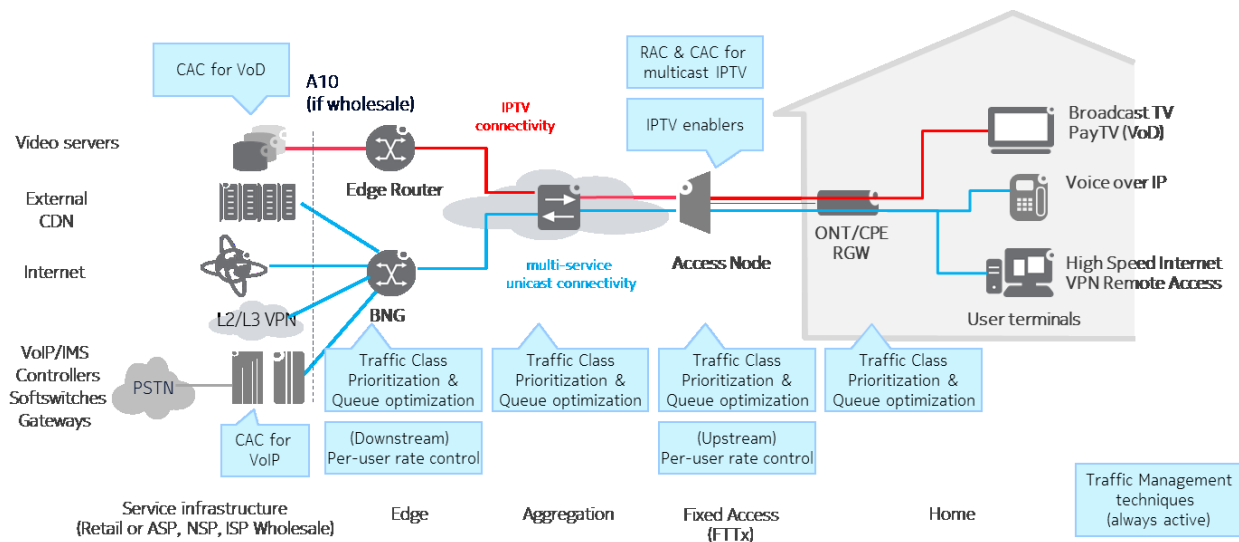


Figure 1 – Typical FTTx Access Network (also shown; aggregation & edge network)

Differentiated Traffic Management in multi-service access networks consists of a classification of the traffic, followed by actions such as active queue management (AQM) with packet drop/remarking, queue scheduling, rate control (shaping, policing), and possibly admission control to the shared resources.

- *Classification* on a per-packet basis is about recognizing certain traffic as being different than other traffic in order to provide it differentiated treatment. It is typically done at the edge of the network, or by the end-user device itself (if that device is trusted by the operator, e.g. the operator’s own Service Box). Classification can be based on a variety of characteristics of the traffic, including origin, destination, transport or application protocol, and the values of specific bits in various protocol headers (e.g., diffserv code point or DSCP in the IP header or priority bits in the Ethernet frame header). Some operators use Traffic Management as part of productizing applications and services.
- Different treatment can be applied to the different traffic classes. *Prioritization* is a practice that causes some traffic to be treated better than other traffic (e.g., packets get sent ahead of other packets, or are not dropped when there is congestion). Different policies can be applied on the different Traffic Classes. The basic mechanisms used for prioritization are queuing and scheduling whereby packets from different traffic classes are served (metered, marked or dropped, and then forwarded by queue shaping and scheduling between queues) in a priority-aware or weighted cyclic way.
- *Rate control* by shaping or policing prevents any particular user from loading the network beyond their commercial service profile (thereby also impacting the service of other users).

- For the same reason, access to the network resources (network capacity, application server capacity) can be controlled for specialized services. *Admission control* can also prevent users from unknowingly degrading their own services by requesting too many instances of applications for their connection rate.

It is important to highlight that congestion management techniques should always be active. It is not sensible to only activate them when congestion is detected, because they can also have a preventive effect. For example, traffic shaping by means of buffering smooths out traffic peaks and avoids hard clipping (policing) of the data; peaks are impossible to predict and hence shaping should always remain enabled.

Finally, note that the network can also improve QoS indirectly by other means;

- By moving the content closer to the end-user. Content Delivery Networks (CDN) help to distribute content globally in a scalable way. In access networks, transparent caches lower the latency, reduce the need for retransmissions, and lower the traffic load at interconnection points. Of course, caching will not address the needs of interactive applications.
- by constructing IP multicast trees as a way of limiting the load in portions of the network compared to an equivalent set of unicast streams, leaving more capacity for other applications and streams, thereby indirectly improving QoS.

Impacts of End-points on end-end QoS

Various mechanisms at the end-points themselves (the peers, or the client and the server) will also impact the end-end QoS and QoE, on top of Traffic Management mechanisms used by the operator in its network elements.

How can we compare the impact of the end-points and the intervening network?

From the *end-point's perspective*, the intermediate network is a black box that can have varying capacity and latency. Although the behavior of the network can be probed by the end-points, it cannot be controlled by them. The end-points can only adapt the way their traffic is sent to the other end by:

- Packet marking (although usually such marking is considered untrusted by operators and overwritten in the network)
- Adaptation of transmission rate to feedback from the other end (e.g. TCP's slow start and congestion avoidance, e.g. Google's QUIC version of UDP)
- Spreading the traffic over multiple parallel TCP flows, which allows the endpoint to use more resources than endpoints trying to use just a single flow, leading to competition between end-points.
- Tweaking the behavior of the server (e.g. running a more aggressive version of TCP at server side at some crucial moments). But this results in competition between flows, some winning and some losing, so it will not lead to a global improvement.

On their own, local measures do influence QoS, but are not sufficient to fully determine the end to end QoS performance.

From the *network perspective*, the local control mechanisms used by the end-points can't be modified, but the network has the great advantage of knowing where congestion happens, and being able to react there. It can apply measures at different levels of granularity (e.g. per-Traffic Class, per-user, per-aggregating link).

Traffic Management in the context of Multi-Service

Applications will differ in terms of their QoS performance requirements. Differentiated treatment of Traffic Classes with prioritization and bandwidth usage control can create a better QoE for sensitive applications without harming the QoE of other applications, provided the lower priority classes are protected from starvation by the higher priority classes.

Congestion still means that those packets that can no longer be served at a given point in time are dropped or delayed. But packets can be dropped without dramatic consequences when done properly, and delays are not always critical. Some applications (such as real-time communication) are very delay sensitive but more loss-tolerant than others, while others (such as video streaming) are intolerant to packet loss but can accommodate larger latencies. This is why it makes sense to group latency-sensitive applications in Traffic Classes with shallow buffers, while grouping more loss-sensitive applications in Traffic Classes with deep buffers. Non-real-time web-based applications carried over TCP are not sensitive to either latency or loss, but large buffering is required to provide fair and steady TCP goodputs to flows of different round trip times (RTT). Without differentiation, any kind of application could be impacted by packet loss or buffer filling. Without bandwidth usage control, some users or applications could be harmed by “misbehaving” other users or applications.

Note that High Speed Internet (HSI) traffic carries any Over-The-Top (OTT) application, from browsing to streaming to real-time communication. HSI is a single, best effort traffic class. Consequently, any OTT application QoS requirements (e.g. video over HTTP) cannot be addressed by the network. This means that even Best Effort HSI traffic should have at least some minimum level of support defined.

Differentiating Traffic Management aims to optimize, per traffic class, the four QoS parameters: throughput, latency, packet delay variation and packet loss². Figure 2 gives an overview of the importance of the quality parameters for HSI, VoIP, IPTV and Mobile Backhaul. It can be noted that although HSI is carried as “Best Effort” there are still some minimum quality expectations (e.g. Internet connectivity must not be completely marginalized by other higher-priority services). With Traffic Management, each Traffic Class can be given a relative priority and the appropriate buffering treatment, allowing co-existence of the different classes on the same infrastructure with limited shared resources.

² Security and availability are more generic expectations related to network design quality aspects such as protection and redundancy, which can impact all services.




Figure 2 – Traffic types and their relevant quality requirements

Differentiated Traffic Management can sometimes provide benefits to customers, for example:

- Prioritization of emergency services over other types of traffic guarantee service availability in order to save lives.
- Adequate queuing and priority scheduling for video traffic (IPTV) reduces packet loss and protects its throughput, and hence can give the possibility to the operator to deliver good quality perception to its customers.
- Business customers typically have strict SLA (Service Level Agreement) with operators, and operators need to meet this SLA without impacting the quality perceived by all the other customers.

As mentioned previously, the cause of congestion and hence QoS degradation is the demand for resources being greater than the resource available, and the techno-economic reality of network dimensioning.

Let us take an example of Traffic differentiation in action. Suppose a managed video conference between two parties and an FTP session via internet access are taking place over the same user connection.

- By applying an overall rate limitation, the total traffic of the user is made to conform to their service subscription rate.
- The operator manages the video conference session. Application layer admission control is applied to verify service subscription.
- The video conference flow is classified as real-time Traffic Class and receives priority over the Best Effort internet traffic. The traffic is buffered in shallow queues, enabling low latency. When congestion occurs on the connection (e.g. due to TCP traffic ramping up), the video conference flow will keep its low latency and throughput, thanks to its shallow queue and its Traffic Class being served by strict priority. Moreover, the traffic volume in

that Traffic Class is controlled (by the managed applications and by rate limitation), to avoid loss due to non-conforming user streams.

- The operator does not manage the FTP session, it is classified as Best Effort, and has lower priority but is served by longer queues, which sustains the TCP throughput. The TCP mechanism will try to use as much of the available bandwidth as possible, up to the point of congestion. When congestion occurs, the TCP session will experience packet drops and back off to a lower rate. As the video call packets receive higher priority, the TCP flow will experience greater congestion when the video call is present and reduce its rate, but as soon as more bandwidth is freed up (e.g. the video call has ended) it will ramp up again to try to fill the newly available capacity.

Future evolutions

The evolution of the telecom industry with novel paradigms requires an ultra-low latency network [14]. Depending on the application, differentiated treatment may or may not be required to give sufficient QoS performance. These new network and service developments include;

- Virtual Reality - this requires very high bandwidth and low latency to avoid an unpleasant user experience.
- The shift to Network Function Virtualization (NFV) and Software Defined Networks (SDN). The SDN/NFV approach opens the possibility for shifting functions from network elements into a cloud-based server NFV infrastructure, and for programmability of the network, by using NETCONF/YANG in the management plane. Concerning Traffic Management, such programmability can cover QoS rules (eg classification rules) and resource sharing or resource reservation between services or virtual operators (eg equipment slicing, or management-based multi-tenancy). The Broadband Forum is actively involved in the definition of YANG models. QoS rules and policies may become more complicated due to this more dynamic environment.
- The “Industrial internet” represents a shift from a human-centric network (with delays on the order of single digit msec) to a machine-compatible network (Internet of Things) with some applications requiring much lower latency (100x less, of the order of μ s).
- Mobile back- and front-hauling will increasingly make use of portions of the same access and aggregation assets as the fixed services. New mobile generations (5G) will pose even more stringent requirements in terms of latency and packet delay variation.
- Eliminating queue delay is a fundamental enabler for such evolutions, however the TCP flavors presently used in the internet need queuing to achieve a steady high link utilization. Novel TCP variants that can operate at quasi zero queue filling have been proposed, such as Data Center TCP (DCTCP). DCTCP can keep queuing delay low without compromising link utilization. But DCTCP is not used in the Internet because it would starve the legacy TCP flavors. Recently a novel AQM which could resolve this problem has been discussed in the IETF (Dual Queue Coupled AQM). Introducing such AQM Traffic Management mechanism in the network nodes would allow the Internet to evolve to support low-latency low-loss TCP service, without compromising the performance of the classic traffic, thus ensuring ‘fairness’ of service, or neutrality in flow performance, irrespective of TCP flavor.

Conclusions

We have explored the value of, and mechanisms for applying some Traffic Management techniques in multi-service access networks. Traffic differentiation is typically used to control the traffic from multiple services at congestion points, with prioritization to balance the QoS needs of the different services, and rate control and access control to manage shared resources in a fair way. Congestion cannot be eliminated as such, due to the limited and shared nature of network resources, and their use by rate-hungry applications and protocols. The benefit of differentiated Traffic Management is to provide Traffic Class specific QoS treatment, which in turn allows sensitive applications to provide their required QoE without blocking other types of applications, even in the event of congestion.

Appropriate Traffic Management is required for both specialized services and HSI, allowing them to co-exist on a shared multi-service network infrastructure. Operators use such Traffic Management to deliver added value services to both residential and business users, including mobile back-hauling (and front-hauling). Although Internet access is generally put in the lowest traffic class in such a scheme, a minimum level of service can be provided by adequate network dimensioning and resource management.

Simply removing access bottlenecks and increasing everyone's line rate will not be the solution without selected Traffic Management techniques. Such techniques are instrumental in optimizing the use of network capacity, but this does not avoid the need for investment in additional capacity when long congestion periods would start to occur.

There is still on-going research work on new Traffic Management mechanisms, so further evolution is expected.

The Broadband Forum is actively working on several topics that linked together would allow Broadband Assured Services for IP services, which includes Traffic Management. A description of the Broadband Forum vision can be found at <https://www.broadband-forum.org/about-the-broadband-forum/about-the-forum/broadband-20-20>.

In conclusion, using some Traffic Management in multi-service access networks is a techno-economic imperative, and it's here to stay for Gigabit access and beyond.

References and Terminology

[1] "Protecting and Promoting the Open Internet", Federal Communications Commission FCC 15-24

[2] BEREC Guidelines on the Implementation by National Regulators of European Net Neutrality Rules BoR (16) 127 of 30 August 2016,

http://berec.europa.eu/eng/document_register/subject_matter/berec/download/0/6160-berec-guidelines-on-theimplementation-b_0.pdf

[3] Regulation (EU) 2015/2120 of the European Parliament and of the Council of 25 November 2015 laying down measures concerning open internet access and amending Directive 2002/22/EC on universal service and users' rights relating to electronic communications networks and services and Regulation (EU) No 531/2012 on roaming on public mobile communications networks within the Union,

<http://eur-lex.europa.eu/legalcontent/EN/TXT/PDF/?uri=CELEX:32015R2120&from=EN>

[4] "Broadband Access Service Attributes and Performance Metrics", Broadband Forum TR-304

[5] "Triple Play Services Quality of Experience (QoE) Requirements", Broadband Forum TR-126

[6] "Differentiated Treatment of Internet Traffic", Broadband Internet Technical Advisory Group Technical Working Group Report – BITAG, October 2015

[7] "Real-time Network Management of Internet Congestion", Broadband Internet Technical Advisory Group Technical Working Group Report – BITAG, October 2013

[8] "Broadband Remote Access Server (BRAS) Requirements Document", Broadband Forum TR-92

[9] "Migration to Ethernet-Based DSL Aggregation", Broadband Forum TR-101

[10] "Using GPON Access in the context of TR-101", Broadband Forum TR-156

[11] "GPON-fed TR-101 Ethernet Access Node", Broadband Forum TR-167, "

[12] "Multi-service Broadband Network Architecture and Nodal Requirements", Broadband Forum TR-178

[13] "Internet Gateway Device Data Model for TR-069", Broadband Forum TR-098

[14] ITU-T Technology Watch Report (August 2014) - The Tactile Internet

Abbreviations

AQM	Active Queue Management
ASP	Application Service Provider
CAC	Call Admission Control
CDN	Content Delivery Networks
CPE	Customer Premises Equipment
DCTCP	Data Center TCP
DiffServ	Differentiated Services

DPI	Deep Packet Inspection
DSL	Digital Subscriber Line
FTP	File Transfer Protocol
FTTx	Fiber To The x
HTTP	HyperText Transfer Protocol
HSI	High Speed Internet
IP	Internet Protocol
IPTV	IP-based delivery of TV
ISP	Internet Service Provider
NSP	Network Service Provider
ONT	Optical Network Termination
OTT	Over The Top
QoE	Quality of Experience
QoS	Quality of Service
QUIC	Quick UDP Internet Connections
RAC	Resource Admission Control
RGW	Residential Gateway
RTT	Round Trip Time
SLA	Service Level Agreement
TC	Traffic Class
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VPN	Virtual Private Network
VoIP	Voice over IP

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Marketing Report has been approved by members of the Forum. This Marketing Report is subject to change. This Marketing Report is copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Marketing Report may be copyrighted by Broadband Forum members.

Intellectual Property

Recipients of this Marketing Report are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of this Marketing Report, or use of any software code normatively referenced in this Marketing Report, and to provide supporting documentation.

Terms of Use

1. License

Broadband Forum hereby grants you the right, without charge, on a perpetual, non-exclusive and worldwide basis, to utilize the Marketing Report for the purpose of developing, making, having made, using, marketing, importing, offering to sell or license, and selling or licensing, and to otherwise distribute, products complying with the Marketing Report, in all cases subject to the conditions set forth in this notice and any relevant patent and other intellectual property rights of third parties (which may include members of Broadband Forum). This license grant does not include the right to sublicense, modify or create derivative works based upon the Marketing Report except to the extent this Marketing Report includes text implementable in computer code, in which case your right under this License to create and modify derivative works is limited to modifying and creating derivative works of such code. For the avoidance of doubt, except as qualified by the preceding sentence, products implementing this Marketing Report are not deemed to be derivative works of the Marketing Report.

2. NO WARRANTIES

THIS MARKETING REPORT IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS MARKETING REPORT SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY IMPLEMENTER OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS MARKETING REPORT.

3. THIRD PARTY RIGHTS

Without limiting the generality of Section 2 above, BROADBAND FORUM ASSUMES NO RESPONSIBILITY TO COMPILE, CONFIRM, UPDATE OR MAKE PUBLIC ANY THIRD

PARTY ASSERTIONS OF PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS THAT MIGHT NOW OR IN THE FUTURE BE INFRINGED BY AN IMPLEMENTATION OF THE MARKETING REPORT IN ITS CURRENT, OR IN ANY FUTURE FORM. IF ANY SUCH RIGHTS ARE DESCRIBED ON THE MARKETING REPORT, BROADBAND FORUM TAKES NO POSITION AS TO THE VALIDITY OR INVALIDITY OF SUCH ASSERTIONS, OR THAT ALL SUCH ASSERTIONS THAT HAVE OR MAY BE MADE ARE SO LISTED.

The text of this notice must be included in all copies of this Marketing Report.

End of Broadband Forum Marketing Report MR-404