MARKETING REPORT

MR-471.2
Maximum IP-Layer Capacity Metric and Measurement

Issue: 1
Issue Date: October 2020

Issue History

Issue Number	Issue Date	Issue Editor	Changes
1	22 October 2020	Al Morton, AT&T	Original

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

Editor Al Morton, AT&T

Work Area Director(s) David Sinicrope, Ericsson

PEAT PS Leader Gregory Mirsky (ZTE)

Contributors Len Ciavattone, AT&T

Ken Kerpez, ASSIA

Table of Contents

List of Figures

Executive Summary

Many assumptions about Internet Access have changed in recent years. While state of the art and some future developments were fairly accurately captured by literature in 2013, some changes were unforeseen. For example, speed/capacity alone was the key performance measure until very recently, but now latency is just as critical. Also, Transmission Control Protocol (TCP) was the dominant reliable transport protocol; however, User Datagram Protocol (UDP) and QUIC have a large role which continues to grow. Emphasis was on test paths with multiple operator networks from a user to content servers, but now the content is moving closer to the user: Content Distribution Network (CDN) or Mobile Edge Computing servers will terminate connections within a single domain. Any one of these shifts in the network environment, along with the needs of Internet Service Providers (ISPs), have significant repercussions on the network performance. Therefore the metrics and measurements used to assess performance must follow quickly.

This document explores the motivation and gives an introduction to the new UDP-based IP Capacity metric and associated measurement methods recently developed in Broadband Forum in coordination with other global organizations. The methods specified in TR-471 [7] aim to solve current TCP-based measurement issues encountered on today's highest-speed Internet access and provide a simple, standardized metric and method where none previously existed.

1 Background

Many assumptions about Internet Access have changed since publication of *Improved Internet speed tests can enhance QoS and QoE* [1] captured state of the art and future developments in 2013, but some changes were unforeseen, as indicated below.

7 years ago:	Today's trends:
User access was the bottleneck;	Mobile Carrier Aggregation & Gbps access
The main emphasis on Speed;	Latency also/more critical
TCP was *the*reliable transport;	UDP with QUIC large & growing
Measure multi-operator paths from a user to content, and	Content moving to the user: CDNs, Mobile Edge Compute
Measure performance across Gateways between Tier 1 Ops	Content everywhere, Less traffic & less congestion at Gateways
You might not see ALL these trends happening in your region today, but the arrival of any one changes the game!	

Figure 1: Designing Measurements: Today's clear trends, from [2]

Metrics and measurements of Access Performance should follow these trends, along with the needs of Internet Service Providers, their subscribers, and equipment suppliers.

2 Terminology

Term	Definition
QUIC	A new transport protocol under development in the IETF (work-in-progress), inspired by [8]

2.1 References

The following references are of relevance to this Marketing Report. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this Marketing Report are therefore encouraged to investigate the possibility of applying the most recent edition of the references listed below.

A list of currently valid Broadband Forum Technical Reports is published at www.broadband-forum.org.

Document	Title	Source	Year
[1]	*Improved Internet speed tests can enhance QoS and QoE*	RAIM-2015 (PQS 2013)	2013
[2]	*Broadband QoS and QoE Assessment, 5-years-on: Major Changes!*	ITU Workshop	2018
[3]	*QUIC: A UDP-Based Multiplexed and Secure Transport (Work-in-progress)*	IETF	2019
[4]	*Using TLS to Secure QUIC (Work-in-progress)*	IETF	2019
[5]	*QUIC Loss Detection and Congestion Control (Work-in-progress)*	IETF	2019
[6]	*Hypertext Transfer Protocol Version 3 (HTTP/3) (Work-in-progress)*	IETF	2019
[7]	*TR-471, Maximum IP-Layer Capacity Metric, Related Metrics, and Measurements*	BBF	2020
[8]	*QUIC: Quick UDP Internet Connections, Multiplexed Stream Transport over UDP*	Jim Roskind, Google Corp.	2013

2.2 Abbreviations

This Marketing Report uses the following abbreviations:

AR	Augmented Reality
BNG	Border Network Gateway
CDN	Content Distribution Network
CPE	Customer Premises Equipment
CUBIC	A widespread form of congestion control for transport-layer protocols
GW	Gateway
IETF	Internet Engineering Task Force
QUIC	A new transport protocol under development in the IETF (work-in-progress) , inspired by [8]
SDO	Standards Development Organization
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VR	Virtual Reality

3 Access Networks and Services Outrun TCP Speed Tests

As the speed of access links increases into the gigabit range for some technologies, the dominant methods of helping users test to discover their actual "Internet Speed" are showing their age. TCP-based methods, which react conservatively to packet loss and round-trip delay, produce a significant underestimate of Maximum IP-Layer Capacity.

Meanwhile, the Industry is seeing a transition to new protocols like QUIC [3,4,5,6,8], that will replace TCP for many applications. These new protocols use UDP as the Transport protocol and encrypt activity above the transport layer. Measuring the IP-Layer Capacity on a user's access link should use the same Transport protocols as applications.

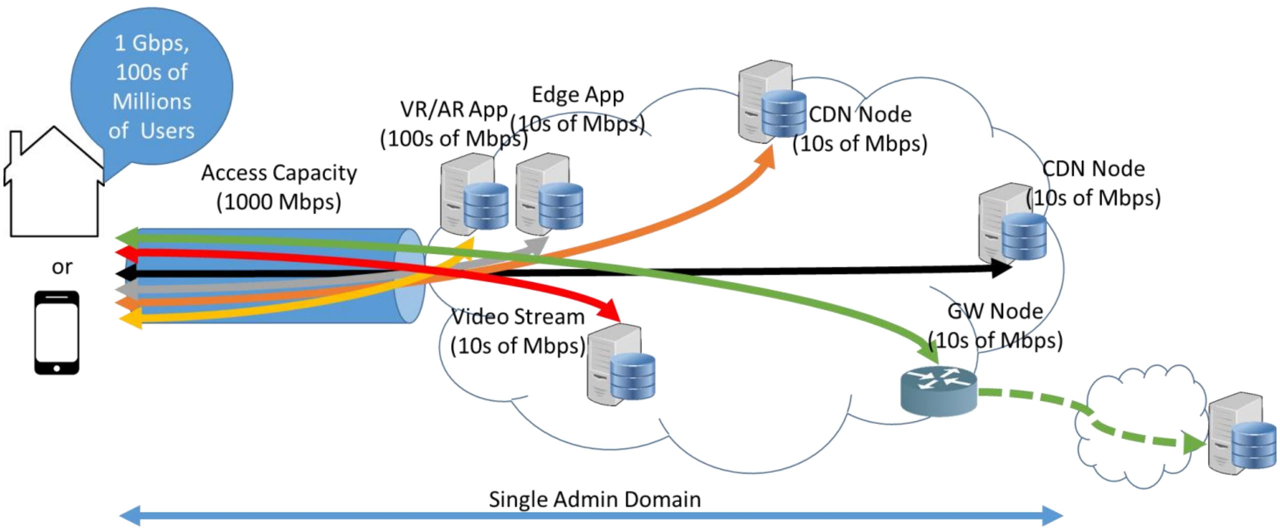

Figure 2: Service Capacity has Access Scope (user paths become diverse beyond), based on [1]

As Figure 2 shows, content is moving closer to users, but the user streams and content to/from servers frequently do not require the user's full access rate. In such cases, latency and not speed becomes the critical issue for many applications (such as Augmented Reality (AR), Virtual Reality (VR), and Games). New metrics and methods of measurement need to reflect these realities.

The metric and methods developed apply to mobile as well as fixed access, with test measurements on an LTE network showing significantly more consistent results than their TCP-based counterparts.

The measurements can also provide consumer confidence in their service, and new dimensions for market messaging to specific user groups (e.g., Gamers emphasize Low Delay). The UDP-based measurement of Maximum IP Capacity simultaneously measures the packet loss, round-trip delay, delay variation, and reordering present. This is superior information to that provided by TCP and Ping measurements made separately.

4 New IP Capacity Metric and Measurement

The Gigabit access services delivered by many service providers today have outstripped the ability of ad-hoc, TCP-based methods to measure their performance. The UDP-based IP Capacity metric and related measurement methods now specified in the Broadband Forum (and other international SDOs) can provide a much more accurate and consistent understanding of network performance.

The Broadband Forum has completed its work on TR-471 [7], a new UDP-based IP Capacity metric and associated measurement methods developed in coordination with other global organizations. This metric and method aim to solve current measurement issues with High-speed Internet access. Choosing UDP transport helps make platform independence and equivalent implementations possible. It anticipates user transition from TCP to UDP and rate control to better take advantage of the Internet's scale in 2020, and it measures other important performance metrics when "speed" is not the only factor considered when choosing service.

At a high-level, the Maximum IP-Layer Capacity Metric and Method comprises:

1. A rate-controlled stream of test packets transmitted from Sender to Receiver.
2. A simple measurement of received bit rate at the Receiver, and measurement of additional packet performance metrics such as loss, delay and delay variation.
3. A measurement feedback channel for Sender packet-rate-control, to conduct an iterative search for the Maximum IP-Layer Capacity

Figure 3 provides a view of the Maximum IP-Layer Capacity measurement at the Receiver, using default values for a test's duration. The total test time is divided into sub-intervals. Many intermediate measurements are made (trials) and used to adjust the Sender's rate to seek the Maximum Capacity, C_{max}, during a test. The search algorithm that controls the Sender's rate is completely different from TCP's congestion control algorithms (even those most recently proposed; this is an active area of research), and avoids many of the drawbacks while allowing more direct control over Sender behavior.

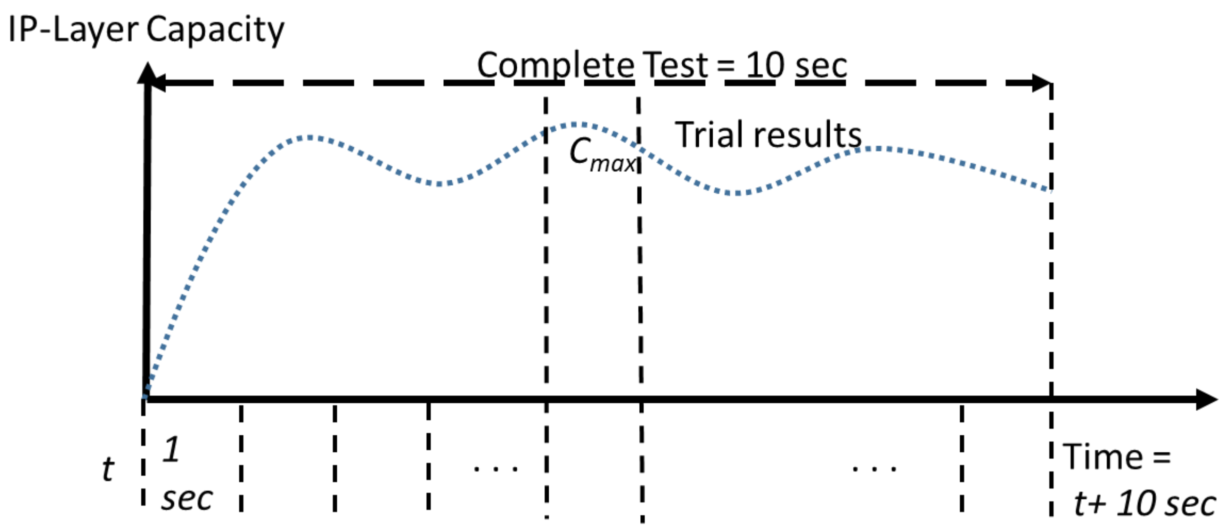

Figure 3: Maximum IP-Layer Capacity measurement at the Receiver

5 New Open Source Implementation of the IP Capacity Metric and Measurement

The new Open Broadband project OB-UDPST introduces the open source version of the udpst utility. A typical test path is shown in Figure 4, with example udpst Server and Client locations identified. Both Downstream and Upstream measurements of IP-Layer Capacity are included.

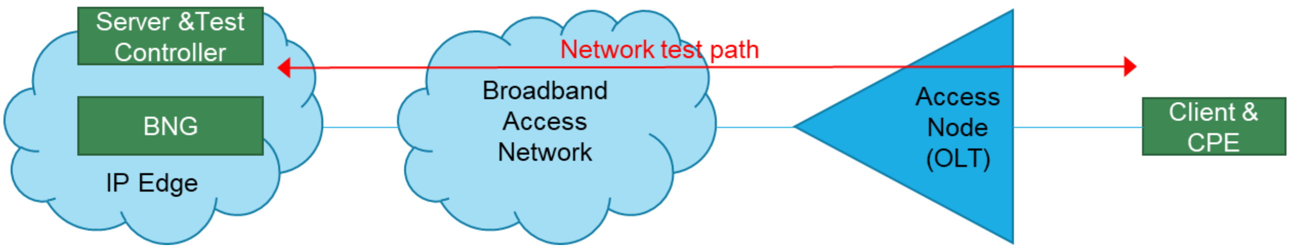

Figure 4: Typical udpst Server and Client test path

The udpst measurement tool runs on the Linux operating system and serves as a working reference for further development. The current project:

- is a utility that can function as a client or server daemon
- is written in C, and built with gcc (release 9.3) and its standard run-time libraries
- works with both IPv4 and IPv6 address families
- includes authentication functionality that accepts a command-line key which is used in the setup request to the server
- allows configuration of most of the measurement parameters described in TR-471 [7].

As Figure 3 indicates, the stream of test packets starts out at a low rate and rapidly increases while the test is in progress, searching for the Maximum IP-Layer Capacity. A load adjustment search algorithm governs this behavior. This algorithm was designed for Capacity assessment (unlike TCP congestion control algorithms), and provides many straightforward ways to control the algorithm behavior. However, the BBF TR-471 [7] specification provides the default settings that produce good results in all tested circumstances.

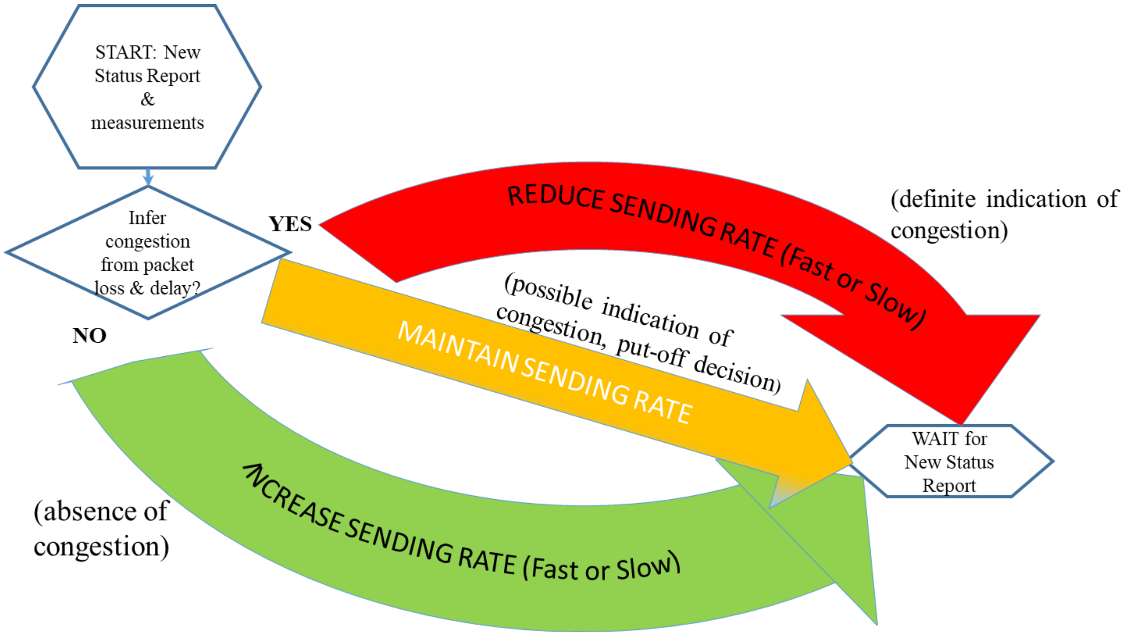

Figure 5: Three main paths through the load adjustment search algorithm

Evaluation is triggered by the arrival of a new status report, and the result is the change (or retention) of the current sending rate, depending on the measurements. Figure 5 illustrates three paths: when feedback indicates measured impairments are absent, or when impairments are first measured and some congestion may be present but sending rate change is deferred, or when measured impairments are confirmed by repeated measurement feedback.

The new open source implementation is available as part of the Open Broadband project. An example of a client's output is provided below.

```
$ udpst -d udp-speedtest.com
UDP Speed Test
Software Ver: 6.4, Protocol Ver: 6, Built: Aug 20 2020 12:08:37
Mode: Client, Jumbo Datagrams: Enabled, Authentication: Available
Downstream Test Interval(sec): 10, DelayVar Thresholds(ms): 30-90 [RTT], Trial Interval(ms): 50,
    SendingRate Index: <Auto>, Congestion Threshold: 2, High-Speed Delta: 10, SeqError Threshold: 0
Sub-Interval(sec):  1, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/1, RTTVar(ms): 0-3, Mbps(L3/IP): 93.40
Sub-Interval(sec):  2, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/1, RTTVar(ms): 0-2, Mbps(L3/IP): 293.37
Sub-Interval(sec):  3, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/1, RTTVar(ms): 0-3, Mbps(L3/IP): 493.30
Sub-Interval(sec):  4, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/1, RTTVar(ms): 0-3, Mbps(L3/IP): 693.47
Sub-Interval(sec):  5, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/2, RTTVar(ms): 0-3, Mbps(L3/IP): 891.30
Sub-Interval(sec):  6, Delivered(%):  99.65, Loss/OoO: 343/0, OWDVar(ms): 2/3/4, RTTVar(ms): 3-5, Mbps(L3/IP): 966.41
Sub-Interval(sec):  7, Delivered(%):  99.78, Loss/OoO: 218/0, OWDVar(ms): 1/3/5, RTTVar(ms): 2-7, Mbps(L3/IP): 965.97
Sub-Interval(sec):  8, Delivered(%):  99.88, Loss/OoO: 112/0, OWDVar(ms): 0/1/3, RTTVar(ms): 1-5, Mbps(L3/IP): 959.28
Sub-Interval(sec):  9, Delivered(%):  99.94, Loss/OoO: 56/0, OWDVar(ms): 2/2/4, RTTVar(ms): 2-6, Mbps(L3/IP): 965.30
Sub-Interval(sec): 10, Delivered(%):  99.78, Loss/OoO: 216/0, OWDVar(ms): 3/5/7, RTTVar(ms): 3-8, Mbps(L3/IP): 967.53
Downstream Summary Delivered(%):  99.90, Loss/OoO: 945/0, OWDVar(ms): 0/1/7, RTTVar(ms): 0-8, Mbps(L3/IP): 728.93
Downstream Minimum One-Way Delay(ms): 3 [w/clock difference], Round-Trip Time(ms): 7
Downstream Maximum Mbps(L3/IP): 967.53, Mbps(L2/Eth): 981.55, Mbps(L1/Eth): 997.12, Mbps(L1/Eth+VLAN): 1000.23

$ udpst -u udp-speedtest.com
UDP Speed Test
Software Ver: 6.4, Protocol Ver: 6, Built: Aug 20 2020 12:08:37
Mode: Client, Jumbo Datagrams: Enabled, Authentication: Available
Upstream Test Interval(sec): 10, DelayVar Thresholds(ms): 30-90 [RTT], Trial Interval(ms): 50,
    SendingRate Index: <Auto>, Congestion Threshold: 2, High-Speed Delta: 10, SeqError Threshold: 0
Sub-Interval(sec):  1, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/2, RTTVar(ms): 0-1, Mbps(L3/IP): 93.63
Sub-Interval(sec):  2, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/1, RTTVar(ms): 0-1, Mbps(L3/IP): 293.39
Sub-Interval(sec):  3, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/1, RTTVar(ms): 0-1, Mbps(L3/IP): 493.49
Sub-Interval(sec):  4, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/1, RTTVar(ms): 0-1, Mbps(L3/IP): 693.82
Sub-Interval(sec):  5, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/3, RTTVar(ms): 0-1, Mbps(L3/IP): 891.94
Sub-Interval(sec):  6, Delivered(%):  99.60, Loss/OoO: 395/0, OWDVar(ms): 2/4/7, RTTVar(ms): 2-6, Mbps(L3/IP): 966.83
Sub-Interval(sec):  7, Delivered(%):  99.83, Loss/OoO: 169/0, OWDVar(ms): 6/7/10, RTTVar(ms): 6-9, Mbps(L3/IP): 967.53
Sub-Interval(sec):  8, Delivered(%):  99.83, Loss/OoO: 163/0, OWDVar(ms): 9/11/13, RTTVar(ms): 9-13, Mbps(L3/IP): 967.75
Sub-Interval(sec):  9, Delivered(%):  99.82, Loss/OoO: 176/0, OWDVar(ms): 9/11/13, RTTVar(ms): 10-13, Mbps(L3/IP): 967.44
Sub-Interval(sec): 10, Delivered(%):  99.66, Loss/OoO: 331/0, OWDVar(ms): 10/12/13, RTTVar(ms): 10-13, Mbps(L3/IP): 966.69
Upstream Summary Delivered(%):  99.87, Loss/OoO: 1234/0, OWDVar(ms): 0/5/13, RTTVar(ms): 0-13, Mbps(L3/IP): 730.25
Upstream Minimum One-Way Delay(ms): 4 [w/clock difference], Round-Trip Time(ms): 7
Upstream Maximum Mbps(L3/IP): 967.75, Mbps(L2/Eth): 981.75, Mbps(L1/Eth): 997.31, Mbps(L1/Eth+VLAN): 1000.42
```

Figure 6: Example udpst 1Gbps Passive Optical Network Service test results

Some key points to note in Figure 5:

- The first 5 lines below the "$ udpst …" command line provide a summary of the test configuration.
- Each line beginning "Sub-Interval(sec):" provides measurements for 1 second of the 10-second complete test in Figure 3.
- Loss and reordering (Out-of-Order packets) are tracked for each sub-interval.
- One-way and round-trip delay variation is measured while testing and reported in status feedback messages.
- Minimum one-way and round-trip absolute delay covers the entire test.
- L3/IP rates are measured, lower-layer protocol rates are calculated.
- L1/Eth/VLAN Calculated values align with relevant provisioned or theoretical maximums.

The source code for the project is publically available at https://github.com/BroadbandForum/obudpst
See https://github.com/BroadbandForum/obudpst/blob/main/CONTRIBUTING.md for details on participating in the project.

It is, of course, necessary to calibrate the equipment performing the IP-layer Capacity measurement, to ensure that the expected capacity can be measured accurately and that equipment choices (processing speed, interface bandwidth, etc.) are suitably matched to the measurement range.

6 Summary

The new Maximum IP-Layer Capacity Metric and Method(s) of Measurement, based on UDP, closes the gap between actual service rates and TCP's estimates under the measured conditions.

The new metric and method can measure the new Gigabit service without the confusing artifacts of TCP performance (such as its throughput sensitivity to packet loss, round-trip time, and its flow control details).

The New Metric and Open Source implementation provides:

- A new and more accurate approach to Internet access measurement.
- An integral part of BBF's specifications, harmonized with other SDOs for broadest applicability and relevance.
- A way to overcome the disadvantages of Ad-Hoc TCP-based methods at Gigabit Internet Access Speeds.

End of Broadband Forum Marketing Report MR-471.2

