International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Series Y
Supplement 60
(10/2021)

SERIES Y: GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS, NEXT-GENERATION NETWORKS, INTERNET OF THINGS AND SMART CITIES

Interpreting ITU-T Y.1540 maximum IP-layer capacity measurements

ITU-T Y-series Recommendations – Supplement 60

ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS, NEXT-GENERATION NETWORKS, INTERNET OF THINGS AND SMART CITIES

For further details, please refer to the list of ITU-T Recommendations.

Supplement 60 to ITU-T Y-series Recommendations

Interpreting ITU-T Y.1540 maximum IP-layer capacity measurements

Summary

This Supplement 60 to the ITU-T Y-series Recommendations provides information on interpreting ITU-T Y.1540 maximum IP-layer capacity measurements as described in Annex A and Annex B of the Recommendation.

This Supplement also provides useful information for those who measure various technologies. Much has been learned as part of extensive testing campaigns thus far, and there is more to learn. Therefore, this Supplement may be updated frequently, and readers are encouraged to ensure that they are using the most recent version.

History

Edition	Recommendation	Approval	Study Group	Unique ID*
1.0	ITU-T Y Suppl. 60	2020-04-24	12	11.1002/1000/14285
2.0	ITU-T Y Suppl. 60	2020-09-11	12	11.1002/1000/14496
3.0	ITU-T Y Suppl. 60	2021-05-13	12	11.1002/1000/14707
4.0	ITU-T Y Suppl. 60	2021-10-21	12	11.1002/1000/14830

Keywords

Fixed access, Internet performance, latency, mobile access, mobile testing, packet delay variation, packet loss, packet reordering.

* To access the Recommendation, type the URL http://handle.itu.int/ in the address field of your web browser, followed by the Recommendation's unique ID. For example, http://handle.itu.int/11.1002/1000/11830-en.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

This is an informative ITU-T publication. Mandatory provisions, such as those found in ITU-T Recommendations, are outside the scope of this publication. This publication should only be referenced bibliographically in ITU-T Recommendations.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this publication may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the publication development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents/software copyrights, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the appropriate ITU-T databases available via the ITU-T website at http://www.itu.int/ITU-T/ipr/.

Table of Contents

Supplement 60 to ITU-T Y-series Recommendations

Interpreting ITU-T Y.1540 maximum IP-layer capacity measurements

1 Scope

This Supplement provides information on methods to interpret ITU-T Y.1540 maximum Internet protocol layer (IP-layer) capacity measurements, and the auxiliary measurements of fundamental IP packet performance parameters collected simultaneously. Specifically, it includes among other content:

– Guidance on the use of Figure A.1 of [ITU-T Y.1540] — Flowchart for offered load adjustment as part of a search algorithm – versus Figure B.1 of [ITU-T Y.1540] – Flowchart for offered load adjustment, Type B search algorithm;

– Examples where maximum IP-layer capacity measurement results need to be carefully interpreted;

– An example of an ITU-T Y.1540 Annex B table of sending rates;

– Guidance for comparing different forms of measurement with the normative methods of [ITU-T Y.1540];

– Test results of note post-dating the publication of [ITU-T Y.1540].

2 References

[ITU-T Y.1540] Recommendation ITU-T Y.1540 (2019), *Internet protocol data communication service – IP packet transfer and availability performance parameters*.

[BBF TR-471] Broadband Forum Technical Report (2020), *TR-471 Maximum IP-Layer Capacity Metric, Related Metrics, and Measurements*.
https://www.broadband-forum.org/technical/download/TR-471.pdf

3 Definitions

3.1 Terms defined elsewhere

None.

3.2 Terms defined in this Supplement

None.

4 Abbreviations and acronyms

This Supplement uses the following abbreviations and acronyms:

ETH Ethernet

ICMP Internet Control Message Protocol

RTT Round-Trip Time

TCP Transmission Control Protocol

UDP User Datagram Protocol

VLAN Virtual Local Area Network

5 Conventions

None.

6 Recommendation for use of Annex B (Type B) search algorithm

This clause provides guidance on the use of Figure A.1 of [ITU-T Y.1540] – Flowchart for offered load adjustment as part of a search algorithm – versus Figure B.1 of [ITU-T Y.1540] – Flowchart for offered load adjustment, Type B search algorithm.

When considering the two normative flowcharts for the two algorithms, it can be seen that they have some common features. As Annex B explains:

> "There are three main paths through the flowchart: 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."

Figure 6-1 illustrates these three paths. 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.

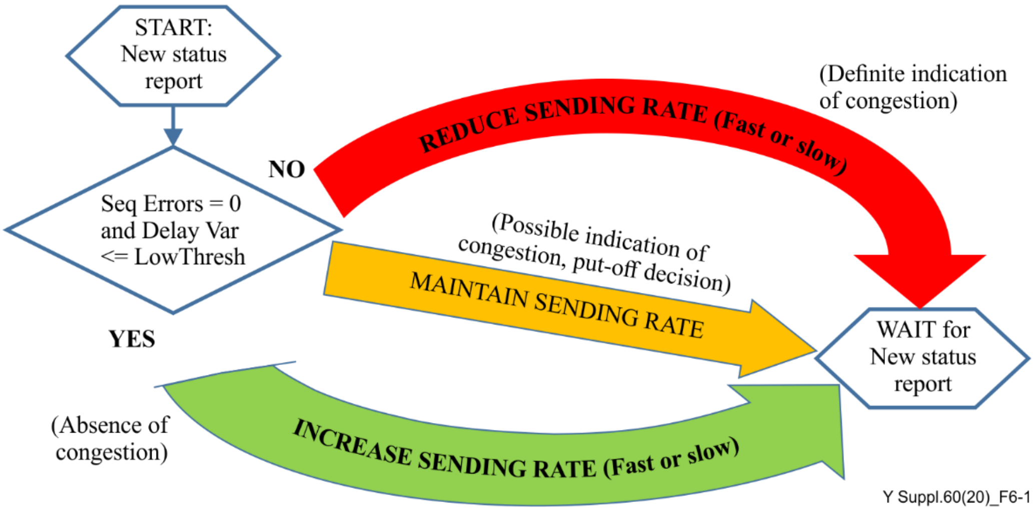

Figure 6-1 – Three main paths through the search algorithm flowcharts

The Annex B (Type B) search adds two supplementary capabilities to the evaluation of measurements and control of the sending rate that are not available in the Annex A algorithm:

1. Initial fast increase of the sending rate from the default starting rate of 500 kbit/s. This saves time when testing access capacities that are in the gigabit range, but allows the same parameter settings to be used over a wide range of IP-layer capacities (it may be some time before cellular rates of 1 Gbit/s are available in many regions).
2. Initial fast decrease of the sending rate when congestion is confirmed, to minimize overshoot and continue a less-aggressive search in the range where measurements indicate that a capacity limit may have been reached.

These features may reduce the time required for measurements of maximum IP-layer capacity and the auxiliary metrics such as loss and round-trip time, as illustrated in Figure 6-2, which shows depictions of time at the receiver.

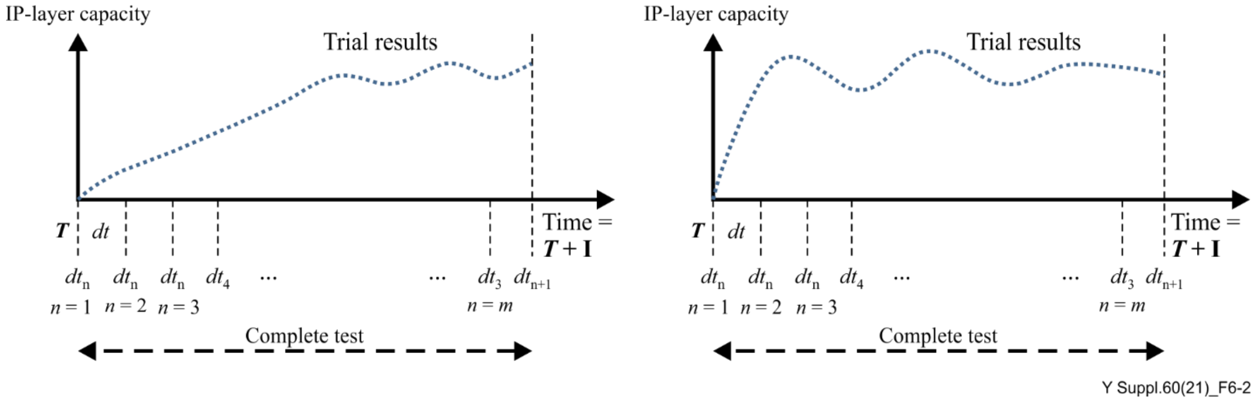

Figure 6-2 – Illustration of relative speed of Annex A (left) and Annex B (right) algorithms

Also, the parameters of the Annex B (Type B) algorithm allow for more control over the behaviour of the search and its performance in specific circumstances. Finally, it is possible to set the parameters of the Annex B (Type B) algorithm so that the resulting search follows the algorithm in Annex A.

It is for these reasons that Figure B.1 of [ITU-T Y.1540] – Flow chart for offered load adjustment, Type B search algorithm – is preferred, and should be implemented by measurement systems measuring maximum IP-layer capacity and auxiliary metrics.

The concluding statements of Annex B concern the topic of testing with parallel connections:

> The current view is that each connection would have its own feedback channel, calculation of measurements and flow chart, and a report of the aggregate results over all connections.

This means that results of each connection for number of bits transferred during *dt* would be computed separately (as well as other metrics), and possibly reported separately for each connection, if requested.

7 Example results requiring careful interpretation

This clause provides examples where the maximum IP-layer capacity measurement results need to be carefully interpreted.

7.1 Measurement results that indicate a bimodal distribution of capacities

Some forms of Internet access include a mode of operation where the early packets of a new connection are allowed to flow at a higher rate than the subsequent packets. The definition of "early" is not specified, and could be time in seconds or total bytes observed, but the intent is to give the appearance of better service to the user's application. For example, a webpage might load faster, or a stream might reach its initial buffering point more quickly. Figure 7-1 illustrates how this might be observed in IP-layer capacity measurements.

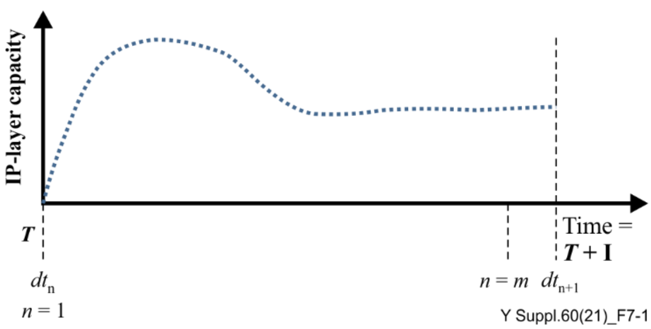

Figure 7-1 – Illustration of multiple modes of operation encountered in measurement

This network behaviour was more prevalent when Internet access speeds were limited by both shaping and technology and is sometimes called "turbo mode". However, this mode of operation was observed in measurements made in the United States of America in 2019 [b-Mathis]. Some applications continue to derive potential performance benefits, so the measurement analyst should be prepared to recognize such cases. The sustained IP-layer capacity will only be observed after the conditions of the "turbo mode" are satisfied, and the shaping parameters return to normal.

The distinguishing feature of "turbo mode" is its reliable and repeatable measurement after a sufficiently long interval where traffic is idle (also, the user may know that they have such a feature enabled, but features and capacities are often added or dropped without the user comprehending the notification in some form).

It is strongly recommended when the bimodal behaviour described above is encountered that the two modes be characterized independently as an enhanced mode and a sustained mode. Near-continuous testing allows more time to measure the IP-layer capacity of the sustained mode.

It is also notable that some conditions may cause an individual measurement to possess bimodal or multimodal characteristics. For example, switches between more or less complicated modulation constellations, cellular modem technology modes (3G, LTE, 5G) or severe weather intervals could cause short-time modes to appear. However, such conditions are unlikely to be exactly replicated in subsequent measurements.

7.2 Measurement results where packet losses occur independently of the send rate

Under some conditions which might be encountered during measurement packet losses may occur independently of the send rate. Such cases include testing, where there is a significant volume of competing traffic from many users, or when the production network performance exhibits a sufficient yet small amount of packet loss that the search algorithms would misinterpret as congestion. The presence of background packet loss will slow the initial advance of the search algorithm, possibly resulting in a measurement result where the true capacity has not been found during the configured test duration. Figure 7-2 illustrates the effect on the progress of search algorithms (the effect is approximately the s same for Annex A or Annex B.

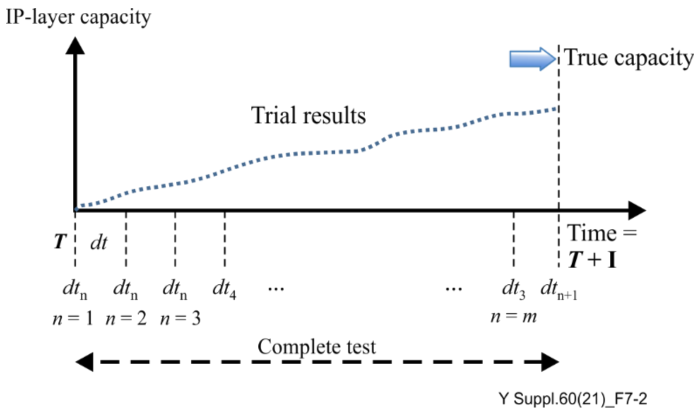

Figure 7-2 – Illustration of the effect of background packet loss encountered in measurement

One factor which distinguishes the background loss condition from normal capacity limits in measurements is the lack of delay variation observed as reported in the frequent status messages.

There are three potential mitigations when encountering background loss:

1. Testing for longer durations: Search progress is being made in Figure 7-2 but doubling the test interval may be helpful to achieve an accurate measurement of capacity in this example.
2. Starting the test at a higher initial sending rate, anticipating the shallow slope of the search.

3. Revising the Annex B (Type B) parameters from the default values, to make the search more aggressive in the presence of packet loss, while retaining the same sensitivity to delay variation. One setting to change is SlowAdjThresh, currently default at two occurrences. Tests have shown that setting the SlowAdjThresh to a very high value avoids both the slow mode of rate increases and the fast recovery branch of the rate decrease branch of the flow chart. As mentioned earlier, this is a key advantage of the Annex B (Type B) algorithm.

In summary, there are several mitigations for the background packet loss effect on accuracy, and a clear way to distinguish the condition: the lack of delay variation present in measurements.

8 Example of an ITU-T Y.1540 Annex B table of sending rates

When describing the system that conforms to Annex B of [ITU-T Y.1540], the text describes a pre-defined table:

> "A table of transmit rates, which are the number of packets sent during each time interval (corresponding to bits per second and a specified protocol layer) and packet sizes. The table has ascending values for offered load rates, between the minimum and maximum supported load rates, inclusive."

The example Table 8-1 begins with the configured initial rate and contains a representation of one table that has been used (with reduced detail).

8.1 Example table

Table 8-1 provides the example. This table can be used with both Annex A and Annex B. Note that this table directly computes the bit rate from packet counts, which was considered to be more efficient.

Table 8-1 – Example table for Annex B

Index No.	Other components, such as sending interval, payload size(s), spacing parameters, back-to-back (burst) lengths	Mbit/s (L3/IP)
0		0.50
1		1.00
2		2.00
3		3.00
4		4.00
…		…
997		997.00
998		998.00
999		999.00
1000		1000.00
…		…

9 Test results of note post-dating the publication of [ITU-T Y.1540]

Appendix X of [ITU-T Y.1540] gives a summary of IP-layer capacity test results up to May 2019. Testing has continued with the udpst utility as a basis (see clause 10).

9.1 New results on 5/100 DOCSIS access

A new set of results are given in Figure 9-1.

```
$ udpst -d udp-speedtest.com
UDP Speed Test
Software Ver: 6.4, Protocol Ver: 6, Built: Aug 21 2020 17:44:21
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): 0/12/61, RTTVar(ms): 0-52, Mbps(L3/IP): 80.47
Sub-Interval(sec):  2, Delivered(%):  89.36, Loss/OoO: 1493/0, OWDVar(ms): 61/122/150, RTTVar(ms): 77-147, Mbps(L3/IP): 122.93
Sub-Interval(sec):  3, Delivered(%):  98.54, Loss/OoO: 192/0, OWDVar(ms): 54/93/125, RTTVar(ms): 62-118, Mbps(L3/IP): 123.10
Sub-Interval(sec):  4, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 4/18/56, RTTVar(ms): 2-60, Mbps(L3/IP): 121.32
Sub-Interval(sec):  5, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 4/28/71, RTTVar(ms): 1-64, Mbps(L3/IP): 123.96
Sub-Interval(sec):  6, Delivered(%):  97.25, Loss/OoO: 361/0, OWDVar(ms): 70/105/125, RTTVar(ms): 69-118, Mbps(L3/IP): 123.12
Sub-Interval(sec):  7, Delivered(%):  99.57, Loss/OoO: 56/0, OWDVar(ms): 37/84/117, RTTVar(ms): 33-110, Mbps(L3/IP): 122.99
Sub-Interval(sec):  8, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 6/12/40, RTTVar(ms): 0-28, Mbps(L3/IP): 122.50
Sub-Interval(sec):  9, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 12/53/101, RTTVar(ms): 10-98, Mbps(L3/IP): 123.47
Sub-Interval(sec): 10, Delivered(%):  96.78, Loss/OoO: 425/0, OWDVar(ms): 100/114/122, RTTVar(ms): 98-138, Mbps(L3/IP): 123.12
Downstream Summary Delivered(%):  98.15, Loss/OoO: 2527/0, OWDVar(ms): 0/64/150, RTTVar(ms): 0-147, Mbps(L3/IP): 118.70
Downstream Minimum One-Way Delay(ms): -562 [w/clock difference], Round-Trip Time(ms): 16
Downstream Maximum Mbps(L3/IP): 123.96, Mbps(L2/Eth): 125.80, Mbps(L1/Eth): 127.85, Mbps(L1/Eth+VLAN): 128.26
```

Figure 9-1 – 5/100 DOCSIS access with Wi-Fi connected client

This is a downstream test, with udpst 6.4 (currently targeted for open source release). The tests were conducted on 5/100 Mbit/s DOCSIS Access, using Wi-Fi connectivity in the home (reported to offer 144 Mbit/s link speed both receive and transmit, apparently sufficient for testing the subscribed IP-layer capacity). Note that calculated values align with relevant provisioned or theoretical maximums.

In the measurement above, the Wi-Fi network was not a handicap to measurement accuracy, as the result with a wired Ethernet (ETH) network shows below:

```
Downstream Maximum Mbps(L3/IP): 123.16, Mbps(L2/Eth): 124.99, Mbps(L1/Eth): 127.03, Mbps(L1/Eth+VLAN):
127.43
```

9.2 Comparison of measurement results with new tools

During July 2020, the Body of European Regulators for Electronic Communications (BEREC) code repository nntool [b-nntool] was announced.

At the ITU-T SG12 meeting in September 2020, it was deemed useful to share a few preliminary test results comparing udpst (see clause 10) and the Linux-nntool's abilities to measure a 1 Gbit/s access link under the circumstances where the comparative measurement was possible. At present, as with the procedures described in Appendix X of [ITU-T Y.1540], the goal is to reliably measure the "ground truth" of maximum IP-layer capacity offered by the service. In this case, both the nntool and udpst user's clients are located on a single host in the United States of America, but two measurement servers are located in the same city in Germany. This allows each tool to communicate with a server that was optimized by the tool designers, but it introduces some path differences at one end within the destination city. The udpst path involved 22 hops from user client to the server running on Amazon Web Services, while the nntool path involved only 11 hops to the peer-ias-de-01.net-neutrality.tools server. The round-trip time (RTT) on each path was ~87.5 ms.

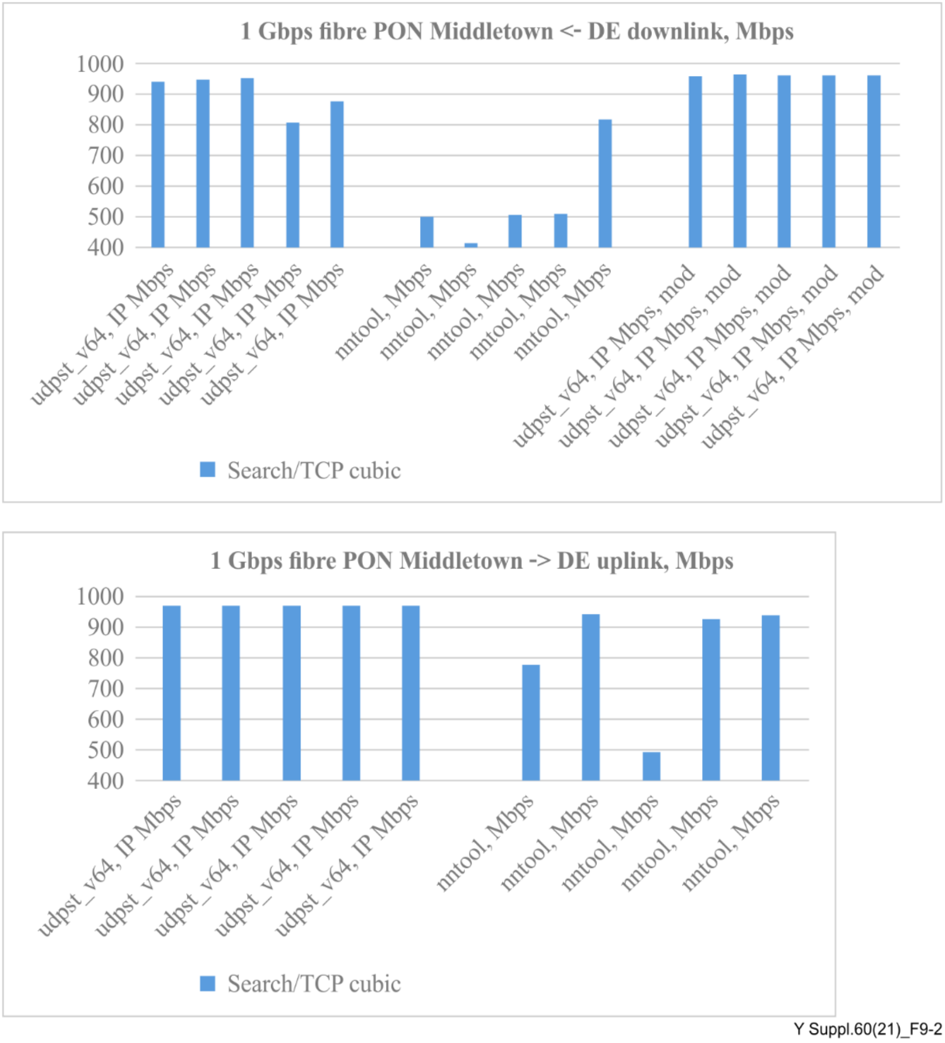

Figure 9-2 – Preliminary comparison of udpst and nntool with 1 Gbit/s access and long RTT

Figure 9-2 provides the preliminary results, using the maximum of values reported by nntool in a 10 second test (similar to udpst). The top graph for downlink results shows that only five repeated tests were needed to illustrate the variability of transmission control protocol (TCP) as a basis for the measurement used in nntool. Default udpst approached the optimum maximum IP-layer 968 Mbit/s in most tests, especially when the rate adjustment algorithm was modified slightly to search more quickly (suffix "mod" in results on the right). In the bottom graph, uplink measurements were very accurate and reliable for five unmodified udpst tests, but nntool results underestimated the capacity and indicated considerable variation when repeated, even when using as many as nine (TCP) threads.

9.3 Results on DOCSIS 3.1 gigabit downlink access

This is a downstream test with udpst 7.1 (in its third public release as of March 5, 2021). The tests were conducted on DOCSIS 3.1 access characterized as "up to 940 Mbit/s", using wired Ethernet (ETH) connectivity in one home as part of a multidwelling unit (uplink capacity is much lower, approximately 24 Mbit/s). In this service offering, the Wi-Fi network does not have sufficient performance to measure the downlink IP-layer capacity accurately, as indicated by the provider.

The udpst 7.1 utility uses all default values for this testing, including 10 second duration, 1 second sub-intervals, zero "sequence errors" and default lower/upper delay variation thresholds. iPerf3.6 also uses as many defaults as possible, except as noted (such as number of connections, as indicated in Figure 9-3 with 1xTCP, 3xTCP, 6xTCP). The test client is a mid-2015 MacBook Pro with eight cores, running Virtual Box and Ubuntu 18.02 LTS VM (2 cores) hosting the test software, with a CUBIC congestion control algorithm.

The test results are given in Figure 9-3.

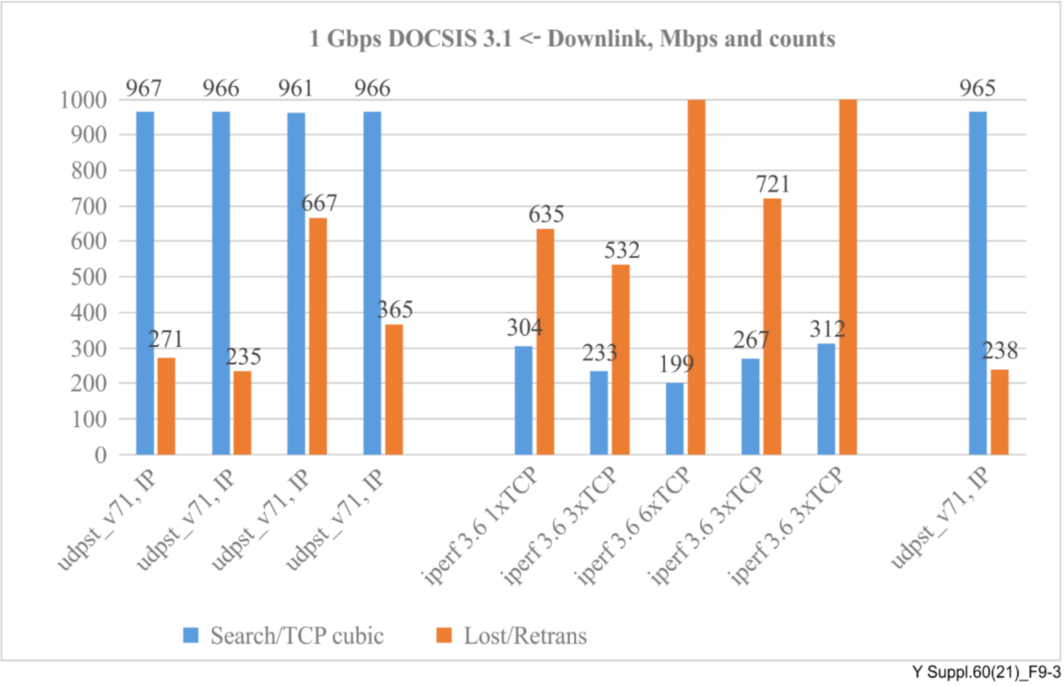

Figure 9-3 – Gigabit DOCSIS 3.1 access with wired connected client in a VM

Figure 9-3 plots both the measured values for udpst 7.1 and iperf 3.6 (using TCP connections), in terms of the measured rates and lost packets during the sub-interval with the maximum IP-layer capacity (or total TCP retransmissions). Traceroute measurements from the client to the server indicated 14 hops and ~35 ms RTT using ICMP.

The measurement results indicate that the udpst load rate adjustment (search) algorithm is primarily responding to packet losses. Delay variation reported for the measurement sub-intervals were below the 30 ms lower delay range threshold. The iPerf3.6 TCP tests also indicate many retransmissions at the end of a test.

The TCP throughput is surprisingly low and increasing the number of parallel connections from 1 to 3 to 6 (where throughput was 199 Mbit/s with 2032 total retransmissions) did not improve the situation. When possible causes for low TCP performance were investigated, it was found that the Ubuntu host settings were sufficient for this test path. Also, omitting the first two seconds of TCP ramp-up *reduced* the overall throughput results. The TCP outcome may be due to the presence of independent packet losses in addition to those induced by TCP's bottleneck probing, also measured with UDP-based udpst 7.2.1 (as discussed in clause 9.4 below).

In the measurements above, the IP-layer maximum is likely to be ~967 Mbit/s because it appears that the Layer 1 transmission may use one virtual local area network (VLAN) tag, as one result and calculation indicate below:

```
Downstream Maximum Mbps(L3/IP): 967.13, Mbps(L2/Eth): 981.11, Mbps(L1/Eth): 996.64, Mbps(L1/Eth+VLAN): 999.74
```

Tests with udpst very nearly attained this maximum, noting the lowest value of capacity was accompanied by the highest user datagram protocol (UDP) packet loss (961 Mbit/s with 667 lost packets).

We repeated the tests with the udpst 7.1.0 client running on a small baremetal host.

The test client is running on a Raspberry Pi 4, Model B with arm64 processor (4 cores) and the raspios-buster 64-bit operating system (Debian GNU/Linux 10 (buster) Kernel: Linux 5.10.17-v8+).

The TCP parameters use the CUBIC congestion control algorithm. The test results are given in Figure 9-4.

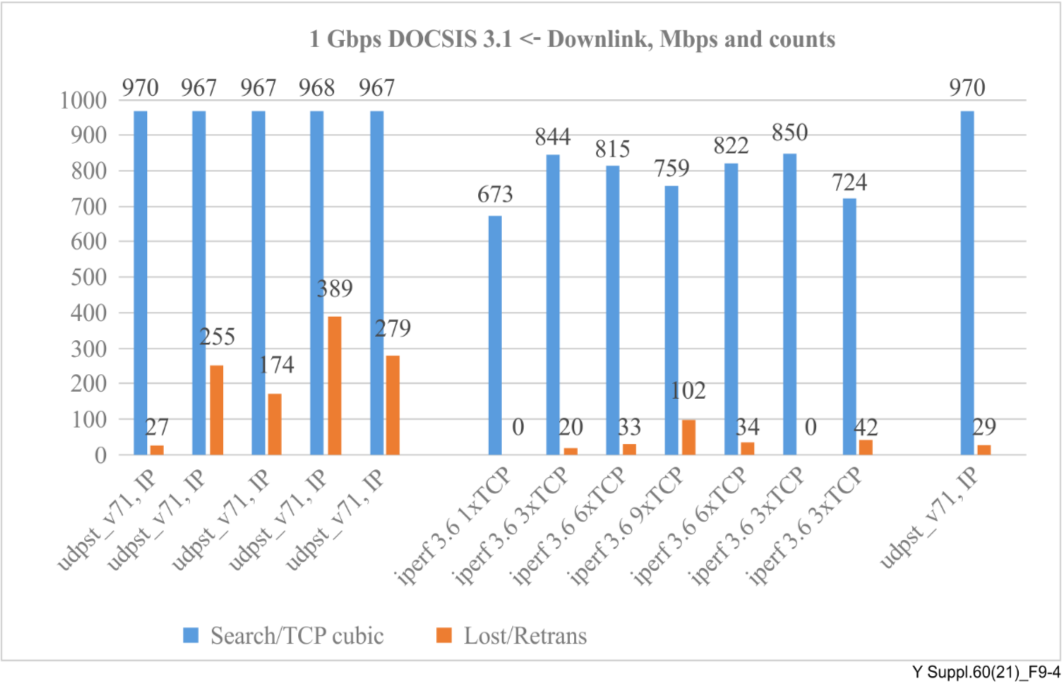

Figure 9-4 – Gigabit DOCSIS 3.1 access with wired connected client baremetal host

Figure 9-4 plots both the measured values for udpst 7.1 and iperf 3.6 (using TCP connections), in terms of the measured rates and lost packets during the sub-interval with the maximum IP-layer capacity (or total TCP retransmissions). Traceroute measurements from the client to the server indicated 14 hops and ~39 ms RTT using Internet control message protocol (ICMP).

The measurement results indicate that the udpst load rate adjustment (search) algorithm is again responding to packet losses. Delay variation reported for the measurement sub-intervals were below the 30 ms lower delay range threshold. The iPerf3.6 TCP tests also indicate various retransmission counts, but considerably fewer than with the VM client. Consequently, the measured rates at the end of the test are much higher for 1, 3, 6, and in this measurement set, 9 connections with the small, baremetal client. However, none of the TCP-based tests achieved 900 Mbit/s on this test path.

Taken overall, the results presented in this clause illustrate TCP-sensitivity to both the path characteristics and the client-host contribution to measurement results; two factors which had little influence over the udpst measurement of maximum IP-layer capacity.

Figure 9-5 shows a Raspberry Pi 4 Model B running udpst 7.1 client as baremetal host.

Figure 9-5 – Raspberry Pi 4 Model B running udpst 7.1 client as baremetal host

9.4 Results with non-congestion-related loss on downlink access

Some testing using 1 Gbit/s access speeds and test paths between client and server that included Internet gateways and servers implemented on virtual hosts prompted some examination of the current set of default values. The test results indicated that higher values for "Seq Errors" and the number of consecutive status messages with impairments would make the "Fast" sending rate ramp-up (see Figure 6-1 above) more robust to the non-congestion-related packet loss that may have been present.

A test using a combination of emulated packet loss in the lab and production access facilities was devised. Figure 9-6 shows the test set-up. Random and burst loss were emulated using the "Linux NetEm Server".

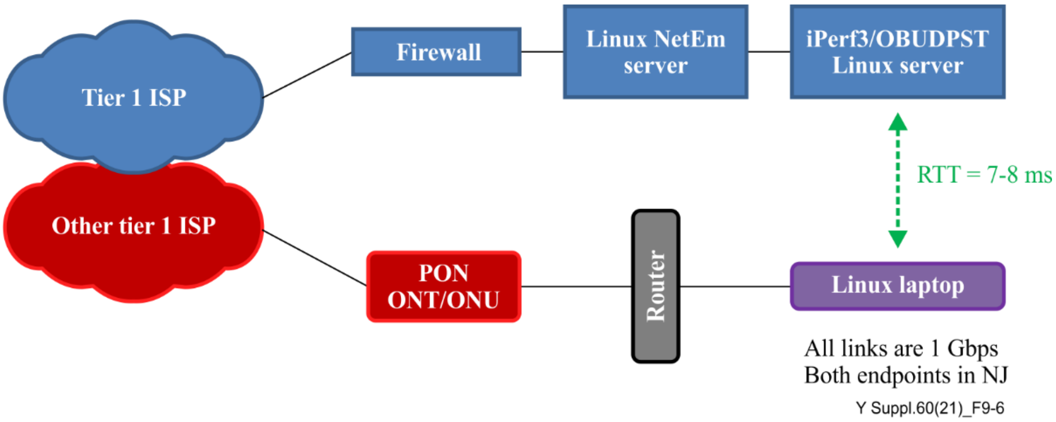

Figure 9-6 – Test set-up for non-congestion-related loss on downlink access

The tests were conducted on fibre PON access characterized many times, using wired ETH connectivity in one home.

This is a downstream test with udpst 7.2.0. After some initial tests with deterministic loss and a range of settings for the parameters under consideration, it was decided to make the following changes:

SeqErrThresh default of 0 errors changes to 10 errors (loss, reordering and duplication)

SlowAdjThresh default of two consecutive status messages to three consecutive status messages

The original udpst performance with the default values identified above is labelled "udp-default" in Figures 9-7 and 9-9, and the changed values identified by the label "udp-q10c3" (which gives a hint to the command-line parameters needed to make the changes with udpst 7.2.0).

The udpst 7.2.0 utility used all other default values for this testing, including 10 second duration, 1 second sub-intervals, and default lower/upper delay variation thresholds.

The testing also included tests with 8 TCP connections supplied by iPerf3.6 using as many defaults as possible. The average of 10 second stable TCP bit stream transfers and the maximum bits per second delivered in any one second sub-interval were plotted in Figures 9-7 and 9-9, as "8xTCP ave" and "8xTCP max", respectively.

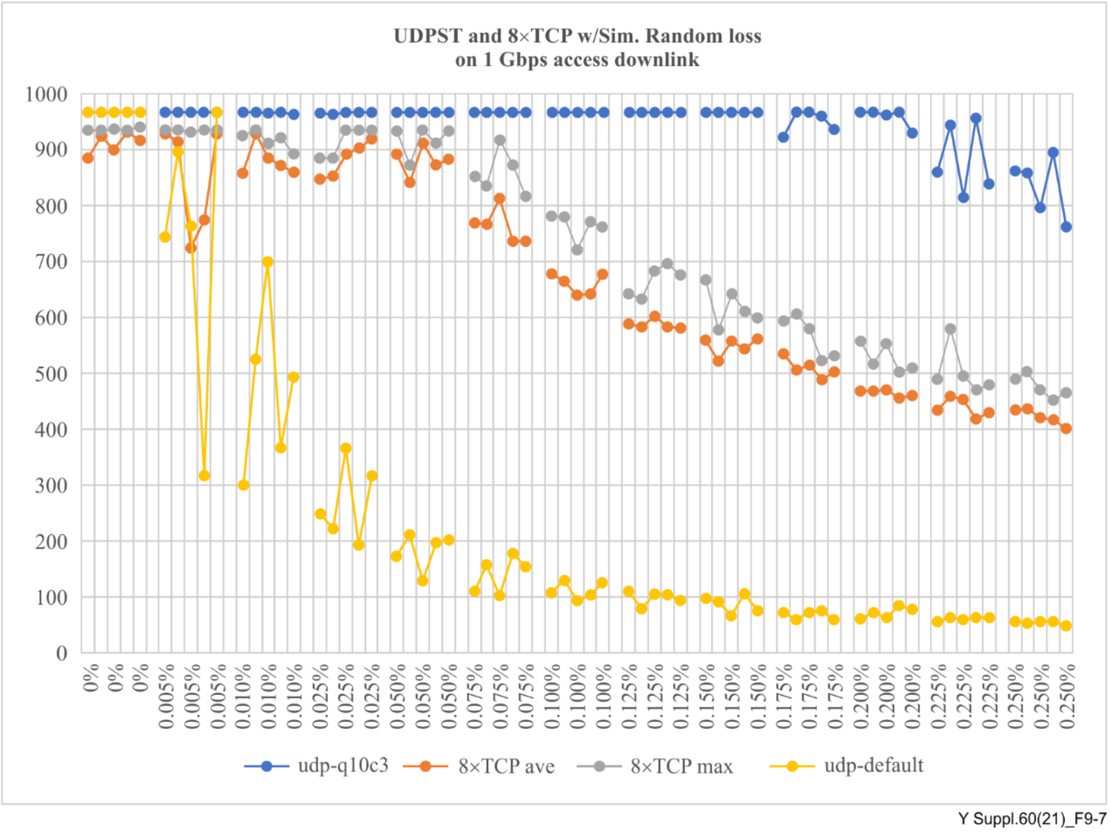

Figure 9-7 – Results for random non-congestion-related loss on downlink access

Figure 9-7 shows the results for four repeated tests at each random loss ratio, from 0% to 0.250%, for each udpst configuration and the two TCP statistics. The measured results are plotted in Mbit/s.

Figure 9-7 clearly indicates the sensitivity of the original default values in udpst to random loss, beginning at 0.005% packet loss ratio. The 8 TCP connections exhibit similar but lower sensitivity overall, especially the 1 second maximum Mbit/s, but trouble is apparent at 0.010%. However, the "udp-q10c3" exhibited robustness to random loss up to 0.175% (1.75×10^{-3}), and represents significant progress toward the goal of maintaining accurate assessment of the maximum IP-layer capacity under these conditions.

This study also examined emulated burst packet loss, generated by NetEm using the Gilbert-Elliot (G-E) model parameters shown in Figure 9-8.

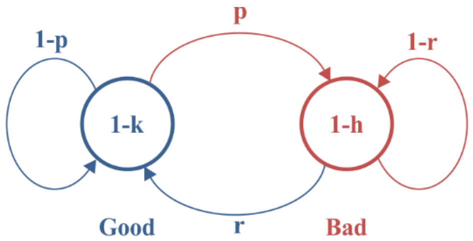

p	r	1-h	1-k
0.001%	1%	95%	0.01%
0.002%	1%	95%	0.01%
0.003%	1%	95%	0.01%

Y Suppl.60(21)_F9-8

Figure 9-8 – Gilbert-Elliot Parameters for Loss Patterns on downlink access

The three parameter sets result in approximately 4, 5 and 6 bursts during a 10 second test.

Figure 9-9 shows the results for four repeated tests at zero loss and three values of "p", the good–bad state transition probability, for each udpst configuration and the two TCP statistics. The measured results are plotted in Mbit/s.

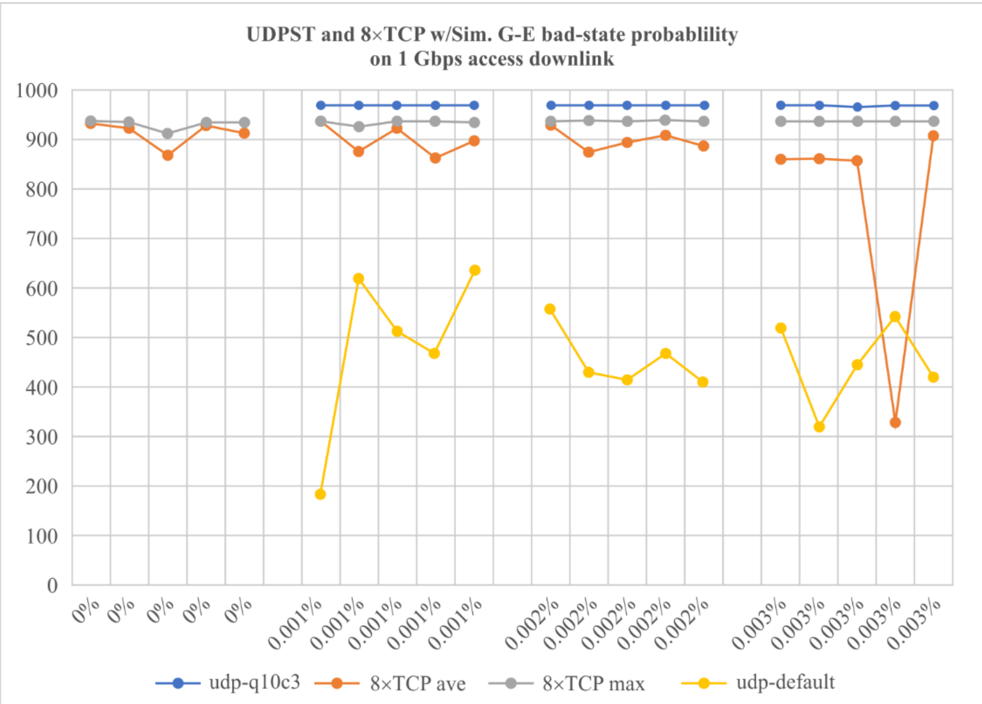

Y Suppl.60(21)_F9-9

Figure 9-9 – Results for bursty non-congestion-related loss on downlink access

Figure 9-9 clearly indicates the sensitivity of the original default values in udpst to the G-E burst loss. The average of 8 TCP connections exhibit similar but lower sensitivity, except for the highest burst transition probability of 0.003%. Note that the loss ratio in the "bad" state is 95%, and some traffic passes successfully.

Again, the "udp-q10c3" exhibited robustness to bursty loss over the range of test conditions, and represents further progress toward the goal of maintaining accurate assessment of the maximum IP-layer capacity. Also, the 1 s maximum TCP was robust under these conditions with 8 connections, but TCP exhibits variation when testing at 1 Gbit/s and zero emulated loss, as the Y.1540 test results and other results in this clause show quite clearly.

On the basis of the testing above and additional investigations in progress, it was decided to revise the default values of the udpst parameters in Supplemental Release 7.2.1 as follows:

SeqErrThresh default of 10 errors (includes loss, reordering and duplication);

SlowAdjThresh default of three consecutive status messages.

The udpst utility is in its fifth public release, 7.2.1, as of August 2021. Additional test results will be included here, in future revisions of this Supplement.

10 Brief description of the udpst utility: an implementation of [BBF TR-471] and [ITU-T Y.1540]

Throughout the development and approval of the maximum IP-layer capacity method of measurement, many tests were conducted to compare the performance of UDP and TCP-based methods of measurement, and various implementations of each transport-layer method. The reference implementation of the method chosen for standardization in [ITU-T Y.1540] and the subsequent development of [BBF TR-471] have now been released in open source form.

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 set-up request to the server;
- Allows the configuration of most of the measurement parameters described in Annexes A and B of [ITU-T Y.1540];
- Allows the user to determine the sub-interval where results for bimodal operation will be split (when the conditions of clause 7.1 have been met).

10.1 udpst results display

An example of udpst client terminal output is illustrated in Figure 10-1.

```
$ udpst -d -a foobar fe80::8639:beff:fe6c:1f90
UDP Speed Test
Software Ver: 6.5, 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, IPv6 TClass: 0
Sub-Interval(sec):  1, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/1, RTTVar(ms): 0-1, Mbps(L3/IP): 96.47
Sub-Interval(sec):  2, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/0, RTTVar(ms): 0-0, Mbps(L3/IP): 299.69
Sub-Interval(sec):  3, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/0, RTTVar(ms): 0-0, Mbps(L3/IP): 502.87
Sub-Interval(sec):  4, Delivered(%): 100.00, Loss/OoO: 0/0, OWDVar(ms): 0/0/0, RTTVar(ms): 0-0, Mbps(L3/IP): 706.02
Sub-Interval(sec):  5, Delivered(%):  99.81, Loss/OoO: 170/0, OWDVar(ms): 0/0/2, RTTVar(ms): 0-2, Mbps(L3/IP): 905.51
Sub-Interval(sec):  6, Delivered(%):  99.88, Loss/OoO: 111/0, OWDVar(ms): 0/0/2, RTTVar(ms): 0-2, Mbps(L3/IP): 970.46
Sub-Interval(sec):  7, Delivered(%):  99.94, Loss/OoO: 61/0, OWDVar(ms): 2/2/2, RTTVar(ms): 2-2, Mbps(L3/IP): 970.81
Sub-Interval(sec):  8, Delivered(%):  99.97, Loss/OoO: 32/0, OWDVar(ms): 2/2/2, RTTVar(ms): 2-2, Mbps(L3/IP): 970.81
Sub-Interval(sec):  9, Delivered(%):  99.96, Loss/OoO: 36/0, OWDVar(ms): 2/2/2, RTTVar(ms): 2-2, Mbps(L3/IP): 971.09
Sub-Interval(sec): 10, Delivered(%):  99.97, Loss/OoO: 30/0, OWDVar(ms): 2/2/2, RTTVar(ms): 2-2, Mbps(L3/IP): 970.80
Downstream Summary Delivered(%):  99.95, Loss/OoO: 440/0, OWDVar(ms): 0/1/2, RTTVar(ms): 0-2, Mbps(L3/IP): 736.45
Downstream Minimum One-Way Delay(ms): 2 [w/clock difference], Round-Trip Time(ms): 0
Downstream Maximum Mbps(L3/IP): 971.09, Mbps(L2/Eth): 984.92, Mbps(L1/Eth): 1000.28, Mbps(L1/Eth+VLAN): 1003.36
```

Figure 10-1 – udpst client results using IPv6 addresses with 1 Gbit/s Ethernet local connection

Description:

- The first five 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 6-2.
- 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.
- Mbps (L3/IP): rates are measured, lower-layer protocol rates are calculated.
- Mbps (L2/Eth), Mbps (L1/Eth), and Mbps (L1/Eth+VLAN) and other rates are calculated values, and usually align with relevant provisioned or theoretical maximums.

10.2 udpst project location

The udpst project is part of the Open Broadband series of projects, and available at the URL below:

https://github.com/BroadbandForum/obudpst

Bibliography

[b-Mathis] Mathis. M. (2019), *How should capacity measurement interact with shaping?* IETF IPPM-List Communications.
https://mailarchive.ietf.org/arch/msg/ippm/KTd7oNve2UXZ3HYuI5cz8eMmdu4/

[b-nntool] net-neutrality-tools/nntool https://github.com/net-neutrality-tools/nntool

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	Tariff and accounting principles and international telecommunication/ICT economic and policy issues
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Environment and ICTs, climate change, e-waste, energy efficiency; construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling, and associated measurements and tests
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	**Global information infrastructure, Internet protocol aspects, next-generation networks, Internet of Things and smart cities**
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2021