

TR-471
**Maximum IP-Layer Capacity Metric, Related Metrics,
and Measurements**

Issue: 4
Date: September 2024

Notice

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

Intellectual Property

Recipients of this Technical 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 Technical Report, or use of any software code normatively referenced in this Technical 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 Technical 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 Technical 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 Technical Report except to the extent this Technical 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 Technical Report are not deemed to be derivative works of the Technical Report.

2. NO WARRANTIES

THIS TECHNICAL REPORT IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT AND ANY IMPLIED WARRANTIES ARE EXPRESSLY DISCLAIMED. ANY USE OF THIS TECHNICAL REPORT SHALL BE MADE ENTIRELY AT THE USER'S OR IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY USER, IMPLEMENTER, OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS TECHNICAL REPORT, INCLUDING BUT NOT LIMITED TO, ANY CONSEQUENTIAL, SPECIAL, PUNITIVE, INCIDENTAL, AND INDIRECT DAMAGES.

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 TECHNICAL REPORT IN ITS CURRENT, OR IN ANY FUTURE FORM. IF ANY SUCH RIGHTS ARE DESCRIBED ON THE TECHNICAL 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.

All copies of this Technical Report (or any portion hereof) must include the notices, legends, and other provisions set forth on this page.

Issue History

Issue Number	Issue Date	Issue Editor	Changes
1	July 2020	Al Morton, AT&T	Original
2	December 2021	Al Morton, AT&T	Expanded configuration (Table 1); Reordering (Table 2); Revised default values in Table 3; New clause 5.2.2; New results specified in Table 4; New Diagnostic Test Metrics in Table 5; Software version and control protocol version now in Table 6.
3	December 2022	Al Morton, AT&T	Expanded configuration (Table 1); UDP Payload Default, Start Sending Rate Index, Start Sending Rate, Table 3); Rate Adj Algorithm, New Annex B.
4	September 2024	Lincoln Lavoie, UNH-IOL	Updates to enable multi-flow and multi-server testing.

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

Editor: Al Morton, AT&T¹

Copy Editors: Lincoln Lavoie, UNH-IOL & Len Ciavattone, AT&T

Work Area Director(s): Jonathan Newton, Vodafone

Project Stream Leader(s): Fabio Giudici, OutSys

¹ In memory of Al Morton, who passed away on June 9, 2023.

Table of Contents

1	Purpose and Scope.....	7
1.1	Purpose	7
1.2	Scope	7
2	References and Terminology	8
2.1	Conventions.....	8
2.2	References	8
2.3	Definitions	9
2.4	Abbreviations	9
3	Technical Report Impact	11
3.1	Energy Efficiency	11
3.2	Security.....	11
3.3	Privacy.....	11
4	Maximum IP-Layer Capacity Metric – related definitions.....	12
4.1	Population of Interest	12
4.2	IP packet sending bit rate Metric Definition	12
4.3	Maximum IP-Layer Capacity Metric Definition.....	12
5	Methods of Measurement for the IP-Layer Capacity Metric	14
5.1	Measurement Reference Architecture.....	14
5.2	Measurement Procedure.....	15
5.2.1	<i>Sending Rate Search Algorithm</i>	24
5.2.2	<i>Lost Status Backoff Timeout and Interaction with the “B” Search Algorithm</i>	27
6	Results	28
Annex A:	Sending Rates Table	35
A.1	Example Table of Sending Rates and Related Information	35
Annex B:	Type C Sending Rate Search Algorithm	36
B.1	Sending Rate Search Algorithm Improvements.....	36
Annex C:	Multi-flow Operation	39
C.1	New Input Parameters for Multi-flow Operation.....	39

Table of Figures

Figure 1: Illustration of Time and Trial results during a single search for Maximum IP-Layer Capacity in the measurement process.	12
Figure 2: Maximum IP Capacity measurement reference architecture	14
Figure 3: IP Edge to CPE measurement of Maximum IP Capacity	15
Figure 4: Flowchart for Offered Load adjustment as part of the “B” Search Algorithm	24

Table of Tables

Table 1: Measurement Variables, Ranges, and Default values	15
Table 2: Parameters for Metrics and Supporting Test Protocols.....	21
Table 3: Flowchart Variables, Descriptions, Ranges, and Default values.....	24
Table 4: Data Model Definition for Results	28
Table A.1: Example Table of Sending Rates.....	35
Table B.2: Table of Additional Flowchart variables, descriptions, ranges and default values	37
Table C.3: Table of Input Variables and Thresholds for Multi-flow Operation.....	39

Executive Summary

As the speed of access links increase into the Gigabit range for some technologies, the dominant methods of helping users test to discover their actual “Internet Speed” are showing their age. These TCP-based methods, even with CUBIC window flow control, reacting conservatively to loss and round-trip delay often produce a significant underestimate of Maximum IP-Layer Capacity.

Meanwhile, the Industry is seeing a transition to new Transport protocols that will replace TCP for many users. They use UDP as the Transport protocol, and encrypt activity above the transport layer (such as IETF QUIC). Measuring the IP-Layer Capacity on a User’s access link should be straightforward and use the same Transport protocols as Users (many use UDP now, more in the future).

This document defines a Maximum IP-layer Capacity Metric and Methods of Measurement whose accuracy has been carefully evaluated with lab and field measurements.

In Issue 2, several enhancements appear in the specification stemming from testing and development in the Open Broadband Project (OB-UDPST) and elsewhere. Some access technologies are configured to offer an initial “fast” mode where the Maximum IP-Layer Capacity delivered is higher than the subsequent Maximum level offered to the same flow later in its lifetime. Repeatable observations of this bimodal behavior are sufficient cause to report the results for each mode separately. Also, Issue 2 includes options whether to process measurements of packet stream sequence discontinuities (such as reordering and duplication/replication) in the load adjustment algorithm, in case these forms of impairment are prevalent in the test stream and contribute a notable fraction of IP-Layer Capacity.

In Issue 3, the main enhancement is the specification of a new load adjustment/search algorithm in an Annex. This optional algorithm, Type C, can be briefly described as “Multiply and Retry”. The “fast” ramp-up is now a multiplicative rate increase to the congestion condition, reaching 1Gbps in ~1 second if the access supports it. The “fast” ramp-up can be re-tried when conditions warrant, to ensure that the Maximum IP-Layer Capacity is reached during a test. Both these attributes are valuable when radio channels are in the path or when other dynamic conditions exist. Type C algorithm supports a search over the full range of rates, even if the subscribed rate is <1Gbps. The Type B algorithm remains the default, for testing that does not benefit from “Multiply and Retry” aspects. Other new features include implementation of the random payload option for paths that use compression, and Server bandwidth management; the server can accept a bandwidth/capacity limit for test admission-control. Finally, the default load adjustment operation now ignores reordered packets during the search for maximum capacity (reordered packets are legitimate contributors to IP-Layer Capacity).

In Issue 4, Annex C of the specification adds the new parameters for Multi-flow and Multi-Server testing capability. This feature is implemented in Release 8.0.0 of OB-UDPST. The body text is updated for this capability in several areas.

1 Purpose and Scope

The Purpose and Scope sections describe the work on Active measurements of Maximum IP-Layer Capacity.

1.1 Purpose

The main goal is to harmonize the specified definitions, metrics and methods across the industry, and this project is the vehicle through which the BBF consensus is captured and communicated to achieve broad agreement, and possibly resulting in changes in similar the specifications of other Standards Development Organizations (SDO).

A goal unique to BBF is to provide efficient test procedures, especially the locations of measurement points, and to recommend results reporting with a view to data models for the results. These procedures are developed with an eye towards future remote control, data collection, and automation through TR-069 or TR-369 (USP).

1.2 Scope

The scope of this TR is to provide definitions, define metrics and the corresponding methods to unambiguously perform active measurements of Maximum IP-Layer Capacity. This includes the definition of the metric's parameters that influence how the test stream is generated and the metric is measured, resulting in an assessment of the Maximum IP-Layer Capacity over the (short) test interval. The scope includes any related IP-layer metrics as required by the method.

2 References and Terminology

2.1 Conventions

In this Technical Report, several words are used to signify the requirements of the specification. These words are always capitalized. More information can be found in RFC 2119 [9].

MUST	This word, or the term “REQUIRED”, means that the definition is an absolute requirement of the specification.
MUST NOT	This phrase means that the definition is an absolute prohibition of the specification.
SHOULD	This word, or the term “RECOMMENDED”, means that there could exist valid reasons in particular circumstances to ignore this item, but the full implications need to be understood and carefully weighed before choosing a different course.
SHOULD NOT	This phrase, or the phrase “NOT RECOMMENDED” means that there could exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications need to be understood and the case carefully weighed before implementing any behavior described with this label.
MAY	This word, or the term “OPTIONAL”, means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option.

2.2 References

The following references are of relevance to this Technical Report. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this Technical 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] TR-069 Amendment 6	CPE WAN Management Protocol	BBF	2019
[2] TR-101 Issue 2	Migration to Ethernet-Based Broadband Aggregation	BBF	2011
[3] TR-143	Enabling Network Throughput Performance Tests and Statistical Monitoring, Issue: 1 Amendment 1 Corrigendum 1	BBF	2015
[4] TR-181	Device Data Model for TR-069	BBF	2015
[5] TR-304	Broadband Access Service Attributes and Performance Metrics	BBF	2015
[6] TR-369	User Services Platform (USP)	BBF	2019
[7] Y.1540	Internet protocol data communication service - IP packet transfer and availability performance parameters (and Amendment 1)	ITU-T	2019
[8] Y.1565	Home network performance parameters	ITU-T	2011

[9] RFC 2119	Key words for use in RFCs to Indicate Requirement Levels	IETF	1997
[10] RFC 2681	A Round-trip Delay Metric for IPPM	IETF	1999
[11] RFC 3393	IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)	IETF	2002
[12] RFC 4737	Packet Reordering Metrics	IETF	2006
[13] RFC 5481	Packet Delay Variation Applicability Statement	IETF	2009
[14] RFC 6335	Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry	IETF	2011
[15] RFC 6576	IP Performance Metrics (IPPM) Standard Advancement Testing (a.k.a. BCP 176)	IETF	2012
[16] RFC 7680	A One-Way Loss Metric for IP Performance Metrics (IPPM) (a.k.a. STD 82)	IETF	2016
[17] RFC 5560	A One-Way Packet Duplication Metric	IETF	2009
[18] RFC 7950	The YANG 1.1 Modeling Language	IETF	2016

2.3 Definitions

The following terminology is used throughout this Technical Report.

Sender	The test endpoint which sends test packets to the Receiver and which receives status messages from the Receiver.
Receiver	The test endpoint which receives test packets from the Sender and which sends status messages to the Sender.
Test Endpoint	The host at either end of the network test path. A test endpoint can be either a Sender or a Receiver for any given test.
Test Controller	A host which configures the test endpoints for a test. The test endpoint which performs the sending rate search algorithm during the test and which reports the measurement results. Either the Sender or the Receiver can be the Test Controller for any given test.

2.4 Abbreviations

This Technical Report uses the following abbreviations:

CUBIC	A widespread form of congestion control for transport-layer protocols
DSCP	Diffserv Code Point
Dst	Destination
IETF	Internet Engineering Task Force
IPDV	IP Packet Delay Variation
IPRR	IP Packet Reordered Ratio
IPSBR	IP Packet Sending Bit Rate
PM	Performance Metrics
QUIC	A new transport protocol under development in the IETF (work-in-progress)
RIPR	Replicated IP Packet Ratio
Src	Source
TCP	Transmission Control Protocol

TR	Technical Report
UDP	User Datagram Protocol
WA	Work Area
WT	Working Text

3 Technical Report Impact

3.1 Energy Efficiency

Similar to TR-143 [3], this Technical Report has no impact on energy efficiency.

3.2 Security

Similar to TR-143 [3], this Technical Report has no additional impact on Security.

3.3 Privacy

Similar to TR-143 [3], this Technical Report has no additional impact on Privacy.

4 Maximum IP-Layer Capacity Metric – related definitions

Supporting definitions and conventions are provided in this section.

4.1 Population of Interest

Population of Interest: All of the performance metrics defined in this document require a specific set of packets called the population of interest. For the end-to-end 1-way measurement case, the population of interest is (usually) the total set of measurement packets being sent from the sender to the receiver. The measurement points in the end-to-end 1-way case are at the sender and receiver. This definition is equivalent to the sum of qualifications used in IETF RFCs developed by the IPPM working group, where Parameters such as the packet stream characteristics and the Type-P describing the make-up of packets in the measurement stream sent from Src (source, or Sender) to Dst (destination, or Receiver) constitute the population of interest. For bi-directional testing, the populations may be different by direction, and require a separate specification applying to each direction.

4.2 IP packet sending bit rate Metric Definition

For a given population of interest, the IP packet sending bit rate (IPPSBR) generated by the Source host is the number of bits in the IP packet headers and payloads resulting in IP packet reference events divided by the time interval duration.

4.3 Maximum IP-Layer Capacity Metric Definition

An illustration of the testing process is given below, as viewed by the Receiver.




Figure 1: Illustration of Time and Trial results during a single search for Maximum IP-Layer Capacity in the measurement process.

The measurement process can be characterized as a search for the Maximum IP-Layer Capacity, where the Sender bit-rate is adjusted according to feedback from the receiver in an effort to discover the Maximum IP-Layer Capacity along with the performance in terms of packet loss and delay. The test packets arrive at the

receiver during a corresponding sub-interval dt_n , (equivalent to $[dt_n, dt_{n+1}]$), of which m sub-intervals comprise the complete measurement interval, $[t, t + \Delta t]$.

During each sub-interval, dt_n , there are many adjustments of the sending rate and subsequent measurement feedback reports from the receiver, called Trials. Each Trial report is the input to a sender load-adjustment algorithm. The search algorithm has a measurement goal, such as minimal loss and delay variation less than a configured value. These adjustments apply to type Search, whereas type Fixed-rate runs at a fixed bit rate.

For each sub-interval, dt_n , the IP-Layer Capacity is calculated as the number of IP-layer bits received (n_0) divided by the time, dt , for a result in bits per second. Corresponding values of Packet loss and delay are summarized and associated with each Capacity.

When the time for a search is completed, the IP-layer Capacity values are evaluated for the maximum in the test interval, and reported with the corresponding values of Packet loss and delay, with the condition that the target threshold for at least one fundamental metric has been met in that sub-interval, dt_n .

Mathematically, we express this process as follows:

For a given population of interest, the maximum IP-layer capacity during time interval $[t, t + \Delta t]$ is:

$$\text{Maximum_C}(t, \Delta t, PM) = \frac{\max_{[t, \Delta t]}(n_0(dt_n, dt_{n+1}))}{dt}$$

Equation 1: Maximum IP Capacity Metric

where:

time interval $[t, t + \Delta t]$ is composed of m equal sub-intervals, dt in length;

PM is a list of fundamental metrics, such as loss, delay, and reordering, and a corresponding Target performance threshold. At least one fundamental metric and Target performance threshold must be supplied (such as One-way IP Packet Loss [RFC7680] equal to zero).

n_0 is the total number of IP-layer header and payload bits that can be transferred over a basic section generating successful IP packet transfer outcomes at the egress measurement point during a specified time interval, $[dt_n, dt_{n+1}]$, and

the Maximum $C(t, \Delta t)$ corresponds to the maximum value of n_0 measured in any sub-interval $[dt_n, dt_{n+1}]$ within time interval $[t, t + \Delta t]$, divided by the duration of the sub-interval.

Note that UDP transport shall be used when assessing the Measurable IP Capacity Metric.

The method of measurement also needs a definition for its sending rate, supplied in section 4.2.

5 Methods of Measurement for the IP-Layer Capacity Metric

Methods of measurement are provided in this section. The Methods have sufficiently detailed specifications for implementation that will produce statistically equivalent results (as described in [IETF RFC 6576] from IPPM work). The Methods are described independently of the data modeling language or the protocols used to organize and transmit information between the different elements in the test architecture.

5.1 Measurement Reference Architecture

The reference architecture for measuring Maximum IP Capacity is shown in Figure 2. The primary functional elements in the architecture are the Sender and Receiver, denoting the transmitting and receiving endpoints of the test path through the network being measured. Sender and Receiver are functional terms defined by both the network path to be tested and the direction of the test, in that the Sender for a test in one direction becomes the Receiver for a test across the same network path in the other direction. A physical endpoint supporting Maximum IP Capacity will support both Sender and Receiver functions.

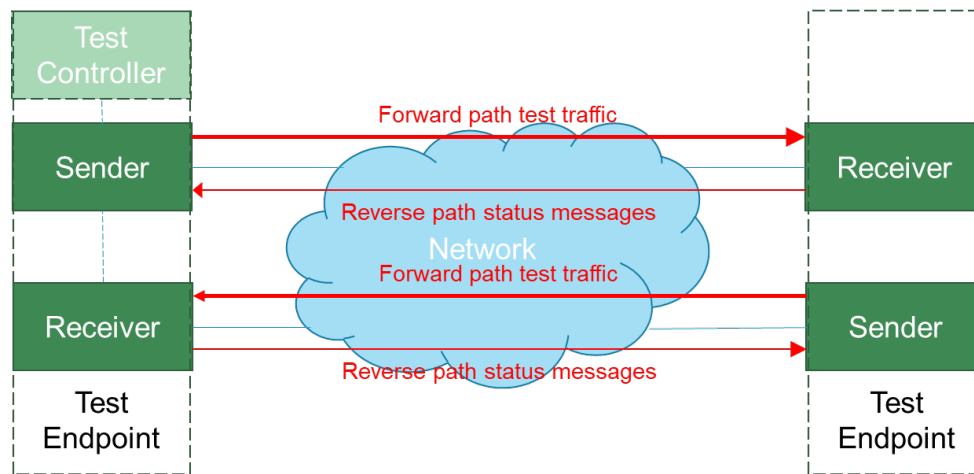


Figure 2: Maximum IP Capacity measurement reference architecture

A Test Controller forms a third element.

The Test Controller receives requests to perform a test from one of the Test Endpoints, providing test parameters prior to test initialization and receiving measurement results at the conclusion of the test. The Test Controller will typically communicate with the same physical endpoint, even as the testing roles are reversed in the course of testing both directions of a network path.

During a test, the Test Controller continually communicates with the Sender and the Receiver, to learn the current measurement results from status messages and increase or decrease the test traffic sending rate. As a result, it is essential to locate the Test Controller sufficiently near one of the Test Endpoints, to reduce the delay in the feedback loop from status message generation to sending rate modification, and to allow more accurate Round-Trip-Time (RTT) measurements.

The reference architecture shown in Figure 1 can be applied at any two nodes that can communicate with each other across an IP-based network path. It can, for example, be applied at two devices within a home network and controlled via USP [TR-369].

Another example, in which a test platform located at the IP edge is used to measure the Maximum IP Capacity between its location and the CPEs served by that access network, is shown in Figure 3 (partly adapted from TR-304). When activating the Multi-flow and/or Multi-server capability, one Test Endpoint (the CPE in Figure 3) may request more than one test flow, and may contact multiple Test Endpoints in the network to share the load of the test flows. See Annex C for details.




Figure 3: IP Edge to CPE measurement of Maximum IP Capacity

5.2 Measurement Procedure

Each measurement of Maximum IP Capacity occurs in a specific direction (for example, IP Edge to CPE) across a network test path. A typical test process between two endpoints may include measuring the Maximum IP Capacity in both directions; however, the measurement procedure in each direction is conducted independently of the measurement procedure in the other direction. Multiple independent one-directional tests **MUST NOT** be performed concurrently between the same endpoint pair.

Before the procedure begins, parameters needed to conduct the test are configured in both the Sender and the Receiver. The test parameters are defined in Table 1.

Table 1: Measurement Variables, Ranges, and Default values

Name	Parameter	Unit	Type; Range	Write	Default value
Connection Category:					
Interface	The IP-layer interface over which the test is to be performed. Example: Device.IP.Interface.1 If an empty string is specified, the device MUST use the interface as directed by its routing policy (Forwarding table entries) to determine the appropriate interface.	NA	string(256)	W	No Default
Role	Configuration for the measurement system's role in the method, either Sender or Receiver	NA	enumeration [Receiver, Sender]	W	No Default
Host	Host name or address of the remote host to perform tests to. This		string(256)		No Default

	parameter is deprecated in Issue 4, see Annex C.				
Port	Port on the remote host to perform tests to. This parameter is deprecated in Issue 4, see Annex C.		unsignedInt-[1:65535]		No Default
JumboFramesPermitted	Configuration for the measurement system permitting use of Jumbo-length Frames	NA	Boolean:[0:1]	W	1(True: permitted for sending rates above 1Gbps) See Note 3
NumberOfConnections (see Note 1)	Number of parallel connections in a test. This parameter is deprecated in Issue 4, see Annex C.	#	unsignedInt; 1 ≤ # ≤ 10, or supported Maximum if less than 10	W	1 connection
EthernetPriority	Ethernet priority code for marking packets transmitted in the test (if applicable).		unsignedInt; [0:7]	W	0
DSCP	DSCP markings for special treatment in the network		unsignedInt; [0:63]	W	0 = Best Effort
ProtocolVersion	Indicates the IP protocol version to be used. The default value SHOULD be Any. Enumeration of: <ul style="list-style-type: none"> Any (Use either IPv4 or IPv6 depending on the system preference) IPv4 (Use IPv4 for the requests) IPv6 (Use IPv6 for the requests) 		string;	W	Any
UDPPayloadMin	Minimum reference size of UDP payload in <i>octets</i> . No default. The implementation will algorithmically determine a value if none is configured. The value SHOULD be set to a value that avoids fragmentation (i.e., using path MTU discovery).	octets	unsignedInt; (35:8972)	W	No default
UDPPayloadMax	Maximum reference size of UDP payload in <i>octets</i> . No default. The implementation will algorithmically	octets	unsignedInt; (35:8972)	W	No default

	determine a value if none is configured. The value SHOULD be set to the largest value that avoids fragmentation (i.e., using path MTU discovery). If JumboFramesPermitted is false, the maximum value MUST be 1472 octets. If JumboFramesPermitted is true, this value can be as large as 8972 octets. UDPPayloadMax MUST be greater than or equal to UDPPayloadMin.				
UDPPayloadContent	<p>UDP Payload Content Type, If there is payload compression in the path and tests intend to characterize a possible advantage due to compression, then payload content SHOULD be supplied by a pseudo-random sequence generator, by using part of a compressed file, or by other means. Payload may also contain the test protocol PDUs.</p> <p>Enumeration of:</p> <ul style="list-style-type: none"> • ones, • zeroes, • alternates0and1 • random 	NA	string;	W	zeroes
PortRange	REQUIRED range of ports supported for test traffic and status feedback messages	#	unsignedInt; The Dynamic Ports, also known as the Private or Ephemeral Ports, from 49152-65535 (never assigned by IANA as per RFC 6335)	W	No default
PortRangeOptional	OPTIONAL range of ports supported for test traffic and status feedback messages	#	unsignedInt; The User Ports, also known as the Registered Ports, from 1024-49151 (assigned by IANA, RFC 6335)	W	No default. Use of Port(s) from the User range and its impact on the

					network needs to be carefully monitored by the testing organization before using the ports in testing, to avoid collision and contention.
Test Category:					
TestType	Type of test, Search or Fixed-Rate		enumeration; search or fixed	W	Search
EnableIPDV	Configuration for the measurement system permitting One-way measurement of IPDV as per [Y.1540]	NA	Boolean;[0:1]	W	0 (False: Use RTT= round-trip delay variation in the load rate adjustment algorithm) (non-default is 1=True EnableIPDV which uses one-way delay variation for the load rate adjustment algorithm)
EnableIPRR	Configuration for the measurement system permitting measurement of IPRR as per [Y.1540]	NA	Boolean;[0:1]	W	0 (False: not enabled)
EnableRIPR	Configuration for the measurement system permitting measurement of RIPR as per [Y.1540]	NA	Boolean;[0:1]	W	0 (False: not enabled)
PreambleDuration (see Note 2)	Duration of active traffic preamble to testing.	seconds	unsignedInt; $0 \leq \text{seconds} \leq 5$	W	2 seconds
StartSendingRateIndex	Configurable initial sending rate index (INDEX to a row of the sending rate table), for the Fixed or Search algorithm. StartSendingRateIndex overrides StartSendingRate if both are present.	#	unsignedInt; $0 \leq \# \leq 11108$ (40 Gbps)	W	default is 0 = nominally 500kbps

StartSendingRate	Configurable initial Sending Rate for a Fixed test or the initial Sending Rate value for a Search test. Value specified in kbps. The default value SHOULD be 500 kbps. StartSendingRateIndex overrides StartSendingRate if both are present.	Kbps, where $b=1$ bit	unsignedInt; $500 \leq \# \leq$ nominally 40 Gbps	W	Nominally 500 Kbps
Δt (TestInterval)	Duration of the test (either downlink or uplink) with search algorithm in use, which serves as the maximum duration of the search process	seconds	unsignedInt; $5 \leq \text{seconds} \leq 60$	-	Computed from m and dt , Type;Range is a constraint on $m * dt = \Delta t$
m (NumberTestSubIntervals)	Number of intermediate measurement intervals, dt , in Δt	#	unsignedInt; $1 \leq \# \leq 100$	W	10
i (NumberFirstModeTestSubIntervals)	Number of measurement intervals, dt , included in the report of the initial Capacity mode (1 and higher). The remaining sub-intervals of the total m are reported separately. "0" is used to replace the EnableBimodal parameter, and means the Bimodal analysis is NOT enabled.	#	unsignedInt; $0 \leq \# < m$	W	Note: m is the practical limit for a consistent test, and 100 is an absolute limit
dt (TestSubInterval)	Duration of intermediate reporting intervals	ms	unsignedInt; $100 \leq \text{ms} \leq 6000$ (max $\Delta t/(m=10)$) in milliseconds) MUST meet Type;Range constraints on $\Delta t = m * dt$.	W	1000 ms
StatusFeedbackInterval	Period of status feedback message (Receiver of offered load returns messages to the sender with the results of the measured metrics)	ms	$20 \leq \text{ms} \leq 250$	W	50 ms
TimeoutNoTestTraffic	Timeout value, no test packets at Receiver since previous test packet	ms	unsignedInt; $500 \leq \text{ms} \leq 1000$	W	1000ms This value is consistent with a 10 sec test duration.

TimeoutNoStatusMessage	Timeout value, no Status Messages at Sender since previous Status Message	ms	unsignedInt; 500 ≤ ms ≤ 1000	W	1000 ms Comment: Take Sender action to reduce rate before stopping test. This value is consistent with a 10 sec test duration.
Tmax	Maximum one-way Waiting time for packets to arrive	ms	unsignedInt; 50 ≤ ms ≤ 3000	W	1000 ms
TmaxRTT	Maximum Round Trip Waiting time for packets to arrive	ms	unsignedInt; 50 ≤ ms ≤ 3000	W	3000 ms
TimestampResolution	This is a function of the measurement protocol, and it is usually determined once the protocol chosen			-	See Table 2 for information on required resolution

Note 1: For NumberOfConnections: Aggregate measurement results are for connections >1.

Note 2: For PreambleDuration: Traffic sent between sender and receiver at the StartingRate.

Note 3: For JumboFramesPermitted: The table of sending rates will change, depending on the option chosen.

In addition, the following should be included in External Model Definitions:

- Max number of supported connections (by a specific Test Endpoint).
- Interface specification: The IP-layer interface over which the test is to be performed.
- Max number of sub intervals (m)

During the measurement procedure, the Sender transmits a stream of test packets to the Receiver at a sending rate that is periodically updated to find the maximum rate at which packets are delivered while meeting a set of predefined performance requirements. The test interval Δt is subdivided into a number of sub-intervals dt , and each sub-interval is further divided into a number of trial intervals. At the end of each trial interval, the Receiver transmits a status message to the Sender with information used to update the sending rate.

The search algorithm that updates the sending rate can run on either the Sender or the Receiver. The endpoint that runs the search algorithm for a specific test is referred to as the Test Controller. In addition to running the search algorithm, the Test Controller calculates the Maximum IP Capacity as well as any secondary metrics and makes the measurement results available at the test's conclusion. At least one endpoint in a given test needs to support the Test Controller functions.

[R-1] The test endpoint MUST support both Sender and Receiver functions.

[R-2] The test endpoint MUST support the test parameters listed in Table 1.

[R-3] At least one test endpoint MUST support Test Controller functions.

Each packet sent during the test includes a sending time stamp and sequence number. The Receiver uses the time of reception for each packet along with its sending time stamp and sequence number to measure received rate, packet loss, and one-way packet delay variation per Recommendation Y.1540. The Receiver then may

run an iteration of the search algorithm (if it is acting as the Test Controller) before sending a status feedback message to the Sender. Depending on which endpoint is acting as the Test Controller:

- If the Receiver is acting as the Test Controller, it populates the status feedback message with the updated sending rate generated by the search algorithm.
- If the Sender is acting as the Test Controller, the Receiver populates the status feedback message with the received rate, packet loss, and one-way packet delay variation measurements calculated for the most recent trial interval. The Sender then uses this information to run the search algorithm to update its sending rate.

In addition to the above metrics, the Sender periodically measures round-trip delay using the following procedure:

- When the Receiver has a status feedback message ready to send, it waits for the next test packet to arrive from the Sender.
- When the test packet is received, the Receiver populates the status message with the sequence number and sending time stamp of the test packet and immediately sends the status message to the Sender.
- When the Sender receives the status message, it uses the sequence number and sending time stamp, along with the received time of the status message to calculate the round trip delay. The calculated round trip delay includes time to prepare the periodic status feedback message by the Receiver. An implementation should strive to minimize that processing time or use other measurement methods to evaluate round-trip delay.

[R-4] The Receiver **MUST** measure (using the sequence number and send and receive timestamps) the following metrics during each trial interval: received rate and loss, specified as per Recommendation Y.1540 [Y.1540].

[R-5] The Sender **MUST** measure sampled Round-trip time (RTT) as per RFC 2681 [RFC2681] or Recommendation Y.1565 [Y.1565], based on periodic status feedback messages from the Receiver.

[R-6] The Receiver **MAY** measure (using the sequence number and send and receive timestamps) the following metrics during each trial interval: 1-way packet delay variation, reordering, and duplication, specified as per Recommendation Y.1540 [Y.1540]], and/or the IETF RFC specifications tabulated below.

Parameter values for metrics in R-4 through R-6 follow:

Table 2: Parameters for Metrics and Supporting Test Protocols

Metric	Parameter	Description	Unit	Range	Default value
IPLR (Loss Ratio) Y.1540, RFC 7680	Tmax	Maximum Waiting time for packets to arrive	ms	$50 \leq \text{ms} \leq 3000$	1000 ms
Sampled RTT Y.1565, RFC 2681: RTT uses feedback status	TmaxRTT	Maximum Round Trip Waiting time for packets to arrive	ms	$50 \leq \text{ms} \leq 3000$	3000 ms

messages from receiver	TimestampResolution (Determined when test protocol is chosen)	Resolution of Timestamps	ms	$0.001 \leq \text{ms} \leq 1$	Suggested for fixed access: .001 (based on current implementation)
IPDV Y.1540, RFC3393, RFC5481 (PDV form is specified here)	Tmax	Maximum Waiting time for packets to arrive	ms	$50 \leq \text{ms} \leq 3000$	1000 ms
	TimestampResolution (Determined when test protocol is chosen)	Resolution of Timestamps	ms	$0.001 \leq \text{ms} \leq 1$	Suggested for fixed access: .001 (based on current implementation)
IPRR Y.1540, RFC4737	Tmax	Maximum Waiting time for packets to arrive	ms	$50 \leq \text{ms} \leq 3000$	1000 ms
	TimestampResolution (Determined when test protocol is chosen)	Resolution of Timestamps	ms	$0.001 \leq \text{ms} \leq 1$	Suggested for fixed access: .001 (based on current implementation)
RIPR Y.1540, RFC5560	Tmax	Maximum Waiting time for packets to arrive	ms	$50 \leq \text{ms} \leq 3000$	1000 ms
	TimestampResolution (Determined when test protocol is chosen)	Resolution of Timestamps	ms	$0.001 \leq \text{ms} \leq 1$	Suggested for fixed access: .001 (based on current implementation)

If the Test Controller is acting as Sender:

- [R-7] The Receiver MUST send a status feedback message to the Sender with the results of the measured metrics at the end of each trial interval.
- [R-8] The Sender MUST update its sending rate per the configured sending rate search algorithm after receiving a status message.
- [R-9] The Sending rate SHOULD be measured using the IP Packet Sending bit Rate (IPSBR) metric defined in section 4.2.

If the Test Controller is acting as Receiver:

- [R-10] The Receiver MUST update the Sender's sending rate per the configured sending rate search algorithm at the end of each trial interval.
- [R-11] The Sending rate SHOULD be measured using the IP Packet Sending bit Rate (IPSBR) metric defined in Annex A/Y.1540.
- [R-12] The Receiver MUST send a status feedback message to the Sender with the updated sending rate.

Two methods for updating the sending rate during a measurement are specified. The first method is the "Type B" sending rate search algorithm specified in Section 5.2.1. The second method is a fixed sending rate that does not change for the duration of the test. The fixed sending rate can be used to confirm the results of a measurement using the search algorithm, for example using a sending rate set to a high percentage (e.g., 99.x%) of the measured Maximum IP Capacity. Additional optional methods for determining the sending rate can be supported in the Test Controller.

- [R-13] The Test Controller MUST support the “Type B” sending rate search algorithm specified in Section 5.2.1.
- [R-14] The Test Controller MUST support a configurable fixed sending rate.
- [R-15] The Test Controller MAY support one or more additional sending rate search algorithms.
- [R-16] The Test Controller MUST support selection of the sending rate search algorithm.

At the end of the test, the Test Controller generates the Maximum IP Capacity measurement result per the metric definition in Equation 1. A minimal Test Controller can generate intermediate results after each trial interval, updating the result each time a new maximum is found and discarding outdated values. Optionally, the Test Controller can save all intermediate trial results, including IP Capacity as well as packet latency, loss and reordering, for post-measurement analysis. Additional results that can be made available by the Test Controller include:

The Maximum IP Capacity includes all bits in the IP header and extension header fields, the IP payload, which is the UDP header and payload fields. The capacity available to the layers above UDP, which includes the UDP payload but not the UDP header fields, is an additional useful metric that can be calculated from the measurement results and test parameters.

Additional calculations are possible within this measurement framework:

- Calculating an average of all measured values (for IP-Layer Capacity) for all Trials in a sub-interval, dt_n , and for the complete measurement interval $[t, t + \Delta t]$.
 - Calculating an average of all measured values (for IP-Layer Capacity) for all Trials in a sub-interval, dt_n , and for all sub intervals in the complete measurement interval $[t, t + \Delta t]$ where the search goal was met.
 - Calculating a maximum of all measured values of (for IP-Layer Capacity) for all Trials in a sub-interval, dt_n , and for all sub intervals in the complete measurement interval $[t, t + \Delta t]$ where the search goal was met.
 - Calculating an average of all measured values of (for IP-Layer Capacity) for all Trials in a sub-interval, dt_n , and for the complete measurement interval $[t, t + \Delta t]$ for all Trials, where a specified result exclusion criteria has been met (e.g., removal of outliers, as determined by specified criteria).
- [R-17] The Test Controller MUST calculate the Maximum IP Capacity measurement result according to Equation 1.
- [R-18] The Test Controller MUST report the Maximum IP Capacity and the IP packet loss ratio corresponding to the sub-interval dt_n in which the Maximum IP Capacity was measured according to Equation 1.
- [R-19] The Test Controller MUST identify the sending rate search algorithm and parameters used when reporting Maximum IP Capacity measurement results.
- [R-20] The Test Controller MAY include one or more secondary measurement results as listed above when reporting Maximum IP Capacity measurement results.
- [R-21] The Test Controller MAY report the UDP Capacity corresponding to the Maximum IP-Layer Capacity in terms of UDP payload bits delivered.

When post-measurement analysis is performed, for example to analyze the consistency of test results, the process to summarize the results should include corresponding post-test analysis to ensure data quality and to detect and exclude data artifacts (where possible). The post-test analysis methods should be published with the analysis results.

Some conditions may temporarily prevent the generation of a valid Maximum IP Capacity result. Senders and Receivers should support mechanisms to detect such conditions where possible and to prevent the generation of invalid results. Specific conditions, and the mechanisms for detecting them, may be node-specific. For example, a network-located test platform may need to check for host resource availability before accepting a test command, while it may be more important for a residential gateway to check for user traffic on the broadband connection before conducting a test.

5.2.1 Sending Rate Search Algorithm

The “Type B” sending rate search algorithm specified below makes use of a table of sender bit rates listed in order of ascending rate from the minimum to the maximum rate supported by the test. Depending on the Sender implementation, each bit rate may correspond to a number of packets of specified sizes sent during a specified time interval (which may correspond to the timer granularity supported by the Sender).

An iteration of the Type B search algorithm is shown in the flowchart in Figure 4. In that flowchart, one step represents a single row increment or decrement in the sending rates table (see Annex A). The flowchart uses many variable names and in some cases, configurable thresholds that determine the flowchart decisions. There are three main paths through the flowchart based on the following conditions:

- Feedback indicates measured impairments are absent,
- Impairments are first measured and some congestion may be present but sending rate change is deferred,
- Measured impairments are confirmed by repeated measurement feedback.

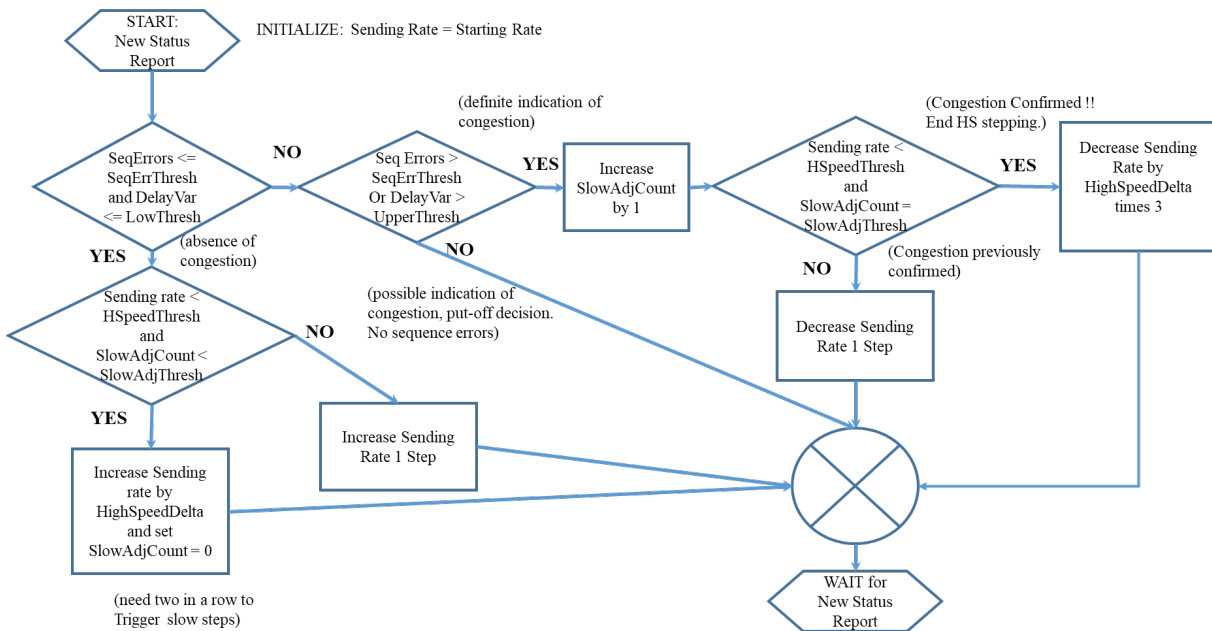


Figure 4: Flowchart for Offered Load adjustment as part of the “B” Search Algorithm

The variables and thresholds used in Figure 2 are explained in Table 3:

Table 3: Flowchart Variables, Descriptions, Ranges, and Default values

Name	Description	Unit	Range	Write	Default value
------	-------------	------	-------	-------	---------------

SendingRate	The current sending rate (equivalent to a row of the table), Initialized at minimum Sending Rate in the Table of Sending Rates	Kbps, where $b=1$ bit	unsignedInt; $500 \leq \# \leq 10,000,000$ (10 Gbps)	-	StartSendingRate
SeqErrors	Measured Count of any of Loss or Reordering or Replication impairments (events where received packet sequence number did not increase by one)	number	unsignedInt; $0 \leq \text{SeqErrors}$	-	NA
SeqErrThresh	Threshold for Loss or Reordering or Replication impairments measured (events where received packet sequence number did not increase by one)	number	unsignedInt; $0 \leq \text{SeqErrThresh} \leq 100$	W	10
ReordDuplIgnoreEnable	When True (enabled) only Loss counts toward received packet sequence number errors, and Reordering and Duplication impairments are ignored. When False, Loss, Reordering and Duplication are all counted as sequence number errors.)	NA	Boolean; [0:1]	W	1 (True: enabled)
DelayVar	Measured Range of Round Trip Time, RTT (or 1-way Packet Delay Variation, above minimum delay when DelayVar 1-way measurements are reliable)	ms	NA	-	NA
LowThresh	Low threshold on the Range of	ms	unsignedInt; $5 \leq \text{ms} \leq 250$	W	30ms default

	Round Trip Time variation, RTT (Range is values above minimum RTT)				
UpperThresh	High threshold on the Range of Round Trip Time variation, RTT (Range is values above minimum RTT)	ms	unsignedInt; $5 \leq \text{ms} \leq 250$	W	90ms default
HighSpeedDelta	The number of rows to move in a single adjustment when initially increasing offered load (to ramp-up quickly)	number of rows	unsignedInt; ≥ 2	W	10 table rows (10 Mbps currently)
RateAdjAlgorithm	Configurable choice of Load Rate Adjustment Algorithm. The Type "B" Algorithm is described in Figure 4 of TR-471.		enumeration; B or C	W	B
SlowAdjCount	Measured Number of consecutive status reports indicating loss and/or delay variation above UpperThreshold.	Count of occurrences	NA	-	See SlowAdjThresh
SlowAdjThresh	Threshold on SlowAdjCount used to infer congestion. Use values >1 to avoid misinterpreting transient loss.	Count of occurrences	unsignedInt; >1	W	3
HSpeedThresh	Threshold for transition between low and high sending rate step sizes (such as 1 Mbps and 100 Mbps). MAY result in use of Jumbo Frames if permitted.	Gbps, where $b=1$ bit	unsignedInt; ≥ 1	W	1 Gbps

The writable parameters above will be used in any data model for configuring IP Capacity testing. Measured output will not be included in reports from the data models (instead they will be included as part of test protocols that support the status message communications).

5.2.2 Lost Status Backoff Timeout and Interaction with the “B” Search Algorithm

A safety control related to the “B” search algorithm is the reduction of sending rate when no feedback is received at the Sender for a time interval called the Lost Status Backoff Timeout.

Calculation of the timeout includes two existing parameters, UpperThresh and StatusFeedbackInterval, plus a variable w that increments on each timeout to further reduce the rate when successive feedback messages are lost.

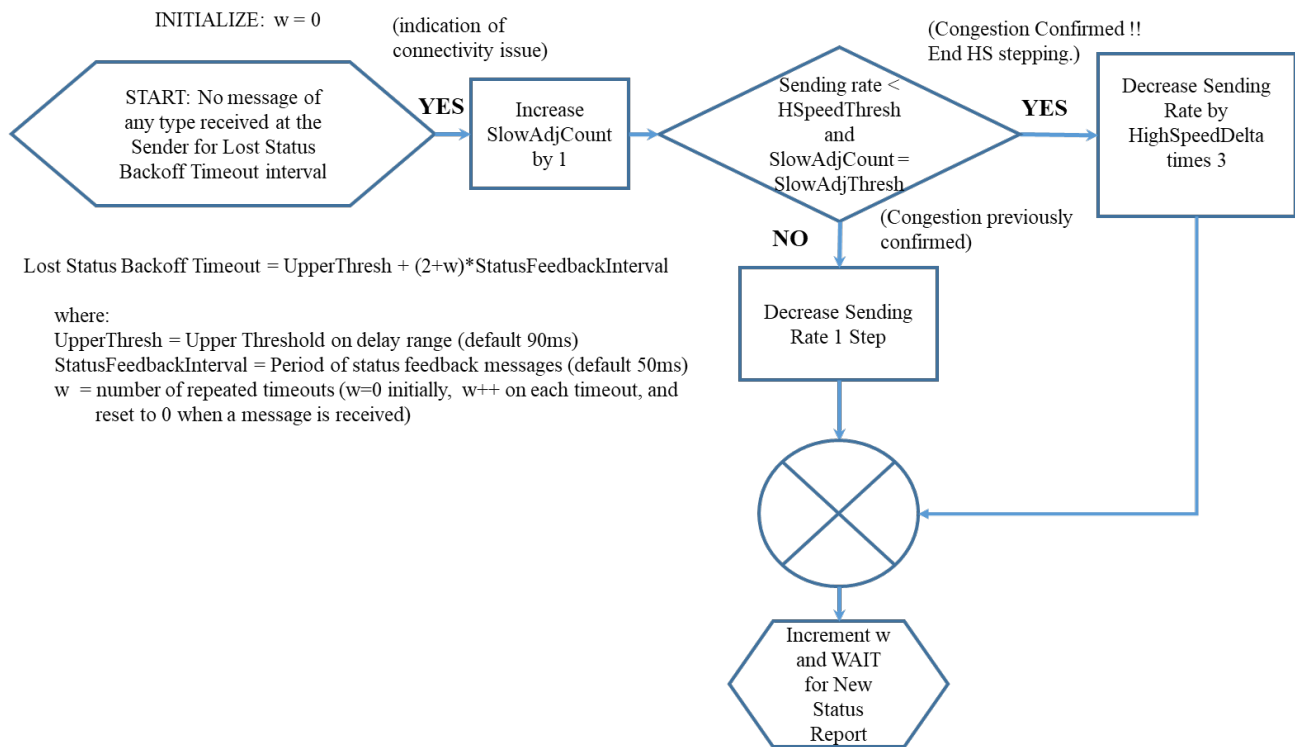


Figure 5: Flowchart for Lost Status Backoff Timeout Interaction with the “B” Search Algorithm

[R-22] If the Type “B” Search Algorithm is used, then the Test Controller MUST update appropriate algorithm parameters and reduce the sending rate as directed when the Sender indicates that the Lost Status Backoff Timeout has expired. The Lost Status Backoff Timeout SHALL be calculated as specified in Figure 5. The rate reduction SHALL be as determined by the flowchart in Figure 5. The RECOMMENDED initial value for variable w is 0.

6 Results

This section specifies details for the Results of measurements using the Maximum IP-Layer Capacity metric.

The primary output parameter is a Summary Statistic (Maximum of individual IP-layer Capacities), calculated at the conclusion of the test.

There are iterations of results that are reported on a per sub-interval basis, and the results associated with the Maximum Capacity measurement, all reported together at the end of the test.

Table 4: Data Model Definition for Results

Name	Description	Unit	Type	Write	Notes
BeginningOfMeasurement	t , the start of a measurement interval, in UTC, which MUST be specified to TimestampResolution precision (Table 2) For example: 2008-04-09T15:01:05.123456Z		dateTime	-	
EndOfMeasurement	$t + \Delta t$, the end of a measurement interval, in UTC, which MUST be specified to TimestampResolution precision (Table 2) For example: 2008-04-09T15:01:05.123456Z		dateTime	-	
TmaxUsed	Configured value of Tmax used in the test	ms	unsignedInt; $50 \leq \text{ms} \leq 3000$	-	
TmaxRTTUsed	Configured value of TmaxRTT used in the test	ms	unsignedInt; $50 \leq \text{ms} \leq 3000$	-	
TestInterval	Measured Duration of the test (either downlink or uplink). This value is expected to equal $dt * m$	seconds	(output)	-	
MaximumIP-LayerCapacity	Results of measurements using the Maximum IP-Layer Capacity metric, see Equation 1	Mbps	decimal64 number with fraction digits = 2, as specified in [Section 9.3 of RFC7950].	-	
TimeOfMaximumIP-LayerCapacity	End Time of the dt_n to dt_{n+1} sub-interval when the Maximum IP-Layer Capacity was measured, in UTC, which MUST be specified to TimestampResolution precision (Table 2). In case of Max Capacity occurring in multiple sub-intervals, report the Time of the earliest sub-interval.		dateTime	-	Note: millisecond resolution would be sufficient!
MaxETHCapacityNoFCS	Results of measurements using the Maximum IP-Layer Capacity metric, see Equation 1, and calculations to estimate the capacity at Layer 2 with Preamble and Inter-frame gap, but no ETH Frame Check Sequence	Mbps	decimal64 number with fraction digits = 2, as specified in [Section 9.3 of RFC7950].	-	
MaxETHCapacityWithFCS	Results of measurements using the Maximum IP-Layer Capacity metric, see Equation 1, and calculations to estimate the capacity at Layer 2 <u>with</u> ETH Frame Check Sequence	Mbps	decimal64 number with fraction digits = 2, as specified in	-	

			[Section 9.3 of RFC7950].		
MaxETHCapacityWithFCSVLAN	Results of measurements using the Maximum IP-Layer Capacity metric, see Equation 1, and calculations to estimate the capacity at Layer 2 <u>with</u> ETH Frame Check Sequence and 1 VLAN tag	Mbps	decimal64 number with fraction digits = 2, as specified in [Section 9.3 of RFC7950].	-	
LossRatioAtMaxCapacity	Ratio of lost to total packets sent during dt_n corresponding to the Max IP-Layer Capacity above), determined at the conclusion of the test.		decimal64 with fraction digits = 9 (see section 9.3 of [RFC6020]) with resolution of 0.0000000001	-	
RTTRangeAtMaxCapacity	<p>The Range of RTT during the dt_n corresponding to the Max IP-Layer Capacity above, determined at the conclusion of the test. The Range of RTT shall be calculated using the conditional distribution of all packets with a finite value of round-trip delay (undefined delays are excluded), a single value as follows:</p> <ul style="list-style-type: none"> See section 4.1 of [RFC3393] for details on the conditional distribution to exclude undefined values of delay, and Section 5 of [RFC6703] for background on this analysis choice. The time value of the result is expressed in units of seconds, as a positive value 	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950]) with resolution of 0.0000000001 seconds (1.0 ns).	-	
PDVRangeAtMaxCapacity	The Range of PDV during the dt_n corresponding to the Max IP-Layer Capacity above, determined at the conclusion of the test. The Range of PDV shall be calculated using the conditional distribution of all packets with a finite value of one-way delay (undefined delays are excluded, as described for other delay measurements).	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950]) with resolution of 0.0000000001 seconds (1.0 ns).	-	Note: PDV is the metric defined in Y.1540 and RFC5481, not inter-packet delay variation. Y.1540 and RFC5481 prefer a calculation of a high (99.9) percentile, rather than Max, and the percentile should be reported where possible.
MinOnewayDelayAtMaxCapacity	The Minimum One-way Delay during the dt_n corresponding to the Max IP-Layer Capacity above, The Minimum One-way Delay is determined at the conclusion of the test. The Minimum One-way Delay shall be calculated using the conditional distribution of all packets with a finite value of one-way	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950]) with resolution of 0.0000000001	-	Note: The Minimum One-way Delay is needed to calculate the Y.1540 and RFC5481 PDV, and

	delay (undefined delays are excluded, as described for other delay measurements).		seconds (1.0 ns).		should be reported separately.
ReorderedRatioAtMaxCapacity	Ratio of Reordered to total packets sent during dt_n corresponding to the Max IP-Layer Capacity above), determined at the conclusion of the test.		decimal64 with fraction digits = 9 (see section 9.3 of [RFC6020]) with resolution of 0.0000000001	-	
ReplicatedRatioAtMaxCapacity	Ratio of Replicated to total packets sent during dt_n corresponding to the Max IP-Layer Capacity above), determined at the conclusion of the test.		decimal64 with fraction digits = 9 (see section 9.3 of [RFC6020]) with resolution of 0.0000000001		
IP-LayerCapacitySummary	Results of measurements using the IP-Layer Capacity metric over the complete TestInterval, $dt * m$, see Equation 1	Mbps	decimal64 number with fraction digits = 2, as specified in [Section 9.3 of RFC7950].	-	
LossRatioSummary	Ratio of lost to total packets sent during the complete TestInterval, $dt * m$, determined at the conclusion of the test.		decimal64 with fraction digits = 9 (see section 9.3 of [RFC6020]) with resolution of 0.0000000001	-	
RTTRangeSummary	<p>The Range of RTT during the complete TestInterval, $dt * m$, determined at the conclusion of the test. The Range of RTT shall be calculated using the conditional distribution of all packets with a finite value of round-trip delay (undefined delays are excluded), a single value as follows:</p> <ul style="list-style-type: none"> See section 4.1 of [RFC3393] for details on the conditional distribution to exclude undefined values of delay, and Section 5 of [RFC6703] for background on this analysis choice. The time value of the result is expressed in units of seconds, as a positive value 	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950]) with resolution of 0.000000001 seconds (1.0 ns).	-	
PDVRangeSummary	The Range of PDV during the complete TestInterval, $dt * m$, determined at the conclusion of the test. The Range of PDV shall be calculated using the conditional distribution of all packets with a finite value of one-way delay (undefined delays are excluded, as described for other delay measurements).	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950]) with resolution of 0.000000001 seconds (1.0 ns).	-	Note: PDV is the metric defined in Y.1540 and RFC5481, not inter-packet delay variation. Y.1540 and RFC5481 prefer a calculation of a high (99.9) percentile,

					rather than Max, and the percentile should be reported where possible.
MinOnewayDelaySummary	The Minimum One-way Delay during the complete TestInterval, $dt * m$. The Minimum One-way Delay is determined at the conclusion of the test. The Minimum One-way Delay shall be calculated using the conditional distribution of all packets with a finite value of one-way delay (undefined delays are excluded, as described for other delay measurements).	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950]) with resolution of 0.000000001 seconds (1.0 ns).	-	Note: The Minimum One-way Delay is needed to calculate the Y.1540 and RFC5481 PDV, and should be reported separately.
MinRTTSummary	The Minimum RTT during the complete TestInterval, $dt * m$. The Minimum RTT is determined at the conclusion of the test. The Minimum RTT shall be calculated using the conditional distribution of all packets with a finite value of RTT (undefined delays are excluded, as described for other delay measurements).	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950]) with resolution of 0.000000001 seconds (1.0 ns).	-	
ReorderedRatioSummary	Ratio of Reordered to total packets sent during the complete TestInterval, $dt * m$, determined at the conclusion of the test.		decimal64 with fraction digits = 9 (see section 9.3 of [RFC6020]) with resolution of 0.000000001	-	
ReplicatedRatioSummary	Ratio of Replicated to total packets sent during the complete TestInterval, $dt * m$, determined at the conclusion of the test.		decimal64 with fraction digits = 9 (see section 9.3 of [RFC6020]) with resolution of 0.000000001		
TimestampResolutionUsed	This is a function of the measurement protocol, and it is usually determined once the protocol chosen. Value specified in microseconds.	microsec	unsignedInt	-	
ModalResult.{i}.	<p>Modal test results. Only returned when bimodal test mode is enabled (NumberFirstModeTestSubIntervals >=1). If returned, it MUST contain 1 or more entries, with instance number 1 corresponding to the second mode and instance number 2 corresponding to the third mode.</p> <p>Results for the Maximum in each mode/instance are calculated based on IncrementalResult.{i} data within the boundary of its corresponding mode.</p> <p>This table's Instance Numbers MUST be 1, 2, 3... (assigned sequentially without gaps).</p>				
IncrementalResult.{i}.	(Incremental Results for all sub-intervals are described below)				

IP-LayerCapacitySubInterval	Results of measurements using the IP-Layer Capacity metric for a single interval dt_n to dt_{n+1} , see Equation 1, where the Capacity is the number of bits received in the subinterval divided by the duration, dt .	Mbps	decimal64 number with fraction digits = 2, as specified in [Section 9.3 of RFC7950].	-	
TimeOfIP-LayerCapacitySubInterval	End Time of the dt_n to dt_{n+1} sub-interval when each of the m IP-Layer Capacity was measured, in UTC, which MUST be specified to TimestampResolution precision (Table 2)		dateTime	-	Note: millisecond resolution would be sufficient!
LossRatioSubInterval	Ratio of lost to total packets sent during dt_n to dt_{n+1} corresponding to each IP-LayerCapacitySubInterval above).		decimal64 with fraction digits = 9 (see section 9.3 of [RFC6020]) with resolution of 0.0000000001	-	
RTTRangeSubInterval	<p>The Range of RTT during dt_n corresponding to packets sent during dt_n to dt_{n+1} corresponding to each IP-LayerCapacitySubInterval above). The Range of RTT shall be calculated using the conditional distribution of all packets with a finite value of round-trip delay (undefined delays are excluded), a single value as follows:</p> <ul style="list-style-type: none"> See section 4.1 of [RFC3393] for details on the conditional distribution to exclude undefined values of delay, and Section 5 of [RFC6703] for background on this analysis choice. The time value of the result is expressed in units of seconds, as a positive value 	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950]) with resolution of 0.000000001 seconds (1.0 ns).	-	
PDVRangeSubInterval	The Range of PDV during dt_n corresponding to packets sent during dt_n to dt_{n+1} corresponding to each IP-LayerCapacitySubInterval above. The Range of PDV shall be calculated using the conditional distribution of all packets with a finite value of one-way delay (undefined delays are excluded, as described for other delay measurements).	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950]) with resolution of 0.000000001 seconds (1.0 ns).	-	Note: PDV is the metric defined in Y.1540 and RFC5481, not inter-packet delay variation. Y.1540 and RFC5481 prefer a calculation of a high (99.9) percentile, rather than Max, and the percentile should be reported where possible.
MinOnewayDelaySubInterval	The Minimum One-way Delay during dt_n corresponding to packets sent during dt_n to dt_{n+1} corresponding to each IP-LayerCapacitySubInterval above. The Minimum One-way Delay is	seconds	decimal64 with fraction digits = 9 (see section 9.3 of [RFC7950])	-	Note: The Minimum One-way Delay is needed to calculate the

	determined at the conclusion of the test. The Minimum One-way Delay shall be calculated using the conditional distribution of all packets with a finite value of one-way delay (undefined delays are excluded, as described for other delay measurements).		with resolution of 0.000000001 seconds (1.0 ns).		Y.1540 and RFC5481 PDV, and should be reported separately.
ReorderedRatioSubInterval	Ratio of Reordered to total packets sent during dt_n to dt_{n+1} corresponding to each IP-LayerCapacitySubInterval above).		decimal64 with fraction digits = 9 (see section 9.3 of [RFC6020]) with resolution of 0.0000000001		
ReplicatedRatioSubInterval	Ratio of Replicated to total packets sent during dt_n to dt_{n+1} corresponding to each IP-LayerCapacitySubInterval above).		decimal64 with fraction digits = 9 (see section 9.3 of [RFC6020]) with resolution of 0.0000000001		

The measurement variable, i , NumberFirstModeTestSubIntervals ≥ 1 (enabling Bimodal reporting) requires that Maximum results are reported separately for two time intervals:

1. Sub-intervals from 1 to i
2. Sub-intervals from $i + 1$ to m

The ModalResult. $\{i\}$ is the section of the data model above where the Second Mode Maximum results (and any additional modes in the future) will be found.

[R-23] When the Client application operates at the CPE in Figure 3, and conducts measurements using an interface that views all traffic during a test (not just test traffic), then the metrics defined in Table 5 **MAY** help to determine when non-test/subscriber traffic was present during a test. These calculations are based on the interface counts (such as TotalBytesReceived) described in TR-143 [3].

Table 5: Additional Metrics for Diagnostic Measurement

Name	Description	Unit	Type	Write	Notes
IncrementalResult. $\{i\}$.					
InterfaceEthMbpsSubInterval	The number of bits observed on the Interface during an IP-Layer Capacity test for a single sub-interval dt_n to dt_{n+1} , divided by the duration, dt , and expressed in Mbps. Measurement direction follows the Role (Sender or Receiver).	Mbps	decimal64 number with fraction digits = 2, as specified in [Section 9.3 of RFC7950].	-	
“AtMaxCapacity” and ModalResult. $\{i\}$					
InterfaceEthMbpsAtMaxCapacity	The number of bits observed on the Interface	Mbps	decimal64 number	-	

	during an IP-Layer Capacity test for a single sub-interval dt_n to dt_{n+1} corresponding to the Max IP-Layer Capacity, divided by the duration, dt , and expressed in Mbps. Measurement direction follows the Role (Sender or Receiver).		with fraction digits = 2, as specified in [Section 9.3 of RFC7950].		
"Summary"					
InterfaceEthMbpsSummary	The number of bits observed on the Interface during an IP-Layer Capacity test for the entire TestInterval Δt , divided by the duration, Δt , and expressed in Mbps. Measurement direction follows the Role (Sender or Receiver).	Mbps	decimal64 number with fraction digits = 2, as specified in [Section 9.3 of RFC7950].	-	

Note that the CPE operating system and hardware implementation will likely determine the layer where these bit counts can be observed as byte counts.

Table 6 provides additional information on the version of software and control protocol running on the configured Endpoint.

[R-24] When the Client application operates at the CPE in Figure 3, then the information defined in Table 6 **MAY** be communicated with the results to confirm the versions used during a test.

Table 6: Supporting Information for Test Context

Name	Description	Unit	Type	Write	Notes
IPLayerCapSupportedSoftwareVersion	Indicates the installed version of the test software. The software version string will be implementation-dependent, and SHOULD identify both the implementation and the version (e.g., UDPST-7.2.1).		string;	-	Currently at version "UDPST-7.3.0" for udpst. This would be a single string for client.
IPLayerCapSupportedControlProtocolVersion	Indicates the control protocol version supported by the test software.		string;	-	Currently at version 8 for udpst. This would be a single string for client.

Annex A: Sending Rates Table

This Annex provides information on the Sending Rates Table.

A.1 Example Table of Sending Rates and Related Information

When describing the system that conforms to this specification, the text in section 5.2.1 describes a table of sending rates. This is a pre-defined 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 sending rates, inclusive. The rows in the table have an index number for simple reference to the flow-chart calculations of the number of rows to move in steps.

The example Table A.1 below begins with the configured Initial Rate, and contains a representation of one table that has been used (with reduced detail). Note that this table results from computation of simple bit rates (in Mbps at Layer 3 or IP Layer) from packet counts, as this was considered efficient.

Table A.1: Example Table of Sending Rates

Index No.	Other Components, such as sending interval, payload size(s), spacing parameters, back-to-back (burst) packet counts, etc.	Mbps (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
...		...

Annex B: Type C Sending Rate Search Algorithm

This Annex specifies the Type C Sending Rate Search Algorithm.

B.1 Sending Rate Search Algorithm Improvements

There are several ways to make the current "Type B" algorithm ramp-up faster:

- Increase the high-speed step from 10 to 20 or 30, or even as high as 50 (although these higher values have consequences in the "congestion-confirmed" path in the flow chart: 50 would result in a back-off of 150 table steps of 1 Mbps)
- decrease the Feedback Interval from 50ms to a lower value, say 20ms
- increase the starting rate from 500kbps (as OB-UDPST Release 7.4.0 allows)

However, the high-speed ramp-up equation can also be changed to a multiplicative increase, instead of additive with steeper slope. There are several advantages to this approach:

- Service rate offerings are often approximately multiplicative (x2): 5, 25, 100, 200, 500, 1000 Mbps...
- The default values in the search algorithm take 5 seconds to reach 1Gbps, leaving 5 seconds to determine the Max Capacity in a default 10 second test. Higher access rates would take longer to reach.
- When the ramp-up takes very little time to reach the Max Capacity, there is less chance that a transient loss burst or competing traffic burst will "knee-cap" the search at a low rate.
- Algorithm changes that decouple the multiplicative (ramp-up) path from the "congestion-confirmed" path have been found useful.

Another feature that improves the search for Maximum Capacity is the logic to allow Retries of Fast-mode, meaning that almost any level of unexpected packet loss/impairments during Fast ramp-up can be mitigated for measurement success. Thus, this new algorithm is called "Multiply and Retry", and designated Type C. The Type C algorithm is first available in the OB-UDPST Release 7.5.0 running code, following many months of testing.

Figure B-1 specifies the flowchart for the Type C Algorithm.

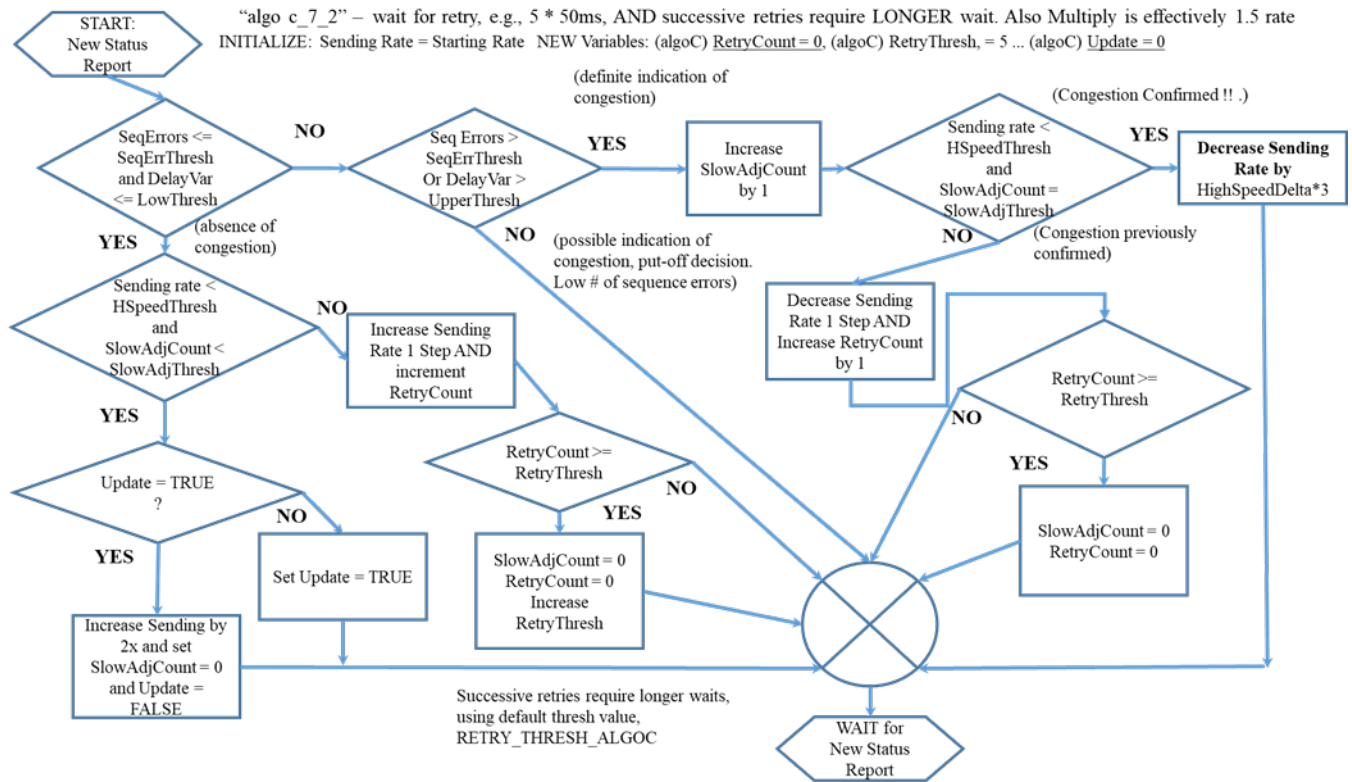


Figure B-1: Flowchart for Offered Load adjustment as part of the Type C Search Algorithm

Comparison of the Type B and Type C algorithms reveals that the Type C design used Type B as a starting point. The Type C flow chart in Figure B-1 above uses additional variable names and in one case, a configurable threshold that determines the flowchart decisions beyond those in the TR-471 Figure 4 Type B flow chart. Three main paths through the flowchart remain: 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. However, the sending rate increases by a multiplicative factor (1.5) when impairments are effectively absent. Further, the multiplicative sending rate increase and “congestion-confirmed” decrease paths and parameters are de-coupled, allowing greater control. The Type C algorithm also introduces automatic retry of the “fast” sending rate increase mode, with increasing intervals between retries as the test proceeds.

The new variables and thresholds used in the Figure B-1 Flowchart are explained below:

Table B.2: Table of Additional Flowchart variables, descriptions, ranges and default values

Category/ Variable Name	Description	Unit	Range	Default value
RetryCount	Number of status reports since previous “fast” mode activation. The counting takes place whether impairments are present or not.	Count of occurrences	NA	RetryThresh
RetryThresh	Threshold on RetryCount used to delay successive activations of “fast” sending	Count of occurrences	> 1	5

	rate increase mode. Use values >1 to increase the delay. Note that the algorithm doubles the threshold each time RetryThresh is exceeded.			
Update	The Update variable operates in “fast” sending rate increase mode to temper the effective increase to a fractional value (1.5). The sending rate increases alternate between multiply by 1 and multiply by 2 for an average of 1.5.		Boolean:[0:1]	0 (False) starting value, the value alternates with each pass through the “fast” sending rate increase mode.

Annex C: Multi-flow Operation

This Annex specifies the use of multiple IP/UDP flows to characterize Maximum IP-Layer Capacity, Latency, Delay variation, and other IP-Layer metrics.

Testing with Multiple flows introduces complexity as well as advantages, such as the possibility of reaching higher sending rates.

Possible benefits include:

- Allowing for load-balancing across server instances
- Permitting for test completion when one or more servers have failed or are too busy to test at the current time.
- Accommodating Active Queue Management that allocates resources on a per-flow basis, enabling the aggregate to achieve the link Capacity limit beyond the per-flow limit.
- Accommodating Link Aggregation Groups, which may be deployed to match the aggregate Capacity present elsewhere in the access architecture.
- Parallel systems may be used to produce the aggregate rate needed using multiple flows.
- Multiple flows may be helpful to saturate the path under test with a single pair of test hosts.
- Additional information could be derived for diagnostic purposes, or to validate the testing process. For example, comparing the data rates on each connection could be informative, where very different data rates might reveal abnormal operation.

C.1 New Input Parameters for Multi-flow Operation

Multi-flow testing is permitted between a single client and

- A single host in the role of server running one or more server instances on one network interface
- A single host running multiple server instances, with flows distributed among those instances. The instances may be distinguished by their control-port, network interface address, or both.
- Multiple geographically diverse hosts running one or more server instances. The instances may be distinguished by their control-port, network interface address, or both.

Each flow has its own control setup, activation, load stream and feedback channel, as well as calculation of measurements and load adjustment via the selected Type (B or C) flowchart.

There must be a report of the aggregate results over all flows, the same as with single flow operation.

One or more servers (distinguished by their control-port, network interface address, or both) intended as testing partners for the client appear in an input list called the ServerList.

The new variables and thresholds used during Multi-flow operation are explained below:

Table C.3: Table of Input Variables and Thresholds for Multi-flow Operation

Category/ Variable Name	Description	Unit	Range	Default value
ServerList	A list of servers intended as testing partners for the client. The client works through the server list attempting individual test flows in a round-robin sequence. One flow is attempted to each server and all flows are required for testing, unless the	Server designations	1 or more	NA

	optional FlowCount and MaximumFlows parameters are specified. The list contains each server's Fully Qualified Domain Name or IP address and listening port if applicable.			
FlowCount	Optional. The number of required flows.	Count of flows	1 or more	NA
MaximumFlows	Optional. The number of flows to attempt.	Count of flows	MaximumFlows must be greater than the size of ServerList or FlowCount.	NA

Table C.2: Table of Output Variables and Thresholds for Multi-flow Operation

Category/ Variable Name	Description	Unit	Range	Default value
ActiveFlows	Number of flows set-up and used for a test.	Count of flows	0 or more	NA

End of Broadband Forum Technical Report TR-471