

Network Working Group
Internet-Draft
Intended status: Informational
Expires: March 10, 2022

A. Morton
AT&T Labs
September 6, 2021

Dream-Pipe or Pipe-Dream: What Do Users Want (and how can we assure it)?
[draft-morton-ippm-pipe-dream-01](#)

Abstract

This memo addresses the problem of defining relevant properties and metrics with the goal of improving Internet access for all users. Where the fundamental metrics are well-defined, a framework to standardize new metrics exists and been used with success. Users consider reliability to be important, as well as latency and capacity; it really depends who you ask and their current experiences.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 10, 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4.e](#) of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. What's happening with users today?	2
1.2. User expectations and how they can change	3
1.3. How fast can you go?	4
2. Dream-Pipe/Pipe-Dream	5
3. Metrics to Assess Fundamental Properties	7
4. Interpreting the Measurements	10
5. Work together, as always!	11
6. Displaying or Reporting Results	11
7. Summary	12
8. Security Considerations	13
9. IANA Considerations	13
10. Acknowledgments	13
11. References	13
11.1. References	13
11.2. More References	15
Author's Address	17

1. Introduction

This memo addresses the problem of defining relevant properties and metrics with the goal of improving Internet access for all users. Much has already been done, and it's important to have a common foundation as we move forward. There is certainly more to understand about the problem and the approaches to a solution.

1.1. What's happening with users today?

Part of the motivation for examining metrics in greater detail at this time is the belief that "Internet speed" is no longer the only service dimension that matters to users. A small sample of recent surveys follows.

In a 2021 UK [[EY-Study](#)] EY study (Summary in Advanced Television), Decoding the digital home, Chapter 2, a survey of 2500 subscribers found that "Fifty-eight per cent of UK households believe broadband reliability is more important than speed". Also, "the appetite for a consistent connection aligns with perceptions that broadband reliability declined during the pandemic - 29 per cent across all households, rising to 46 per cent in households with children aged up to 11 years." If the family is on-line more-often, they will notice more outages.

Reliability is surely important, but other factors come into play: "...nearly half (47 per cent) don't think upgrading to higher-speed packages is worth the cost. Meanwhile, 29 per cent say they don't understand what broadband speed means in practice." [EY-Study] Price trade-off and knowledge about the communication details play a role when discussing speed. Reliability issues are experienced by everyone.

All the findings above are predicated on the current "speeds" being delivered. For example, "Fifty-two percent of rural users are frustrated that the fastest speeds are not available in their area, and only 52 per cent believe they are getting value for money from their current broadband package." [EY-Study] The proportion of Urban users was 10-15% higher when asked about value. Broadband speed available often defines the development gap in world-wide surveys, and users in under-served areas are likely glad to see increases.[N-Africa]

Another survey described in [DontKnow] found that 36% of Americans *don't* know their Internet speed. A higher percentage of females did not know (47%) than males (25%). Older Americans (55+) and those in lower income households are other demographics where service speed was an above-overall-average gap in their Internet knowledge. Those who did not know their speed tended to be satisfied with the speed they receive, about 20% of these indicated dissatisfaction.

Latency wasn't mentioned in [EY-Study], but "While the survey indicates that most households ultimately want 'the basics' to work well, those that do consider additional features as part of a broadband bundle favour privacy and security (48 per cent), reflecting wider anxieties and concerns relating to data protection experienced during the pandemic." Many users want their service provider to provide all the services.

1.2. User expectations and how they can change

User expectations are greatly influenced by their current experiences, and by what is technologically feasible at the time. Users have a view of the levels of availability, quality, and utility that constitute the overall experience in their situation of use (e.g., stationary in their home).

When we insert a communication network as an integral part of an activity (a task or form of entertainment), then the expectation doesn't change or make allowances without a trade-off between the new features and the new situation For example, take the home but add the ability to travel: a stretch limousine fills some of the need but

with reduced capabilities (the refrigerator is smaller, no beds or lavatory; there is less of everything). But we have a TV!

Let's assume the TV uses analog transmission for this discussion. The TV is smaller and experiences reception outages, so we are suddenly aware of radio characteristics like fading and multipath (especially with earlier analog transmissions, there is no buffering at all). The limo is really nicely appointed, but it's not all we dreamed-of (or possibly even expected), especially the TV reception because we wanted to watch a play-off game on this trip!

Where did they go wrong providing the TV in the Limo? They made the communication channel a much more obvious part of the viewing experience. And the first-time users didn't expect it. They inserted unexpected impairments in the communication channel by adding mobility. The TV designer might not have been aware of the moving use case; their "portable TV" means you can pick-up the TV easily, not watch TV at high speed, so needed features were not provided (a tracking antenna, to start with...).

Since the goal of the workshop is to improve Internet Access for *all* users, then we have set a difficult task. Some users might want a dedicated pipe to communicate in the ways they choose or access content that they have identified. Other users are willing to share a pool of communication resources to communicate only when and where they want, for the potential benefits of being able to communicate more widely, spontaneously, and possibly for less cost than the dedicated pipe. It seems that we should focus on the subset of performance attributes that will benefit as many users as possible, and over-which our constituent organizations have some control.

1.3. How fast can you go?

To some, the maximum bit rate remains the primary goal.

The Internet Speed record previously held by the University College of London (178 Tbps) was topped in July 2021 by a team of researchers from the National Institute of Information and Communications Technology (NICT) in Japan. The new world record for Internet speed, is 319 Tbps, using 4-core fiber [[WorldRecord](#)].

Higher link rates and subscriber rates may not be everything to users, but there can be a cross-over dependency to latency performance. Packet serialization time is reduced at higher link speeds, directly proportional to the increased rate. Bursts of large packets arriving for one stream affect the buffer time for packets in other streams that arrive behind them in a single FIFO queue, but

again the problem is relieved by using higher link rates. For example, early VoIP used very low bit rate codecs: 8kbps was common, and so were 10 Mbps LANs. But mixing bursty and periodic traffic meant unreliable delivery and delay variation for the latter. Packet marking, prioritization of multiple queues and queue management methods helped. More rate headroom and adaptation for the network impairments was also welcome.

Instead of providing more capacity for a single user, today's Gigabit services support more users efficiently on the same access service. So higher rates can improve the other important dimensions of performance.

2. Dream-Pipe/Pipe-Dream

Perhaps the ideal view of point-to-point communications is a pipe that illustrates many fundamental communication properties. A dream-pipe, so to speak.

- o The pipe is always ready to support communications when the user desires.
- o Material that enters the pipe leaves the pipe in a sufficiently unadulterated form.
- o The pipe has sufficient and dedicated capacity to deliver information at the same rate at which it entered (once the diameter of the pipe has been chosen).

The apparent rigidity of the pipe model helps identify additional properties that are needed:

- o The capacity requirement may change/increase over time, even from session to session, so a user might anticipate a growing need for communications or anticipate use by more than one user by choosing a larger capacity pipe than they currently need.
- o When we say, "delivered in sufficiently unadulterated form", we could mean:
 1. a perfect reproduction (the communication of all messages with the same timing and order in which they started), or
 2. a reliable byte stream delivery, or
 3. communication of most messages with sufficient timeliness to reconstruct the original messages, or

4. even a perceptual interpretation of a portion of the decoded messages delivered in a sufficiently timely way (for human-to-human voice communication: the users understood what was said, the conversation was sufficiently interactive as though they were in the same room (imperceptible latency), and each user can identify the remote speaker's voice. But the overall tolerance for imperfections in the pipe depends on the particular use case and many design choices.
- o The pipe itself must be ready for use, and the systems at the source and destination ends of the pipe must be equally ready and operational; all the systems required between users are responsible for success.
 - o The pipe may be invisible! Radio access has its own advantages and challenges.

So, a short list of network properties that contribute to good user experience are:

1. Available always when needed
2. Sufficient Capacity when needed
3. No apparent loss
4. No apparent latency (which implies both low and consistent latency).

Networking and geographic reality tells us that we are unlikely to see all properties at once, for all time, AOE (anywhere on Earth); that's the extreme Pipe-Dream.

But attaining a good user experience level in different communication activities/scenarios likely implies different demands among the four properties above, and places different relative importance on each of these properties.

Author's Note: I'm positing a rigid pipe as the object of an idealistic idea, or "dream pipe". I hope there won't be any confusion about the play on words with "pipe dream; the dream pipe is a pipe dream, especially when the dream includes zero latency between any two points. I'm not talking about a hose or flexible pipe, either.

3. Metrics to Assess Fundamental Properties

The IETF has been working the problem of standardizing metrics for the Internet and the communication streams it transports for well over 20 years. Many other organizations have been successfully working in this area as well, and hopefully they will identify their literature and key results for review.

The IETF's efforts to define IP-Layer and Transport-Layer performance metrics and methods have largely been carried-out in the IPPM working group (IP-Performance Metrics, and later, called IP-Performance Measurements). Beginning as a somewhat joint effort with the performance-focused Benchmarking Methodology Working Group (BMWG), IPPM was chartered individually in 1997. IPPM has extensive literature relevant to Internet measurement. The IPPM working group has a strong foundation in its Framework [RFC2330] that has been updated over time, with [RFC7312] and [RFC8468]. What we can *standardize and measure*, we have as a basis to evaluate and determine whether we have made it better (or not).

The problem that initiated the IPPM work turned-out to be the most difficult to solve (Bulk Transfer Capacity, BTC [RFC3148]), and has taken the longest. Meanwhile, the standards for fundamental metrics other than BTC turned-out to be sufficiently challenging.

Here is a list of fundamental packet transfer metrics, specified in RFCs:

- o Connectivity [RFC2678]
- o One-Way Loss [RFC7679], STD 81
- o One-Way Delay [RFC7680], STD 82
- o Round-Trip Delay [RFC2681]
- o Round-Trip Loss [RFC6673]
- o Reordering [RFC4737]
- o Duplication [RFC5560]

The metrics and methods above were specified with considerable flexibility, so that they could be applied in a range of specific circumstances.

One of the most flexible metrics is IP Packet Delay Variation, [RFC3393] which is a "derived metric", in that it requires One-Way

Delay measurements for assessment. The powerful feature of [RFC3393] is the selection function, which permits comparing the delays of any pair of packets in the stream. Fortunately, the performance community predominantly uses one of two forms of delay variation, the inter-packet delay variation and the packet delay variation forms. These forms are defined and compared in efficacy for measurement uses in [RFC5481] along with many other considerations and measurement forms/processing.

It is possible to create new derived metrics at the IP-layer, and to measure similar quantities (loss, delay, reordering) at other layers [ForAll] [RFC6390] [RFC6076].

Another example of a derived metric uses loss instead of delay. This metric that does not have a direct parallel in the IETF literature is the Stream Block metric found in ITU-T Recommendation Y.1540 [Y.1540] (virtually all the IP-layer metrics are found in one standard). This metric assigns consecutive packets into multi-packet blocks, and assesses the number of lost packets in a block as a surrogate for a higher-layer process's ability to maintain good communication. For example, a Forward Error Correction process might be able to replace 2 lost packets in any order, but not 3. There is a parallel to retransmission rate limits when attempting to maintain a continued loss-free ratio with buffering to allow for the retransmission time.

Network and Bulk Transport Capacity have been chartered and progressed over twenty years. The performance community has seen development of Informative definitions in [RFC3148] for Framework for Bulk Transport Capacity (BTC), [RFC5136] for Network Capacity, Maximum IP-layer Capacity (in RFC9097-to-be), and the Experimental metric definitions and methods in [RFC8337], Model-Based Metrics for BTC.

One quantity that could be measured without too much controversy is the Maximum IP-Layer Capacity (judging by the adoption in standards bodies, RFC9097-to-be) [I-D.ietf-ippm-capacity-metric-method]. This is the basis for many service specifications (and there is a technology or configuration-limited "ground truth" for the measurement), and can be tested simply with minimal reliance on end systems. The method deploys a feedback channel from the receiver to control the sender's transmission rate in near-real-time, and search for the maximum.

The "invisible" radio dream pipe presents a challenge in terms of results variability for all metrics and adds at least one critical input parameter: location. It may be that the variability with location and time are key metrics to help users understand radio

coverage (in addition to signal strength portrayed by the "number of "bars").

A network property that is very high on the list in [Section 2](#) is Availability. There is treatment of this important property as a metric in IETF (Connectivity [[RFC2678](#)]) and somewhat different details in an alternate definition (ITU-T Point-to-Point IP Service Availability [[Y.1540](#)]), but not much deployment for such an important pre-requisite to the rest of the metrics. Both Connectivity and Availability rely on packet loss measurements, in fact they can be considered *derived metrics*, adding time constraints and/or loss ratio thresholds to the fundamental loss measurements on a packet stream.

Sometimes the end systems decide whether the path is available or not. In one on-line gaming system, reception of ICMP Destination Unreachable caused complete and immediate termination of the session with loss of all progress and resources accumulated. Customer service centers sometimes experienced call-in overload from users seeking to restore their player environment. But the Destination Unreachable condition was temporary and resolved by routing updates in a few seconds. The ultimate fix was to delay the session's reaction to Destination Unreachable for a few seconds, and recovery was automatic. So, when end systems play a role in the definition of connectivity or availability, they must also be cognizant that all automatic failure detection and restoration requires some amount of time. Since failures are inevitable, the dream-pipe heals itself, too (and doesn't confuse end-systems or users with error messages that "sound final").

Most measurement systems begin their process with a Source-Destination packet exchange prior to actual measurements. If this pre-measurement exchange fails, then the test is not conducted (and re-tried later). But the most useful information to assess continuous connectivity/Availability is the record of test set-up success or failure over time. Measurement systems that make this info readily available do a more complete job of network characterization than others.

We cannot leave the topic of metrics without mentioning the equally important topic of measurement streams for Active Metrics and Measurements [[RFC7799](#)]. When attempting to measure characteristics of VoIP streams, the IETF agreed on ways to produce periodic streams in an acceptable way [[RFC3432](#)]. Many measured results completely depend on the stream characteristics, and inter-packet delay variation is a great example. If the stream contains packet bursts, are the bursts preserved in transit? We would ask later-on whether preserving burst spacing matters to the communication quality or to

the user's experience. The answer is likely a matter of degree, and dependent on the communications activity itself.

When we consider the topic of new test streams in the context of Gigabit and higher access speeds intended to support multiple users, we might consider the notion of a "standard single user's stream set". Then we might measure how many simultaneous standard users can be supported with sufficient network performance: an indicator of each user's experience in the "dream-pipe". Of course, that details of a standard user's traffic would change over time, so we couldn't argue over the current year's definition for a year... The facility to register the new standard user test streams would be a key part of such a solution.

4. Interpreting the Measurements

If we had measured only the fundamental metrics, we might ask, "we saw a small proportion of packet losses; did this matter to users? Did losses affect their satisfaction during their activity in any way? Did any users experience an outage at this loss level?"

The likely role for derived metrics is to improve results interpretation by measuring a new quantity, possibly in the context of a newly-defined packet stream, thereby making the process of results-interpretation easier to perform.

For example, Delay Variation metrics had many possible formulations, but two main forms emerged. [RFC 5481 \[RFC5481\]](#) compared the IPDV and PDV forms with tasks that network and application designers were facing at the time. The RFC describes measurement considerations and results interpretation from a purely objective point of view. The most useful result interpretation was to show how PDV (a characterization of the delay distribution from minimum to a high percentile) could be used to determine the size of the de-jitter buffer needed for the tested path.

For other fundamental IP-Layer metrics, there is some amount of discussion of best practices and interpretation in each of the IETF IPPM Metric RFCs.

Many researchers (working in ITU-T Study Group 12 and elsewhere) take the information that can be derived from packet-layer measurements, plus higher layers when available, and produce objective estimates of user satisfaction by modeling user Mean Opinion Score (MOS) variation over a range of conditions. The process determines the user MOS through formal subjective testing in laboratories (or more recently prompted by COVID-19 conditions, in crowd-sourced scenarios). The corresponding objective models of user satisfaction are often

determined through competition among several candidates, where the goal is to seek the most accurate model possible at the time. The modeling efforts often produce new derived metrics that facilitate automated interpretation. The main drawback is that the process described above takes significant time when conducted in the context of an industry standards body. So, user activities that are not particularly demanding of network performance do not receive much attention from researchers doing modeling; their performance is assumed to come-along for the ride (but in a system with multiple queues or other categorizations, each activities' requirements need to be quantified). Nevertheless, a process to take-in network measurements and produce a measure of user satisfaction is well-understood and used. The output of these objective models can rightly be called Quality of Experience (QoE), because a set of users' opinions is an inherent part of the result (without users, you don't have QoE! That's how QoE is defined and differs from QoS and network performance).

5. Work together, as always!

"What are the best ways to communicate these properties to service providers and network operators?" Work together with service providers and network operators. Everyone has a stake in our future.

There are quite a few networking professionals whose day-job is network operation, and who are participating now.

6. Displaying or Reporting Results

Whether we present a single figure of merit, or a set of relevant measurements on a dashboard summary, each number requires a frame of reference. This is true for everyone, not just everyday users.

One example of a solid reference for results comes from the fundamental benchmarking specification, [RFC2544]. The measured Throughput (as defined in [RFC2544]) must be compared to the maximum theoretical frame rate on the layer-2 technology (accounting for frame size, inter-frame gap, preamble, etc.). Tests with small frame size may not achieve the maximum frame rate due to header processing rate limitations, and tabulating the maximum with the results makes this fact very clear. Other reference levels can be made available, such as the capacity required to support 4k video (~25Mbps), etc.

Some people will simply want to know whether the measurement result is good, bad, or somewhere in-between. We can follow common practice here to use colors (green, red, yellow in-between), or present the number on a gauge with suitable color cues. But we need to know the use case or the service specification accurately to do this.

In fact, a portion of user testing is prompted by subscribing to a new service provider, a new level of service (higher speed?), or a perception that poor-performance and trouble-shooting may be necessary; even an apparent outage may prompt test attempts using alternate devices and networks.

This last testing scenario is the most interesting: how can we help users when they encounter a problem? It's usually most important to isolate the problem in the complex network, and when user to network host results are failing, can the next step remove some of the components and check them in isolation (while keeping in mind that the acceptance level for a sub-network is a part of the end-to-end budget for performance!). Is it impossible to reach a particular far-end host? Does the access network appear to be unavailable, or is the problem related to interference on the WiFi radio network? With test hosts placed at strategic points in the path, it may be possible to segment the problem a user is experiencing.

7. Summary

When all components that comprise a user's activity work well-together, then there are no surprises. Given that a large proportion of user expectations are met at some point in time, then the metrics and performance levels that characterize the network's contribution to overall satisfaction are what we want to describe and maintain. Users consider reliability to be important, as well as latency and capacity; it really depends who you ask and their current experiences. End-system designers have a role to play in the process by recognizing the realities of packet networks and compensating for them: the dream-pipe absolutes are still a pipe-dream today.

There are many fundamental metrics already-defined. But we might find that we need new metrics that make interpreting the results easier! The notion of **derived metrics** has been applied successfully. Test streams with a known bias toward a particular class of user streams can also be useful basis for performance measurement. Where the fundamental metrics are well-defined, a framework to standardize new metrics and active test streams exists and been used with success. Metrics can be defined that immediately improve our understanding of the performance presented to users, but to understand user satisfaction requires that user opinions are part of the development process.

All users, both knowledgeable and newcomers, need a frame of reference to understand what numerical measurements are telling them. The clues from the expected measurement range, the results from the recent past, or the theoretical maximum value all have their place. If users are willing, measurements should help them isolate their

current issue to one or more networks and/or components in the user-to-X path.

If we break the problem down by specific communications activities and look for specific metrics for each one, it could take a long time to complete. Perhaps a categorization of the performance metrics, numeric criteria, and reliability of a "pseudo-dream-pipe" for a set of communication activities that have similar needs is a way to move ahead.

8. Security Considerations

Active metrics and measurements have a long history of security considerations. The security considerations that apply to any active measurement of live paths are relevant here. See [RFC4656] and [RFC5357].

When considering privacy of those involved in measurement or those whose traffic is measured, the sensitive information available to potential observers is greatly reduced when using active techniques which are within this scope of work. Passive observations of user traffic for measurement purposes raise many privacy issues. We refer the reader to the privacy considerations described in the Large Scale Measurement of Broadband Performance (LMAP) Framework [RFC7594], which covers active and passive techniques.

9. IANA Considerations

This memo makes no requests of IANA.

10. Acknowledgments

Thanks to all the folks who have worked on performance metric development, many of whom are the authors of the references. Many more have provided their insights along the way. It's rewarding to travel with others (even for a short time), and to meet new people on the journey.

11. References

11.1. References

[I-D.ietf-ippm-capacity-metric-method]
Morton, A., Geib, R., and L. Ciavattone, "Metrics and Methods for One-way IP Capacity", [draft-ietf-ippm-capacity-metric-method-12](#) (work in progress), June 2021.

- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, DOI 10.17487/RFC2330, May 1998, <<https://www.rfc-editor.org/info/rfc2330>>.
- [RFC2678] Mahdavi, J. and V. Paxson, "IPPM Metrics for Measuring Connectivity", RFC 2678, DOI 10.17487/RFC2678, September 1999, <<https://www.rfc-editor.org/info/rfc2678>>.
- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, DOI 10.17487/RFC2681, September 1999, <<https://www.rfc-editor.org/info/rfc2681>>.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, DOI 10.17487/RFC3393, November 2002, <<https://www.rfc-editor.org/info/rfc3393>>.
- [RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, DOI 10.17487/RFC3432, November 2002, <<https://www.rfc-editor.org/info/rfc3432>>.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, DOI 10.17487/RFC4656, September 2006, <<https://www.rfc-editor.org/info/rfc4656>>.
- [RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", RFC 4737, DOI 10.17487/RFC4737, November 2006, <<https://www.rfc-editor.org/info/rfc4737>>.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarez, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, DOI 10.17487/RFC5357, October 2008, <<https://www.rfc-editor.org/info/rfc5357>>.
- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, DOI 10.17487/RFC5481, March 2009, <<https://www.rfc-editor.org/info/rfc5481>>.
- [RFC5560] Uijterwaal, H., "A One-Way Packet Duplication Metric", RFC 5560, DOI 10.17487/RFC5560, May 2009, <<https://www.rfc-editor.org/info/rfc5560>>.

- [RFC6076] Malas, D. and A. Morton, "Basic Telephony SIP End-to-End Performance Metrics", RFC 6076, DOI 10.17487/RFC6076, January 2011, <<https://www.rfc-editor.org/info/rfc6076>>.
- [RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, DOI 10.17487/RFC6390, October 2011, <<https://www.rfc-editor.org/info/rfc6390>>.
- [RFC6673] Morton, A., "Round-Trip Packet Loss Metrics", RFC 6673, DOI 10.17487/RFC6673, August 2012, <<https://www.rfc-editor.org/info/rfc6673>>.
- [RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, RFC 7679, DOI 10.17487/RFC7679, January 2016, <<https://www.rfc-editor.org/info/rfc7679>>.
- [RFC7680] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Loss Metric for IP Performance Metrics (IPPM)", STD 82, RFC 7680, DOI 10.17487/RFC7680, January 2016, <<https://www.rfc-editor.org/info/rfc7680>>.

11.2. More References

[DontKnow]

<https://advanced-television.com/2021/06/25/survey-just-36-of-us-know-their-broadband-speed/>, "Survey: Just 36% of US know their broadband speed (title is incorrect)", June 2021, <<https://advanced-television.com/2021/06/25/survey-just-36-of-us-know-their-broadband-speed/>>.

[EY-Study]

<https://advanced-television.com/2021/06/25/study-uk-homes-favour-broadband-reliability-over-speed/>, "Study: UK homes favour broadband reliability over speed", June 2021, <<https://advanced-television.com/2021/06/25/study-uk-homes-favour-broadband-reliability-over-speed/>>.

[ForAll]

Morton, A., "'Performance Metrics for All," in IEEE Internet Computing, vol. 13, no. 4, pp. 82-86, doi: 10.1109/MIC.2009.87.", July 2009, <<https://ieeexplore.ieee.org/document/5167273>>.

[N-Africa]

<https://www.libyaherald.com/2021/06/08/internet-speeds-improved-in-north-africa-over-last-four-quarters/>,
"Internet speeds improved in North Africa over last four quarters", June 2021,
<<https://www.libyaherald.com/2021/06/08/internet-speeds-improved-in-north-africa-over-last-four-quarters/>>.

[RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, DOI 10.17487/RFC2544, March 1999, <<https://www.rfc-editor.org/info/rfc2544>>.

[RFC3148] Mathis, M. and M. Allman, "A Framework for Defining Empirical Bulk Transfer Capacity Metrics", RFC 3148, DOI 10.17487/RFC3148, July 2001, <<https://www.rfc-editor.org/info/rfc3148>>.

[RFC5136] Chimento, P. and J. Ishac, "Defining Network Capacity", RFC 5136, DOI 10.17487/RFC5136, February 2008, <<https://www.rfc-editor.org/info/rfc5136>>.

[RFC7312] Fabini, J. and A. Morton, "Advanced Stream and Sampling Framework for IP Performance Metrics (IPPM)", RFC 7312, DOI 10.17487/RFC7312, August 2014, <<https://www.rfc-editor.org/info/rfc7312>>.

[RFC7594] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A Framework for Large-Scale Measurement of Broadband Performance (LMAP)", RFC 7594, DOI 10.17487/RFC7594, September 2015, <<https://www.rfc-editor.org/info/rfc7594>>.

[RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.

[RFC8337] Mathis, M. and A. Morton, "Model-Based Metrics for Bulk Transport Capacity", RFC 8337, DOI 10.17487/RFC8337, March 2018, <<https://www.rfc-editor.org/info/rfc8337>>.

[RFC8468] Morton, A., Fabini, J., Elkins, N., Ackermann, M., and V. Hegde, "IPv4, IPv6, and IPv4-IPv6 Coexistence: Updates for the IP Performance Metrics (IPPM) Framework", RFC 8468, DOI 10.17487/RFC8468, November 2018, <<https://www.rfc-editor.org/info/rfc8468>>.

[WorldRecord]

https://www.fibre-systems.com/news/internet-speed-record-319tbs-reached?utm_source=Adestra&utm_medium=email&utm_content=Internet%20speed%20record%20of%20319Tb%2Fs%20reached&utm_campaign=FS%20July%202021%20Newsline&utm_term=Fibre%20Systems, "Internet speed record of 319Tb/s reached", July 2021, <https://www.fibre-systems.com/news/internet-speed-record-319tbs-reached?utm_source=Adestra&utm_medium=email&utm_content=Internet%20speed%20record%20of%20319Tb%2Fs%20reached&utm_campaign=FS%20July%202021%20Newsline&utm_term=Fibre%20Systems>.

[Y.1540] Y.1540, I. R., "Internet protocol data communication service - IP packet transfer and availability performance parameters", December 2019, <<https://www.itu.int/rec/T-REC-Y.1540-201912-I/en>>.

Author's Address

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown, NJ 07748
USA

Phone: +1 732 420 1571
Fax: +1 732 368 1192
Email: acm@research.att.com