

Workgroup: Network Working Group  
Internet-Draft: draft-iab-mnqeu-report-00  
Published: 10 November 2021  
Intended Status: Informational  
Expires: 14 May 2022  
Authors: W. Hardaker  
USC/ISI

## **IAB workshop report: Measuring Network Quality for End-Users**

### **Abstract**

The Measuring Network Quality for End-Users workshop was held virtually by the Internet Architecture Board (IAB) in September, 2021. This workshop summarizes the workshop, the topics discussed and some preliminary conclusions drawn at the end of the workshop.

### **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 14 May 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](#)
  - [1.1. Problem space](#)
- [2. Workshop Agenda](#)
- [3. Position Papers](#)
- [4. Discussions](#)
  - [4.1. Introduction and overviews](#)
  - [4.2. Metrics considerations](#)
  - [4.3. Cross-layer considerations](#)
  - [4.4. Synthesis](#)
- [5. Conclusions](#)
  - [5.1. General statements](#)
  - [5.2. Specific statements about detailed protocols/techniques](#)
  - [5.3. Problem statements and concerns](#)
  - [5.4. No-consensus reached statements](#)
- [6. Follow on work](#)
- [7. Security considerations](#)
- [8. References](#)
  - [8.1. Normative References](#)
  - [8.2. Informative References](#)
- [Appendix A. Participants List](#)
- [Appendix B. IAB Members at the Time of Approval](#)
- [Appendix C. Acknowledgements](#)
  - [C.1. Workshop Chairs](#)
  - [C.2. Program Committee](#)
- [Appendix D. Github Version of this document](#)
- [Author's Address](#)

## 1. Introduction

The Internet Architecture Board (IAB) holds occasional workshops designed to consider long-term issues and strategies for the Internet, and to suggest future directions for the Internet architecture. This long-term planning function of the IAB is complementary to the ongoing engineering efforts performed by working groups of the Internet Engineering Task Force (IETF).

The Measuring Network Quality for End-Users workshop was held virtually by the Internet Architecture Board (IAB) in September, 2021. This workshop summarizes the workshop, the topics discussed and some preliminary conclusions drawn at the end of the workshop.

### 1.1. Problem space

The Internet in 2021 is quite different from what it was 10 years ago. Today, it is a crucial part of everyone's daily life. People use the Internet for their social life, for their daily jobs, for routine shopping, and for keeping up with major events. An

increasing number of people can access a Gigabit connection, which would be hard to imagine a decade ago. And, thanks to improvements in security, people trust the Internet for both planning their finances and for everyday payments.

At the same time, some aspects of end-user experience have not improved as much. Many users have typical connection latency that remains at decade-old levels. Despite significant reliability improvements in data center environments, end users often see interruptions in service. Despite algorithmic advances in the field of control theory, one can often find that the queuing delay in the last-mile equipment exceeds the accumulated transit delay. Transport improvements, such as QUIC, Multipath TCP, and TCP Fast Open are still not fully supported in some networks. Likewise, various advances in the security and privacy of user data are not widely supported, such as encrypted DNS to the local resolver.

Some of the major factors behind this lack of progress is the popular perception that throughput is the often sole measure of the quality of Internet connectivity. With such narrow focus, the workshop aimed to discuss various questions:

- \*What is the latency under typical working conditions?

- \*How reliable is the connectivity across longer time periods?

- \*Does the network allow the use of a broad range of protocols?

- \*What services can be run by clients of the network?

- \*What kind of IPv4, NAT or IPv6 connectivity is offered, and are there firewalls?

- \*What security mechanisms are available for local services, such as DNS?

- \*To what degree are the privacy, confidentiality, integrity and authenticity of user communications guarded?

- \*Improving these aspects of network quality will likely depend on measurement and exposing metrics to all involved parties, including to end users in a meaningful way. Such measurements and exposure of the right metrics will allow service providers and network operators to focus on the aspects that impacts the users' experience most and at the same time empowers users to choose the Internet service that will give them the best experience.

- \*What are the fundamental properties of a network that contribute to good user experience?

- \*What metrics quantify these properties, and how to collect such metrics in a practical way?
- \*What are the best practices for interpreting those metrics, and incorporating those in a decision making process?
- \*What are the best ways to communicate these properties to service providers and network operators?
- \*How can these metrics be displayed to users in a meaningful way?

## **2. Workshop Agenda**

The Measuring Network Quality for End-Users for divided into the following main topic areas:

- \*Introduction overviews and a keynote by Vint Cerf
- \*Metrics considerations
- \*Cross-layer considerations
- \*Synthesis
- \*Group conclusions

## **3. Position Papers**

- \*Stuart Cheshire. "The Internet is a Shared Network"
- \*Jana Iyengar. "The Internet Exists In Its Use"
- \*Yaakov (J) Stein. "The Futility of QoS"
- \*Keynote by Vint Cerf
- \*Pedro Casas. "10 Years of Internet-QoE Measurements. Video, Cloud, Conferencing, Web and Apps. What do we need from the Network Side?"
- \*Lucas Pardue, Sreeni Tellakula. "Lower layer performance not indicative of upper layer success"
- \*Ahmed Aldabbagh. "Regulatory perspective on measuring network quality for end users"
- \*Michael Welzl. "A Case for Long-Term Statistics"
- \*Joachim Fabini. "Objective and subjective network quality"

- \*Matt Mathis. "Preliminary Longitudinal Study of Internet Responsiveness"
- \*Brandon Schlinker. "Internet's performance from Facebook's edge"
- \*Jonathan Foulkes. "Metrics helpful in assessing Internet Quality"
- \*Vijay Sivaraman, Sharat Madanapalli, Himal Kumar. "Measuring Network Experience Meaningfully, Accurately, and Scalably"
- \*Dave Reed, Levi Perigo. "Measuring ISP Performance in Broadband America: a Study of Latency Under Load"
- \*Kyle MacMillian, Nick Feamster. "Beyond Speed Test: Measuring Latency Under Load Across Different Speed Tiers"
- \*Gregory Mirsky, Xiao Min, Gyan Mishra, Liuyan Han. "Error Performance Measurement in Packet-Switched Networks"
- \*Gino Dion. "Focusing on latency, not throughput, to provide better internet experience and network quality"
- \*Praveen Balasubramanian. "Transport Layer Statistics for Network Quality"
- \*Jari Arkko, Mirja Kuehlewind. "Observability is needed to improve network quality"
- \*Robin Marx, Joris Herbots. "Merge Those Metrics: Towards Holistic (Protocol) Logging"
- \*Rajat Ghai. "Measuring & Improving QoE on the Xfinity Wi-Fi Network"
- \*Koen De Schepper, Olivier Tilmans, Gino Dion. "Challenges and opportunities of hardware support for Low Queuing Latency without Packet Loss"
- \*Olivier Bonaventure, Francois Michel. "Packet delivery time as a tie-breaker for assessing Wi-Fi access points"
- \*Ken Kerpez, Jinous Shafiei, John Cioffi, Pete Chow, Djamel Bousaber. "State of Wi-Fi Reporting"
- \*Mikhail Liubogoshchev. "Cross-layer Cooperation for Better Network Service"
- \*Sandor Laki, Szilveszter Nadas, Balazs Varga, Luis M. Contreras. "Incentive-Based Traffic Management and QoS Measurements"

\*Satadal Segupta, Hyojoon Kim, Jennifer Rexford. "Fine-Grained RTT Monitoring Inside the Network"

\*Al Morton. "Dream-Pipe or Pipe-Dream: What Do Users Want (and how can we assure it)?"

\*Kalevi Kilkki, Benajamin Finley. "In Search of Lost QoS"

\*Neil Davies, Peter Thompson. "Measuring Network Impact on Application Outcomes using Quality Attenuation"

\*Mingrui Zhang, Vidhi Goel, Lisong Xu. "User-Perceived Latency to measure CCAs"

\*Discussion

\*Break

\*Christoph Paasch, Randall Meyer, Stuart Cheshire, Omer Shapira. "Responsiveness under Working Conditions"

\*Bob Briscoe, Greg White, Vidhi Goel and Koen De Schepper. "A single common metric to characterize varying packet delay"

\*Christoph Paasch, Kristen McIntyre, Randall Meyer, Stuart Cheshire, Omer Shapira. "An end-user approach to the Internet Score"

## **4. Discussions**

The three day workshop was broken into four separate sections, including introductory material and conclusions, that each played a role in framing the discussions.

### **4.1. Introduction and overviews**

The Introduction section allowed participants to introduce and discuss the problem space, existing mechanisms for QoS and QoE measurements. Also discussed was the interaction between multiple users within the Network, as well as the interaction between multiple layers of the OSI stack. Some existing measurement works were presented. Vint Cerf provided a key note support describing the history and importance of the topic.

### **4.2. Metrics considerations**

The Metrics section of the workshop concentrated on both defining new and existing measures and how they might apply to different sections of the Internet. The need for improvements to latency and

its measurements was heavily discussed, especially for certain classes of users such as live, collaborative content and gaming.

#### **4.3. Cross-layer considerations**

In the Cross-layer section participants present material and discussed how accurately measuring exactly where problems occur is difficult when many components of a network connection can affect the measurement. Discussion centered especially on the differences between physically wired and wireless connections and the difficulties of accurately determining problem spots when multiple different network types are responsible the quality.

#### **4.4. Synthesis**

Finally, in the Synthesis section presentations and discussions concentrated on the next steps likely needed to make forward progress. Of particular concern is how to bring forward measurements that can make sense to end users trying to make subscription decisions.

### **5. Conclusions**

During the final hour of the workshop we gathered statements that group thought were summary statements from the 3 day event. We later discarded any that were in contention (listed further below for completeness). For this document, the editor took the original list and divided it into rough categories, applied some suggested edits discussed on the mailing list and further edited for clarity and to provide context.

#### **5.1. General statements**

1. Bandwidth is necessary but not alone sufficient
2. In many cases, Internet users don't need more bandwidth, but rather need "better bandwidth" - i.e., they need other improvements to their connectivity.
3. We need both active and passive measurements - passive measurements can provide historical debugging.
4. We need passive measurements to be continuous and archivable and queriable - include reliability/connectivity measurements.
5. A really meaningful metric for users is whether their application will work properly or fail because of a lack of a network with sufficient characteristics.

6. An useful metric for goodness must actually incentive goodness  
- good metrics should actionable to help drive industries toward improvement.
7. A lower latency internet, however achieved would benefit all end users.

## **5.2. Specific statements about detailed protocols/techniques**

1. Round trips Per Minute (RPM) is a useful, consumable metric
2. We need a usable tool that fills the current gap between network reachability, latency and speed tests.
3. End-users that want to be involved in QoS decisions should be able to voice their needs and desires.
4. Applications are needed that can perform and report good quality measurements in order to identify insufficient points in network access.
5. Research done by regulators indicate that users/consumers prefer a simple metric per application, which frequently resolves to whether the application will work properly not.
6. New measurements and QoS or QoE techniques should not rely only or depend on reading TCP headers.
7. It is clear from developers of interactive applications and from network operators that lower latency is a strong factor in user QoE. However, metrics are lacking to support this statement directly.

## **5.3. Problem statements and concerns**

1. Latency mean and medians are distractions from better measurements.
2. It is frustrating to only measure network services without simultaneously improving those services.
3. Stakeholder incentives aren't aligned for easy wins in this space. Incentives are needed to motivate improvements in public network access. Measurements may be one step toward driving competitive market incentive.
4. For future-proof networking, measuring ecological impact of measuring material and energy usage is important.



5. We do not have incontrovertible evidence that any one metric (e.g. latency or speed) is more important than others to persuade device vendors to concentrate on any one optimization.

#### **5.4. No-consensus reached statements**

Additional statements were recorded that did not have consensus of the group at the time, but we list here for completeness about the fact they were discussed:

1. We do not have incontrovertible evidence that buffer bloat is a prevalent problem
2. The measurement needs to support reporting localization in order to find problems. Specifically:

- \*Detecting a problem is not sufficient if you can't find the location

- \*Need more than just english - different localization concerns

3. Stakeholder incentives aren't aligned for easy wins in this space

#### **6. Follow on work**

There was discussion during the workshop about where future work should be performed. The group agreed that some work could be done more immediately within existing IETF working groups, while other longer-term research may be needed in IRTF groups.

#### **7. Security considerations**

A few security relevant topics were discussed at the workshop, including but not limited to:

- \*What prioritization techniques can work without invading the privacy of the communicating parties.

- \*How oversubscribed networks can essentially be viewed as a DDoS attack.

#### **8. References**

##### **8.1. Normative References**

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/

RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005, <<https://www.rfc-editor.org/info/rfc4035>>.

[RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", RFC 5155, DOI 10.17487/RFC5155, March 2008, <<https://www.rfc-editor.org/info/rfc5155>>.

## 8.2. Informative References

[RFC1111] Postel, J., "Request for comments on Request for Comments: Instructions to RFC authors", RFC 1111, DOI 10.17487/RFC1111, August 1989, <<https://www.rfc-editor.org/info/rfc1111>>.

## Appendix A. Participants List

The following is a list of participants attended the workshop over a remote connection:

Ahmed Aldabbagh  
Al Morton  
Alexander Clemm  
Alvaro Retana  
Anna Brunstrom  
Balazs Varga  
Bjørn Ivar Teigen  
Bob Briscoe  
Brandon Schlinker  
Bren Tully Walsh  
Christoph Paasch  
Cindy Morgan  
Cullen Jennings  
Dan Siemon  
Dave Taht  
David Reed  
David Schinazi  
Djamel Bousaber  
Eve Schooler  
Evgeny Khorov  
François Michel  
Gavin Young  
Geoff Huston  
Gino Dion  
Gorry Fairhurst  
Greg Mirsky  
Greg White  
Jana Iyengar  
Jared Mauch  
Jari Arkko  
Jason Livingood  
Jiankang Yao  
Jim Gettys  
Jinous Shafiei  
Joachim Fabini  
John Cioffi  
Jonathan Foulkes  
Joon Kim  
Joris Herbots  
Kalevi Kilkki  
Karthik Sundaresan  
Kathleen Nichols  
Keith Winstein  
Ken Kerpez  
Kenjiro Cho  
Koen De Schepper  
Kristen McIntyre  
Kyle MacMillan  
Lai Yi Ohlsen

Lars Eggert  
Levi Perigo  
Lisong Xu  
Lucas Pardue  
Luis M. Contreras  
Mat Ford  
Matt Mathis  
Michael Welzl  
Mikhail Liubogoshchev  
Mingrui Zhang  
Neil Davies  
Nick Feamster  
Nicolas (Tessares)  
Olivier Bonaventure  
Omer Shapira  
Pedro Casas  
Peter Thompson  
Praveen Balasubramanian  
Rajat Ghai  
Randall Meyer  
Rich Brown  
Rick Taylor  
Roberto  
Robin Marx  
Russ White  
Sam Crawford  
Satadal Sengupta  
Shapelez  
Sharat Madanapalli  
Steve Christianson  
Stuart Cheshire  
Szilveszter Nadas  
Toerless Eckert  
Toke Høiland-Jørgensen  
Tommy Pauly  
Vesna Manojlovic  
Vidhi Goel  
Vijay Sivaraman  
Vint Cerf  
Wes Hardaker  
Zhenbin Li

## **Appendix B. IAB Members at the Time of Approval**

Internet Architecture Board members at the time this document was approved for publication were:

Jari Arkko

Deborah Brungard

Ben Campbell

Lars Eggert

Wes Hardaker

Cullen Jennings

Mirja Kühlewind

Zhenbin Li

Jared Mauch

Tommy Pauly

Colin Perkins

David Schinazi

Russ White

Jiankang Yao

## **Appendix C. Acknowledgements**

The authors would like to thank the workshop participants, the members of the IAB, and the program committee for creating and participating in many interesting discussions.

### **C.1. Workshop Chairs**

The workshop chairs consisted of:

Evgeny Khorov

Omer Shapira

Wes Hardaker

### **C.2. Program Committee**

The program committee consisted of:

Christoph Paasch  
Cullen Jennings  
Geoff Huston  
Greg White  
Jari Arkko  
Jason Livingood  
Jim Gettys  
Katarzyna Kosek-Szott  
Kathleen Nichols  
Keith Winstein  
Matt Mathis  
Mirja Kuehlewind  
Nick Feamster  
Olivier Bonaventure  
Randall Meyer  
Sam Crowford  
Stuart Cheshire  
Toke Hoiland-Jorgensen  
Tommy Pauly  
Vint Cerf

#### **Appendix D. Github Version of this document**

While this document is under development, it can be viewed, tracked, fill here:

<https://github.com/intarchboard/network-quality-workshop-report>

#### **Author's Address**

Wes Hardaker  
USC/ISI

Email: [ietf@hardakers.net](mailto:ietf@hardakers.net)