

## M-Lab draft response to the FCC's submission to the Broadband Forum

This document comprises targeted suggestions from the network research community concerning the scope and content of the FCC's proposal to the Broadband Forum

**Date:** August 21, 2012

**Authors:**

- Nick Feamster  
Author of the BISmark measurement tool, Professor, University of Maryland Institute for Advanced Computer Studies
- Matt Mathis  
Author of the NPAD network measurement tool, Network Research Scientist, Google
- Thomas Gideon  
Technical Director, Open Technology Institute at the New America Foundation

On behalf of the M-Lab research consortium, we submit these comments for consideration by the FCC and the Broadband Forum.

It's important to emphasize at the outset of this measurement architecture standardization effort that M-Lab provides the current de-facto standard for collecting open, scientifically verifiable measurement on a global scale. *We strongly encourage the commission and those involved in this effort at the broadband forum to include M-Lab's platform as a core component of any measurement effort.* M-Lab relies on PlanetLab, and its architecture incorporates the combined decades of experience from across the network research spectrum. Building on this base, M-Lab has created a robust, extensible, operationalized platform that provides the benefits of open, accessible, scientifically verifiable broadband performance data on a global scale.

Scientific verifiability is M-Lab's core value, and is the foundation of every decision we make. We urge all efforts in this space to adopt the same.

Scientific verification relies on researchers' ability to replicate published results. Verification and replication, in turn, rely on openness and transparency at every step. In the measurement space, the measurement tools, methodology, and raw data must be open to scrutiny, such that the scientific community can bring their broad insights to bear on understanding the infinite complexities, dependencies, and massive scale of global networks. Only through the process of deep re-evaluation of open, raw data, can we discover subtle artifacts in the data that might affect the measurements, and the conclusions drawn from these measurements.

Tools or data that are private and privately collected are potentially subject to undetected systematic measurement bias, simply because a small pool of researchers is unlikely to notice all of their own assumptions and blind spots. As such, even if one can assume that no malice is present, data and methods that can not be independently verified are unlikely to be relied upon in the broader community.

Only an iterative process that includes developing tools and evaluating them in an open scientific forum with access to the raw the data, the measurement methodologies (in the form of open source tools and openly-documented infrastructure) can result in measurement that truly addresses the concerns of all stakeholders.

***The authors, representing the academic community behind the M-Lab effort, believe strongly that no standard should move forward without explicitly provisioning for openness and transparency, especially a standard that is set to provide the foundation for national and global assumptions on the state of broadband networks. We urge the FCC to revise its comment to the Broadband Forum with this in mind, and we urge the Broadband Forum working groups concerned with this issue to found their effort on these principles.***

Adherence to the following tenants can ensure scientifically verifiable measurement. We present these as the foundation on which any effort concerning broadband measurement standardization should be built:

- **Openly documented, extensible infrastructure and hardware.**  
It is important, furthermore, that the requirements not be prohibitive. For example, implementation of an “open” standard should not require ownership of the network infrastructure itself.
- **Open-source tools (test methodologies) running active measurements.**  
We specify “active measurements” so as to ensure that data can be made truly open without a risk of leaking user PII. Furthermore, active measurements (as opposed to passive monitoring) allows true reproducibility and sampling calibration that isn’t possible using aggregate data.
- **All data collected, in its raw form, is made available to the public.**  
The assumption for any standardization effort should be the production of open, accessible raw measurement data. Open, raw data benefits the research community, and allows true verifiability. An implementation of a standard delivered by a national regulator, academic institution, or ISP, should provide access to all data on which public claims of performance are based.
- **Statistical assumptions and methodologies are made publicly available.**  
While not the purview of an infrastructure standardization effort specifically, the means by which raw data are transformed into concrete, meaningful numbers and assumptions must be clear, and insofar as infrastructure plays a part in this computation, this consideration must be taken into account in constructing infrastructure standards.
- **The research community must be actively involved.**  
The management and operations of a broad platform should fall to the research community. This structure avoids the issue of perceived conflict of interest, and allows consistent and adaptive practices to scale across a global platform. The M-Lab platform should be taken as a model here.

#### **Notes, questions, and suggested revisions of the FCC’s BBF proposal:**

To provide targeted, expedient comments, we have addressed specific sections of the FCC’s comments to the Broadband Forum below. Overall, we strongly encourage the Commission to review and revise the document to ensure that the tenets of open, verifiable measurement are provisioned for explicitly.

1. Background
  - a. The document needs to clearly define what’s meant by “build[ing] measurement capability into existing network elements rather than through the provision of additional equipment.” In particular, it’s important that this specification not imply ownership of the network infrastructure itself.

- b. This section stresses “consistent” measurement. However, if we’re talking about a cursory repurposing of existing network infrastructure as a “measurement network” consistency is nearly impossible. It’s important that this be clearly defined. A measurement platform must be carefully implemented, consistently managed, and openly documented.
  - c. It’s not clear how end-to-end openness will be achieved if (as implied) the measurement network integrates closed infrastructure.
  - d. It is not clear what specifically is meant by a “more integrated approach.”?
2. Standards Proposal
- a. This item discusses the way in which standardization could serve multiple constituencies (ISPs, regulators, academic), and reduce costs. We support this overall sentiment. In fact, M-Lab exists currently as a completely extensible de-facto global standard, serving these three constituencies at incredibly low cost. We strongly encourage the Commission, and the Broadband Forum, to consider M-Lab as an existing core component of any proposed standard.
3. Data Collection Infrastructure Goals

*End to end model*

Note that to the IETF and Internet community “end-to-end” means all the way from a content provider to an end-user. Passive end-to-end measurement could be interpreted as advocating that a given content provider deploy instrumentation in their servers and client devices. While this is an interesting proposal, and would no doubt result in interesting data, *it will not address the need for open data and open tools*. Such measurements could not be validated by all parties, and thus wouldn’t be scientifically verifiable or, ultimately, credible in the broader community.

- a. It’s important that the implementation details on measurement from a given end-point to an application provider are supplied, and that it’s clear whether this could be implemented openly, and by means of active measurement.
- b. If this is a specification calling for measurement of live applications, we suggest instead implementing M-Lab’s model of simulated application protocols. These enable open, transparent measurements that can be replicated, as specified in the tenets above. It is crucial that the correct assumptions -- assumptions calling for open, verifiable measurement -- be those on which this standards effort is founded.
- c. We understand the statement “data test endpoints to include any desired point in an end-to-end model” to infer a model of “sectionalizing” -- testing a long path section by section to understand its component parts. It’s important that there be a clear provision to implement the final form openly, and that the specification also ensure that implementation be possible for entities apart from network operators.
- d. Specifically, it’s important to be clear about who would (or, potentially, could) control the “point-to-point” measurements along this path. We encourage a clarity on exactly how this implementation can be made openly.

*Comparable results*

- e. There is mention that test data needs to be consistent across a range of “different technologies.” This should be self-evident: for example, metrics intended to characterize the user’s experience should be comparable across technologies such that they can bring data to bear on the relative benefits of

the user's choice of technology. Consistency in data collection also requires consistent, consistently managed measurement infrastructure, and this should be taken into account.

#### *Adaptability*

- f. It's important to square "variety of physical implementations allowing the greatest degree of freedom in implementation" with a commitment to openness and consistent measurement. Much more detail on the exact nature of this freedom of implementation is necessary, as well as a commitment to strict provisions that will ensure that measurements are open and consistent.

#### *Multiuser Support*

- g. We note that the two defined user groups consist of ISPs and national regulators, with academics as a secondary consideration. It's important to stress that for the former groups to succeed at any such endeavor, participation and close involvement by the research and academic community is crucial. This goes for direct involvement, as well as indirect involvement in which the global scientific community is able to review and reproduce results achieved by an ISP or an FCC program.
- h. With the above in mind, an open implementation is imperative; any data that is used as the basis for publicly stated assumptions (from policy, to ISPs' claims to consumers) needs to be verifiable.
- i. The document mentions "how tests are scheduled." It is again important to stress that an independent entity needs to be involved in operating and managing the platform, and that if tests are going to be scheduled on this open platform, the methods and assumptions involved must be open to review.
- j. The phrase, "how data access is obtained" needs to be clarified. We contend strongly that data must be open. *The assumption of any global measurement standard should be the support of truly open, publicly accessible data.*
- k. To this end, storage, access, and an automated pipeline should be a part of any design. We encourage the commission to revisit its draft and include these components explicitly.

#### *Consumer privacy*

- l. M-Lab's model is consistent with the goal of both providing a truly open measurement ecosystem, and protecting consumer privacy. Again, it should be considered as a core component that any standardization effort integrates and builds on, with the participation of the M-Lab research collaborative.