



**Before the  
Federal Communications Commission  
Washington, D.C. 20554**

In the Matter of: )  
)  
Comments - Mobile Broadband Measurement ) CG Docket No. 09-158  
) CC Docket No. 98-170  
) WC Docket No. 04-36  
)

**COMMENTS OF  
INFORMATION TECHNOLOGY AND INNOVATION FOUNDATION**

Richard Bennett

Information Technology and Innovation Foundation<sup>1</sup>  
1101 K St N.W.  
Suite 610  
Washington, DC 20005

---

<sup>1</sup> ITIF is a nonprofit, non-partisan public policy think tank committed to articulating and advancing a pro-productivity, pro-innovation and pro-technology public policy agenda internationally, in Washington and in the states. Through its research, policy proposals, and commentary, ITIF is working to advance and support public policies that boost innovation, e-transformation and productivity.

The Information Technology and Innovation Foundation (ITIF) is pleased to offer the following comments on *The Measurement of Mobile Broadband Network Performance and Coverage*. ITIF advocates full disclosure of network performance and reliability metrics that impact the user experience on the broadband Internet in both its wireline and wireless forms. Customers should not have to guess about the general performance parameters that a given network provides when buying a device and committing to a contract, and networks should perform as well as their operators claim they do. Measuring the performance of mobile IP networks is an exercise fraught with difficulty, however. Therefore, we recommend a system of operator-supplied coverage maps, primarily verified by consumers through extensive refund rights when user experience does not conform to the claims implicit in the coverage map. The data collected for the typical FCC report is too stale to be of much use by the time the reports are released, but operator-constructed maps are dynamic and up-to-date. A secondary level of verification can be supplied by opt-in systems of measurement at the handset and the web site level, as described in the following.

### **The Internet is Inherently Unpredictable**

Measuring the performance of non-deterministic, packet-switched networks such as those comprising the Internet is much more difficult challenge than is generally appreciated, however. As currently operated, the Internet provides no performance guarantees, relying on a “best-effort” system of packet transfer across facilities shared by a large number of users – some 500 million systems are attached to the Internet presently – operating under wildly different loading scenarios. Many advocates argue that this “best-effort” system represents an ideal state of affairs, and are offended by the notion that network operators might supplement basic service with a more deterministic, for-fee system with bounded performance guarantees.

At a high level, the Internet as a whole exhibits marked variations of performance in diurnal, daily, and seasonal cycles, and the systems that provide World-Wide Web and streaming audio/video services are subject to radical fluctuations in load on very short time scales. In fact, the most extreme variations in the Internet user experience are caused by variations in the demand for particular pieces of content. Shocking photographs or amusing videos can flat-line servers so completely that the typical user can’t distinguish the phenomenon from a disconnected cable. Cloud computing platforms are also highly variable with respect to performance, as a recent study by Bitcurrent cautions: “the data in this report should serve only as a guideline for further testing: **your mileage will vary greatly.**”<sup>2</sup> Any meaningful measurement of the Internet experience requires a large number of samples taken under a variety of scenarios and in a variety of locations.

---

<sup>2</sup> Emphasis in original. “Cloud Computing Performance: A Bitcurrent study on the performance of cloud computing platforms” (Bitcurrent, June 2010), [http://www.webmetrics.com/landingpage/bitcurrentcloud/The\\_Performance\\_of\\_Clouds\\_Complete.pdf](http://www.webmetrics.com/landingpage/bitcurrentcloud/The_Performance_of_Clouds_Complete.pdf).

## **Mobile Networks are Inherently Unpredictable**

Mobile broadband networks add an additional level of non-deterministic behavior to the Internet's intrinsically statistical nature. Mobile radio networks are inherently non-deterministic for three reasons:

- 1) They share radio resources such as frequencies and digital codes among users;
- 2) The communications medium they employ – air – is more susceptible to noise, interference, and signal fade than is wire; and
- 3) Users roam about the network in ways that the network operator must attempt to predict but can't control.

Consequently, variations in transient conditions are more radical on mobile networks than on wireline networks. Internet services over mobile radio networks is therefore multiply non-deterministic, so the basic premise of network performance measurement – that past performance predicts future results – is questionable. Internet service over mobile broadband is highly variable with respect to location, so measurements taken on one part of a network will not even predict performance on other parts of the same network at the very same time. In addition to the requirement for frequent sampling, mobile performance sampling needs to be fair and geographically representative as well.

## **There is a Risk in Over-Generalizing Survey Data**

Data that appears authoritative – because it's issued by a federal agency – but happens to be partial and incomplete can lead to unfortunate results in a “data-driven policy” exercise. A poorly-crafted Internet performance framework will support inconsistent policy preferences, such as the demand that the non-deterministic mobile Internet should provide a deterministic Quality of Service more consistent with wireline circuit switching than with cost-effective and powerful mobile packet switching. Because the FCC has been engaged in the regulation of wireline circuit-switched networks since 1934, it risks bringing technological and historical bias to the exercise without even knowing it.

Sophisticated test and measurement tools exist that allow the skilled tester to measure packet latency and throughput to the sub-millisecond level; such tools are a routine part of the engineering Quality Assurance process employed by all major vendors of handsets, network equipment, and network services. The measurements obtained by such tools are taken under carefully controlled conditions, however, that can't be duplicated in the wild. So the questions for meaningful consumer disclosure aren't so much those that concern what can be measured – many aspects of the mobile Internet can be measured in a variety of scenarios with microsecond granularity – as much as what aspects of measurement provide the user with meaningful insight. In other words, we need to be clear at the outset what we are measuring and why we are measuring it, and what the practical limits of the testing scenario are.

## **Applications, Users, and Handsets Strongly Influence Performance**

At a high level, performance measurement of the wireline Internet reflects the application mix that's active on the portion of the Internet being measured, in proportion to the

capacity of that portion of the Internet. In practice, bandwidth demand and supply are indistinguishable, and communication bandwidth is indistinguishable from server capacity (often called “bandwidth,” coincidentally.) On the wireless Internet, measurements taken from a number of handsets with different performance and usage characteristics tell us more, in the aggregate, about the application and handset mix on a given network than about the network’s intrinsic performance. Performance measurements tend to be snapshots of the performance of the measurement application on a portion of a network at a moment in time that are notoriously non-repeatable. Given that observation repeatability is one of the hallmarks of the scientific method, the significance of non-repeatable data is very, very limited.

The FCC’s Notice suggests a number of parameters such as *typical data throughput, signal strength, accessibility, retainability, latency, other quality of service parameters*<sup>3</sup> that are more strongly influenced by the handset itself than by the network, in most scenarios. These factors are also strongly influenced by the orientation of the handset with respect to the connected cell tower, the distance between the cell tower and the user, the nature of the physical barriers between handset and tower, and the activities of other users of the tower, as well as the state of the Internet as a whole. The assumption that we can simply run measurement applications on arbitrary handsets at arbitrary times and collect meaningful data about the network that can then be aggregated in a meaningful way across the entire ecosystem doesn’t hold up in real life.

## Taking Meaningful Measurements

What measurements, then, will provide users with meaningful insight about the performance of a given mobile network? We believe that a simple model can be constructed around vendor claims and user experience that can be measured in a repeatable, replicable way. This model is driven by the claims that network operators make for the services they provide pertaining to the classic mobile applications, which at this point are telephone calls and web surfing. It deliberately omits leading-edge applications whose performance is most strongly affected by new handsets and the capacity upgrades in progress on all mobile networks in the United States at present; operators are enhancing networks with fourth-generation LTE and “third-generation plus” technologies such as HSPA+ and EV-DO Revision B and handset vendors are following suit with compatible devices, but we’re in mid-transition to these new technologies. The measurements we suggest are the following:

1. Phone calls – the telephone application, typically charged against buckets of minutes, uses the semi-deterministic, circuit-switched network service and should be relatively immune from the effects of packet-switched applications on most networks. Users increasingly rely on mobile phones as their primary means of direct interpersonal communication, and the performance of calls is easy to observe and measure. The dimensions of calling are well understood:
  - a. In what areas can the user obtain a dial tone?

---

<sup>3</sup> Public Notice, page 2.

- b. Do calls made within the service area to wireline phones get connected?<sup>4</sup>
- c. Does the call remain connected until the user decides to disconnect?

We could go on to measure call quality, but it's highly variable, somewhat subjective, and often dependent on Bluetooth-connected headsets of variable quality. When call quality becomes untenable, a disconnection is likely in any event. Phone call performance is variable across cities for reasons that network operators can't always control, such as local zoning restrictions on cell tower siting and antenna size, so this data must be localized. Phone call data, including location, can be obtained from calling logs and doesn't require network modification. This data needs to be obtained per random sample and not from groups of reporters with any sort of agenda; the dynamics of reporting should be similar to opinion polls. Access to user call logs has obvious implications for privacy and security, however; overcoming these barriers while maintaining impartiality in test subjects will be difficult.

So the questions that need to be answered about calling data – and this will be repeated when we talk about web data – is how to ensure that statistically representative samples are taken and how user privacy is protected. There is no easy answer to this pair of questions, which ultimately form the heart of this inquiry. One possible approach is to implement an opt-in system of voluntary sharing of data on call volume, call success, and call failure.

2. Web access – the ability to access the variety of sites with the user's chosen degree of restriction is an important dimension of mobile broadband, independent of performance. Contracts may allow users to invoke parental controls and other content restrictions, or they may be more permissive. Operators may also make voluntary decisions to block certain networks or web sites in order to protect against Denial of Service Attacks, identity theft, or malware, and these activities are fully legitimate. All these factors notwithstanding, the ability to access the whole web, or the portion the user has chosen to access, within reasonable limits<sup>5</sup> is an important element of the mobile Internet experience.
3. Web performance – Claims about the ability to access the web and any claims as to performance should be verified, bearing in mind the fact that web site loading times can be highly variable depending on server location and whether the web site is hosted on a Content Delivery Network. Some operators have invested more than others in core Internet capacity, which makes far-away web sites load nearly as fast as local ones. Those operators who have invested in end-to-end capacity should be recognized, of course, which doesn't happen in testing scenarios that run test servers in the nearest Internet Exchange.

---

<sup>4</sup> Calls made from mobile phones on one network to mobile phones on other networks don't provide meaningful service measurements as we don't know which network to blame if the call fails.

<sup>5</sup> At any given time, mitigation of Denial of Service attacks will make some domains inaccessible to some Internet users. Network operators should not be penalized for DoS mitigation.

Web site page load performance measurement ideally needs to consider a variety of web sites at locations around the country and overseas; the group of the same web sites should be measured from all mobile networks, but they should not be made public. As handsets influence performance, this variable needs to be isolated, as does time of day and user location.

Web sites permit a wide range of content interactions, ranging from loading text pages, graphic images, audio and video files, and plug-ins for rich media formats such as YouTube and Netflix video streams. Web performance measurements should not include streaming performance yet; once the 4G transition is more advanced, this option can be reviewed.

Rather than going into detailed capture of web site load times, the test should set a loose target that corresponds to typical consumer expectation of web site load time and isn't overwhelmed by extraneous factors such as handset performance and transient radio conditions. The general human factors goal for the wireline Internet is to have web pages load in one second or less, but it's not unusual to experience load times that are closer to a minute on handsets with minimal web capability. A reasonable target for successful web site loading on state-of-the-art handsets such as the iPhone and the Nexus One would be on the order of 5 seconds; for lesser devices it might be closer to ten seconds. We certainly can't take web access that requires 30 seconds or more to load a web page as successful, however.

Taking the 5-10 second goal as the measure of success, it becomes possible to measure web access on a pass-fail basis and produce a web access coverage area map similar to a cellular call service area map. These measurements can be conducted by web sites themselves, with no particular cooperation by the handset. Web sites are informed about the nature of the mobile browser and other device characteristics as a matter of course, and may choose to share representative data with the Commission. Some of the most meaningful measurements of Internet performance are taken by Content Delivery Networks such as Akamai, and the value of these measurements is that they're taken as a matter of course rather than under a self-selected testing scenario. Network operators can suggest a set of data capture scenarios that make sense. These data should be reported by handset and by network.

Web site loading data raise a similar set of concerns as call log data: how to ensure the data are representative and respectful of user privacy. These are questions with no easy answers, but most would urge voluntary participation by web sites and anonymizing data before reporting. Measuring web site load times at the web site itself also entails distinguishing origin networks by IP address prefix, so there are twin risks of misidentifying the origin network and reporting the full IP address, which ultimately becomes a personal identifier. Until an acceptable solution is found to these issues, the best course is to rely solely on

operator-supplied coverage maps, but many do not find that course acceptable.

A potential solution to the problems of participation and privacy is to construct an opt-in system for web sites with a robust set of privacy protections. Similar systems, albeit with less detail, are already in use by Akamai and others to measure the state of the Internet as a whole.

The general notion here is restraint from attempts to measure leading-edge network conditions in favor of providing mainstream consumers with meaningful data that relates to the patterns of usage most important to their day-to-day experience with the mobile Internet.

### **Advanced Measurements**

The comments that have been offered to the Commission on this Public Notice describe a number of testing scenarios that generally fall into two categories:

1. End-to-end measurement of user experience
2. Network element testing

End-to-end testing obviously captures the most realistic view of what consumers can expect, but it's fraught with the difficulties mentioned in terms of isolating variables over which the network operator has control. Network element testing aims to isolate specific performance factors, but it's too intrusive with respect to trade secret operator practice and unlikely to be interpreted correctly by the press.

These types of testing are the norm for the wireline Internet, so there's considerable intellectual inertia urging the Commission to adopt them. One major issue with applying such methods to the mobile Internet that hasn't been discussed is the dynamic nature of mobile networks with respect to technology upgrades.

Wireline broadband networks are upgraded on a regular basis as advances in technology permit higher speeds over particular systems of wiring, operators increase backhaul capacity, and peering agreements among network operators overcome the "meet-me switch" bottleneck in Internet Exchanges. Similarly, mobile networks are upgraded on a regular basis for similar reasons, but the generational changes that improve mobile networks are more frequent and more dramatic than their counterparts in the wireline world. DOCSIS 3, the leading-edge cable modem standard which appeared in 2006, was the first major speed upgrade for cable modems since DOCSIS was initially standardized in 1997.<sup>6</sup> DOCSIS 3 increased aggregate cable network speed from 38 to 152 Mbps, a factor of four. During the same period, the 3GPP standard for mobile networks increased data capacity from 14.4 Kbps to as much as 56 Mbps, a factor of 4,000. There is clearly a major difference in the rate of change between these two technology families.

---

<sup>6</sup> "DOCSIS," in *Wikipedia*, <http://en.wikipedia.org/wiki/DOCSIS>.

Consequently, a culture has developed in which it's acceptable to take measurements from a network and then spend a year or more analyzing them before publishing results. The FCC's reports on broadband network performance aren't published within the same year that their data were collected; the last Broadband Deployment Report, published this July, was based on data collected in 2008. A delay as long as this completely invalidates any reports on mobile network performance.

The wireline Internet is also much more stable and predictable than the mobile Internet with respect to application mix and perceived performance. Laptop and desktop computers steadily improve in performance, but they haven't been significant performance bottlenecks for popular Internet applications for a very long time. The marketplace for mobile applications and devices is very different today than it was at the advent of the iPhone in 2007 and again at the launch of Apple's app store a year later. The iPhone set a new standard for consumer expectations of mobile performance, but its applications only began to appear after the data for the FCC's latest Mobile Competition Report were collected.

Consequently, any gathering of data and publishing of reports on mobile performance has to be conducted in a much timelier manner than the norm for wireline performance, and it must emphasize rates of change, deployment of new network technology, and trend lines in investment. Given the dizzying rate of mobile Internet technology upgrades, it would be more meaningful to report the current year's investment figures than to publicize performance data from the year before last.

Consumer choices predicated on investment are more likely to be predictive than those made on the basis of stale end-to-end and network element measurements.

## **Conclusion**

We live in a time in which "data-driven policy" has become a watchword. While the slogan is appealing, it raises the temptation to collect only those data that are likely to support a given policy conclusion; when this happens, we have "data driven by policy" instead. Every exercise in data collection can continue, in principle, without any limit. It's not appropriate to burden network operators, device manufacturers, and application developers with data collection efforts that don't have real significance for consumers, so we urge restraint on the FCC with respect to its data-gathering ambitions.

Ultimately, the factors that matter most to consumers are their personal satisfaction with the devices, applications, and networks they purchase, which is assured over the long term by the rate of improvement in underlying infrastructures and technologies. These factors can be assessed by largely non-technical means, so any forays into technical data gathering have to be taken cautiously and with full awareness that they're unlikely to add much to the consumer experience.

Until the issues with sampling and privacy are resolved, it's best to rely primarily on operator-supplied coverage maps for voice and basic web access, the accuracy of which

can be assured by both consumers and independent means. Network operators should certainly be encouraged to publish continually-updated and quite precise coverage maps, and consumers should be entitled to refunds and to release from contracts where inaccurate maps have induced a buying decision. The mobile Internet is too dynamic for the FCC to move beyond this basic kind of reporting at present, however, and we should hope that it remains that way.