FCC's "Measuring Broadband America" Report Tells Only Half the Story

By Peter Sevcik
November 2011

In August of 2011, the FCC published its *Measuring Broadband America* report [1], which publicizes measurements of US ISPs' broadband performance. The report summarizes initial results of an unprecedented, ongoing performance measurement project to enable consumers to validate ISP broadband performance claims and compare performance across ISPs. NetForecast examined the FCC report and the underlying data.

In NetForecast's view, the FCC report is the most thorough and accurate account to date of the true performance delivered by US ISPs. It does, however, fall short in several respects, notably that some of the test servers introduced data inconsistencies and errors; there is insufficient data analysis transparency; and the FCC incorrectly uses the terms 'speed' and 'bandwidth' interchangeably.

Most importantly, the report covers only half of the performance equation—the supply half. NetForecast believes it is time for the FCC to expand its investigation to include the demand half of the equation—the prodigious growth in resource consumption by major websites, browsers, and behind-the-scenes services. A report on the rest of the equation will reveal the need for an Energy Star equivalent for Internet applications such as Google and Facebook.

NetForecast Assessment of the SamKnows-based Test Methodology

The FCC's broadband performance monitoring project involves measuring Internet connection performance at consumers' homes using on-site performance measurement 'white boxes' provided by the British firm SamKnows. At the other end of each consumer test site is an array of target test servers.

Based on our analysis of the FCC test methodology and our experience with alternative methodologies, we conclude that the SamKnows consumer premises measurement boxes produce high-quality test results.

The SamKnows boxes are well suited to the FCC testing because they are minimally affected by consumers' home network, devices, and traffic and they provide:

- Better accuracy than alternative measurement solutions
- Controlled performance measurements
- Consistent and repeated tests
- Broad ISP subscriber participation
- Tests of many critical performance parameters
- Reliable customer premise devices

The array of servers on the other end of the consumer tests; however, are sources of data inconsistencies and even errors because their locations are not subject to due diligence, and they are managed by a hodgepodge of organizations and administrative controls.

On the data analysis front, we note insufficient transparency in explaining the data analysis methodology and we note inconsistencies between the stated data analysis plan and its execution. More descriptive information is needed to determine and validate how data calculations behind the FCC results were made.

These criticisms should be viewed as areas for future improvement. On the whole, the FCC report is the most thorough and accurate documentation of the true performance

NetForecast Report 5108

©2011 NetForecast, Inc.

delivered by US ISPs, and is a commendable start to the FCC's ongoing ISP performance monitoring program.

'Bandwidth' Is Not 'Speed'

The FCC erroneously uses the terms 'bandwidth' and 'speed' interchangeably. In the 360-page *FCC National Broadband Plan*, the FCC uses the word speed 360 times and the word bandwidth only 40 times. ISPs claim to deliver high speeds if you buy their higher bandwidth services. This misleads consumers into assuming that a higher bandwidth connection will automatically deliver a faster user experience. In fact, bandwidth is but one of a half dozen factors that affect user response time (a.k.a. speed).

The correct way to refer to the bandwidth is as capacity, NOT speed. Yet the FCC has chosen to adopt the ISP marketing term for bandwidth. Here is a simple equation NetForecast developed to explain the difference between speed and bandwidth. The equation is annotated showing who/what is responsible for each factor in this "speed" equation. ISPs are primarily responsible for just one factor.

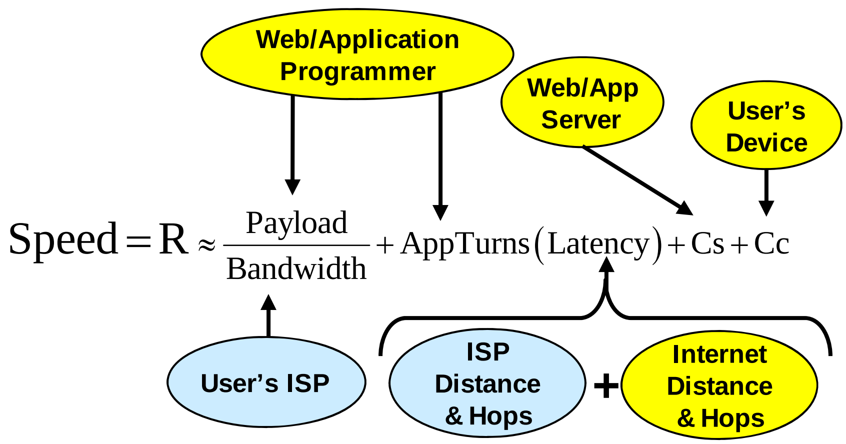

$$\text{Speed} = R \approx \frac{\text{Payload}}{\text{Bandwidth}} + \text{AppTurns}(\text{Latency}) + Cs + Cc$$

Where:

Task response time (R) is the most useful measure of the speed of a user's experience. A task is each user interaction with the application during a session or process. Task time is measured from the time the user enters an application query, command, function, etc. requiring a server response, to the moment the user receives the response and can proceed. Some call this "user wait time"—or in the case of the Web—"page load time" (PLT). The aggregation of these individual task completion times defines application "responsiveness" perceived by the user.

Payload is information content in bits (8xBytes) that must be delivered to/from the user's device. It is determined by the application or website developer.

Bandwidth is the minimum capacity (in bits per second) across all network links between the user and the application server. The slowest link is typically the user's access line to the Internet that is supplied by the user's ISP. Effectively useable link bandwidth may be reduced by congestion and protocol inefficiency (e.g., small TCP window size).

AppTurns are the application client-server software interactions (turn count) needed to generate a user-level system response or task. Turns are determined by those who program an application or website. They do not include two-way

TCP interactions (e.g., ACKs). Newer browsers try to mitigate AppTurns by operating multiple parallel connections to the server(s). In practice this modestly improves overall response time.

Latency is the round-trip-time (in seconds) between the user and the application server. Latency is determined by the physical user-server distance which is limited by the speed of light and the number of switching hops that each packet must traverse. Latency due to distance is governed by the speed of light. Latency due to hops is determined by network topology. The distributed nature of the Internet creates the need for many hops between two points (all users cannot be one hop away from all sources). Latency measured between a user and server is actually the sum of two paths:

- ISP – the distance and hops required to traverse the user's local ISP until the user's packets reach the public Internet. ISPs can make this path efficient.
- Internet – the distance and hops over the Internet required to reach the server. Neither the ISP nor the user can have any effect on the round trip time (RTT) of this path.

Cc (Compute Client) is the total processing time (in seconds) required by the client device.

Cs (Compute Server) is the total processing time (seconds) required by the server(s).

Note that the AppTurn-Latency product is most often the largest factor affecting response time and broadband ISPs have essentially no control over it. Today's typical web pages often require more than 100 AppTurns to load.

It is unhelpful, confusing, and we suggest just plain wrong, to use bandwidth and speed interchangeably. We suggest that the FCC and ISPs begin referring to bandwidth as capacity, and to define it as one of many factors that can contribute to a faster user experience. For proper clarity, ISPs should tout 'high capacity' rather than 'high speed' Internet connections.

The formula presented above quantifies response time or 'speed' of transactional applications like the web, online games, email, etc. It does not apply to streaming applications such as video, music, online radio, etc. Higher bandwidth will improve the quality of these applications by permitting higher resolution video and better fidelity audio. More bandwidth will also enable a household to operate more video/audio streams at the same time (more users watching different programming). The important point is that in this application category bandwidth also does not translate to speed; it directly translates to capacity.

However you look at it, bandwidth is capacity that enables applications and devices to do more of what they do. Digital devices and applications consume bandwidth capacity just like a car consumes gasoline and a refrigerator consumes electricity. Putting more gas in the tank does not make the car go faster.

Factors Affecting 'Speed' (a.k.a. Response Time)

Like most systems, the Internet is subject to the laws of resource supply and demand. Demand for Internet resources that affect speed is most heavily influenced and driven by the way websites are designed and operated. As Figure 1 shows, the primary factors influencing the demand for resources that affect page load time are payload size, and AppTurn count. Payload consumes bandwidth resources, and AppTurns consume latency.

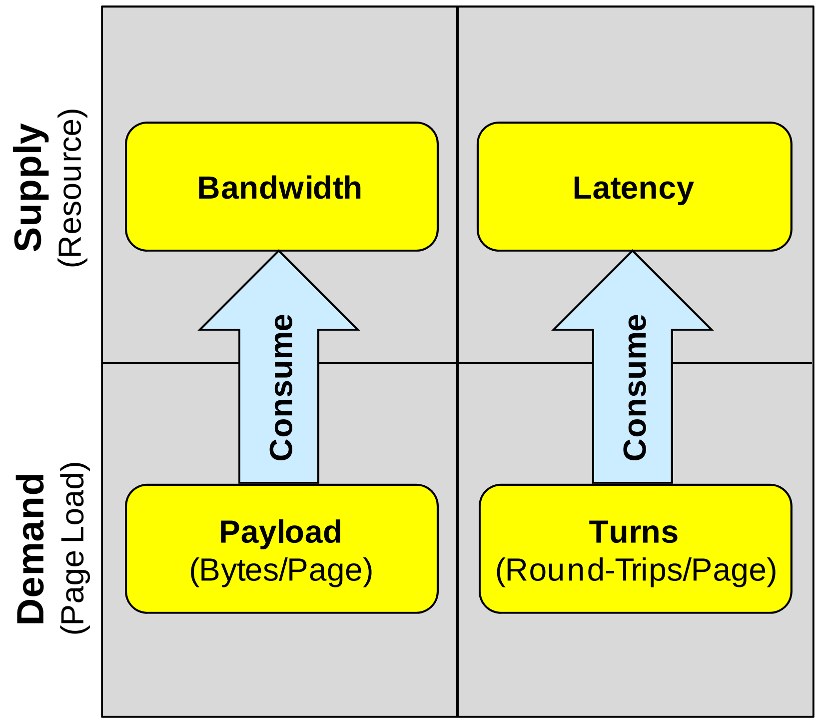

Figure 1. Supply and Demand for "Speed-limiting" Internet Resources

It is important to look at these factors because payload size and turn counts are growing at incredible rates. Figure 2 shows that payload growth is exponential. NetForecast has measured average payload size for typical business-oriented websites since 1995, and during that time payload size has experienced a compound annual growth rate (CAGR) of 20% from 50K Bytes in 1995 to 960K Bytes per page today. NetForecast started measuring when a Fortune 50 company asked us to study the effect of the web on their corporate network. Later, we used the Keynote Business 40 as the proxy for typical business websites. Many consumer websites generate much more traffic per page.

A recent phenomenon contributing to payload traffic growth is a dramatic increase in what we call "clandestine traffic" or behind-the-scenes traffic of which users are largely unaware. This includes cookies, content prefetching, updates, etc. For example, website 'background' traffic accounts for about 15K Bytes per second and Google Chrome "features" add another 6K Bytes per second. This hidden traffic is *not* shown in Figure 2, but if added to the natural growth of web page size it nearly *doubles* the CAGR to 37%.

The hidden traffic of Google Chrome is alarming. For example, just launching a Chrome browser generates 36K Bytes of traffic in the first 60 seconds.

Google may argue that the improved Chrome user experience justifies the various hidden traffic sources that accumulate over a day. Perhaps, but the consumer is not told of the cost, and bandwidth is not free.

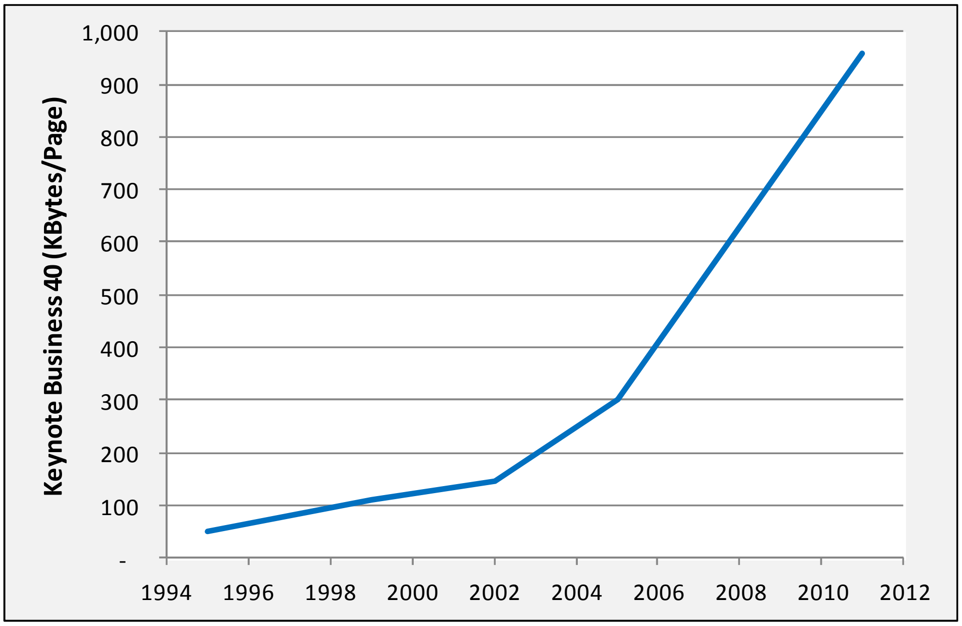

Figure 2. Payload Growth from 1995 to 2011

AppTurn count (number of network round trips) has also dramatically risen (see Figure 3) from 20 in 1995, to 124 today, a compound annual growth rate increase of 12%. Given that the speed of light is a fundamental barrier, the more turns in a web page or application, the greater the latency and the slower the user experience.

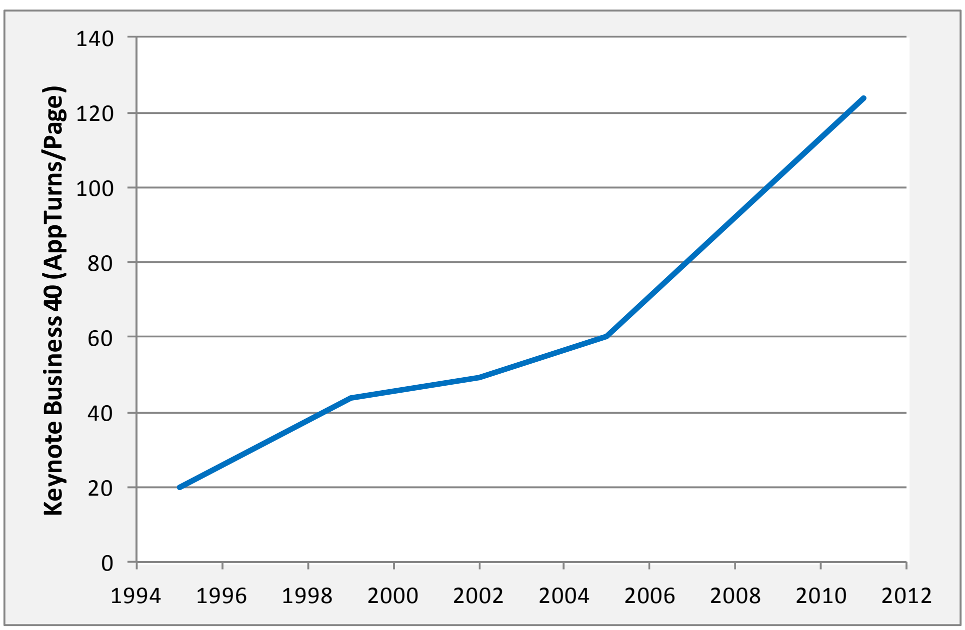

Figure 3. AppTurn Count Growth from 1995 to 2011

Do Consumers Get the Additional Bandwidth Capacity Paid For?

Given that the performance discussion should be about bandwidth capacity, for which demand by all applications is increasing, NetForecast examined the FCC data to determine if consumers in aggregate across all the ISPs measured are receiving the bandwidth they pay for. We found that the higher the bandwidth purchased, the better the chance that the bandwidth delivered meets or exceeds advertised levels (see Figure 4).

The mean bandwidth values for each service tier from all the ISPs in the FCC study vary greatly. Figure 4 consolidates the values by showing the 95th percentile confidence range of two pertinent measurement periods: all the hours in the month, and the peak Internet usage period of 7 to 11 PM local time Monday-Friday. The resulting confidence interval trend lines (lower and upper) for each ISP service tier are highlighted in Figure 4.

Lower bandwidth tiers on the left of Figure 4 receive on average less than the advertised bandwidth; however, the opposite is true for higher service tiers. Consumers who purchase more than 12 Mbps service are highly likely to experience bandwidth levels that meet or exceed advertised levels. The performance delivered falls for all service tiers during the peak usage period on weekdays between 7PM and 11PM.

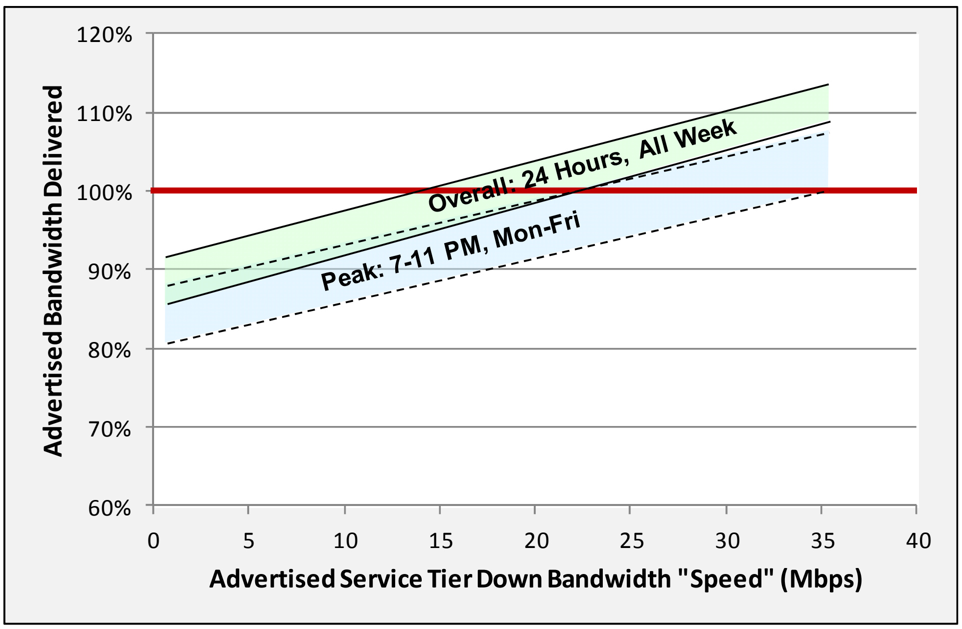

Figure 4. Advertised vs. Delivered Bandwidth

The takeaway is that, at lower service tiers, consumers generally receive less bandwidth capacity than advertised, while at higher tiers consumers tend to receive more than the advertised bandwidth capacity, except during peak usage periods.

Another takeaway is that the FCC report documents a favorable picture of broadband delivery in the United States. The overall confidence interval ranges are centered on the 100% line. This is in stark contrast to the very poor ISP results shown in UK Ofcom studies that used the same SamKnows measurement methodology. Broadband services in the US are significantly better than in the UK. This calls into question various rankings that place US Internet service as vastly inferior to the rest of the world.

The FCC measurement of ISP performance is a critical first step in understanding Internet resource economics. We expect that the FCC's ongoing scrutiny will improve performance and the value received by consumers from ISPs over time.

Time to Address Content Provider-Generated Resource Demand

The FCC *Measuring Broadband America* report only covers half of the performance equation: the supply half (see Figure 5). It is time for the FCC to investigate the demand half of the equation: the prodigious growth in resource consumption by major websites, browsers, and behind-the-scenes services.

"Clandestine demand" for bandwidth, combined with the proliferation of digital devices, means that every US household is continually in need of more bandwidth. The typical consumer is not only unaware of the hidden demand, but also lacks the means to control it, mistaking free or very cheap applications for a free lunch. They will pay for those "lunches" eventually in the form of more expensive broadband service. Meanwhile, purveyors of devices, applications, and content count on the consumer's willingness to pay in order to succeed at invisible activities with high bandwidth demands, such as the traffic required to determine which ads will appear on users' screens at any moment. The latest high-demand trend among content providers is to pre-load content that the user may or may not want to see. NetForecast believes that this economic model will lead to an unsustainable imbalance between Internet resource supply and demand.

The FCC should implement an Energy Star equivalent for bandwidth consumption by Internet applications such as Google and Facebook. This is a good idea whose time has come.

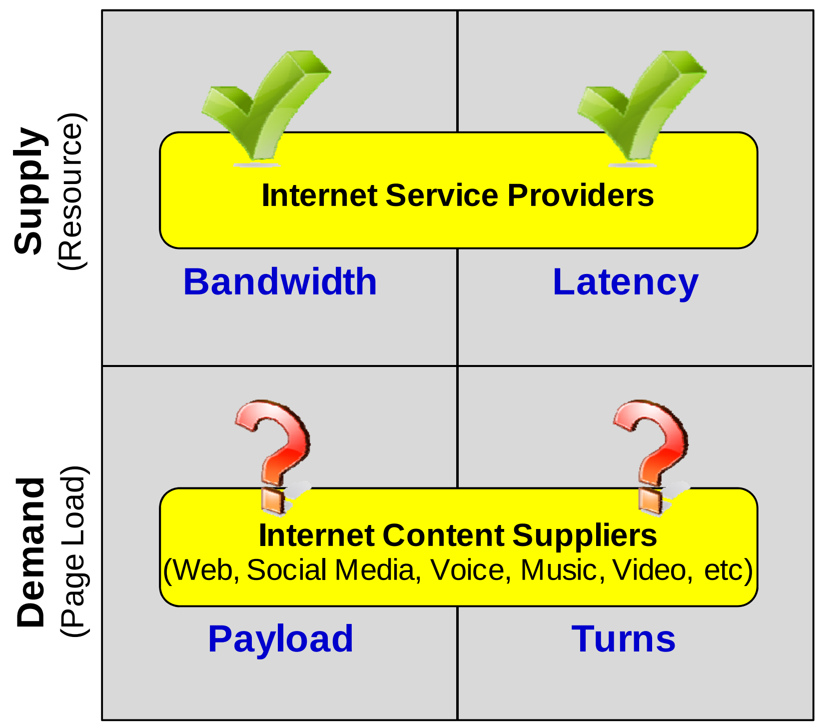

Figure 5. The FCC Broadband Report Covers only Half the Equation

A Proposal

The FCC should undertake a new initiative that follows the consensus model it used to measure ISP performance. The FCC can create a new collaborative forum of technologists, content providers, device makers, ISPs and academics to develop a plan to document digital demand. The following project elements should be considered.

- Profile major application-use scenarios in a lab environment. There are many tools that can be used to document what actually transpires on the wire.
- Measure aggregate traffic demand at a number of households. This effort could leverage the existing SamKnows network of white boxes.
- Have major content providers contribute data that quantifies usage needs. The content providers will also be key to vetting the results from the other elements.

This demand-side project will require consensus on metrics, measurement techniques, analysis, and visualization of results. However, there will also be policy issues to examine such as economic impact, fostering innovation, and consumer education.

The FCC's broadband measurement program is evidence that the FCC can successfully lead a similar initiative to fill in the missing half of the economic picture. The FCC made a step in the right direction when it recently hosted the OECD Broadband Metrics Workshop that covered supply-side topics in the context of the SamKnows measurement system currently being implemented in more than two dozen countries.

The Internet is an economic ecosystem. It is impossible to promote its health by focusing solely on either the supply or the demand sides of the system. Responsible stewardship requires a balanced approach that includes both sides of the equation.

References

[1] *Measuring Broadband America* presents the results of the first nationwide performance study of residential wireline broadband service in the United States. The study examined service offerings from 13 of the largest wireline broadband providers using automated, direct measurements of broadband performance delivered to the homes of thousands of volunteers during March 2011. The report and supporting data can be accessed at:

http://www.fcc.gov/measuring-broadband-america

About the Author

Peter Sevcik is President of NetForecast and is a leading authority on Internet traffic, performance, and technology. Peter has contributed to the design of more than 100 networks, including the Internet, and holds the patent on application response-time prediction. He can be reached at peter@netforecast.com.

NetForecast helps enterprises understand and improve the performance of networked applications.

Additional information is available at: www.netforecast.com

NetForecast and the curve-on-grid logo are registered trademarks of NetForecast, Inc.