APPENDIX D

APPX. D-1: NINTH MEASURING BROADBAND AMERICA REPORT AND TECHNICAL APPENDIX

APPX. D-2: TENTH MEASURING BROADBAND AMERICA REPORT AND TECHNICAL APPENDIX

Ninth

Measuring Broadband America

Fixed Broadband Report

A Report on Consumer Fixed Broadband Performance

in the United States

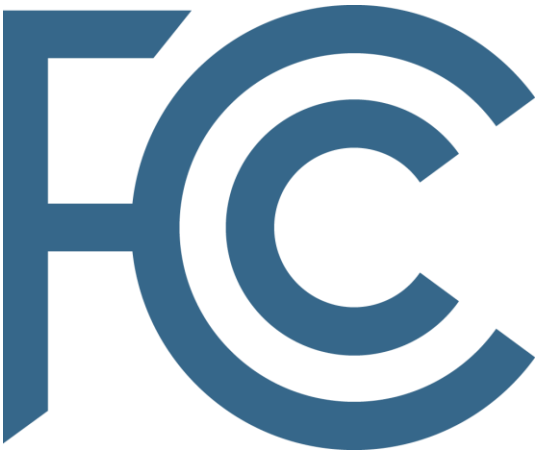

Federal Communications Commission

Office of Engineering and Technology

TABLE OF CONTENTS

List of Charts

List of Tables

1. Executive Summary

The Ninth Measuring Broadband America Fixed Broadband Report ("Ninth Report" or "Report") contains validated data collected in September and October 2018[1] from fixed Internet Service Providers (ISPs) as part of the Federal Communication Commission's (FCC) Measuring Broadband America (MBA) program. This program is an ongoing, rigorous, nationwide study of consumer broadband performance in the United States. The goal of this program is to measure the network performance delivered on selected service tiers to a representative sample set of the population. Thousands of volunteer panelists are drawn from subscribers of Internet Service Providers serving over 80% of the residential marketplace[2].

The initial Measuring Broadband America Fixed Broadband Report was published in August 2011,[3] and presented the first broad-scale study of directly measured consumer broadband performance throughout the United States. As part of an open data program, all methodologies used in the program are fully documented, and all data collected is published for public use without any restrictions. Including this current Report, nine reports have now been issued.[4] These reports provide a snapshot of fixed broadband Internet access service performance in the United States. These reports present analysis of broadband information in a variety of ways and have evolved to make the information more understandable and useful, as well as, to reflect the evolving applications supported by the nation's broadband infrastructure.

C. MAJOR FINDINGS OF THE NINTH REPORT

The key findings of this report are:

- The maximum advertised download speeds amongst the service tiers offered by ISPs and measured by the FCC ranged from 24 Mbps to 1 Gbps for the period covered by this report.
- The weighted average advertised speed of the participating ISPs was 123.3 Mbps, representing a 96% increase from the previous year.
- For most of the major broadband providers that were tested, measured download speeds were 100% or better than advertised speeds during the peak hours (7 p.m. to 11 p.m. local time).

---

[1] The actual dates used for measurements for this Ninth Report were September 25 – October 25, 2018 (inclusive).

[2] At the request of and with the assistance of the State of Hawaii Department of Commerce and Consumer Affairs (DCCA) the state of Hawaii was added to the MBA program in 2017. The ISPs whose performance were measured in the State of Hawaii were Hawaiian Telcom and Oceanic Time Warner Cable (which is now a part of Charter Spectrum).

[3] All reports can be found at https://www.fcc.gov/general/measuring-broadband-america.

[4] The First Report (2011) was based on measurements taken in March 2011, the Second Report (2012) on measurements taken in April 2012, and the Third (2013) through Eighth (2018) Reports on measurements taken in September of the year prior to the reports' release dates. In order to avoid confusion between the date of release of the report and the measurement dates we have shifted last year to numbering the reports. Thus, this year's report is termed the Ninth MBA Report instead of the 2019 MBA Report. Going forward we will continue with a numbered approach and the next report will be termed as the Tenth Report.

- Eleven ISPs were evaluated in this report. Of these AT&T, Cincinnati Bell, Frontier and Verizon employed multiple different broadband technologies across the USA. Overall 14 different ISP/technology configurations were evaluated in this report and ten performed at or better than their advertised speed and only one performed below 90% for actual-to-advertised download speed.
- In addition to providing download and upload speed measurements of ISPs, this report also provides a measure of consistency of advertised speeds of ISPs with the use of our "80/80" metric. The 80/80 metric measures the percentage of the advertised speed that at least 80% of subscribers experience at least 80% of the time over peak periods. Ten of the 14 ISP/technologies configurations provide better than 70% of advertised speed to at least 80% of panelists for at least 80% of the time.

These and other findings are described in greater detail within this report.

D. SPEED PERFORMANCE METRICS

Speed (both download and upload) performance continues to be one of the key metrics reported by the MBA. The data presented includes ISP broadband performance as a median[5] of speeds experienced by panelists within a specific service tier. These reports mainly focus on common service tiers used by an ISP's subscribers.[6]

Additionally, consistent with previous Reports, we also compute ISP performance by weighting the median speed for each service tier by the number of subscribers in that tier. Similarly, in calculating the overall average speed of all ISPs in a specific year, the median speed of each ISP is used and weighted by the number of subscribers of that ISP as a fraction of the total number of subscribers across all ISPs.

In calculating these weighted medians, we have drawn on two sources for determining the number of subscribers per service tier. ISPs may voluntarily contribute their data per surveyed service tier as the most recent and authoritative data. Many ISPs have chosen to do so.[7] When such information has not been provided by an ISP, we instead rely on the FCC's Form 477 data.[8] All facilities-based broadband providers are required to file data with the FCC twice a year (Form 477) regarding deployment of

[5] We first determine the mean value over all the measurements for each individual panelist's "whitebox." (Panelists are sent "whiteboxes" that run pre-installed software on off-the-shelf routers that measure thirteen broadband performance metrics, including download speed, upload speed, and latency.) Then for each ISP's speed tiers, we choose the median of the set of mean values for all the panelists/whiteboxes. The median is that value separating the top half of values in a sample set with the lower half of values in that set; it can be thought of as the middle (*i.e.*, most typical) value in an ordered list of values. For calculations involving multiple speed tiers, we compute the weighted average of the medians for each tier. The weightings are based on the relative subscriber numbers for the individual tiers.

[6] Only tiers that contribute to the top 80% of an ISPs total subscribership are included in this report.

[7] The ISPs that provided SamKnows, the FCC's contractor supporting the MBA program, with weights for each of their tiers were: Cincinnati Bell, CenturyLink, Charter, Comcast, Cox Frontier, Hawaiian Telcom, Optimum, and Verizon.

[8] For an explanation of Form 477 filing requirements and required data see: https://transition.fcc.gov/form477/477inst.pdf (Last accessed 5/2/2018).

broadband services, including subscriber counts. For this report, we used the June 2018 Form 477 data. It should be noted that the Form 477 subscriber data values are for a month that generally lags the reporting month, and therefore, there are likely to be small inaccuracies in the tier ratios. It is for this reason that we encourage ISPs to provide us with subscriber numbers for the measurement month.

As in our previous reports, we found that for most ISPs the actual speeds experienced by subscribers either nearly met or exceeded advertised service tier speeds. However, since we started our MBA program, consumers have changed their Internet usage habits. In 2011, consumers mainly browsed the web and downloaded files; thus, we reported average broadband speeds since these average speeds were likely to closely mirror user satisfaction. By contrast, in September-October 2018 (the measurement period for this report) consumer internet usage had become dominated by video consumption, with consumers regularly streaming video for entertainment and education.[9] Both the median measured speed and consistency in service are likely to influence the perception and usefulness of Internet access service. Therefore, our network performance analytics have been expanded to better capture this.

Specifically, we use two kinds of metrics to reflect the consistency of service delivered to the consumer: First, we report the percentage of advertised speed experienced by at least 80% of panelists during at least 80% of the daily peak usage period ("80/80 consistent speed" measure). Second, we show the fraction of consumers who obtain median speeds greater than 95%, between 80% and 95%, and less than 80% of advertised speeds.

E. USE OF OTHER PERFORMANCE METRICS

Although download and upload speeds remain the network performance metric of greatest interest to the consumer, we also spotlight two other key network performance metrics in this report: latency and packet loss. These metrics can significantly affect the overall quality of Internet applications.

Latency is the time it takes for a data packet to travel across a network from one point on the network to another. High latencies may affect the perceived quality of some interactive services such as phone calls over the Internet, video chat and video conferencing, or online multiplayer games. All network access technologies have a minimum latency that is largely determined by the technology. In addition, network congestion will lead to an increase in measured latency. Technology-dependent latencies are typically small for terrestrial broadband services and are thus unlikely to affect the perceived quality of applications. Additionally, for certain applications the user experience is not necessarily affected by high latencies. As an example, when using entertainment video streaming applications, because the data can be cached prior to display, the user experience is likely to be unaffected by relatively high latencies.

Packet loss measures the fraction of data packets sent that fail to be delivered to the intended destination. Packet loss may affect the perceived quality of applications that do not request retransmission of lost packets, such as phone calls over the Internet, video chat, some online multiplayer games, and some video streaming. High packet loss also degrades the achievable throughput of download and streaming applications. However, packet loss of a few tenths of a percent are unlikely to significantly affect the

[9] The sum of all forms of IP video, which includes Internet video, IP video-on-demand (VoD), video files exchanged through file sharing, video-streamed gaming, and video conferencing, will continue to be in the range of 80 to 90 percent of total IP traffic. Globally, IP video traffic will account for 82 percent of traffic by 2022. *See Cisco Visual Networking Index: Forecast and Methodology, 2017-2022 White Paper*, https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html (Last accessed Dec. 12, 2019).

perceived quality of most Internet applications and are common. During network congestion, both latency and packet loss typically increase.

The Internet is continuing to evolve in its architectures, performances, and services. Accordingly, we will continue to adapt our measurement and analysis methodologies to help consumers understand the performance characteristics of their broadband Internet access service, and thus make informed choices about their use of such services.

2. Summary of Key Findings

A. MOST POPULAR ADVERTISED SERVICE TIERS

A list of the ISP download and upload speed service tiers that were measured in this report are shown in Table 1. It should be noted that while upload and downloads speeds are measured independently and shown separately, they are typically offered by an ISP in a paired configuration. Together, these plans serve the majority of Internet users of the participating ISPs. The service tiers that are included for reporting represent the top 80% of an ISP's set of tiers based on subscriber numbers.

Table 1: List of ISP service tiers whose broadband performance was measured in this report

Tech-nology	Company	Speed Tiers (Download)									Speed Tiers (Upload)									
DSL	AT&T IPBB	6	12*	18	24*	25*	45*	50*	100*	1000*	0.768*	1	1.5	3	5*	6*	10*	50*	100*	1000*
	CenturyLink	1.5	3	7	10	12	20	25	40		0.512*	0.768	0.896	2	5					
	Cincinnati Bell DSL	5	30								0.768	3								
	Frontier DSL	3	6	12	18	24*					0.768	1	1.5*							
	Hawaiian Telcom DSL	7*	11*	21*	50*	100*	300*	500*			1*	3*	50*	300*	500					
	Verizon DSL	(1.1-3)									(0.384 - 0.768)									
	Windstream	1.5*	3	6*	10	12	25				0.384	0.768	1.5							
Cable	Altice Optimum	60*	100	200							25	35								
	Charter	60	100	200							5	10	20							
	Comcast	60	100*	150	250	400*					5	10								
	Cox	30	100	150	300						3	10	30							
	Mediacom	60	100								5	10								
Fiber	Cincinnati Bell Fiber	50	250								10	100								
	Frontier Fiber	50	75	100	150						50	75	100	150						
	Hawaiian Telecom Fiber	500*									300*									
	Verizon Fiber	50	75	100	940**						50	75	100	880**						

*Tiers that lack sufficient panelists to meet the program's target sample size.

** Although Verizon Fiber's 940/880 Mbps service tier was amongst the top 80% of Verizon's offered tiers by subscription numbers, it is not included in the report charts because technical procedures for measuring high speed rates near Gigabit and above have not yet been established for the MBA program.

Chart 1 (below) displays the weighted (by subscriber numbers) mean of the top 80% advertised download speed tiers for each participating ISP for September-October 2018 as well as September 2017, grouped by the access technology used to offer the broadband Internet access service (DSL, cable or fiber). In September-October 2018, the weighted average advertised download speed was 123.3 Mbps among the measured ISPs, which represents an 96% increase compared to the average in September 2017 which was 62.9Mbps.

Chart 1: Weighted average advertised download speed among the top 80% service tiers offered by each ISP

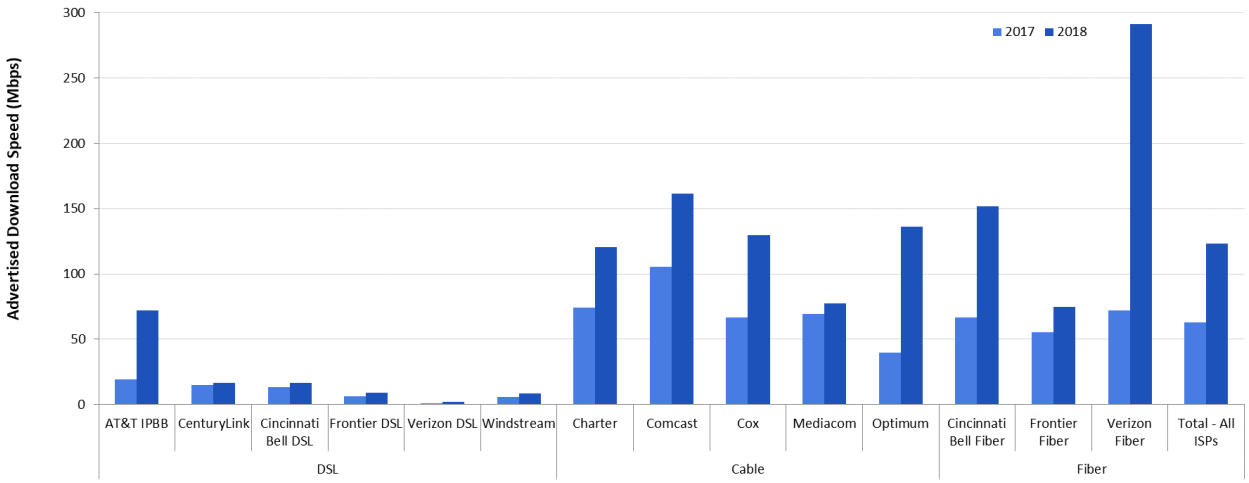

Among participating broadband ISPs, only AT&T IPBB[10], Cincinnati Bell, Hawaiian Telecom fiber, Frontier, and Verizon use fiber as the access technology for a substantial number of their customers and their maximum speed offerings range from 150 Mbps to 1 Gbps. A key difference between the fiber vendors and other technology vendors is that (with the exception of Cincinnati Bell), most fiber vendors advertise generally symmetric upload and download speeds. This is in sharp contrast to the asymmetric offerings for all the other technologies, where the upload advertised speeds are typically 5 to 10 times below the download advertised speeds.

It should be noted that there is also considerable difference between the offered average weighted speed tier by technology. Chart 2 plots the weighted average of the top 80% ISP tiers by technology both for September 2017 as well as September-October 2018. As can be seen in this chart, all technologies showed increases in the set of advertised download speeds by ISPs. For the September-October 2018 period, the weighted mean advertised speeds for DSL technology was 50 Mbps which lagged considerably behind the weighted mean advertised download speeds for cable and fiber technologies, which were 139 Mbps and 251 Mbps respectively. Fiber technology showed the greatest increase in speed offerings in 2018 compared to 2017 with a weighted mean going up from 70 Mbps to 251 Mbps representing a 258% increase. In comparison, DSL and cable technologies showed 96%, and 64% increase from 2017 to 2018.

[10] Although AT&T IPBB has been characterized here as a DSL technology it actually includes a mix of ADSL2+, VDSL2, G.Fast and Ethernet technologies delivered over a hybrid of fiber optic and copper facilities.

Chart 2: Weighted average advertised download speed among the top 80% service tiers based on technology.

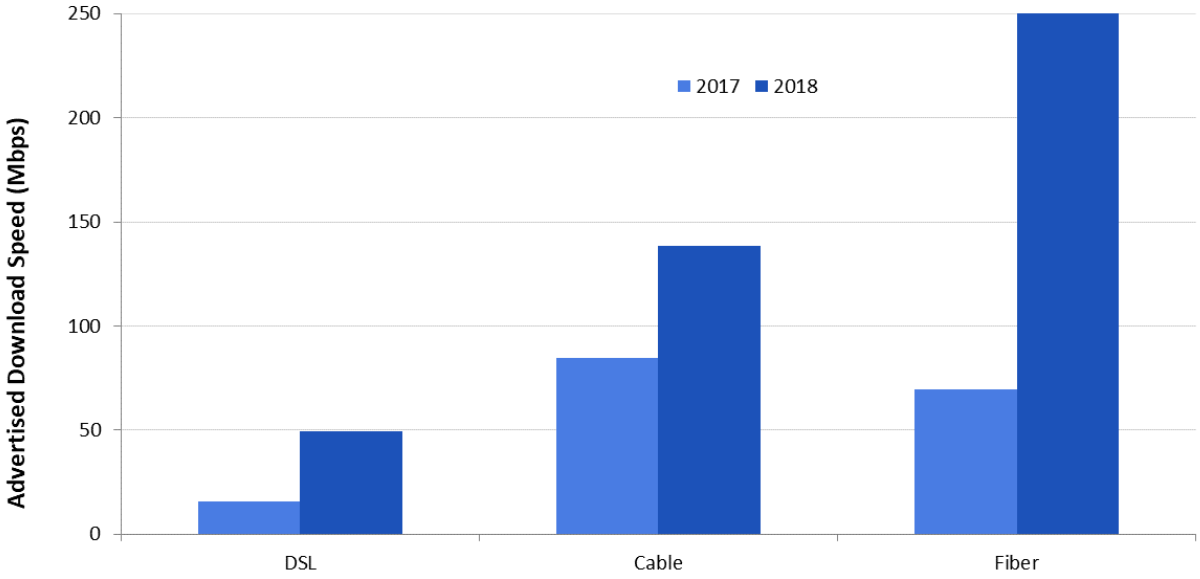

Chart 3 plots the migration of panelists to a higher service tier based on their access technology.[11] Specifically, the horizontal axis of Chart 3 partitions the September 2017 panelists by the advertised download speed of the service tier to which they were subscribed. For each such set of panelists who also participated in the September-October 2018 collection of data,[12] the vertical axis of Chart 3 displays the percentage of panelists that migrated by September-October 2018 to a service tier with a higher advertised download speed. There are two ways that such a migration could occur: (1) if a panelist changed their broadband plan during the intervening year to a service tier with a higher advertised download speed, or (2) if a panelist did not change their broadband plan but the panelist's ISP increased the advertised download speed of the panelist's subscribed plan.[13]

Chart 3 shows that the percentage of panelists subscribed in September 2017 who moved to higher tiers in September-October 2018 was between 3% to 67% for DSL subscribers, 22% to 100% for cable

[11] Where several technologies are plotted at the same point in the chart, this is identified as "Multiple Technologies."

[12] Of the 4,545 panelists who participated in the September 2017 collection of data, 4,355 panelists continued to participate in the September-October 2018 collection of data.

[13] We do not attempt here to distinguish between these two cases.

subscribers and 8% to 80% for fiber subscribers. In addition, 1% to 13% subscribers migrated to a higher speed tier using a different technology from what they had in September 2017.

Chart 3: Consumer migration to higher advertised download speeds

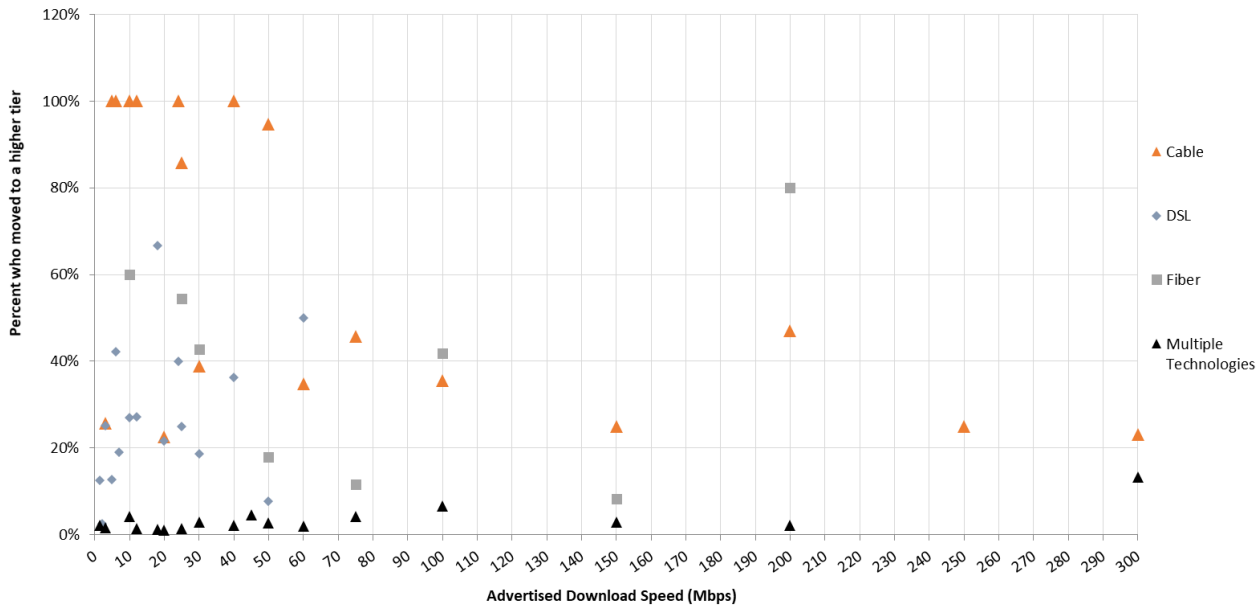

B. MEDIAN DOWNLOAD SPEEDS

Advertised download speeds may differ from the speeds that subscribers actually experience. Some ISPs more consistently meet network service objectives than others or meet them unevenly across their geographic coverage area. Also, speeds experienced by a consumer may vary during the day if the network cannot carry the aggregate user demand during busy hours. Unless stated otherwise, all actual speeds were measured only during peak usage periods, which we define as 7 p.m. to 11 p.m. local time.

To compute the average ISP performance, we determine the ratio of the median speed for each tier to the advertised tier speed and then calculate the weighted average of these based on the subscriber count per tier. Subscriber counts for the weightings were provided from the ISPs themselves or, if unavailable, from FCC Form 477 data.

Chart 4 shows the ratio of the median download and upload speeds experienced by an ISP's subscribers to that ISP's advertised download and upload speeds weighted by the subscribership numbers for the tiers. The actual speeds experienced by most ISPs' subscribers are close to or exceed the advertised speeds. However, DSL broadband ISPs continue to advertise "up-to" speeds that on average exceed the actual speeds experienced by their subscribers. Verizon, instead, advertises a speed range for DSL performance and has requested that we include this range in relevant charts; we indicate this speed range by shading on all bar charts describing Verizon's DSL performance. Out of the 14 ISP/technology configurations shown, 10 met or exceeded their advertised download speed and three more reached at least 90% of their advertised download speed. Only Cincinnati-DSL (at 81%) performed below 90% of its advertised download speed.

Chart 4: The ratio of weighted median speed (download and upload) to advertised speed for each ISP. Note Verizon advertises a speed range for both its download and upload DSL tier and hence appears as a range in this and other charts.

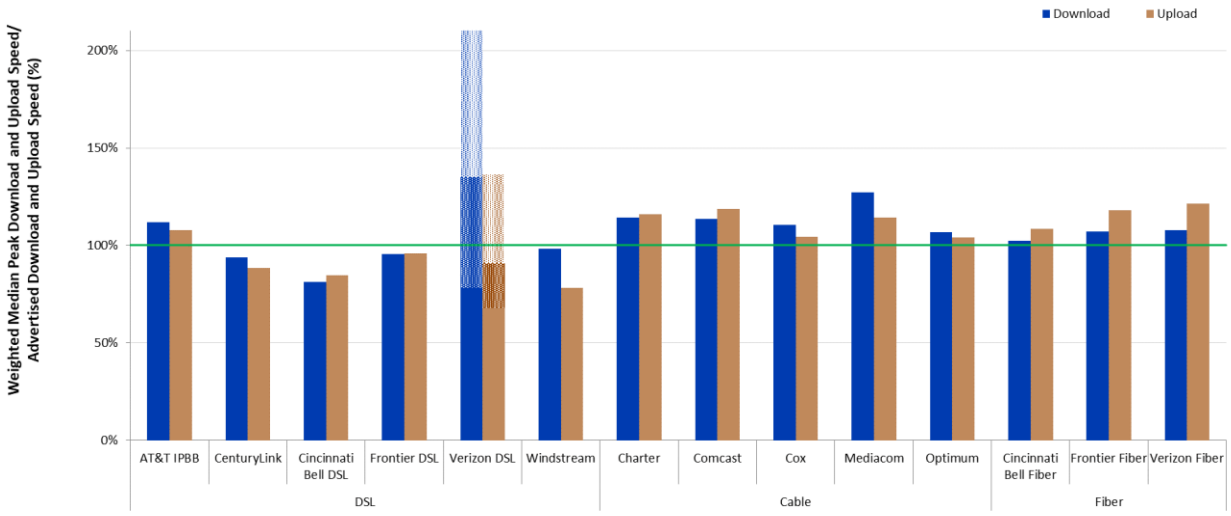

C. VARIATIONS IN SPEEDS

As discussed earlier, actual speeds experienced by individual consumers may vary by location and time of day. Chart 5 shows, for each ISP, the percentage of panelists who experienced a median download speed (averaged over the peak usage period during our measurement period) that was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed.

Chart 5: The percentage of consumers whose median download speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed

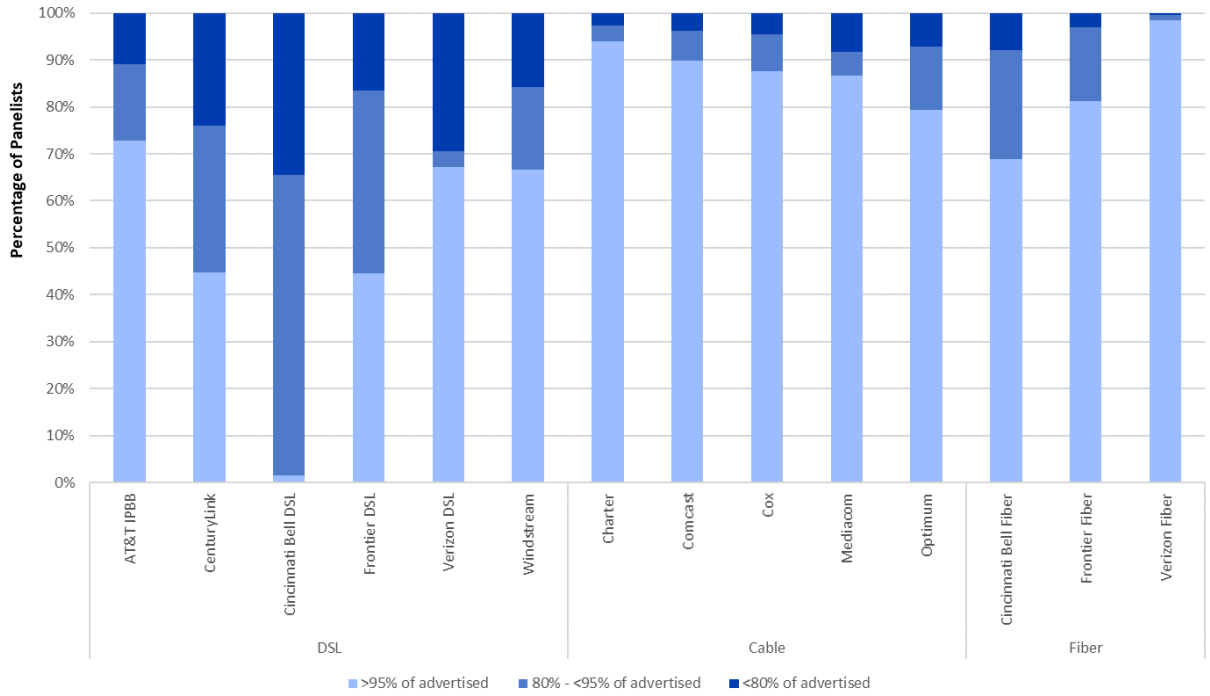

ISPs using DSL technology had between 2% to 73% of their subscribers getting greater than or equal to 95% of their advertised download speeds during peak hours. ISPs using cable technology and fiber technology had between 79%-94% and between 69%-98%, respectively, of their subscribers getting equal to or better than 95% of their advertised download speeds.

Though the median download speeds experienced by most ISPs' subscribers nearly met or exceeded the advertised download speeds, there are some customers of each ISP for whom the median download speed fell significantly short of the advertised download speed. Relatively few subscribers of cable or fiber broadband service experienced this. The best performing ISPs, when measured by this metric, are Charter, Comcast, Cox, Mediacom, Frontier-Fiber and Verizon-Fiber; more than 80% of their panelists were able to attain an actual median download speed of at least 95% of the advertised download speed.

In addition to variations based on a subscriber's location, speeds experienced by a consumer may fluctuate during the day. This is typically caused by increased traffic demand and the resulting stress on different parts of the network infrastructure. To examine this aspect of performance, we use the term "80/80 consistent speed." This metric is designed to assess temporal and spatial variations in measured values of a user's download speed.[14] While consistency of speed is in itself an intrinsically valuable service characteristic, its impact on consumers will hinge on variations in usage patterns and needs. As an example, a good consistency of speed measure is likely to indicate a higher quality of service experience for internet users consuming video content.

Chart 6 summarizes, for each ISP, the ratio of 80/80 consistent median download speed to advertised download speed, and, for comparison, the ratio of median download speed to advertised download speed

[14] For a detailed definition and discussion of this metric, please refer to the Technical Appendix.

shown previously in Chart 4. The ratio of 80/80 consistent median download speed to advertised download speed is less than the ratio of median download speed to advertised download speed for all participating ISPs due to congestion periods when median download speeds are lower than the overall average. When the difference between the two ratios is small, the median download speed is fairly insensitive to both geography and time. When the difference between the two ratios is large, there is a greater variability in median download speed, either across a set of different locations or across different times during the peak usage period at the same location.

Chart 6: The ratio of 80/80 consistent median download speed to advertised download speed.

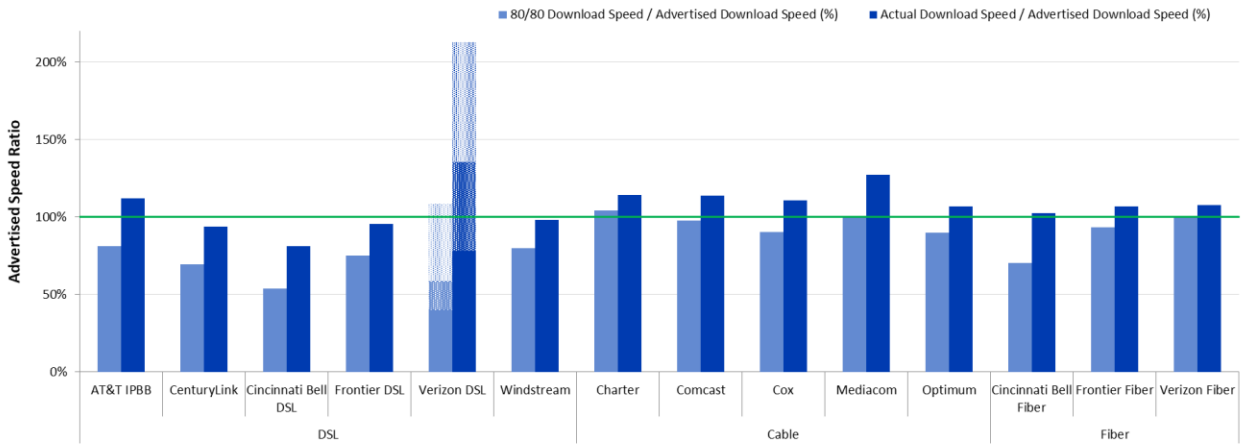

Customers of Charter, Comcast, Cox, Mediacom, Optimum, Frontier Fiber and Verizon Fiber (FiOS) experienced median download speeds that were very consistent; *i.e.*, they provided greater than 90% of the advertised speed during peak usage period to more than 80% of panelists for more than 80% of the time. As can be seen in Chart 6, except for AT&T-IPBB, cable and fiber ISPs performed better than DSL ISPs with respect to their 80/80 consistent speeds. For example, for September-October 2018, the 80/80 consistent download speed for Cincinnati Bell DSL was 54% of the advertised speed.

D. LATENCY

Latency is the time it takes for a data packet to travel from one point to another in a network. It has a fixed component that depends on the distance, the transmission speed, and transmission technology between the source and destination, and a variable component that increases as the network path congests with traffic. The MBA program measures latency by measuring the round-trip time from the consumer's home to the closest measurement server and back.

Chart 7 shows the median latency for each participating ISP. In general, higher-speed service tiers have lower latency, as it takes less time to transmit each packet. The median latencies ranged from 9.5 ms to 36 ms in our measurements (with the exception of Verizon DSL which had a median latency of 42 ms).

Chart 7: Latency by ISP

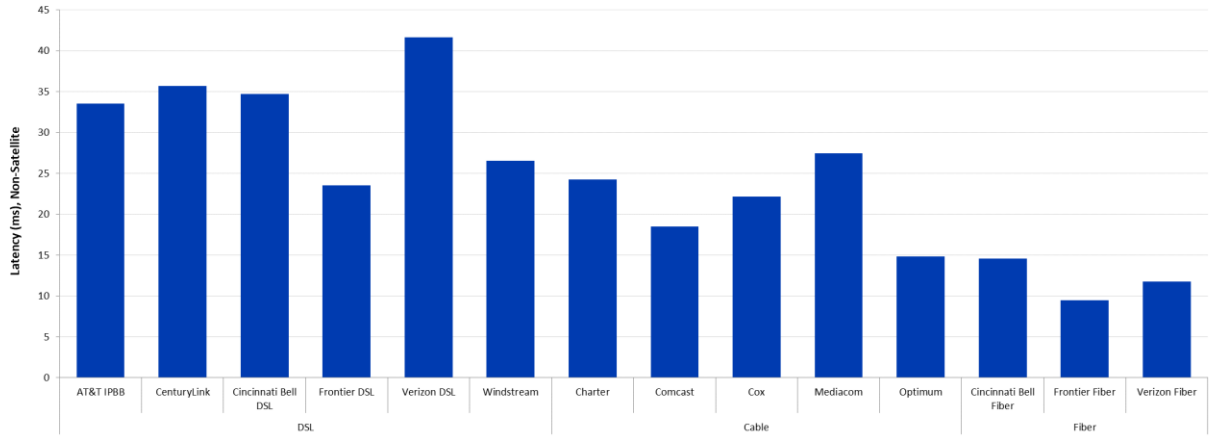

DSL latencies (between 24 ms to 42 ms) were slightly higher than those for cable (15 ms to 27 ms). Fiber ISPs showed the lowest latencies (10 ms to 15 ms). The differences in median latencies among terrestrial-based broadband services are relatively small and are unlikely to affect the perceived quality of highly interactive applications.

E. PACKET LOSS

Packet loss is the percentage of packets that are sent by a source but not received at the intended destination. The most common causes of packet loss are high latency or encountered congestion along the network route. A small amount of packet loss is expected, and indeed packet loss is commonly used by some Internet protocols to infer Internet congestion and to adjust the sending rate to mitigate for the congestion. The MBA program considers a packet lost if the packet's round-trip latency exceeds 3 seconds.

Chart 8 shows the average peak-period packet loss for each participating ISP, grouped into bins. We have broken the packet loss performance into three bands, allowing a more granular view of the packet loss performance of the ISP network. The breakpoints for the three bins used to classify packet loss have been chosen with an eye towards balancing commonly accepted packet loss standards and provider packet loss Service Level Agreements (SLAs). Specifically, the 1% standard for packet loss is commonly accepted as the point at which highly interactive applications such as VoIP experience significant degradation and quality according to international documents.[15] The 0.4% breakpoint was chosen as a generic breakpoint between the highly desired performance of 0% packet loss described in many documents and the 1% unacceptable limit on the high side. The specific value of 0.4% is based upon a compromised value between those two limits and is generally supported by many SLAs and major ISPs for network performance. Indeed, most SLAs support 0.1% to 0.3% SLA packet loss guarantees,[16] but these are generally for enterprise level services which generally have more stringent requirements for higher-level performance.

[15] See: https://www.voip-info.org/wiki/view/QoS and http://www.ciscopress.com/articles/article.asp?p=357102.

[16] See: http://www.itu.int/dms_pubrec/itu-r/rec/m/r-rec-m.1079-2-200306-i!!msw-e.doc.

Chart 8: Percentage of consumers whose peak-period packet loss was less than 0.4%, between 0.4% to 1%, and greater than 1%.

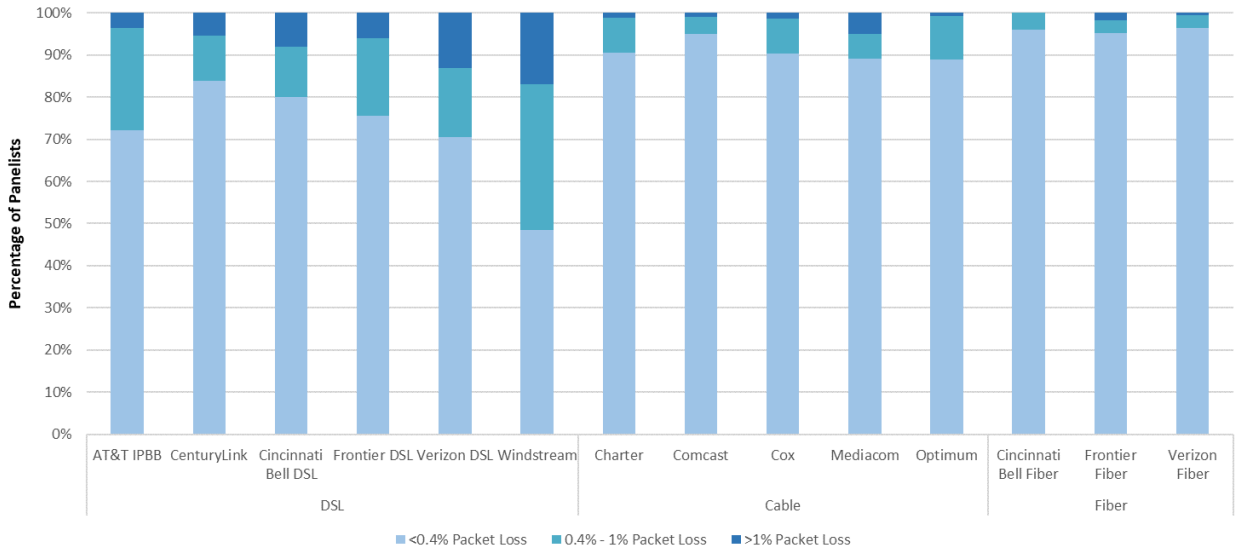

Chart 8 shows that ISPs using fiber technology have the lowest packet loss, and that ISPs using DSL technology tends to have the highest packet loss. Within a given technology class, packet loss also varies among ISPs.

F. WEB BROWSING PERFORMANCE

The MBA program also conducts a specific test to gauge web browsing performance. The web browsing test accesses nine popular websites that include text and images, but not streaming video. The time required to download a webpage depends on many factors, including the consumer's in-home network, the download speed within an ISP's network, the web server's speed, congestion in other networks outside the consumer's ISP's network (if any), and the time required to look up the network address of the webserver. Only some of these factors are under control of the consumer's ISP. Chart 9 displays the average webpage download time as a function of the advertised download speed. As shown by this chart, webpage download time decreases as download speed increases, from about 9.3 seconds at 1.5 Mbps download speed to about 1.4-1.7 seconds for 30 Mbps download speed. Subscribers to service tiers exceeding 25 Mbps experience slightly smaller webpage download times decreasing to 1.1 second at 300 Mbps. These download times assume that only a single user is using the Internet connection when the webpage is downloaded, and does not account for more common scenarios, where multiple users within a household are simultaneously using the Internet connection for viewing web pages, as well as other applications such as real-time gaming or video streaming.

Chart 9: Average webpage download time, by advertised download speed.

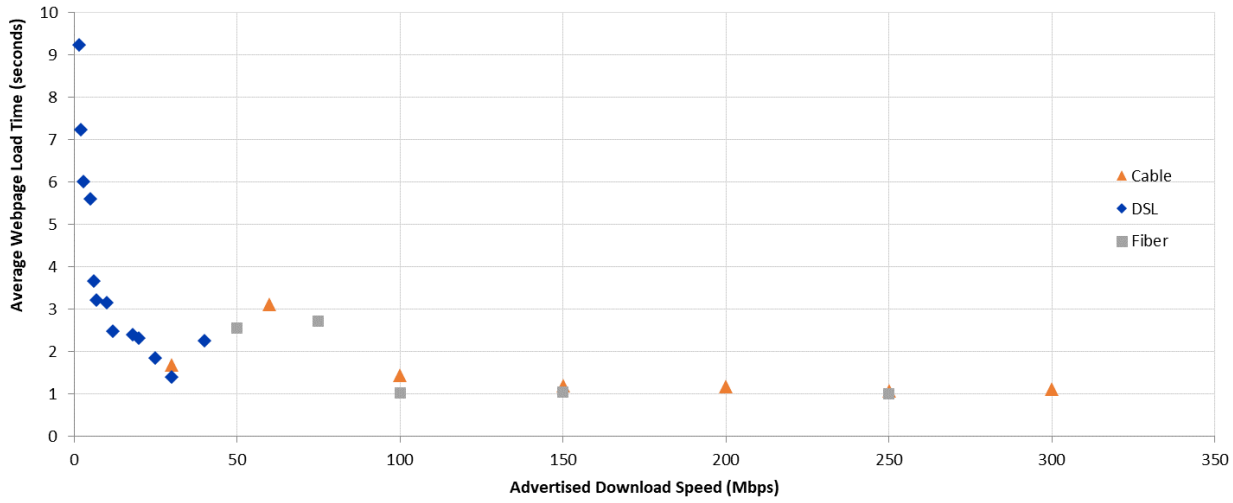

3. Methodology

A. PARTICIPANTS

Eleven ISPs participated in the Fixed MBA program in September-October 2018.[17] They were:

- CenturyLink
- Charter Communications
- Cincinnati Bell
- Comcast
- Cox Communications
- Frontier Communications Company
- Hawaiian Telcom
- Mediacom Communications Corporation
- Optimum
- Verizon
- Windstream Communications

The methodologies and assumptions underlying the measurements described in this Report are reviewed at meetings that are open to all interested parties and documented in public ex parte letters filed in the GN Docket No. 12-264. Policy decisions regarding the MBA program were discussed at these meetings prior to adoption, and involved issues such as inclusion of tiers, test periods, mitigation of operational issues affecting the measurement infrastructure, and terms-of-use notifications to panelists. Participation in the MBA program is open and voluntary. Participants include members of academia, consumer equipment vendors, telecommunications vendors, network service providers, consumer policy groups as well as our contractor for this project, SamKnows. In 2018-2019, participants at these meetings (collectively and informally referred to as "the broadband collaborative"), included all eleven participating ISPs and the following additional organizations:

- Level 3 Communications ("Level 3"), now part of CenturyLink
- Massachusetts Institute of Technology ("MIT")
- Measurement Lab (M-Lab)
- NCTA – The Internet & Television Association ("NCTA")
- New America Foundation
- Princeton University
- United States Telecom Association ("US Telecom")
- University of California - Santa Cruz

[17] Both AT&T and Hughes Network Systems left the program as participating ISPs this year, bringing the total number of participating ISPs to eleven. We continued to evaluate AT&T's sets of tiers with sufficient numbers of panelists despite the fact that AT&T did not participate this year, so the total number of ISPs evaluated in this report was twelve. As of the Eighth Report (previous year's report), Viasat, operating under the brand name Exede internet, left the program as a participating ISP the prior year and consequently no longer provides panelists with an increased data allowance to offset the data used by the MBA measurements. We, however, continue reporting raw data results for ViaSat/Exede and Hughes Network Systems tiers by using lightweight tests aimed at reducing the data burden on these panelists. These tests are described in greater detail in the accompanying Technical Appendix to this Ninth MBA Report.

Participants have contributed in important ways to the integrity of this program and have provided valuable input to FCC decisions for this program. Initial proposals for test metrics and testing platforms were discussed and critiqued within the broadband collaborative. M-Lab and Level 3 contributed their core network testing infrastructure, and both parties continue to provide invaluable assistance in helping to define and implement the FCC testing platform. We thank all the participants for their continued contributions to the MBA program.

B. MEASUREMENT PROCESS

The measurements that provided the underlying data for this report were conducted between MBA measurement clients and MBA measurement servers. The measurement clients (*i.e.*, whiteboxes) were situated in the homes of 5,855 panelists each of whom received service from one of the 12 evaluated ISPs. The evaluated ISPs collectively accounted for over 80% of U.S. residential broadband Internet connections. After the measurement data was processed (as described in greater detail in the Technical Appendix), test results from 3,192 panelists were used in this report.

The measurement servers used by the MBA program were hosted by M-Lab and Level 3 Communications, and were located in eleven cities (often with multiple locations within each city) across the United States near a point of interconnection between the ISP's network and the network on which the measurement server resided.

The measurement clients collected data throughout the year, and this data is available as described below. However, only data collected from September 25 through October 25, 2018, referred to throughout this report as the "September-October 2018" reporting period, were used to generate the charts in this Report.[18]

Broadband performance varies with the time of day. At peak hours, more people tend to use their broadband Internet connections, giving rise to a greater potential for network congestion and degraded user performance. Unless otherwise stated, this Report focuses on performance during peak usage period, which is defined as weeknights between 7:00 p.m. to 11:00 p.m. local time at the subscriber's location. Focusing on peak usage period provides the most useful information because it demonstrates what performance users can expect when the Internet in their local area experiences the highest demand from users.

Our methodology focuses on the network performance of each of the participating ISPs. The metrics discussed in this Report are derived from active measurements, *i.e.*, test-generated traffic flowing between a measurement client, located within the modem/router within a panelist's home, and a measurement server, located outside the ISP's network. For each panelist, the tests automatically choose the measurement server that has the lowest latency to the measurement client. Thus, the metrics measure performance along the path followed by the measurement traffic within each ISP's network, through a point of interconnection between the ISP's network and the network on which the chosen

[18] This proposed time period avoids the dates in early September when parts of North Carolina and Florida were affected by Hurricanes Florence and Michael. It also avoided the increased traffic resulting from latest iOS release which also took place in early September. Omitting dates during these periods was done consistent with the FCC's data collection policy for fixed MBA data. *See* FCC, Measuring Fixed Broadband, Data Collection Policy, https://www.fcc.gov/general/measuring-broadband-america-measuring-fixed-broadband (explaining that the FCC has developed policies to deal with impairments in the data collection process with potential impact for the validity of the data collected).

measurement server is located. However, the service performance that a consumer experiences could differ from our measured values for several reasons.

First, as noted, in the course of each test instance we measure performance only to a single measurement server rather than to multiple servers. This is consistent with the approach chosen by most network measurement tools. As a point of comparison, the average web page may load its content from a multiplicity of end points.

In addition, bottlenecks or congestion points in the full path traversed by consumer application traffic might also impact a consumer's perception of Internet service performance. These bottlenecks may exist at various points: within the ISP's network, beyond its network (depending on the network topology encountered *en route* to the traffic destination), in the consumer's home, on the Wi-Fi used to access the in-home access router, or from a shortfall of capacity at the far end point being accessed by the application. The MBA tests explore how a service performs from the point at which a fixed ISP's Internet service is delivered to the home on fixed infrastructure (deliberately excluding Wi-Fi, due to the many confounding factors associated with it) to the point at which the test servers are located. As MBA tests are designed to focus on the access to the ISP's network, they will not include phenomena at most interconnection points or transit networks that consumer traffic may traverse.

To the extent possible[19] the MBA focuses on performance within an ISP's network. It should be noted that the overall performance a consumer experiences with their service can also be affected by congestion such as may arise at other points in the path potentially taken by consumer traffic (*e.g.*, in-home Wi-Fi, peering points, transit networks *etc.*) but this does not get reflected in MBA measurements.

A consumer's home network, rather than the ISP's network, may be the bottleneck with respect to network congestion. We measure the performance of the ISP's service delivered to the consumer's home network, but this service is often shared simultaneously among multiple users and applications within the home. In-home networks, which typically include Wi-Fi, may not have sufficient capacities to support peak loads.[20]

In addition, consumers' experience of ISP performance is manifested through the set of applications they utilize. The overall performance of an application depends not only on the network performance (*i.e.*, raw speed, latency or packet loss) but also on the application's architecture and implementation and on the operating system and hardware on which it runs. While network performance is considered in this Report, application performance is generally not.

---

[19] The MBA program uses test servers that are both neutral (*i.e.*, operated by third parties that are not ISP-operated or owned) and located as close as practical, in terms of network topology, to the boundaries of the ISP networks under study. As described earlier in this section, a maximum of two interconnection points and one transit network may be on the test path. If there is congestion on such paths to the test server, it may impact the measurement, but the cases where it does so are detectable by the test approach followed by the MBA program, which uses consistent longitudinal measurements and comparisons with averaged results. Details of the methodology used in the MBA program are given in the Technical Appendix to this report.

[20] Independent research, drawing on the FCC's MBA test platform [numerous instances of research supported by the fixed MBA test platform are described at https://www.fcc.gov/general/mba-assisted-research-studies, suggests that home networks are a significant source of end-to-end service congestion. *See* Srikanth Sundaresan et al., *Home Network or Access Link? Locating Last-Mile Downstream Throughput Bottlenecks*, PAM 2016 - Passive and Active Measurement Conference, at 111-123, March 2016.

C. MEASUREMENT TESTS AND PERFORMANCE METRICS

This Report is based on the following measurement tests:

- Download speed: This test measures the download speed of each whitebox over a 10-second period, once per hour during peak hours (7 p.m. to 11 p.m.) and once during each of the following periods: midnight to 6 a.m., 6 a.m. to noon, and noon to 6 p.m. The download speed measurement results from each whitebox are then averaged across the measurement month; and the median value for these average speeds across the entire set of whiteboxes is used to determine the *median download speed* for a service tier. The overall ISP download speed is computed as the weighted median for each service tier, using the subscriber counts for the tiers as weights.
- Upload speed: This test measures the upload speed of each whitebox over a 10-second period, which is the same measurement interval as the download speed. The upload speed measured in the last five seconds of the 10-second interval is retained, the results of each whitebox are then averaged over the measurement period, and the median value for the average speed taken over the entire set of whiteboxes is used to determine the *median upload speed* for a service tier. The ISP upload speed is computed in the same manner as the download speed.
- Latency and packet loss: These tests measure the round-trip times for approximately 2,000 packets per hour sent at randomly distributed intervals. Response times less than three seconds are used to determine the mean latency. If the whitebox does not receive a response within three seconds, the packet is counted as lost.
- Web browsing: The web browsing test measures the total time it takes to request and receive webpages, including the text and images, from nine popular websites and is performed once every hour. The measurement includes the time required to translate the web server name (URL) into the webserver's network (IP) address.

This Report focuses on three key performance metrics of interest to consumers of broadband Internet access service, as they are likely to influence how well a wide range of consumer applications work: download and upload speed, latency, and packet loss. Download and upload speeds are also the primary network performance characteristic advertised by ISPs. However, as discussed above, the performance observed by a user in any given circumstance depends not only on the actual speed of the ISP's network, but also on the performance of other parts of the Internet and on that of the application itself.

The standard speed tests use TCP with 8 concurrent TCP sessions. In 2017 we also introduced a single TCP speed test (termed as Lightweight tests), which ran less frequently and thereby provided less strain on consumer accounts that are data-capped. The Lightweight tests are used exclusively to provide broadband performance results for satellite ISPs. The Technical Appendix to this Report describes each test in more detail, including additional tests not contained in this Report.

D. AVAILABILITY OF DATA

The Validated Data Set[21] on which this Report is based, as well as the full results of all tests, are available at http://www.fcc.gov/measuring-broadband-america. To encourage additional research, we also provide raw data for the reference month and other months. Previous reports of the MBA program, as well as the data used to produce them, are also available there.

[21] The September-October 2018 data set was validated to remove anomalies that would have produced errors in the Report. This data validation process is described in the Technical Appendix.

Both the Commission and SamKnows, the Commission's contractor for this program, recognize that, while the methodology descriptions included in this document provide an overview of the project, interested parties may be willing to contribute to the project by reviewing the software used in the testing. SamKnows welcomes review of its software and technical platform, consistent with the Commission's goals of openness and transparency for this program.[22]

[22] The software that was used for the MBA program will be made available for noncommercial purposes. To apply for noncommercial review of the code, interested parties may contact SamKnows directly at team@samknows.com, with the subject heading "Academic Code Review."

4. Test Results

A. MOST POPULAR ADVERTISED SERVICE TIERS

Chart 1 above summarizes the weighted average of the advertised download speeds[23] for each participating ISP, for September-October 2018 and September 2017 where the weighting is based upon the number of subscribers to each tier, grouped by the access technology used to offer the broadband Internet access service (DSL, cable, or fiber). Only the top 80% tiers (by subscriber number) of each ISP were included. Chart 10 below shows the corresponding weighted average of the advertised upload speeds among the measured ISPs. The computed weighted average of the advertised upload speed of all the ISPs is 27 Mbps representing a 141% increase over the previous year's value of 11 Mbps.

Chart 10: Weighted average advertised upload speed among the top 80% service tiers offered by each ISP.

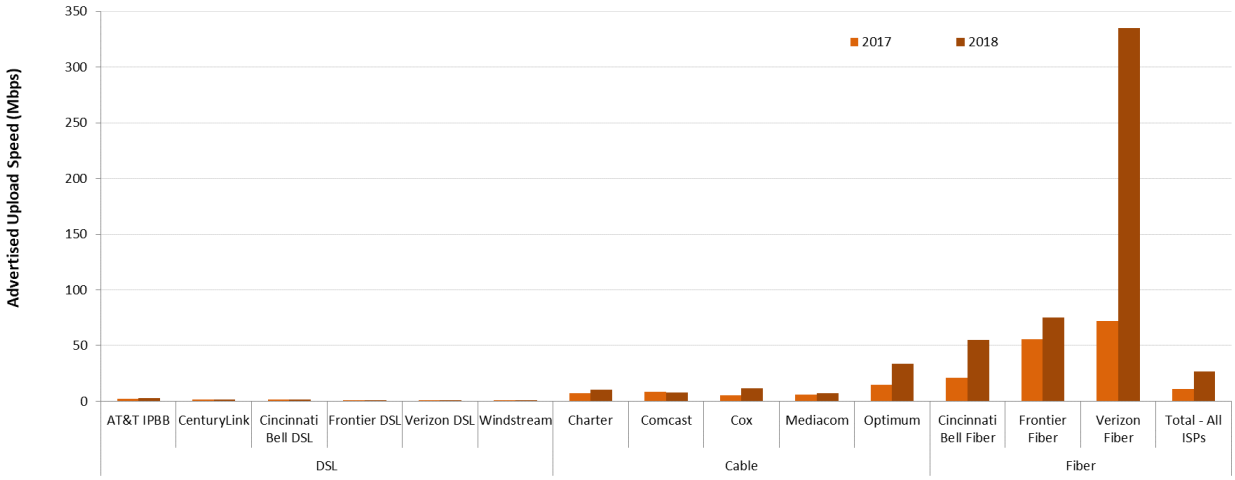

Chart 11 compares the weighted average of the advertised upload speeds by technology both for September 2017 and September-October 2018. As can be seen in this chart, all technologies showed increased rates in 2018 as compared to 2017. However, the rates of increase were not the same for all technologies. The rate of increase in the weighted average of Fiber technology was 308% compared to DSL and Cable which were 22% and 31 % respectively.

Observing both the download and upload speeds, it is clear that fiber service tiers are generally symmetric in their actual upload and download speeds. This results from the fact that fiber technology has significantly more capacity than other technologies and it can be engineered to have symmetric upload and download speeds. For other technologies with more limited capacity, higher capacity is usually allocated to download speeds than to upload speeds, typically in ratios ranging from 5:1 to 10:1. This resulting asymmetry in download/upload speeds is reflective of actual usage because consumers typically download significantly more data than they upload.

[23] Measured service tiers were tiers which constituted the top 80% of an ISP's broadband subscriber base.

Chart 11: Weighted average advertised upload speed among the top 80% service tiers based on technology.

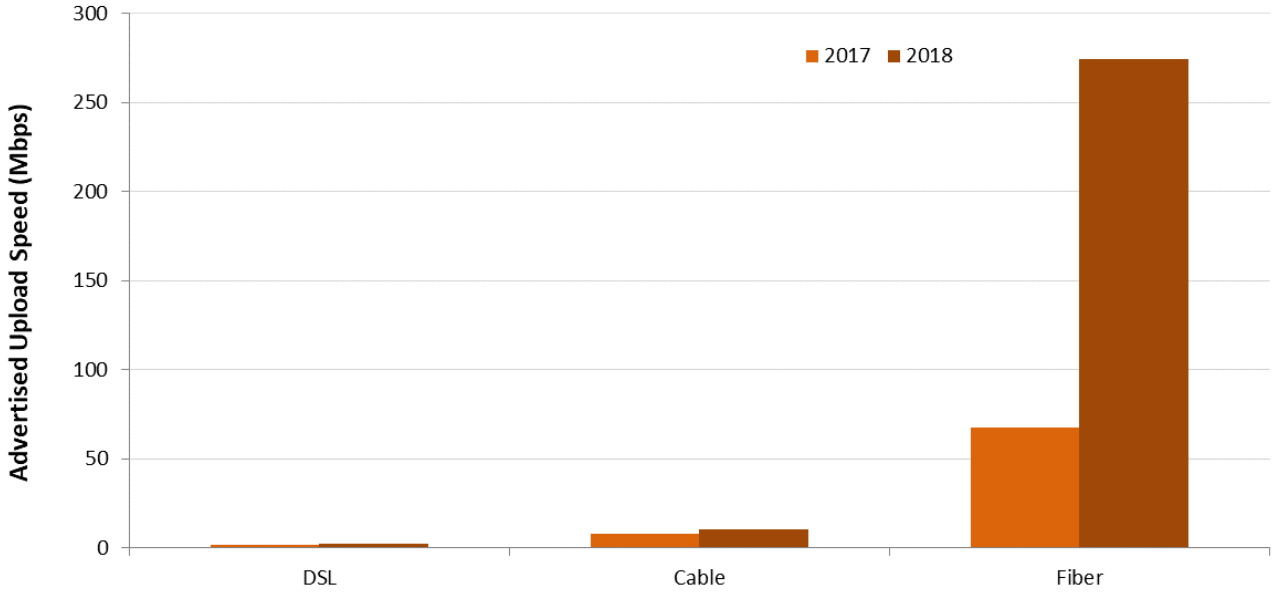

B. OBSERVED MEDIAN DOWNLOAD AND UPLOAD SPEEDS

Chart 4 (in Section 2.B) shows the ratio in September-October 2018 of the weighted median of both download and upload speeds of each ISP's subscribers to advertised speeds. Charts 12.1 and 12.2 below show the same ratios separately for download speed and for upload speed.[24] The median download speeds of most ISPs' subscribers have been close to, or have exceeded, the advertised speeds. Exceptions to this were the following DSL providers: CenturyLink, Cincinnati Bell DSL, Frontier DSL and Windstream with respective ratios of 94%, 81%, 96% and 98%.

Chart 12.1: The ratio of median download speed to advertised download speed.

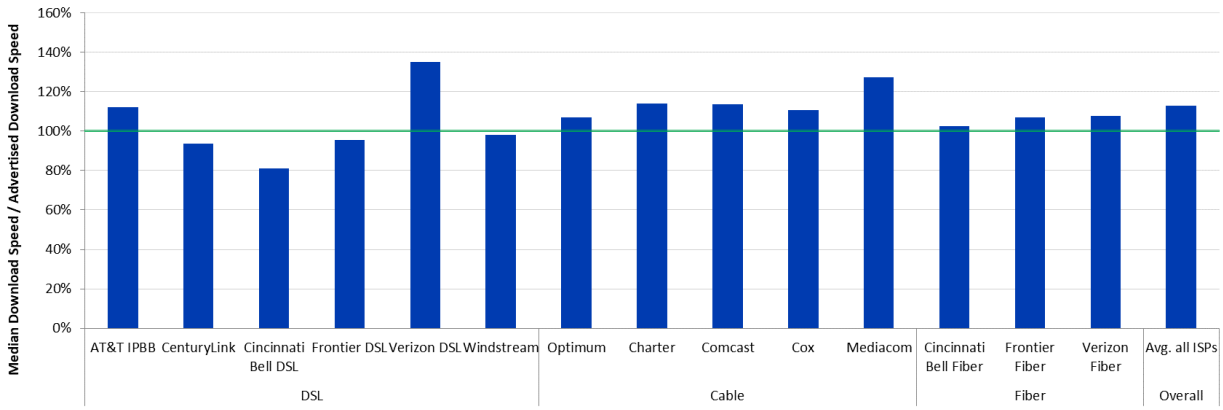

[24] In these charts, we show Verizon's median speed as a percentage of the mid-point between their lower and upper advertised speed range.

Chart 12.2 shows the median upload speed as a percentage of the advertised speed. As was the case with download speeds most ISPs met or exceeded the advertised rates except for a number of DSL providers: CenturyLink, Cincinnati Bell DSL, Frontier DSL, Verizon DSL and Windstream which had respective ratios of 88%, 85%, 96%, 91%, and 78%.

Chart 12.2: The ratio of median upload speed to advertised upload speed.

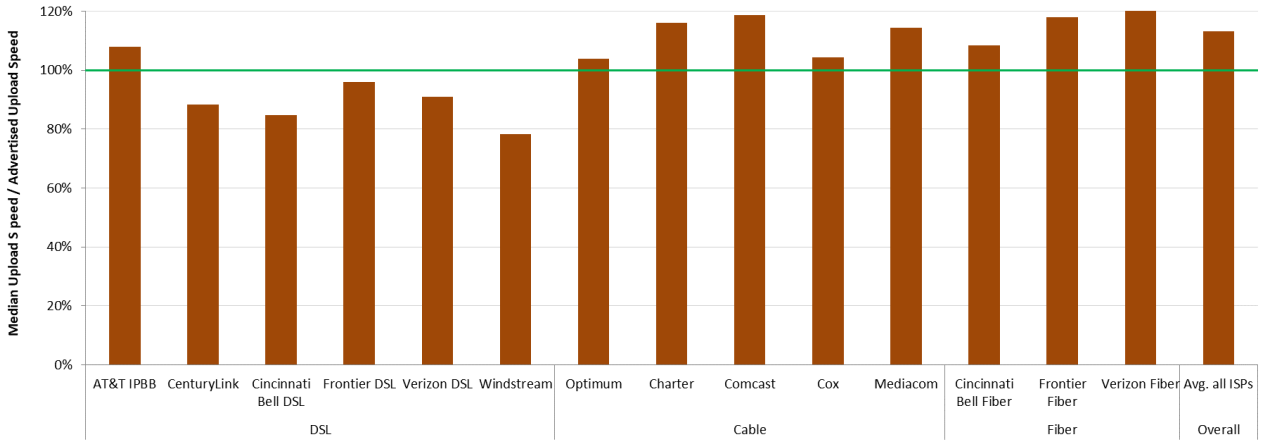

C. VARIATIONS IN SPEEDS

Median speeds experienced by consumers may vary based on location and time of day. Chart 5 above showed, for each ISP, the percentage of consumers (across the ISP's service territory) who experienced a median download speed over the peak usage period that was either greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed. Chart 13 below shows the corresponding percentage of consumers whose median upload speed fell in each of these ranges. With the exception of AT&T IPBB, ISPs using DSL technology had between 20% and 49% of their subscribers getting greater than or equal to 95% of their advertised upload speeds during peak hoursISPs using cable or fiber technology had between 90% - 99% of their subscribers getting equal to or better than 95% of their advertised upload speeds.

Chart 13: The percentage of consumers whose median upload speed was (a) greater than 95%, (b) between 80% and 95%, or (c) less than 80% of the advertised upload speed.

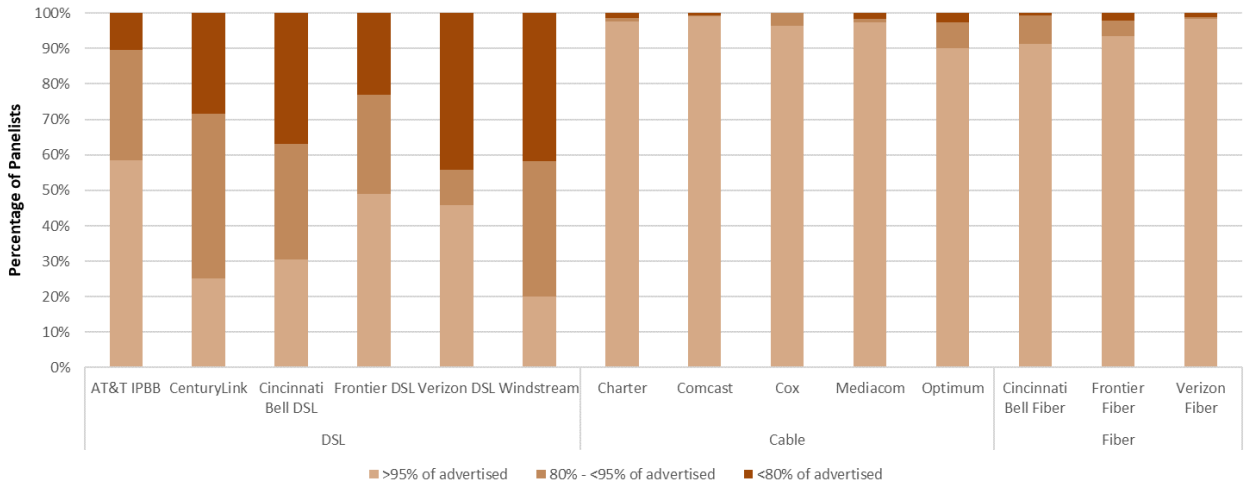

Though the median upload speeds experienced by most subscribers were close to or exceeded the advertised upload speeds there were some subscribers, for each ISP, whose median upload speed fell significantly short of the advertised upload speed. This issue was most prevalent for ISPs using DSL technology. On the other hand, ISPs using cable and fiber technology generally showed very good consistency based on this metric.

We can learn more about the variation in network performance by separately examining variations across geography and across time. We start by examining the variation across geography within each participating ISP's service territory. For each ISP, we first calculate the ratio of the median download speed (over the peak usage period) to the advertised download speed for each panelist subscribing to that ISP. We then examine the distribution of this ratio across the ISP's service territory.

Charts 14.1 and 14.2 show the complementary cumulative distribution of the ratio of median download speed (over the peak usage period) to advertised download speed for each participating ISP. For each ratio of actual to advertised download speed on the horizontal axis, the curves show the percentage of panelists subscribing to each ISP that experienced at least this ratio.[25] For example, the Cincinnati Bell fiber curve in Chart 14.1 shows that 90% of its subscribers experienced a median download speed exceeding 83% of the advertised download speed, while 70% experienced a median download speed exceeding 95% of the advertised download speed, and 50% experienced a median download speed exceeding 102% of the advertised download speed.

[25] In Reports prior to the 2015 MBA Report, for each ratio of actual to advertised download speed on the horizontal axis, the cumulative distribution function curves showed the percentage of measurements, rather than panelists subscribing to each ISP, that experienced at least this ratio. The methodology used since then, *i.e.*, using panelists subscribing to each ISP, more accurately illustrates ISP performance from a consumer's point of view.

Chart 14.1: Complementary cumulative distribution of the ratio of median download speed to advertised download speed.

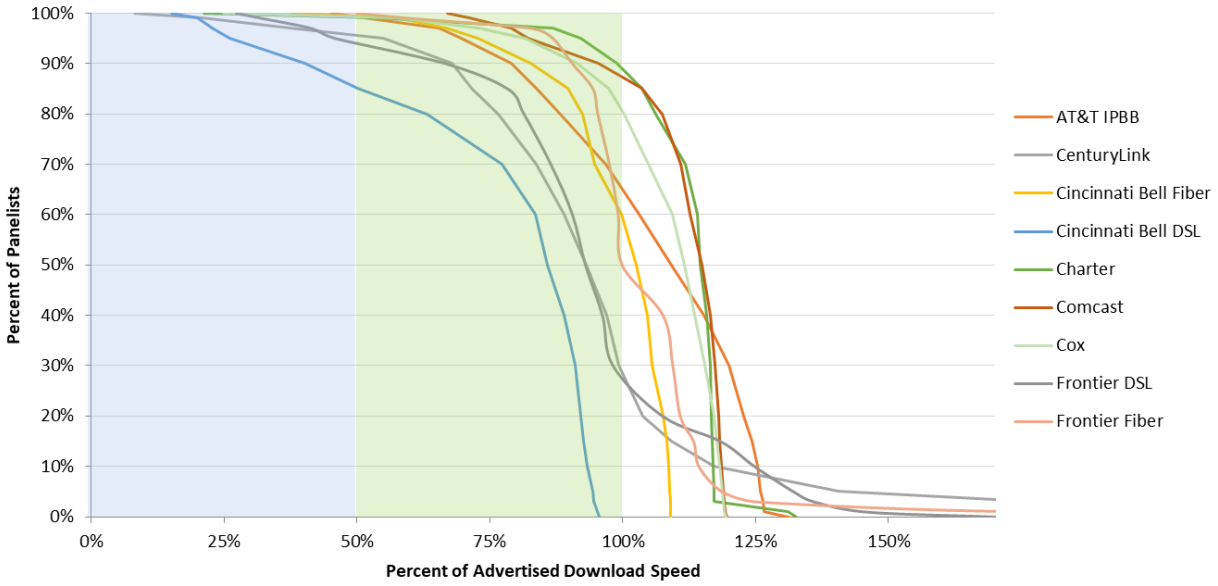

Chart 14.2: Complementary cumulative distribution of the ratio of median download speed to advertised download speed (continued).

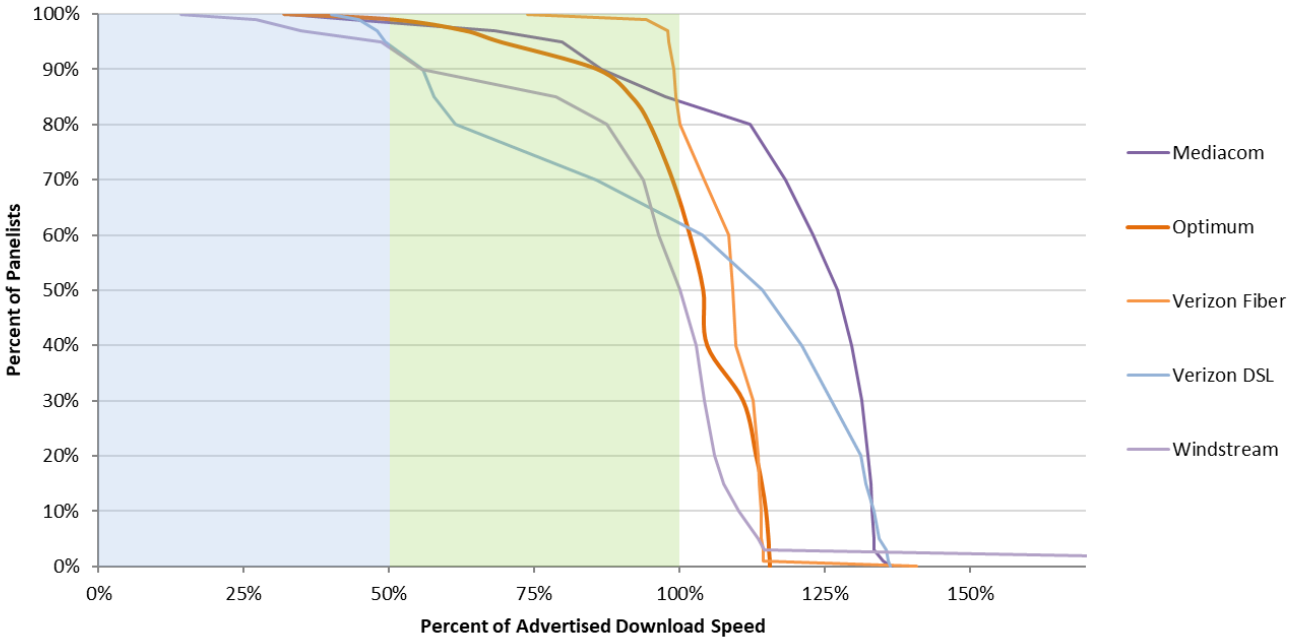

The curves for cable-based broadband and fiber-based broadband are steeper than those for DSL-based broadband. This can be seen more clearly in Chart 14.3, which plots aggregate curves for each technology. Approximately 80% of subscribers to cable and 60% of subscribers to fiber-based technologies experience

median download speeds exceeding the advertised download speed. In contrast, only 30% of subscribers to DSL-based services experience median download speeds exceeding the advertised download speed.[26]

Chart 14.3: Complementary cumulative distribution of the ratio of median download speed to advertised download speed, by technology.

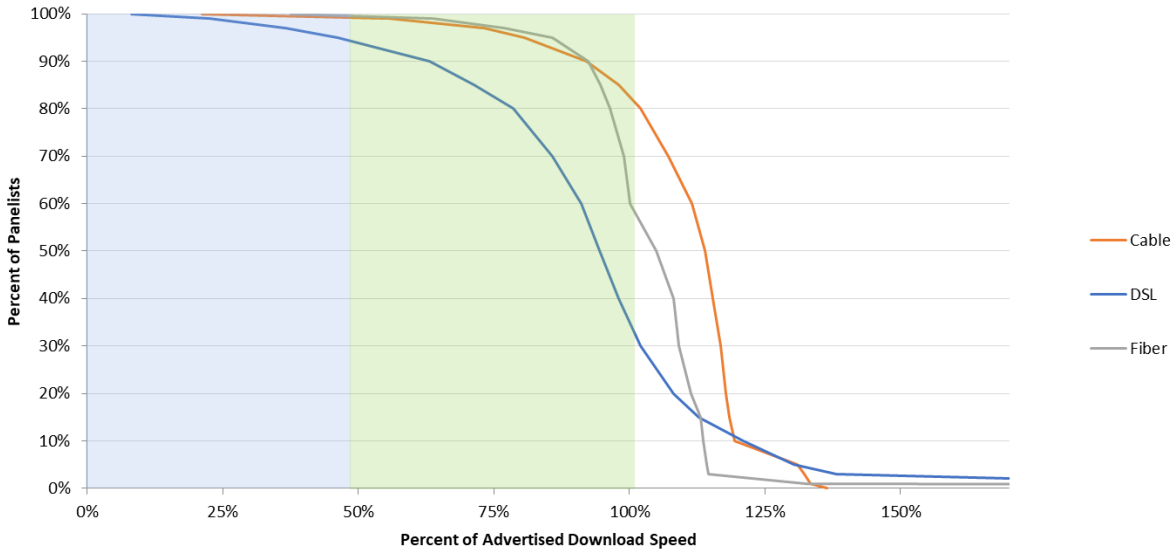

Charts 14.4 to 14.6 show the complementary cumulative distribution of the ratio of median upload speed (over the peak usage period) to advertised upload speed for each participating ISP (Charts 14.4 and 14.5) and by access technology (Chart 14.6).

[26] The speed achievable by DSL depends on the distance between the subscriber and the central office. Thus, the complementary cumulative distribution function will fall slowly unless the broadband ISP adjusts its advertised rate based on the subscriber's location. (Chart 16 illustrates that the performance during non-busy hours is similar to the busy hour, making congestion less likely as an explanation.)

Chart 14.4: Complementary cumulative distribution of the ratio of median upload speed to advertised upload speed.

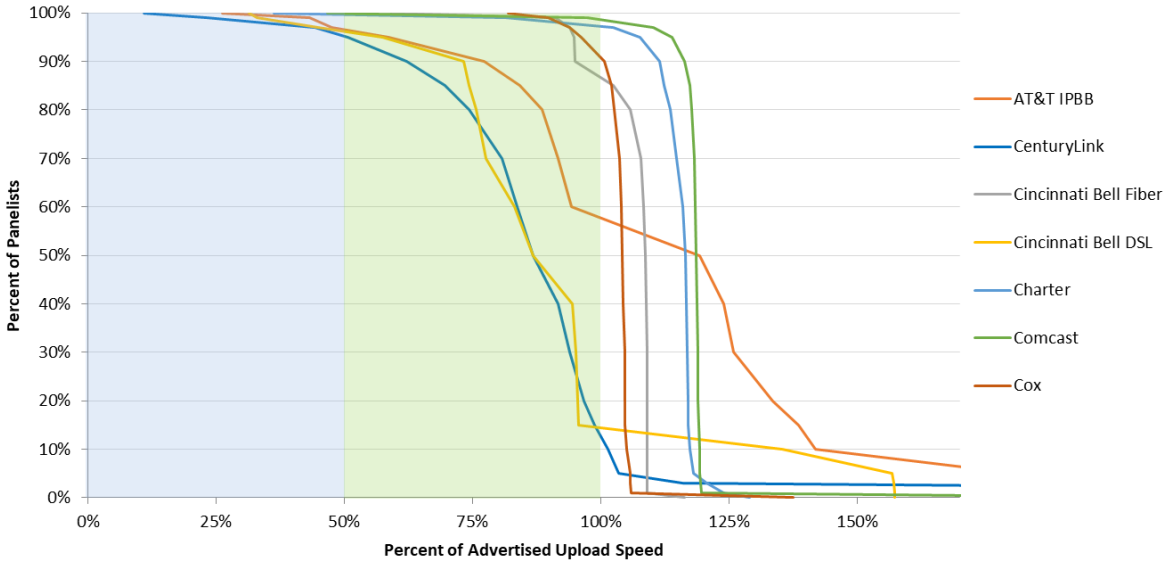

Chart 14.5: Complementary cumulative distribution of the ratio of median upload speed to advertised upload speed (continued).

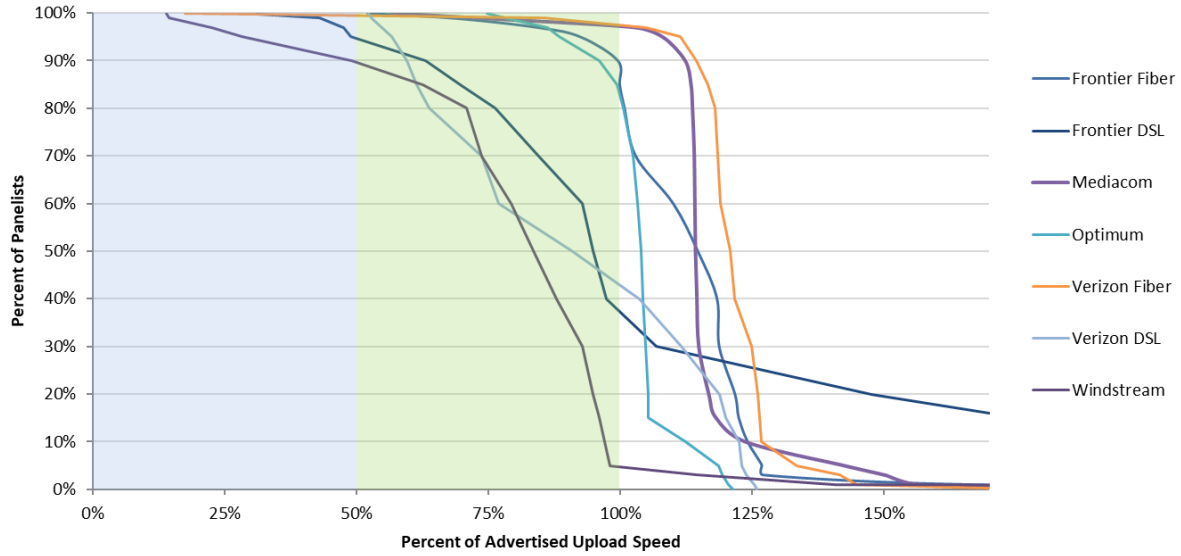

Chart 14.6: Complementary cumulative distribution of the ratio of median upload speed to advertised upload speed, by technology.

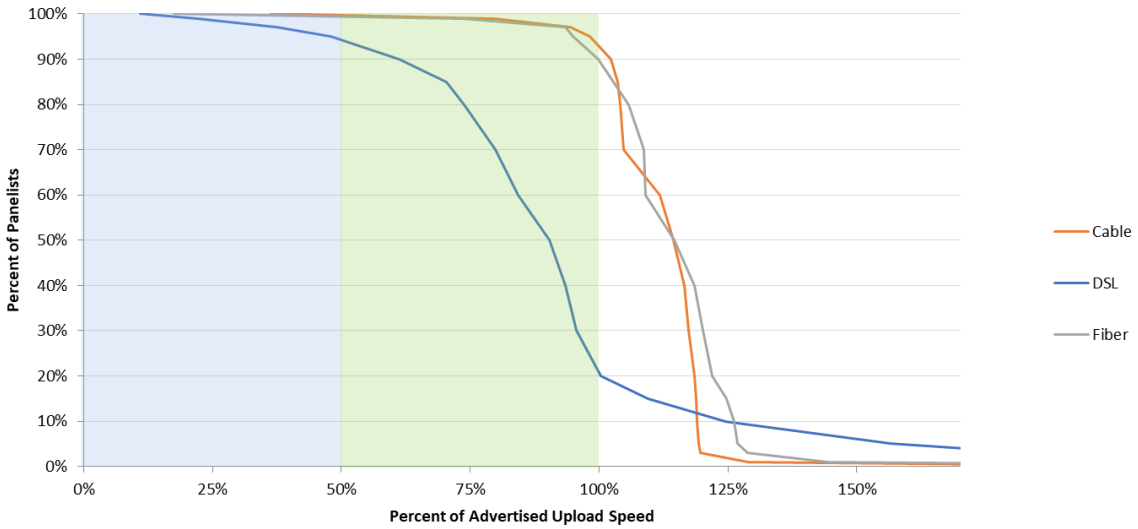

All actual speeds discussed above were measured during peak usage periods. In contrast, Charts 15.1 and 15.2 below compare the ratio of actual download and upload speeds to advertised download and upload speeds during peak and off-peak times.[27] Charts 15.1 and 15.2 show that most ISP subscribers experience only a slight degradation from off-peak to peak hour performance.

Chart 15.1: The ratio of weighted median download speed to advertised download speed, peak hours versus off-peak hours.

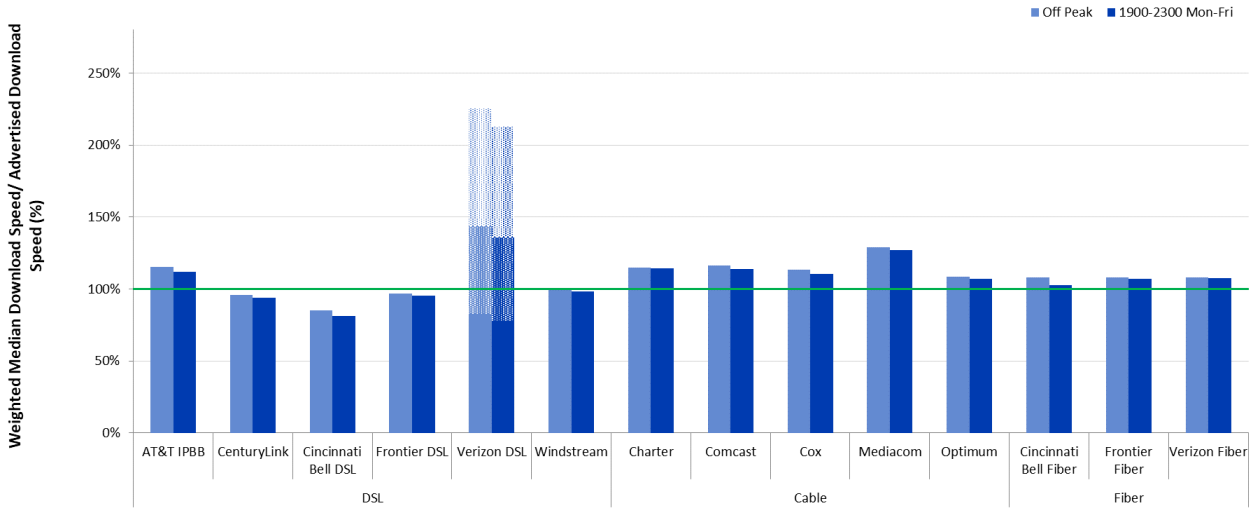

[27] As described earlier, Verizon DSL download and upload results are shown as a range since Verizon advertises its DSL speed as a range rather than as a specific speed.

Chart 15.2: The ratio of weighted median upload speed to advertised upload speed, peak versus off-peak.

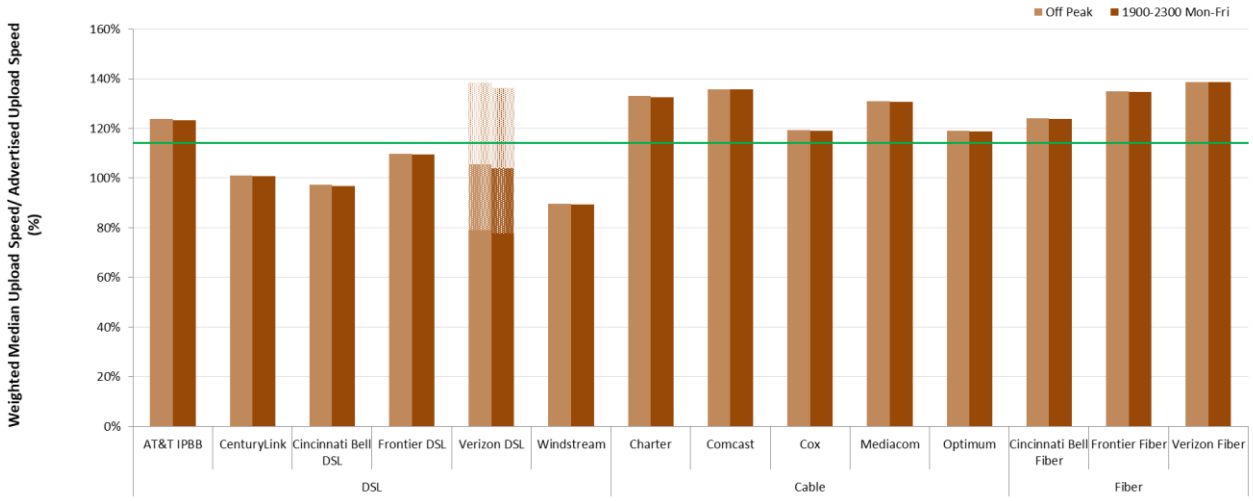

Chart 16[28] below shows the actual download speed to advertised speed ratio in each two-hour time block during weekdays for each ISP. The ratio is lowest during the busiest four-hour time block (7:00 p.m. to 11:00 p.m.).

[28] In this chart, we have shown the median download speed of Verizon-DSL as a percentage of the midpoint of the advertised speed range for its tier.

Chart 16: The ratio of median download speed to advertised download speed, Monday-to-Friday, two-hour time blocks, terrestrial ISPs.

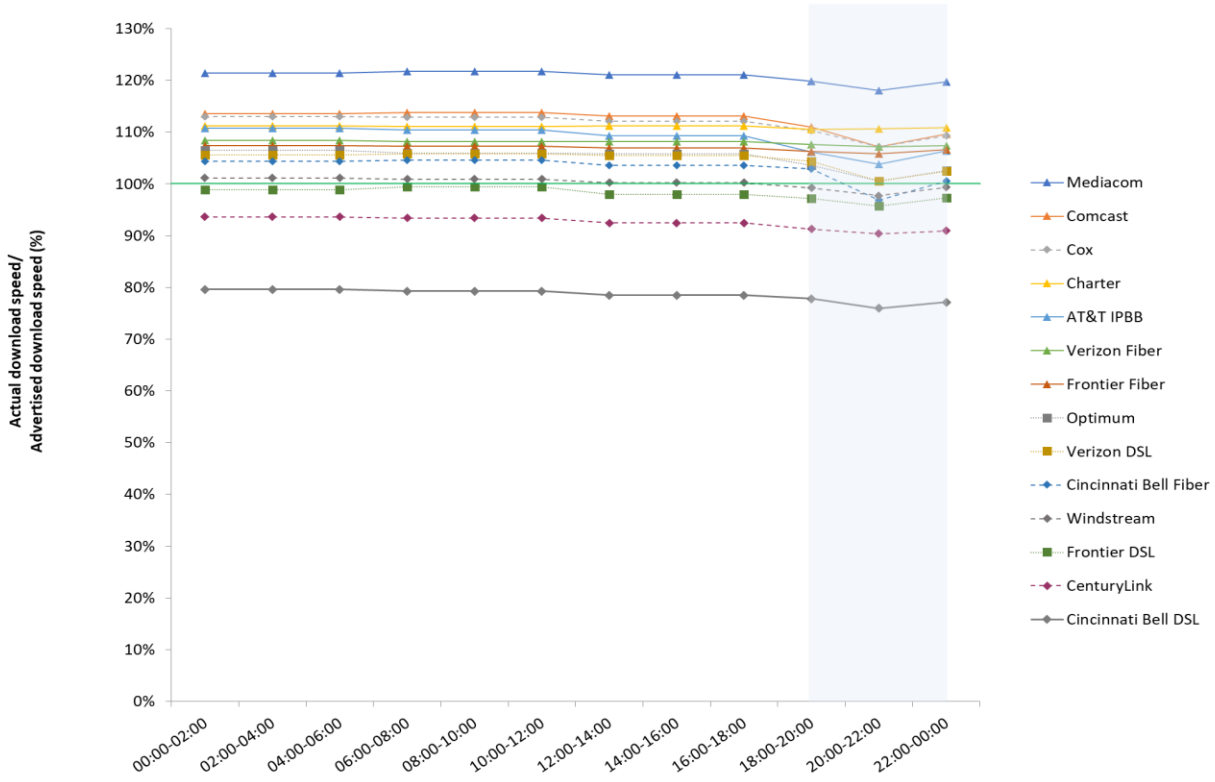

For each ISP, Chart 6 (in section 2.C) showed the ratio of the 80/80 consistent median download speed to advertised download speed, and for comparison, Chart 4 showed the ratio of median download speed to advertised download speed.

Chart 17.1 illustrates information concerning 80/80 consistent upload speeds. While all the upload 80/80 speeds were slightly lower than the median speed the differences were more marked in DSL. Charts 6 and 17.1 make it clear that cable and fiber technologies behaved more consistently than DSL technology both for download as well as upload speeds.

Chart 17.1: The ratio of 80/80 consistent upload speed to advertised upload speed.

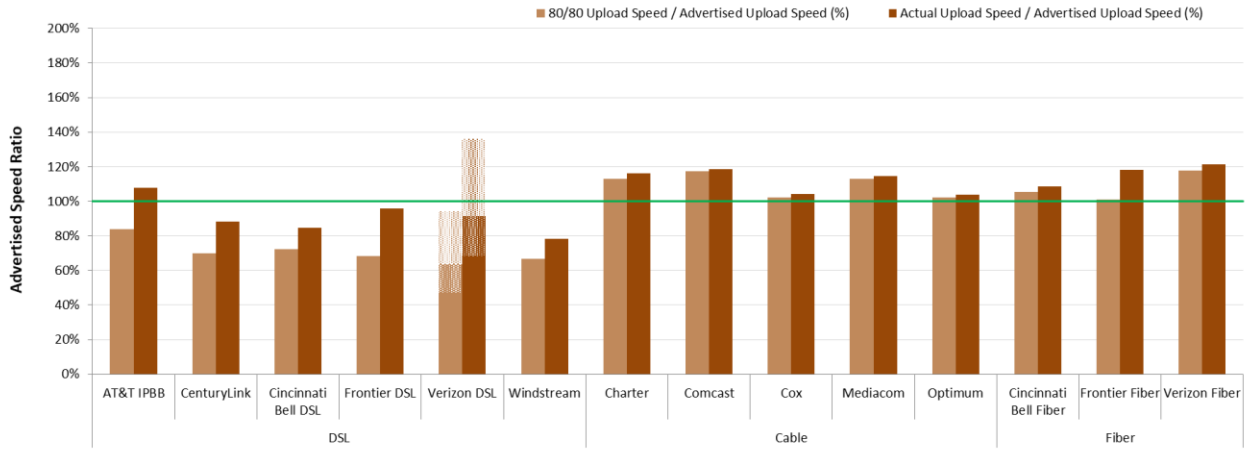

Charts 17.2 and 17.3 below illustrate similar consistency metrics for 70/70 consistent download and upload speeds, *i.e.*, the minimum download or upload speed (as a percentage of the advertised download or upload speed) experienced by at least 70% of panelists during at least 70% of the peak usage period. The ratios for 70/70 consistent speeds as a percentage of the advertised speed are higher than the corresponding ratios for 80/80 consistent speeds. In fact, for many ISPs, the 70/70 consistent download or upload speed is close to the median download or upload speed. Once again, ISPs using DSL technology showed a considerably smaller value for the 70/70 download and upload speeds as compared to the download and upload median speeds, respectively.

Chart 17.2: The ratio of 70/70 consistent download speed to advertised download speed.

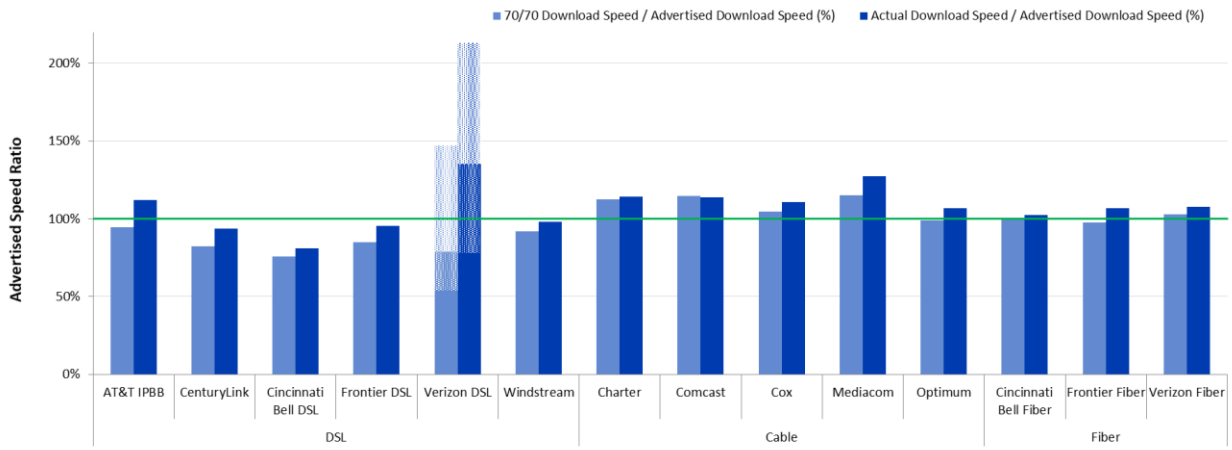

Chart 17.3: The ratio of 70/70 consistent upload speed to advertised upload speed.

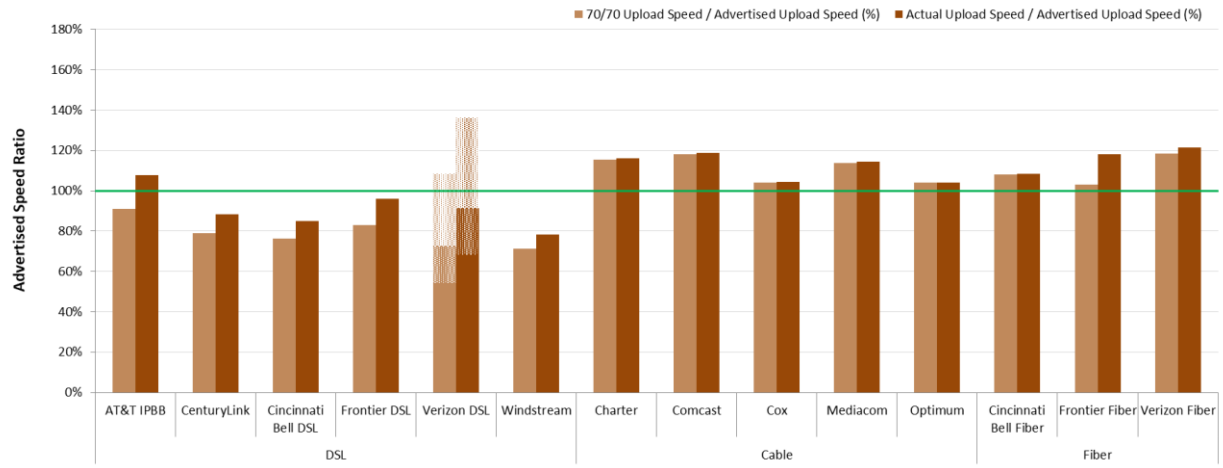

D. LATENCY

Chart 18 below shows the weighted median latencies, by technology and by advertised download speed for terrestrial technologies. For all terrestrial technologies, latency varied little with advertised download speed. DSL service typically had higher latencies, and lower latency was better correlated with advertised download speed, than with either cable or fiber. Cable latencies ranged between 18ms to 24ms, fiber latencies between 5ms to 12ms, and DSL between 27ms to 55ms.

Chart 18: Latency for Terrestrial ISPs, by technology, and by advertised download speed.

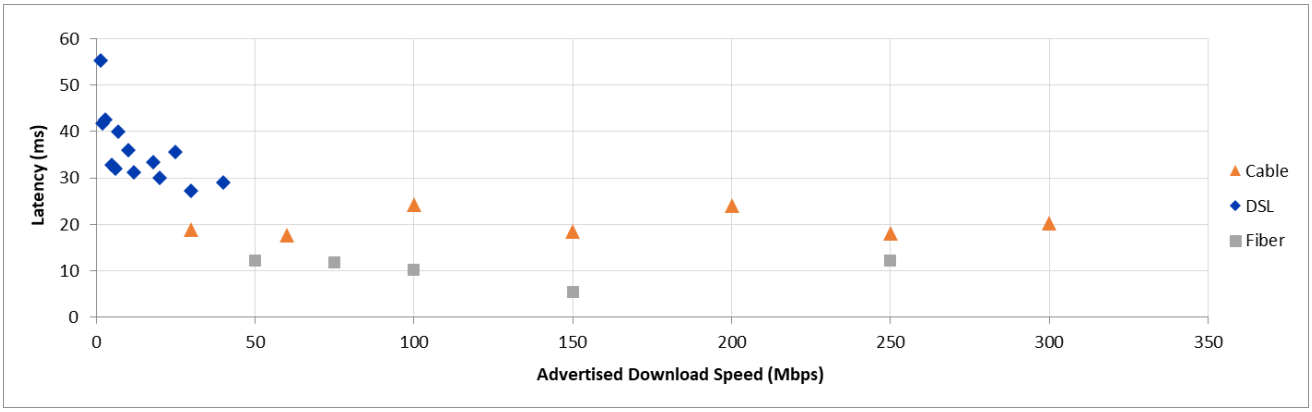

5. ADDITIONAL TEST RESULTS

A. ACTUAL SPEED, BY SERVICE TIER

As shown in Charts 19.1-19.8, peak usage period performance varied by service tier among participating ISPs during the September-October 2018 period. On average, during peak periods, the ratio of median download speed to advertised download speed for all ISPs was 57% or better, and 90% or better for most ISPs. However, the ratio of median download speed to advertised download speed varies among service tiers. It should be noted that for Verizon-DSL, which advertises a range of speeds, we have calculated a range of values corresponding to its advertised range. Out of the 44 speed tiers that were measured a large majority (41) showed that they at least achieved 90% of the advertised speed and 24 of the 44 tiers either met or exceeded the advertised speed.

Chart 19.1: The ratio of median download speed to advertised download speed, by ISP (1-5 Mbps).

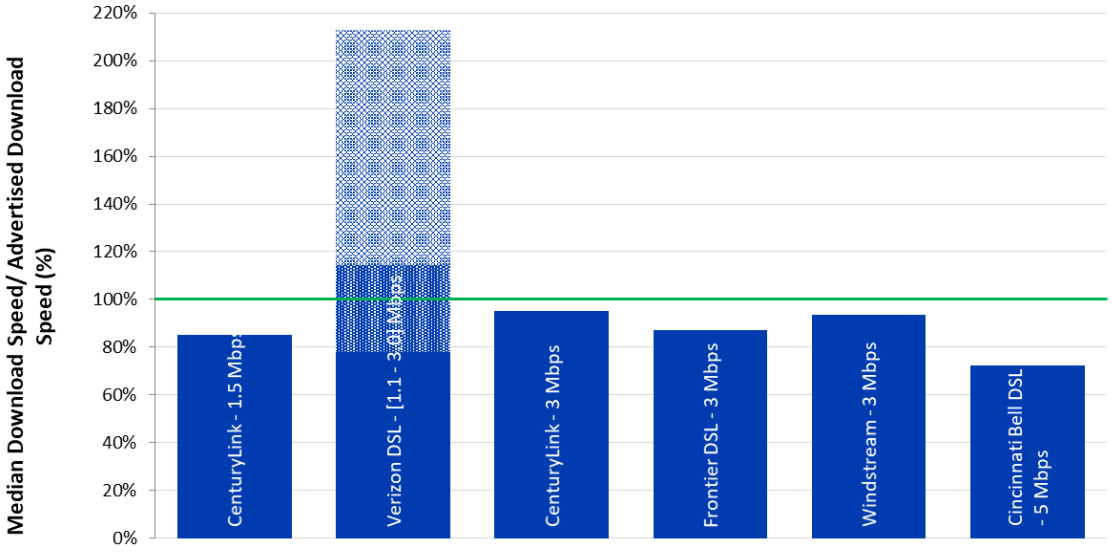

Chart 19.2: The ratio of median download speed to advertised download speed, by ISP (6-10 Mbps).

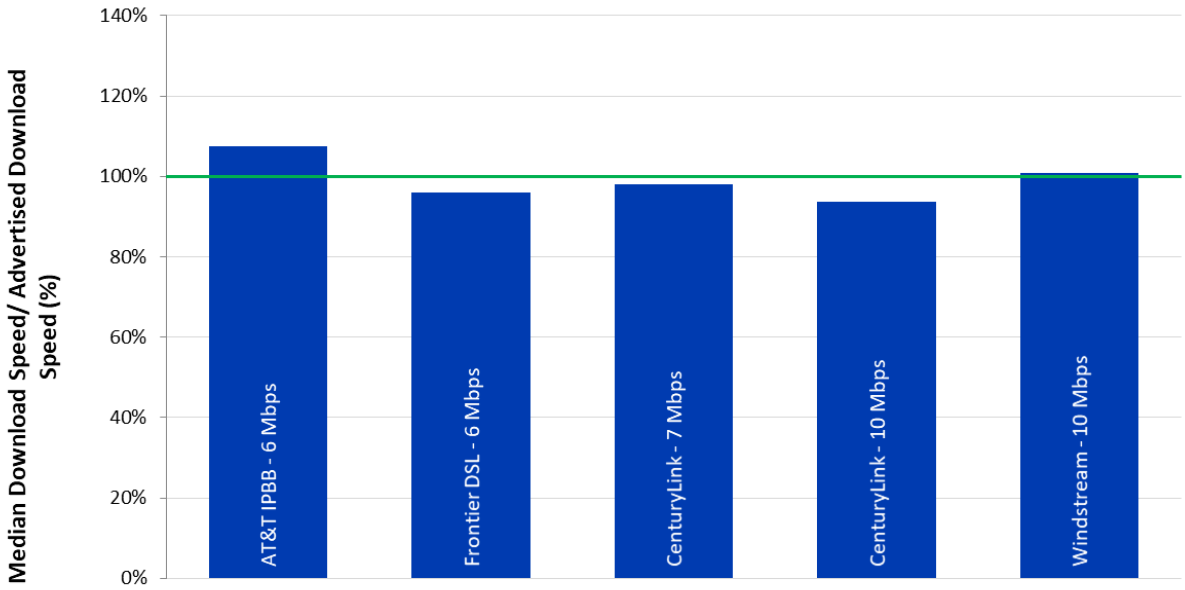

Chart 19.3: The ratio of median download speed to advertised download speed, by ISP (12-20 Mbps).

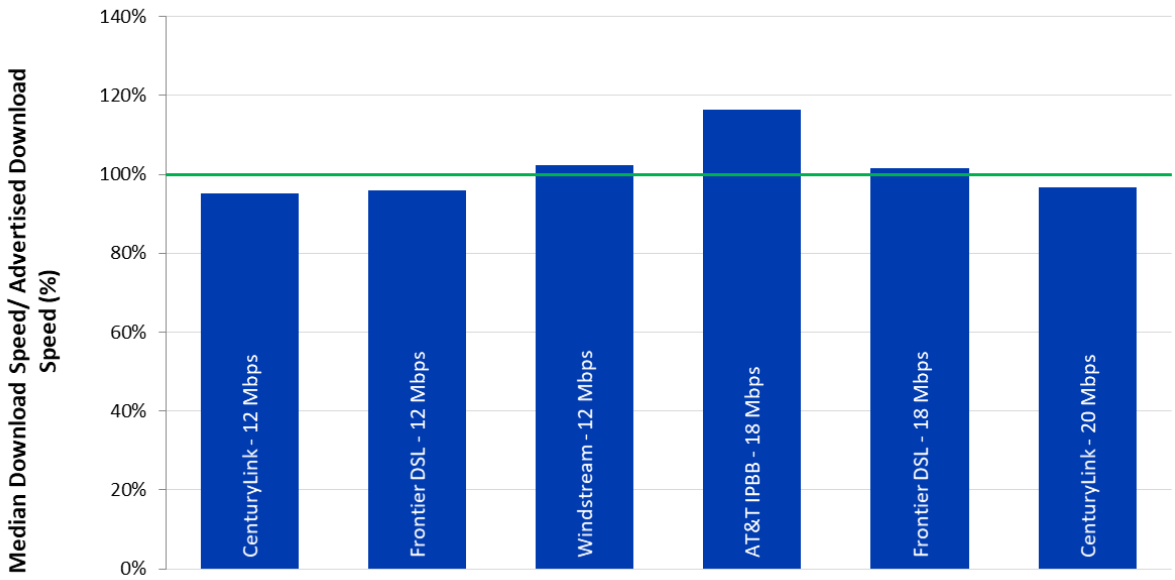

Chart 19.4: The ratio of median download speed to advertised download speed, by ISP (25-30 Mbps).

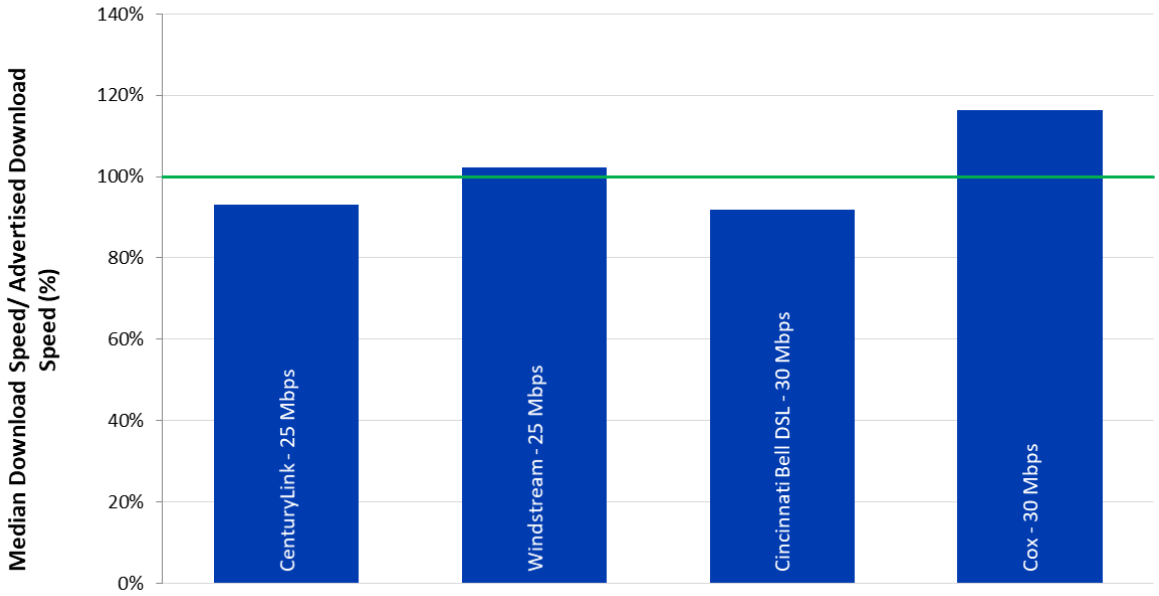

Chart 19.5: The ratio of median download speed to advertised download speed, by ISP (40-50 Mbps).

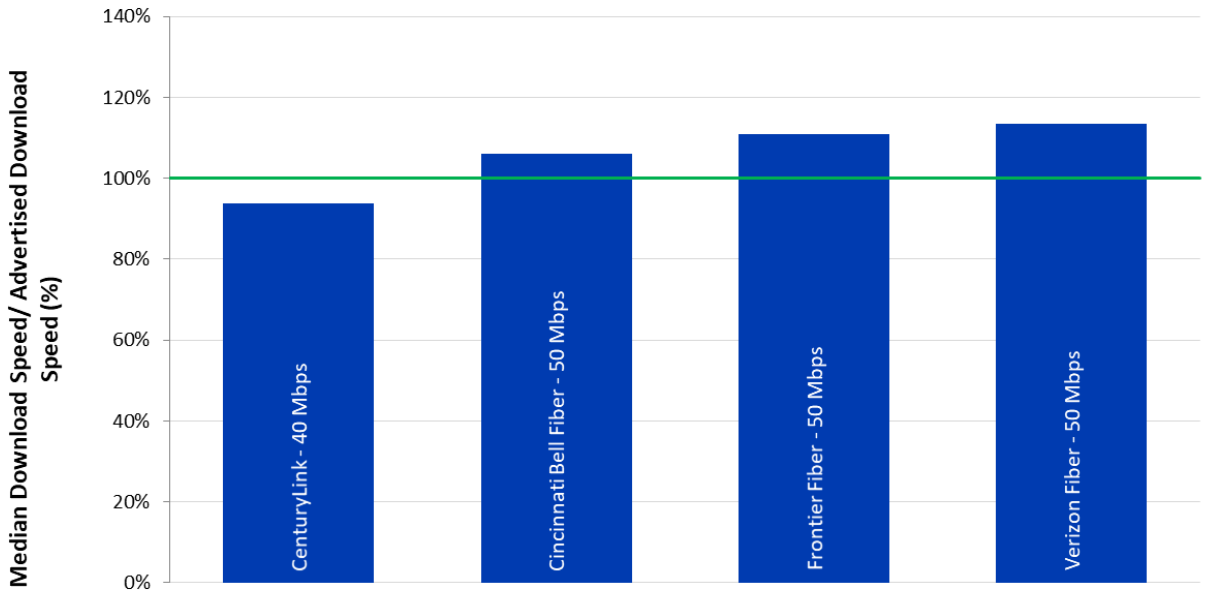

Chart 19.6: The ratio of median download speed to advertised download speed, by ISP (60-75 Mbps).

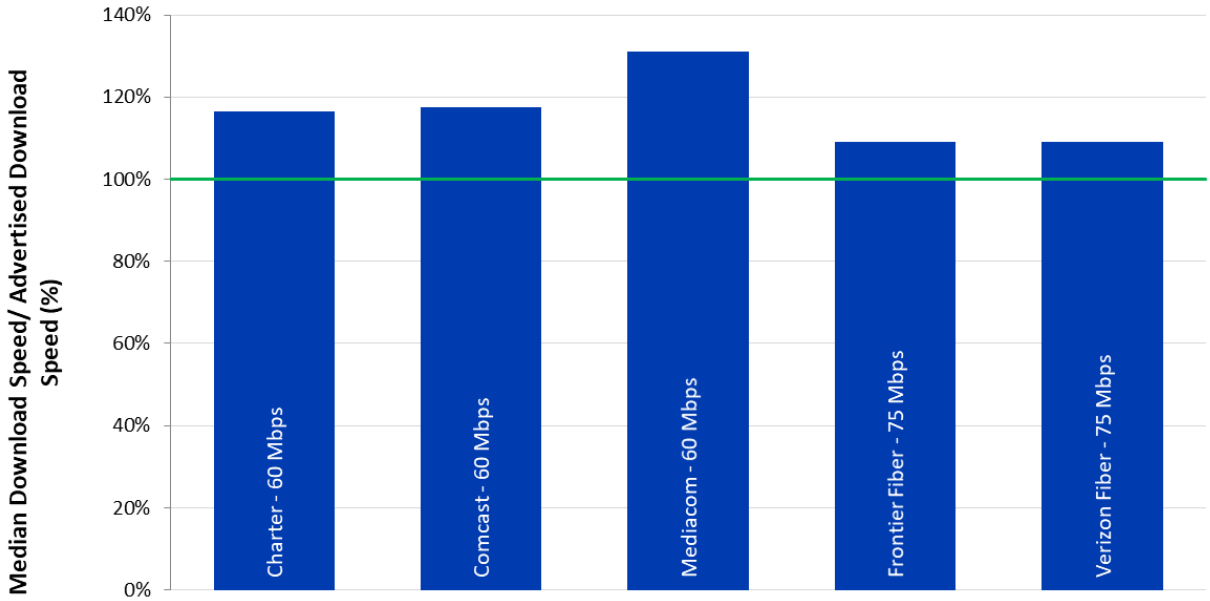

Chart 19.7: The ratio of median download speed to advertised download speed, by ISP (100-150 Mbps).

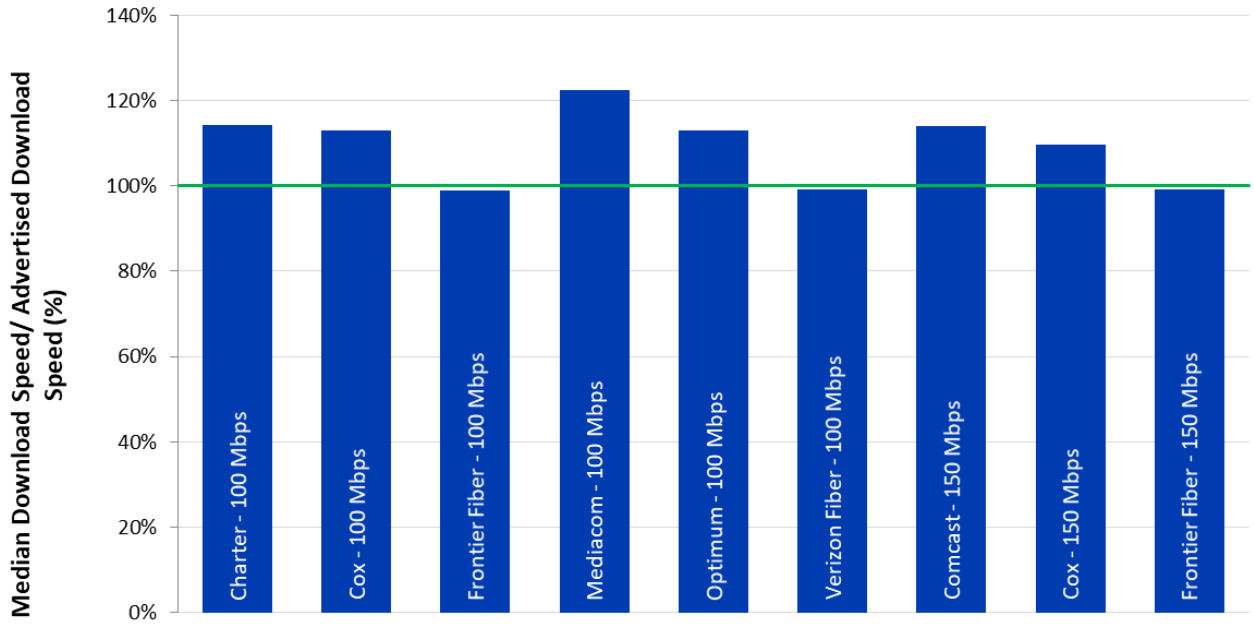

Chart 19.8: The ratio of median download speed to advertised download speed, by ISP (200-300 Mbps).

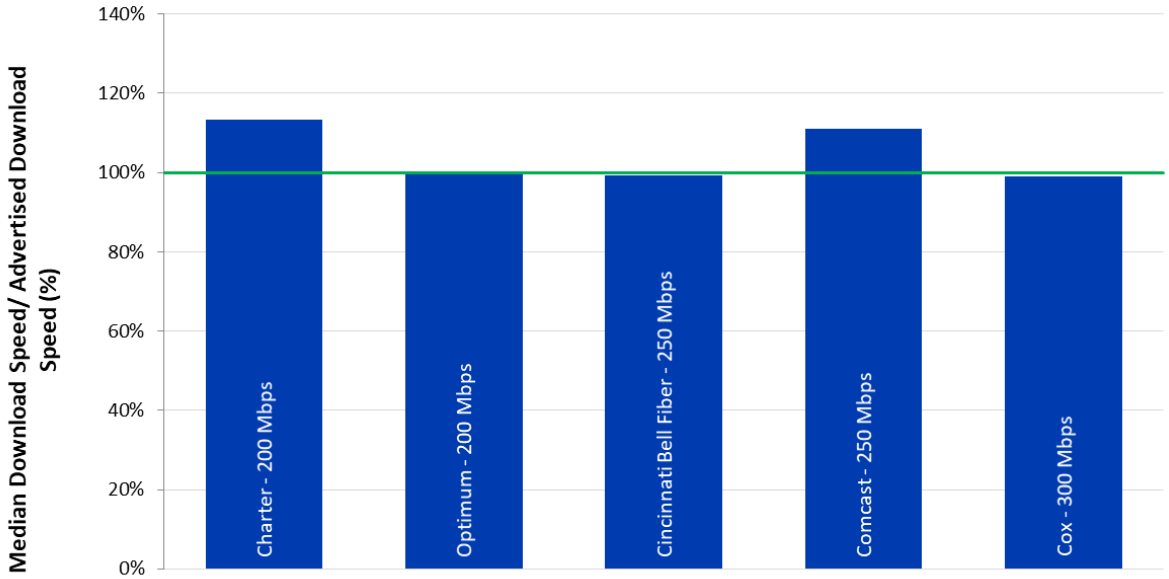

Charts 20.1 – 20.6 depict the ratio of median upload speeds to advertised upload speeds for each ISP by service tier.

Chart 20.1: The ratio of median upload speed to advertised upload speed, by ISP (0.384 - 0.768 Mbps).

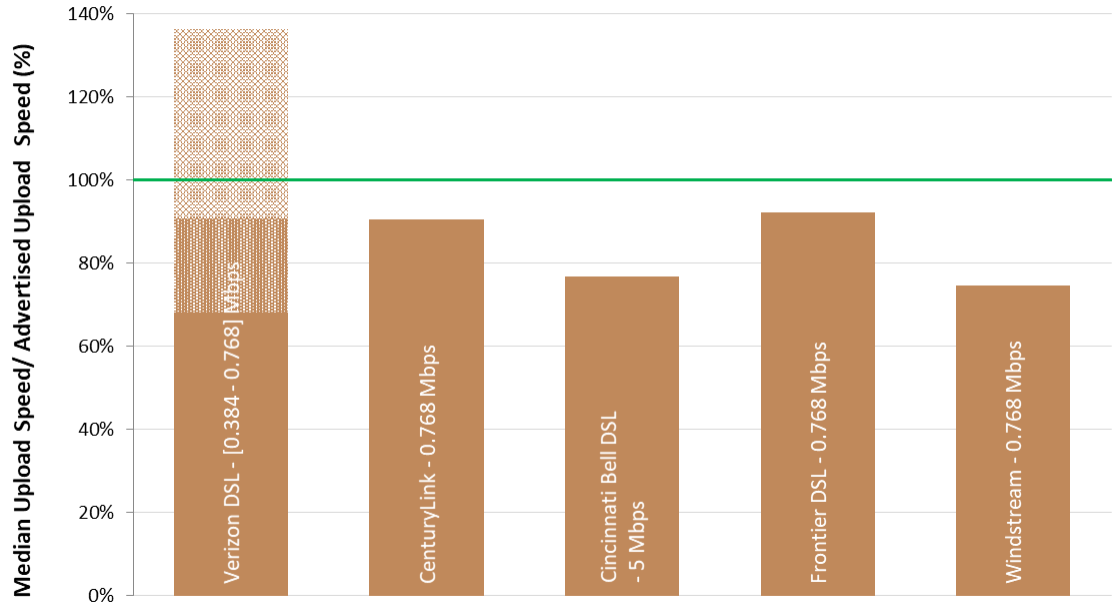

Chart 20.2: The ratio of median upload speed to advertised upload speed, by ISP (0.896 – 1.5 Mbps).

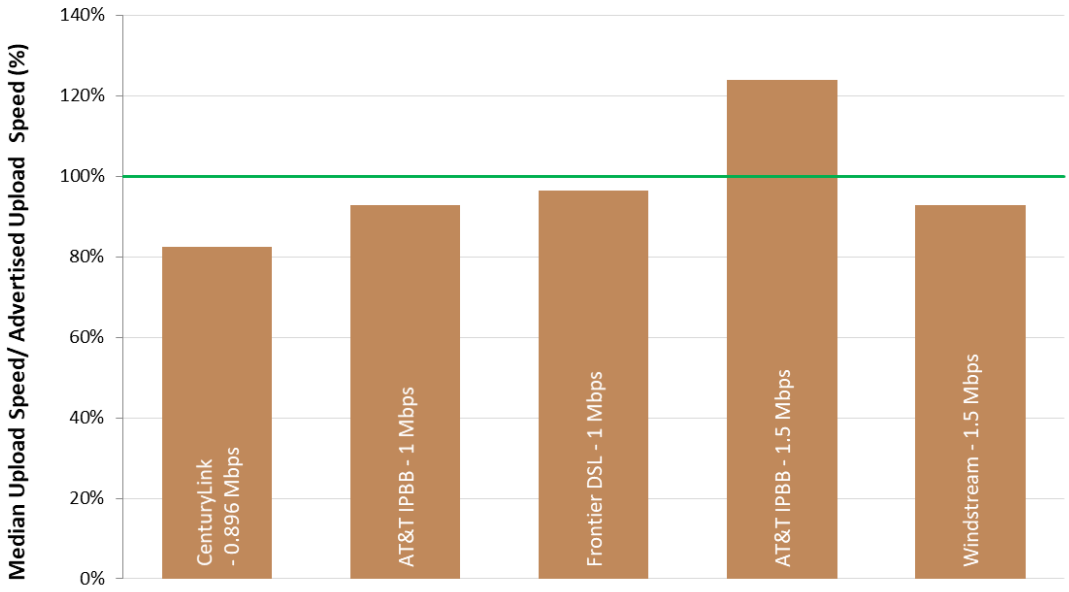

Chart 20.3: The ratio of median upload speed to advertised upload speed, by ISP (2-5 Mbps).

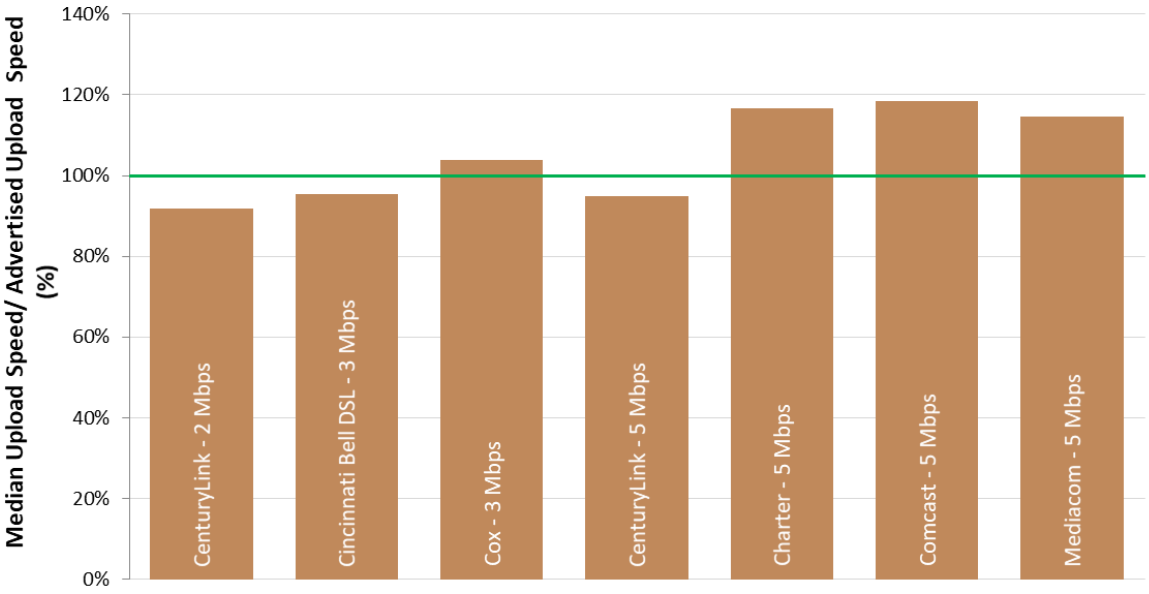

Chart 20.4: The ratio of median upload speed to advertised upload speed, by ISP (10 - 20 Mbps).

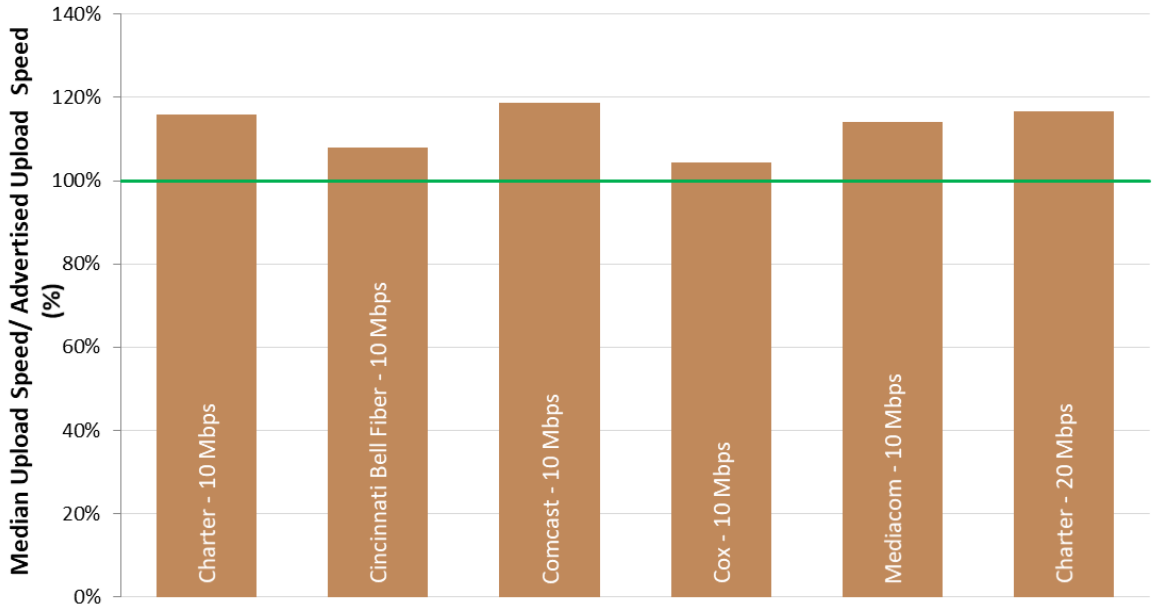

Chart 20.5: The ratio of median upload speed to advertised upload speed, by ISP (30 - 75 Mbps).

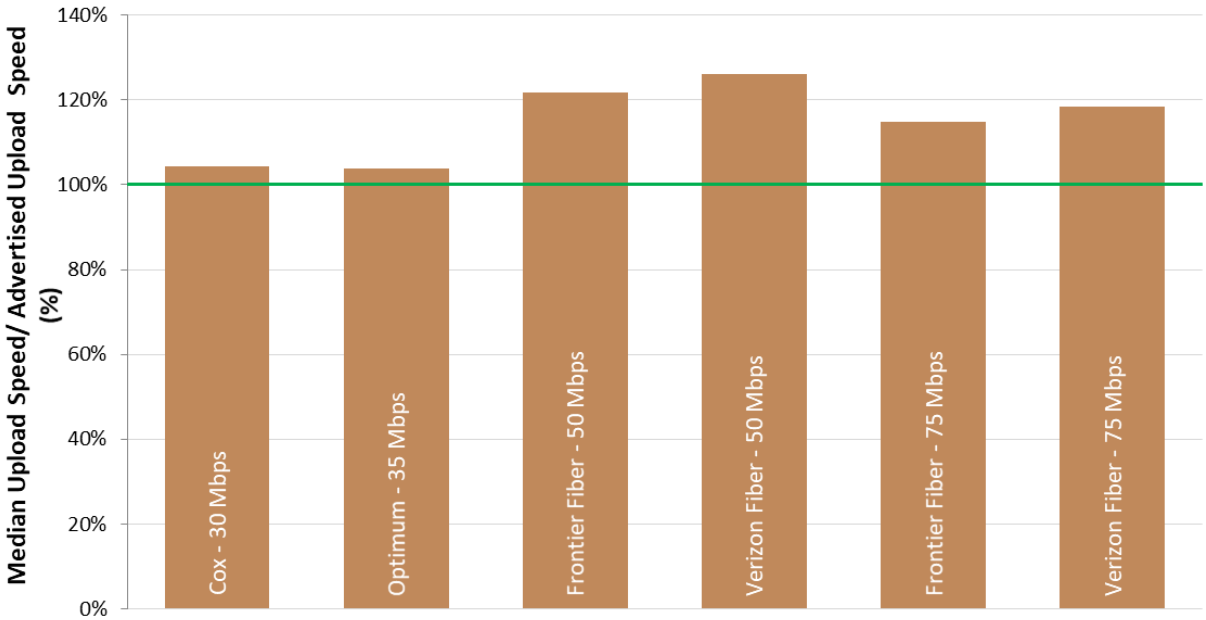

Chart 20.6: The ratio of median upload speed to advertised upload speed, by ISP (100-150 Mbps).

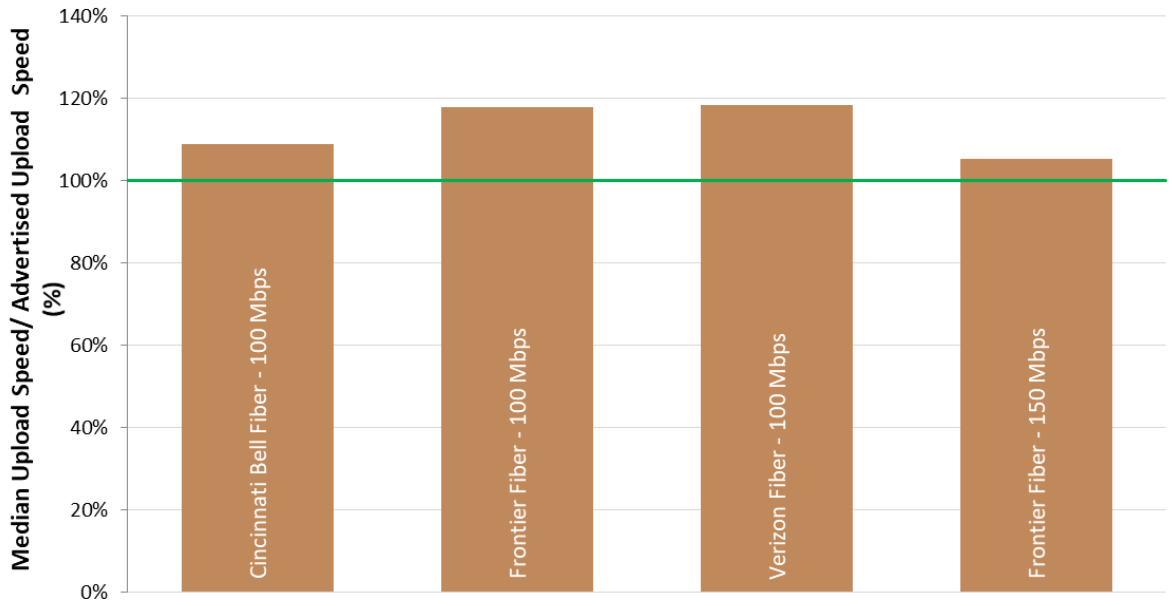

Table 2 lists the advertised download service tiers included in this study. For each tier, an ISP's advertised download speed is compared with the median of the measured download speed results. As we noted in the past reports, the download speeds listed here are based on national averages and may not represent the performance experienced by any particular consumer at any given time or place.

Table 2: Peak period median download speed, sorted by actual download speed

Download Median Speed (Mbps)	Advertised Download Speed (Mbps)	ISP	Actual Speed / Advertised Speed (%)
1.28	1.5	CenturyLink	85.2
2.34	1.1 - 3	Verizon DSL	114.2% (78.1% - 212.9%)
2.86	3	CenturyLink	95.2
2.61	3	Frontier DSL	87
2.81	3	Windstream	93.7
3.62	5	Cincinnati Bell DSL	72.4
6.45	6	AT&T IPBB	107.6
5.75	6	Frontier DSL	95.9

6.86	7	CenturyLink	98.0
9.35	10	CenturyLink	93.6
10.08	10	Windstream	100.8
11.42	12	CenturyLink	95.1
11.53	12	Frontier DSL	96.1
12.29	12	Windstream	102.4
20.93	18	AT&T IPBB	116.3
18.27	18	Frontier DSL	101.5
19.35	20	CenturyLink	96.8
23.28	25	CenturyLink	93.1
25.55	25	Windstream	102.2
27.54	30	Cincinnati Bell DSL	91.8
34.90	30	Cox	116.3
37.50	40	CenturyLink	93.8
52.99	50	Cincinnati Bell Fiber	106
55.45	50	Frontier Fiber	110.9
56.75	50	Verizon Fiber	113.5
69.90	60	Charter	116.5
70.43	60	Comcast	117.4
78.64	60	Mediacom	131.1
81.82	75	Frontier Fiber	109.1
81.73	75	Verizon Fiber	109
114.26	100	Charter	114.3
112.98	100	Cox	113.0
98.82	100	Frontier Fiber	98.8
122.35	100	Mediacom	122.4
112.89	100	Optimum	112.9
99.17	100	Verizon Fiber	99.2
170.92	150	Comcast	114.0
164.37	150	Cox	109.6
148.61	150	Frontier Fiber	99.1

226.53	200	Charter	113.3
199.47	200	Optimum	99.7
248.06	250	Cincinnati Bell Fiber	99.2
277.39	250	Comcast	111.0
296.67	300	Cox	98.9

B. VARIATIONS IN SPEED

In Section 3.C above, we present speed consistency metrics for each ISP based on test results averaged across all service tiers. In this section, we provide detailed speed consistency results for each ISP's individual service tiers. Consistency of speed is important for services such as video streaming. A significant reduction in speed for more than a few seconds can force a reduction in video resolution or an intermittent loss of service.

Charts 21.1 – 21.3 below show the percentage of consumers that achieved greater than 95%, between 85% and 95%, or less than 80% of the advertised download speed for each ISP speed tier. Consistent with past performance, ISPs using DSL technology frequently fail to deliver advertised service rates. ISPs quote a single 'up-to' speed, but the actual speed of DSL depends on the distance between the subscriber and the serving central office.

Cable companies and fiber-based systems, in general, showed a high consistency of speed.

Chart 21.1: The percentage of consumers whose median download speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed, by service tier (DSL).

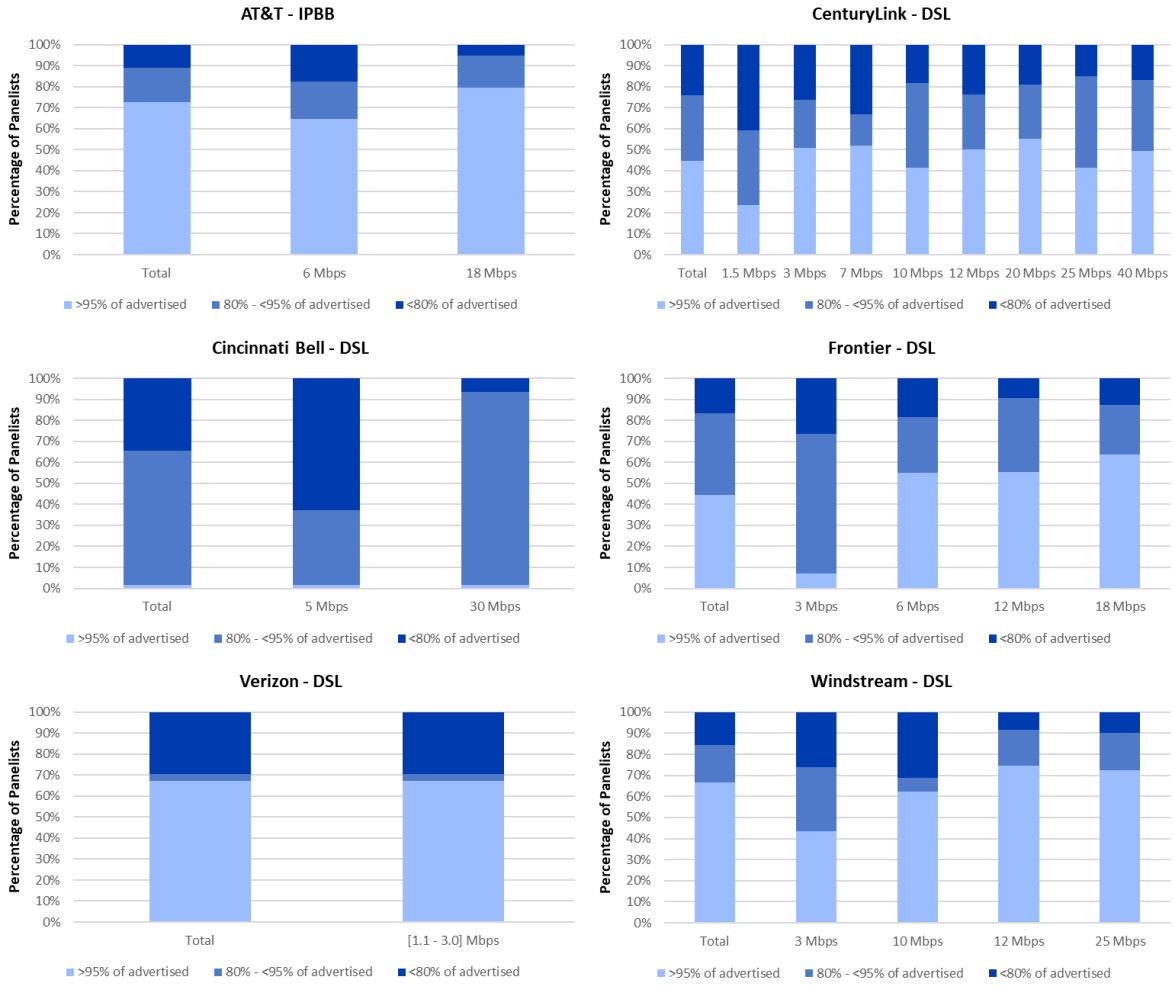

Chart 21.2: The percentage of consumers whose median download speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed (cable).

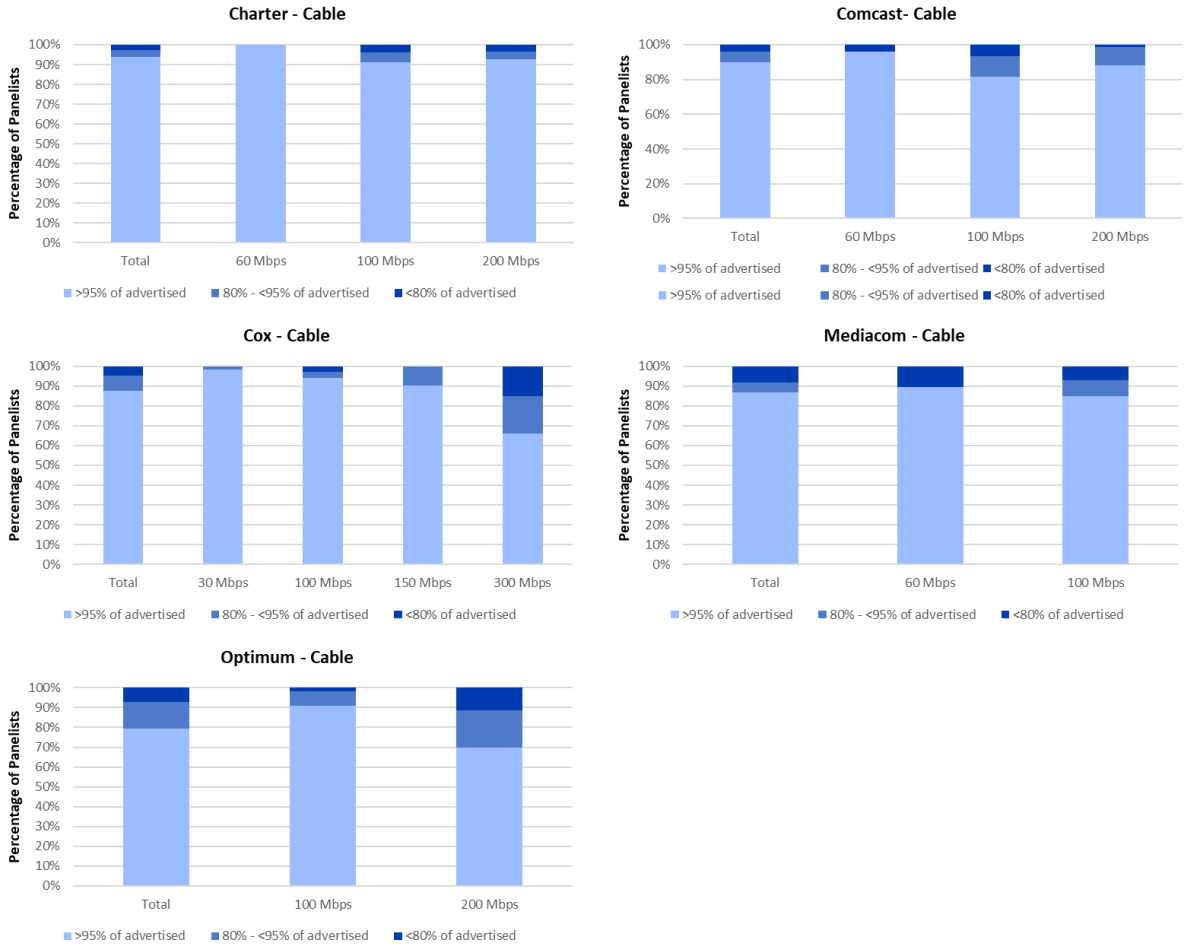

Chart 21.3: The percentage of consumers whose median download speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed (fiber).

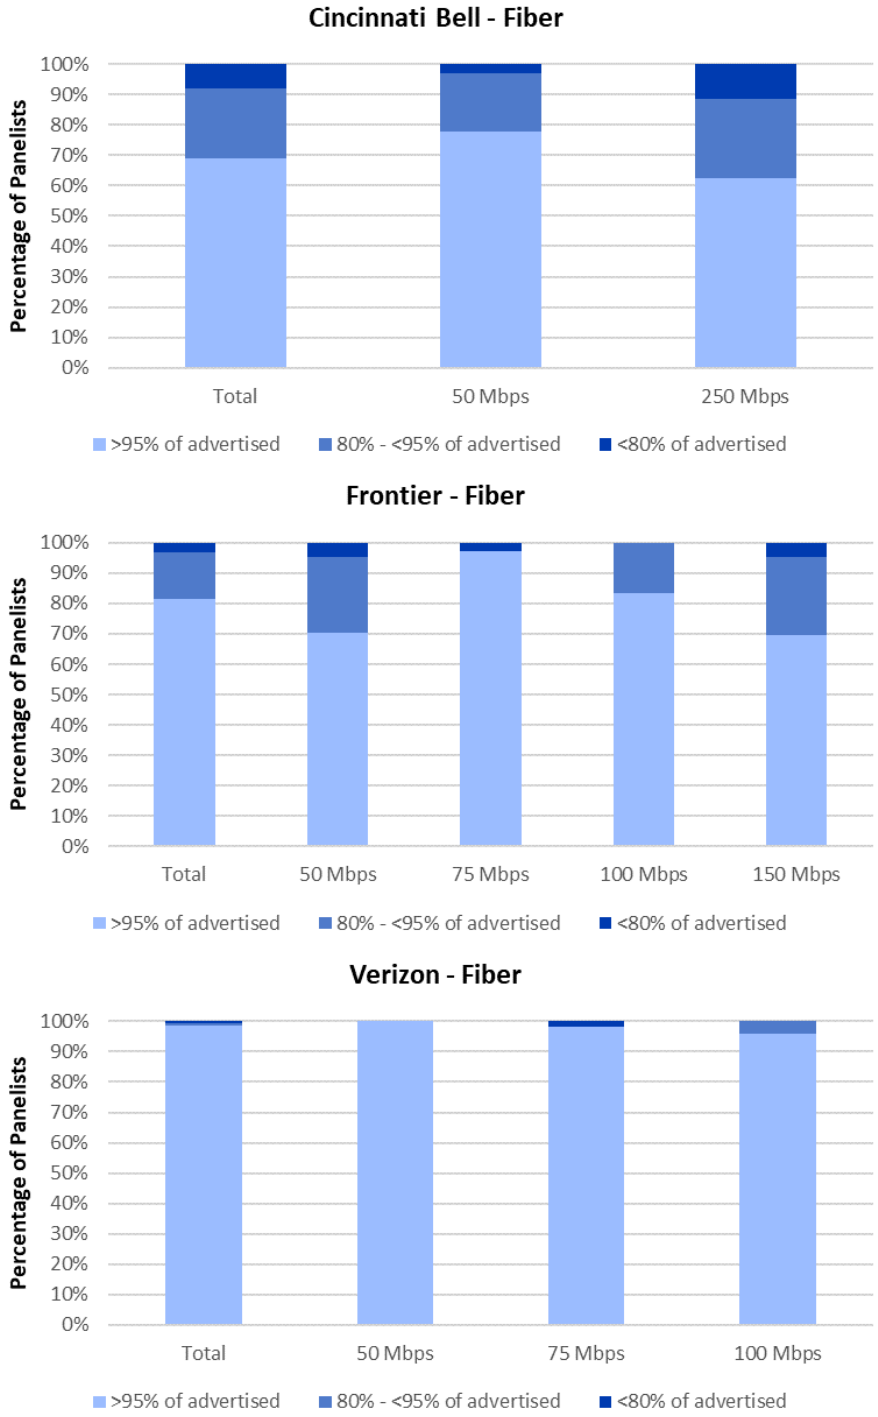

Similarly, Charts 22.1 to 22.3 show the percentage of consumers that achieved greater than 95%, between 85% and 95%, or less than 80% of the advertised upload speed for each ISP speed tier.

Chart 22.1: The percentage of consumers whose median upload speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised upload speed (DSL).

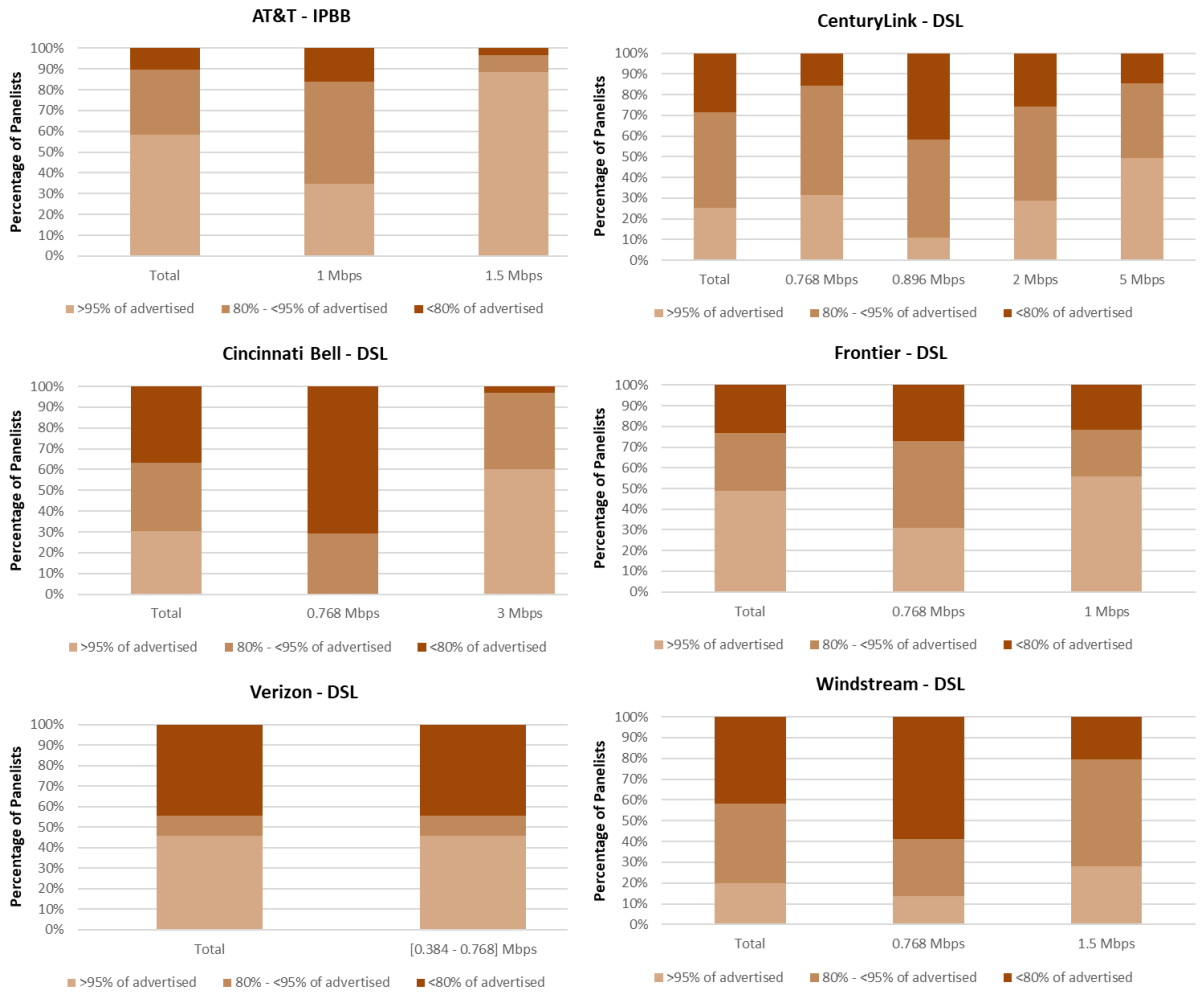

Chart 22.2: The percentage of consumers whose median upload speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised upload speed (cable).

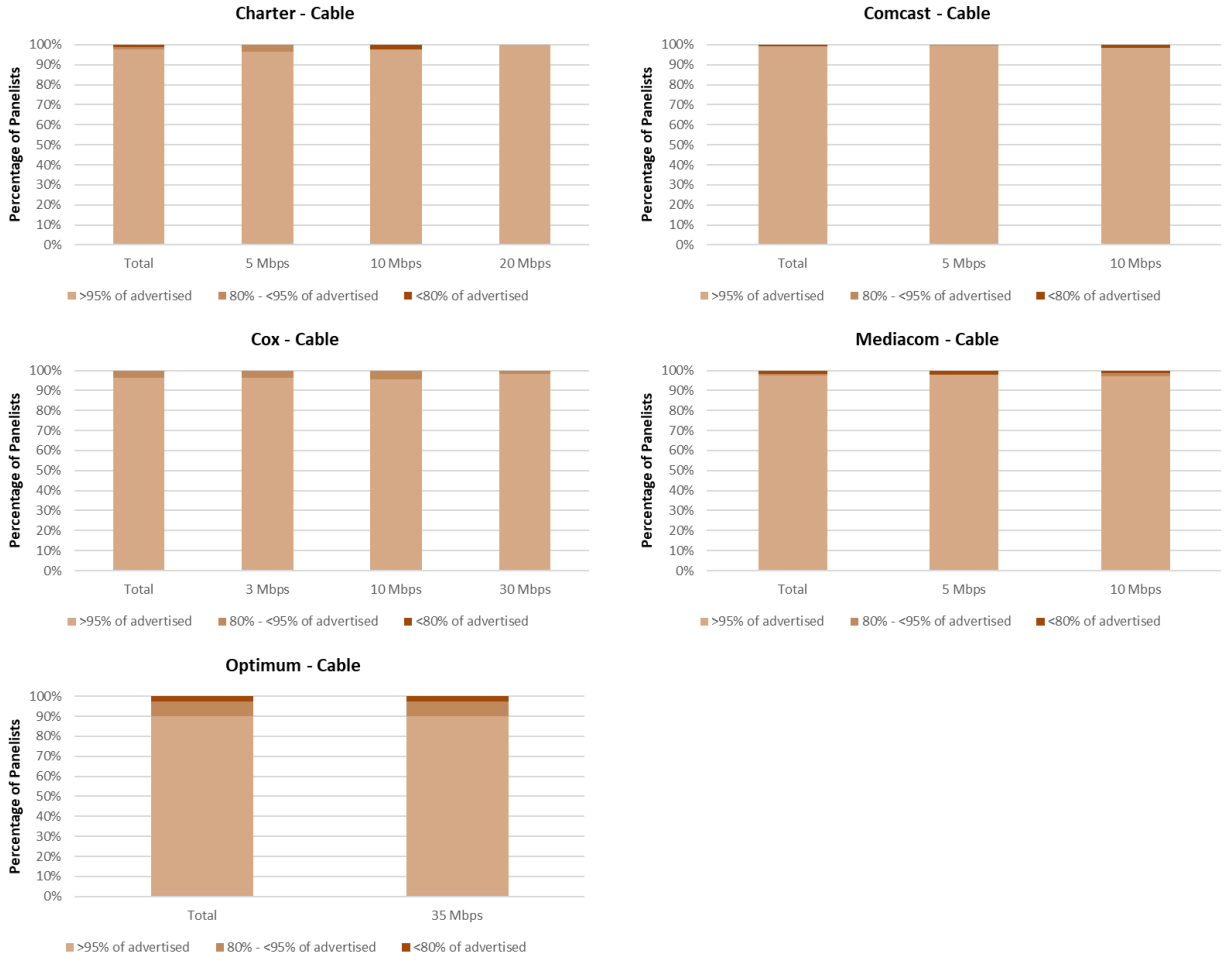

Chart 22.3: The percentage of consumers whose median upload speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised upload speed (fiber).

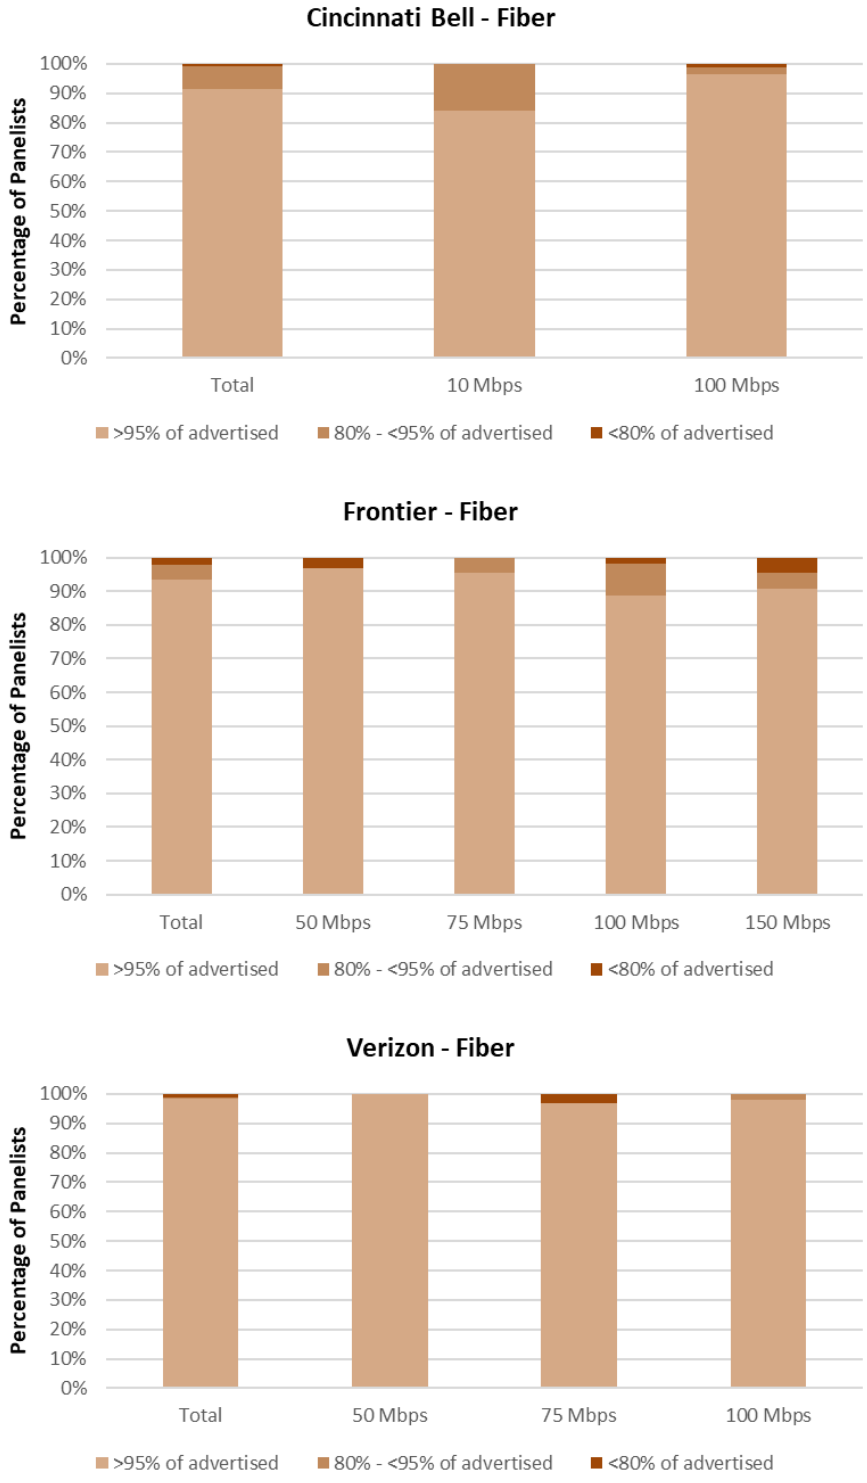

In Section 3.C above, we present complementary cumulative distributions for each ISP based on test results across all service tiers. Below, we provide tables showing selected points on these distributions by each individual ISP. In general, DSL technology showed performance between 26% and 55% of advertised speed for at least 95% of their subscribers. Among cable-based companies, the average download speeds that at least 95% of their subscribers received were between 69% and 92% of advertised rates. Fiber-based services provided a range from 73% to 98% of advertised download speeds for at least 95% of subscribers.

Table 3: Complementary cumulative distribution of the ratio of median download speed to advertised download speed by ISP

ISP	20%	50%	70%	80%	90%	95%
AT&T IPBB	123.0%	109.6%	97.1%	88.4%	79.2%	69.8%
CenturyLink	103.7%	93.1%	83.7%	76.5%	67.8%	55.0%
Cincinnati Bell Fiber	107.5%	102.5%	94.7%	92.5%	82.6%	72.8%
Cincinnati Bell DSL	92.1%	85.8%	77.2%	63.1%	40.2%	26.0%
Charter	116.7%	114.6%	111.8%	106.2%	98.8%	92.1%
Comcast	118.1%	114.9%	110.9%	107.5%	95.4%	82.2%
Cox	117.0%	111.5%	104.6%	99.6%	91.3%	81.8%
Frontier Fiber	110.7%	99.8%	97.4%	95.2%	90.7%	86.8%
Frontier DSL	107.5%	93.0%	86.5%	81.4%	66.7%	46.0%
Mediacom	132.4%	127.2%	118.2%	112.2%	86.6%	79.7%
Optimum	113.1%	104.0%	98.8%	94.8%	85.9%	69.0%
Verizon Fiber	113.5%	109.1%	104.2%	100.0%	98.9%	98.1%
Verizon DSL	131.2%	114.2%	85.7%	61.4%	55.8%	49.3%
Windstream	106.1%	100.1%	93.7%	87.4%	55.6%	48.6%

Table 4: Complementary cumulative distribution of the ratio of median upload speed to advertised upload speed by ISP

ISP	20%	50%	70%	80%	90%	95%
AT&T IPBB	137.3%	120.8%	91.7%	88.7%	78.3%	59.5%
CenturyLink	96.8%	86.8%	80.8%	74.3%	62.1%	50.7%
Cincinnati Bell Fiber	109.1%	108.7%	107.7%	105.7%	95.0%	94.7%

Cincinnati Bell DSL	95.5%	86.9%	77.7%	75.7%	73.2%	57.4%
Charter	117.0%	116.4%	114.7%	113.5%	111.4%	107.6%
Comcast	119.0%	118.6%	118.2%	117.7%	116.3%	113.9%
Cox	104.7%	104.2%	103.6%	102.5%	100.6%	96.1%
Frontier Fiber	121.8%	114.8%	103.0%	100.9%	99.6%	92.6%
Frontier DSL	147.7%	94.8%	84.6%	76.4%	63.1%	48.9%
Mediacom	116.8%	114.3%	114.1%	113.7%	112.3%	108.0%
Optimum	105.3%	103.9%	102.5%	100.7%	96.0%	88.4%
Verizon Fiber	126.1%	121.0%	118.5%	118.1%	114.5%	111.5%
Verizon DSL	118.9%	90.8%	73.8%	63.7%	59.7%	56.7%
Windstream	94.9%	83.6%	73.7%	70.9%	49.2%	28.3%

C. WEB BROWSING PERFORMANCE, BY SERVICE TIER

Below, we provide the detailed results of the webpage download time for each individual service tier of each ISP. Generally, website loading time decreased steadily with increasing tier speed until a tier speed of 15 Mbps and does not change markedly above that speed.

Chart 23.1: Average webpage download time, by ISP (1.1-5 Mbps).

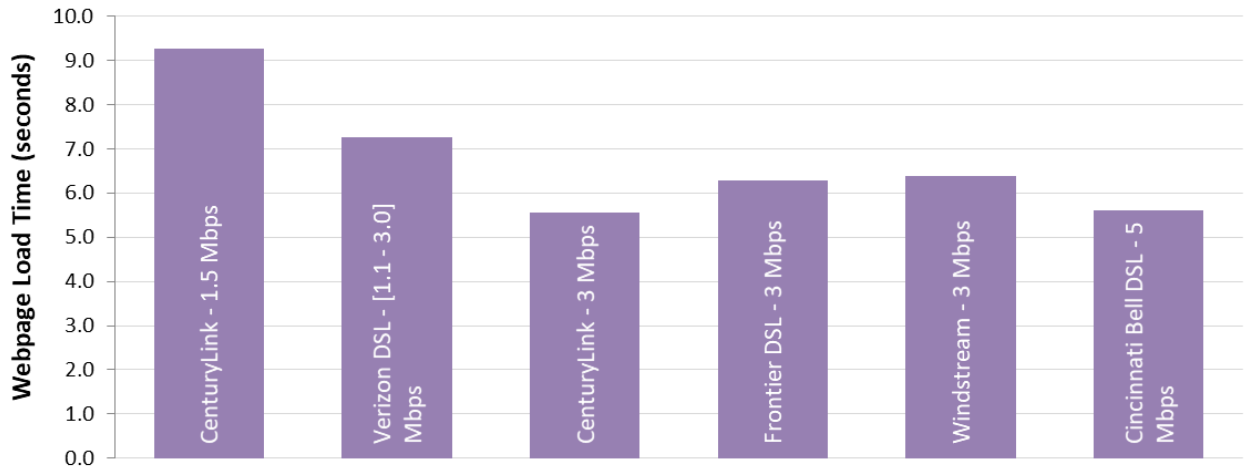

Chart 23.2: Average webpage download time, by ISP (6-10 Mbps),

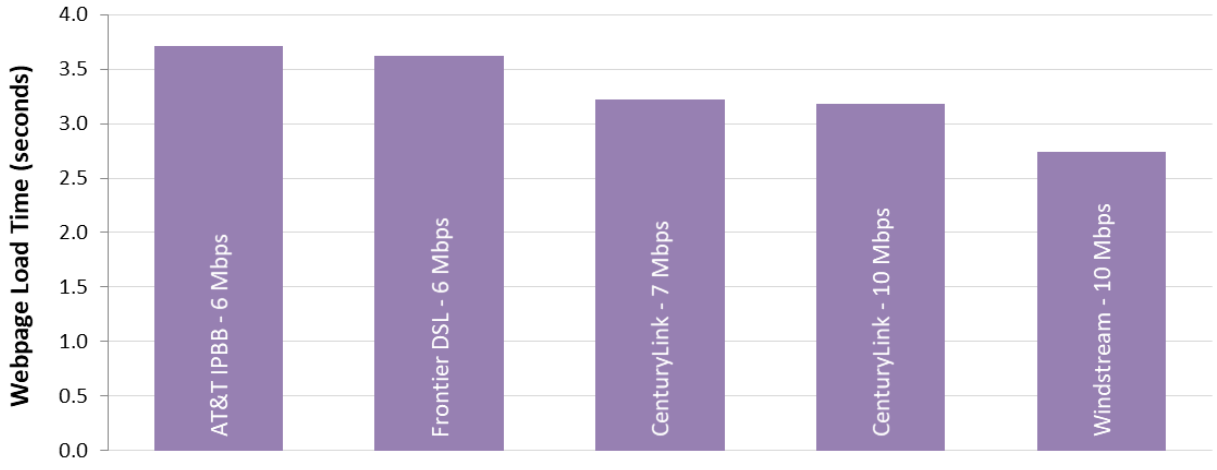

Chart 23.3: Average webpage download time, by ISP (12-20 Mbps).

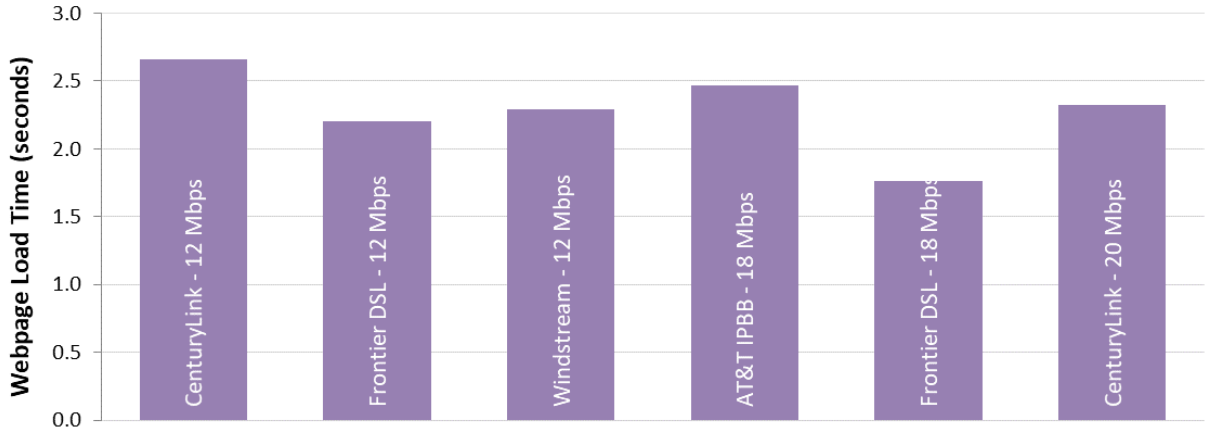

Chart 23.4: Average webpage download time, by ISP (25-30 Mbps).

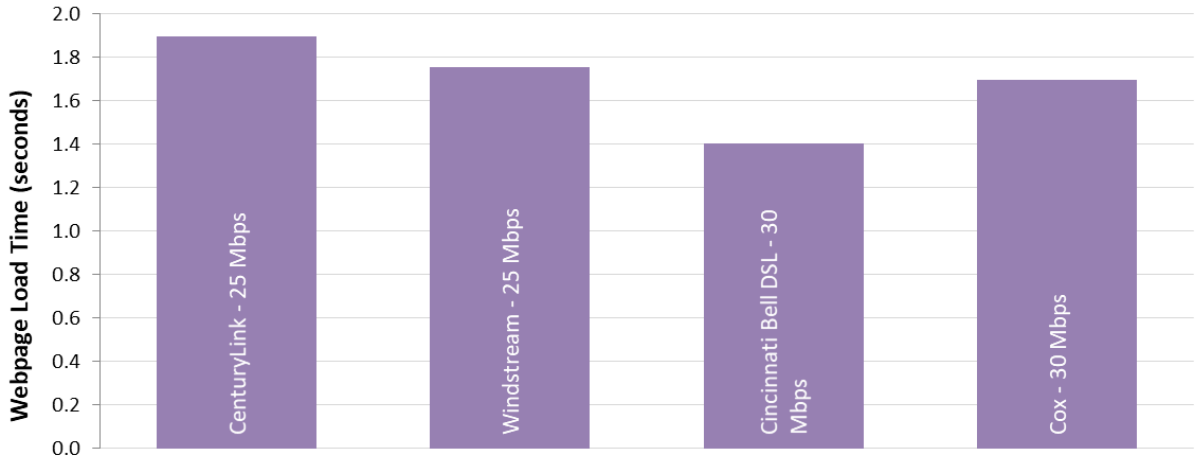

Chart 23.5: Average webpage download time, by ISP (40-50 Mbps).

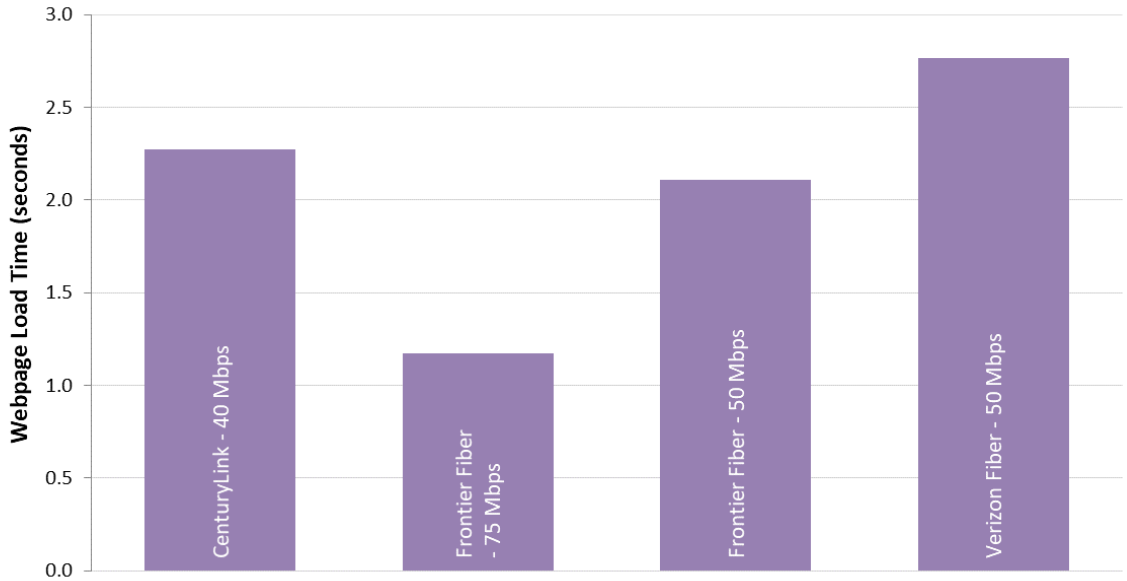

Chart 23.6: Average webpage download time, by ISP (60-75 Mbps).

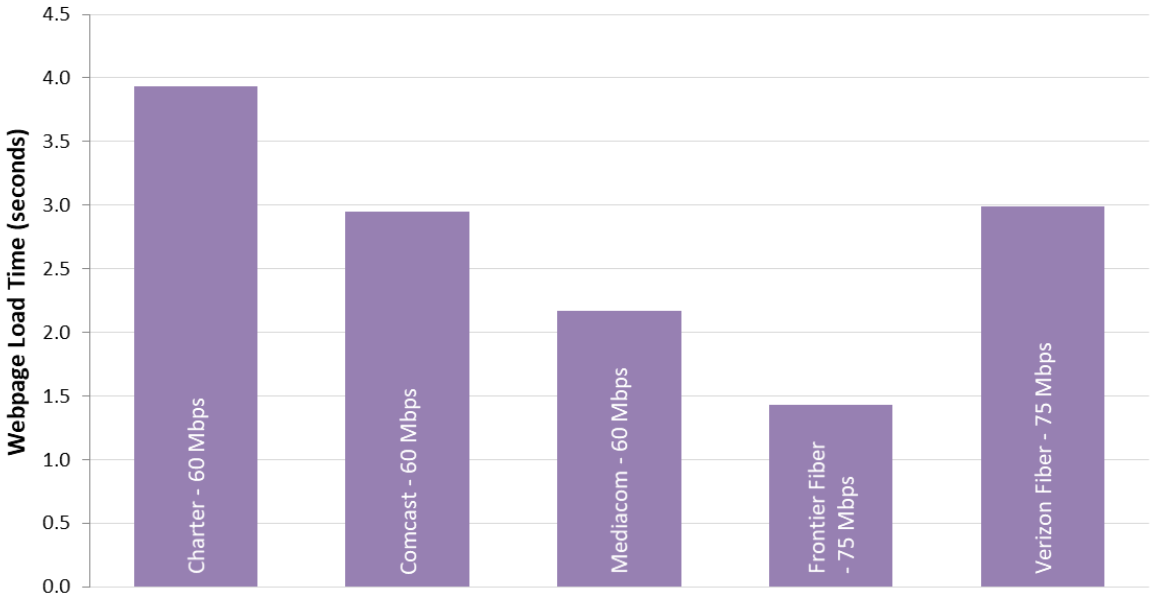

Chart 23.7: Average webpage download time, by ISP (100-150 Mbps).

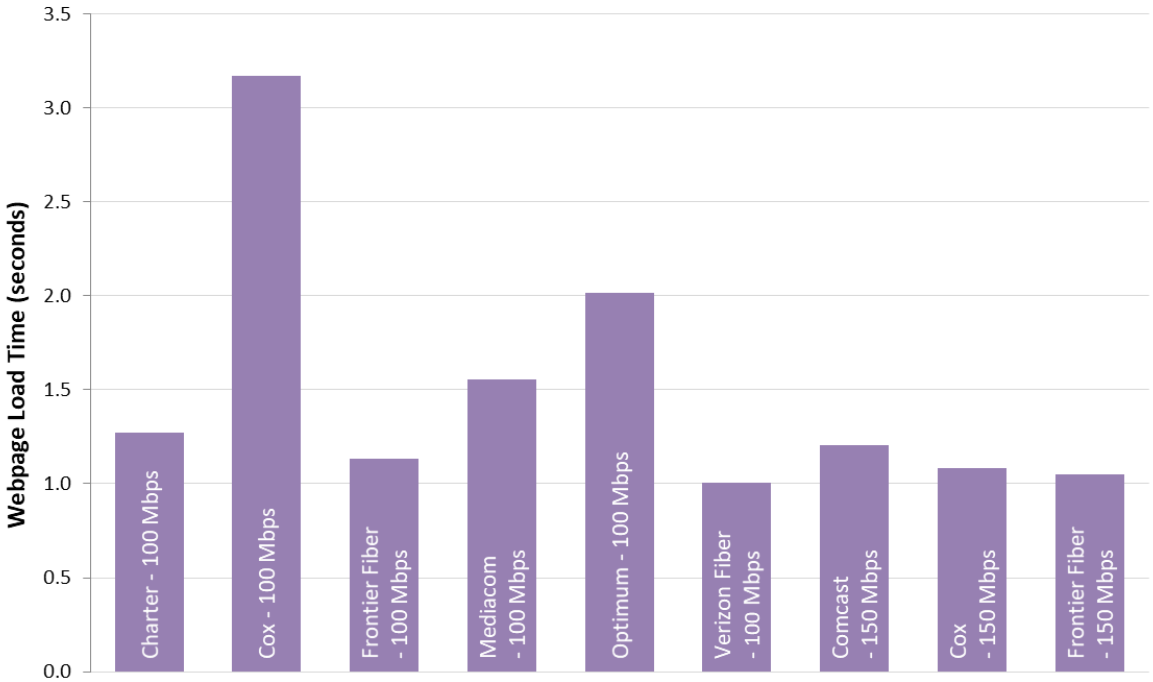

Chart 23.8: Average webpage download time, by ISP (200-300 Mbps).

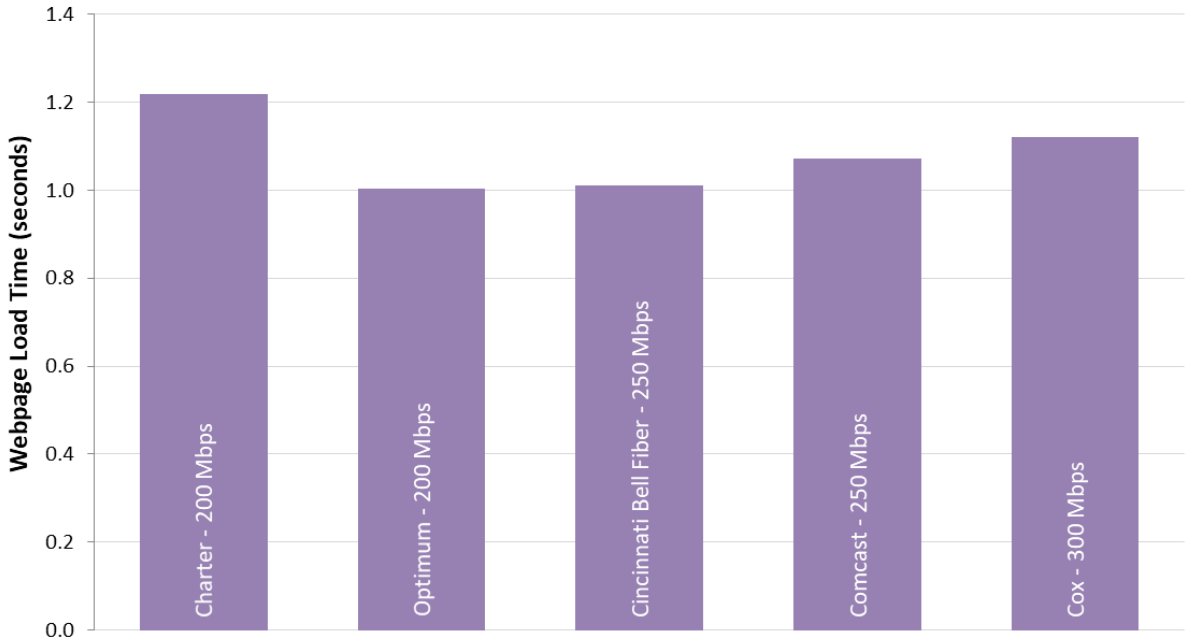

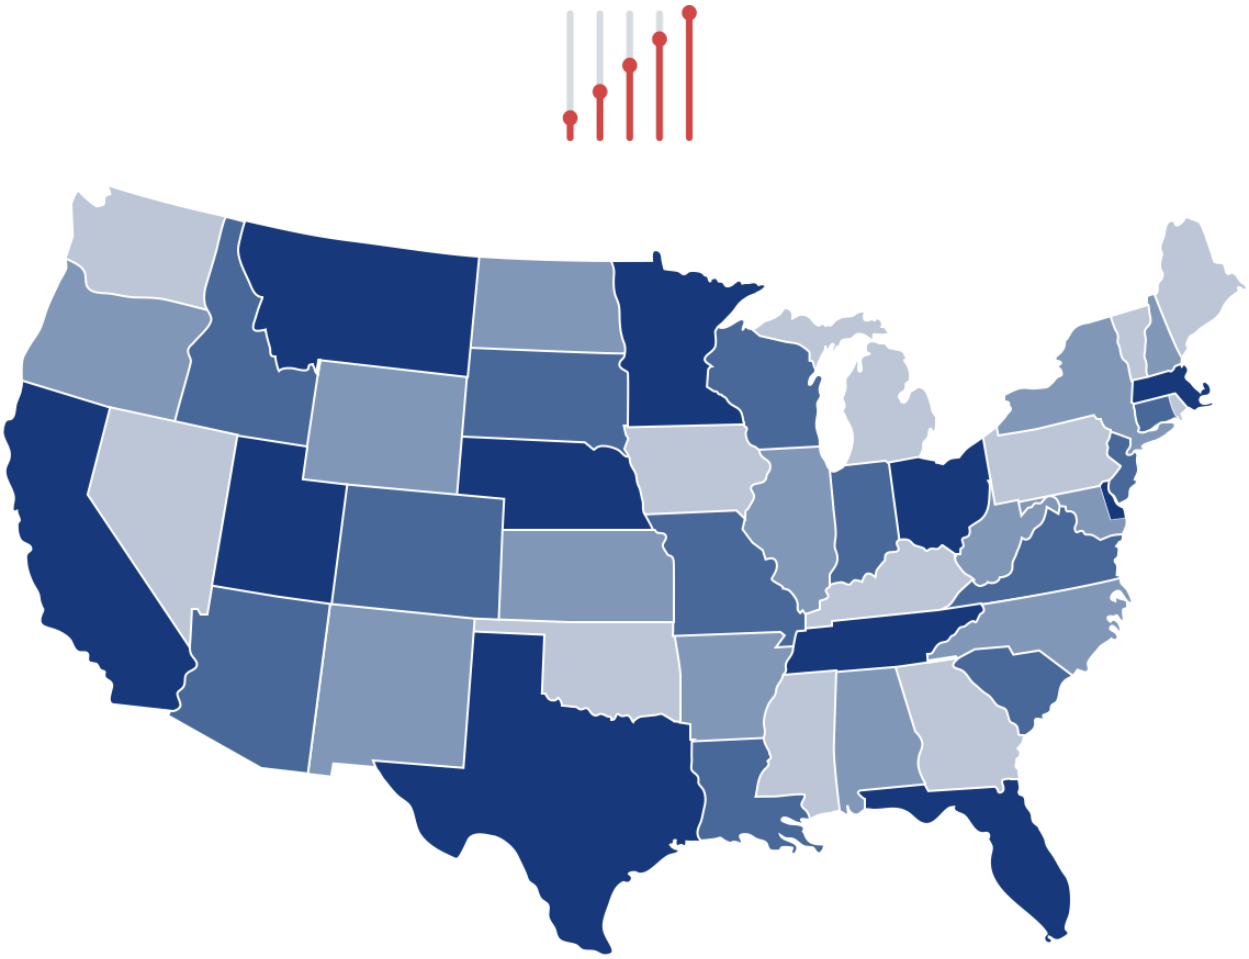

Measuring Broadband America

Technical Appendix to the Ninth MBA Report

FCC's Office of Engineering and Technology

Table of Contents

LIST OF TABLES

LIST OF FIGURES

1 - INTRODUCTION AND SUMMARY

This Appendix to the Ninth Measuring Broadband America Report,[1] a report on consumer wireline broadband performance in the United States, provides detailed technical background information on the methodology that produced the Report. It covers the process by which the panel of consumer participants was originally recruited and selected for the August 2011 MBA Report, and maintained and evolved over the last nine years. This Appendix also discusses the testing methodology used for the Report and describes how the test data was analyzed.

2 - PANEL CONSTRUCTION

This section describes the background of the study, as well as the methods employed to design the target panel, select volunteers for participation, and manage the panel to maintain the operational goals of the program.

The study aims to measure fixed broadband service performance in the United States as delivered by an Internet Service Provider (ISP) to the consumer's broadband modem. Many factors contribute to end-to-end broadband performance, only some of which are under the control of the consumer's ISP. The methodology outlined here is focused on the measurement of broadband performance within the scope of an ISP's network, and specifically focuses on measuring performance from the consumer Internet access point, or consumer gateway, to a close major Internet gateway point. The actual quality of experience seen by consumers depends on many other factors beyond the consumer's ISP, including the performance of the consumer's in-home network, transit providers, interconnection points, content distribution networks (CDN) and the infrastructure deployed by the providers of content and services. The design of the study methodology allows it to be integrated with other technical measurement approaches that focus on specific aspects of broadband performance (i.e., download speed, upload speed, latency, packet loss), and in the future, could focus on other aspects of broadband performance.

[1] The First Report (2011) was based on measurements taken in March 2011, the Second Report (2012) on measurements taken in April 2012, and the Third (2013) through this, the Ninth (2018) Reports on measurements taken in September of the year prior to the reports' release dates.

2.1 - USE OF AN ALL VOLUNTEER PANEL

During a 2008 residential broadband speed and performance test in the United Kingdom,[2] SamKnows[3] had determined that attrition rates of an all-volunteer panel was lower than a panel maintained with an incentive scheme of monthly payments. Consequently, in designing the methodology for this broadband performance study, the Commission had decided to rely entirely on volunteer consumer broadband subscribers. Volunteers are selected from a large pool of prospective participants according to a plan designed to generate a representative sample of desired consumer demographics, including geographical location, ISP, and speed tier. As an incentive for participation, volunteers are given access to a personal dashboard which allows them to monitor the performance of their broadband service. They are also provided with a measurement device referred to in the study as a "Whitebox," consisting of an off-the-shelf commodity router configured to run custom SamKnows software.[4]

2.2 - SAMPLE SIZE AND VOLUNTEER SELECTION

The Ninth MBA Report relies on data gathered from 3,192 volunteer panelists across the United States. The methodological factors and considerations that influenced the selection of the sample size and makeup include proven practices originating from the first MBA report and test period, and adaptations beyond the first period. Both are described below:

- The panel of U.S. broadband subscribers was initially drawn from a pool of over 175,000 volunteers during a recruitment campaign that ran in May 2010. Since then, to manage attrition and accommodate the evolving range of subscriber demographics (*i.e.*, tiers, technology, population), additional panelists have been recruited through email solicitations by the ISPs as well as through press releases, a web page,[5] social media outreach and blog posts.

---

[2] *See* https://files.samknows.com/~fcc_public/PM_Summer_08.pdf, (last accessed June 21, 2016).

[3] SamKnows is a company that specializes in broadband availability measurement and was retained under contract by the FCC to assist in this study. *See* http://www.samknows.com/.

[4] The Whiteboxes are named after the appearance of the first hardware implementation of the measurement agent. The Whiteboxes remain in consumer homes and continue to run the tests described in this report. Participants may remain in the measurement project as long as it continues, and may retain their Whitebox when they end their participation.

[5] https://www.measuringbroadbandamerica.com/.

- The volunteer sample was originally organized with a goal of covering major ISPs in the 48 contiguous states across five broadband technologies: DSL, cable, fiber-to-the-home, fixed terrestrial wireless, and satellite.[6]
- Target numbers for volunteers were set across the four Census Regions—Northeast, Midwest, South, and West—to help ensure geographic diversity in the volunteer panel and compensate for differences in networks across the United States.[7]
- A target plan for allocation of Whiteboxes was developed based on the market share of participating ISPs. Initial market share information was based principally on FCC Form 477[8] data filed by participating ISPs for December 2018. This data is further enhanced by the ISPs who brief SamKnows on new products and changes in subscribership numbers which may have occurred after the submission of the 477 data. Speed tiers that comprise the top 80% of a Participating ISP's subscriber base are included. This threshold ensures that we are measuring the ISP's most popular speed tiers and that it is possible to recruit sufficient panelists.
- An initial set of prospective participants was selected from volunteers who had responded directly to SamKnows as a result of media solicitations, as described in detail in Section 2.3. Where gaps existed in the sample plan, SamKnows worked with participating ISPs via email solicitations targeted at underrepresented tiers.
- Since the initial panel was created in 2011, participating ISPs have contacted random subsets of their subscribers by email to replenish cells that were falling short of their desired panel size. Additional recruitment via social media, press releases and blog posts has also taken place.

The sample plan is designed prior to the reporting period and is sent to each ISP by SamKnows. ISPs review this and respond directly to SamKnows with feedback on speed tiers that ought to be included based on the threshold criteria stated above. SamKnows will include all relevant tiers in the final report, assuming a target sample size is available. As this may not be known until after the reporting period is over, a final sample description containing all included tiers is produced and shared with the FCC and ISPs once the reporting period has finished and the data

---

[6] At the request of, and with the cooperation of the Department of Commerce and Consumer Affairs, Hawaii, we have begun to collect data from the state of Hawaii. Data from Hawaii has been included in this year's report.

[7] Although the Commission's volunteer recruitment was guided by Census Region to ensure the widest possible distribution of panelists throughout the United States, as discussed below, a sufficient number of testing devices were not deployed to enable, in every case, the evaluation of regional differences in broadband performance. The States associated with each Census Region are described in Table 4.

[8] The FCC Form 477 data collects information about broadband connections to end user locations, wired and wireless local telephone services, and interconnected Voice over Internet Protocol (VoIP) services. *See* https://www.fcc.gov/general/broadband-deployment-data-fcc-form-477 for further information.

has been processed. Test results from a total of 3,192 panelists were used in the Ninth MBA Report. This figure includes only panelists that are subscribed to the tiers that were tested as part of the sample plan.

The recruitment campaign resulted in the coverage needed to ensure balanced representation of users across the United States. Table 1 shows the number of volunteers with reporting Whiteboxes for the months of September/October 2018 listed by ISP, as well as the percentage of total volunteers subscribed to each ISP. Tables 2 and 3 shows the distributions of the Whiteboxes by State and by Region respectively. This can be compared with the percentage of subscribers per state or region.[9]

Table 1: ISPs, Sample Sizes and Percentages of Total Volunteers

ISP	Sample Size	% of Total Volunteers
AT&T	174	5.45%
CenturyLink	564	17.67%
Charter	238	7.46%
Cincinnati Bell DSL	126	3.95%
Cincinnati Bell Fiber	155	4.86%
Comcast	316	9.90%
Cox	224	7.02%
Frontier DSL	270	8.46%
Frontier Fiber	233	7.30%
Mediacom	123	3.85%
Optimum	155	4.86%
Verizon DSL	62	1.94%
Verizon Fiber	201	6.30%

[9] Subscriber data in the Ninth MBA Report is based on the FCC's Internet Access Services Report with data current to June 30, 2017. See *Internet Access Services: Status as of June 30, 2017*, Wireline Competition Bureau, Industry Analysis and Technology Division (rel. Nov. 2018), available at https://docs.fcc.gov/public/attachments/DOC-355166A1.pdf.

Windstream	351	11.0%
Total	**3,192**	**100%**

Table 2: Distribution of Whiteboxes by State

State	Total Boxes	% of Total Boxes	% of Total US Broadband
Alabama	26	0.8%	1.4%
Arizona	112	3.5%	2.1%
Arkansas	26	0.8%	0.8%
California	246	7.7%	11.6%
Colorado	94	2.9%	1.9%
Connecticut	73	2.3%	1.2%
Delaware	10	0.3%	0.3%
District of Columbia	5	0.2%	0.2%
Florida	146	4.6%	7.1%
Georgia	119	3.7%	3.1%
Hawaii	23	0.7%	0.4%
Idaho	24	0.8%	0.5%
Illinois	57	1.8%	3.9%
Indiana	42	1.3%	2.0%
Iowa	128	4.0%	1.0%
Kansas	21	0.7%	0.9%
Kentucky	130	4.1%	1.3%
Louisiana	21	0.7%	1.3%
Maine	2	0.1%	0.5%
Maryland	53	1.7%	2.0%
Massachusetts	51	1.6%	2.4%
Michigan	45	1.4%	3.1%
Minnesota	85	2.7%	1.8%
Mississippi	8	0.3%	0.7%
Missouri	63	2.0%	1.8%
Montana	5	0.2%	0.3%

Nebraska	27	0.8%	0.6%
Nevada	28	0.9%	0.9%
New Hampshire	6	0.2%	0.5%
New Jersey	111	3.5%	3.1%
New Mexico	47	1.5%	0.6%
New York	150	4.7%	6.4%
North Carolina	96	3.0%	3.2%
North Dakota	1	0.0%	0.3%
Ohio	303	9.5%	3.7%
Oklahoma	37	1.2%	1.1%
Oregon	77	2.4%	1.4%
Pennsylvania	157	4.9%	4.2%
Rhode Island	6	0.2%	0.4%
South Carolina	14	0.4%	1.5%
South Dakota	1	0.0%	0.3%
Tennessee	22	0.7%	1.9%
Texas	156	4.9%	7.6%
Utah	21	0.7%	0.8%
Vermont	1	0.0%	0.3%
Virginia	114	3.6%	2.6%
Washington	122	3.8%	2.5%
West Virginia	17	0.5%	0.6%
Wisconsin	61	1.9%	1.9%
Wyoming	2	0.1%	0.2%
	3,192		

The distribution of Whiteboxes by Census Region is found in the table on the next page.

Table 3: Distribution of Whiteboxes by Census Region

Census Region	Total Boxes	% Total Boxes	% Total U.S. Broadband Subscribers
Midwest	834	26.1%	21%
Northeast	557	17.5%	19%
South	1,000	31.3%	36%
West	801	25.1%	24%

The distribution of states associated with the four Census Regions used to define the panel strata are included in the table below.

Table 4: Panelists States Associated with Census Regions

Census Region	States
Northeast	CT MA ME NH NJ NY PA RI VT
Midwest	IA IL IN KS MI MN MO ND NE OH SD WI
South	AL AR DC DE FL GA KY LA MD MS NC OK SC TN TX VA WV
West	AK AZ CA CO HI ID MT NM NV OR UT WA WY

2.3 - PANELIST RECRUITMENT PROTOCOL

Panelists are recruited in the 2011- 2018 panels using the following method:

- Recruitment has evolved since the start of the program. At that time, (2011) several thousand volunteers were initially recruited through an initial public relations and social media campaign led by the FCC. This campaign included discussion on the FCC website and on technology blogs, as well as articles in the press. Currently volunteers are drafted with the help of a recruitment website[10] which keeps them informed about the MBA program and allows them to view MBA data on a dashboard. The composition of the panel is reviewed each year to identify any deficiencies with regard to the sample plan described above. Target demographic goals are set for volunteers based on ISP, speed tier, technology type, and region. Where the pool of volunteers falls short of the desired goal, ISPs send out email messages to their customers asking them to participate in the MBA program. The messages direct interested volunteers to contact SamKnows to request participation in the trial. The ISPs do not know which of the email recipients volunteer. In almost all cases, this ISP outreach allows the program to meet its desired demographic targets.

The mix of panelists recruited using the above methodologies varies by ISP.

A multi-mode strategy was used to qualify volunteers for the 2018 testing period. The key stages of this process were as follows:

1. Volunteers were directed to complete an online form which provided information on the study and required volunteers to submit a small amount of information.
2. Volunteers were selected from respondents to this follow-up email based on the target requirements of the panel. Selected volunteers were then asked to agree to the *User Terms and Conditions* that outlined the permissions to be granted by the volunteer in key areas such as privacy.[11]
3. From among the volunteers who agreed to the User Terms and Conditions, SamKnows selected the panel of participants,[12] each of whom received a Whitebox for self-installation. SamKnows provided full support during the Whitebox installation phase.

The graphic in Figure 1 illustrates the study recruitment methodology.

[10] The Measuring Broadband America recruitment website is: https://www.measuringbroadbandamerica.com/.

[11] The *User Terms and Conditions* is found in the Reference Documents at the end of this Appendix.

[12] Over 23,000 Whiteboxes have been shipped to targeted volunteers since 2011, of which 5,855 were online and reporting data used in the Ninth Report from the months of September/October 2018.

Figure 1: Panelist Recruitment Protocol

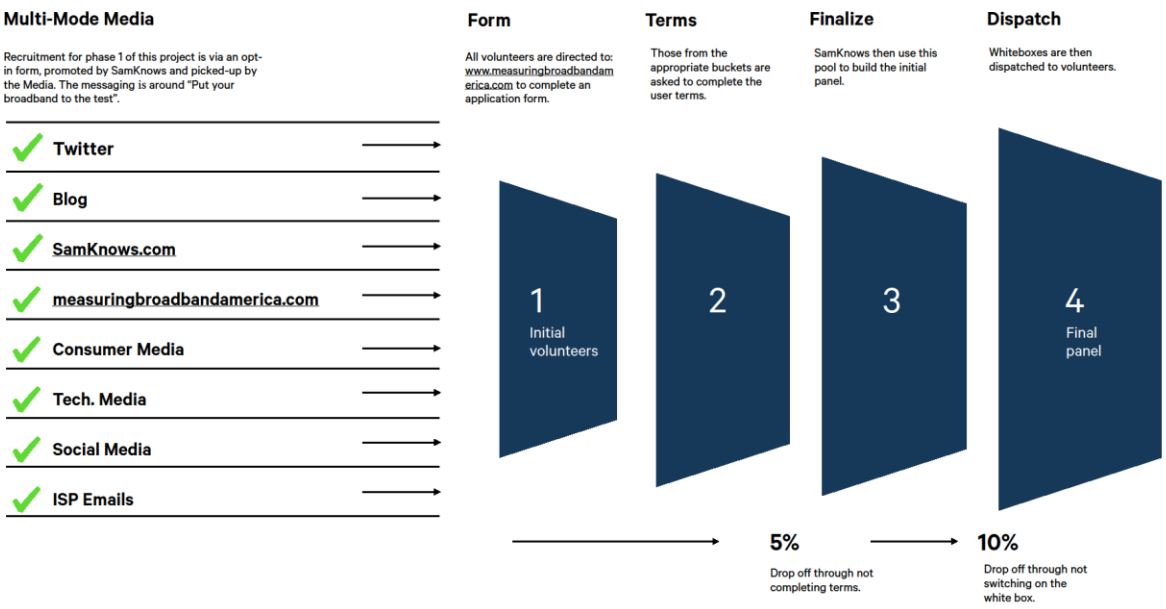

2.4 - VALIDATION OF VOLUNTEERS' SERVICE TIER

The methodology employed in this study included verifying each panelist's service tier and ISP against the customer records of participating ISPs.[13] Initial throughput tests were used to confirm reported speeds.

The broadband service tier reported by each panelist was validated as follows:

- When the panelist installed the Whitebox, the device automatically ran an IP address test to check that the ISP identified by the volunteer was correct.
- The Whitebox also ran an initial test which flooded each panelist's connection in order to accurately detect the throughput speed when their deployed Whitebox connected to a test node.
- Each ISP was asked to confirm the broadband service tier reported by each selected panelist.
- SamKnows then took the validated speed tier information that was provided by the ISPs and compared this to both the panelist-provided information, and the actual test results obtained, in order to ensure accurate tier validation.

SamKnows manually completed the following four steps for each panelist:

- Verified that the IP address was in a valid range for those served by the ISP.
- Reviewed data for each panelist and removed data where speed changes such as tier upgrade or downgrade appeared to have occurred, either due to a service change on the part of the consumer or a network change on the part of the ISP.
- Identified panelists whose throughput appeared inconsistent with the provisioned service tier. Such anomalies were re-certified with the consumer's ISP.[14]
- Verified that the resulting downstream-upstream test results corresponded to the ISP-provided speed tiers, and updated accordingly if required.

---

[13] Past FCC studies found that a high rate of consumers could not reliably report information about their broadband service, and the validation of subscriber information ensured the accuracy of expected speed and other subscription details against which observed performance was measured. *See* John Horrigan and Ellen Satterwhite, *Americans' Perspectives on Online Connection Speeds for Home and Mobile Devices*, 1 (FCC 2010), available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-298516A1.doc (finding that 80 percent of broadband consumers did not know what speed they had purchased).

[14] For example, when a panelist's upload or download speed was observed to be significantly higher than that of the rest of the tier, it could be inferred that a mischaracterization of the panelist's service tier had occurred. Such anomalies, when not resolved in cooperation with the service provider, were excluded from the Ninth Report, but will be included in the raw bulk data set.

Of the more than 23,000 Whiteboxes that were shipped to panelists since 2011, 5,855[15] units reported sufficient data in September/October 2018, with the participating ISPs validating 4,766 for the reporting period. Of the validated units, 16 percent were reallocated to a different tier following the steps listed above. A total of 3,192 validated units were part of download or upload tiers included in the sample plan and were ultimately included in this report.

A total of 2,663 boxes were excluded for the following reasons:

- 1,241 belonged to users subscribed to plans that were not included in this study
- 546 belonged to users whose details could not be successfully validated by the ISP
- 439 Whiteboxes were legacy models that could not fully support the plan speeds
- 357 were excluded due to other legacy equipment, such as modems or ethernet links that could not fully support subscribed speeds
- 45 were excluded as download and/or upload test speeds were significantly different from the product validated by the ISP
- 33 belonged to either ISP employees, or were connected to non-residential plans
- 2 were excluded due to known connection issues reported by the ISPs

2.5 - PROTECTION OF VOLUNTEERS' PRIVACY

Protecting the panelists' privacy is a major concern for this program. The panel was comprised entirely of volunteers who knowingly and explicitly opted in to the testing program. For audit purposes, we retain the correspondence with panelists documenting their opt-in.

All personal data was processed in conformity with relevant U.S. law and in accordance with policies developed to govern the conduct of the parties handling the data. The data were processed solely for the purposes of this study and are presented here and in all online data sets with all personally identifiable information (PII) removed.

A set of materials was created both to inform each panelist regarding the details of the trial, and to gain the explicit consent of each panelist to obtain subscription data from the participating ISPs. These documents were reviewed by the Office of General Counsel of the FCC and the participating ISPs and other stakeholders involved in the study.

[15] This figure represents the total number of boxes reporting during September/October 2018, the month chosen for the Ninth Report. Shipment of boxes continued in succeeding months and these results will be included in the raw bulk data set.

3 - BROADBAND PERFORMANCE TESTING METHODOLOGY

This section describes the system architecture and network programming features of the tests, and other technical aspects of the methods employed to measure broadband performance during this study.

3.1 - RATIONALE FOR HARDWARE-BASED MEASUREMENT APPROACH

Either a hardware or software approach can be used to measure broadband performance. Software approaches are by far the most common and allow for measurements to easily and cost-effectively include a very large sample size. Web-based speed tests fall into this category and typically use Flash applets, Java applets or JavaScript that execute within the user's web browser. These clients download content from remote web servers and measure the throughput. Some web-based performance tests also measure upload speed or round-trip latency.

Other, less common, software-based approaches to performance measurement install applications on the user's computer. These applications run tests periodically while the computer is on.

All software solutions implemented on a consumer's computer, smart phone, or other device connected to the Internet suffer from the following disadvantages:

- The software and computing platform running the software may not be capable of reliably recording the higher speed service tiers currently available.
- The software typically cannot know if other devices on the home network are accessing the Internet when the measurements are being taken. The lack of awareness as to other, non-measurement related network activity can produce inconsistent and misleading measurement data.
- Software measurements may be affected by the performance, quality and configuration of the device.
- Potential bottlenecks, such as Wi-Fi networks and other in-home networks, are generally not accounted for and may result in unreliable data.
- If the device hosting the software uses in-home WIFI access to fixed broadband service, differing locations in the home may impact measurements.
- The tests can only run when the computer is turned on, limiting the ability to provide a 24-hour profile.

- If software tests are performed manually, panelists might only run tests when they experience problems and thus bias the results.

In contrast, the hardware approach used in the MBA program requires the placement of the previously described Whitebox inside the user's home, directly connected to the consumer's service interconnection device (router), via Ethernet cable. The measurement device therefore directly accesses fixed Internet service to the home over this dedicated interface and periodically runs tests to remote targets over the Internet. The use of hardware devices avoids the disadvantages listed earlier with the software approach. However, hardware approaches are much more expensive than the software alternative, are thus more constrained in the achievable panel size, and require correct installation of the device by the consumer or a third party. This is still subject to unintentional errors due to misconfigurations, *i.e.*, connecting the Whitebox incorrectly but these can often be detected in the validation process that follows installation. The FCC chose the hardware approach since its advantages far outweigh these disadvantages.

3.2 - DESIGN OBJECTIVES AND TECHNICAL APPROACH

For this test of broadband performance, as in previous Reports, the FCC used design principles that were previously developed by SamKnows in conjunction with their study of broadband performance in the U.K. The design principles comprise 17 technical objectives:

Table 5: Design Objectives and Methods

#	Technical Objectives	Methodological Accommodations
1	The Whitebox measurement process must not change during the monitoring period.	The Whitebox measurement process is designed to provide automated and consistent monitoring throughout the measurement period.
2	Must be accurate and reliable.	The hardware solution provides a uniform and consistent measurement of data across a broad range of participants.
3	Must not interrupt or unduly degrade the consumer's use of the broadband connection.	The volume of data produced by tests is controlled to avoid interfering with panelists' overall broadband experience, and tests only execute when consumer is not making heavy use of the connection.
4	Must not allow collected data to be distorted by any use of the broadband connection by other applications on the host PC and other devices in the home.	The hardware solution is designed not to interfere with the host PC and is not dependent on that PC.

5	Must not rely on the knowledge, skills and participation of the consumer for its ongoing operation once installed.	The Whitebox is "plug-and-play." Instructions are graphics-based and the installation process has been substantially field tested. Contacts for support are also provided and the outreach once a Whitebox has been dispatched and activated.
6	Must not collect data that might be deemed to be personal to the consumer without consent.	The data collection process is explained in plain language and consumers are asked for their consent regarding the use of their personal data as defined by any relevant data protection legislation.
7	Must be easy for a consumer to completely remove any hardware and/or software components if they do not wish to continue with the MBA program.	Whiteboxes can be disconnected at any time from the home network. As soon as the Whitebox is reconnected the reporting is resumed as before.
8	Must be compatible with a wide range of DSL, cable, satellite and fiber-to-the-home modems.	Whiteboxes can be connected to all modem types commonly used to support broadband services in the U.S., either in a routing or bridging mode, depending on the model.
9	Where applicable, must be compatible with a range of computer operating systems, including, without limitation, Windows XP, Windows Vista, Windows 7, Mac OS and Linux.	Whiteboxes are independent of the PC operating system and therefore able to provide testing with all devices regardless of operating system.
10	Must not expose the volunteer's home network to increased security risk, *i.e.*, it should not be susceptible to viruses, and should not degrade the effectiveness of the user's existing firewalls, antivirus and spyware software.	The custom software in the Whitebox is hardened for security and cannot be accessed without credentials only available to SamKnows. Most user firewalls, antivirus and spyware systems are PC-based. The Whitebox is plugged in to the broadband connection "before" the PC. Its activity is transparent and does not interfere with those protections.
11	Must be upgradeable remotely if it contains any software or firmware components.	The Whitebox can be completely controlled remotely for updates without involvement of the consumer, providing the Whitebox is switched on and connected.
12	Must identify when a user changes broadband provider or package (*e.g.*, by a reverse look up of the consumer's IP address to check provider, and by capturing changes in modem connection speed to identify changes in package).	Ensures regular data pool monitoring for changes in speed, ISP, IP address or performance, and flags when a panelist should notify and confirm any change to their broadband service since the last test execution.

13	Must permit, in the event of a merger between ISPs, separate analysis of the customers of each of the merged ISP's predecessors.	Data are stored based on the ISP of the panelist, and therefore can be analyzed by individual ISP or as an aggregated dataset.
14	Must identify if the consumer's computer is being used on a number of different fixed networks (*e.g.*, if it is a laptop).	The Whiteboxes are broadband dependent, not PC or laptop dependent.
15	Must identify when a specific household stops providing data.	The Whitebox needs to be connected and switched on to push data. If it is switched off or disconnected its absence is detected at the next data push process.
16	Must not require an amount of data to be downloaded which may materially impact any data limits, usage policy, or traffic shaping applicable to the broadband service.	The data volume generated by the information collected does not exceed any policies set by ISPs. Panelists with bandwidth restrictions can have their tests set accordingly.
17	Must limit the possibility for ISPs to identify the broadband connections which form their panel and therefore potentially "game" the data by providing different quality of service to the panel members and to the wider customer base.	ISPs signed a Code of Conduct[16] to protect against gaming test results. While the identity of each panelist was made known to the ISP as part of the speed tier validation process, the actual Unit ID for the associated Whitebox was not released to the ISP so specific test results were not directly assignable against a specific panelist. Moreover, most ISPs had hundreds, and some had more than 1,000, participating subscribers spread throughout their service territory, making it difficult to improve service for participating subscribers without improving service for all subscribers.

[16] Signatories to the Code of Conduct are: AT&T, CenturyLink, Charter, Cincinnati Bell, Comcast, Cox, Frontier, Hughes, Level3, Measurement Lab, Mediacom, NCTA, Optimum, Time Warner Cable, Verizon, ViaSat, and Windstream. A copy of the Code of Conduct is included as a Reference Document attached to this Appendix.

3.3 - TESTING ARCHITECTURE

Overview of Testing Architecture

As illustrated in Figure 2, the performance monitoring system comprises a distributed network of Whiteboxes in the homes of members of the volunteer consumer panel. The Whiteboxes are controlled by a cluster of servers, which hosts the test scheduler and the reporting database. The data was collated on the reporting platform and accessed via a reporting interface[17] and secure FTP site. The system also included a series of speed-test servers, which the Whiteboxes called upon according to the test schedule.

Figure 2: Testing Architecture

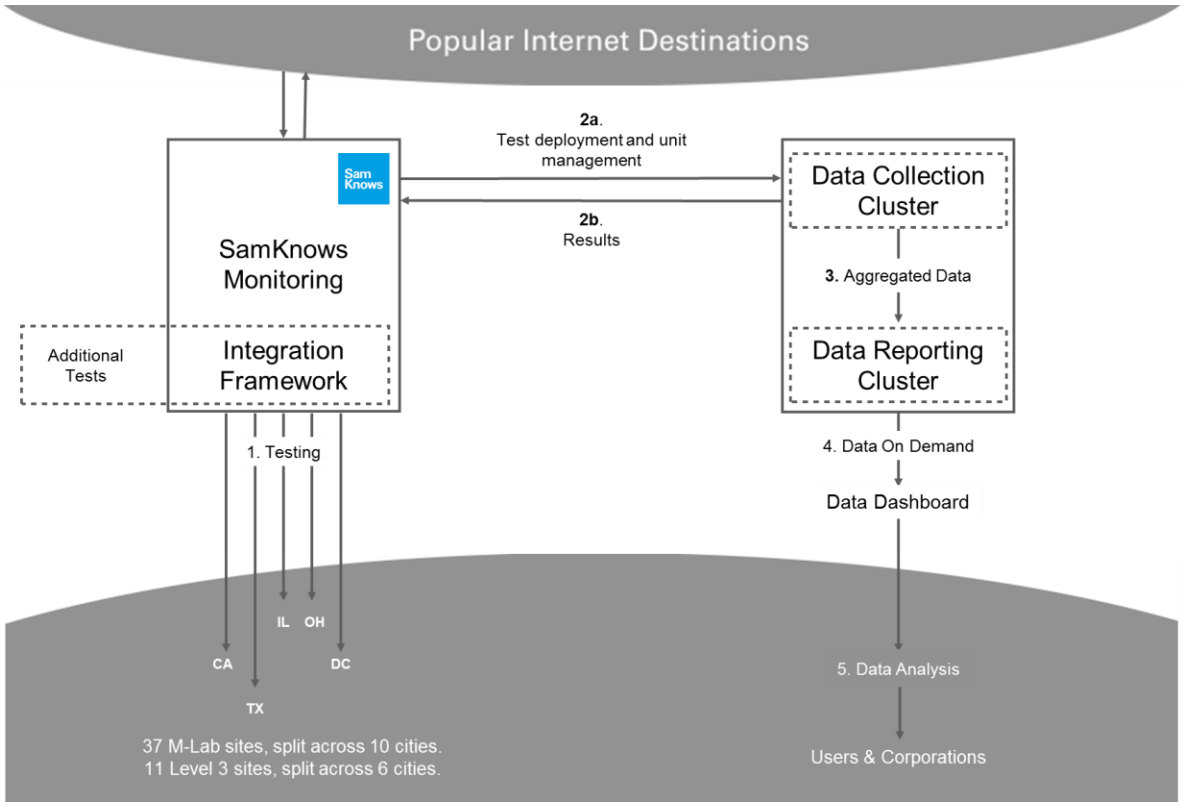

[17] Each reporting interface included a data dashboard for the consumer volunteers, which provided performance metrics associated with their Whitebox.

Approach to Testing and Measurement

Any network monitoring system needs to be capable of monitoring and executing tests 24 hours a day, seven days a week. Similar to the method used by the television audience measurement industry, each panelist is equipped with a Whitebox, which is self-installed by each panelist and conducts the performance measurements. Since 2011, the project has used three different hardware platforms, described below. The software on each of the Whiteboxes was programmed to execute a series of tests designed to measure key performance indicators (KPIs) of a broadband connection. The tests comprise a suite of applications, written by SamKnows in the programming language C, which were rigorously tested by the ISPs and other stakeholders. The Ninth Report incorporates data from all three types of Whiteboxes and we use the term Whitebox generically. Testing has found that they produce results that are indistinguishable.

During the initial testing period in 2011, the Whitebox provided used hardware manufactured by NETGEAR, Inc. (NETGEAR) and operated as a broadband router. It was intended to replace the panelist's existing router and be directly connected to the cable or DSL modem, ensuring that tests could be run at any time the network was connected and powered, even if all home computers were switched off. Firmware for the Whitebox routers was developed by SamKnows with the cooperation of NETGEAR. In addition to running the latest versions of the SamKnows testing software, the routers retained all of the native functionality of the NETGEAR consumer router.

Following the NETGEAR Whitebox new models were introduced starting with the 2012 testing period. These versions were based upon hardware produced by TP-Link and then later manufactured by SamKnows and operate as a bridge rather than as a router. It connects to the customer's existing router, rather than replacing it, and all hardwired home devices connect to LAN ports on the TP-Link Whitebox. The TP-Link Whitebox / SamKnows Whitebox passively monitors wireless network activity in order to determine when the network is active and defer measurements. It runs a modified version of OpenWrt, an open source router platform based on Linux. All Whiteboxes deployed since 2012 use the TP-Link or SamKnows hardware.

SamKnows Whiteboxes (Whitebox 8.0), introduced in August 2016, have been shown to provide accurate information about broadband connections with throughput rates of up to 1 Gbps.

Home Deployment of the NETGEAR Based Whitebox

This study was initiated by using existing NETGEAR firmware, and all of its features were intended to allow panelists to replace their existing routers with the Whitebox. If the panelist did not have an existing router and used only a modem, they were asked to install the Whitebox according to the usual NETGEAR instructions.

However, this architecture could not easily accommodate scenarios where the panelist had a combined modem/router supplied by their ISP that had specific features that the Whitebox could not provide. For example, some Verizon FiOS gateways connect via a MoCA (Multimedia over Cable) interface and AT&T IPBB gateways provide U-Verse specific features, such as IPTV.

In these cases, the Whitebox was connected to the existing router/gateway and all home devices plugged into the Whitebox. In order to prevent a double-NAT configuration, in which multiple routers on the same network perform network address translation (NAT) and make access to the SamKnows router difficult, the Whitebox was set to dynamically switch to operate as a transparent Ethernet bridge when deployed in these scenarios. All consumer configurations were evaluated and tested by participating ISPs to confirm their suitability.[18]

Home Deployment of the TP-Link Based Whitebox

The TP-Link-based Whitebox, which operates as a bridge, was introduced in response to the increased deployment of integrated modem/gateway devices. To use the TP-Link-based Whitebox, panelists are required to have an existing router. Custom instructions guided these panelists to connect the Whitebox to their existing router and then connect all of their home devices to the Whitebox. This allows the Whitebox to measure traffic volumes from wired devices in the home and defer tests accordingly. As an Ethernet bridge, the Whitebox does not provide services such as network address translation (NAT) or DHCP.

Home Deployment of the SamKnows Whitebox 8.0

The Whitebox 8.0 was manufactured by SamKnows and deployed starting in August 2016. Like the TP-Link device, this Whitebox works as a bridge, rather than a router, and operates in a similar manner. Unlike the NETGEAR and TP-Link hardware, it can handle bandwidths of up to 1 Gbps.

Internet Activity Detection

No tests are performed if the Whiteboxes detect wired or wireless traffic beyond a defined bandwidth threshold. This ensures both that testing does not interfere with consumer use of

[18] The use of legacy equipment has the potential to impede some panelists from receiving the provisioned speed from their ISP, and this impact is captured by the survey.

their Internet service and that any such use does not interfere with testing or invalidate test results.

Panelists were not asked to change their wireless network configurations. Since the TP-Link Whiteboxes and Whitebox 8.0 attach to the panelist's router that may contain a built-in wireless (Wi-Fi) access point, these devices measure the strongest wireless signal. Since they only count packets, they do not need access to the Wi-Fi encryption keys and do not inspect packet content.

Test Nodes (Off-Net and On-Net)

For the tests in this study, SamKnows employed fifty-four core measurement servers as test nodes that were distributed geographically across eleven locations, outside the network boundaries of the participating ISPs. These so-called off-net measurement points were supplemented by additional measurement points located within the networks of some of the ISPs participating in this study, called on-net servers. The core measurement servers were used to measure consumers' broadband performance between the Whitebox and an available reference point that was closest in roundtrip time to the consumer's network address. The distribution of off-net primary reference points operated by M-Lab and Level 3 and on-net secondary reference points operated by broadband providers provided additional validity checks and insight into broadband service performance within an ISP's network. In total, the following 133 measurement servers were deployed for the Ninth Report:

Table 6: Overall Number of Testing Servers

Operated By	Number of Servers
AT&T	6
CenturyLink (inc Qwest)	7
Charter (inc TWC)	4
Comcast	36
Cox	2
Frontier	5
Hawaiian Telecom	1

Level 3 (off-net)	11
M-Lab (off-net)	45
Mediacom	1
Optimum	1
Time Warner Cable (now part of Charter)	7
Uhnet (Hawaii)	1
Verizon	2
Windstream	4

Test Node Locations

Off-Net Test Nodes

The M-Lab test nodes were located in the following major U.S. Internet peering locations:

- New York City, New York (three locations)
- Chicago, Illinois (five locations)
- Atlanta, Georgia (four locations)
- Miami, Florida (five locations)
- Washington, DC (four locations)
- Mountain View, California (five locations)
- Seattle, Washington (five locations)
- Los Angeles, California (four locations)
- Dallas, Texas (four locations)
- Denver, Colorado (four locations)

The Level 3 nodes were located in the following major U.S. Internet peering locations:

- Chicago, Illinois (two locations)
- Dallas, Texas
- New York City, New York (two locations)
- San Jose, California (two locations)
- Washington D.C. (two locations)
- Los Angeles, California (two locations)

On-Net Test Nodes

In addition to off-net nodes, some ISPs deployed their own on-net servers to cross-check the results provided by off-net nodes. Whiteboxes were instructed to test against the off-net M-Lab and Level 3 nodes and the on-net ISP nodes, when available.

The following ISPs provided on-net test nodes:

- AT&T
- CenturyLink[19]
- Charter[20]
- Cincinnati Bell
- Comcast
- Cox
- Frontier
- Mediacom
- Optimum
- Verizon
- Windstream

The same suite of tests was scheduled for these on-net nodes as for the off-net nodes and the same server software developed by SamKnows was used regardless of whether the Whitebox

[19] QWest was reported separately from Centurylink in reports prior to 2016. The entities completed merging their test infrastructure in 2016.

[20] Time Warner Cable was reported separately from Charter in reports prior to the Eighth report. The entities completed merging their test infrastructure in early 2018.

was interacting with on-net or off-net nodes. Off-net test nodes are continually monitored for load and congestion.

While these on-net test nodes were included in the testing, the results from these tests were used as a control set; the results presented in the Report are based only on tests performed using off-net nodes. Results from both on-net and off-net nodes are included in the raw bulk data set that will be released to the public.

Test Node Selection

Each Whitebox fetches a complete list of off-net test nodes and on-net test nodes hosted by the serving ISP from a SamKnows server and measures the round-trip time to each. This list of test servers is loaded at startup and refreshed daily. It then selects the on-net and off-net test nodes with lowest round trip time to test against. The selected nodes may not be the geographically closest node.

Technical details for the minimum requirements for hardware and software, connectivity, and systems and network management are available in the **5.3 - Test Node Briefing** provided in the Reference Document section of this Technical Appendix.

3.4 - TESTS METHODOLOGY

Each deployed Whitebox performs the following tests.[21] All tests are conducted with both the on-net and off-net servers except as noted, and are described in more detail in the next section.

Table 7: List of Tests Performed by SamKnows[22]

Metric	Primary Metric(s)
Download speed	Throughput in Megabits per second (Mbps) utilizing three concurrent TCP connections
Upload speed	Throughput in Mbps utilizing three concurrent TCP connections
Web browsing	Total page fetch time and all its embedded resources from a popular website
UDP latency	Average round trip time of a series of randomly transmitted UDP packets distributed over a long timeframe
UDP packet loss	Fraction of UDP packets lost from UDP latency test
Voice over IP	Upstream packet loss, downstream packet loss, upstream jitter, downstream jitter, round trip latency
DNS resolution	Time taken for the ISP's recursive DNS resolver to return an A record[23] for a popular website domain name
DNS failures	Percentage of DNS requests performed in the DNS resolution test that failed
ICMP latency	Round trip time of five evenly spaced ICMP packets
ICMP packet loss	Percentage of packets lost in the ICMP latency test
UDP Latency under load	Average round trip time for a series of evenly spaced UDP packets sent during downstream/upstream sustained tests
Lightweight download speed	Downstream throughput in Megabits per second (Mbps) utilizing a burst of UDP datagrams
Lightweight upload speed	Upstream throughput in Megabits per second (Mbps) utilizing a burst of UDP datagrams

[21] Specific questions on test procedures may be addressed to team@samknows.com.

[22] Other tests may be run on the MBA panel; this list outlines the published tests in the report.

[23] An "A record" is the numeric IP address associated with a domain address such as https://www.fcc.gov.

3.5 - TEST DESCRIPTIONS

The following sub-sections detail the methodology used for the individual tests. As noted earlier, all tests only measure the performance of the part of the network between the Whitebox and the target (*i.e.*, a test node). In particular, the VoIP tests can only approximate the behavior of real applications and do not reflect the impact of specific consumer hardware, software, media codecs, bandwidth adjustment algorithms, Internet backbones and in-home networks.

Download Speed and Upload Speed

These tests measure the download and upload throughput by performing multiple simultaneous HTTP GET and HTTP POST requests to a target test node.

Binary, non-zero content—herein referred to as the payload—is hosted on a web server on the target test node. The test operates for a fixed duration of 10 seconds. It records the average throughput achieved during this 10 second period. The client attempts to download as much of the payload as possible for the duration of the test.

The test uses three concurrent TCP connections (and therefore three concurrent HTTP requests) to ensure that the line is saturated. Each connection used in the test counts the numbers of bytes transferred and is sampled periodically by a controlling thread. The sum of these counters (a value in bytes) divided by the time elapsed (in microseconds) and converted to Mbps is taken as the total throughput of the user's broadband service.

Factors such as TCP slow start and congestion are taken into account by repeatedly transferring small chunks (256 kilobytes, or kB) of the target payload before the real testing begins. This "warm-up" period is completed when three consecutive chunks are transferred at within 10 percent of the speed of one another. All three connections are required to have completed the warm-up period before the timed testing begins. The warm-up period is excluded from the measurement results.

Downloaded content is discarded as soon as it is received, and is not written to the file system. Uploaded content is generated and streamed on the fly from a random source.

The test is performed for both IPv4 and IPv6, where available, but only IPv4 results are reported.

Web Browsing

The test records the averaged time taken to sequentially download the HTML and referenced resources for the home page of each of the target websites, the number of bytes transferred, and the calculated rate per second. The primary measure for this test is the total time taken to download the HTML front page for each web site and all associated images, JavaScript, and stylesheet resources. This test does not measure against the centralized testing nodes; instead

it tests against actual websites, ensuring that the effects of content distribution networks and other performance enhancing factors can be taken into account.

Each Whitebox tests against the following nine websites:[24]

- http://www.edition.cnn.com/
- http://www.bing.com/
- http://www.msn.com/
- http://www.bbc.com/
- http://www.apple.com/
- http://www.ebay.com/
- http://www.m.imdb.com/help/
- http://www.google.com/policies/

The results include the time needed for DNS resolution. The test uses up to eight concurrent TCP connections to fetch resources from targets. The test pools TCP connections and utilizes persistent connections where the remote HTTP server supports them.

The client advertises the user agent as Microsoft Internet Explorer 10. Each website is tested in sequence and the results summed and reported across all sites.

UDP Latency and Packet Loss

These tests measure the round-trip time of small UDP packets between the Whitebox and a target test node.

Each packet consists of an 8-byte sequence number and an 8-byte timestamp. If a response packet is not received within three seconds of sending, it is treated as being lost. The test records the number of packets sent each hour, the average round trip time and the total number of packets lost. The test computes the summarized minimum, maximum, standard deviation and mean from the lowest 99 percent of results, effectively trimming the top (*i.e.*, slowest) 1 percent of outliers.

The test operates continuously in the background. It is configured to randomly distribute the sending of the requests over a fixed interval of one hour (using a Poisson distribution), reporting the summarized results once the interval has elapsed. Approximately two thousand packets are sent within a one-hour period, with fewer packets sent if the line is not idle.

This test is started when the Whitebox boots and runs permanently as a background test. The test is performed for both IPv4 and IPv6, where available, but only IPv4 results are reported.

[24] These websites were chosen based on a list by Alexa, http://www.alexa.com/, of the top twenty websites in October 2010.

Voice over IP

The Voice over IP (VoIP) test operates over UDP and utilizes bidirectional traffic, as is typical for voice calls.

The Whitebox handshakes with the server, and each initiates a UDP stream with the other. The test uses a 64 kbps stream with the same characteristics and properties (*i.e.*, packet sizes, delays, bitrate) as the G.711 codec. 160 byte packets are used. The test measures jitter, delay, and loss.

Jitter is calculated using the Packet Delay Variation (PDV) approach described in section 4.2 of RFC 5481. The 99th percentile is recorded and used in all calculations when deriving the PDV.

DNS Resolutions and DNS Failures

These tests measure the DNS resolution time of an A record query for the domains of the websites used in the web browsing test, and the percentage of DNS requests performed in the DNS resolution test that failed.

The DNS resolution test is targeted directly at the ISP's recursive resolvers. This circumvents any caching introduced by the panelist's home equipment (such as another gateway running in front of the Whitebox) and also accounts for panelists that might have configured the Whitebox (or upstream devices) to use non-ISP provided DNS servers. ISPs provide lists of their recursive DNS servers for the purposes of this study.

ICMP Latency and Packet Loss

These tests measure the round-trip time (RTT) of ICMP echo requests in microseconds from the Whitebox to a target test node. The client sends five ICMP echo requests of 56 bytes to the target test node, waiting up to three seconds for a response to each. Packets that are not received in response are treated as lost. The mean, minimum, maximum, and standard deviation of the successful results are recorded. The number of packets sent and received are recorded too.

Latency Under Load

The latency under load test operates for the duration of the 10-second downstream and upstream speed tests, with results for upstream and downstream recorded separately. While the speed tests are running, the latency under load test sends UDP datagrams to the target server and measures the round-trip time and number of packets lost. Packets are spaced five hundred milliseconds (ms) apart, and a three second timeout is used. The test records the mean, minimum, and maximum round trip times in microseconds. The number of lost UDP packets is also recorded.

This test represents an updated version of the methodology used in the initial August 2011 Report and aligns it with the methodology for the regular latency and packet loss metrics.

Traceroute

A traceroute client is used to send UDP probes to each hop in the path between client and destination. Three probes are sent to each hop. The round-trip times, the standard deviation of the round-trip times of the responses from each hop and the packet loss are recorded. The open source traceroute client "mtr" (https://github.com/traviscross/mtr) is used for carrying out the traceroute measurements.

Lightweight Capacity Test

This test measures the instantaneous capacity of the link using a small number of UDP packets. The test supports both downstream and upstream measurements, conducted independently.

In the downstream mode, the test client handshakes with the test server over TCP, requesting a fixed number of packets to be transmitted back to the client. The client specifies the transmission rate, number of packets and packet size in this handshake. The client records the arrival times of each of the resulting packets returns to it.

In the upstream mode, the client again handshakes with the test server, this time informing it of the characteristics of the stream it is about to transmit. The client then transmits the stream to the server, and the server locally records the arrival times of each packet. At the conclusion of this stream, the client asks the server for its summary of the arrival time of each packet.

With this resulting set of arrival times, the test client calculates the throughput achieved. This throughput may be divided into multiple windows, and an average taken across those, in order to smooth out buffering behavior.

This test uses approximately 99% less data than the TCP speed test and completes in a fraction of the time (100 milliseconds versus 10 seconds). The lightweight capacity test achieves results are within 1% deviation from the existing speed test results on fixed-line connections tested on average.

Table 8: Estimated Total Traffic Volume Generated by Test

The standard test schedule, below, was used across all ISPs, with the exception of Viasat. In 2017, Viasat opted to no longer provide panelists with an increased data allowance to offset the amount of data used by the measurements. This meant that the standard test schedule could no longer be used on Viasat, so a lighter weight test schedule was developed for them.

Standard Test Schedule

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
Web Browsing	9 popular US websites	Every 2 hours, 24x7	Est. 30 seconds	80 MB
Voice over IP	1 off-net test node	Hourly, 24x7	Fixed 10 seconds at 64k	1.8 MB
	1 on-net test node	Hourly, 24x7	Fixed 10 seconds at 64k	1.8 MB
Download Speed (Capacity – 8x parallel TCP connections)	1 off-net test node	Once 12 am - 6 am Once 6 am - 12 pm Once 12 pm - 6 pm Hourly thereafter	Fixed 10 seconds	107 MB at 10 Mbps
	1 on-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Once 6pm-8pm, Once 8pm-10pm, Once 10pm-12am	Fixed 10 seconds	70 MB at 10 Mbps
Download Speed (Single TCP connection)	1 off-net test node 1 on-net test node	Once in peak hours, once in off-peak hours	Fixed 10 seconds	46 MB at 10 Mbps
Upload Speed (Capacity – 8x parallel TCP connections on terrestrial, 3x on satellite)	1 off-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Hourly thereafter	Fixed 10 seconds	11 MB at 1 Mbps
	1 on-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Once 6pm-8pm, Once 8pm-10pm, Once 10pm-12am	Fixed 10 seconds	7 MB at 1 Mbps

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
Upload Speed (Single TCP connection)	1 off-net test node 1 on-net test node	Once in peak hours, once in off-peak hours	Fixed 10 seconds	6 MB at 1 Mbps
UDP Latency	2 off-net test nodes (Level3/MLab) 1 on-net test node	Hourly, 24x7 Hourly, 24x7	Permanent Permanent	5.8 MB 2.9 MB
UDP Packet Loss	2 off-net test node 1 on-net test nodes	Hourly, 24x7 Hourly, 24x7	Permanent Permanent	N/A (uses above) N/A (uses above)
Consumption	N/A	24x7	N/A	N/A
DNS Resolution	10 popular US websites	Hourly, 24x7	Est. 3 seconds	0.3 MB
ICMP Latency	1 off-net test node 1 on-net test node	Hourly, 24x7	Est. 5 seconds	0.3 MB
ICMP Packet loss	1 off-net test node 1 on-net test node	Hourly, 24x7	N/A (As IMCP latency)	N/A (uses above)
Traceroute	1 off-net test node 1 on-net test node	Three times a day, 24x7	N/A	N/A
Download Speed IPv6^^	1 off-net test node	Three times a day	Fixed 10 seconds	180 MB at 50 Mbps 72 MB at 20 Mbps 11 MB at 3 Mbps 5.4 MB at 1.5 Mbps
Upload Speed IPv6^^	1 off-net test node	Three times a day	Fixed 10 seconds	172 MB at 2 Mbps 3.6MB at 1 Mbps 1.8MB at 0.5 Mbps
UDP Latency / Loss IPv6^^	2 off-net test nodes (Level3/MLab)	Hourly, 24x7	Permanent	5.8 MB
Lightweight Capacity Test – Download (UDP)	1 off-net test node	Once 12am-6am,	Fixed 1000 packets	9MB

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
		Once 6am-12pm, Once 12pm-6pm, Hourly thereafter		
Lightweight capacity test – Upload (UDP)	1 off-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Hourly thereafter	Fixed 1000 packets	9MB

Lightweight test schedule (currently Viasat only)

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
Web Browsing	9 popular US websites	Once 8pm-10-pm	Est. 30 seconds	7MB
Download Speed (Capacity – 8x parallel TCP connections)	1 off-net test node	Once 8pm-10-pm	Fixed 10 seconds	30MB at 10Mbps
Upload Speed (Capacity – 8x parallel TCP connections on terrestrial, 3x on satellite)	1 off-net test node	Once 8pm-10-pm	Fixed 10 seconds	3MB at 1Mbps
UDP Latency	1 off-net test node	Hourly, 24x7	Permanent	1MB
UDP Latency	1 on-net test node	Hourly, 24x7	Permanent	1MB
UDP Packet loss	1 off-net test node	Hourly, 24x7	Permanent	N/A (uses above)
UDP Packet loss	1 on-net test node	Hourly, 24x7	Permanent	N/A (uses above)
Consumption	N/A	24x7	N/A	N/A
DNS Resolution	10 popular US websites	Hourly, 24x7	Est. 3 seconds	0.3MB
ICMP Latency	1 off-net test node 1 on-net test node	Hourly, 24x7	Est. 5 seconds	0.3MB
ICMP Packet Loss	1 off-net test node 1 on-net test node	Hourly, 24x7	N/A (As IMCP latency)	N/A (uses above)
Traceroute	1 off-net test node	Three times a day, 24x7	N/A	N/A

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
	1 on-net test node			
CDN Performance	Amazon, Apple, Microsoft, Google, Cloudflare, Akamai	Every 2 hours, 24x7	5 seconds	3MB
UDP Latency / Loss **IPv6^**	1 off-net test node	Hourly, 24x7	Permanent	1MB
Lightweight Capacity Test – Download (UDP)	1 off-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Hourly thereafter	Fixed 1000 packets	9MB
Lightweight Capacity Test – Upload (UDP)	1 off-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Hourly thereafter	Fixed 1000 packets	9MB

**Download/upload daily volumes are estimates based upon likely line speeds. All tests will operate at maximum line rate so actual consumption may vary.

^Currently in beta testing.

^^Only carried out on broadband connections that support IPv6.

Tests to the off-net destinations alternate randomly between Level3 and M-Lab, except that latency and loss tests operate continuously to both Level3 and M-Lab off-net servers. All tests are also performed to the closest on-net server, where available.

Consumption

This test was replaced by the new data usage test. A technical description for this test is outlined here: https://transition.fcc.gov/oet/mba/Data-Usage-Technical-Methodology_2018-08-24_Final-v1.3.pdf

Cross-Talk Testing and Threshold Manager Service

In addition to the tests described above, for 60 seconds prior to and during testing, a ”threshold manager” service on the Whitebox monitors the inbound and outbound traffic across the WAN interface to calculate if a panelist is actively using the Internet connection. The threshold for

traffic is set to 64 kbps downstream and 32 kbps upstream. Metrics are sampled and computed every 10 seconds. If either of these thresholds is exceeded, the test is delayed for a minute and the process repeated. If the connection is being actively used for an extended period of time, this pause and retry process continues for up to five times before the test is abandoned.

4 - DATA PROCESSING AND ANALYSIS OF TEST RESULTS

This section describes the background for the categorization of data gathered for the Ninth Report, and the methods employed to collect and analyze the test results.

4.1 - BACKGROUND

Time of Day

Most of the metrics reported in the Ninth Report draw on data gathered during the so-called peak usage period of 7:00 p.m. to 11:00 p.m. local time[25]. This time period is generally considered to experience the highest amount of Internet usage under normal circumstances.

ISP and Service Tier

A sufficient sample size is necessary for analysis and the ability to robustly compare the performance of specific ISP speed tiers. In order for a speed tier to be considered for the fixed line MBA Report, it must meet the following criteria:

(a) The speed tier must make up the top 80% of the ISP's subscriber base;
(b) There must be a minimum of 45 panelists that are recruited for that tier who have provided valid data for the tier within the validation period; and
(c) Each panelist must have a minimum of five days of valid data within the validation period.

The study achieved target sample sizes for the following download and upload speeds[26] (listed in alphabetical order by ISP):

Download Speeds:

AT&T IP-BB: 6 and 18 Mbps tiers;

[25] This period of time was agreed to by ISP participants in open meetings conducted at the beginning of the program.

[26] Due to the large number of different combinations of upload/download speed tiers supported by ISPs where, for example, a single download speed might be offered paired with multiple upload speeds or vice versa, upload and download test results were analyzed separately.

CenturyLink: 1.5, 3, 7, 10, 12, 20, 25, and 40 Mbps tiers;
Charter: 60, 100, and 200 Mbps tiers;
Cincinnati Bell DSL: 5 and 30 Mbps tiers;
Cincinnati Bell Fiber: 50 and 250 Mbps tier;
Comcast: 60, 150, and 250 Mbps tiers;
Cox: 30, 100, 150, and 300 Mbps tiers;
Frontier DSL: 3, 6, 12, 18 Mbps tiers;
Frontier Fiber: 50, 75, 100, 150 Mbps tiers;
Hughes: 25 Mbps tier;
Mediacom: 60 and 100 Mbps tiers;
Optimum: 100 and 200 Mbps tiers;
Verizon DSL: [1.1 - 3.0] Mbps tier;
Verizon Fiber: 50, 75, 100, and 1 Gbps tiers;[27]
Windstream: 3, 10, 12, and 25 Mbps tiers.

Upload Speeds:

AT&T IP-BB: 1 and 1.5 Mbps tiers;
CenturyLink: 0.768, 0.896, 2, and 5 Mbps tiers;
Charter: 5, 10, and 20 Mbps tiers;
Cincinnati Bell DSL: 0.768 and 3 Mbps tiers;
Cincinnati Bell Fiber: 10 and 100 Mbps tiers;
Comcast: 5 and 10 Mbps tiers;
Cox: 3, 10, and 30 Mbps tiers;
Frontier DSL: 0.768 and 1 Mbps tiers;
Frontier Fiber: 50, 75, 100, and 150 Mbps tiers;
Hughes: 1 and 3 Mbps tiers;
Mediacom: 5, and 10 Mbps tiers;
Optimum: 35 Mbps tier;
Verizon DSL: [0.384 – 0.768] Mbps tier;
Verizon Fiber: 50, 75, 100, and 1 Gbps tiers;[28]
Windstream: 0.768 and 1.5 Mbps tiers.

[27] Verizon's 1 Gbps tier was not included in the final report. 1Gbps tiers may be included in a separate/subsequent report focusing on faster speeds.
[28] Verizon's 1 Gbps tier was not included in the final report. *Id* at n. 27.

A file containing averages for each metric from the validated September/October 2018 data can be found on FCC's Measuring Broadband America website.[29] Some charts and tables are divided into speed bands, to group together products with similar levels of advertised performance. The results within these bands are further broken out by ISP and service tier. Where an ISP does not offer a service tier within a specific band or a representative sample could not be formed for tier(s) in that band, the ISP will not appear in that speed band.

[29] See: https://www.fcc.gov/general/measuring-broadband-america-measuring-fixed-broadband.

4.2 - DATA COLLECTION AND ANALYSIS METHODOLOGY

Data Integrity

To ensure the integrity of the data collected, the following validity checks were developed:

1. *Change of ISP intra-month*: By checking the WHOIS results once a day for the user's IP address, we found units that changed ISP during the month. We only kept data for the ISP where the panelist was active the most.
2. *Change of service tier intra-month*: This validity check found units that changed service tier intra-month by comparing the average sustained throughput observed for the first three days in the reporting period against that for the final three days in the reporting period. If a unit was not online at the start or end of that period, we used the first or final three days when they were actually online. If this difference was over 50 percent, the downstream and upstream charts for this unit were individually reviewed. Where an obvious step change was observed (*e.g.*, from 1 Mbps to 3 Mbps), the data for the shorter period was flagged for removal.
3. *Removal of any failed or irrelevant tests*: This validity check removed any failed or irrelevant tests by removing measurements against any nodes other than the US-based off-net nodes. We also removed measurements using any off-net server that showed a failure rate of 10 percent or greater during a specific one-hour period, to avoid using any out-of-service test nodes.
4. *Removal of any problem Whiteboxes*: We removed measurements for any Whitebox that exhibited greater than or equal to 10 percent failures in a particular one-hour period. This removed periods when the Whitebox was unable to reach the Internet.

Legacy Equipment

In previous reports, we discussed the challenges ISPs face in improving network performance where equipment under the control of the subscriber limits the end-to-end performance achievable by the subscriber.[30] Simply, some consumer-controlled equipment may not be capable of operating fully at new, higher service tiers. Working in open collaboration with all service providers we developed a policy permitting changes in ISP panelists when their installed modems were not capable of meeting the delivered service speed that included several conditions on participating ISPs. First, proposed changes in consumer panelists would only be considered where an ISP was offering free upgrades for modems they owned and leased to the consumer. Second, each ISP needed to disclose its policy regarding the treatment of legacy modems and its efforts to inform consumers regarding the impact such modems may have on

[30] See pgs. 8-9, 2014 Report, pg. 8 of the 2013 Report, as well as endnote 14. http://www.fcc.gov/measuring-broadband-america/2012/july.

their service.

While the issue of DOCSIS 3 modems and network upgrades affect the cable industry today, we may see other cases in the future where customer premises equipment affects the achievable network performance.

In accordance with the above stated policy, 95 Whiteboxes connected to legacy modems were identified and removed from the final data set in order to ensure that the study would only include equipment that would be able to meet its advertised speed. The 95 excluded Whiteboxes were connected to Charter, Comcast, and Cox.

Collation of Results and Outlier Control

All measurement data were collated and stored for analysis purposes as monthly trimmed averages during three time intervals (24 hours, 7:00 p.m. to 11:00 p.m. local time Monday through Friday, 12:00 a.m. to 12:00 a.m. local time Saturday and Sunday). Only participants who provided a minimum of five days of valid measurements and had valid data in each of the three time intervals were included in the September / October 2018 test results. In addition, the top and bottom 1 percent of measurements were trimmed to control for outliers that may have been anomalous or otherwise not representative of actual broadband performance. All results were computed on the trimmed data.[31]

Data was only charted when results from at least 45 separate Whiteboxes was available for individual ISP download speed tiers. Service tiers of 50 or fewer Whiteboxes were noted for possible future panel augmentation.

The resulting final validated sample of data for September/October 2018 included in the MBA Ninth Report was collected from 3,355 participants.

Peak Hours Adjusted to Local Time

Peak hours were defined as weekdays (Mondays through Fridays) between 7:00 p.m. to 11:00 p.m. (inclusive) for the purposes of the study. All times were adjusted to the panelist's local time zone. Since some tests are performed only once every two hours on each Whitebox, the duration of the peak period had to be a multiple of two hours.

Congestion in the Home Not Measured

Download, upload, latency, and packet loss measurements were taken between the panelist's home gateway and the dedicated test nodes provided by M-Lab and Level 3. Web browsing measurements were taken between the panelist's home gateway and nine popular United States-hosted websites. Any congestion within the user's home network is, therefore, not measured by this study. The web browsing measurements are subject to possible congestion at the content provider's side, although the choice of eight popular websites configured to serve high traffic loads reduced that risk.

Traffic Shaping Not Studied

The effect of traffic shaping is not studied in the Ninth Report, although test results were subject to any bandwidth management policies put in place by ISPs. The effects of bandwidth management policies, which may be used by ISPs to maintain consumer traffic rates within

[31] These methods were reviewed with statistical experts by the participating ISPs.

advertised service tiers, may be most readily seen in those charts in the 2016 Report that show performance over 24-hour periods, where tested rates for some ISPs and service tiers flatten for periods at a time.

Analysis of PowerBoost and Other "Enhancing" Services

The use of transient speed enhancing services marketed under names such as "PowerBoost" on cable connections presented a technical challenge when measuring throughput. These services will deliver a far higher throughput for the earlier portion of a connection, with the duration varying by ISP, service tier, and potentially other factors. For example, a user with a contracted 6 Mbps service tier may receive 18 Mbps for the first 10 MB of a data transfer. Once the "burst window" is exceeded, throughput will return to the contracted rate, with the result that the burst speed will have no effect on very long sustained transfers.

Existing speed tests transfer a quantity of data and divide this quantity by the duration of the transfer to compute the transfer rate, typically expressed in Mbps. Without accounting for burst speed techniques, speed tests employing the mechanism described here will produce highly variable results depending on how much data they transfer or how long they are run. Burst speed techniques will have a dominant effect on short speed tests: a speed test running for two seconds on a connection employing burst speed techniques would likely record the burst speed rate, whereas a speed test running for two hours will reduce the effect of burst speed techniques to a negligible level.

The earlier speed test configuration employed in this study isolated the effects of transient performance enhancing burst speed techniques from the long-term sustained speed by running for a fixed 30 seconds and recording the average throughput at 5 second intervals. The throughput at the 0-5 second interval is referred to as the burst speed and the throughput at the 25-30 second interval is referred to as the actual speed. Testing was conducted prior to the start of trial to estimate the length of time during which the effects of burst speed techniques might be seen. Even though the precise parameters used for burst-speed techniques are not known, their effects were no longer observable in testing after 20 seconds of data transfer.

In the Sixth report we noted that the use of this technology by providers was on the decline. For the Seventh, Eighth, and Ninth reports, we no longer provide the results of burst-speed since these techniques are now rarely used. The speed test configuration has been altered to shorten the test duration to 10 seconds, as there is no need to run it for 30 seconds any more.

Consistency of Speed Measurements

In addition to reporting on the median speed of panelists, the MBA Report also provides a measure of the consistency of speed that panelists experience in each tier. For purposes of discussion we use the term "80/80 consistent speed" to refer to the minimum speed that was experienced by at least 80% of panelists for at least 80% of the time during the peak periods. The process used in defining this metric for a specific ISP tier is to take each panelist's set of download or upload speed data during the peak period across all the days of the validated measurement

period and arrange it in increasing order. The speed that corresponds to the 20th percentile represents the minimum speed that the panelist experienced at least 80% of the time. The 20 percentile values of all the panelists on a specific tier are then arranged in an increasing order. The speed that corresponds to the 20th percentile now represents the minimum speed that at least 80% of panelists experienced 80% of the time. This is the value reported as the 80/80 consistent speed for that ISP's tier. We also report on the 70/70 consistent speed for an ISP's tier, which is the minimum speed that at least 70% of the panelists experience at least 70% of the time. We typically report the 70/70 and the 80/80 consistent speeds as a percentage of the advertised speed.

When reporting on these values for an ISP, we weigh the 80/80 or 70/70 consistent speed results (as a percentage of the advertised speed) of each of the ISP's tier based on the number of subscribers to that tier; so as to get a weighted average across all the tiers for that ISP.

Latencies Attributable to Propagation Delay

The speeds at which signals can traverse networks are limited at a fundamental level by the speed of light. While the speed of light is not believed to be a significant limitation in the context of the other technical factors addressed by the testing methodology, a delay of approximately 5ms per 1000 km of distance traveled can be attributed solely to the speed of light (depending on the transmission medium). The geographic distribution and the testing methodology's selection of the nearest test servers are believed to minimize any significant effect. However, propagation delay is not explicitly accounted for in the results.

Limiting Factors

A total of 8,417,695,058 measurements were taken across 144,636,223 unique tests.

All scheduled tests were run, aside from when monitoring units detected concurrent use of bandwidth.

Schedules were adjusted when required for specific tests to avoid triggering data usage limits applied by some ISPs.

4.3 DATA PROCESSING OF RAW AND VALIDATED DATA

The data collected in this program are made available as open data for review and use by the public. Raw and processed data sets, mobile testing software, and the methodologies used to process and analyze data are freely and publicly available. Researchers and developers interested in working with measurement data in raw form will need skills in database management, SQL programming, and statistics, depending on the analysis. A developer FAQ for database configuration and data importing instructions for MySQL and PostgreSQL are available at https://www.fcc.gov/general/database-setup-and-importing-measuring-broadband-america-data-april-2012.

The process flow below describes how the raw collected data was processed for the production of the *Measuring Broadband America Report*. Researchers and developers interested in replicating or extending the results of the Report are encouraged to review the process below and supporting files that provide details.

Raw Data:	Raw data for the chosen period is collected from the measurement database. The ISPs and products that panelists were on are exported to a "unit profile" file, and those that changed during the period are flagged. 2018 Raw Data Links
Validated Data Cleansing:	Data is cleaned. This includes removing measurements when a user changed ISP or tier during the period. Anomalies and significant outliers are also removed at this point. A data cleansing document describes the process in detail. 2018 Data Cleansing Document Link
SQL Processing:	Per-unit results are generated for each metric. Time-of-day averages are computed and a trimmed median is calculated for each metric. The SQL scripts used here are contained in SQL processing scripts available with the release of each report. 2018 SQL Processing Links
Unit Profile:	This document identifies the various details of each test unit, including ISP, technology, service tier, and general location. Each unit represents one volunteer panelists. The unit ID's were randomly generated, which served to protect the anonymity of the volunteer panelists. 2018 Unit Profile link
Excluded Units:	A listing of units excluded from the analysis due to insufficient sample size for that particular ISP's speed tier. 2018 Excluded Units Link
Unit Census Block:	This step identifies the census block (for blocks containing more than 1,000 people) in which each unit running tests is located. Census block is from 2010 census and is in the FIPS code format. We have used block FIPS codes for blocks that contains more than 1,000 people. For blocks with fewer than 1,000 people we have aggregated to the next highest level, *i.e.*, tract, and used the Tract FIPS code, provided there are more than 1,000 people in the tract. In cases where there are less than 1,000 people in a tract we have aggregated to Regional level. 2018 Unit Census Block Link.
Excel Tables & Charts:	Summary data tables and charts in Excel are produced from the averages. These are used directly in the report 2018 Statistical Averages Links

The raw data collected for each active metric is made available by month in tarred gzipped files. The files in the archive containing active metrics are described in table 9.

Table 9: Test to Data File Cross-Reference List

Test	Validated Data File Name
Download Speed	curr_httpgetmt.csv — IPv4 Tests curr_httpgetmt6.csv — IPv6 Tests
Upload Speed	curr_httppostmt.csv — IPv4 Tests curr_httppostmt6.csv — IPv6 Tests
Web Browsing	curr_webget.csv
UDP Latency	curr_udplatency.csv — IPv4 Tests curr_udplatency6.csv — IPv6 Tests
UDP Packet Loss	curr_udplatency.csv — IPv4 Tests curr_udplatency6.csv — IPv6 Tests
Voice over IP	curr_udpjitter.csv
DNS Resolution	curr_dns.csv
DNS Failures	curr_dns.csv
ICMP Latency	curr_ping.csv
ICMP Packet Loss	curr_ping.csv
Latency under Load	curr_dlping.csv – Downstream latency under load results curr_ulping.csv – Upstream latency under load results
Traceroute	curr_traceroute.csv

Table 10: Validated Data Files - Dictionary

The following Data Dictionary file describes the schema for each active metric test for row level results stored in the files described in table 9.[32] All dtime entries are in the UTC timezone. All durations are in microseconds unless otherwise noted. The location_id field should be ignored.

curr_dlping.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address

[32] This data dictionary is also available on the FCC Measuring Broadband America website, located with the other validated data files available for download.

rtt_avg	Average RTT
rtt_min	Minimum RTT
rtt_max	Maximum RTT
rtt_std	Standard deviation in measured RTT
successes	Number of successes
failiures	Number of failures
location_id	Internal key mapping to unit profile data
curr_dns.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
nameserver	Name server used to handle the DNS request
lookup_host	Hostname to be resolved
response_ip	Field currently unused
rtt	DNS resolution time
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Internal key mapping to unit profile data
curr_httpgetmt.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
address	The IP address of the server (resolved by the client's DNS)
fetch_time	Time the test ran for
bytes_total	Total bytes downloaded across all connections
bytes_sec	Running total of throughput, which is sum of speeds measured for each stream (in bytes/sec), from the start of the test to the current interval
bytes_sec_interval	Throughput at this specific interval (*e.g.*, Throughput between 25-30 seconds)
warmup_time	Time consumed for all the TCP streams to arrive at optimal window size
warmup_bytes	Bytes transferred for all the TCP streams during the warm-up phase

sequence	The interval that this row refers to (*e.g.*, in the US, sequence=0 implies result is for 0-5 seconds of the test)
threads	The number of concurrent TCP connections used in the test
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Internal key mapping to unit profile data
curr_httppostmt.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
address	The IP address of the server (resolved by the client's DNS)
fetch_time	Time the test ran for
bytes_total	Total bytes downloaded across all connections
bytes_sec	Running total of throughput, which is sum of speeds measured for each stream (in bytes/sec), from the start of the test to the current interval
bytes_sec_interval	Throughput at this specific interval (*e.g.*, throughput between 25-30 seconds)
warmup_time	Time consumed for all the TCP streams to arrive at optimal window size
warmup_bytes	Bytes transferred for all the TCP streams during the warm-up phase.
sequence	The interval that this row refers to (*e.g.*, in the US, sequence=0 implies result is for 0-5 seconds of the test)
threads	The number of concurrent TCP connections used in the test
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Internal key mapping to unit profile data
curr_ping.csv	ICMP based
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address

rtt_avg	Average RTT
rtt_min	Minimum RTT
rtt_max	Maximum RTT
rtt_std	Standard deviation in measured RTT
successes	Number of successes
failiures	Number of failures
location_id	Internal key mapping to unit profile data
curr_udpjitter.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
packet_size	Size of each UDP Datagram (bytes)
stream_rate	Rate at which the UDP stream is generated (bits/sec)
duration	Total duration of test
packets_up_sent	Number of packets sent in upstream (measured by client)
packets_down_sent	Number of packets sent in downstream (measured by server)
packets_up_recv	Number of packets received in upstream (measured by server)
packets_down_recv	Number of packets received in downstream (measured by client)
jitter_up	Upstream Jitter measured
jitter_down	Downstream Jitter measured
latency	99th percentile of round trip times for all packets
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Internal key mapping to unit profile data
curr_udplatency.csv	UDP based
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
rtt_avg	Average RTT
rtt_min	Minimum RTT

rtt_max	Maximum RTT
rtt_std	Standard deviation in measured RTT
successes	Number of successes (note: use failures/(successes + failures)) for packet loss)
failiures	Number of failures (packets lost)
location_id	Internal key mapping to unit profile data
curr_ulping.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
rtt_avg	Average RTT
rtt_min	Minimum RTT
rtt_max	Maximum RTT
rtt_std	Standard deviation in measured RTT
successes	Number of successes
failures	Number of failures
location_id	Internal key mapping to unit profile data
curr_webget.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	URL to fetch
address	IP address used to fetch content from initial URL
fetch_time	Sum of time consumed to download HTML content and then concurrently download all resources
bytes_total	Sum of HTML content size and all resources size (bytes)
bytes_sec	Average speed of downloading HTML content and then concurrently downloading all resources (bytes/sec)
objects	Number of resources (images, CSS, …) downloaded
threads	Maximum number of concurrent threads allowed
requests	Total number of HTTP requests made
connections	Total number of TCP connections established
reused_connections	Number of TCP connections re-used

lookups	Number of DNS lookups performed
request_total_time	Total duration of all requests summed together, if made sequentially
request_min_time	Shortest request duration
request_avg_time	Average request duration
request_max_time	Longest request duration
ttfb_total_time	Total duration of the time-to-first-byte summed together, if made sequentially
ttfb_min_time	Shortest time-to-first-byte duration
ttfb_avg_time	Average time-to-first-byte duration
ttfb_max_time	Longest time-to-first-byte duration
lookup_total_time	Total duration of all DNS lookups summed together, if made sequentially
lookup_min_time	Shortest DNS lookup duration
lookup_avg_time	Average DNS lookup duration
lookup_max_time	Longest DNS lookup duration
successes	Number of successes
failures	Number of failures
location_id	Internal key mapping to unit profile data
curr_netusage.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
wan_rx_bytes	Total bytes received via the WAN interface on the unit (incl. Ethernet and IP headers)
wan_tx_bytes	Total bytes transmitted via the WAN interface on the unit (incl. Ethernet and IP headers)
sk_rx_bytes	Bytes received as a result of active performance measurements
sk_tx_bytes	Bytes transmitted as a result of active performance measurements
location_id	Internal key mapping to unit profile data

curr_lct_dl.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC

target	Target hostname
address	Target IP address
packets_received	Total number of packets received
packets_sent	Total number of packets sent
packet_size	Packet size
bytes_total	Total number of bytes
duration	Duration of the test in microseconds
bytes_sec	Throughput in bytes/sec
error_code	An internal error code from the test.
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Please ignore (this is an internal key mapping to unit profile data)

curr_lct_ul.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	Target hostname
address	Target IP address
packets_received	Total number of packets received
packets_sent	Total number of packets sent
packet_size	Packet size
bytes_total	Total number of bytes
duration	Duration of the test in microseconds
bytes_sec	Throughput in bytes/sec
error_code	An internal error code from the test.
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Please ignore (this is an internal key mapping to unit profile data)

5 - Reference Documents

5.1 - USER TERMS AND CONDITIONS

The following document was agreed to by each volunteer panelist who agreed to participate in the broadband measurement study:

End User License Agreement

PLEASE READ THESE TERMS AND CONDITIONS CAREFULLY. BY APPLYING TO BECOME A PARTICIPANT IN THE BROADBAND COMMUNITY PANEL AND/OR INSTALLING THE WHITEBOX, YOU ARE AGREEING TO THESE TERMS AND CONDITIONS.

YOUR ATTENTION IS DRAWN PARTICULARLY TO CONDITIONS 3.5 (PERTAINING TO YOUR CONSENT TO YOUR ISPS PROVIDING CERTAIN INFORMATION AND YOUR WAIVER OF CLAIMS), 6 (LIMITATIONS OF LIABILITY) AND 7 (DATA PROTECTION).

1. Interpretation

1.1. The following definitions and rules of interpretation apply to these terms & conditions.

Connection: the Participant's own broadband internet connection, provided by an Internet Service Provider ("ISP").

Connection Equipment: the Participant's broadband router or cable modem, used to provide the Participant's Connection.

Intellectual Property Rights: all patents, rights to inventions, utility models, copyright and related rights, trademarks, service marks, trade, business and domain names, rights in trade dress or get-up, rights in goodwill or to sue for passing off, unfair competition rights, rights in designs, rights in computer software, database right, moral rights, rights in confidential information (including know-how and trade secrets) and any other intellectual property rights, in each case whether registered or unregistered and including all applications for and renewals or extensions of such rights, and all similar or equivalent rights or forms of protection in any part of the world.

ISP: the company providing broadband internet connection to the Participant during the term of this Program.

Participant/You/Your: the person who volunteers to participate in the Program, under these terms and conditions. The Participant must be the named account holder on the Internet service account with the ISP.

Open Source Software: the software in the Whitebox device that is licensed under an open source license (including the GPL).

Participant's Equipment: any equipment, systems, cabling or facilities provided by the Participant and used directly or indirectly in support of the Services, excluding the Connection Equipment.

Parties: both the Participant and SamKnows.

Party: one of either the Participant or SamKnows.

Requirements: the requirements specified by SamKnows as part of the sign-up process that the Participant must fulfil in order to be selected to receive the Services.

SamKnows/We/Our: the organization providing the Services and conducting the Program, namely:

SamKnows Limited (Co. No. 6510477) of 25 Harley Street, London W1G 9BR

Services / Program: the performance and measurement of certain broadband and Internet services and research program (Broadband Community Panel), as sponsored by the Federal Communications Committee (FCC), in respect of measuring broadband Internet Connections.

Software: the software that has been installed and/or remotely uploaded onto the Whitebox, by SamKnows as updated by SamKnows, from time to time, but not including any Open Source Software.

Test Results: Information concerning the Participant's ISP service results.

Whitebox: the hardware supplied to the Participant by SamKnows with the Software.

1.2. Headings in these terms and conditions shall not affect their interpretation.

1.3. A person includes a natural person, corporate or unincorporated body (whether or not having separate legal personality).

1.4. The schedules form part of these terms and conditions.

1.5. A reference to writing or written includes faxes and e-mails.

1.6.Any obligation in these terms and conditions on a person not to do something includes, without limitation, an obligation not to agree, allow, permit or acquiesce in that thing being done.

2. SamKnows' Commitment to You

2.1 Subject to the Participant complying fully with these terms and conditions, SamKnows shall use reasonable care to:

(a) provide the Participant with the Measurement Services under these terms and conditions;

(b) supply the Participant with the Whitebox and instructions detailing how it should be connected to the Participant's Connection Equipment; and

(c) if requested, SamKnows will provide a pre-paid postage label for the Whitebox to be returned.

(d) comply with all applicable United States, European Union, and United Kingdom privacy laws and directives, and will access, collect, process and distribute the information according to the following principles:

Fairness: We will process data fairly and lawfully;

Specific purpose: We will access, collect, process, store and distribute data for the purposes and reasons specified in this agreement and not in ways incompatible with those purposes;

Restricted: We will restrict our data collection and use practices to those adequate and relevant, and not excessive in relation to the purposes for which we collect the information;

Accurate: We will work to ensure that the data we collect is accurate and up-to-date, working with Participant and his/her ISP;

Destroyed when obsolete: We will not maintain personal data longer than is necessary for the purposes for which we collect and process the information;

Security: We will collect and process the information associated with this trial with adequate security through technical and organizational measures to protect personal data against destruction or loss, alteration, unauthorized disclosure or access, in particular where the processing involves the transmission of data over a network.

2.2 In addition, SamKnows shall:

(a) provide Participant with access to a Program-specific customer services email address, which the Participant may use for questions and to give feedback and comments;

(b) provide Participant with a unique login and password in order to access to an online reporting system for access to Participant's broadband performance statistics.

(c) provide Participant with a monthly email with their specific data from the Program or notifying Participant that their individual data is ready for viewing;

(d) provide Participant with support and troubleshooting services in case of problems or issues with their Whitebox;

(e) notify Participant of the end of the FCC-sponsored Program and provide a mechanism for Participant to opt out of any further performance/measuring services and research before collecting any data after termination of the Program;

(f) use only data generated by SamKnows through the Whitebox, and not use any Participant data for measuring performance without Participant's prior written consent; and

(g) not monitor/track Participant's Internet activity without Participant's prior written consent.

2.3 While SamKnows will make all reasonable efforts to ensure that the Services cause no disruption to the performance of the Participant's broadband Connection, including only running tests when there is no concurrent network activity generated by users at the Participant's location. The Participant acknowledges that the Services may occasionally impact the performance of the Connection and agrees to hold SamKnows and their ISP harmless for any impact the Services may have on the performance of their Connection.

3. Participant's Obligations

3.1 The Participant is not required to pay any fee for the provision of the Services by SamKnows or to participate in the Program.

3.2 The Participant agrees to use reasonable endeavors to:

(a) connect the Whitebox to their Connection Equipment within 14 days of receiving it;

(b) not to unplug or disconnect the Whitebox unless (i) they will be absent from the property in which it is connected for more than 3 days and/or (ii) it is reasonably necessary for maintenance of the Participant's Equipment and the Participant agrees that they shall use reasonable endeavors to minimize the length of time the Whitebox is unplugged or disconnected;

(c) in no way reverse engineer, tamper with, dispose of or damage the Whitebox, or attempt to do so;

(d) notify SamKnows within 7 days in the event that they change their ISP or their Connection tier or package (for example, downgrading/upgrading to a different broadband package), to the email address provided by SamKnows;

(e) inform SamKnows of a change of postal or email address by email; within 7 days of the change, to the email address provided by SamKnows;

(f) agrees that the Whitebox may be upgraded to incorporate changes to the Software and/or additional tests at the discretion of SamKnows, whether by remote uploads or otherwise;

(g) on completion or termination of the Services, return the Whitebox to SamKnows by mail, if requested by SamKnows. SamKnows will provide a pre-paid postage label for the Whitebox to be returned;

(h) be an active part of the Program and as such will use all reasonable endeavors to complete the market research surveys received within a reasonable period of time;

(i) not publish data, give press or other interviews regarding the Program without the prior written permission of SamKnows; and

(k) contact SamKnows directly, and not your ISP, in the event of any issues or problems with the Whitebox, by using the email address provided by SamKnows.

3.3 You will not give the Whitebox or the Software to any third party, including (without limitation) to any ISP. You may give the Open Source Software to any person in accordance with the terms of the relevant open source licence.

3.4 The Participant acknowledges that he/she is not an employee or agent of, or relative of, an employee or agent of an ISP or any affiliate of any ISP. In the event that they become one, they will inform SamKnows, who at its complete discretion may ask for the immediate return of the Whitebox.

3.5 THE PARTICIPANT'S ATTENTION IS PARTICULARLY DRAWN TO THIS CONDITION. The Participant expressly consents to having their ISP provide to SamKnows and the Federal Communications (FCC) information about the Participant's broadband service, for example: service address, speed tier, local loop length (for DSL customers), equipment identifiers and other similar information, and hereby waives any claim that its ISPs disclosure of such information to SamKnows or the FCC constitutes a violation of any right or any other right or privilege that the Participant may have under any federal, state or local statute, law, ordinance, court order, administrative rule, order or regulation, or other applicable law, including, without limitation, under 47 U.S.C**.** §§ 222 and 631 (each a "Privacy Law"). If notwithstanding Participant's consent under this Section 3.5, Participant, the FCC or any other party brings any claim or action against any ISP under a Privacy Law, upon the applicable ISPs request SamKnows promptly shall cease collecting data from such Participant and remove from its records all data collected with respect to such Participant prior to the date of such request, and shall not provide such data in any form to the FCC. The Participant further consents to transmission of information from this Program Internationally, including the information provided by the Participant's ISP, specifically the transfer of this information to SamKnows in the United Kingdom, SamKnows' processing of it there and return to the United States.

4. Intellectual Property Rights

4.1 All Intellectual Property Rights relating to the Whitebox are the property of its manufacturer. The Participant shall use the Whitebox only to allow SamKnows to provide the Services.

4.2 As between SamKnows and the Participant, SamKnows owns all Intellectual Property Rights in the Software. The Participant shall not translate, copy, adapt, vary or alter the Software. The Participant shall use the Software only for the purposes of SamKnows providing the Services and shall not disclose or otherwise use the Software.

4.3 Participation in the Broadband Community Panel gives the participant no Intellectual Property Rights in the Test Results. Ownership of all such rights is governed by Federal Acquisition Regulation Section 52.227-17, which has been incorporated by reference in the relevant contract between SamKnows and the FCC. The Participant hereby acknowledges and agrees that SamKnows may make such use of the Test Results as is required for the Program.

4.4 Certain core testing technology and aspects of the architectures, products and services are developed and maintained directly by SamKnows. SamKnows also implements various technical features of the measurement services using particular technical components from a variety of vendor partners including: NetGear, Measurement Lab, TP-Link.

5. SamKnows' Property

The Whitebox and Software will remain the property of SamKnows. SamKnows may at any time ask the Participant to return the Whitebox, which they must do within 28 days of such a request being sent. Once SamKnows has safely received the Whitebox, SamKnows will reimburse the Participant's reasonable postage costs for doing so.

6. Limitations of Liability - THE PARTICIPANT'S ATTENTION IS PARTICULARLY DRAWN TO THIS CONDITION

6.1 This condition 6 sets out the entire financial liability of SamKnows (including any liability for the acts or omissions of its employees, agents, consultants, and subcontractors) to the Participant, including and without limitation, in respect of:

(a) any use made by the Participant of the Services, the Whitebox and the Software or any part of them; and

(b) any representation, statement or tortious act or omission (including negligence) arising under or in connection with these terms and conditions.

6.2 All implied warranties, conditions and other terms implied by statute or other law are, to the fullest extent permitted by law, waived and excluded from these terms and conditions.

6.3 Notwithstanding the foregoing, nothing in these terms and conditions limits or excludes the liability of SamKnows:

(a) for death or personal injury resulting from its negligence or willful misconduct;

(b) for any damage or liability incurred by the Participant as a result of fraud or fraudulent misrepresentation by SamKnows;

(c) for any violations of U.S. consumer protection laws;

(d) in relation to any other liabilities which may not be excluded or limited by applicable law.

6.4 Subject to condition 6.2 and condition 6.3, SamKnows' total liability in contract, tort (including negligence or breach of statutory duty), misrepresentation, restitution or otherwise arising in connection with the performance, or contemplated performance, of these terms and conditions shall be limited to $100.

6.5 In the event of any defect or modification in the Whitebox, the Participant's sole remedy shall be the repair or replacement of the Whitebox at SamKnows' reasonable cost, provided that the defective Whitebox is safely returned to SamKnows, in which case SamKnows shall pay the Participant's reasonable postage costs.

6.6 The Participant acknowledges and agrees that these limitations of liability are reasonable in all the circumstances, particularly given that no fee is being charged by SamKnows for the Services or participation in the Program.

6.7 It is the Participant's responsibility to pay all service and other charges owed to its ISP in a timely manner and to comply with all other ISP applicable terms. The Participant shall ensure that their broadband traffic, including the data pushed by SamKnows during the Program, does not exceed the data allowance included in the Participant's broadband package. If usage allowances are accidentally exceeded and the Participant is billed additional charges from the ISP as a result, SamKnows is not under any obligation to cover these charges although it may choose to do so at its discretion.

7. Data protection - the participation's attention is particularly drawn to this condition.

7.1 The Participant acknowledges and agrees that his/her personal data, such as service tier, address and line performance, will be processed by SamKnows in connection with the program.

7.2 Except as required by law or regulation, SamKnows will not provide the Participant's personal data to any third party without obtaining Participant's prior consent. However, for the avoidance of doubt, the Participant acknowledges and agrees that subject to the privacy polices discussed below, the specific technical characteristics of tests and other technical features associated with the Internet Protocol environment of architecture, including the client's IP address, may be shared with third parties as necessary to conduct the Program and all aggregate statistical data produced as a result of the Services (including the Test Results) may be provided to third parties.

7.3 You acknowledge and agree that SamKnows may share some of Your information with Your ISP, and request information about You from Your ISP so that they may confirm Your service tiers and other information relevant to the Program. Accordingly You hereby expressly waive claim that any disclosure by Your ISP to SamKnows constitutes a violation of any right or privilege that you may have under any law, wherever it might apply.

8. Term and Termination

8.1 This Agreement shall continue until terminated in accordance with this clause.

8.2 Each party may terminate the Services immediately by written notice to the other party at any time. Notice of termination may be given by email. Notices sent by email shall be deemed to be served on the day of transmission if transmitted before 5.00 pm Eastern Time on a working day, but otherwise on the next following working day.

8.3 On termination of the Services for any reason:

(a) SamKnows shall have no further obligation to provide the Services; and

(b) the Participant shall safely return the Whitebox to SamKnows, if requested by SamKnows, in which case SamKnows shall pay the Participant's reasonable postage costs.

8.4 Notwithstanding termination of the Services and/or these terms and conditions, clauses 1, 3.3 and 4 to 14 (inclusive) shall continue to apply.

9. Severance

If any provision of these terms and conditions, or part of any provision, is found by any court or other authority of competent jurisdiction to be invalid, illegal or unenforceable, that provision or part-provision shall, to the extent required, be deemed not to form part of these terms and conditions, and the validity and enforceability of the other provisions these terms and conditions shall not be affected.

10. Entire agreement

10.1 These terms and conditions constitute the whole agreement between the parties and replace and supersede any previous agreements or undertakings between the parties.

10.2 Each party acknowledges that, in entering into these terms and conditions, it has not relied on, and shall have no right or remedy in respect of, any statement, representation, assurance or warranty.

11. Assignment

11.1 The Participant shall not, without the prior written consent of SamKnows, assign, transfer, charge, mortgage, subcontract all or any of its rights or obligations under these terms and conditions.

11.2 Each party that has rights under these terms and conditions acknowledges that they are acting on their own behalf and not for the benefit of another person.

12. No Partnership or Agency

Nothing in these terms and conditions is intended to, or shall be deemed to, constitute a partnership or joint venture of any kind between any of the parties, nor make any party the agent of another party for any purpose. No party shall have authority to act as agent for, or to bind, the other party in any way.

13. Rights of third parties

Except for the rights and protections conferred on ISPs under these Terms and Conditions which they may defend, a person who is not a party to these terms and conditions shall not have any rights under or in connection with these Terms and Conditions.

14. Privacy and Paperwork Reduction Acts

14.1 For the avoidance of doubt, the release of IP protocol addresses of client's Whiteboxes are not PII for the purposes of this program and the client expressly consents to the release of IP address and other technical IP protocol characteristics that may be gathered within the context of the testing architecture. SamKnows, on behalf of the FCC, is collecting and storing broadband performance information, including various personally identifiable information (PII) such as the street addresses, email addresses, sum of data transferred, and broadband performance information, from those individuals who are participating voluntarily in this test. PII not necessary to conduct this study will not be collected. Certain information provided by or collected from you will be confirmed with a third party, including your ISP, to ensure a representative study and otherwise shared with third parties as necessary to conduct the program. SamKnows will not release, disclose to the public, or share any PII with any outside entities, including the FCC, except as is consistent with the SamKnows privacy policy or these Terms and Conditions. See https://www.measuringbroadbandamerica.com/privacy/. The broadband performance

information that is made available to the public and the FCC, will be in an aggregated form and with all PII removed. For more information, see the Privacy Act of 1974, as amended (5 U.S.C. § 552a), and the SamKnows privacy policy.

14.2 The FCC is soliciting and collecting this information authorized by OMB Control No. 3060-1139 in accordance with the requirements and authority of the Paperwork Reduction Act, Pub. L. No. 96-511, 94 Stat. 2812 (Dec. 11, 1980); the Broadband Data Improvement Act of 2008, Pub. L. No. 110-385, Stat 4096 § 103(c)(1); American Reinvestment and Recovery Act of 2009 (ARRA), Pub. L. No. 111-5, 123 Stat 115 (2009); and Section 154(i) of the Communications Act of 1934, as amended.

14.3 *Paperwork Reduction Act of 1995 Notice*. We have estimated that each Participant of this study will assume a one hour time burden over the course of the Program. Our estimate includes the time to sign-up online, connect the Whitebox in the home, and periodic validation of the hardware. If you have any comments on this estimate, or on how we can improve the collection and reduce the burden it causes you, please write the Federal Communications Commission, Office of Managing Director, AMD-PERM, Washington, DC 20554, Paperwork Reduction Act Project (3060-1139). We will also accept your comments via the Internet if you send an e-mail to PRA@fcc.gov. Please DO NOT SEND COMPLETED APPLICATION FORMS TO THIS ADDRESS. You are not required to respond to a collection of information sponsored by the Federal government, and the government may not conduct or sponsor this collection, unless it displays a currently valid OMB control number and provides you with this notice. This collection has been assigned an OMB control number of 3060-1139. THIS NOTICE IS REQUIRED BY THE PAPERWORK REDUCTION ACT OF 1995, PUBLIC LAW 104-13, OCTOBER 1, 1995, 44 U.S.C. SECTION 3507. This notice may also be found at https://www.measuringbroadbandamerica.com/paperwork-reduction-act/.

15. Jurisdiction

These terms and conditions shall be governed by the laws of the state of New York.

SCHEDULE

THE SERVICES

Subject to the Participant complying with its obligations under these terms and conditions, SamKnows shall use reasonable endeavors to test the Connection so that the following information is recorded:

1. Web browsing
2. Video streaming
3. Voice over IP
4. Download speed
5. Upload speed
6. UDP latency
7. UDP packet loss
8. Consumption

9. Availability
10. DNS resolution
11. ICMP latency
12. ICMP packet loss

In performing these tests, the Whitebox will require a variable download capacity and upload capacity per month, which will be available to the Participant in motion 2.3. The Participant acknowledges that this may impact on the performance of the Connection.

1. SamKnows will perform tests on the Participant's Connection by using SamKnows' own data and will not monitor the Participant's content or internet activity. The purpose of this study is to measure the Connection and compare this data with other consumers to create a representative index of US broadband performance.

5.2 - CODE OF CONDUCT

The following Code of Conduct, available at http://data.fcc.gov/download/measuring-broadband-america/2017/Code-of-Conduct-fixed.pdf*, was signed by ISPs and other entities participating in the study:*

FCC MEASURING BROADBAND AMERICA PROGRAM

FIXED TESTING AND MEASUREMENT

STAKEHOLDERS CODE OF CONDUCT

WHEREAS the Federal Communications Commission of the United States of America (FCC) is conducting a Broadband Testing and Measurement Program, with support from its contractor SamKnows, the purpose of which is to establish a technical platform for the Measuring Broadband America Program Fixed Broadband Testing and Measurement and further to use that platform to collect data;

WHEREAS volunteer panelists have been recruited, and in so doing have agreed to provide broadband performance information measured on their Whiteboxes to support the collection of broadband performance data; and steps have been taken to protect the privacy of panelists to the program's effort to measure broadband performance. WE, THE UNDERSIGNED, as participants and stakeholders in that Fixed Broadband Testing and Measurement, do hereby agree to be bound by and conduct ourselves in accordance with the following principles and shall:

1. At all times act in good faith;
2. Not act, nor fail to act, if the intended consequence of such act or omission is inconsistent with the privacy policies of the program;
3. Not act, nor fail to act, if the intended consequence of such act or omission is to enhance, degrade, or tamper with the results of any test for any individual panelist or broadband provider, except that:

3.1. It shall not be a violation of this principle for broadband providers to:

3.1.1. Operate and manage their business, including modifying or improving services delivered to any class of subscribers that may or may not include panelists among them, provided that such actions are consistent with normal business practices, and

3.1.2. Address service issues for individual panelists at the request of the panelist or based on information not derived from the trial;

3.2. It shall not be a violation of this principle for academic and research purposes to simulate or observe tests and components of the testing architecture, provided that no impact to MBA data or the Internet Service of the subscriber volunteer panelist occurs; and

4. Not publish any data generated by the tests, nor make any public statement based on such data, until such time as the FCC releases data, or except where expressly permitted by the FCC; and

5. Not publish or make use of any test data or testing infrastructure in a manner that would significantly reduce the anonymity of collected data, compromise panelists privacy, or compromise the MBA privacy policy governing collection and analysis of data except that:

5.1. It shall not be a violation of this principle for stakeholder signatories under the direction of the FCC to:

5.1.1. Make use of test data or testing infrastructure to support the writing of FCC fixed Measuring Broadband America Reports;

5.1.2. Make use of test data or testing infrastructure to support various aspects of the testing and architecture for the program including to facilitate data processing or analysis;

5.1.3. Make use of test data or testing infrastructure to support the analysis of collected data or testing infrastructure for privacy risks or concerns, and plan for future measurement efforts;

6. Ensure that their employees, agents, and representatives, as appropriate, act in accordance with this Code of Conduct.

Signatories: _____

Printed: _____

Date: _____

5.3 - TEST NODE BRIEFING

Test Node Briefing

DOCUMENT REFERENCE:

SQ302-002-EN

TEST NODE BRIEFING

Technical information relating to

the SamKnows test nodes

August 2013

Important Notice

1 - SamKnows Test Nodes

In order to gauge an Internet Service Provider's broadband performance at a User's access point, the SamKnows Whiteboxes need to measure the service performance (*e.g.*, upload/download speeds, latency, *etc.*) from the Whitebox to a specific test node. SamKnows supports a number of "test nodes" for this purpose.

The test nodes run special software designed specifically for measuring the network performance when communicating with the Whiteboxes.

It is critical that these test nodes be deployed near to the customer (and their Whitebox). The further the test node is from the customer, the higher the latency and the greater the possibility that third-party networks may need to be traversed, making it difficult to isolate the individual ISP's performance. This is why SamKnows operates so many test nodes all around the world—locality to the customer is critical.

1.1 Test node definition

When referring to "test nodes," we are specifically referring to either the dedicated servers that are under SamKnows' control, or the virtual machines that may be provided to us. In the case of virtual machines provided by Measurement-Lab, Level3, and others, the host operating system is under the control of and maintained by these entities and not by SamKnows.

1.2 Test node selection

The SamKnows Whiteboxes select the nearest node by running round-trip latency checks to all test nodes before measurement begins. Note that when we use the term "nearest" we are referring to the test node nearest to the Whitebox from the point of view of network delay, which may not necessarily always be the one nearest geographically.

Alternatively, it is possible to override test node selection based on latency and implement a static configuration so that the Whitebox will only test against the test node chosen by the Administrator. This is so that the Administrator can choose to test any particular test node that is of interest to the specific project and also to maintain configuration consistency. Similarly, test node selection may be done on a scheduled basis, alternating between servers, to collect test data from multiple test nodes for comparison purposes.

1.3 Test node positioning—on-net versus off-net

It is important that measurements collected by the test architecture support the comparison of ISP performance in an unbiased manner. Measurements taken from using the standardized set of "off-net" measurement test nodes (off-net here refers to a test node located outside a specific ISP's network) ensure that the performance of all ISPs can be measured under the same conditions and would avoid artificially biasing results for any one ISP over another. Test nodes located on a particular ISP's network ("on-net" test nodes), might introduce bias with respect to the ISP's own network performance. Thus data to be used to compare ISP performance are collected using "off-net" test nodes, because they reside outside the ISP network.

However, it is also very useful to have test nodes inside the ISP network ("on-net" test nodes). This allows us to:

- Determine what degradation in performance occurs when traffic leaves the ISP network; and
- Check that the off-net test nodes are performing properly (and vice versa).
- By having both on-net and off-net measurement data for each Whitebox, we can have a great deal of confidence in the quality of the data.

2.3 Data that is stored on test nodes

No measurement data collected by SamKnows is stored on test nodes.[33] The test nodes provide a "dumb" endpoint for the Whiteboxes to test against. All measurement performance results are recorded by the Whiteboxes, which are then transmitted from the Whitebox to data collection servers managed by SamKnows.

Note that Measurement-Lab run sidestream measurements for all TCP connections against their test nodes, and publish this data in accordance with their data embargo policy.

[33] Note that Measurement-Lab runs sidestream measurements for all TCP connections against their test nodes and publishes these data in accordance with their data embargo policy.

2 - Test Node Hosting and Locations

SamKnows test nodes reside in major peering locations around the world. Test nodes are carefully sited to ensure optimal connectivity on a market-by-market basis. SamKnows' test infrastructure utilizes nodes made available by Level3, Measurement-Lab and various network operators, as well as under contract with select hosting providers.

2.1 Global Test Nodes

Level3 has provided SamKnows with 11 test nodes to use for the FCC's Measuring Broadband America Program. These test nodes are virtual servers meeting SamKnows specifications. Similarly, Measurement-Lab has also provided SamKnows with test nodes in various cities and countries for use with the Program's fixed measurement efforts. Measurement-Lab provides location hosting for at least three test nodes per site. Furthermore, SamKnows maintains its own test nodes, which are separate from the test nodes provided by Measurement-Lab and Level3.

Table 1 below shows the locations of the SamKnows test node architecture supporting the Measuring Broadband America Program.[34] All of these listed test nodes reside outside individual ISP networks and therefore are designated as off-net test nodes. Note, that in many locations there are multiple test nodes installed which may be connected to different providers.

Location	SamKnows	Level3	Measurement-Lab
Atlanta, Georgia			✓
Chicago, Illinois		✓	✓
Dallas, Texas		✓	✓
Los Angeles, California	✓	✓	✓
Miami, Florida			✓
Mountain View, California			✓

[34] In addition to the test nodes used to support the Measuring Broadband America Program, SamKnows utilizes a diverse fleet of nodes in locations around the globe for other international programs.

New York City, New York	✓	✓	✓
San Jose, California		✓	
Seattle, Washington			✓
Washington D.C	✓	✓	
Washington, Virginia			✓
Denver, Colorado			✓

Table 1: Test Node Locations

SamKnows also has access to many test nodes donated by ISPs around the world. These particular test nodes reside within individual ISP networks and are therefore considered on-net test nodes.

ISPs have the advantage of measuring to both on-net and off-net test nodes, which allows them to segment end-to-end network performance and determine the performance of their own network versus third party networks. For example, an ISP can see what impact third party networks have on their end-users Quality of Experience ('QoE') by placing test nodes within their own network and at major National and International peering locations.

Diagram 1 below shows this set-up.

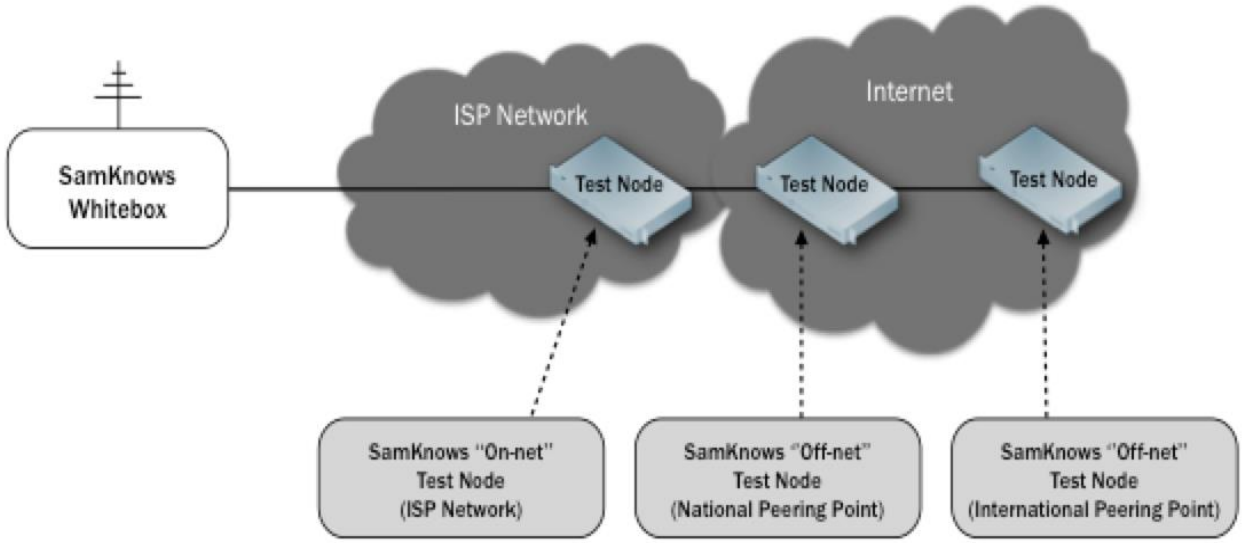

Diagram 1: On-net and Off-net Testing

Both the on-net and off-net test nodes are monitored by SamKnows as part of the global test node fleet. Test node management is explained in more detail within the next section of this document.

3 - Test Node Management

SamKnows test node infrastructure is a critical element of the SamKnows global measurement platform and includes extensive monitoring in place. SamKnows uses a management tool to control and configure the test nodes, while the platform is closely scrutinized using the Nagios monitoring application. System alerts are also in place to ensure the test node infrastructure is always available and operating well within expected threshold bounds.

The SamKnows Operations team continuously checks all test nodes to monitor capacity and overall health. Also included is data analysis to safeguard data accuracy and integrity. This level of oversight not only helps to maintain a healthy, robust platform but also allows us to spot and flag actual network issues and events as they happen. Diagnostic information also supports the Program managers' decision-making process for managing the impact of data accuracy and integrity incidents. This monitoring and administration is fully separate from any monitoring and administration of operating systems and platforms that may be necessary by hosting entities with which SamKnows may be engaged.

3.1 Seamless Test Node Management

SamKnows controls its network of test nodes via a popular open-source management tool called Puppet (**https://puppetlabs.com**). Puppet allows the SamKnows Operations team to easily

manage hundreds of test nodes and ensure that each group of test nodes is configured properly as per each project requirement. Coded in Python, Puppet uses a low-overhead agent installed on each test node that regularly communicates with the controlling SamKnows server to check for updates and ensure the integrity of the configuration.

This method of managing our test nodes allows us to deal with the large number of test nodes without affecting the user's performance in any way. We are also able to quickly and safely make changes to large parts of our test node fleet while ensuring that only the relevant test nodes are updated. This also allows us to keep a record of changes and rapidly troubleshoot any potential problems.

3.2 Proactive Test Node Monitoring

While Puppet handles the configuration and management of the test nodes, Nagios (the most popular online monitoring application) is used by SamKnows to monitor the test nodes. Each test node is configured to send Nagios regular status updates on core metrics such as CPU usage, disk space, free memory, and SamKnows-specific applications. Nagios will also perform active checks of each test nodes where possible, providing us with connectivity information—both via "ping" and connections to any webserver that may be running on the target host.

4 - Test Node Specification and Connectivity

SamKnows maintains a standard specification for all test nodes to ensure consistency and accuracy across the fleet.

4.1 SamKnows test node specifications

All dedicated test nodes must meet the following minimum specifications:

- CPU: Dual core Xeon (2 GHz+)
- RAM: 4 GB
- Disk: 80 GB
- Operating System: CentOS/RHEL 6.x
- Connectivity: Gigabit Ethernet connectivity, with gigabit upstream link.

4.2 Level3 test node specifications

All test nodes provided by level3 meet the following minimum specifications:

- CPU: 2.2 GHz Dual Core
- RAM: 4GB
- Disk: 10 GB

- Operating System: CentOS 6 (64bit)
- Connectivity: 4x1 Gigabit Ethernet (LAG protocol)

4.3 Measurement-Lab Test Node Specifications

All test nodes provided by Measurement-Lab meet the following minimum specifications:

- CPU: 2 GHz 8-core CPU
- RAM: 8 GB
- Disk: 2x100 GB
- OS: CentOS 6.4
- Connectivity: some locations 1 Gbps, some locations 10 Gbps

4.4 Test Node Connectivity

Measurement test nodes must be connected to a Tier-1 or equivalently neutral peering point. Each test node must be able to sustain 1 Gbps throughput.

At minimum, one publicly routable IPv4 address must be provisioned per-test node. The test node must not be presented with a NAT'd address. It is highly preferable for any new test nodes to also be provisioned with an IPv6 address at installation time.

It is preferred that the test nodes do not sit behind a firewall. If a firewall is used, then care must be taken to ensure that it can sustain the throughput required above.

4.5 Test Node Security

Each of the SamKnows test nodes is firewalled using the IPTables linux firewall. We close any ports that are not required, restrict remote administration to SSH only, and ensure access is only granted from a limited number of specified IP addresses. Only ports that require access from the outside world—for example TCP Port 80 on a webserver—would have that port fully open. SamKnows regularly checks its rulesets to ensure that there are no outdated rules and that the access restriction is up to date.

SamKnows accounts on each test node are restricted to the systems administration team by default. When required for further work, an authorized SamKnows employee will have an account added.

5 - Test Node Provisioning

SamKnows also has a policy of accepting test nodes provided by network operators providing that

- The test node meets the specifications outlined earlier

- Minimum of 1 Gbps upstream is provided and downstream connectivity to national peering locations

Please note that donated test nodes may also be subject to additional local requirements.

5.1 Installation and Qualification

ISPs are requested to complete an information form for each test node they wish to provision. This will be used by SamKnows to configure the test node on the management system.

SamKnows will then provide an installation script and an associated installation guide. This will require minimal effort from the ISPs involved and will take a very similar form to the package used on existing test nodes.

Once the ISP has completed installation, SamKnows will verify the test node meets performance requirements by running server-to-server tests from known-good servers. These server-to-server measurements will be periodically repeated to verify performance levels.

5.2 Test Node Access and Maintenance

ISPs donating test nodes are free to maintain and monitor the test nodes using their existing toolsets, providing that these do not interfere with the SamKnows measurement applications or system monitoring tools. ISPs must not run resource intensive processes on the test nodes (*e.g.*, packet captures), as this may affect measurements.

ISPs donating test nodes must ensure that these test nodes are only accessed by maintenance staff when absolutely necessary.

SamKnows requests SSH access to the test nodes, with sudo abilities. sudo is a system administration tool that allows elevated privileges in a controlled granular manner. This has greatly helped diagnosis of performance issues with ISP-provided test nodes historically and would enable SamKnows to be far more responsive in investigating issues.

[DOCUMENT ENDS]

Tenth

Measuring Broadband America

Fixed Broadband Report

A Report on Consumer Fixed Broadband Performance in the United States

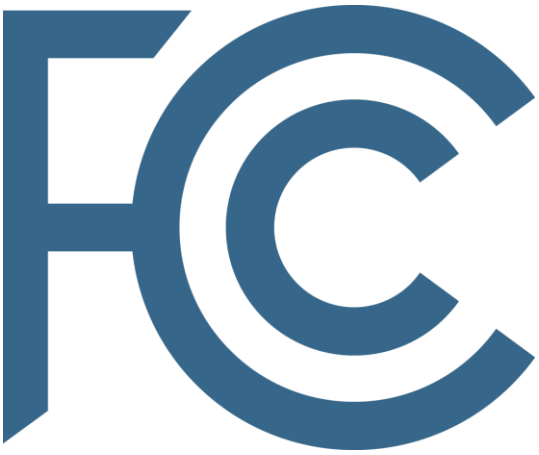

Federal Communications Commission

Office of Engineering and Technology

TABLE OF CONTENTS

List of Charts

List of Tables

1. Executive Summary

The Tenth Measuring Broadband America Fixed Broadband Report ("Tenth Report" or "Report") presents perspectives on empirical performance for data collected in September and October 2019[1] from fixed Internet Service Providers (ISPs), as part of the Federal Communication Commission's (FCC) Measuring Broadband America (MBA) program. This program is an ongoing, rigorous, nationwide study of consumer broadband performance in the United States. The goal of this program is to measure the network performance realized on a representative sample of service offerings and the residential broadband consumer demographic across the country.[2] This representative sample is referred to as the MBA 'panel'. Thousands of volunteer panelists are drawn from the subscriber bases of ISPs which collectively serve a large percentage of the residential marketplace.[3]

The initial Measuring Broadband America Fixed Broadband Report was published in August 2011,[4] and presented the first broad-scale study of directly measured consumer broadband performance throughout the United States. As part of an open data program, all methodologies used in the program are fully documented, and all data collected is published for public use without any restrictions. Including this current Report, ten reports have now been issued.[5] These reports provide a snapshot of fixed broadband Internet access service performance in the United States utilizing a comprehensive set of performance metrics. The resulting performance data is analyzed in a variety of ways that has evolved to make the information more understandable and useful.

A. MAJOR FINDINGS OF THE TENTH REPORT

The key findings of this report are:

- The maximum advertised download speeds amongst the service tiers offered by ISPs and measured by the FCC ranged from 24 Mbps to 940 Mbps for the period covered by this report.

[1] The actual dates used for measurements for this Tenth Report were September 6 – October 3, 2019 (inclusive) plus October 8 – 9, 2019 (inclusive). An isolated server outage forced the exclusion of data from October 4 to 7 to avoid anomalous results.

[2] The sample is representative in that it aims to include those tiers that constitute the top 80% of the subscriber base per ISP. Some tiers accordingly are not included. As with any sample, budget and sample constitution constraints limit completeness of coverage.

[3] At the request of and with the assistance of the State of Hawaii Department of Commerce and Consumer Affairs (DCCA) the state of Hawaii was added to the MBA program in 2017. The ISPs whose performance were measured in the State of Hawaii were Hawaiian Telcom and Oceanic Time Warner Cable (which is now a part of Charter Spectrum).

[4] All reports can be found at https://www.fcc.gov/general/measuring-broadband-america.

[5] The First Report (2011) was based on measurements taken in March 2011, the Second Report (2012) on measurements taken in April 2012, and the Third (2013) through Ninth (2019) Reports on measurements taken in September of the year prior to the reports' release dates. In order to avoid confusion between the date of release of the report and the measurement dates we have shifted last year to numbering the reports. Thus, this year's report is termed the Tenth MBA Report instead of the 2020 MBA Report. Going forward we will continue with a numbered approach and the next report will be termed as the Eleventh Report.

- The weighted average advertised speed of the participating ISPs was 146.1 Mbps, representing an 8% increase from the previous year (Ninth Report) and over 100% increase from two years prior (Eighth Report).
- For most of the major broadband providers that were tested, measured download speeds were 100% or better than advertised speeds during the peak hours (7 p.m. to 11 p.m. local time).
- Ten ISPs were evaluated in this report. Of these Cincinnati Bell and Frontier employed multiple different broadband technologies across the USA. Overall 12 different ISP/technology configurations were evaluated in this report and eight performed at or better than their advertised speed during the peak hours. Only one performed below 90% for actual-to-advertised download speed during the peak hours.
- In addition to providing download and upload speed measurements of ISPs, this report also provides a measure of consistency of measured to advertised speeds of ISPs with the use of our "80/80" metric. The 80/80 metric measures the percentage of the advertised speed that at least 80% of subscribers experience at least 80% of the time over peak periods. Ten of the 12 ISP/technology configurations provide better than 75% of advertised speed to at least 80% of panelists for at least 80% of the time.

These and other findings are described in greater detail within this report.

B. SPEED PERFORMANCE METRICS

Speed (both download and upload) performance continues to be one of the key metrics reported by the MBA. The data presented includes ISP broadband performance as a median[6] of speeds experienced by panelists within a specific service tier. These reports mainly focus on common service tiers used by an ISP's subscribers.[7]

Additionally, consistent with previous Reports, we also compute average per-ISP performance by weighting the median speed for each service tier by the number of subscribers in that tier. Similarly, in calculating the composite average speed taking into account all ISPs in a specific year, the median speed of each ISP is used and weighted by the number of subscribers of that ISP as a fraction of the total number of subscribers across all ISPs.

In calculating these weighted medians, we draw on two sources for determining the number of subscribers per service tier. ISPs may voluntarily contribute subscription demographics per surveyed service tier as the most recent and authoritative data. Many ISPs have chosen to do so.[8] When such

[6] We first determine the mean value over all the measurements for each individual panelist's "whitebox." (Panelists are sent "whiteboxes" that run pre-installed software on off-the-shelf routers that measure thirteen broadband performance metrics, including download speed, upload speed, and latency.) Then for each ISP's speed tiers, we choose the median of the set of mean values for all the panelists/whiteboxes. The median is that value separating the top half of values in a sample set with the lower half of values in that set; it can be thought of as the middle (*i.e.*, most typical) value in an ordered list of values. For calculations involving multiple speed tiers, we compute the weighted average of the medians for each tier. The weightings are based on the relative subscriber numbers for the individual tiers.

[7] Only tiers that contribute to the top 80% of an ISPs total subscribership are included in this report.

[8] The ISPs that provided SamKnows, the FCC's contractor supporting the MBA program, with weights for each of their tiers were: Cincinnati Bell, CenturyLink, Charter, Comcast, Cox Frontier, Optimum, and Windstream.

information has not been provided by an ISP, we instead rely on the FCC's Form 477 data.[9] All facilities-based broadband providers are required to file data with the FCC twice a year (Form 477) regarding deployment of broadband services, including subscriber counts. For this report, we used the June 2019 Form 477 data. It should be noted that the Form 477 subscriber data values generally lag the reporting month, and therefore, there are likely to be small inaccuracies in the tier ratios. It is for this reason that we encourage ISPs to provide us with subscriber numbers for the measurement month.

As in our previous reports, we found that for most ISPs the actual speeds experienced by subscribers either nearly met or exceeded advertised service tier speeds. However, since we started our MBA program, consumers have changed their Internet usage habits. In 2011, consumers mainly browsed the web and downloaded files; thus, we reported mean broadband speeds since these statistics were likely to closely mirror user experience. By contrast, in September-October 2019 (the measurement period for this report) consumer internet usage had become dominated by video consumption, with consumers regularly streaming video for entertainment and education.[10] Therefore, our network performance analytics have been expanded by using consistency in service metrics to better capture the shift in usage patterns. Both the median measured speed metric and consistency in service metrics help to better reflect the consumer's perception and usefulness of Internet access service.

Specifically, we use two kinds of metrics to reflect the consistency of service delivered to the consumer: First, we report the percentage of advertised speed experienced by at least 80% of panelists during at least 80% of the daily peak usage period ("80/80 consistent speed" measure). Second, we show the fraction of consumers who obtain median speeds greater than 95%, between 80% and 95%, and less than 80% of advertised speeds.

A. USE OF OTHER PERFORMANCE METRICS

Although download and upload speeds remain the network performance metric of greatest interest to the consumer, we also spotlight two other key network performance metrics in this report: latency and packet loss. These metrics can significantly affect the overall quality of Internet applications.

Latency is the time it takes for a data packet to travel across a network from one point on the network to another. High latencies may affect the perceived quality of some interactive services such as phone calls over the Internet, video chat and video conferencing, or online multiplayer games. All network access technologies have a minimum latency that is largely determined by the technology There are many other factors that affect latency though, including the location of the server you're communicating with, the route taken to the server, and whether or not there is any congestion on that route. Technology-

[9] For an explanation of Form 477 filing requirements and required data see: https://transition.fcc.gov/form477/477inst.pdf (Last accessed 8/10/2020).

[10] "It is important to track the changing mix of devices and connections and growth in multidevice ownership as it affects traffic patterns. Video devices, in particular, can have a multiplier effect on traffic. An Internet-enabled HD television that draws couple - three hours of content per day from the Internet would generate as much Internet traffic as an entire household today, on an average. Video effect of the devices on traffic is more pronounced because of the introduction of Ultra-High-Definition (UHD), or 4K, video streaming. This technology has such an effect because the bit rate for 4K video at about 15 to 18 Mbps is more than double the HD video bit rate and nine times more than Standard-Definition (SD) video bit rate. We estimate that by 2023, two-thirds (66 percent) of the installed flat-panel TV sets will be UHD, up from 33 percent in 2018" *See Cisco Annual Internet Report (2018-2023) White Paper* , https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html (Last accessed Aug. 8, 2020).

dependent latencies are typically small for terrestrial broadband services and are thus unlikely to affect the perceived quality of applications. Additionally, for certain applications the user experience is not necessarily affected by high latencies. As an example, when using entertainment video streaming applications, because the data can be cached prior to display, the user experience is likely to be unaffected by relatively high latencies

Packet loss measures the fraction of data packets sent that fail to be delivered to the intended destination. Packet loss may affect the perceived quality of applications that do not incorporate retransmission of lost packets, such as phone calls over the Internet, video chat, some online multiplayer games, and some video streaming. High packet loss also degrades the achievable throughput of download and streaming applications. However, packet loss of a few tenths of a percent are unlikely to significantly affect the perceived quality of most Internet applications and are common. During network congestion, both latency and packet loss typically increase.

The Internet continually evolves in its architecture, performance, and services. Accordingly, we will continue to adapt our measurement and analysis methodologies to further improve the collective understanding of performance characteristics of broadband Internet access. By doing so we aim to help the community of interest across the board, from consumers to technologists, service providers and regulators.

2. Summary of Key Findings

A. MOST POPULAR ADVERTISED SERVICE TIERS

A list of the ISP download and upload speed service tiers that were measured in this report are shown in Table 1. It should be noted that while upload and download speeds are measured independently and shown separately, they are typically offered by an ISP in a paired configuration. The service tiers that are included for reporting represent the top 80% (therefore 'most popular') of an ISP's set of tiers based on subscriber numbers. Taken in aggregate, these plans serve the majority of the subscription base of the participating ISPs.

Table 1: List of ISP service tiers whose broadband performance was measured in this report

Tech-nology	Company	Speed Tiers (Download)									Speed Tiers (Upload)					
DSL	CenturyLink	1.5	3	7	8*	10	12	20	25*	40	0.512*	0.768	0.896	2	5	10*
	Cincinnati Bell DSL	5	30*								0.768	3*				
	Frontier DSL	3	6	12	24*						0.768	1	1.5*			
	Windstream	3	6	10	12*	15*	25	50*	100*		0.768*	1	1.5	4*		
Cable	Altice Optimum	100	200	300*							35					
	Charter	100	200	400							10	20				
	Comcast	60	150	250							5	10				
	Cox	30	100*	150*	300						3	10	30			
	Mediacom	60	100	200							5	10	20			
Fiber	Cincinnati Bell Fiber	50	250	500							10	100	125			
	Frontier Fiber	50	75	100	150	200					50	75	100	150	200	
	Verizon Fiber	50*	75	100	940**						50*	75	100	880**		

*Tiers that lack sufficient panelists to meet the program's target sample size.

** Although Verizon Fiber's 940/880 Mbps service tier was amongst the top 80% of Verizon's offered tiers by subscription numbers, it is not included in the report charts because technical methodologies for measuring high speed rates near Gigabit and above have not yet been established for the MBA program.

Chart 1.1 (below) displays the weighted (by subscriber numbers) mean of the top 80% advertised download speed tiers for each participating ISP for the last three years (September 2017 to September-October 2019) grouped by the access technology used to offer the broadband Internet access service (DSL, cable, or fiber). It should be noted that this chart does not reflect the actual performance of the ISPs and only provides the weighted average of the ISP's advertised speeds. In September-October 2019, the weighted average advertised download speed was 146.1 Mbps among the measured ISPs, which represents a 100% increase from 2017 and a 8% increase compared to the average in September-October 2018 which was 135.7 Mbps.[11]

[11] Please note that this average for September-October 2018 and September 2017 represents the average advertised download speed with AT&T tiers removed. We did this to have a fairer comparison between the years since AT&T is no longer an active participant in the MBA program. The actual weighted average advertised download speed (with AT&T included) for September-October 2018, as reported in the Ninth MBA Report is 123.3 Mbps.

Chart 1.1: Weighted average advertised download speed among the top 80% service tiers offered by each ISP

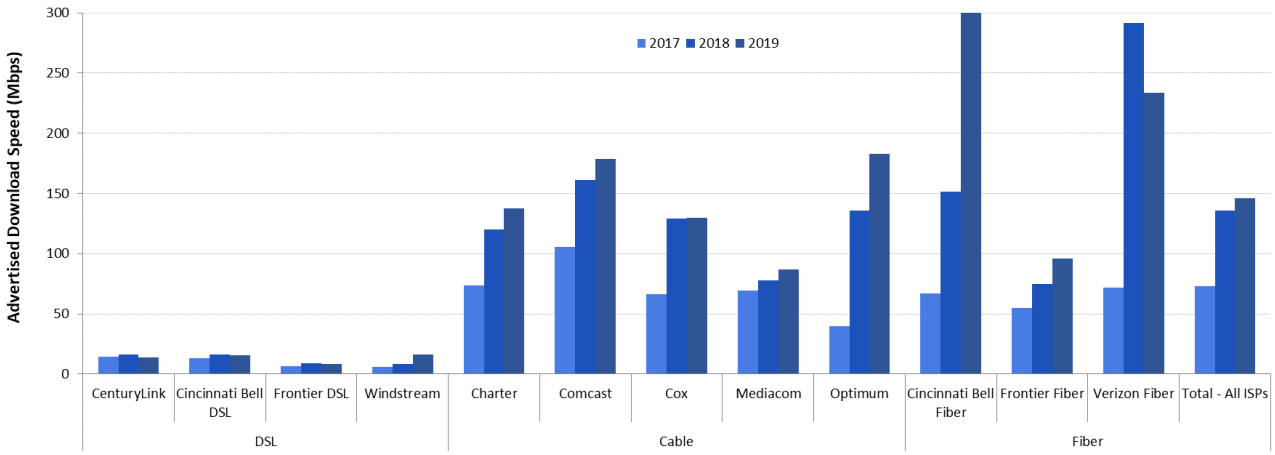

All of the ISPs, except Verizon, showed higher weighted-averages of advertised speeds in September-October 2019 as compared to September 2018. Verizon-fiber showed a slight decrease in 2019 compared to 2018 which was not due to any reduction in the service speed offerings but arose from changes in weighting due to relative shifts in subscriber numbers on the advertised tiers from 2018 to 2019.

It can be seen from Chart 1.1 that the DSL speeds lag far behind the speed of other technologies. In order to better compare the DSL speed offerings by the various ISPs we have added a separate Chart 1.2 drawn to a scale that makes their relative speeds more discernable.

Chart 2.2: Weighted average advertised download speed among the DSL ISPs

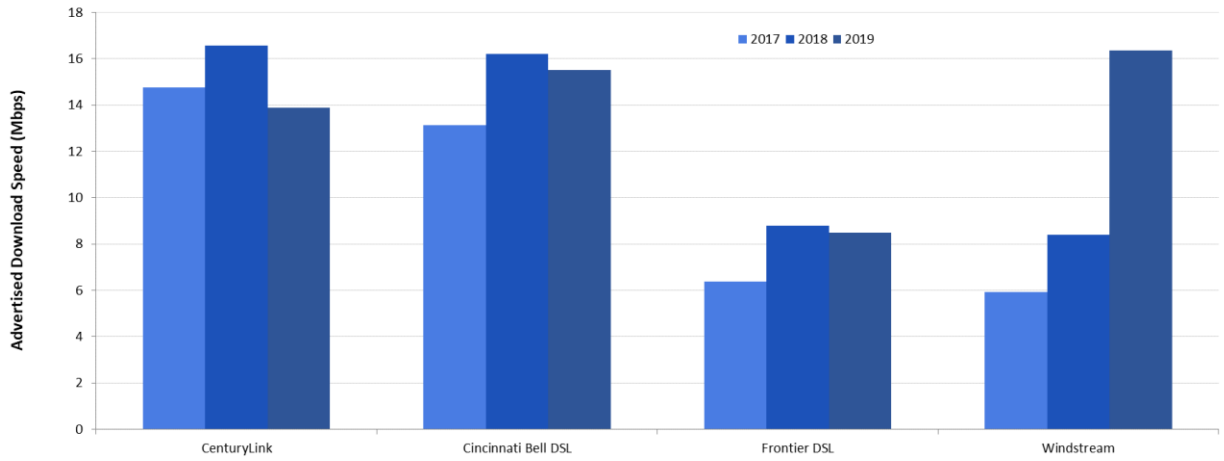

Among participating broadband ISPs, only Cincinnati Bell, Frontier, and Verizon use fiber as the access technology for a substantial number of their customers and their maximum speed offerings range from 200 Mbps to 940 Mbps. A key difference between the fiber vendors and other technology vendors is that (with the exception of Cincinnati Bell), most fiber vendors advertise generally symmetric upload and download speeds. This is in sharp contrast to the asymmetric offerings for all the other technologies, where the upload advertised speeds are typically 5 to 10 times below the download advertised speeds.

As can be seen in Chart 1.1, there is considerable difference between the offered average weighted speed tier by technology. Chart 2 plots the weighted average of the top 80% ISP tiers by technology for the last three years.[12] As can be seen in this chart, most technologies showed increases in the set of advertised download speeds by ISPs. For the September-October 2019 period, the weighted mean advertised speeds for DSL technology was 13 Mbps which lagged considerably behind the weighted mean advertised download speeds for cable and fiber technologies, which were 155 Mbps and 208 Mbps, respectively. Fiber technology showed the greatest increase in speed offerings in 2019 compared to 2017 with a weighted mean going up from 70 Mbps to 208 Mbps representing a nearly 200% increase. This year's (2019) average advertised speed for fiber, however, showed a slight decrease by 17% from last year's (2018) speed. DSL technology speed increased from 11 Mbps to 13 Mbps from 2017 to 2019, a 16% increase overall (though it did show a small 1% decrease in speed this year compared to last year). In comparison, cable technology showed an 12% increase from 2018 to 2019 and an overall 83% increase from 2017 to 2019

Chart 2: Weighted average advertised download speed among the top 80% service tiers based on technology.

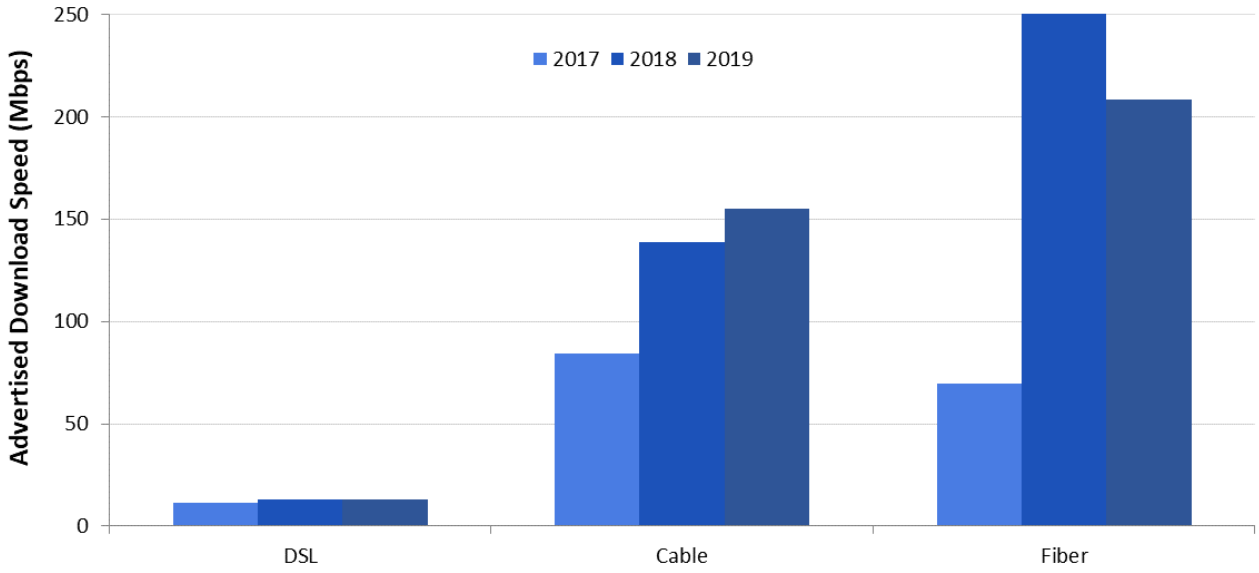

Chart 3 plots the migration of panelists to a higher service tier based on their access technology.[13] Specifically, the horizontal axis of Chart 3 partitions the September 2018 panelists by the advertised download speed of the service tier to which they were subscribed. For each such set of panelists who

[12] Since AT&T is no longer actively participating in the Measuring Broadband America program, we have removed it from previous years' results in Charts 1 and 2. This allows a proper comparison to be made between the results for this year as compared to previous years. It should also be noted that although AT&T IPBB had been characterized in previous reports as a DSL technology it actually included a mix of ADSL2+, VDSL2, G.Fast and Ethernet technologies delivered over a hybrid of fiber optic and copper facilities.

[13] Where several technologies are plotted at the same point in the chart, this is identified as "Multiple Technologies."

also participated in the September-October 2019 collection of data,[14] the vertical axis of Chart 3 displays the percentage of panelists that migrated by September-October 2019 to a service tier with a higher advertised download speed. There are two ways that such a migration could occur: (1) if a panelist changed their broadband plan during the intervening year to a service tier with a higher advertised download speed, or (2) if a panelist did not change their broadband plan but the panelist's ISP increased the advertised download speed of the panelist's subscribed plan.[15]

Chart 3 shows that the percentage of panelists subscribed in September-October 2018 who moved to higher tiers in September-October 2019 was between 3% to 26% for DSL subscribers, 4% to 100% for cable subscribers and 16% to 50% for fiber subscribers. In addition, 1% to 8% subscribers migrated to a higher speed tier using a different technology from what they had in September 2018.

Chart 3: Consumer migration to higher advertised download speeds

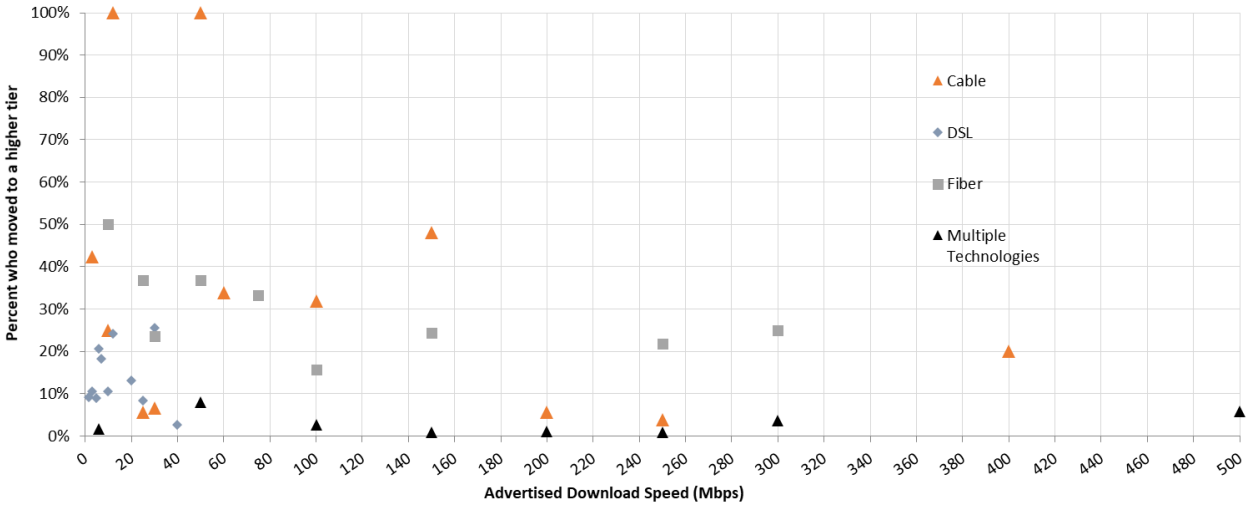

B. MEDIAN DOWNLOAD SPEEDS

Advertised download speeds may differ from the speeds that subscribers actually experience. Some ISPs more consistently meet network service objectives than others or meet them unevenly across their geographic coverage area. Also, speeds experienced by a consumer may vary during the day if the aggregate user demand during busy hours causes network congestion. Unless stated otherwise, all actual speeds were measured only during peak usage periods, which we define as 7 p.m. to 11 p.m. local time.

To compute the average ISP performance, we determine the ratio of the median speed for each tier to the advertised tier speed and then calculate the weighted average of these based on the subscriber count per tier. Subscriber counts for the weightings were provided from the ISPs themselves or, if unavailable, from FCC Form 477 data.

Chart 4 shows the ratio of the measured median download and upload speeds experienced by an ISP's subscribers to that ISP's advertised download and upload speeds weighted by the subscribership numbers for the tiers. The actual speeds experienced by most ISPs' subscribers are close to or exceed the

[14] Of the 5,855 panelists who participated in the September 2018 collection of data, 4,246 panelists continued to participate in the September-October 2019 collection of data.

[15] We do not attempt here to distinguish between these two cases.

advertised speeds. However, DSL broadband ISPs continue to advertise "up-to" speeds that on average exceed the actual speeds experienced by their subscribers. Out of the 12 ISP/technology configurations shown, 8 met or exceeded their advertised download speed and three more reached at least 90% of their advertised download speed. Only Cincinnati-DSL (at 79%) performed below 90% of its advertised download speed.

Chart 4: The ratio of weighted median speed (download and upload) to advertised speed for each ISP.

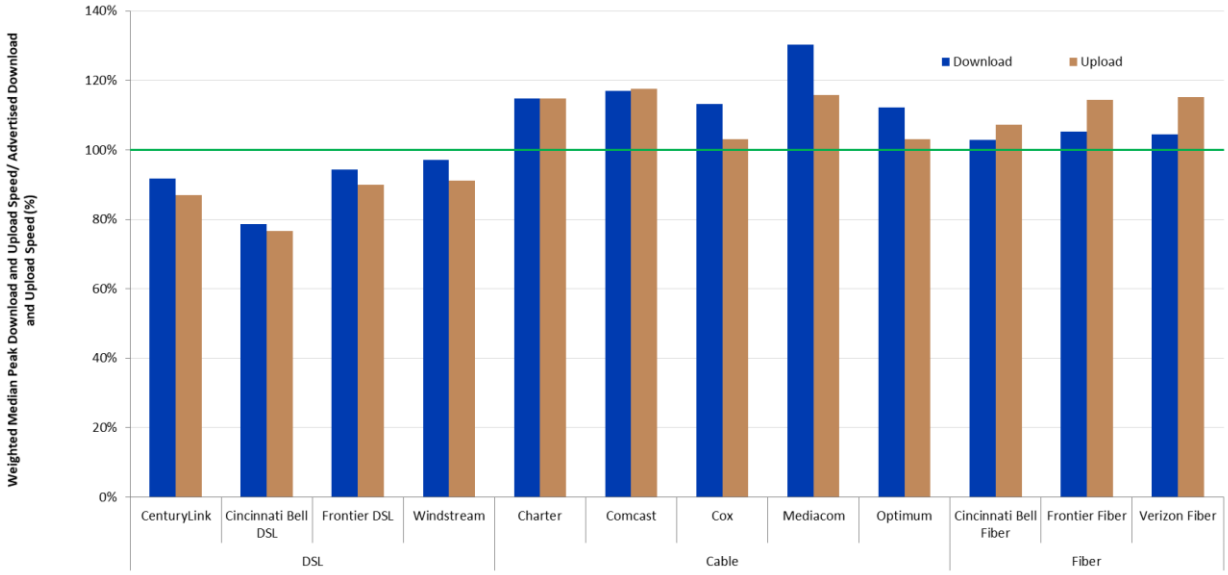

C. VARIATIONS IN SPEEDS

As discussed earlier, actual speeds experienced by individual consumers may vary by location and time of day. Chart 5 shows, for each ISP, the percentage of panelists who experienced a median download speed (averaged over the peak usage period during our measurement period) that was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed.

Chart 5: The percentage of consumers whose median download speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed

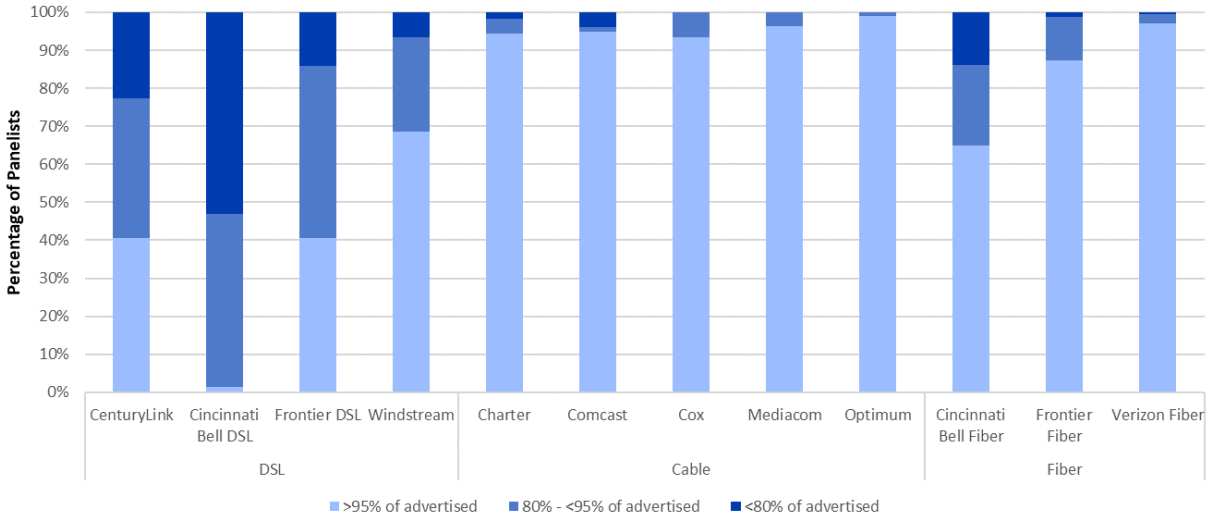

ISPs using DSL technology had between 2% to 69% of their subscribers getting greater than or equal to 95% of their advertised download speeds during peak hours. ISPs using cable technology and fiber technology had between 93%-99% and between 65%-97%, respectively, of their subscribers getting equal to or better than 95% of their advertised download speeds.

Though the median download speeds experienced by most ISPs' subscribers nearly met or exceeded the advertised download speeds, there are some customers of each ISP for whom the median download speed fell significantly short of the advertised download speed. Relatively few subscribers of cable service experienced this. The best performing ISPs, when measured by this metric, are Charter, Comcast, Cox, Mediacom, Optimum, Frontier-Fiber and Verizon-Fiber; more than 80% of their panelists were able to attain an actual median download speed of at least 95% of the advertised download speed.

In addition to variations based on a subscriber's location, speeds experienced by a consumer may fluctuate during the day. This is typically caused by increased traffic demand and the resulting stress on different parts of the network infrastructure. To examine this aspect of performance, we use the term "80/80 consistent speed." This metric is designed to assess temporal and spatial variations in measured values of a user's download speed.[16] While consistency of speed is in itself an intrinsically valuable service characteristic, its impact on consumers will hinge on variations in usage patterns and needs. As an example, a good consistency of speed measure is likely to indicate a higher quality of service experience for internet users consuming video content.

Chart 6 summarizes, for each ISP, the ratio of 80/80 consistent median download speed to advertised download speed, and, for comparison, the ratio of median download speed to advertised download speed shown previously in Chart 4. The ratio of 80/80 consistent median download speed to advertised download speed is less than the ratio of median download speed to advertised download speed for all participating ISPs due to congestion periods when median download speeds are lower than the overall average. When the difference between the two ratios is small, the median download speed is fairly insensitive to both geography and time. When the difference between the two ratios is large, there is a

[16] For a detailed definition and discussion of this metric, please refer to the Technical Appendix.

greater variability in median download speed, either across a set of different locations or across different times during the peak usage period at the same location.

Chart 6: The ratio of 80/80 consistent median download speed to advertised download speed.

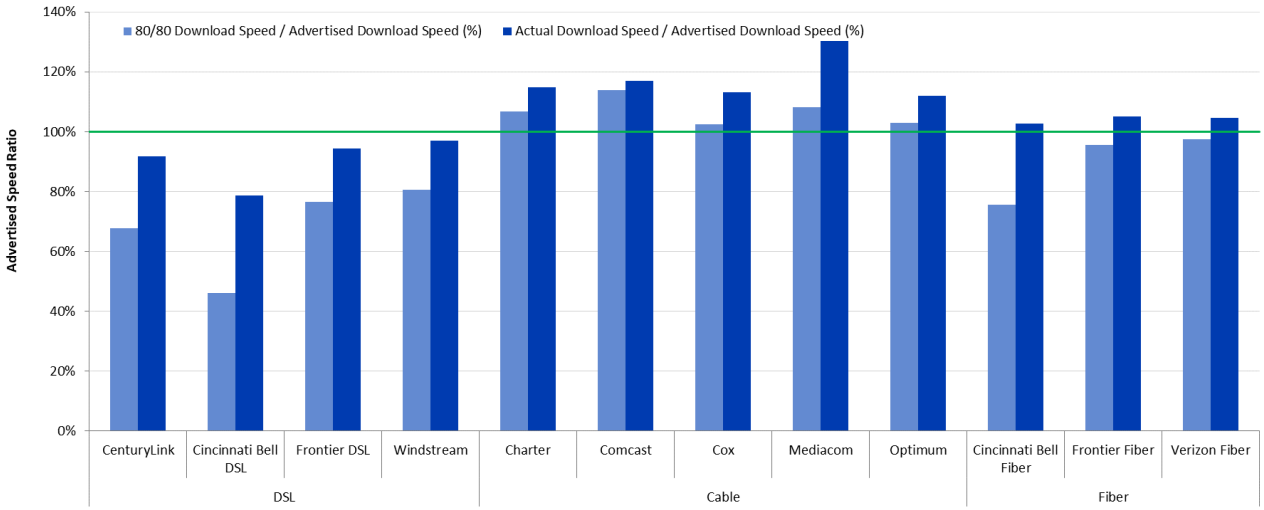

Customers of Charter, Comcast, Cox, Mediacom and Optimumexperienced median download speeds that were very consistent; *i.e.*, they provided greater than 100% of the advertised speed during peak usage period to more than 80% of panelists for more than 80% of the time. As can be seen in Chart 6 cable and fiber ISPs performed better than DSL ISPs with respect to their 80/80 consistent speeds. For example, for September-October 2019, the 80/80 consistent download speed for Cincinnati Bell DSL was 46% of the advertised speed.

D. LATENCY

The latency between any two points in the network is the time it takes for a packet to travel from one point to the other. It has a fixed component that depends on the distance, the transmission speed, and transmission technology between the source and destination, and a variable component due to queuing delay that increases as the network path congests with traffic. The MBA program measures latency by measuring the round-trip time between the consumer's home and the closest measurement server.

Chart 7 shows the median latency for each participating ISP. In general, higher-speed service tiers have lower latency, as it takes less time to transmit each packet. The median latencies ranged from 10 ms to 27 ms in our measurements (with the exception of CenturyLink DSL and Cincinnati Bell DSL which had median latencies of 40 ms and 34 ms, respectively).

Chart 7: Latency by ISP

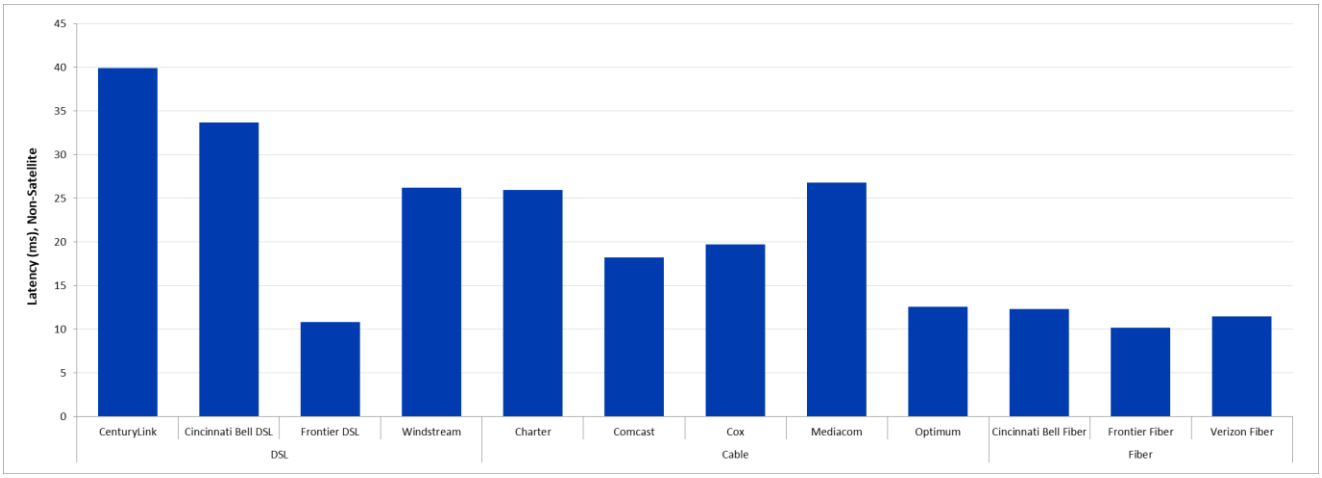

DSL latencies (between 11 ms to 40 ms) were slightly higher than those for cable (13 ms to 27 ms). Fiber ISPs showed the lowest latencies (10 ms to 12 ms). The differences in median latencies among terrestrial-based broadband services are relatively small and are unlikely to affect the perceived quality of highly interactive applications.

E. PACKET LOSS

Packet loss is the percentage of packets that are sent by a source but not received at the intended destination. The most common causes of packet loss are congestion leading to buffer overflows or active queue management along the network path. Alternatively, high latency might lead to a packet being counted as lost if it does not arrive within a specified window. A small amount of packet loss is expected, and indeed packet loss is commonly used by some Internet protocols such as TCP to infer Internet congestion and to adjust the sending rate to mitigate the offered load, thus lessening the contribution to congestion and the risk of lost packets. The MBA program uses an active UDP-based packet loss measurement method and considers a packet lost if it is not returned within 3 seconds.

Chart 8 shows the average peak-period packet loss for each participating ISP, grouped into bins. We have broken the packet loss performance into three bands, allowing a more granular view of the packet loss performance of the ISP network. The breakpoints for the three bins used to classify packet loss have been chosen with an eye towards balancing commonly accepted packet loss thresholds for specific services and provider packet loss Service Level Agreements (SLAs) for enterprise services, as consumer offerings are not typically accompanied by SLAs. Specifically, the 1% standard for packet loss is commonly accepted as the point at which highly interactive applications such as VoIP experience significant degradation in quality according to industry publications and international (ITU) standards.[17] The 0.4% breakpoint was chosen as middle ground between the highly desirable performance of 0% packet loss described in many documents (for Voice over Internet Protocol (VoIP)) and the 1% unacceptable limit on the high side. The specific value of 0.4% is also generally supported by major ISP SLAs for network performance. Indeed,

[17] See: http://www.itu.int/dms_pubrec/itu-r/rec/m/r-rec-m.1079-2-200306-i!!msw-e.doc.

most SLAs support 0.1% to 0.3% packet loss guarantees,[18] but these are generally for enterprise level services which entail business-critical applications that require some service guarantees.

Chart 8: Percentage of consumers whose peak-period packet loss was less than 0.4%, between 0.4% to 1%, and greater than 1%.

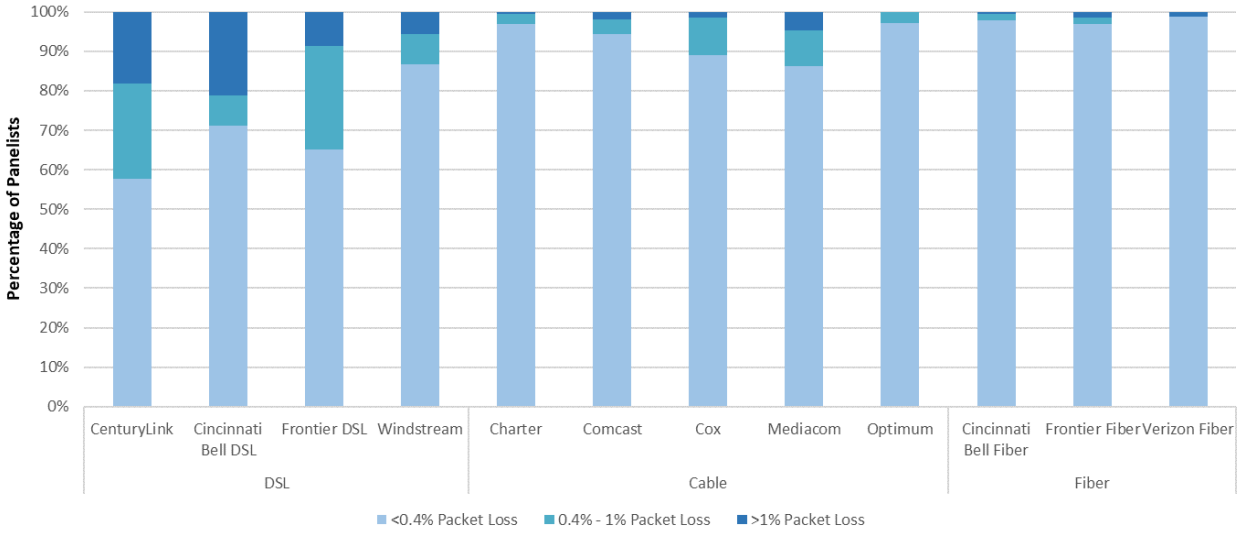

Chart 8 shows that ISPs using fiber technology have the lowest packet loss, and that ISPs using DSL technology tend to have the highest packet loss. As shown in this chart, 6% to 21% of DSL subscribers experience 1% or greater packet loss. The corresponding numbers for cable and fiber are 0% to 5% and 0% to 1.5%, respectively. Within a given technology class, packet loss also varies among ISPs.

F. WEB BROWSING PERFORMANCE

The MBA program also conducts a specific test to gauge web browsing performance. The web browsing test accesses nine popular websites that include text and images, but not streaming video. The time required to download a webpage depends on many factors, including the consumer's in-home network, the download speed within an ISP's network, the web server's speed, congestion in other networks outside the consumer's ISP's network (if any), and the time required to look up the network address of the webserver. Only some of these factors are under control of the consumer's ISP. Chart 9 displays the average webpage download time as a function of the advertised download speed. As shown by this chart, webpage download time decreases as download speed increases, from about 9.8 seconds at 1.5 Mbps download speed to about 1.5 seconds for 25 Mbps download speed. Subscribers to service tiers exceeding 25 Mbps experience slightly smaller webpage download times decreasing to 1 – 1.25 seconds at 150 Mbps. Beyond 150 Mbps, the webpage download times decrease only by minor amounts. These download times assume that only a single user is using the Internet connection when the webpage is downloaded, and does not account for more common scenarios, where multiple users within a household are simultaneously using the Internet connection for viewing web pages, as well as other applications such as real-time gaming or video streaming.

[18] See: https://www.voip-info.org/wiki/view/QoS and http://www.ciscopress.com/articles/article.asp?p=357102.

Chart 9: Average webpage download time, by advertised download speed.

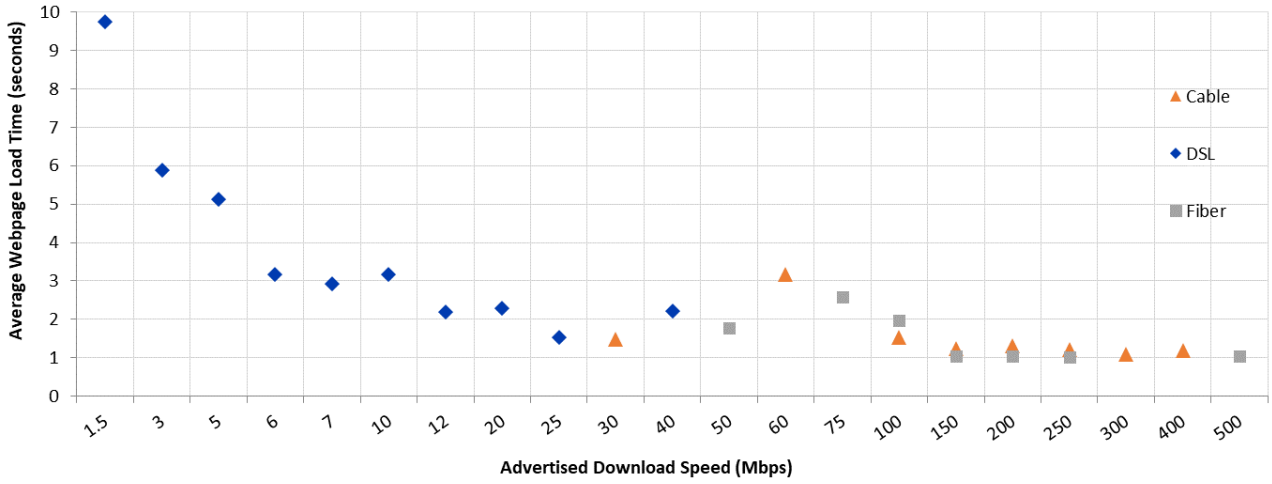

3. Methodology

A. PARTICIPANTS

Ten ISPs actively participated in the Fixed MBA program in September-October 2019.[19] They were:

- CenturyLink
- Charter Communications
- Cincinnati Bell
- Comcast
- Cox Communications
- Frontier Communications Company
- Mediacom Communications Corporation
- Optimum
- Verizon
- Windstream Communications

The methodologies and assumptions underlying the measurements described in this Report are reviewed at meetings that are open to all interested parties and documented in public ex parte letters filed in the GN Docket No. 12-264. Policy decisions regarding the MBA program were discussed at these meetings prior to adoption, and involved issues such as inclusion of tiers, test periods, mitigation of operational issues affecting the measurement infrastructure, and terms-of-use notifications to panelists. Participation in the MBA program is open and voluntary. Participants include members of academia, consumer equipment vendors, telecommunications vendors, network service providers, consumer policy groups, as well as our contractor for this project, SamKnows. In 2019-2020, participants at these meetings (collectively and informally referred to as "the broadband collaborative"), included all eleven participating ISPs and the following additional organizations:

- Level 3 Communications ("Level 3"), now part of CenturyLink
- Massachusetts Institute of Technology ("MIT")
- Measurement Lab (M-Lab)
- StackPath
- NCTA – The Internet & Television Association ("NCTA")
- New America Foundation
- Princeton University
- United States Telecom Association ("US Telecom")
- University of California - Santa Cruz

Participants have contributed in important ways to the integrity of this program and have provided valuable input to FCC decisions for this program. Initial proposals for test metrics and testing platforms were discussed and critiqued within the broadband collaborative. M-Lab and Level 3 contributed their core network testing infrastructure, and both parties continue to provide invaluable assistance in helping to define and implement the FCC testing platform. We thank all the participants for their continued contributions to the MBA program.

[19]While Hawaiian Telcom participated in the Fixed MBA program, we did not report on it since we did not have sufficient number of panelists on Hawaiian Telcom tiers to have a statistically valid dataset.

B. MEASUREMENT PROCESS

The measurements that provided the underlying data for this report were conducted between MBA measurement clients and MBA measurement servers. The measurement clients (*i.e.*, whiteboxes) were situated in the homes of 6,006 panelists each of whom received service from one of the 11 evaluated ISPs. The evaluated ISPs collectively accounted for over 80% of U.S. residential broadband Internet connections. After the measurement data was processed (as described in greater detail in the Technical Appendix), test results from 3,075 panelists were used in this report.

The measurement servers used by the MBA program were hosted by StackPath, M-Lab, and Level 3 Communications, and were located in thirteen cities (often with multiple locations within each city) across the United States near a point of interconnection between the ISP's network and the network on which the measurement server resided.

The measurement clients collected data throughout the year, and this data is available as described below. However, only data collected from September 6 – October 3, 2019 (inclusive) plus October 8 – 9, 2019 (inclusive), referred to throughout this report as the "September-October 2019" reporting period, were used to generate the charts in this Report.[20]

Broadband performance varies with the time of day. At peak hours, more people tend to use their broadband Internet connections, giving rise to a greater potential for network congestion and degraded user performance. Unless otherwise stated, this Report focuses on performance during peak usage period, which is defined as weeknights between 7:00 p.m. to 11:00 p.m. local time at the subscriber's location. Focusing on peak usage period provides the most useful information because it demonstrates what performance users can expect when the Internet in their local area experiences the highest demand from users.

Our methodology focuses on the network performance of each of the participating ISPs. The metrics discussed in this Report are derived from active measurements, *i.e.*, test-generated traffic flowing between a measurement client, located within the modem/router within a panelist's home, and a measurement server, located outside the ISP's network. For each panelist, the tests automatically choose the measurement server that has the lowest latency to the measurement client. Thus, the metrics measure performance along the path followed by the measurement traffic within each ISP's network, through a point of interconnection between the ISP's network and the network on which the chosen measurement server is located. However, the service performance that a consumer experiences could differ from our measured values for several reasons.

First, as noted, in the course of each test instance we measure performance only to a single measurement server rather than to multiple servers. This is consistent with the approach chosen by most network measurement tools. As a point of comparison, the average web page may load its content from a multiplicity of end points.

---

[20] This proposed time period avoids the dates in early September when parts of North Carolina and Florida were affected by Hurricanes Florence and Michael. It also avoided the increased traffic resulting from latest iOS release which also took place in early September. Omitting dates during these periods was done consistent with the FCC's data collection policy for fixed MBA data. *See* FCC, Measuring Fixed Broadband, Data Collection Policy, https://www.fcc.gov/general/measuring-broadband-america-measuring-fixed-broadband (explaining that the FCC has developed policies to deal with impairments in the data collection process with potential impact for the validity of the data collected).

In addition, bottlenecks or congestion points in the full path traversed by consumer application traffic might also impact a consumer's perception of Internet service performance. These bottlenecks may exist at various points: within the ISP's network, beyond its network (depending on the network topology encountered *en route* to the traffic destination), in the consumer's home, on the Wi-Fi used to access the in-home access router, or from a shortfall of capacity at the far end point being accessed by the application. The MBA tests explore how a service performs from the point at which a fixed ISP's Internet service is delivered to the home on fixed infrastructure (deliberately excluding Wi-Fi, due to the many confounding factors associated with it) to the point at which the test servers are located. As MBA tests are designed to focus on the access to the ISP's network, they will not include phenomena at most interconnection points or transit networks that consumer traffic may traverse.

To the extent possible[21] the MBA focuses on performance within an ISP's network. It should be noted that the overall performance a consumer experiences with their service can also be affected by congestion such as may arise at other points in the path potentially taken by consumer traffic (*e.g.*, in-home Wi-Fi, peering points, transit networks, *etc.*) but this does not get reflected in MBA measurements.

A consumer's home network, rather than the ISP's network, may be the bottleneck with respect to network congestion. We measure the performance of the ISP's service delivered to the consumer's home network, but this service is often shared simultaneously among multiple users and applications within the home. In-home networks, which typically include Wi-Fi, may not have sufficient capacities to support peak loads.[22]

In addition, consumers' experience of ISP performance is manifested through the set of applications they utilize. The overall performance of an application depends not only on the network performance (*i.e.*, raw speed, latency, or packet loss), but also on the application's architecture and implementation and on the operating system and hardware on which it runs. While network performance is considered in this Report, application performance is generally not.

C. MEASUREMENT TESTS AND PERFORMANCE METRICS

This Report is based on the following measurement tests:

- Download speed: This test measures the download speed of each whitebox over a 10-second period, once per hour during peak hours (7 p.m. to 11 p.m.) and once during each of the following periods: midnight to 6 a.m., 6 a.m. to noon, and noon to 6 p.m. The download speed measurement results from each whitebox are then averaged across the measurement month;

[21] The MBA program uses test servers that are both neutral (*i.e.*, operated by third parties that are not ISP-operated or owned) and located as close as practical, in terms of network topology, to the boundaries of the ISP networks under study. As described earlier in this section, a maximum of two interconnection points and one transit network may be on the test path. If there is congestion on such paths to the test server, it may impact the measurement, but the cases where it does so are detectable by the test approach followed by the MBA program, which uses consistent longitudinal measurements, comparisons with control servers located on-net and trend analyses of averaged results. Details of the methodology used in the MBA program are given in the Technical Appendix to this report.

[22] Independent research, drawing on the FCC's MBA test platform, suggests that home networks are a significant source of end-to-end service congestion. *See* Srikanth Sundaresan et al., *Home Network or Access Link? Locating Last-Mile Downstream Throughput Bottlenecks*, PAM 2016 - Passive and Active Measurement Conference, at 111-123 (Mar. 2016). Numerous instances of research supported by the fixed MBA test platform are described at https://www.fcc.gov/general/mba-assisted-research-studies.

and the median value for these average speeds across the entire set of whiteboxes on a given tier is used to determine the *median measured download speed* for that tier. The overall ISP download speed is computed as the weighted median for each service tier, using the subscriber counts for the tiers as weights.

- Upload speed: This test measures the upload speed of each whitebox over a 10-second period, which is the same measurement interval as the download speed. The upload speed measured in the last five seconds of the 10-second interval is retained, the results of each whitebox are then averaged over the measurement period, and the median value for the average speed taken over the entire set of whiteboxes is used to determine the *median upload speed* for a service tier. The ISP upload speed is computed in the same manner as the download speed.
- Latency and packet loss: These tests measure the round-trip times for approximately 2,000 packets per hour sent at randomly distributed intervals. Response times less than three seconds are used to determine the mean latency. If the whitebox does not receive a response within three seconds, the packet is counted as lost.
- Web browsing: The web browsing test measures the total time it takes to request and receive webpages, including the text and images, from nine popular websites and is performed once every hour. The measurement includes the time required to translate the web server name (URL) into the webserver's network (IP) address.

This Report focuses on three key performance metrics of interest to consumers of broadband Internet access service, as they are likely to influence how well a wide range of consumer applications work: download and upload speed, latency, and packet loss. Download and upload speeds are also the primary network performance characteristic advertised by ISPs. However, as discussed above, the performance observed by a user in any given circumstance depends not only on the actual speed of the ISP's network, but also on the performance of other parts of the Internet and on that of the application itself.

The standard speed tests use TCP with 8 concurrent TCP sessions. In 2017 we also introduced a less-data intensive throughput test, which both generated less traffic and ran less frequently and thereby provided less strain on consumer accounts that are data-capped. The Lightweight tests are used exclusively to provide broadband performance results for satellite ISPs. The Technical Appendix to this Report describes each test in more detail, including additional tests not contained in this Report.

D. AVAILABILITY OF DATA

The MBA panel sample used in the reporting period is validated (i.e., upload and download tiers of the whiteboxes are verified with providers) and the measurement results are carefully inspected to eliminate misleading outliers. This leads to a 'validated data set' that accompanies each report. The Validated Data Set[23] on which this Report is based, as well as the full results of all tests, are available at http://www.fcc.gov/measuring-broadband-america. For interested parties, as tests are run 24x7x365, we also provide raw data (referred to as such because cross-checks are not done except in the test period used for the report, thus subscriber tier changes may be missed) for the reference month and other months. Previous reports of the MBA program, as well as the data used to produce them, are also available there.

Both the Commission and SamKnows, the Commission's contractor for this program, recognize that, while the methodology descriptions included in this document provide an overview of the project, interested

[23] The September-October 2019 data set was validated to remove anomalies that would have produced errors in the Report. This data validation process is described in the Technical Appendix.

parties may be willing to contribute to the project by reviewing the software used in the testing. SamKnows welcomes review of its software and technical platform, consistent with the Commission's goals of openness and transparency for this program.[24]

[24] The software that was used for the MBA program will be made available for noncommercial purposes. To apply for noncommercial review of the code, interested parties may contact SamKnows directly at team@samknows.com, with the subject heading "Academic Code Review."

4. Test Results

A. MOST POPULAR ADVERTISED SERVICE TIERS

Chart 1 above summarizes the weighted average of the advertised download speeds[25] for each participating ISP, for the last 3 years (September 2017 to September-October 2019) where the weighting is based upon the number of subscribers to each tier, grouped by the access technology used to offer the broadband Internet access service (DSL, cable, or fiber). Only the top 80% tiers (by subscriber number) of each ISP were included. Chart 10 below shows the corresponding weighted average of the advertised upload speeds among the measured ISPs. The computed weighted average of the advertised upload speed of all the ISPs is 30.5 Mbps representing a 133% increase compared to 13.1 Mbps in 2017. However, the computed average weighted upload speed decreased slightly this year by 4% over the previous year's value of 31.9 Mbps.[26]

Chart 10.1: Weighted average advertised upload speed among the top 80% service tiers offered by each ISP.

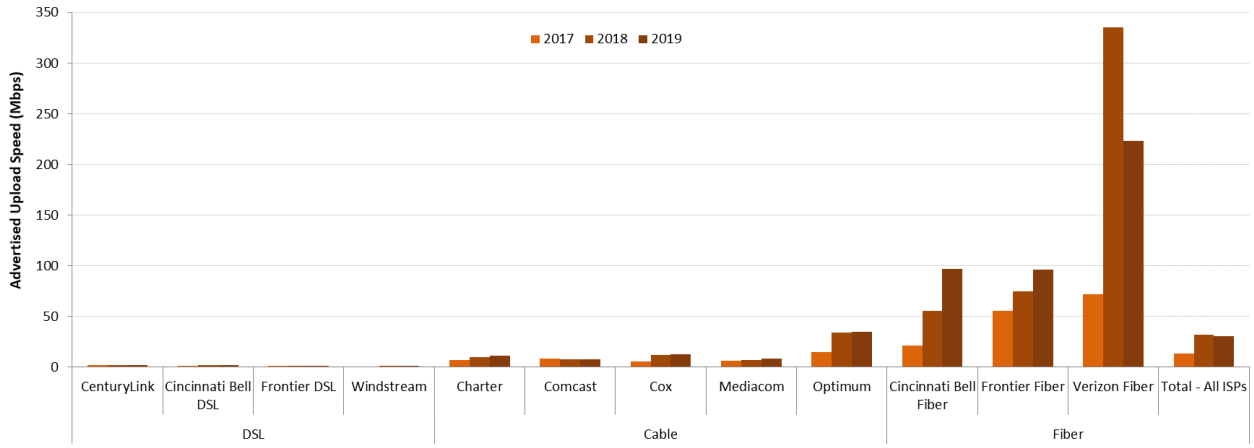

Due to the relatively high upload speeds for optical technology, it is difficult to discern the variations in speed for both DSL and cable technologies when drawn to the same scale. Separate Charts 10.1 and 10.2 are included here that provide the weighted-average upload speeds for ISPs using DSL and cable technologies, respectively.

Chart 10.2: Weighted average advertised upload speed offered by ISPs using DSL technology.

[25] Measured service tiers were tiers which constituted the top 80% of an ISP's broadband subscriber base.

[26] Please note that this average for Sept-Oct 2018 represents the average advertised upload speed with AT&T tiers removed. We did this to have a fairer comparison between the years since AT&T is no longer an active participant in the MBA program. The actual weighted average upload speed for September-October 2018, as reported in the Ninth MBA Report, is 27.4 Mbps.

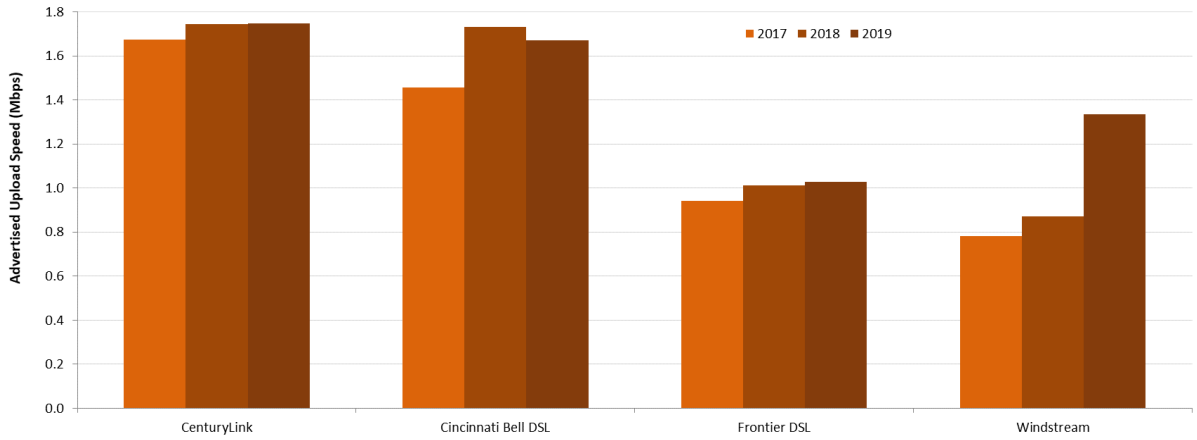

Chart 10.3: Weighted average advertised upload speed offered by ISPs using Cable technology.

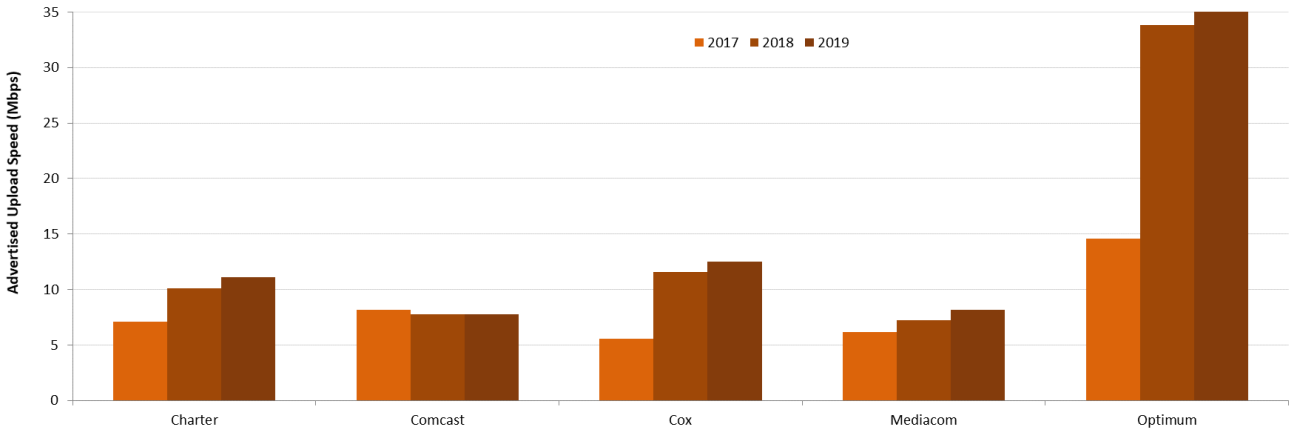

Chart 11 compares the weighted average of the advertised upload speeds by technology for the last 3 years (September 2017 to September-October 2019). As can be seen in this chart, all technologies showed increased rates in 2019 as compared to 2017. However, the rates of increase were not the same for all technologies. The rate of increase in the weighted average of Fiber technology was 189% compared to DSL and Cable which were 11% and 43%, respectively. Comparing the 2019 results with the previous year's (2018) results, we see an increase of offered upload speeds in DSL by 6% to 1.5 Mbps and an increase in cable of 9% to 11 Mbps. However, Fiber upload speed decreased by 29% in 2019 as compared with 2018. This drop in fiber upload speed is due to relative shifts in the number of subscribers to the tiers rather than lowering of offered upload tier speeds. Despite this drop, the advertised fiber upload speeds (194 Mbps) were still far higher than for other technologies.

Observing both the download and upload speeds, it is clear that fiber service tiers are generally symmetric in their actual upload and download speeds. This results from the fact that fiber technology has

significantly more capacity than other technologies and it can be engineered to have symmetric upload and download speeds. For other technologies with more limited capacity, higher capacity is usually allocated to download speeds than to upload speeds, typically in ratios ranging from 5:1 to 10:1. This resulting asymmetry in download/upload speeds is reflective of actual usage because consumers typically download significantly more data than they upload.

Chart 11: Weighted average advertised upload speed among the top 80% service tiers based on technology.

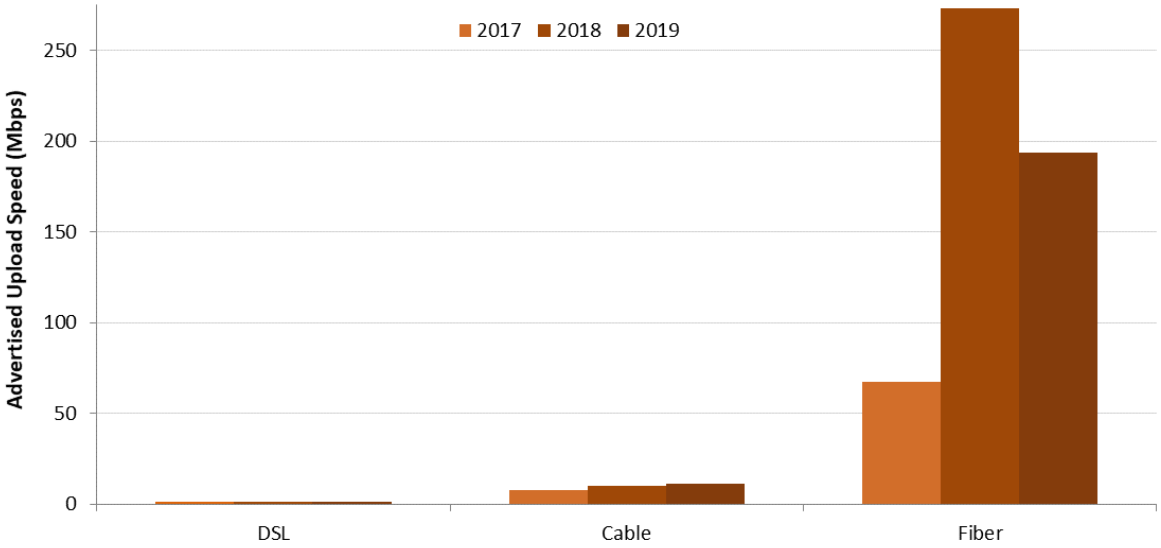

B. OBSERVED MEDIAN DOWNLOAD AND UPLOAD SPEEDS

Chart 4 (in Section 2.B) shows the ratio in September-October 2019 of the weighted median of both download and upload speeds of each ISP's subscribers to advertised speeds. Charts 12.1 and 12.2 below show the same ratios separately for download speed and for upload speed. The median download speeds of most ISPs' subscribers have been close to, or have exceeded, the advertised speeds. Exceptions to this were the following DSL providers: CenturyLink, Cincinnati Bell DSL, Frontier DSL and Windstream with respective ratios of 92%, 79%, 94% and 97%.

Chart 12.1: The ratio of median download speed to advertised download speed.

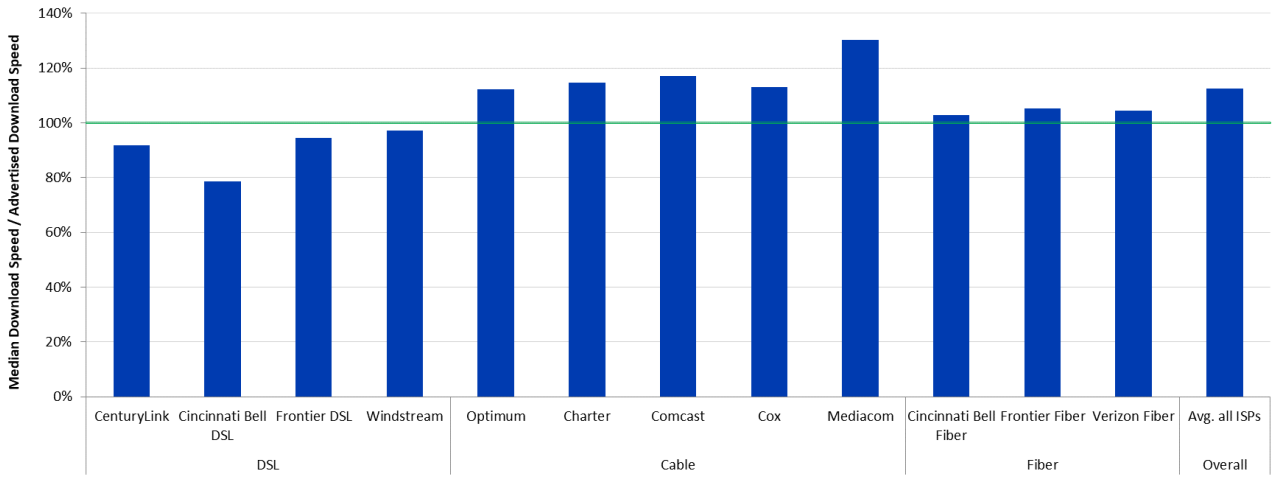

Chart 12.2 shows the median upload speed as a percentage of the advertised speed. As was the case with download speeds most ISPs met or exceeded the advertised rates except for a number of DSL providers: CenturyLink, Cincinnati Bell DSL, Frontier DSL and Windstream which had respective ratios of 87%, 77%, 90%, and 91%.

Chart 12.2: The ratio of median upload speed to advertised upload speed.

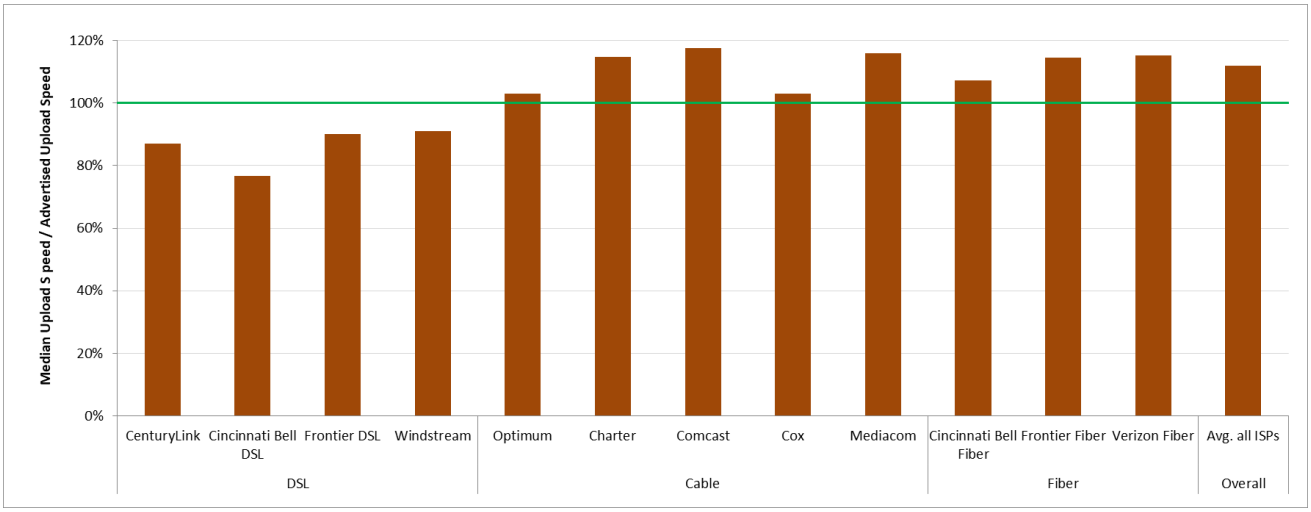

C. VARIATIONS IN SPEEDS

Median speeds experienced by consumers may vary based on location and time of day as the network architectures and traffic patterns may differ. Chart 5 in Section 2 above showed, for each ISP, the percentage of consumers (across the ISP's service territory) who experienced a median download speed over the peak usage period that was either greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed. Chart 13 below shows the corresponding percentage of consumers whose median upload speed fell in each of these ranges. ISPs using DSL technology had only between 0% to 36% of their subscribers getting greater than or equal to 95% of their advertised upload speeds during peak hours. In contrast, ISPs using cable or fiber technology had between 92% - 100% of their subscribers getting equal to or better than 95% of their advertised upload speeds.

Chart 13: The percentage of consumers whose median upload speed was (a) greater than 95%, (b) between 80% and 95%, or (c) less than 80% of the advertised upload speed.

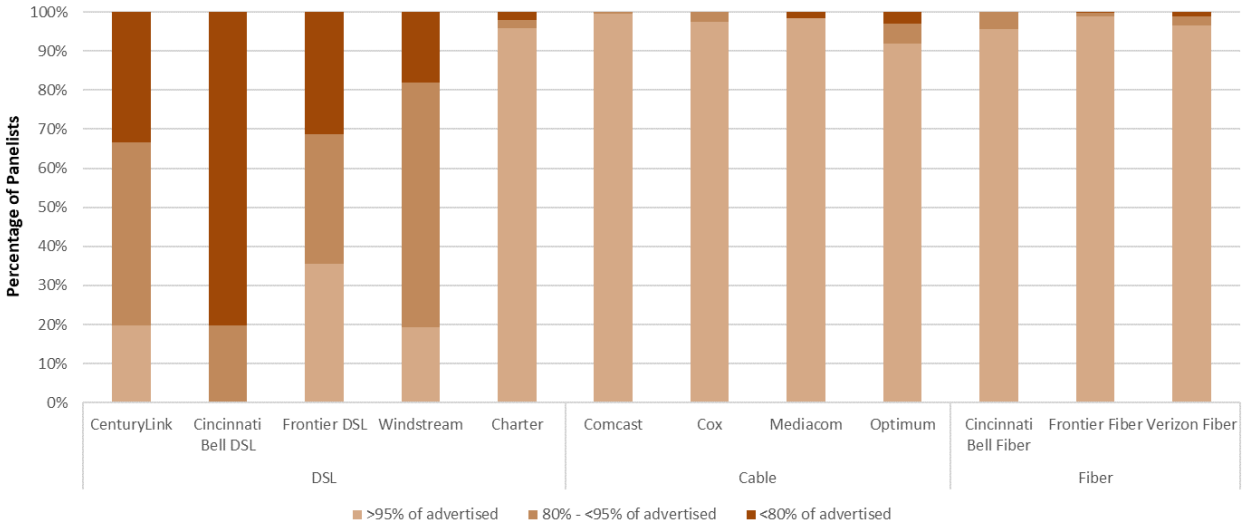

Though the median upload speeds experienced by most subscribers were close to or exceeded the advertised upload speeds there were some subscribers, for each ISP, whose median upload speed fell significantly short of the advertised upload speed. This issue was most prevalent for ISPs using DSL technology. On the other hand, ISPs using cable and fiber technology generally showed very good consistency based on this metric.

We can learn more about the variation in network performance by separately examining variations across geography and across time. We start by examining the variation across geography within each participating ISP's service territory. For each ISP, we first calculate the ratio of the median download speed (over the peak usage period) to the advertised download speed for each panelist subscribing to that ISP. We then examine the distribution of this ratio across the ISP's service territory.

Charts 14.1 and 14.2 show the complementary cumulative distribution of the ratio of median download speed (over the peak usage period) to advertised download speed for each participating ISP. For each ratio of actual to advertised download speed on the horizontal axis, the curves show the percentage of panelists subscribing to each ISP that experienced at least this ratio.[27] For example, the Cincinnati Bell fiber curve in Chart 14.1 shows that 90% of its subscribers experienced a median download speed exceeding 76% of the advertised download speed, while 70% experienced a median download speed exceeding 92% of the advertised download speed, and 50% experienced a median download speed exceeding 107% of the advertised download speed.

[27] In Reports prior to the 2015 MBA Report, for each ratio of actual to advertised download speed on the horizontal axis, the cumulative distribution function curves showed the percentage of measurements, rather than panelists subscribing to each ISP, that experienced at least this ratio. The methodology used since then, *i.e.*, using panelists subscribing to each ISP, more accurately illustrates ISP performance from a consumer's point of view.

Chart 14.1: Complementary cumulative distribution of the ratio of median download speed to advertised download speed.

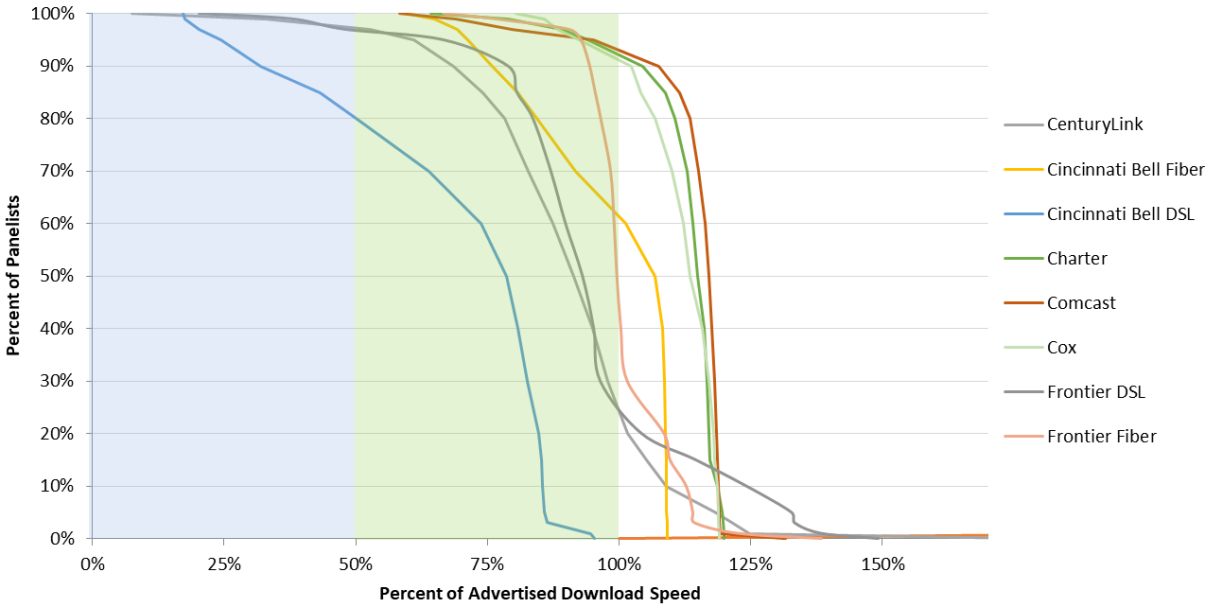

Chart 14.2: Complementary cumulative distribution of the ratio of median download speed to advertised download speed (continued).

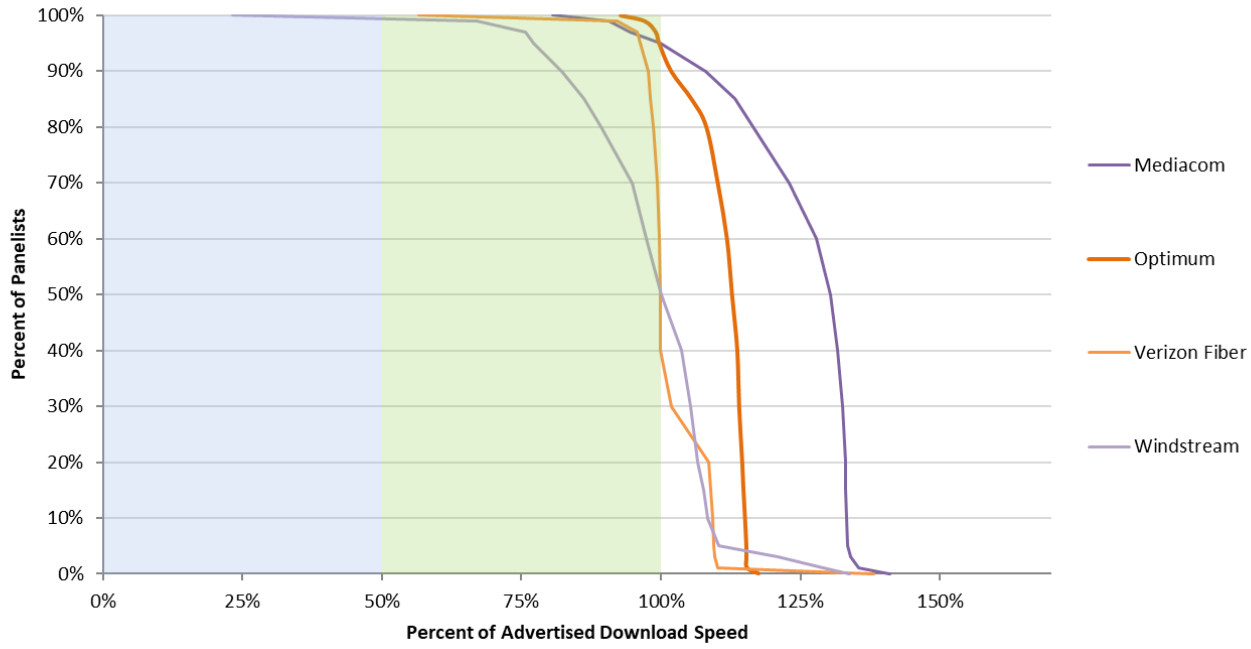

The curves for cable-based broadband and fiber-based broadband are steeper than those for DSL-based broadband. This can be seen more clearly in Chart 14.3, which plots aggregate curves for each technology. Approximately 90% of subscribers to cable and 50% of subscribers to fiber-based technologies experience

median download speeds exceeding the advertised download speed. In contrast, less than 30% of subscribers to DSL-based services experience median download speeds exceeding the advertised download speed.[28]

Chart 14.3: Complementary cumulative distribution of the ratio of median download speed to advertised download speed, by technology.

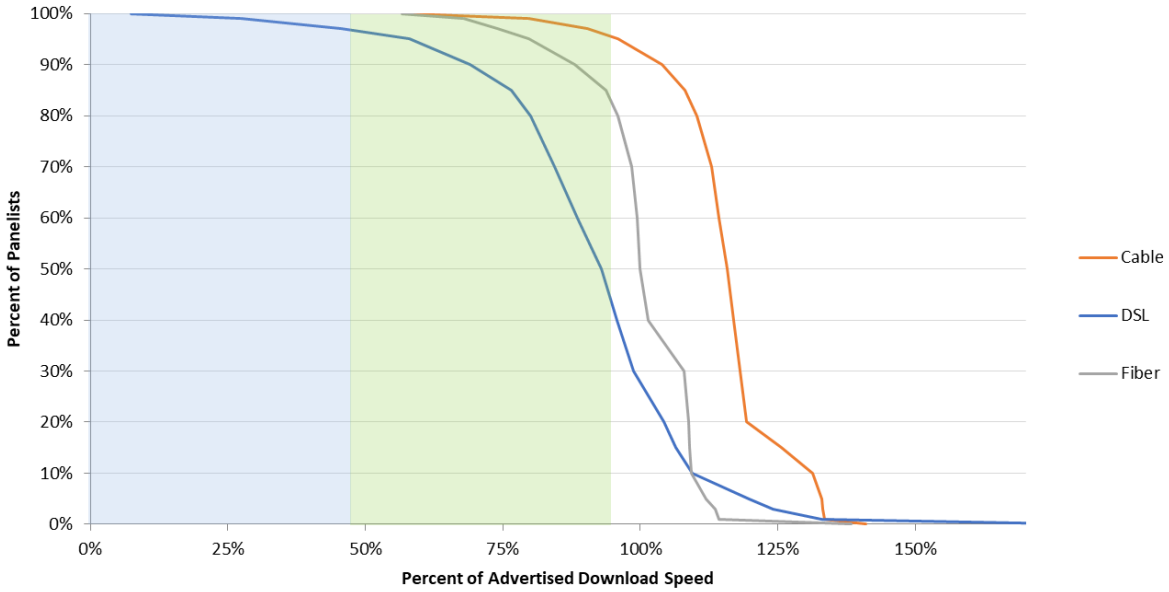

Charts 14.4 to 14.6 show the complementary cumulative distribution of the ratio of median upload speed (over the peak usage period) to advertised upload speed for each participating ISP (Charts 14.4 and 14.5) and by access technology (Chart 14.6).

[28] The speed achievable by DSL depends on the distance between the subscriber and the central office. Thus, the complementary cumulative distribution function will fall slowly unless the broadband ISP adjusts its advertised rate based on the subscriber's location. (Chart 16 illustrates that the performance during non-busy hours is similar to the busy hour, making congestion less likely as an explanation.)

Chart 14.4: Complementary cumulative distribution of the ratio of median upload speed to advertised upload speed.

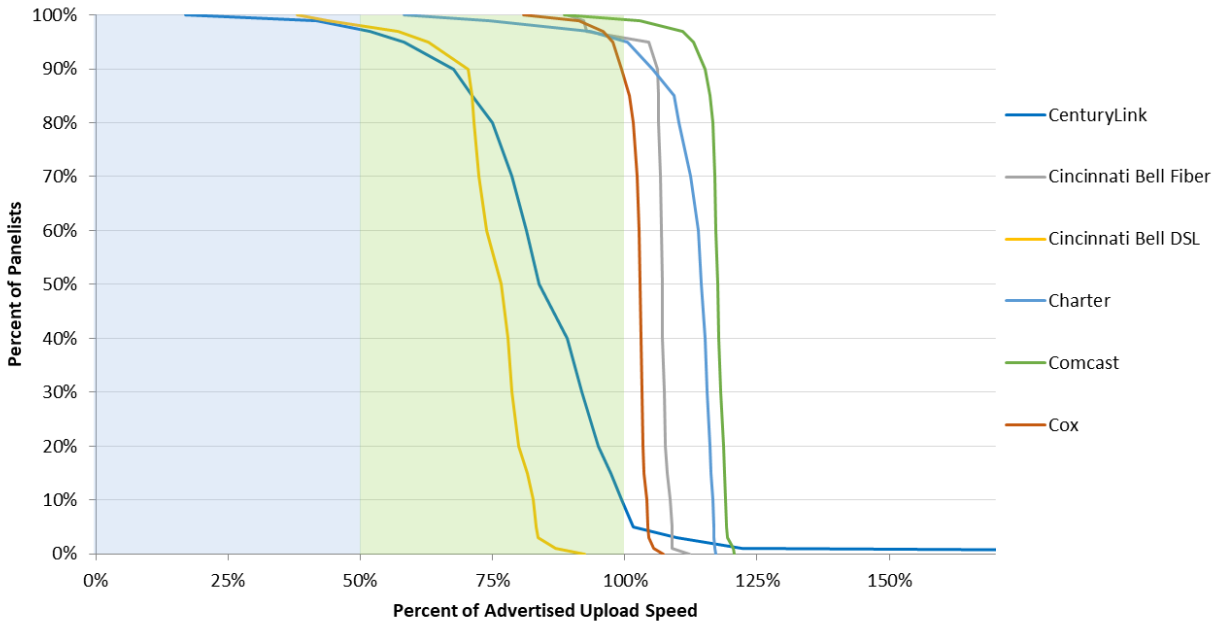

Chart 14.5: Complementary cumulative distribution of the ratio of median upload speed to advertised upload speed (continued).

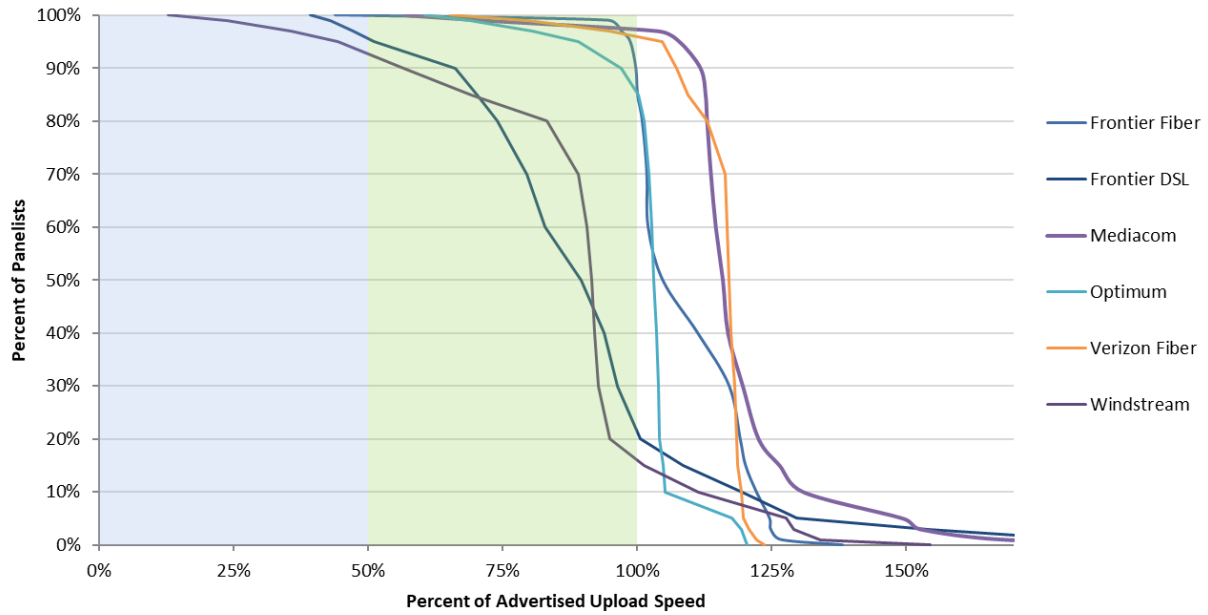

Chart 14.6: Complementary cumulative distribution of the ratio of median upload speed to advertised upload speed, by technology.

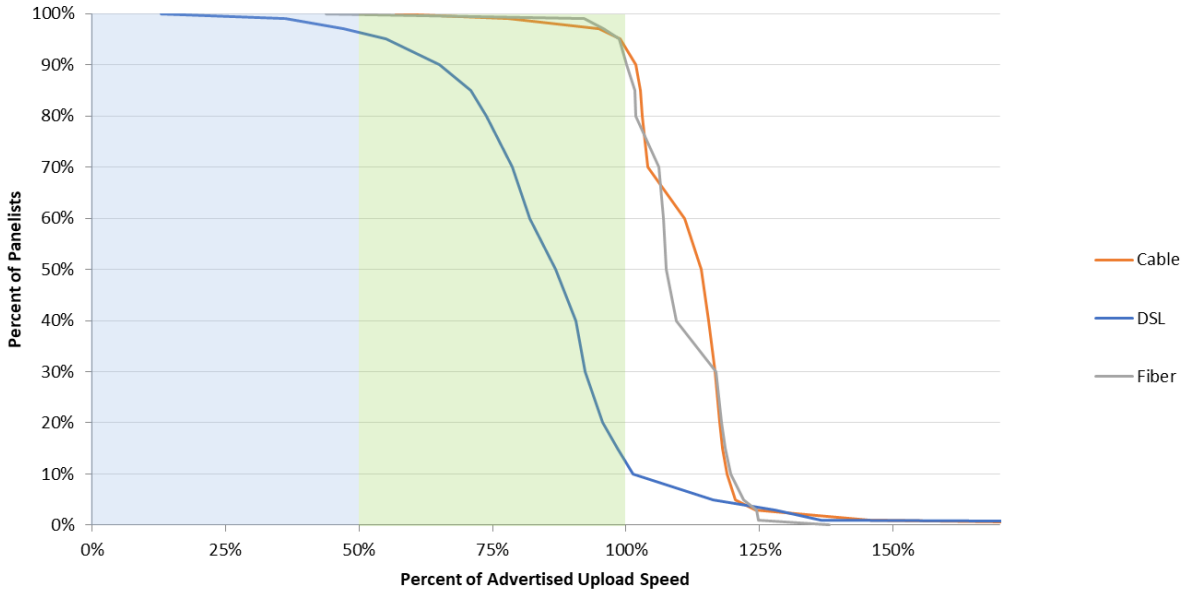

All actual speeds discussed above were measured during peak usage periods. In contrast, Charts 15.1 and 15.2 below compare the ratio of actual download and upload speeds to advertised download and upload speeds during peak and off-peak times. Charts 15.1 and 15.2 show that most ISP subscribers experience only a slight degradation from off-peak to peak hour performance.

Chart 15.1: The ratio of weighted median download speed to advertised download speed, peak hours versus off-peak hours.

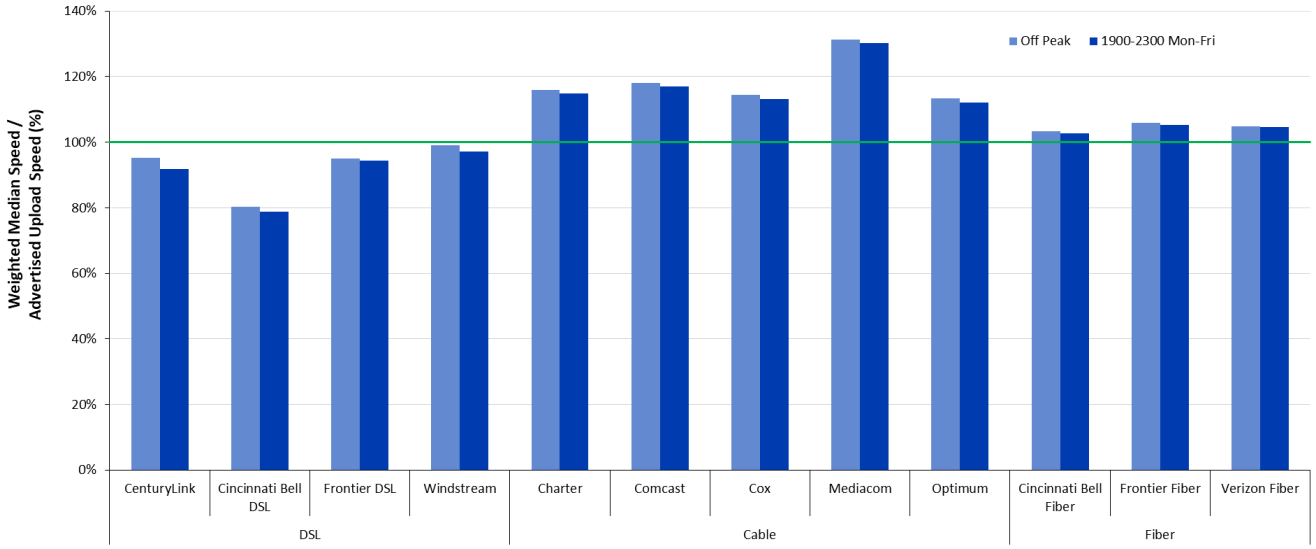

Chart 15.2: The ratio of weighted median upload speed to advertised upload speed, peak versus off-peak.

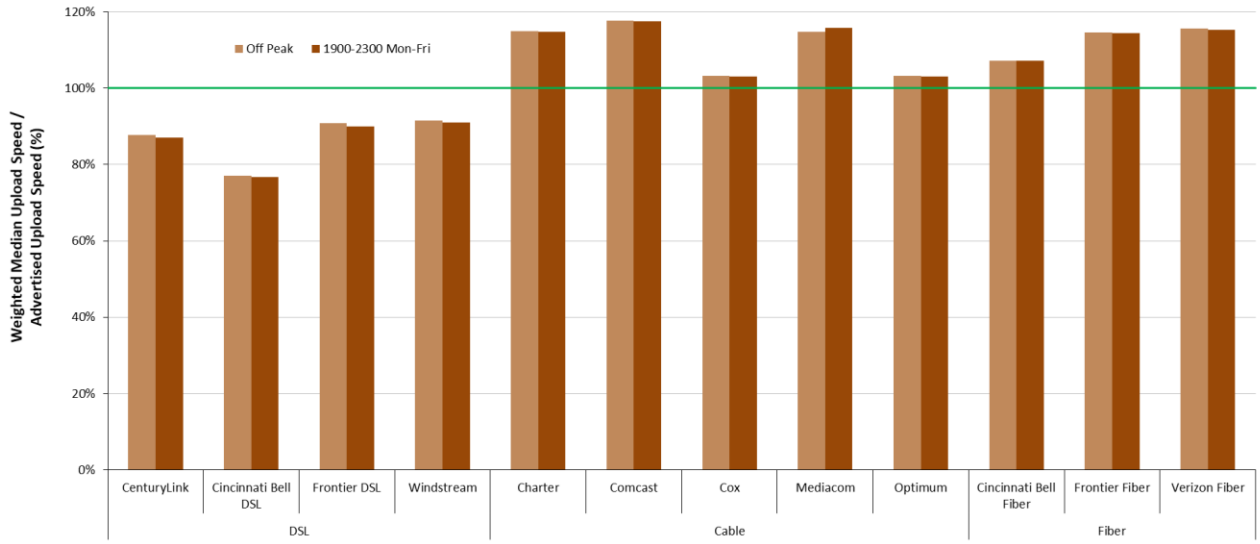

Chart 16 below shows the actual download speed to advertised speed ratio in each two-hour time block during weekdays for each ISP. The ratio is lowest during the busiest four-hour time block (7:00 p.m. to 11:00 p.m.).

Chart 16: The ratio of median download speed to advertised download speed, Monday-to-Friday, two-hour time blocks, terrestrial ISPs.

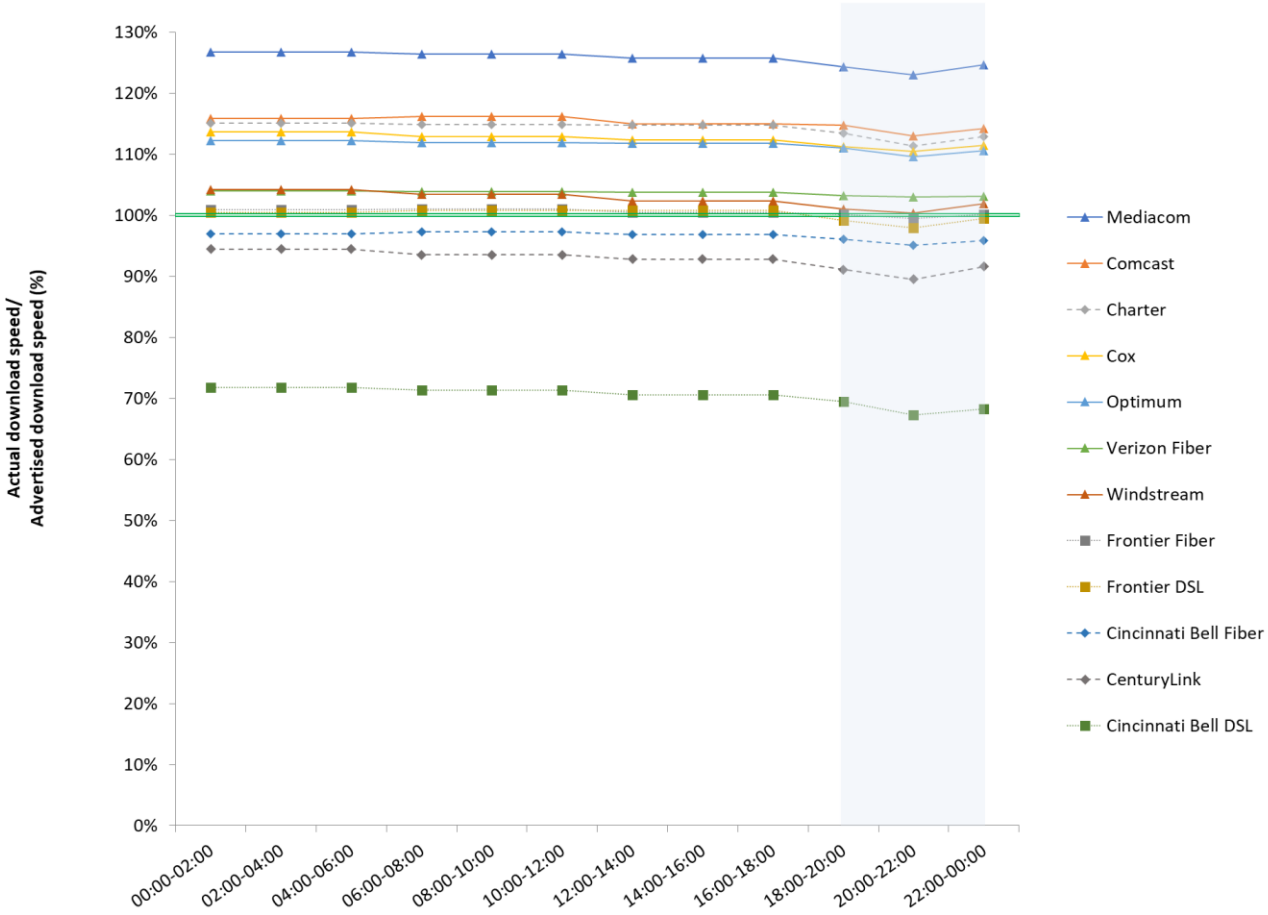

For each ISP, Chart 6 (in Section 2.C) showed the ratio of the 80/80 consistent median download speed to advertised download speed, and for comparison, Chart 4 showed the ratio of median download speed to advertised download speed.

Chart 17.1 illustrates information concerning 80/80 consistent upload speeds. While all the upload 80/80 speeds were slightly lower than the median speed the differences were more marked in DSL. Charts 6 and 17.1 make it clear that cable and fiber technologies behaved more consistently than DSL technology both for download as well as upload speeds.

Chart 17.1: The ratio of 80/80 consistent upload speed to advertised upload speed.

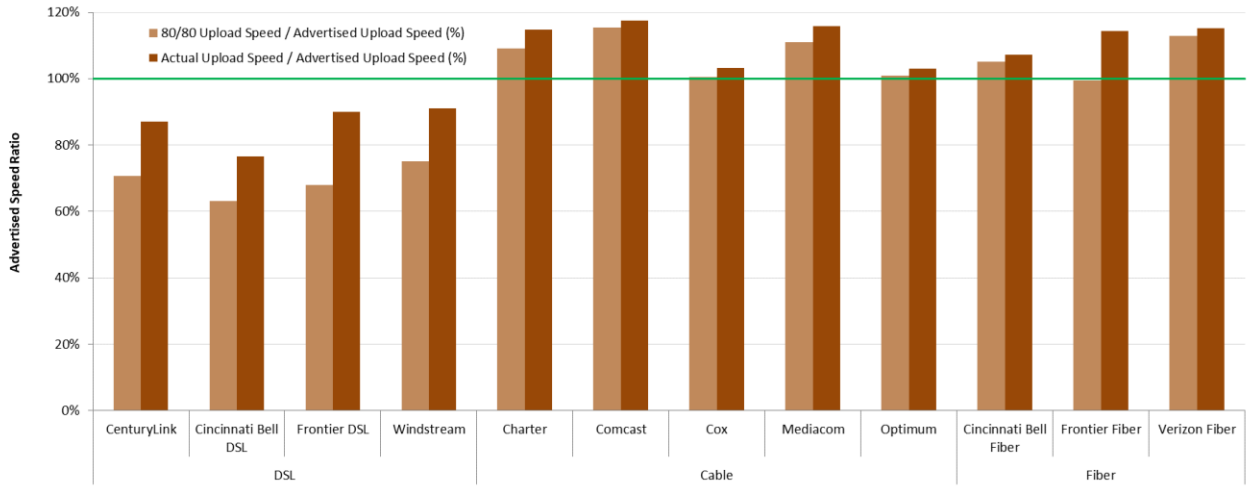

Charts 17.2 and 17.3 below illustrate similar consistency metrics for 70/70 consistent download and upload speeds, *i.e.*, the minimum download or upload speed (as a percentage of the advertised download or upload speed) experienced by at least 70% of panelists during at least 70% of the peak usage period. The ratios for 70/70 consistent speeds as a percentage of the advertised speed are higher than the corresponding ratios for 80/80 consistent speeds. In fact, for many ISPs, the 70/70 consistent download or upload speed is close to the median download or upload speed. Once again, ISPs using DSL technology showed a considerably smaller value for the 70/70 download and upload speeds as compared to the download and upload median speeds, respectively.

Chart 17.2: The ratio of 70/70 consistent download speed to advertised download speed.

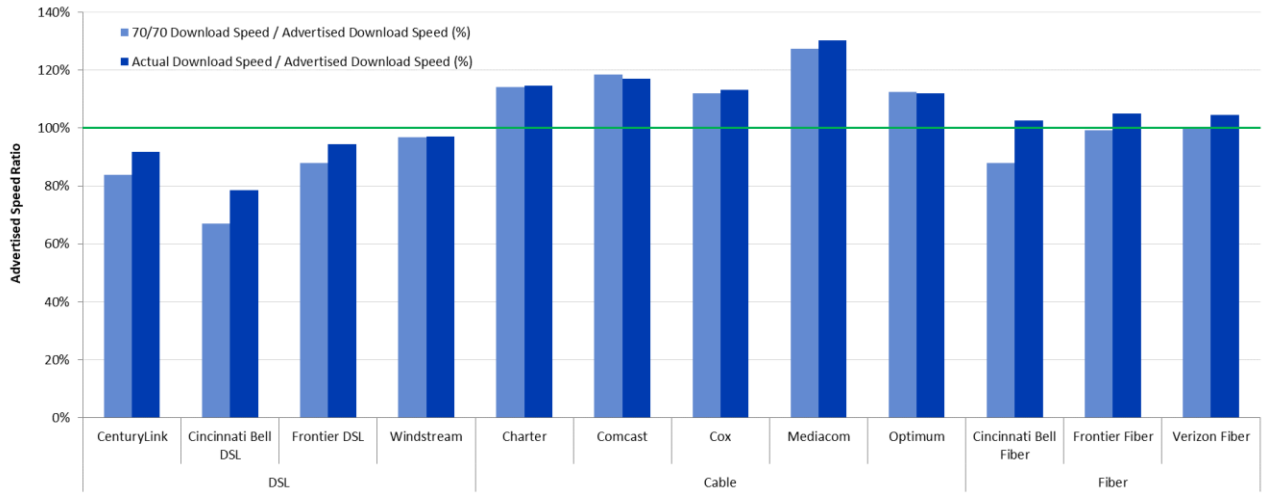

Chart 17.3: The ratio of 70/70 consistent upload speed to advertised upload speed.

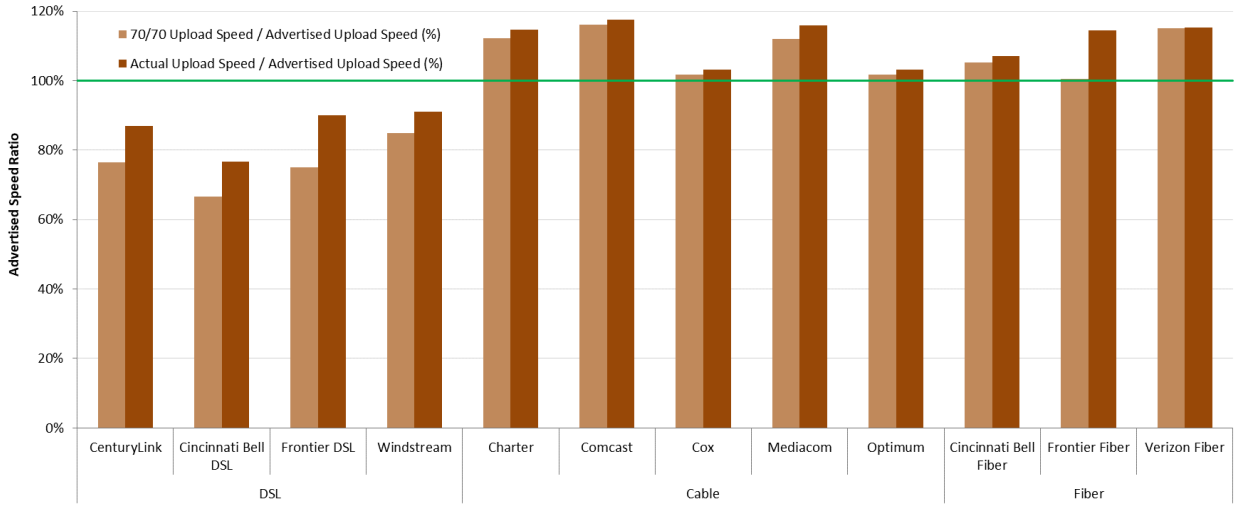

D. LATENCY

Chart 18 below shows the weighted median latencies, by technology and by advertised download speed for terrestrial technologies. For all terrestrial technologies, latency varied little with advertised download speed. DSL service typically had higher latencies, and lower latency was better correlated with advertised download speed, than with either cable or fiber. Cable latencies ranged between 16ms to 28ms, fiber latencies between 5ms to 11ms, and DSL between 21ms to 61ms.

Chart 18: Latency for Terrestrial ISPs, by technology, and by advertised download speed.

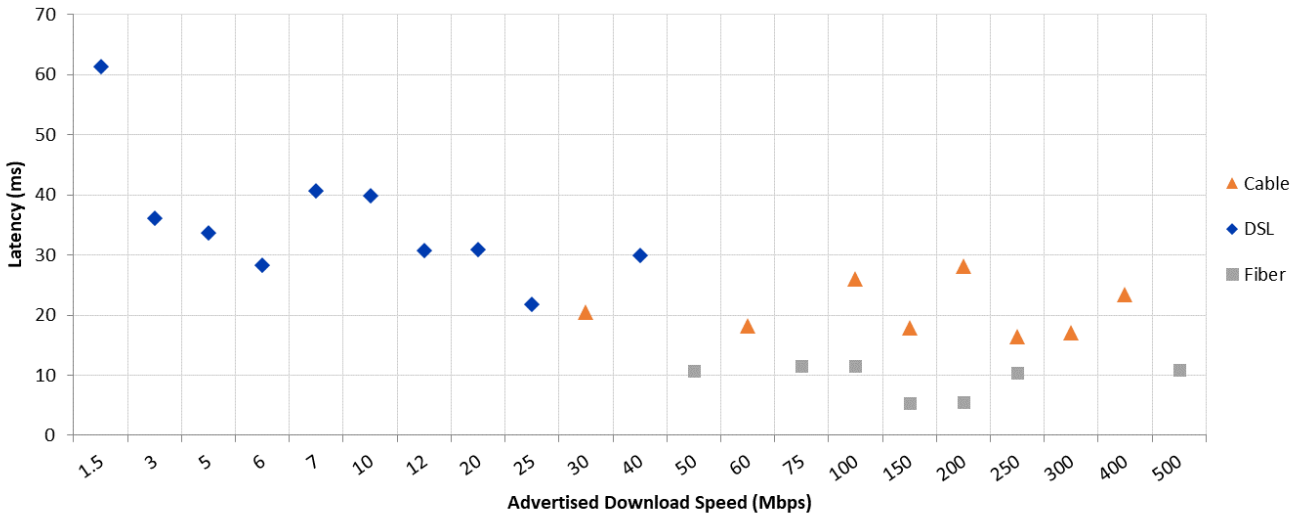

5. ADDITIONAL TEST RESULTS

A. ACTUAL SPEED, BY SERVICE TIER

As shown in Charts 19.1-19.8, peak usage period performance varied by service tier among participating ISPs during the September-October 2019 period. On average, during peak periods, the ratio of median download speed to advertised download speed for all ISPs was 79% or better, and 90% or better for most ISPs. However, the ratio of median download speed to advertised download speed varies among service tiers. Out of the 37 speed tiers that were measured a large majority (32) showed that they at least achieved 90% of the advertised speed and 23 of the 37 tiers either met or exceeded the advertised speed.

Chart 19.1: The ratio of median download speed to advertised download speed, by ISP (1-5 Mbps).

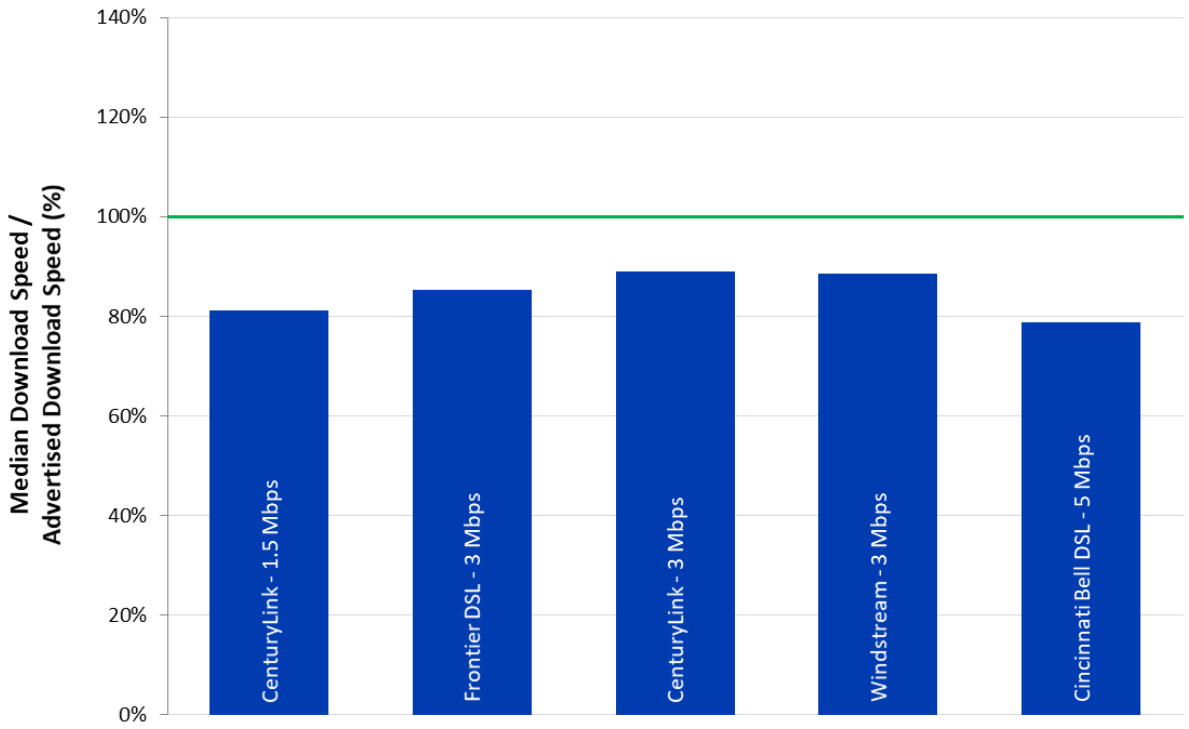

Chart 19.2: The ratio of median download speed to advertised download speed, by ISP (6-10 Mbps).

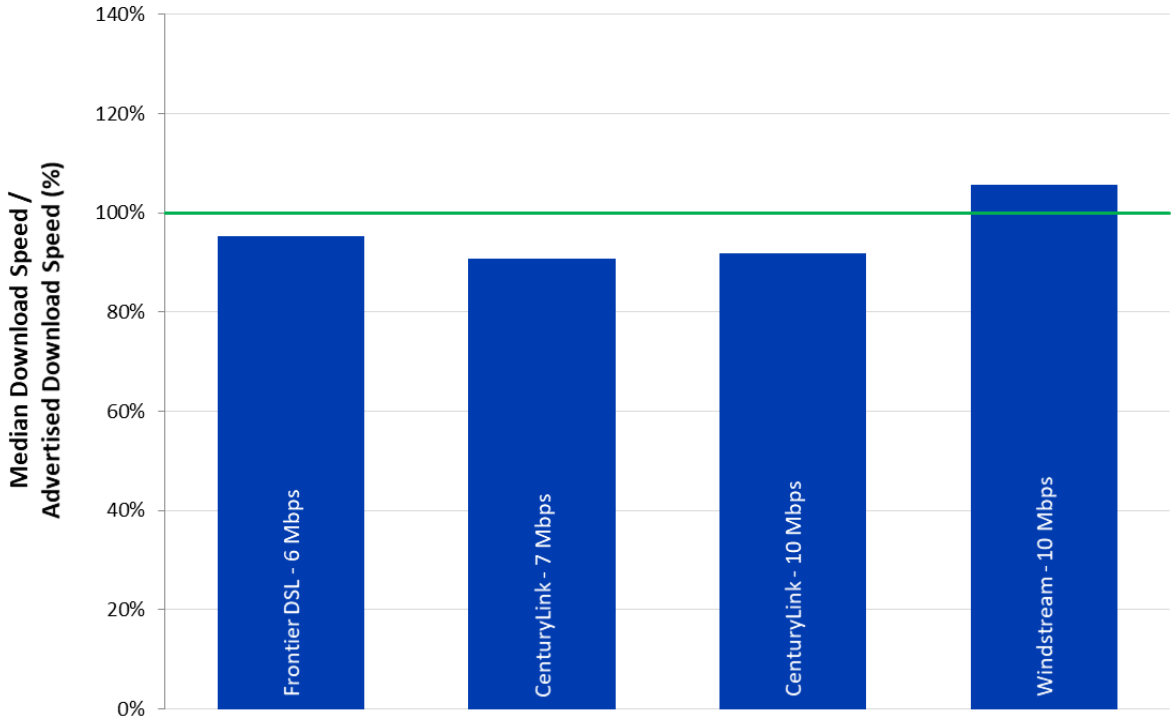

Chart 19.3: The ratio of median download speed to advertised download speed, by ISP (12-25 Mbps).

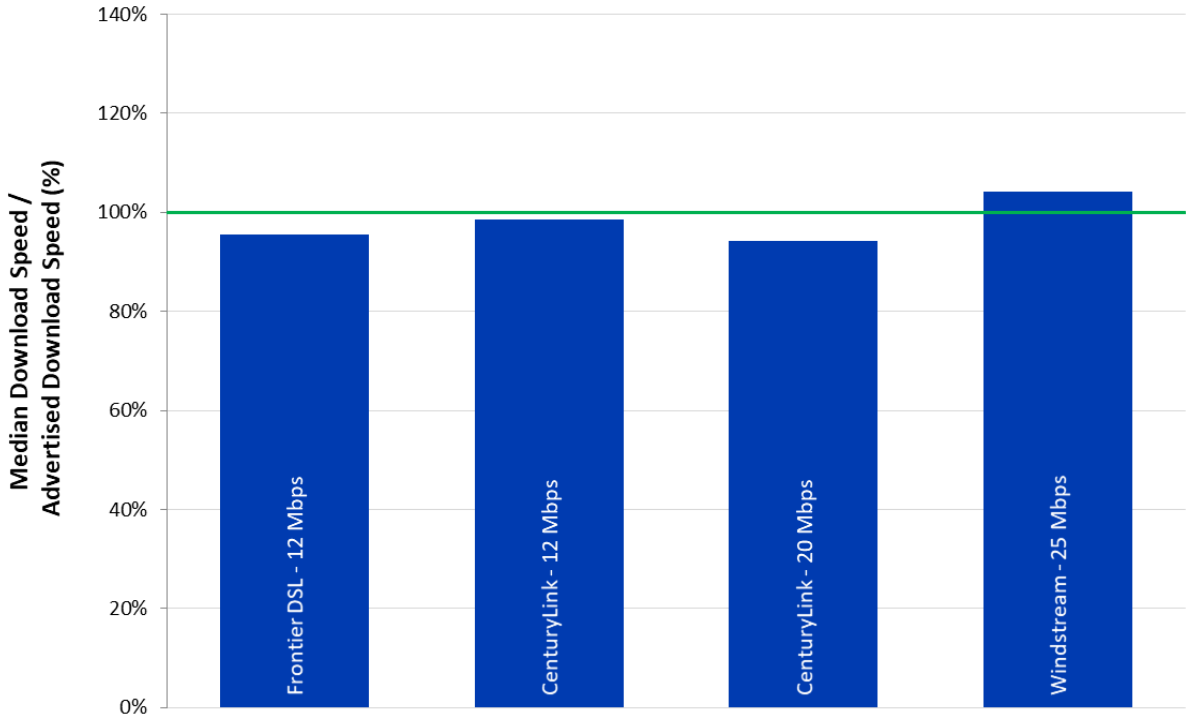

Chart 19.4: The ratio of median download speed to advertised download speed, by ISP (30-60 Mbps).

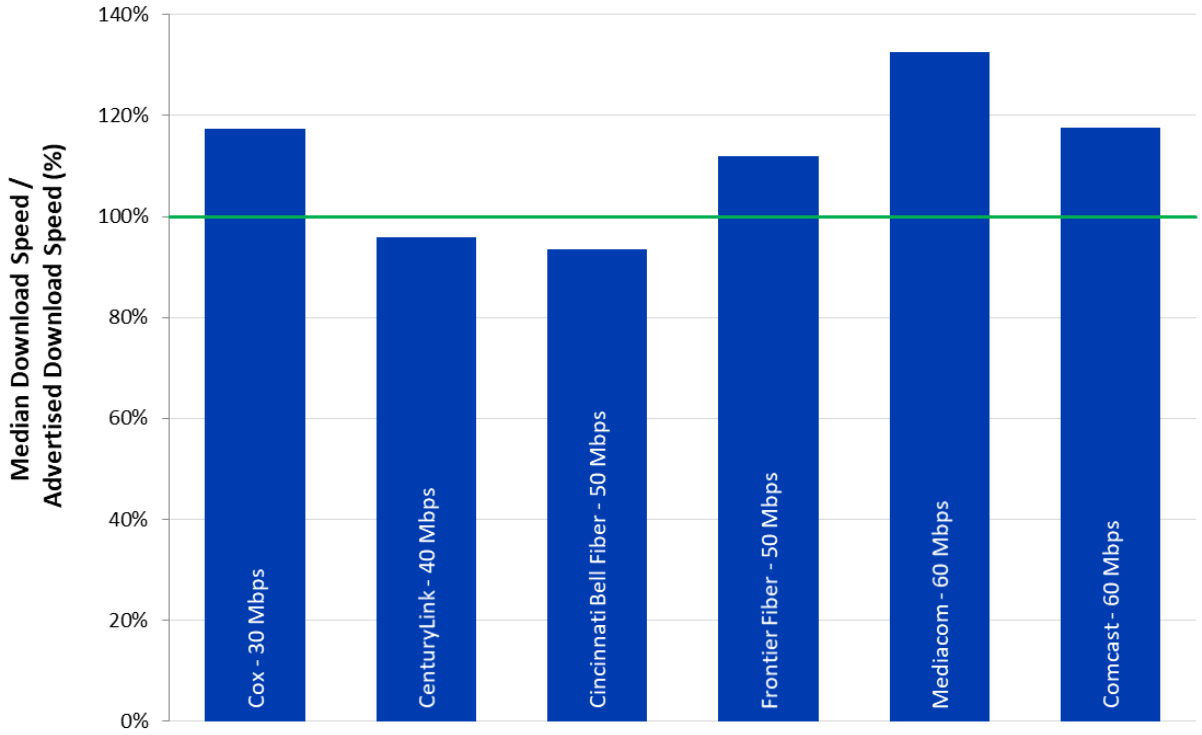

Chart 19.5: The ratio of median download speed to advertised download speed, by ISP (75-100Mbps).

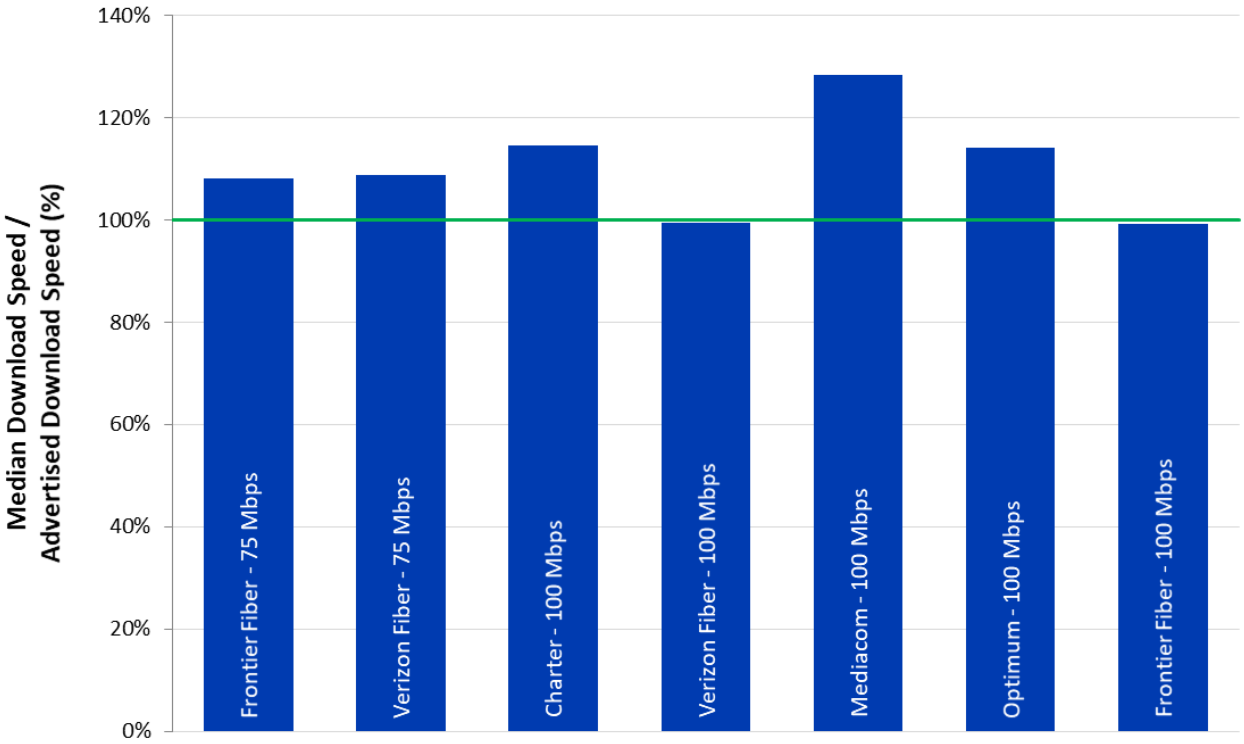

Chart 19.6: The ratio of median download speed to advertised download speed, by ISP (150-200 Mbps).

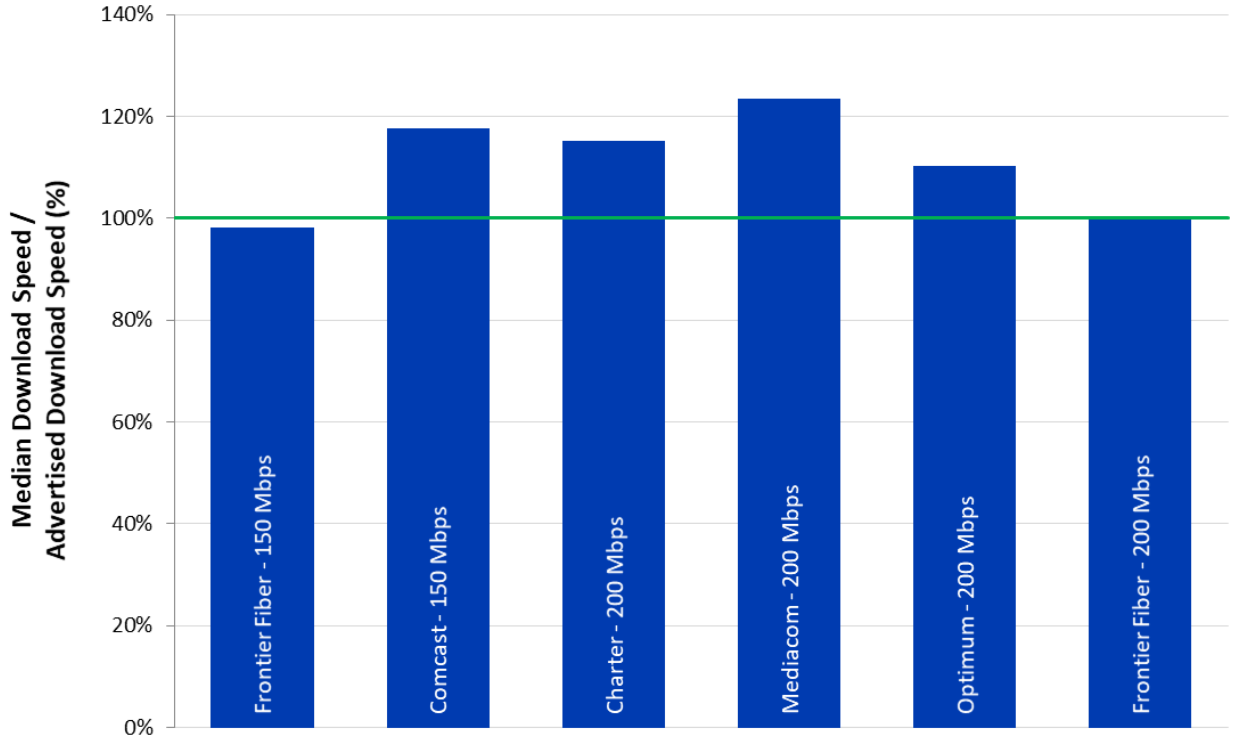

Chart 19.7: The ratio of median download speed to advertised download speed, by ISP (250-500 Mbps).

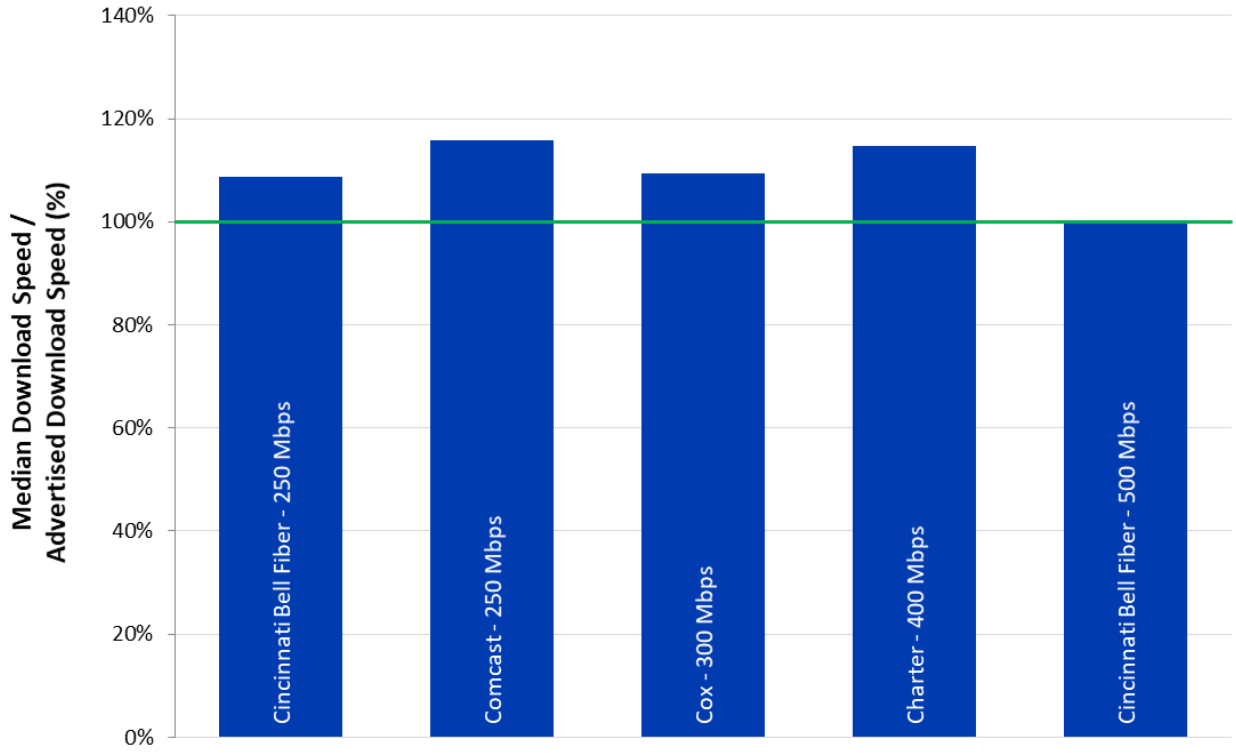

Charts 20.1 – 20.6 depict the ratio of median upload speeds to advertised upload speeds for each ISP by service tier. Out of the 30 upload speed tiers that were measured a large majority (25) showed that they at least achieved 90% of the advertised upload speed, and 21 of the 30 tiers either met or exceeded the advertised upload speed.

Chart 20.1: The ratio of median upload speed to advertised upload speed, by ISP (0.768 - 1 Mbps).

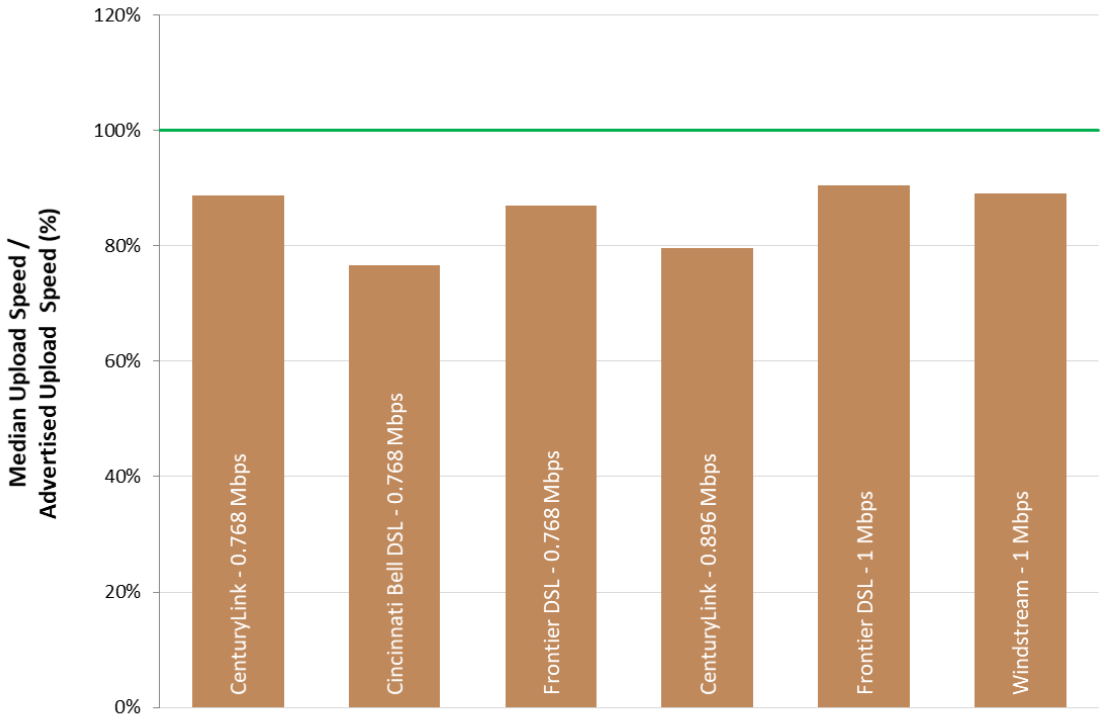

Chart 20.2: The ratio of median upload speed to advertised upload speed, by ISP (1.5-5 Mbps).

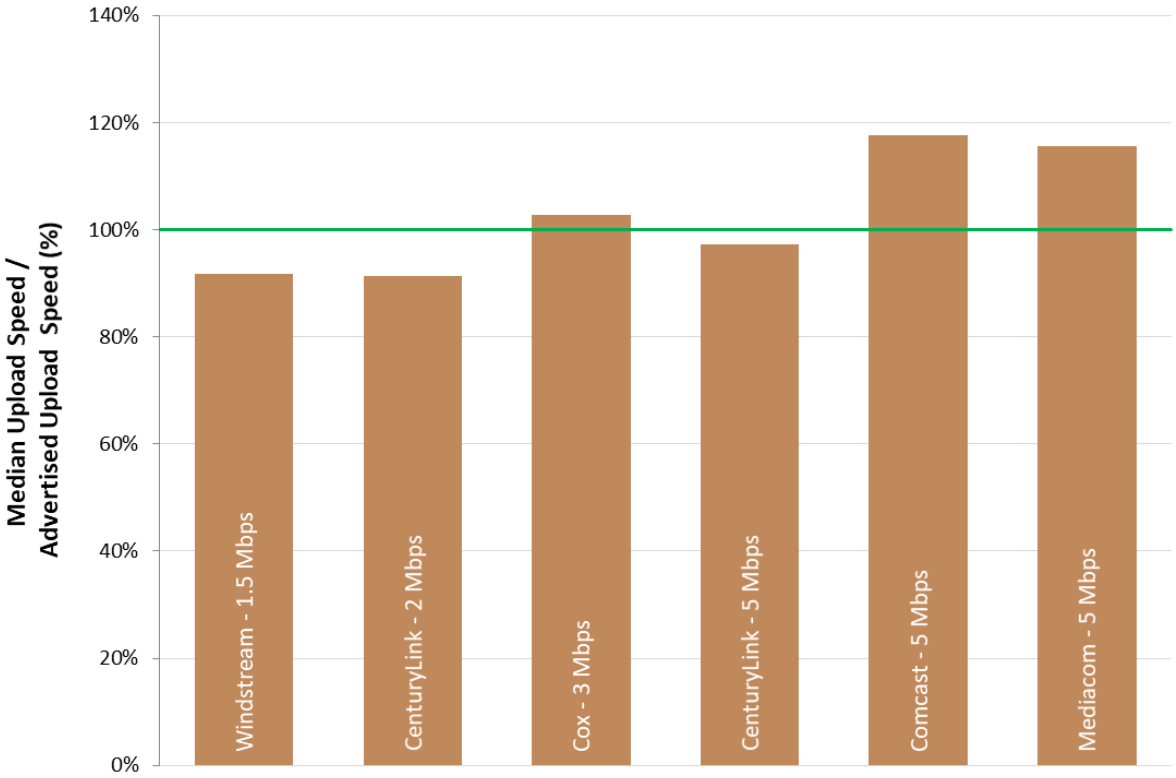

Chart 20.3: The ratio of median upload speed to advertised upload speed, by ISP (10 -20 Mbps).

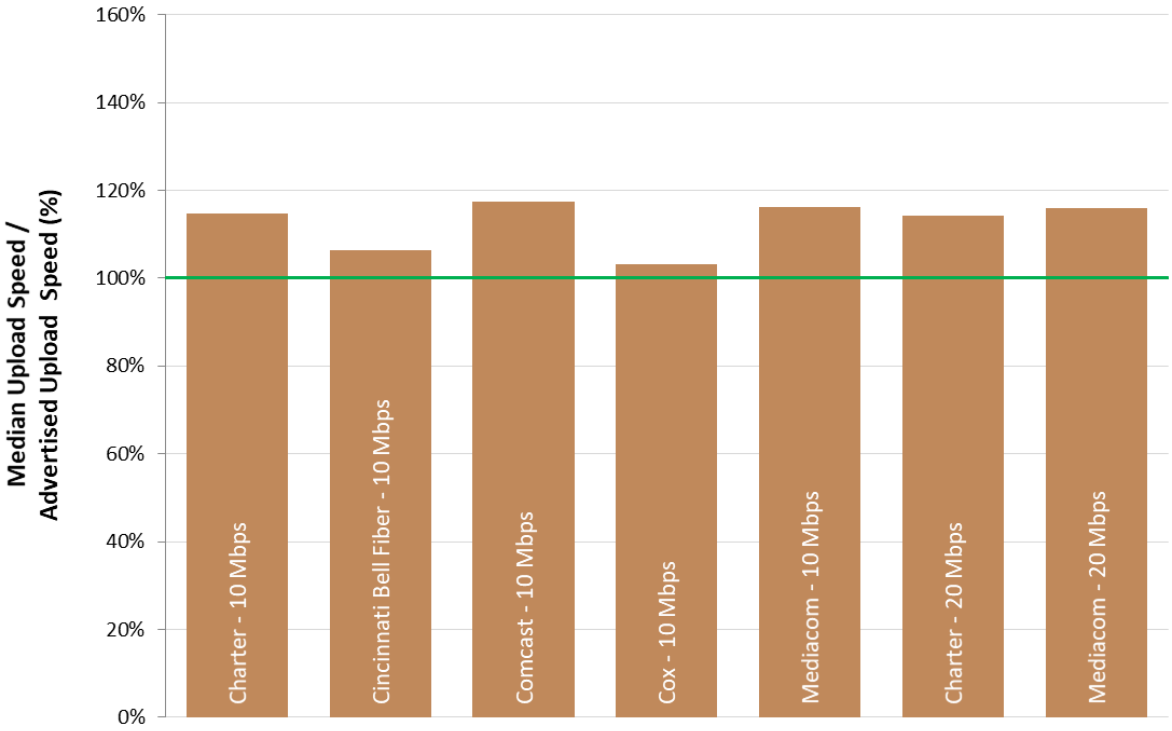

Chart 20.4: The ratio of median upload speed to advertised upload speed, by ISP (30-75 Mbps).

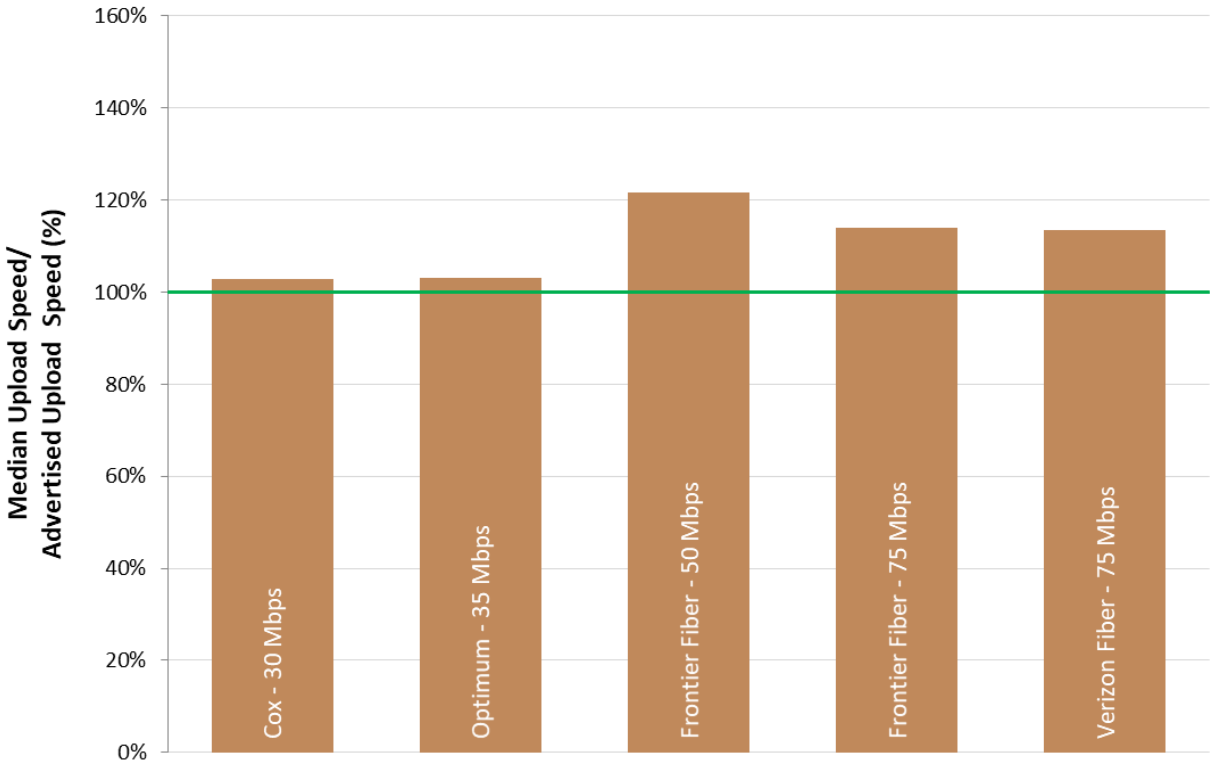

Chart 20.5: The ratio of median upload speed to advertised upload speed, by ISP (100–200 Mbps).

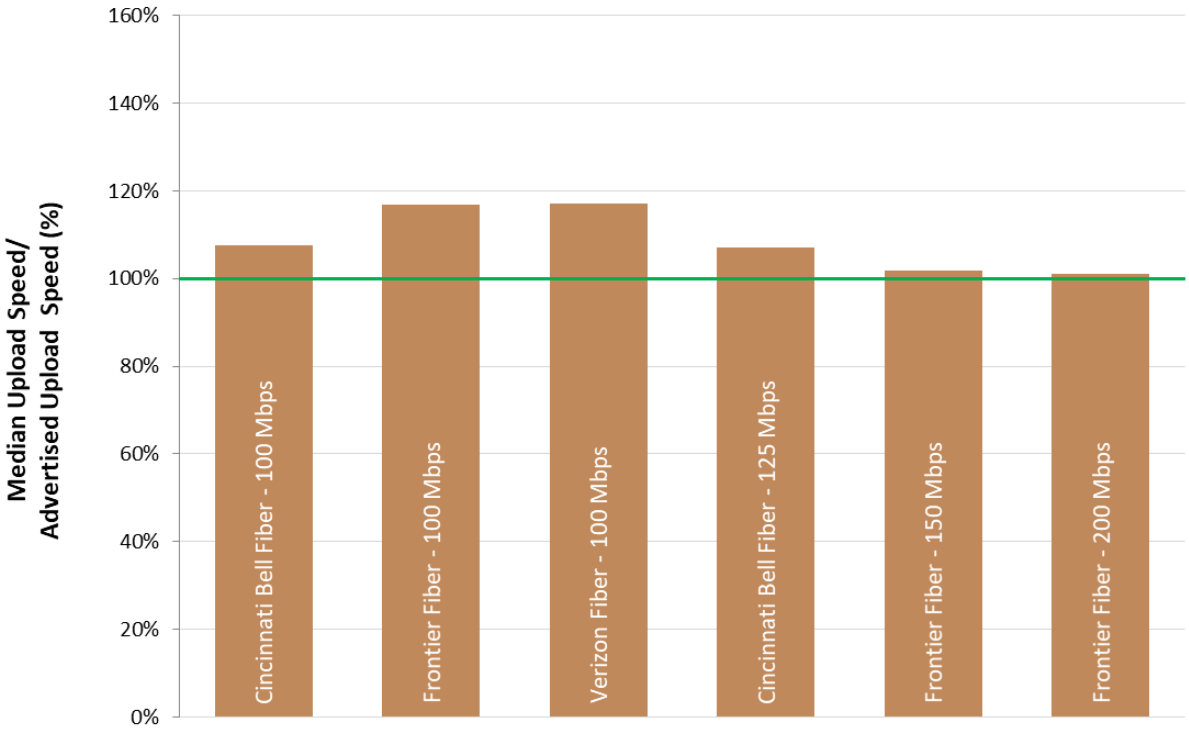

Table 2 lists the advertised download service tiers included in this study. For each tier, an ISP's advertised download speed is compared with the median of the measured download speed results. As we noted in the past reports, the download speeds listed here are based on national averages and may not represent the performance experienced by any particular consumer at any given time or place.

Table 2: Peak period median download speed, sorted by actual download speed

ISP	Advertised Download Speed (Mbps)	Download Median Speed (Mbps)	Actual Speed / Advertised Speed (%)
CenturyLink	1.5	1.22	81.26%
Frontier DSL	3	2.56	85.23%
CenturyLink	3	2.67	88.97%
Windstream	3	2.66	88.54%
Cincinnati Bell DSL	5	3.94	78.71%
Frontier DSL	6	5.72	95.36%
CenturyLink	7	6.35	90.71%

CenturyLink	10	9.17	91.75%
Windstream	10	10.57	105.71%
Frontier DSL	12	11.47	95.54%
CenturyLink	12	11.83	98.57%
CenturyLink	20	18.86	94.31%
Windstream	25	26.02	104.07%
Cox	30	35.22	117.39%
CenturyLink	40	38.37	95.92%
Cincinnati Bell Fiber	50	46.76	93.52%
Frontier Fiber	50	55.98	111.97%
Mediacom	60	79.53	132.55%
Comcast	60	70.61	117.69%
Frontier Fiber	75	81.03	108.04%
Verizon Fiber	75	81.65	108.87%
Charter	100	114.58	114.58%
Verizon Fiber	100	99.53	99.53%
Mediacom	100	128.36	128.36%
Optimum	100	114.21	114.21%
Frontier Fiber	100	99.10	99.10%
Frontier Fiber	150	147.23	98.15%
Comcast	150	176.36	117.57%
Charter	200	230.37	115.18%
Mediacom	200	246.89	123.45%
Optimum	200	220.63	110.32%
Frontier Fiber	200	200.07	100.04%
Cincinnati Bell Fiber	250	271.85	108.74%
Comcast	250	289.35	115.74%
Cox	300	327.92	109.31%
Charter	400	458.87	114.72%
Cincinnati Bell Fiber	500	499.15	99.83%

E. VARIATIONS IN SPEED

In Section 3.C above, we present speed consistency metrics for each ISP based on test results averaged across all service tiers. In this section, we provide detailed speed consistency results for each ISP's individual service tiers. Consistency of speed is important for services such as video streaming. A significant reduction in speed for more than a few seconds can force a reduction in video resolution or an intermittent loss of service.

Charts 21.1 – 21.3 below show the percentage of consumers that achieved greater than 95%, between 85% and 95%, or less than 80% of the advertised download speed for each ISP speed tier. Consistent with past performance, ISPs using DSL technology frequently fail to deliver advertised service rates. ISPs quote a single 'up-to' speed, but the actual speed of DSL depends on the distance between the subscriber and the serving central office.

Cable companies, in general, showed a high consistency of speed.

Chart 21.1: The percentage of consumers whose median download speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed, by service tier (DSL).

S

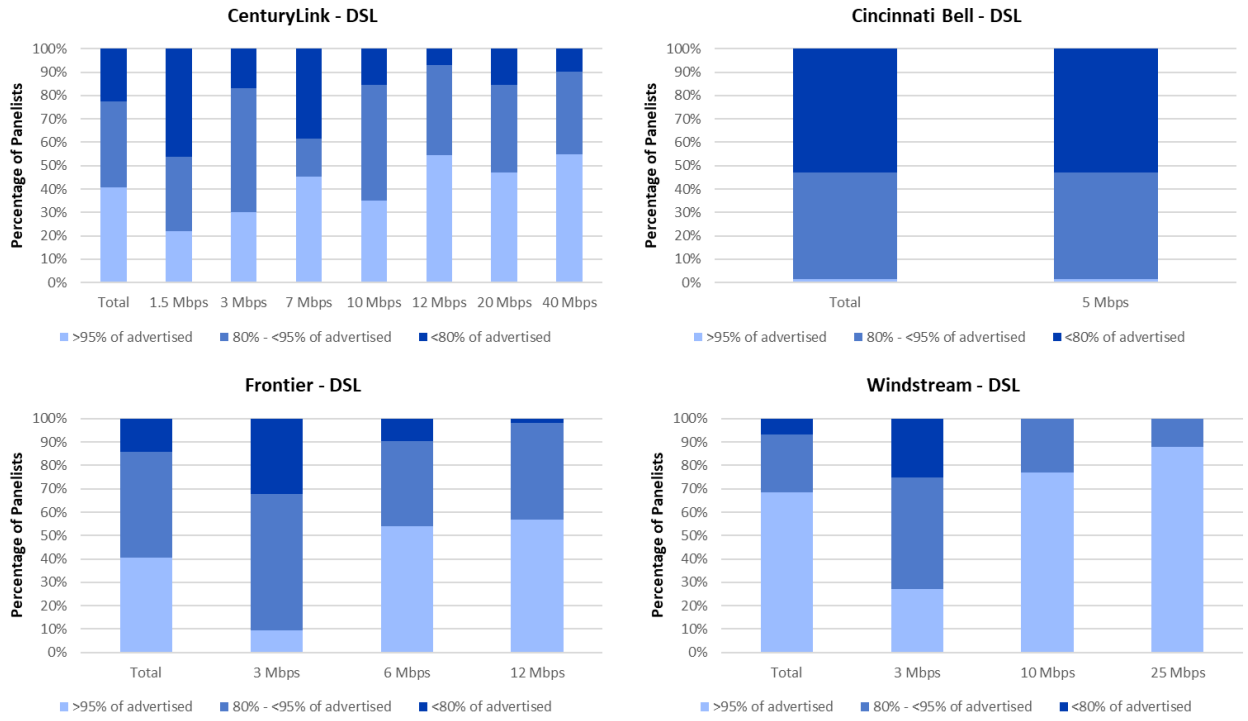

Chart 21.2: The percentage of consumers whose median download speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed (cable).

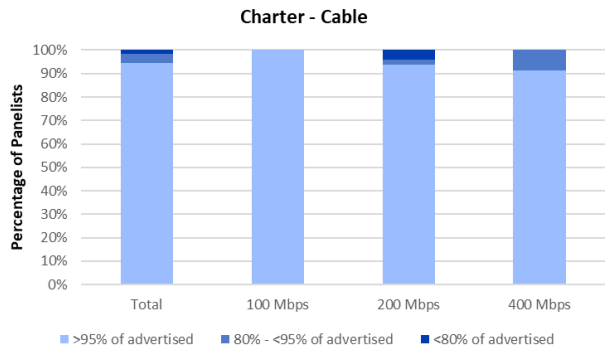

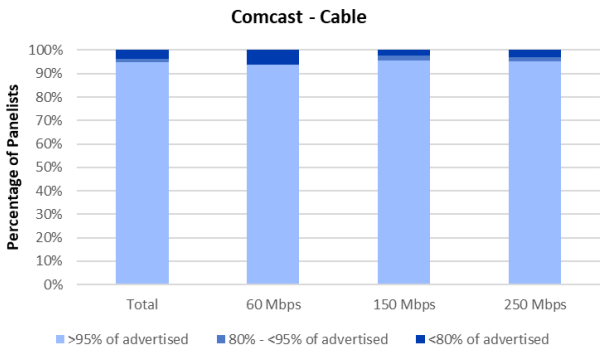

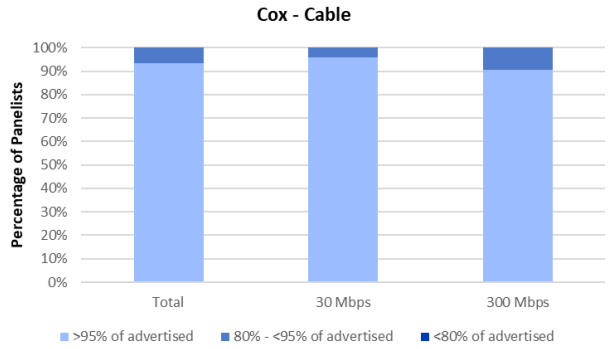

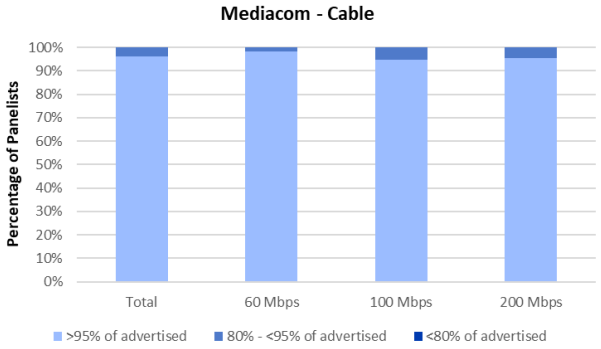

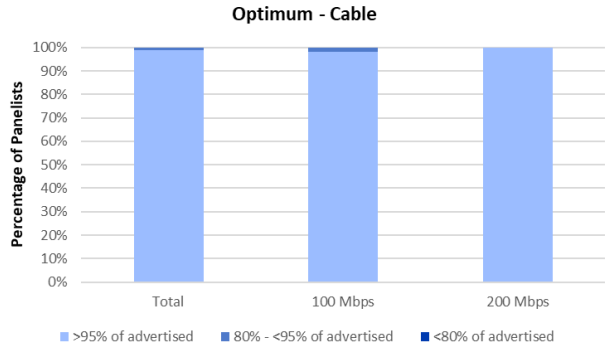

Chart 21.3: The percentage of consumers whose median download speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised download speed (fiber).

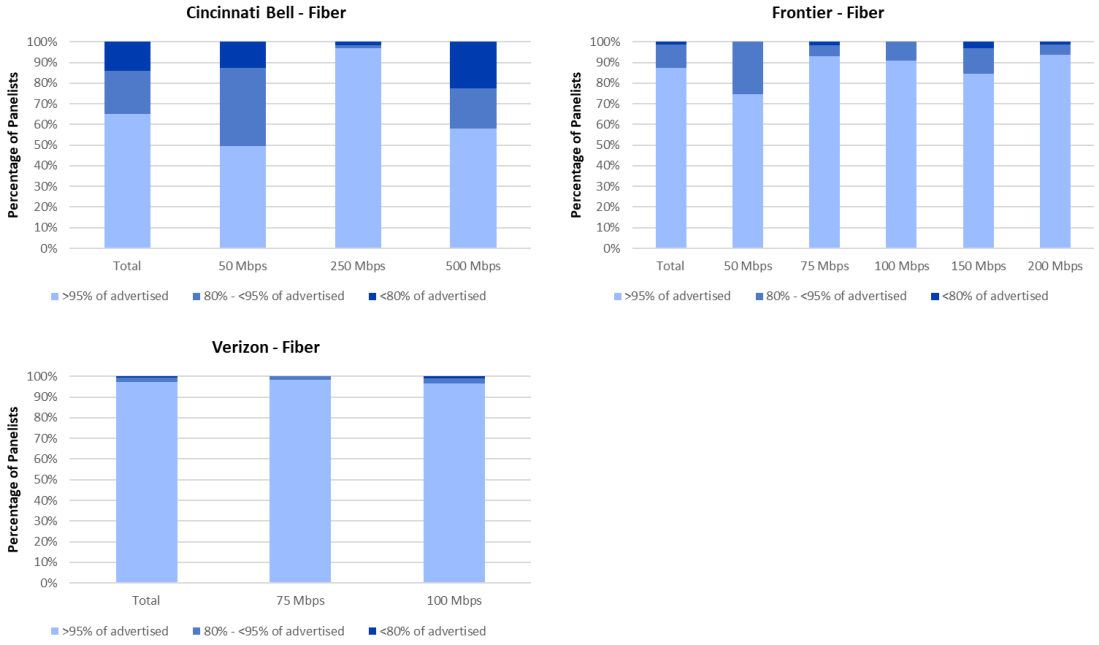

Similarly, Charts 22.1 to 22.3 show the percentage of consumers that achieved greater than 95%, between 85% and 95%, or less than 80% of the advertised upload speed for each ISP speed tier.

Chart 22.1: The percentage of consumers whose median upload speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised upload speed (DSL).

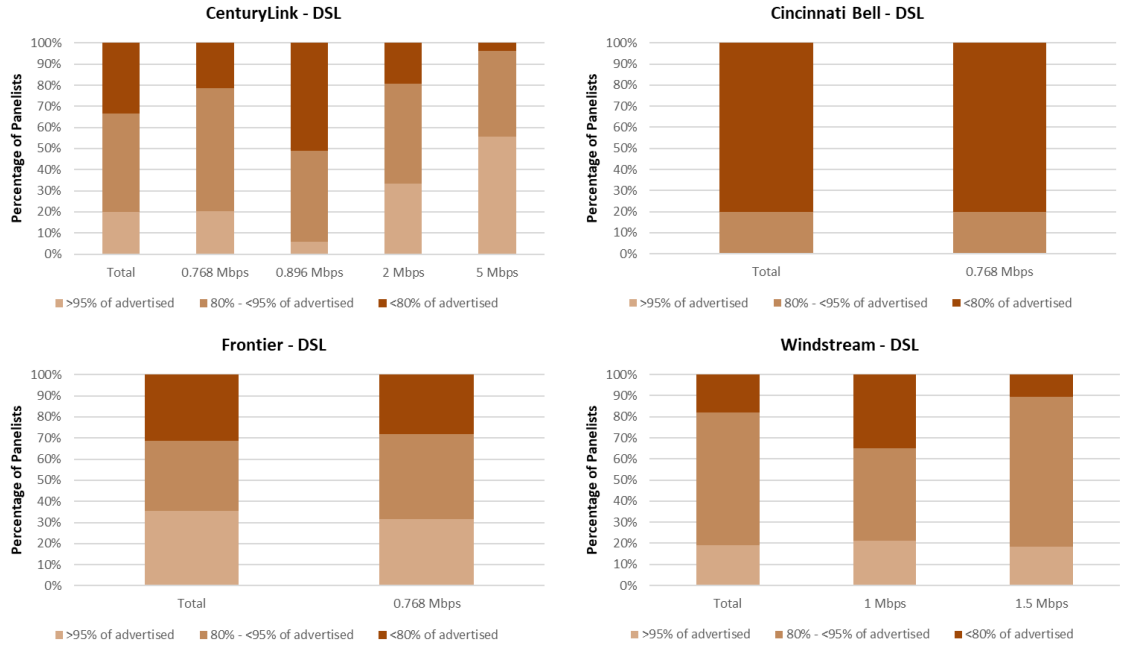

Chart 22.2: The percentage of consumers whose median upload speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised upload speed (cable).

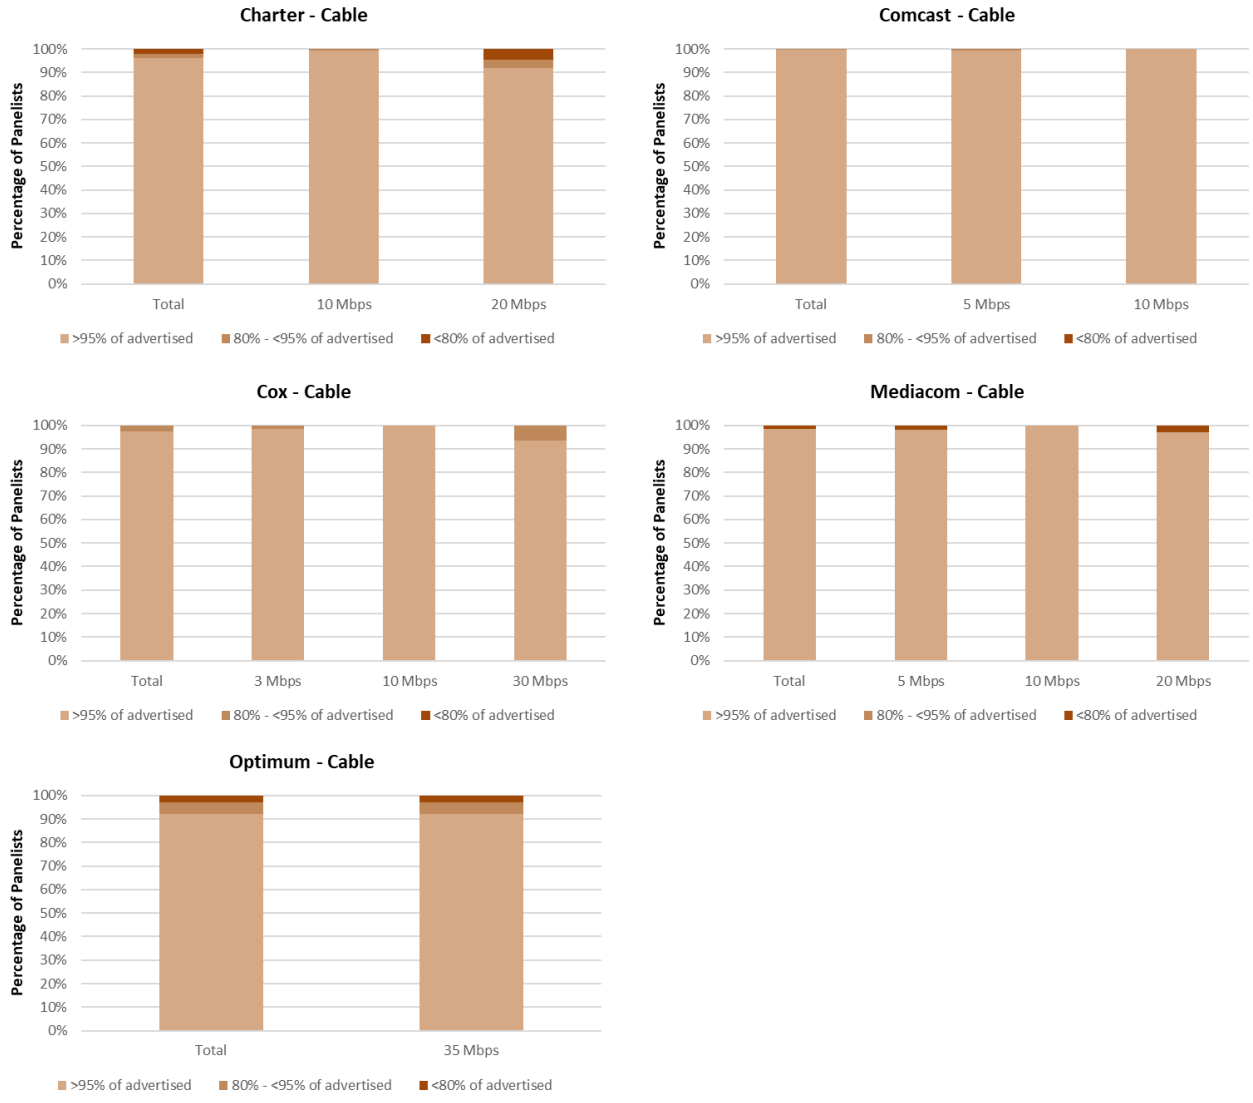

Chart 22.3: The percentage of consumers whose median upload speed was greater than 95%, between 80% and 95%, or less than 80% of the advertised upload speed (fiber).

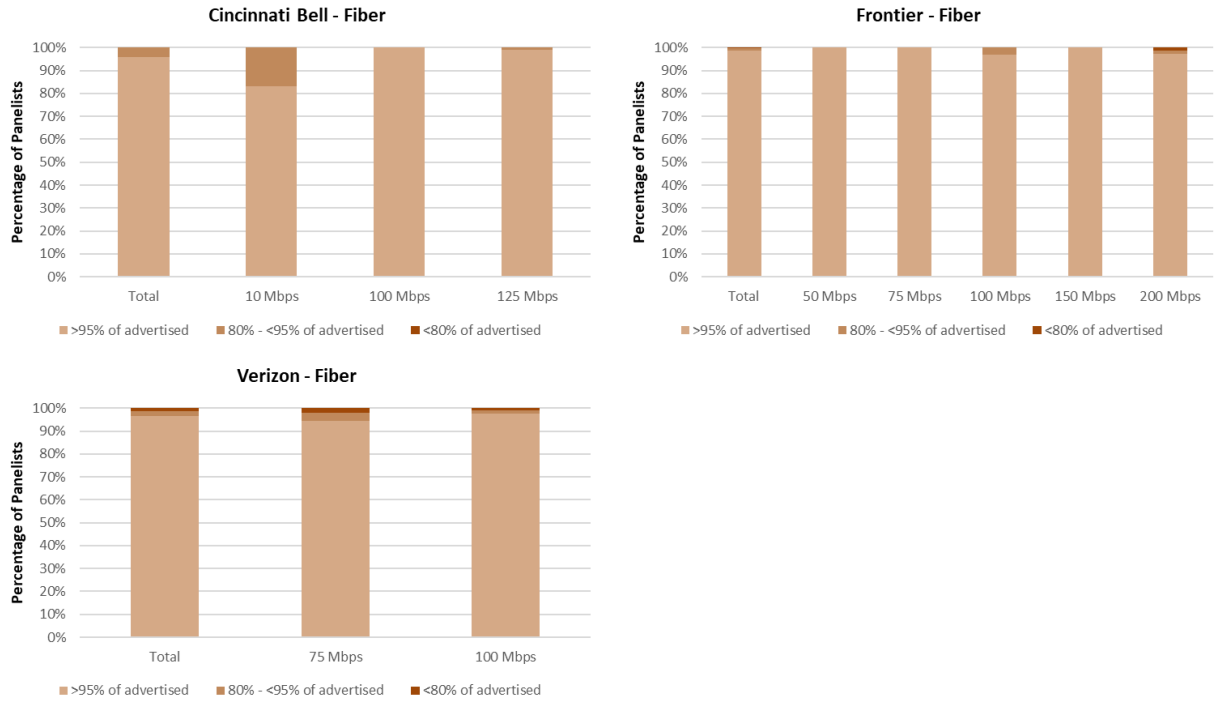

In Section 3.C above, we present complementary cumulative distributions for each ISP based on test results across all service tiers. Below, we provide tables showing selected points on these distributions by each individual ISP. In general, DSL technology showed performance between 25% and 77% of advertised speed for at least 95% of their subscribers. Among cable-based companies, the average download speeds that at least 95% of their subscribers received were between 92% and 100% of advertised rates. Fiber-based services provided a range from 71% to 96% of advertised download speeds for at least 95% of subscribers.

Table 3: Complementary cumulative distribution of the ratio of median download speed to advertised download speed by ISP.

ISP	20%	50%	70%	80%	90%	95%
CenturyLink	101.5%	91.3%	82.7%	78.2%	68.6%	60.9%
Cincinnati Bell Fiber	108.8%	106.8%	91.7%	84.4%	75.9%	71.3%
Cincinnati Bell DSL	84.7%	78.7%	63.9%	50.1%	31.9%	24.5%
Charter	117.0%	115.0%	113.0%	110.6%	104.5%	94.1%
Comcast	118.5%	117.1%	115.0%	113.5%	107.5%	95.1%
Cox	117.9%	113.5%	110.0%	106.8%	102.4%	92.4%
Frontier Fiber	108.5%	99.5%	98.3%	96.4%	94.3%	92.7%

Frontier DSL	104.2%	93.0%	87.0%	83.6%	78.9%	66.6%
Mediacom	133.0%	130.3%	123.1%	116.5%	108.1%	99.8%
Optimum	114.6%	112.8%	110.2%	108.2%	101.9%	99.7%
Verizon Fiber	108.6%	99.9%	99.3%	98.6%	97.7%	96.3%
Windstream	106.6%	100.1%	94.9%	89.2%	82.2%	77.2%

Table 4: Complementary cumulative distribution of the ratio of median upload speed to advertised upload speed by ISP.

ISP	20%	50%	70%	80%
CenturyLink	94.9%	83.7%	78.6%	75.1%
Cincinnati Bell Fiber	107.6%	107.0%	106.7%	106.4%
Cincinnati Bell DSL	79.9%	76.7%	72.4%	71.5%
Charter	116.1%	114.5%	112.5%	110.3%
Comcast	118.6%	117.5%	117.1%	116.6%
Cox	103.4%	103.0%	102.3%	101.5%
Frontier Fiber	119.1%	104.7%	101.7%	100.9%
Frontier DSL	100.7%	89.7%	79.5%	74.0%
Mediacom	122.6%	115.9%	113.7%	113.0%
Optimum	104.1%	103.1%	102.2%	101.3%
Verizon Fiber	118.5%	117.0%	116.4%	113.0%
Windstream	94.9%	91.5%	89.1%	83.3%

F. WEB BROWSING PERFORMANCE, BY SERVICE TIER

Below, we provide the detailed results of the webpage download time for each individual service tier of each ISP. Generally, website loading time decreased steadily with increasing tier speed until a tier speed of 25 Mbps and does not change markedly above that speed.

Chart 23.1: Average webpage download time, by ISP (1.5-5 Mbps).

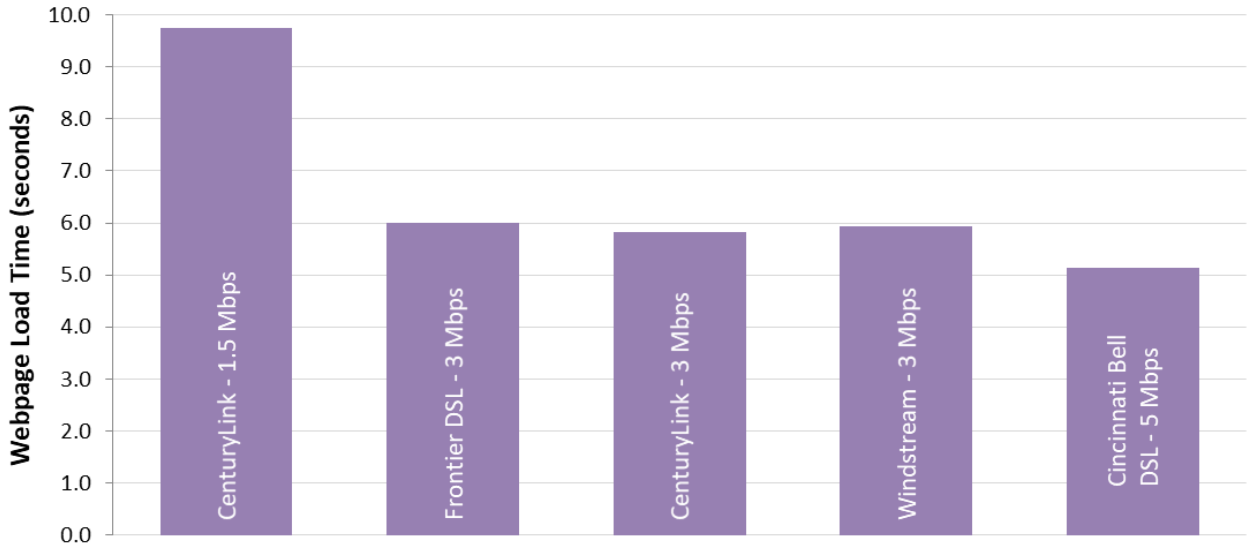

Chart 23.2: Average webpage download time, by ISP (6-10 Mbps).

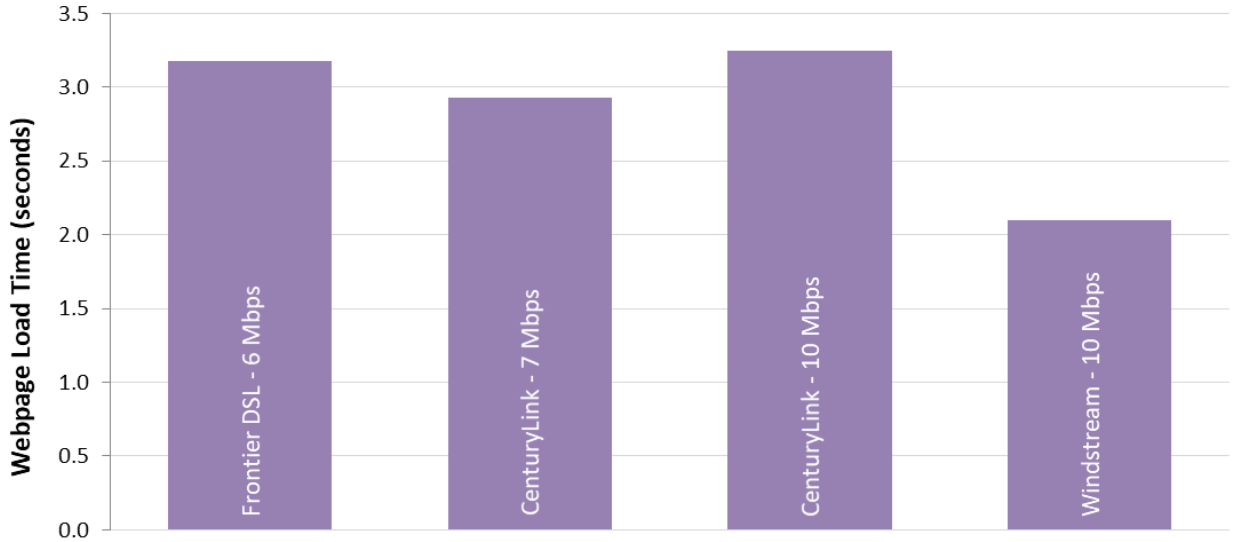

Chart 23.3: Average webpage download time, by ISP (12-25 Mbps).

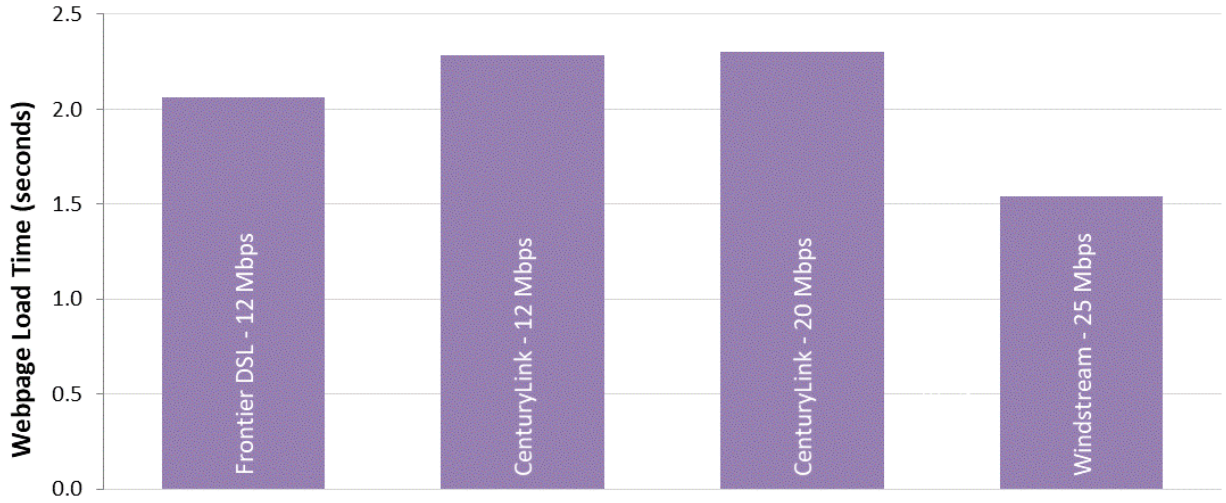

Chart 23.4: Average webpage download time, by ISP (30-60Mbps).

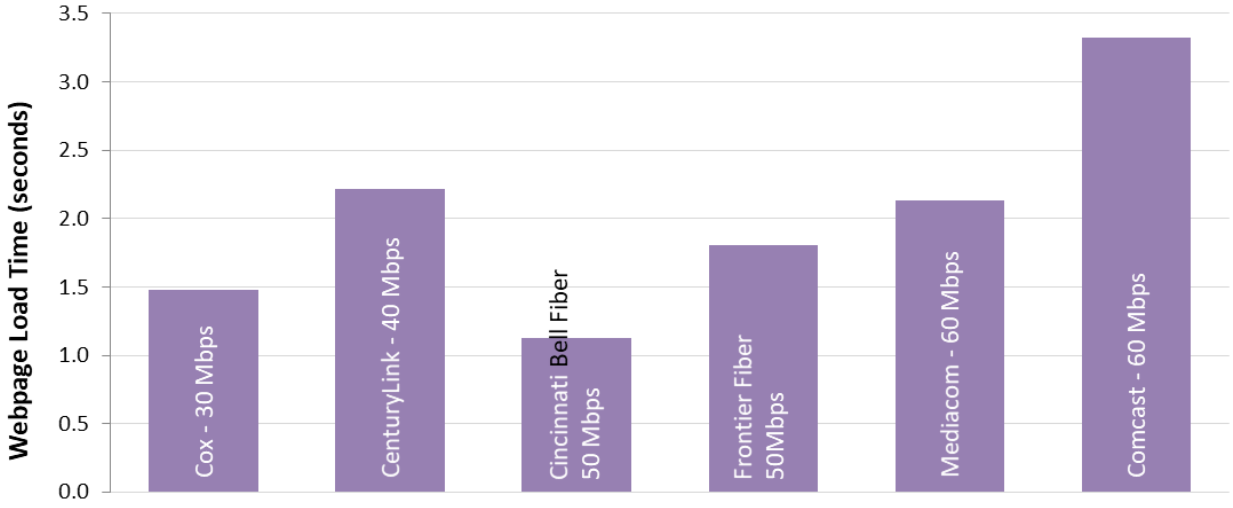

Chart 23.5: Average webpage download time, by ISP (75 - 100 Mbps).

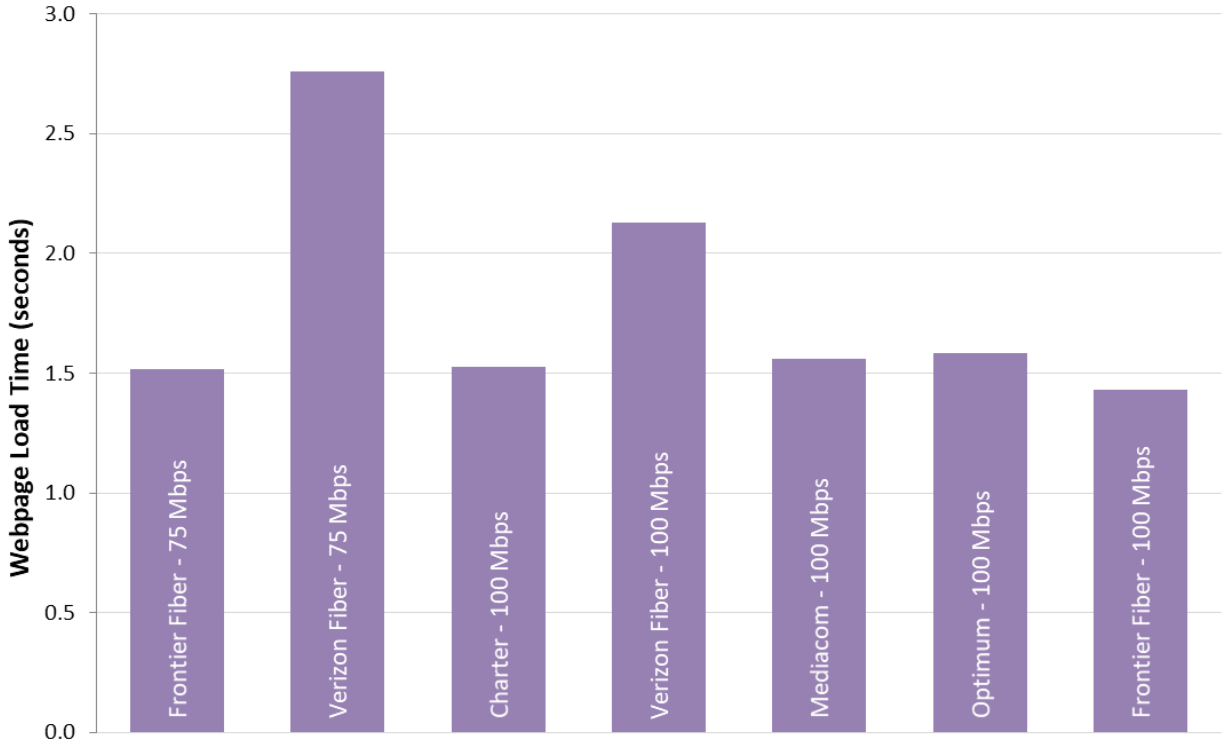

Chart 23.6: Average webpage download time, by ISP (150 - 200 Mbps).

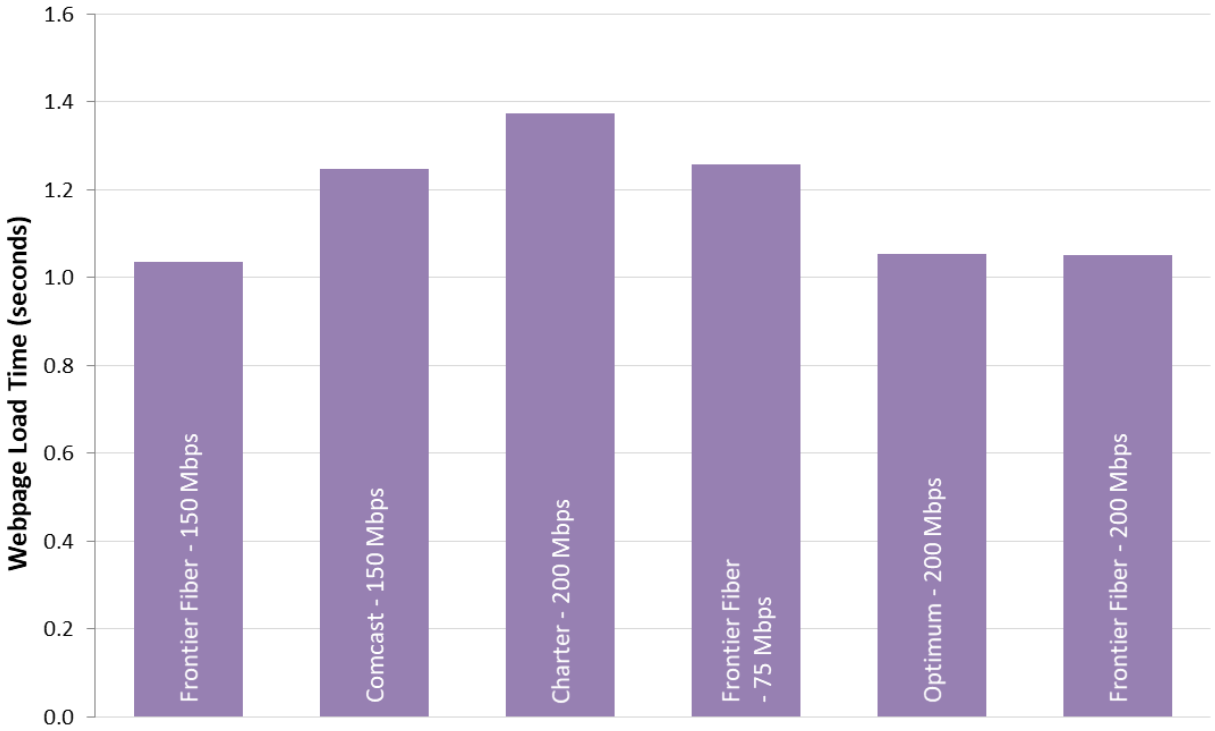

Chart 23.7: Average webpage download time, by ISP (250 - 500 Mbps).

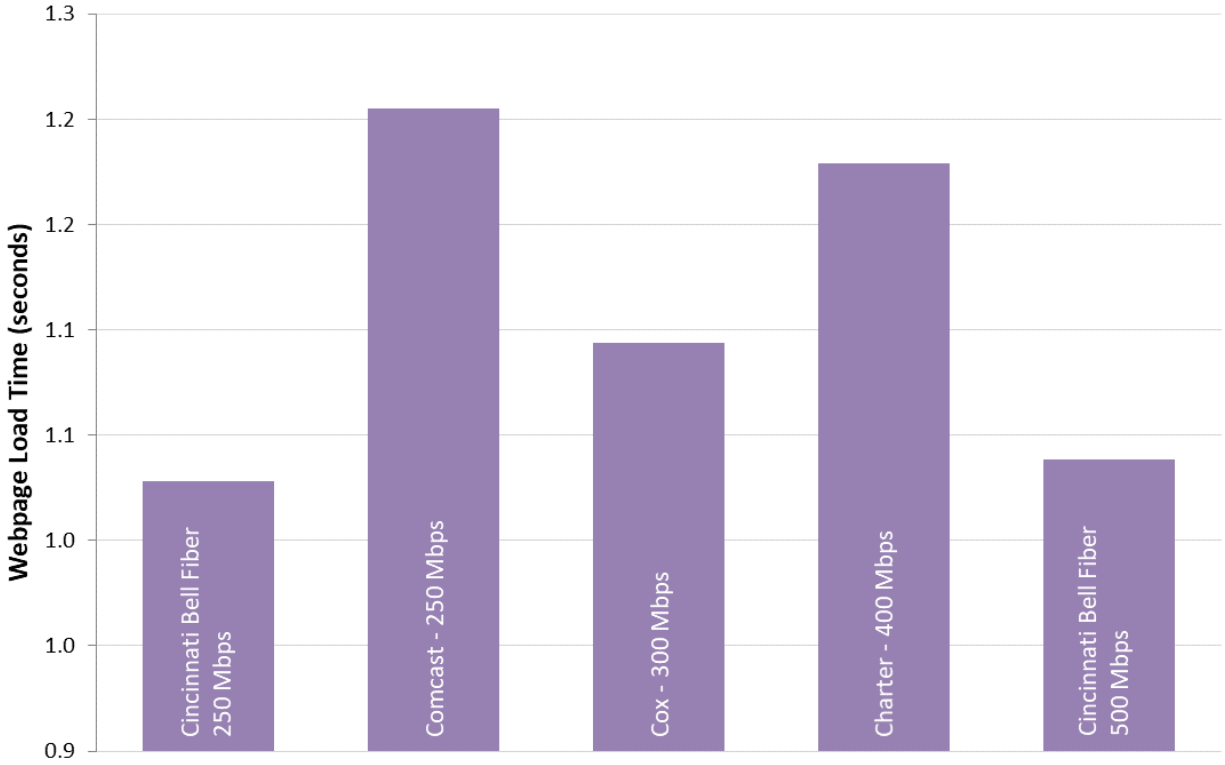

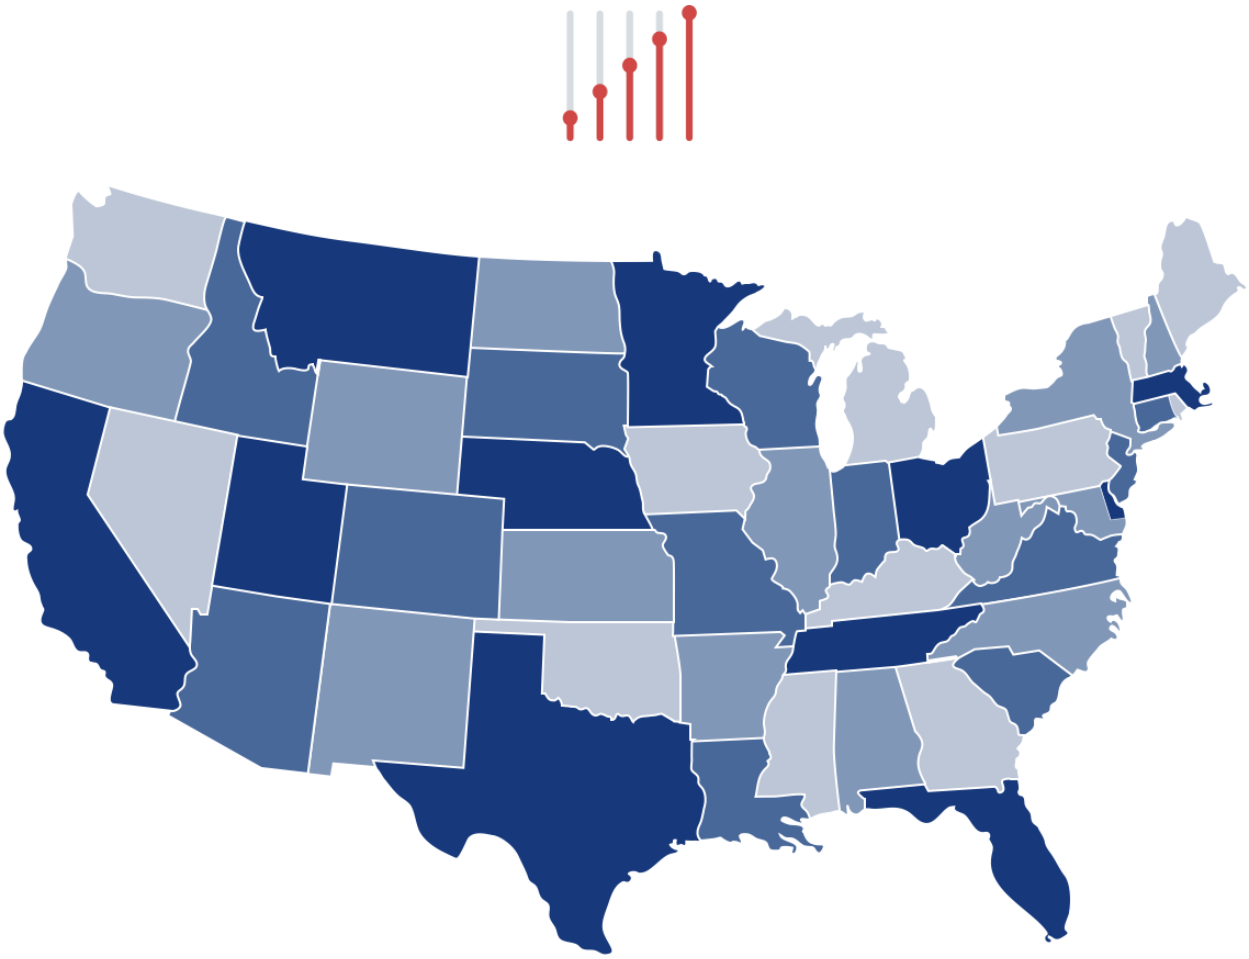

Measuring Broadband America

Technical Appendix to the Tenth MBA Report

FCC's Office of Engineering and Technology

Table of Contents

LIST OF TABLES

LIST OF FIGURES

1 - INTRODUCTION AND SUMMARY

This Appendix to the Tenth Measuring Broadband America Report,[1] a report on consumer wireline broadband performance in the United States, provides detailed technical background information on the methodology that produced the Report. It covers the process by which the panel of consumer participants was originally recruited and selected for the August 2011 MBA Report, and maintained and evolved over the last ten years. This Appendix also discusses the testing methodology used for the Report and describes how the test data was analyzed.

2 - PANEL CONSTRUCTION

This section describes the background of the study, as well as the methods employed to design the target panel, select volunteers for participation, and manage the panel to maintain the operational goals of the program.

The study aims to measure fixed broadband service performance in the United States as delivered by an Internet Service Provider (ISP) to the consumer's broadband modem. Many factors contribute to end-to-end broadband performance, only some of which are under the control of the consumer's ISP. The methodology outlined here is focused on the measurement of broadband performance within the scope of an ISP's network, and specifically focuses on measuring performance from the consumer Internet access point, or consumer gateway, to a close major Internet gateway point. The actual quality of experience seen by consumers depends on many other factors beyond the consumer's ISP, including the performance of the consumer's in-home network, transit providers, interconnection points, content distribution networks (CDN) and the infrastructure deployed by the providers of content and services. The design of the study methodology allows it to be integrated with other technical measurement approaches that focus on specific aspects of broadband performance (i.e., download speed, upload speed, latency, packet loss), and in the future, could focus on other aspects of broadband performance.

[1] The First Report (2011) was based on measurements taken in March 2011, the Second Report (2012) on measurements taken in April 2012, and the Third (2013) through this, the Tenth (2020) Reports on measurements taken in September of the year prior to the reports' release dates.

2.1 - USE OF AN ALL VOLUNTEER PANEL

During a 2008 residential broadband speed and performance test in the United Kingdom,[2] SamKnows[3] had determined that attrition rates of an all-volunteer panel was lower than a panel maintained with an incentive scheme of monthly payments. Consequently, in designing the methodology for this broadband performance study, the Commission had decided to rely entirely on volunteer consumer broadband subscribers. Volunteers are selected from a large pool of prospective participants according to a plan designed to generate a representative sample of desired consumer demographics, including geographical location, ISP, and speed tier. As an incentive for participation, volunteers are given access to a personal dashboard which allows them to monitor the performance of their broadband service. They are also provided with a measurement device referred to in the study as a "Whitebox," consisting of an off-the-shelf commodity router configured to run custom SamKnows software.[4]

2.2 - SAMPLE SIZE AND VOLUNTEER SELECTION

The Tenth MBA Report relies on data gathered from 2,931 volunteer panelists across the United States. The methodological factors and considerations that influenced the selection of the sample size and makeup include proven practices originating from the first MBA report and test period, and adaptations beyond the first period. Both are described below:

- The panel of U.S. broadband subscribers was initially drawn from a pool of over 175,000 volunteers during a recruitment campaign that ran in May 2010. Since then, to manage attrition and accommodate the evolving range of subscriber demographics (*i.e.*, tiers, technology, population), additional panelists have been recruited through email solicitations by the ISPs as well as through press releases, a web page,[5] social media outreach and blog posts.

[2] *See* https://files.samknows.com/~fcc_public/PM_Summer_08.pdf, (last accessed June 21, 2016).

[3] SamKnows is a company that specializes in broadband availability measurement and was retained under contract by the FCC to assist in this study. *See* http://www.samknows.com/.

[4] The Whiteboxes are named after the appearance of the first hardware implementation of the measurement agent. The Whiteboxes remain in consumer homes and continue to run the tests described in this report. Participants may remain in the measurement project as long as it continues and may retain their Whitebox when they end their participation.

[5] https://www.measuringbroadbandamerica.com/

- The volunteer sample was originally organized with a goal of covering major ISPs in the 48 contiguous states across five broadband technologies: DSL, cable, fiber-to-the-home, fixed terrestrial wireless, and satellite.[6]
- Target numbers for volunteers were set across the four Census Regions—Northeast, Midwest, South, and West—to help ensure geographic diversity in the volunteer panel and compensate for differences in networks across the United States.[7]
- A target plan for allocation of Whiteboxes was developed based on the market share of participating ISPs. Initial market share information was based principally on FCC Form 477[8] data filed by participating ISPs for December 2018. This data is further enhanced by the ISPs who brief SamKnows on new products and changes in subscribership numbers which may have occurred after the submission of the 477 data. Speed tiers that comprise the top 80% of a Participating ISP's subscriber base are included. This threshold ensures that we are measuring the ISP's most popular speed tiers and that it is possible to recruit sufficient panelists.
- An initial set of prospective participants was selected from volunteers who had responded directly to SamKnows as a result of media solicitations, as described in detail in Section 2.3. Where gaps existed in the sample plan, SamKnows worked with participating ISPs via email solicitations targeted at underrepresented tiers.
- Since the initial panel was created in 2011, participating ISPs have contacted random subsets of their subscribers by email to replenish cells that were falling short of their desired panel size. Additional recruitment via social media, press releases and blog posts has also taken place.

The sample plan is designed prior to the reporting period and is sent to each ISP by SamKnows. ISPs review this and respond directly to SamKnows with feedback on speed tiers that ought to be included based on the threshold criteria stated above. SamKnows will include all relevant tiers in the final report, assuming a target sample size is available. As this may not be known until after the reporting period is over, a final sample description containing all included tiers is produced and shared with the FCC and ISPs once the reporting period has finished and the data has been processed. Test results from a total of 2,931 panelists were used in the Tenth MBA Report. This

---

[6] At the request of, and with the cooperation of the Department of Commerce and Consumer Affairs, Hawaii, we are also collecting data from the state of Hawaii.

[7] Although the Commission's volunteer recruitment was guided by Census Region to ensure the widest possible distribution of panelists throughout the United States, as discussed below, a sufficient number of testing devices were not deployed to enable, in every case, the evaluation of regional differences in broadband performance. The States associated with each Census Region are described in Table 4.

[8] The FCC Form 477 data collects information about broadband connections to end user locations, wired and wireless local telephone services, and interconnected Voice over Internet Protocol (VoIP) services. *See* https://www.fcc.gov/general/broadband-deployment-data-fcc-form-477 for further information.

figure includes only panelists that are subscribed to the tiers that were tested as part of the sample plan.

The recruitment campaign resulted in the coverage needed to ensure balanced representation of users across the United States. Table 1 shows the number of volunteers with reporting Whiteboxes for the months of September/October 2019 listed by ISP, as well as the percentage of total volunteers subscribed to each ISP. Tables 2 and 3 shows the distributions of the Whiteboxes by State and by Region respectively. This can be compared with the percentage of subscribers per state or region.[9]

Table 1: ISPs, Sample Sizes and Percentages of Total Volunteers

ISP	Sample Size	% of Total Volunteers
CenturyLink	571	19.48%
Charter	250	8.53%
Cincinnati Bell DSL	66	2.25%
Cincinnati Bell Fiber	243	8.29%
Comcast	276	9.42%
Cox	197	6.72%
Frontier DSL	222	7.57%
Frontier Fiber	333	11.36%
Mediacom	188	6.41%
Optimum	162	5.53%
Verizon Fiber	177	6.04%
Windstream	246	8.39%
Total	**2,931**	**100%**

[9] Subscriber data in the Tenth MBA Report is based on the FCC's Internet Access Services Report with data current to December 31, 2017. See *Internet Access Services: Status as of Dec 30, 2017*, Wireline Competition Bureau, Industry Analysis and Technology Division (rel. Nov. 2018), available at https://docs.fcc.gov/public/attachments/DOC-359342A1.pdf.

Table 2: Distribution of Whiteboxes by State

State	Total Boxes	% of Total Boxes	% of Total US Broadband
Alabama	25	0.9%	1.50%
Alaska	0	0.0%	0.23%
Arizona	121	4.1%	1.97%
Arkansas	15	0.5%	0.86%
California	179	6.1%	12.17%
Colorado	86	2.9%	1.71%
Connecticut	56	1.9%	1.13%
Delaware	9	0.3%	0.30%
District of Columbia	4	0.1%	0.27%
Florida	182	6.2%	6.56%
Georgia	85	2.9%	3.18%
Hawaii	17	0.6%	0.47%
Idaho	23	0.8%	0.48%
Illinois	47	1.6%	3.92%
Indiana	43	1.5%	1.92%
Iowa	146	5.0%	0.90%
Kansas	14	0.5%	1.21%
Kentucky	106	3.6%	1.35%
Louisiana	15	0.5%	1.41%
Maine	0	0.0%	0.42%
Maryland	37	1.3%	1.91%
Massachusetts	48	1.6%	2.27%
Michigan	33	1.1%	2.98%
Minnesota	97	3.3%	1.68%
Mississippi	2	0.1%	0.86%
Missouri	67	2.3%	1.78%

Montana	9	0.3%	0.32%
Nebraska	24	0.8%	0.54%
Nevada	24	0.8%	0.91%
New Hampshire	8	0.3%	0.42%
New Jersey	103	3.5%	2.86%
New Mexico	45	1.5%	0.58%
New York	151	5.2%	6.51%
North Carolina	75	2.6%	3.02%
North Dakota	2	0.1%	0.24%
Ohio	316	10.8%	3.55%
Oklahoma	18	0.6%	1.09%
Oregon	78	2.7%	1.26%
Pennsylvania	109	3.7%	3.93%
Rhode Island	7	0.2%	0.31%
South Carolina	15	0.5%	1.45%
South Dakota	3	0.1%	0.26%
Tennessee	13	0.4%	2.01%
Texas	166	5.7%	8.16%
Utah	17	0.6%	0.83%
Vermont	2	0.1%	0.20%
Virginia	87	3.0%	2.44%
Washington	126	4.3%	2.31%
West Virginia	21	0.7%	0.46%
Wisconsin	52	1.8%	1.64%
Wyoming	3	0.1%	0.19%
	2,931		

The distribution of Whiteboxes by Census Region is found in the table on the next page.

Table 3: Distribution of Whiteboxes by Census Region

Census Region	Total Boxes	% Total Boxes	% Total U.S. Broadband Subscribers
Midwest	844	28.8%	21%
Northeast	484	16.5%	18%
South	875	29.9%	37%
West	728	24.8%	24%

The distribution of states associated with the four Census Regions used to define the panel strata are included in the table below.

Table 4: Panelists States Associated with Census Regions

Census Region	States
Northeast	CT MA ME NH NJ NY PA RI VT
Midwest	IA IL IN KS MI MN MO ND NE OH SD WI
South	AL AR DC DE FL GA KY LA MD MS NC OK SC TN TX VA WV
West	AK AZ CA CO HI ID MT NM NV OR UT WA WY

2.3 - PANELIST RECRUITMENT PROTOCOL

Panelists are recruited in the 2011- 2019 panels using the following method:

- Recruitment has evolved since the start of the program. At that time, (2011) several thousand volunteers were initially recruited through an initial public relations and social media campaign led by the FCC. This campaign included discussion on the FCC website and on technology blogs, as well as articles in the press. Currently volunteers are drafted with the help of a recruitment website[10] which keeps them informed about the MBA program and allows them to view MBA data on a dashboard. The composition of the panel is reviewed each year to identify any deficiencies with regard to the sample plan described above. Target demographic goals are set for volunteers based on ISP, speed tier, technology type, and region. Where the pool of volunteers falls short of the desired goal, ISPs send out email messages to their customers asking them to participate in the MBA program. The messages direct interested volunteers to contact SamKnows to request participation in the trial. The ISPs do not know which of the email recipients volunteer. In almost all cases, this ISP outreach allows the program to meet its desired demographic targets.

The mix of panelists recruited using the above methodologies varies by ISP.

A multi-mode strategy was used to qualify volunteers for the 2019 testing period. The key stages of this process were as follows:

1. Volunteers were directed to complete an online form which provided information on the study and required volunteers to submit a small amount of information.
2. Volunteers were selected from respondents to this follow-up email based on the target requirements of the panel. Selected volunteers were then asked to agree to the *User Terms and Conditions* that outlined the permissions to be granted by the volunteer in key areas such as privacy.[11]
3. From among the volunteers who agreed to the User Terms and Conditions, SamKnows selected the panel of participants,[12] each of whom received a Whitebox for self-installation. SamKnows provided full support during the Whitebox installation phase.

The graphic in Figure 1 illustrates the study recruitment methodology.

[10] The Measuring Broadband America recruitment website is: https://www.measuringbroadbandamerica.com/.

[11] The *User Terms and Conditions* is found in the Reference Documents at the end of this Appendix.

[12] Over 23,000 Whiteboxes have been shipped to targeted volunteers since 2011, of which 6,006 were online and reporting data from the months of September/October 2019.

Figure 1: Panelist Recruitment Protocol

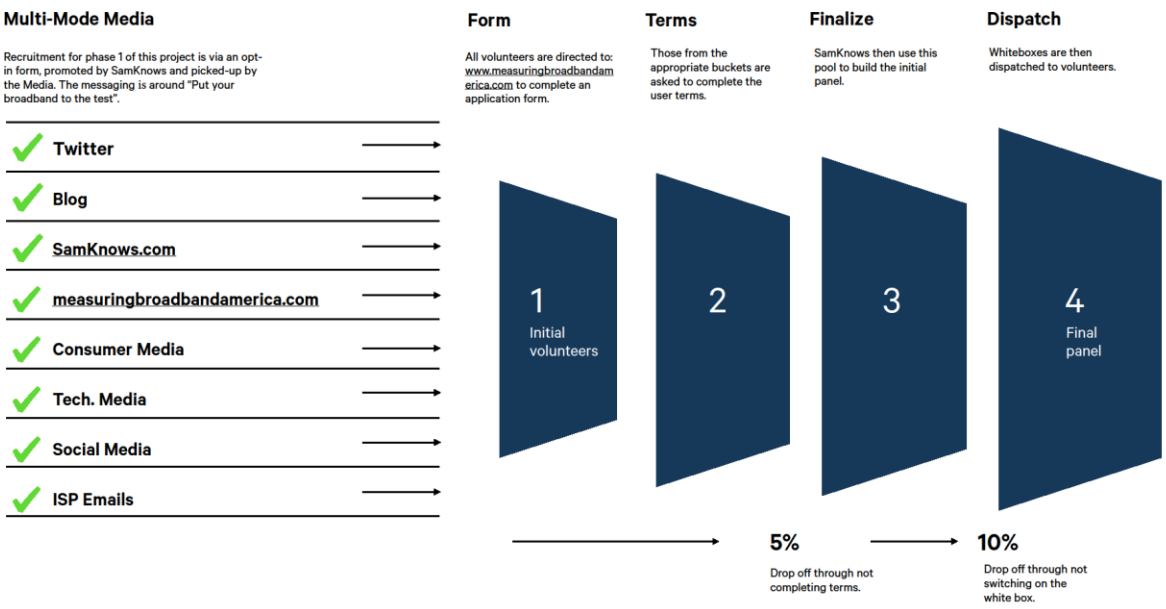

2.4 - VALIDATION OF VOLUNTEERS' SERVICE TIER

The methodology employed in this study included verifying each panelist's service tier and ISP against the customer records of participating ISPs.[13] Initial throughput tests were used to confirm reported speeds.

The broadband service tier reported by each panelist was validated as follows:

- When the panelist installed the Whitebox, the device automatically ran an IP address test to check that the ISP identified by the volunteer was correct.
- The Whitebox also ran an initial test which flooded each panelist's connection in order to accurately detect the throughput speed when their deployed Whitebox connected to a test node.
- Each ISP was asked to confirm the broadband service tier reported by each selected panelist.
- SamKnows then took the validated speed tier information that was provided by the ISPs and compared this to both the panelist-provided information, and the actual test results obtained, in order to ensure accurate tier validation.

SamKnows manually completed the following four steps for each panelist:

- Verified that the IP address was in a valid range for those served by the ISP.
- Reviewed data for each panelist and removed data where speed changes such as tier upgrade or downgrade appeared to have occurred, either due to a service change on the part of the consumer or a network change on the part of the ISP.
- Identified panelists whose throughput appeared inconsistent with the provisioned service tier. Such anomalies were re-certified with the consumer's ISP.[14]
- Verified that the resulting downstream-upstream test results corresponded to the ISP-provided speed tiers and updated accordingly if required.

[13] Past FCC studies found that a high rate of consumers could not reliably report information about their broadband service, and the validation of subscriber information ensured the accuracy of expected speed and other subscription details against which observed performance was measured. *See* John Horrigan and Ellen Satterwhite, *Americans' Perspectives on Online Connection Speeds for Home and Mobile Devices*, 1 (FCC 2010), available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-298516A1.doc (finding that 80 percent of broadband consumers did not know what speed they had purchased).

[14] For example, when a panelist's upload or download speed was observed to be significantly higher than that of the rest of the tier, it could be inferred that a mischaracterization of the panelist's service tier had occurred. Such anomalies, when not resolved in cooperation with the service provider, were excluded from the Tenth Report, but will be included in the raw bulk data set.

Of the more than 23,000 Whiteboxes that were shipped to panelists since 2011, 6,006[15] units reported sufficient data in September/October 2019, with the participating ISPs validating 4,964 for the reporting period. Of the validated units, 17 percent were reallocated to a different tier following the steps listed above. A total of 2,931 validated units were part of download or upload tiers included in the sample plan and were ultimately included in this report.

A total of 3,075 boxes were excluded for the following reasons:

- 1,763 belonged to users subscribed to plans that were not included in this study
- 263 were excluded due to legacy equipment such as modem that could not fully support the subscribed speeds
- 291 Whiteboxes were legacy models that could not fully support the plan speeds
- 293 belonged to users whose details or subscribed tier could not be successfully validated by the ISP
- 142 Whiteboxes were excluded due to ethernet limitations
- 23 were connected to non-residential plans
- 1 Whitebox was a test unit not to be included in the program
- 7 belonged to employees of ISPs taking part in the MBA program
- And a further 292 were excluded as the test speed profile did not match the product validated by the ISP.

2.5 - PROTECTION OF VOLUNTEERS' PRIVACY

Protecting the panelists' privacy is a major concern for this program. The panel was comprised entirely of volunteers who knowingly and explicitly opted into the testing program. For audit purposes, we retain the correspondence with panelists documenting their opt-in.

All personal data was processed in conformity with relevant U.S. law and in accordance with policies developed to govern the conduct of the parties handling the data. The data were processed solely for the purposes of this study and are presented here and in all online data sets with all personally identifiable information (PII) removed.

A set of materials was created both to inform each panelist regarding the details of the trial, and to gain the explicit consent of each panelist to obtain subscription data from the participating ISPs. These documents were reviewed by the Office of General Counsel of the FCC and the participating ISPs and other stakeholders involved in the study.

[15] This figure represents the total number of boxes reporting during September/October 2019, the month chosen for the Tenth Report. Shipment of boxes continued in succeeding months and these results will be included in the raw bulk data set.

3 - BROADBAND PERFORMANCE TESTING METHODOLOGY

This section describes the system architecture and network programming features of the tests, and other technical aspects of the methods employed to measure broadband performance during this study.

3.1 - RATIONALE FOR HARDWARE-BASED MEASUREMENT APPROACH

Either a hardware or software approach can be used to measure broadband performance. Software approaches are by far the most common and allow for measurements to easily and cost-effectively include a very large sample size. Web-based speed tests fall into this category and typically use Flash applets, Java applets or JavaScript that execute within the user's web browser. These clients download content from remote web servers and measure the throughput. Some web-based performance tests also measure upload speed or round-trip latency.

Other, less common, software-based approaches to performance measurement install applications on the user's computer. These applications run tests periodically while the computer is on.

All software solutions implemented on a consumer's computer, smart phone, or other device connected to the Internet suffer from the following disadvantages:

- The software and computing platform running the software may not be capable of reliably recording the higher speed service tiers currently available.
- The software typically cannot know if other devices on the home network are accessing the Internet when the measurements are being taken. The lack of awareness as to other, non-measurement related network activity can produce inconsistent and misleading measurement data.
- Software measurements may be affected by the performance, quality and configuration of the device.
- Potential bottlenecks, such as Wi-Fi networks and other in-home networks, are generally not accounted for and may result in unreliable data.
- If the device hosting the software uses in-home WIFI access to fixed broadband service, differing locations in the home may impact measurements.
- The tests can only run when the computer is turned on, limiting the ability to provide a 24-hour profile.

- If software tests are performed manually, panelists might only run tests when they experience problems and thus bias the results.

In contrast, the hardware approach used in the MBA program requires the placement of the previously described Whitebox inside the user's home, directly connected to the consumer's service interconnection device (router), via Ethernet cable. The measurement device therefore directly accesses fixed Internet service to the home over this dedicated interface and periodically runs tests to remote targets over the Internet. The use of hardware devices avoids the disadvantages listed earlier with the software approach. However, hardware approaches are much more expensive than the software alternative, are thus more constrained in the achievable panel size, and require correct installation of the device by the consumer or a third party. This is still subject to unintentional errors due to misconfigurations, *i.e.*, connecting the Whitebox incorrectly but these can often be detected in the validation process that follows installation. The FCC chose the hardware approach since its advantages far outweigh these disadvantages.

3.2 - DESIGN OBJECTIVES AND TECHNICAL APPROACH

For this test of broadband performance, as in previous Reports, the FCC used design principles that were previously developed by SamKnows in conjunction with their study of broadband performance in the U.K. The design principles comprise 17 technical objectives:

Table 5: Design Objectives and Methods

#	Technical Objectives	Methodological Accommodations
1	The Whitebox measurement process must not change during the monitoring period.	The Whitebox measurement process is designed to provide automated and consistent monitoring throughout the measurement period.
2	Must be accurate and reliable.	The hardware solution provides a uniform and consistent measurement of data across a broad range of participants.
3	Must not interrupt or unduly degrade the consumer's use of the broadband connection.	The volume of data produced by tests is controlled to avoid interfering with panelists' overall broadband experience, and tests only execute when consumer is not making heavy use of the connection.
4	Must not allow collected data to be distorted by any use of the broadband connection by other applications on the host PC and other devices in the home.	The hardware solution is designed not to interfere with the host PC and is not dependent on that PC.

5	Must not rely on the knowledge, skills and participation of the consumer for its ongoing operation once installed.	The Whitebox is "plug-and-play." Instructions are graphics-based and the installation process has been substantially field tested. Contacts for support are also provided and the outreach once a Whitebox has been dispatched and activated.
6	Must not collect data that might be deemed to be personal to the consumer without consent.	The data collection process is explained in plain language and consumers are asked for their consent regarding the use of their personal data as defined by any relevant data protection legislation.
7	Must be easy for a consumer to completely remove any hardware and/or software components if they do not wish to continue with the MBA program.	Whiteboxes can be disconnected at any time from the home network. As soon as the Whitebox is reconnected the reporting is resumed as before.
8	Must be compatible with a wide range of DSL, cable, satellite and fiber-to-the-home modems.	Whiteboxes can be connected to all modem types commonly used to support broadband services in the U.S., either in a routing or bridging mode, depending on the model.
9	Where applicable, must be compatible with a range of computer operating systems, including, without limitation, Windows XP, Windows Vista, Windows 7, Mac OS and Linux.	Whiteboxes are independent of the PC operating system and therefore able to provide testing with all devices regardless of operating system.
10	Must not expose the volunteer's home network to increased security risk, *i.e.*, it should not be susceptible to viruses, and should not degrade the effectiveness of the user's existing firewalls, antivirus and spyware software.	The custom software in the Whitebox is hardened for security and cannot be accessed without credentials only available to SamKnows. Most user firewalls, antivirus and spyware systems are PC-based. The Whitebox is plugged in to the broadband connection "before" the PC. Its activity is transparent and does not interfere with those protections.
11	Must be upgradeable remotely if it contains any software or firmware components.	The Whitebox can be completely controlled remotely for updates without involvement of the consumer, providing the Whitebox is switched on and connected.
12	Must identify when a user changes broadband provider or package (*e.g.*, by a reverse look up of the consumer's IP address to check provider, and by capturing changes in modem connection speed to identify changes in package).	Ensures regular data pool monitoring for changes in speed, ISP, IP address or performance, and flags when a panelist should notify and confirm any change to their broadband service since the last test execution.

13	Must permit, in the event of a merger between ISPs, separate analysis of the customers of each of the merged ISP's predecessors.	Data are stored based on the ISP of the panelist, and therefore can be analyzed by individual ISP or as an aggregated dataset.
14	Must identify if the consumer's computer is being used on a number of different fixed networks (*e.g.*, if it is a laptop).	The Whiteboxes are broadband dependent, not PC or laptop dependent.
15	Must identify when a specific household stops providing data.	The Whitebox needs to be connected and switched on to push data. If it is switched off or disconnected its absence is detected at the next data push process.
16	Must not require an amount of data to be downloaded which may materially impact any data limits, usage policy, or traffic shaping applicable to the broadband service.	The data volume generated by the information collected does not exceed any policies set by ISPs. Panelists with bandwidth restrictions can have their tests set accordingly.
17	Must limit the possibility for ISPs to identify the broadband connections which form their panel and therefore potentially "game" the data by providing different quality of service to the panel members and to the wider customer base.	ISPs signed a Code of Conduct[16] to protect against gaming test results. While the identity of each panelist was made known to the ISP as part of the speed tier validation process, the actual Unit ID for the associated Whitebox was not released to the ISP so specific test results were not directly assignable against a specific panelist. Moreover, most ISPs had hundreds, and some had more than 1,000, participating subscribers spread throughout their service territory, making it difficult to improve service for participating subscribers without improving service for all subscribers.

[16] Signatories to the Code of Conduct are: CenturyLink, Charter, Cincinnati Bell, Comcast, Cox, Frontier, Level3, Measurement Lab, Mediacom, NCTA, Optimum, Time Warner Cable, Verizon and Windstream. A copy of the Code of Conduct is included as a Reference Document attached to this Appendix.

3.3 - TESTING ARCHITECTURE

Overview of Testing Architecture

As illustrated in Figure 2, the performance monitoring system comprises a distributed network of Whiteboxes in the homes of members of the volunteer consumer panel. The Whiteboxes are controlled by a cluster of servers, which hosts the test scheduler and the reporting database. The data was collated on the reporting platform and accessed via a reporting interface[17] and secure FTP site. The system also included a series of speed-test servers, which the Whiteboxes called upon according to the test schedule.

Figure 2: Testing Architecture

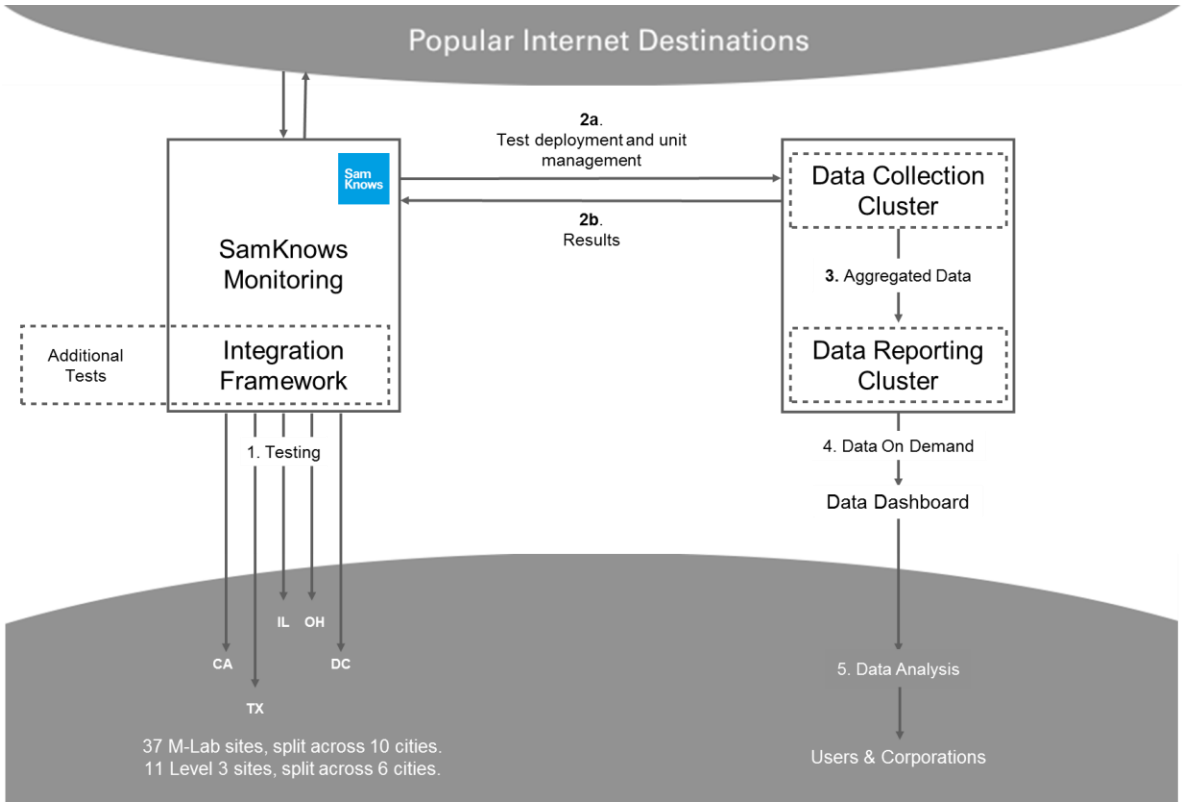

[17] Each reporting interface included a data dashboard for the consumer volunteers, which provided performance metrics associated with their Whitebox.

Approach to Testing and Measurement

Any network monitoring system needs to be capable of monitoring and executing tests 24 hours a day, seven days a week. Similar to the method used by the television audience measurement industry, each panelist is equipped with a Whitebox, which is self-installed by each panelist and conducts the performance measurements. Since 2011, the project has used three different hardware platforms, described below. The software on each of the Whiteboxes was programmed to execute a series of tests designed to measure key performance indicators (KPIs) of a broadband connection. The tests comprise a suite of applications, written by SamKnows in the programming language C, which were rigorously tested by the ISPs and other stakeholders. The Tenth Report incorporates data from all three types of Whiteboxes and we use the term Whitebox generically. Testing has found that they produce results that are indistinguishable.

During the initial testing period in 2011, the Whitebox provided used hardware manufactured by NETGEAR, Inc. (NETGEAR) and operated as a broadband router. It was intended to replace the panelist's existing router and be directly connected to the cable or DSL modem, ensuring that tests could be run at any time the network was connected and powered, even if all home computers were switched off. Firmware for the Whitebox routers was developed by SamKnows with the cooperation of NETGEAR. In addition to running the latest versions of the SamKnows testing software, the routers retained all of the native functionality of the NETGEAR consumer router.

Following the NETGEAR Whitebox new models were introduced starting with the 2012 testing period. These versions were based upon hardware produced by TP-Link and then later manufactured by SamKnows and operate as a bridge rather than as a router. It connects to the customer's existing router, rather than replacing it, and all hardwired home devices connect to LAN ports on the TP-Link Whitebox. The TP-Link Whitebox / SamKnows Whitebox passively monitors wireless network activity in order to determine when the network is active and defer measurements. It runs a modified version of OpenWrt, an open source router platform based on Linux. All Whiteboxes deployed since 2012 use the TP-Link or SamKnows hardware.

SamKnows Whiteboxes (Whitebox 8.0), introduced in August 2016, have been shown to provide accurate information about broadband connections with throughput rates of up to 1 Gbps.

Home Deployment of the NETGEAR Based Whitebox

This study was initiated by using existing NETGEAR firmware, and all of its features were intended to allow panelists to replace their existing routers with the Whitebox. If the panelist did not have an existing router and used only a modem, they were asked to install the Whitebox according to the usual NETGEAR instructions.

However, this architecture could not easily accommodate scenarios where the panelist had a combined modem/router supplied by their ISP that had specific features that the Whitebox could not provide. For example, some Verizon FiOS gateways connect via a MoCA (Multimedia over Cable) interface and AT&T IPBB gateways provide U-Verse specific features, such as IPTV.

In these cases, the Whitebox was connected to the existing router/gateway and all home devices plugged into the Whitebox. In order to prevent a double-NAT configuration, in which multiple routers on the same network perform network address translation (NAT) and make access to the SamKnows router difficult, the Whitebox was set to dynamically switch to operate as a transparent Ethernet bridge when deployed in these scenarios. All consumer configurations were evaluated and tested by participating ISPs to confirm their suitability.[18]

Home Deployment of the TP-Link Based Whitebox

The TP-Link-based Whitebox, which operates as a bridge, was introduced in response to the increased deployment of integrated modem/gateway devices. To use the TP-Link-based Whitebox, panelists are required to have an existing router. Custom instructions guided these panelists to connect the Whitebox to their existing router and then connect all of their home devices to the Whitebox. This allows the Whitebox to measure traffic volumes from wired devices in the home and defer tests accordingly. As an Ethernet bridge, the Whitebox does not provide services such as network address translation (NAT) or DHCP.

Home Deployment of the SamKnows Whitebox 8.0

The Whitebox 8.0 was manufactured by SamKnows and deployed starting in August 2016. Like the TP-Link device, this Whitebox works as a bridge, rather than a router, and operates in a similar manner. Unlike the NETGEAR and TP-Link hardware, it can handle bandwidths of up to 1 Gbps.

Internet Activity Detection

No tests are performed if the Whiteboxes detect wired or wireless traffic beyond a defined bandwidth threshold. This ensures both that testing does not interfere with consumer use of

[18] The use of legacy equipment has the potential to impede some panelists from receiving the provisioned speed from their ISP, and this impact is captured by the survey.

their Internet service and that any such use does not interfere with testing or invalidate test results.

Panelists were not asked to change their wireless network configurations. Since the TP-Link Whiteboxes and Whitebox 8.0 attach to the panelist's router that may contain a built-in wireless (Wi-Fi) access point, these devices measure the strongest wireless signal. Since they only count packets, they do not need access to the Wi-Fi encryption keys and do not inspect packet content.

Test Nodes (Off-Net and On-Net)

For the tests in this study, SamKnows employed fifty-four core measurement servers as test nodes that were distributed geographically across eleven locations, outside the network boundaries of the participating ISPs. These off-net measurement points were supplemented by additional measurement points located within the networks of some of the ISPs participating in this study, called on-net servers. The core measurement servers were used to measure consumers' broadband performance between the Whitebox and an available reference point that was closest in roundtrip time to the consumer's network address. The distribution of off-net primary reference points operated by M-Lab, Level 3 and Stackpath[19].

On-net secondary reference points operated by broadband providers provided additional validity checks and insight into broadband service performance within an ISP's network. In total, the following 133 measurement servers were deployed for the Tenth Report:

Table 6: Overall Number of Testing Servers

Operated By	Number of Servers
AT&T	9
CenturyLink (inc Qwest)	14
Charter (inc TWC)	18
Comcast	37
Cox	2

[19] Stackpath was added to the list of hosting providers for the MBA project to provide further resilience for the testing platform. Stackpath servers have a minimum 10Gbps – 200Gbps transit / peering links and have been located in the major US cities as per the other hosting providers used for the program.

Frontier	5
Hawaiian Telecom	1
Level 3 (off-net)	13
M-Lab (off-net)	51
Mediacom	1
Optimum	3
Time Warner Cable (now part of Charter)	18
Uhnet (Hawaii)	1
Verizon	2
Windstream	4
Stackpath	10

Test Node Locations

Off-Net Test Nodes

The M-Lab test nodes were located in the following major U.S. Internet peering locations:

- New York City, New York (five locations)
- Chicago, Illinois (five locations)
- Atlanta, Georgia (five locations)
- Miami, Florida (five locations)
- Washington, DC (five locations)
- Mountain View, California (six locations)

- Seattle, Washington (six locations)
- Los Angeles, California (five locations)
- Dallas, Texas (five locations)
- Denver, Colorado (four locations)

The Level 3 nodes were located in the following major U.S. Internet peering locations:

- Chicago, Illinois (two locations)
- Dallas, Texas (two locations)
- New York City, New York (two locations)
- San Jose, California (two locations)
- Washington D.C. (two locations)
- Los Angeles, California (three locations)

The Stackpath nodes were located in the following major U.S. Internet peering locations:

- Ashburn, Virginia (one location)
- Atlanta, Georgia (one location)
- Chicago, Illinois (one location)
- Dallas, Texas (one location)
- Los Angeles, California (one location)
- New York City, New York (one location)
- San Jose, California (one location)
- Seattle, Washington (one location)
- Denver, Colorado (one location)
- Miami, Florida (one location)

On-Net Test Nodes

In addition to off-net nodes, some ISPs deployed their own on-net servers to cross-check the results provided by off-net nodes. Whiteboxes were instructed to test against the off-net M-Lab, Stackpath and Level 3 nodes and the on-net ISP nodes, when available.

The following ISPs provided on-net test nodes:

- CenturyLink[20]
- Charter[21]
- Cincinnati Bell
- Comcast
- Cox
- Frontier
- Mediacom
- Optimum
- Verizon
- Windstream

The same suite of tests was scheduled for these on-net nodes as for the off-net nodes and the same server software developed by SamKnows was used regardless of whether the Whitebox was interacting with on-net or off-net nodes. Off-net test nodes are continually monitored for load and congestion.

While these on-net test nodes were included in the testing, the results from these tests were used as a control set; the results presented in the Report are based only on tests performed using off-net nodes. Results from both on-net and off-net nodes are included in the raw bulk data set that will be released to the public.

Test Node Selection

Each Whitebox fetches a complete list of off-net test nodes and on-net test nodes hosted by the serving ISP from a SamKnows server and measures the round-trip time to each. This list of test servers is loaded at startup and refreshed daily. It then selects the on-net and off-net test nodes with lowest round trip time to test against. The selected nodes may not be the geographically closest node.

Technical details for the minimum requirements for hardware and software, connectivity, and systems and network management are available in the **5.3 - Test Node Briefing** provided in the Reference Document section of this Technical Appendix.

[20] QWest was reported separately from Centurylink in reports prior to 2016. The entities completed merging their test infrastructure in 2016.

[21] Time Warner Cable was reported separately from Charter in reports prior to the Eighth report. The entities completed merging their test infrastructure in early 2018.

3.4 - TESTS METHODOLOGY

Each deployed Whitebox performs the following tests.[22] All tests are conducted with both the on-net and off-net servers except as noted, and are described in more detail in the next section.

Table 7: List of Tests Performed by SamKnows[23]

Metric	Primary Metric(s)
Download speed	Throughput in Megabits per second (Mbps) utilizing three concurrent TCP connections
Upload speed	Throughput in Mbps utilizing three concurrent TCP connections
Web browsing	Total page fetch time and all its embedded resources from a popular website
UDP latency	Average round trip time of a series of randomly transmitted UDP packets distributed over a long timeframe
UDP packet loss	Fraction of UDP packets lost from UDP latency test
Voice over IP	Upstream packet loss, downstream packet loss, upstream jitter, downstream jitter, round trip latency
DNS resolution	Time taken for the ISP's recursive DNS resolver to return an A record[24] for a popular website domain name
DNS failures	Percentage of DNS requests performed in the DNS resolution test that failed
ICMP latency	Round trip time of five evenly spaced ICMP packets
ICMP packet loss	Percentage of packets lost in the ICMP latency test
UDP Latency under load	Average round trip time for a series of evenly spaced UDP packets sent during downstream/upstream sustained tests
Lightweight download speed	Downstream throughput in Megabits per second (Mbps) utilizing a burst of UDP datagrams
Lightweight upload speed	Upstream throughput in Megabits per second (Mbps) utilizing a burst of UDP datagrams

[22] Specific questions on test procedures may be addressed to team@samknows.com.

[23] Other tests may be run on the MBA panel; this list outlines the published tests in the report.

[24] An "A record" is the numeric IP address associated with a domain address such as https://www.fcc.gov.

3.5 - TEST DESCRIPTIONS

The following sub-sections detail the methodology used for the individual tests. As noted earlier, all tests only measure the performance of the part of the network between the Whitebox and the target (*i.e.*, a test node). In particular, the VoIP tests can only approximate the behavior of real applications and do not reflect the impact of specific consumer hardware, software, media codecs, bandwidth adjustment algorithms, Internet backbones and in-home networks.

Download Speed and Upload Speed

These tests measure the download and upload throughput by performing multiple simultaneous HTTP GET and HTTP POST requests to a target test node.

Binary, non-zero content—herein referred to as the payload—is hosted on a web server on the target test node. The test operates for a fixed duration of 10 seconds. It records the average throughput achieved during this 10 second period. The client attempts to download as much of the payload as possible for the duration of the test.

The test uses three concurrent TCP connections (and therefore three concurrent HTTP requests) to ensure that the line is saturated. Each connection used in the test counts the numbers of bytes transferred and is sampled periodically by a controlling thread. The sum of these counters (a value in bytes) divided by the time elapsed (in microseconds) and converted to Mbps is taken as the total throughput of the user's broadband service.

Factors such as TCP slow start and congestion are taken into account by repeatedly transferring small chunks (256 kilobytes, or kB) of the target payload before the real testing begins. This "warm-up" period is completed when three consecutive chunks are transferred at within 10 percent of the speed of one another. All three connections are required to have completed the warm-up period before the timed testing begins. The warm-up period is excluded from the measurement results.

Downloaded content is discarded as soon as it is received, and is not written to the file system. Uploaded content is generated and streamed on the fly from a random source.

The test is performed for both IPv4 and IPv6, where available, but only IPv4 results are reported.

Web Browsing

The test records the averaged time taken to sequentially download the HTML and referenced resources for the home page of each of the target websites, the number of bytes transferred, and the calculated rate per second. The primary measure for this test is the total time taken to download the HTML front page for each web site and all associated images, JavaScript, and stylesheet resources. This test does not measure against the centralized testing nodes; instead

it tests against actual websites, ensuring that the effects of content distribution networks and other performance enhancing factors can be taken into account.

Each Whitebox tests against the following nine websites:[25]

- **http://www.edition.cnn.com/**
- **http://www.bing.com/**
- **http://www.msn.com/**
- **http://www.bbc.com/**
- **http://www.apple.com/**
- **http://www.ebay.com/**
- **http://www.m.imdb.com/help/**
- **http://www.google.com/policies/**

The results include the time needed for DNS resolution. The test uses up to eight concurrent TCP connections to fetch resources from targets. The test pools TCP connections and utilizes persistent connections where the remote HTTP server supports them.

The client advertises the user agent as Microsoft Internet Explorer 10. Each website is tested in sequence and the results summed and reported across all sites.

UDP Latency and Packet Loss

These tests measure the round-trip time of small UDP packets between the Whitebox and a target test node.

Each packet consists of an 8-byte sequence number and an 8-byte timestamp. If a response packet is not received within three seconds of sending, it is treated as being lost. The test records the number of packets sent each hour, the average round trip time and the total number of packets lost. The test computes the summarized minimum, maximum, standard deviation and mean from the lowest 99 percent of results, effectively trimming the top (*i.e.*, slowest) 1 percent of outliers.

The test operates continuously in the background. It is configured to randomly distribute the sending of the requests over a fixed interval of one hour (using a Poisson distribution), reporting the summarized results once the interval has elapsed. Approximately two thousand packets are sent within a one-hour period, with fewer packets sent if the line is not idle.

This test is started when the Whitebox boots and runs permanently as a background test. The test is performed for both IPv4 and IPv6, where available, but only IPv4 results are reported.

[25] These websites were chosen based on a list by Alexa, http://www.alexa.com/, of the top twenty websites in October 2010.

Voice over IP

The Voice over IP (VoIP) test operates over UDP and utilizes bidirectional traffic, as is typical for voice calls.

The Whitebox handshakes with the server, and each initiates a UDP stream with the other. The test uses a 64 kbps stream with the same characteristics and properties (*i.e.*, packet sizes, delays, bitrate) as the G.711 codec. 160 byte packets are used. The test measures jitter, delay, and loss.

Jitter is calculated using the Packet Delay Variation (PDV) approach described in section 4.2 of RFC 5481. The 99th percentile is recorded and used in all calculations when deriving the PDV.

DNS Resolutions and DNS Failures

These tests measure the DNS resolution time of an A record query for the domains of the websites used in the web browsing test, and the percentage of DNS requests performed in the DNS resolution test that failed.

The DNS resolution test is targeted directly at the ISP's recursive resolvers. This circumvents any caching introduced by the panelist's home equipment (such as another gateway running in front of the Whitebox) and also accounts for panelists that might have configured the Whitebox (or upstream devices) to use non-ISP provided DNS servers. ISPs provide lists of their recursive DNS servers for the purposes of this study.

ICMP Latency and Packet Loss

These tests measure the round-trip time (RTT) of ICMP echo requests in microseconds from the Whitebox to a target test node. The client sends five ICMP echo requests of 56 bytes to the target test node, waiting up to three seconds for a response to each. Packets that are not received in response are treated as lost. The mean, minimum, maximum, and standard deviation of the successful results are recorded. The number of packets sent and received are recorded too.

Latency Under Load

The latency under load test operates for the duration of the 10-second downstream and upstream speed tests, with results for upstream and downstream recorded separately. While the speed tests are running, the latency under load test sends UDP datagrams to the target server and measures the round-trip time and number of packets lost. Packets are spaced five hundred milliseconds (ms) apart, and a three second timeout is used. The test records the mean, minimum, and maximum round trip times in microseconds. The number of lost UDP packets is also recorded.

This test represents an updated version of the methodology used in the initial August 2011 Report and aligns it with the methodology for the regular latency and packet loss metrics.

Traceroute

A traceroute client is used to send UDP probes to each hop in the path between client and destination. Three probes are sent to each hop. The round-trip times, the standard deviation of the round-trip times of the responses from each hop and the packet loss are recorded. The open source traceroute client "mtr" (https://github.com/traviscross/mtr) is used for carrying out the traceroute measurements.

Lightweight Capacity Test

This test measures the instantaneous capacity of the link using a small number of UDP packets. The test supports both downstream and upstream measurements, conducted independently.

In the downstream mode, the test client handshakes with the test server over TCP, requesting a fixed number of packets to be transmitted back to the client. The client specifies the transmission rate, number of packets and packet size in this handshake. The client records the arrival times of each of the resulting packets returns to it.

In the upstream mode, the client again handshakes with the test server, this time informing it of the characteristics of the stream it is about to transmit. The client then transmits the stream to the server, and the server locally records the arrival times of each packet. At the conclusion of this stream, the client asks the server for its summary of the arrival time of each packet.

With this resulting set of arrival times, the test client calculates the throughput achieved. This throughput may be divided into multiple windows, and an average taken across those, in order to smooth out buffering behavior.

This test uses approximately 99% less data than the TCP speed test and completes in a fraction of the time (100 milliseconds versus 10 seconds). The lightweight capacity test achieves results are within 1% deviation from the existing speed test results on fixed-line connections tested on average.

Table 8: Estimated Total Traffic Volume Generated by Test

The standard test schedule, below, was used across all ISPs, with the exception of Viasat. In 2017, Viasat opted to no longer provide panelists with an increased data allowance to offset the amount of data used by the measurements. This meant that the standard test schedule could no longer be used on Viasat, so a lighter weight test schedule was developed for them.

Standard Test Schedule

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
Web Browsing	9 popular US websites	Every 2 hours, 24x7	Est. 30 seconds	80 MB
Voice over IP	1 off-net test node	Hourly, 24x7	Fixed 10 seconds at 64k	1.8 MB
	1 on-net test node	Hourly, 24x7	Fixed 10 seconds at 64k	1.8 MB
Download Speed (Capacity – 8x parallel TCP connections)	1 off-net test node	Once 12 am - 6 am Once 6 am - 12 pm Once 12 pm - 6 pm Hourly thereafter	Fixed 10 seconds	107 MB at 10 Mbps
	1 on-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Once 6pm-8pm, Once 8pm-10pm, Once 10pm-12am	Fixed 10 seconds	70 MB at 10 Mbps
Download Speed (Single TCP connection)	1 off-net test node 1 on-net test node	Once in peak hours, once in off-peak hours	Fixed 10 seconds	46 MB at 10 Mbps
Upload Speed (Capacity – 8x parallel TCP connections on terrestrial, 3x on satellite)	1 off-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Hourly thereafter	Fixed 10 seconds	11 MB at 1 Mbps
	1 on-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Once 6pm-8pm, Once 8pm-10pm, Once 10pm-12am	Fixed 10 seconds	7 MB at 1 Mbps

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
Upload Speed (Single TCP connection)	1 off-net test node 1 on-net test node	Once in peak hours, once in off-peak hours	Fixed 10 seconds	6 MB at 1 Mbps
UDP Latency	2 off-net test nodes (Level3/MLab) 1 on-net test node	Hourly, 24x7 Hourly, 24x7	Permanent Permanent	5.8 MB 2.9 MB
UDP Packet Loss	2 off-net test node 1 on-net test nodes	Hourly, 24x7 Hourly, 24x7	Permanent Permanent	N/A (uses above) N/A (uses above)
Consumption	N/A	24x7	N/A	N/A
DNS Resolution	10 popular US websites	Hourly, 24x7	Est. 3 seconds	0.3 MB
ICMP Latency	1 off-net test node 1 on-net test node	Hourly, 24x7	Est. 5 seconds	0.3 MB
ICMP Packet loss	1 off-net test node 1 on-net test node	Hourly, 24x7	N/A (As IMCP latency)	N/A (uses above)
Traceroute	1 off-net test node 1 on-net test node	Three times a day, 24x7	N/A	N/A
Download Speed IPv6^^	1 off-net test node	Three times a day	Fixed 10 seconds	180 MB at 50 Mbps 72 MB at 20 Mbps 11 MB at 3 Mbps 5.4 MB at 1.5 Mbps
Upload Speed IPv6^^	1 off-net test node	Three times a day	Fixed 10 seconds	172 MB at 2 Mbps 3.6MB at 1 Mbps 1.8MB at 0.5 Mbps
UDP Latency / Loss IPv6^^	2 off-net test nodes (Level3/MLab)	Hourly, 24x7	Permanent	5.8 MB
Lightweight Capacity Test – Download (UDP)	1 off-net test node	Once 12am-6am,	Fixed 1000 packets	9MB

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
		Once 6am-12pm, Once 12pm-6pm, Hourly thereafter		
Lightweight capacity test – Upload (UDP)	1 off-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Hourly thereafter	Fixed 1000 packets	9MB

Lightweight test schedule (currently Viasat only)

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
Web Browsing	9 popular US websites	Once 8pm-10-pm	Est. 30 seconds	7MB
Download Speed (Capacity – 8x parallel TCP connections)	1 off-net test node	Once 8pm-10-pm	Fixed 10 seconds	30MB at 10Mbps
Upload Speed (Capacity – 8x parallel TCP connections on terrestrial, 3x on satellite)	1 off-net test node	Once 8pm-10-pm	Fixed 10 seconds	3MB at 1Mbps
UDP Latency	1 off-net test node	Hourly, 24x7	Permanent	1MB
UDP Latency	1 on-net test node	Hourly, 24x7	Permanent	1MB
UDP Packet loss	1 off-net test node	Hourly, 24x7	Permanent	N/A (uses above)
UDP Packet loss	1 on-net test node	Hourly, 24x7	Permanent	N/A (uses above)
Consumption	N/A	24x7	N/A	N/A
DNS Resolution	10 popular US websites	Hourly, 24x7	Est. 3 seconds	0.3MB
ICMP Latency	1 off-net test node 1 on-net test node	Hourly, 24x7	Est. 5 seconds	0.3MB
ICMP Packet Loss	1 off-net test node 1 on-net test node	Hourly, 24x7	N/A (As IMCP latency)	N/A (uses above)
Traceroute	1 off-net test node	Three times a day, 24x7	N/A	N/A

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
	1 on-net test node			
CDN Performance	Amazon, Apple, Microsoft, Google, Cloudflare, Akamai	Every 2 hours, 24x7	5 seconds	3MB
UDP Latency / Loss **IPv6^**	1 off-net test node	Hourly, 24x7	Permanent	1MB
Lightweight Capacity Test – Download (UDP)	1 off-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Hourly thereafter	Fixed 1000 packets	9MB
Lightweight Capacity Test – Upload (UDP)	1 off-net test node	Once 12am-6am, Once 6am-12pm, Once 12pm-6pm, Hourly thereafter	Fixed 1000 packets	9MB

**Download/upload daily volumes are estimates based upon likely line speeds. All tests will operate at maximum line rate so actual consumption may vary.

^Currently in beta testing.

^^Only carried out on broadband connections that support IPv6.

Tests to the off-net destinations use the nearest (in terms of latency) server from the Level3, M-Lab and StackPath list of test servers. The one exception is the latency and packet loss tests, which operate continuously to Level3, M-Lab and StackPath off-net servers. All tests are also performed to the closest on-net server, where available.

Consumption

This test was replaced by the new data usage test. A technical description for this test is outlined here: https://transition.fcc.gov/oet/mba/Data-Usage-Technical-Methodology_2018-08-24_Final-v1.3.pdf

Cross-Talk Testing and Threshold Manager Service

In addition to the tests described above, for 60 seconds prior to and during testing, a ”threshold manager” service on the Whitebox monitors the inbound and outbound traffic across the WAN interface to calculate if a panelist is actively using the Internet connection. The threshold for

traffic is set to 64 kbps downstream and 32 kbps upstream. Metrics are sampled and computed every 10 seconds. If either of these thresholds is exceeded, the test is delayed for a minute and the process repeated. If the connection is being actively used for an extended period of time, this pause and retry process continues for up to five times before the test is abandoned.

4 - DATA PROCESSING AND ANALYSIS OF TEST RESULTS

This section describes the background for the categorization of data gathered for the Tenth Report, and the methods employed to collect and analyze the test results.

4.1 - BACKGROUND

Time of Day

Most of the metrics reported in the Tenth Report draw on data gathered during the so-called peak usage period of 7:00 p.m. to 11:00 p.m. local time[26]. This time period is generally considered to experience the highest amount of Internet usage under normal circumstances.

ISP and Service Tier

A sufficient sample size is necessary for analysis and the ability to robustly compare the performance of specific ISP speed tiers. In order for a speed tier to be considered for the fixed line MBA Report, it must meet the following criteria:

(a) The speed tier must make up the top 80% of the ISP's subscriber base;
(b) There must be a minimum of 45 panelists that are recruited for that tier who have provided valid data for the tier within the validation period; and
(c) Each panelist must have a minimum of five days of valid data within the validation period.

The study achieved target sample sizes for the following download and upload speeds[27] (listed in alphabetical order by ISP):

Download Speeds:

CenturyLink: 1.5, 3, 7, 10, 12, 20, and 40 Mbps tiers;

[26] This period of time was agreed to by ISP participants in open meetings conducted at the beginning of the program.

[27] Due to the large number of different combinations of upload/download speed tiers supported by ISPs where, for example, a single download speed might be offered paired with multiple upload speeds or vice versa, upload and download test results were analyzed separately.

Charter: 60, 100, and 200 Mbps tiers;
Cincinnati Bell DSL: 5 Mbps tier;
Cincinnati Bell Fiber: 50, 250 and 500 Mbps tier;
Comcast: 60, 150, and 250 Mbps tiers;
Cox: 30, 100, 150, and 300 Mbps tiers;
Frontier DSL: 3, 6, and 12 Mbps tiers;
Frontier Fiber: 50, 75, 100, 150 and 200 Mbps tiers;
Mediacom: 60, 100 and 200 Mbps tiers;
Optimum: 100 and 200 Mbps tiers;
Verizon Fiber: 75, 100 and 1000 Mbps tiers;[28]
Windstream: 3, 10 and 25 Mbps tiers.

Upload Speeds:

CenturyLink: 0.768, 0.896, 2, and 5 Mbps tiers;
Charter: 10, and 20 Mbps tiers;
Cincinnati Bell DSL: 0.768 and 3 Mbps tiers;
Cincinnati Bell Fiber: 10, 100 and 125 Mbps tiers;
Comcast: 5 and 10 Mbps tiers;
Cox: 3, 10, and 30 Mbps tiers;
Frontier DSL: 0.768 Mbps tier;
Frontier Fiber: 50, 75, 100, 150 and 200 Mbps tiers;
Mediacom: 5, 10 and 20 Mbps tiers;
Optimum: 35 Mbps tier;
Verizon Fiber: 75 and 100 Mbps tiers;[29]
Windstream: 1 and 1.5 Mbps tiers.

A file containing averages for each metric from the validated September/October 2019 data can be found on FCC's Measuring Broadband America website.[30] Some charts and tables are divided into speed bands, to group together products with similar levels of advertised performance. The results within these bands are further broken out by ISP and service tier. Where an ISP does not offer a service tier within a specific band or a representative sample could not be formed for tier(s) in that band, the ISP will not appear in that speed band.

[28] Verizon's 1 Gbps tier was not included in the final report. 1Gbps tiers may be included in a separate/subsequent report focusing on faster speeds.
[29] Verizon's 1 Gbps tier was not included in the final report. *Id* at n. 28.

[30] See: https://www.fcc.gov/general/measuring-broadband-america-measuring-fixed-broadband.

Results from tests run on speed tiers of 1Gbps were not included in the Tenth Report. This was due to concerns from ISPs that the Whitebox 8.0 could not measure these speeds accurately. An investigation was conducted to establish if this was the case, or if speeds of 1Gbps could be reliably reported on.

Following investigation and testing with one of the ISPs which takes part in the program this conclusion was reached:

The network of the ISP concerned was quite “bursty” in nature, with servers on a 1Gbps network sometimes bursting to 3Gbps. This caused small amounts of packet loss which negatively affected overall speed test results. However once implementing new traffic shaping rules restricting traffic from the server to 1Gbps consistent high speeds were recorded by the Whitebox. The other solution to this specific problem was seen when using a very large number of parallel TCP connections. This investigation established that there is not an issue with the Whitebox 8.0 measuring speeds up to 1Gbps consistently.

4.2 - DATA COLLECTION AND ANALYSIS METHODOLOGY

Data Integrity

To ensure the integrity of the data collected, the following validity checks were developed:

1. *Change of ISP intra-month*: By checking the WHOIS results once a day for the user's IP address, we found units that changed ISP during the month. We only kept data for the ISP where the panelist was active the most.
2. *Change of service tier intra-month*: This validity check found units that changed service tier intra-month by comparing the average sustained throughput observed for the first three days in the reporting period against that for the final three days in the reporting period. If a unit was not online at the start or end of that period, we used the first or final three days when they were actually online. If this difference was over 50 percent, the downstream and upstream charts for this unit were individually reviewed. Where an obvious step change was observed (*e.g.*, from 1 Mbps to 3 Mbps), the data for the shorter period was flagged for removal.
3. *Removal of any failed or irrelevant tests*: This validity check removed any failed or irrelevant tests by removing measurements against any nodes other than the US-based off-net nodes. We also removed measurements using any off-net server that showed a failure rate of 10 percent or greater during a specific one-hour period, to avoid using any out-of-service test nodes.
4. *Removal of any problem Whiteboxes*: We removed measurements for any Whitebox that exhibited greater than or equal to 10 percent failures in a particular one-hour period. This removed periods when the Whitebox was unable to reach the Internet.

Legacy Equipment

In previous reports, we discussed the challenges ISPs face in improving network performance where equipment under the control of the subscriber limits the end-to-end performance achievable by the subscriber.[31] Simply, some consumer-controlled equipment may not be capable of operating fully at new, higher service tiers. Working in open collaboration with all service providers we developed a policy permitting changes in ISP panelists when their installed modems were not capable of meeting the delivered service speed that included several conditions on participating ISPs. First, proposed changes in consumer panelists would only be considered where an ISP was offering free upgrades for modems they owned and leased to the consumer. Second, each ISP needed to disclose its policy regarding the treatment of legacy modems and its efforts to inform consumers regarding the impact such modems may have on

[31] See pgs. 8-9, 2014 Report, pg. 8 of the 2013 Report, as well as endnote 14. http://www.fcc.gov/measuring-broadband-america/2012/july.

their service.

While the issue of DOCSIS 3 modems and network upgrades affect the cable industry today, we may see other cases in the future where customer premises equipment affects the achievable network performance.

In accordance with the above stated policy, 135 Whiteboxes connected to legacy modems were identified and removed from the final data set in order to ensure that the study would only include equipment that would be able to meet its advertised speed. The 95 excluded Whiteboxes were connected to Charter, Comcast, and Cox.

Collation of Results and Outlier Control

All measurement data were collated and stored for analysis purposes as monthly trimmed averages during three time intervals (24 hours, 7:00 p.m. to 11:00 p.m. local time Monday through Friday, 12:00 a.m. to 12:00 a.m. local time Saturday and Sunday). Only participants who provided a minimum of five days of valid measurements and had valid data in each of the three time intervals were included in the September / October 2019 test results. In addition, the top and bottom 1 percent of measurements were trimmed to control for outliers that may have been anomalous or otherwise not representative of actual broadband performance. All results were computed on the trimmed data.[32]

Data was only charted when results from at least 45 separate Whiteboxes was available for individual ISP download speed tiers. Service tiers of 50 or fewer Whiteboxes were noted for possible future panel augmentation.

The resulting final validated sample of data for September/October 2019 included in the MBA Tenth Report was collected from 2,931 participants.

Peak Hours Adjusted to Local Time

Peak hours were defined as weekdays (Mondays through Fridays) between 7:00 p.m. to 11:00 p.m. (inclusive) for the purposes of the study. All times were adjusted to the panelist's local time zone. Since some tests are performed only once every two hours on each Whitebox, the duration of the peak period had to be a multiple of two hours.

Congestion in the Home Not Measured

Download, upload, latency, and packet loss measurements were taken between the panelist's home gateway and the dedicated test nodes provided by M-Lab and Level 3. Web browsing measurements were taken between the panelist's home gateway and nine popular United States-hosted websites. Any congestion within the user's home network is, therefore, not measured by this study. The web browsing measurements are subject to possible congestion at the content provider's side, although the choice of eight popular websites configured to serve high traffic loads reduced that risk.

Traffic Shaping Not Studied

The effect of traffic shaping is not studied in the Tenth Report, although test results were subject to any bandwidth management policies put in place by ISPs. The effects of bandwidth management policies, which may be used by ISPs to maintain consumer traffic rates within

[32] These methods were reviewed with statistical experts by the participating ISPs.

advertised service tiers, may be most readily seen in those charts in the 2016 Report that show performance over 24-hour periods, where tested rates for some ISPs and service tiers flatten for periods at a time.

Analysis of PowerBoost and Other "Enhancing" Services

The use of transient speed enhancing services marketed under names such as "PowerBoost" on cable connections presented a technical challenge when measuring throughput. These services will deliver a far higher throughput for the earlier portion of a connection, with the duration varying by ISP, service tier, and potentially other factors. For example, a user with a contracted 6 Mbps service tier may receive 18 Mbps for the first 10 MB of a data transfer. Once the "burst window" is exceeded, throughput will return to the contracted rate, with the result that the burst speed will have no effect on very long sustained transfers.

Existing speed tests transfer a quantity of data and divide this quantity by the duration of the transfer to compute the transfer rate, typically expressed in Mbps. Without accounting for burst speed techniques, speed tests employing the mechanism described here will produce highly variable results depending on how much data they transfer or how long they are run. Burst speed techniques will have a dominant effect on short speed tests: a speed test running for two seconds on a connection employing burst speed techniques would likely record the burst speed rate, whereas a speed test running for two hours will reduce the effect of burst speed techniques to a negligible level.

The earlier speed test configuration employed in this study isolated the effects of transient performance enhancing burst speed techniques from the long-term sustained speed by running for a fixed 30 seconds and recording the average throughput at 5 second intervals. The throughput at the 0-5 second interval is referred to as the burst speed and the throughput at the 25-30 second interval is referred to as the actual speed. Testing was conducted prior to the start of trial to estimate the length of time during which the effects of burst speed techniques might be seen. Even though the precise parameters used for burst-speed techniques are not known, their effects were no longer observable in testing after 20 seconds of data transfer.

In the Sixth report we noted that the use of this technology by providers was on the decline. For the Seventh, Eighth, Ninth and Tenth reports, we no longer provide the results of burst-speed since these techniques are now rarely used. The speed test configuration has been altered to shorten the test duration to 10 seconds, as there is no need to run it for 30 seconds any more.

Consistency of Speed Measurements

In addition to reporting on the median speed of panelists, the MBA Report also provides a measure of the consistency of speed that panelists experience in each tier. For purposes of discussion we use the term "80/80 consistent speed" to refer to the minimum speed that was experienced by at least 80% of panelists for at least 80% of the time during the peak periods. The process used in defining this metric for a specific ISP tier is to take each panelist's set of download or upload speed data during the peak period across all the days of the validated measurement

period and arrange it in increasing order. The speed that corresponds to the 20th percentile represents the minimum speed that the panelist experienced at least 80% of the time. The 20 percentile values of all the panelists on a specific tier are then arranged in an increasing order. The speed that corresponds to the 20th percentile now represents the minimum speed that at least 80% of panelists experienced 80% of the time. This is the value reported as the 80/80 consistent speed for that ISP's tier. We also report on the 70/70 consistent speed for an ISP's tier, which is the minimum speed that at least 70% of the panelists experience at least 70% of the time. We typically report the 70/70 and the 80/80 consistent speeds as a percentage of the advertised speed.

When reporting on these values for an ISP, we weigh the 80/80 or 70/70 consistent speed results (as a percentage of the advertised speed) of each of the ISP's tier based on the number of subscribers to that tier; so as to get a weighted average across all the tiers for that ISP.

Latencies Attributable to Propagation Delay

The speeds at which signals can traverse networks are limited at a fundamental level by the speed of light. While the speed of light is not believed to be a significant limitation in the context of the other technical factors addressed by the testing methodology, a delay of approximately 5ms per 1000 km of distance traveled can be attributed solely to the speed of light (depending on the transmission medium). The geographic distribution and the testing methodology's selection of the nearest test servers are believed to minimize any significant effect. However, propagation delay is not explicitly accounted for in the results.

Limiting Factors

A total of 8,417,695,058 measurements were taken across 144,636,223 unique tests.

All scheduled tests were run, aside from when monitoring units detected concurrent use of bandwidth.

Schedules were adjusted when required for specific tests to avoid triggering data usage limits applied by some ISPs.

4.3 DATA PROCESSING OF RAW AND VALIDATED DATA

The data collected in this program are made available as open data for review and use by the public. Raw and processed data sets, mobile testing software, and the methodologies used to process and analyze data are freely and publicly available. Researchers and developers interested in working with measurement data in raw form will need skills in database management, SQL programming, and statistics, depending on the analysis. A developer FAQ for database configuration and data importing instructions for MySQL and PostgreSQL are available at https://www.fcc.gov/general/database-setup-and-importing-measuring-broadband-america-data-april-2012.

The process flow below describes how the raw collected data was processed for the production of the *Measuring Broadband America Report*. Researchers and developers interested in replicating or extending the results of the Report are encouraged to review the process below and supporting files that provide details.

Raw Data:	Raw data for the chosen period is collected from the measurement database. The ISPs and products that panelists were on are exported to a "unit profile" file, and those that changed during the period are flagged. 2020 Raw Data Links
Validated Data Cleansing:	Data is cleaned. This includes removing measurements when a user changed ISP or tier during the period. Anomalies and significant outliers are also removed at this point. A data cleansing document describes the process in detail. 2020 Data Cleansing Document Link
SQL Processing:	Per-unit results are generated for each metric. Time-of-day averages are computed and a trimmed median is calculated for each metric. The SQL scripts used here are contained in SQL processing scripts available with the release of each report. 2020 SQL Processing Links
Unit Profile:	This document identifies the various details of each test unit, including ISP, technology, service tier, and general location. Each unit represents one volunteer panelists. The unit ID's were randomly generated, which served to protect the anonymity of the volunteer panelists. 2020 Unit Profile link
Excluded Units:	A listing of units excluded from the analysis due to insufficient sample size for that particular ISP's speed tier. 2020 Excluded Units Link
Unit Census Block:	This step identifies the census block (for blocks containing more than 1,000 people) in which each unit running tests is located. Census block is from 2010 census and is in the FIPS code format. We have used block FIPS codes for blocks that contains more than 1,000 people. For blocks with fewer than 1,000 people we have aggregated to the next highest level, *i.e.*, tract, and used the Tract FIPS code, provided there are more than 1,000 people in the tract. In cases where there are less than 1,000 people in a tract we have aggregated to Regional level. 2020 Unit Census Block Link.
Excel Tables & Charts:	Summary data tables and charts in Excel are produced from the averages. These are used directly in the report. 2020 Statistical Averages Links

The raw data collected for each active metric is made available by month in tarred gzipped files. The files in the archive containing active metrics are described in table 9.

Table 9: Test to Data File Cross-Reference List

Test	Validated Data File Name
Download Speed	curr_httpgetmt.csv — IPv4 Tests curr_httpgetmt6.csv — IPv6 Tests
Upload Speed	curr_httppostmt.csv — IPv4 Tests curr_httppostmt6.csv — IPv6 Tests
Web Browsing	curr_webget.csv
UDP Latency	curr_udplatency.csv — IPv4 Tests curr_udplatency6.csv — IPv6 Tests
UDP Packet Loss	curr_udplatency.csv — IPv4 Tests curr_udplatency6.csv — IPv6 Tests
Voice over IP	curr_udpjitter.csv
DNS Resolution	curr_dns.csv
DNS Failures	curr_dns.csv
ICMP Latency	curr_ping.csv
ICMP Packet Loss	curr_ping.csv
Latency under Load	curr_dlping.csv – Downstream latency under load results curr_ulping.csv – Upstream latency under load results
Traceroute	curr_traceroute.csv

Table 10: Validated Data Files - Dictionary

The following Data Dictionary file describes the schema for each active metric test for row level results stored in the files described in table 9.[33] All dtime entries are in the UTC timezone. All durations are in microseconds unless otherwise noted. The location_id field should be ignored.

curr_dlping.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address

[33] This data dictionary is also available on the FCC Measuring Broadband America website, located with the other validated data files available for download.

rtt_avg	Average RTT
rtt_min	Minimum RTT
rtt_max	Maximum RTT
rtt_std	Standard deviation in measured RTT
successes	Number of successes
failiures	Number of failures
location_id	Internal key mapping to unit profile data
curr_dns.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
nameserver	Name server used to handle the DNS request
lookup_host	Hostname to be resolved
response_ip	Field currently unused
rtt	DNS resolution time
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Internal key mapping to unit profile data
curr_httpgetmt.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
address	The IP address of the server (resolved by the client's DNS)
fetch_time	Time the test ran for
bytes_total	Total bytes downloaded across all connections
bytes_sec	Running total of throughput, which is sum of speeds measured for each stream (in bytes/sec), from the start of the test to the current interval
bytes_sec_interval	Throughput at this specific interval (*e.g.*, Throughput between 25-30 seconds)
warmup_time	Time consumed for all the TCP streams to arrive at optimal window size
warmup_bytes	Bytes transferred for all the TCP streams during the warm-up phase

sequence	The interval that this row refers to (*e.g.*, in the US, sequence=0 implies result is for 0-5 seconds of the test)
threads	The number of concurrent TCP connections used in the test
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Internal key mapping to unit profile data
curr_httppostmt.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
address	The IP address of the server (resolved by the client's DNS)
fetch_time	Time the test ran for
bytes_total	Total bytes downloaded across all connections
bytes_sec	Running total of throughput, which is sum of speeds measured for each stream (in bytes/sec), from the start of the test to the current interval
bytes_sec_interval	Throughput at this specific interval (*e.g.*, throughput between 25-30 seconds)
warmup_time	Time consumed for all the TCP streams to arrive at optimal window size
warmup_bytes	Bytes transferred for all the TCP streams during the warm-up phase.
sequence	The interval that this row refers to (*e.g.*, in the US, sequence=0 implies result is for 0-5 seconds of the test)
threads	The number of concurrent TCP connections used in the test
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Internal key mapping to unit profile data
curr_ping.csv	ICMP based
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address

rtt_avg	Average RTT
rtt_min	Minimum RTT
rtt_max	Maximum RTT
rtt_std	Standard deviation in measured RTT
successes	Number of successes
failiures	Number of failures
location_id	Internal key mapping to unit profile data
curr_udpjitter.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
packet_size	Size of each UDP Datagram (bytes)
stream_rate	Rate at which the UDP stream is generated (bits/sec)
duration	Total duration of test
packets_up_sent	Number of packets sent in upstream (measured by client)
packets_down_sent	Number of packets sent in downstream (measured by server)
packets_up_recv	Number of packets received in upstream (measured by server)
packets_down_recv	Number of packets received in downstream (measured by client)
jitter_up	Upstream Jitter measured
jitter_down	Downstream Jitter measured
latency	99th percentile of round trip times for all packets
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Internal key mapping to unit profile data
curr_udplatency.csv	UDP based
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
rtt_avg	Average RTT
rtt_min	Minimum RTT

rtt_max	Maximum RTT
rtt_std	Standard deviation in measured RTT
successes	Number of successes (note: use failures/(successes + failures)) for packet loss)
failiures	Number of failures (packets lost)
location_id	Internal key mapping to unit profile data
curr_ulping.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	Target hostname or IP address
rtt_avg	Average RTT
rtt_min	Minimum RTT
rtt_max	Maximum RTT
rtt_std	Standard deviation in measured RTT
successes	Number of successes
failures	Number of failures
location_id	Internal key mapping to unit profile data
curr_webget.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
target	URL to fetch
address	IP address used to fetch content from initial URL
fetch_time	Sum of time consumed to download HTML content and then concurrently download all resources
bytes_total	Sum of HTML content size and all resources size (bytes)
bytes_sec	Average speed of downloading HTML content and then concurrently downloading all resources (bytes/sec)
objects	Number of resources (images, CSS, …) downloaded
threads	Maximum number of concurrent threads allowed
requests	Total number of HTTP requests made
connections	Total number of TCP connections established
reused_connections	Number of TCP connections re-used

lookups	Number of DNS lookups performed
request_total_time	Total duration of all requests summed together, if made sequentially
request_min_time	Shortest request duration
request_avg_time	Average request duration
request_max_time	Longest request duration
ttfb_total_time	Total duration of the time-to-first-byte summed together, if made sequentially
ttfb_min_time	Shortest time-to-first-byte duration
ttfb_avg_time	Average time-to-first-byte duration
ttfb_max_time	Longest time-to-first-byte duration
lookup_total_time	Total duration of all DNS lookups summed together, if made sequentially
lookup_min_time	Shortest DNS lookup duration
lookup_avg_time	Average DNS lookup duration
lookup_max_time	Longest DNS lookup duration
successes	Number of successes
failures	Number of failures
location_id	Internal key mapping to unit profile data
curr_netusage.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished
wan_rx_bytes	Total bytes received via the WAN interface on the unit (incl. Ethernet and IP headers)
wan_tx_bytes	Total bytes transmitted via the WAN interface on the unit (incl. Ethernet and IP headers)
sk_rx_bytes	Bytes received as a result of active performance measurements
sk_tx_bytes	Bytes transmitted as a result of active performance measurements
location_id	Internal key mapping to unit profile data

curr_lct_dl.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC

target	Target hostname
address	Target IP address
packets_received	Total number of packets received
packets_sent	Total number of packets sent
packet_size	Packet size
bytes_total	Total number of bytes
duration	Duration of the test in microseconds
bytes_sec	Throughput in bytes/sec
error_code	An internal error code from the test.
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Please ignore (this is an internal key mapping to unit profile data)

curr_lct_ul.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	Target hostname
address	Target IP address
packets_received	Total number of packets received
packets_sent	Total number of packets sent
packet_size	Packet size
bytes_total	Total number of bytes
duration	Duration of the test in microseconds
bytes_sec	Throughput in bytes/sec
error_code	An internal error code from the test.
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Please ignore (this is an internal key mapping to unit profile data)

5 - Reference Documents

5.1 - USER TERMS AND CONDITIONS

The following document was agreed to by each volunteer panelist who agreed to participate in the broadband measurement study:

End User License Agreement

PLEASE READ THESE TERMS AND CONDITIONS CAREFULLY. BY APPLYING TO BECOME A PARTICIPANT IN THE BROADBAND COMMUNITY PANEL AND/OR INSTALLING THE WHITEBOX, YOU ARE AGREEING TO THESE TERMS AND CONDITIONS.

YOUR ATTENTION IS DRAWN PARTICULARLY TO CONDITIONS 3.5 (PERTAINING TO YOUR CONSENT TO YOUR ISPS PROVIDING CERTAIN INFORMATION AND YOUR WAIVER OF CLAIMS), 6 (LIMITATIONS OF LIABILITY) AND 7 (DATA PROTECTION).

1. Interpretation

1.1. The following definitions and rules of interpretation apply to these terms & conditions.

Connection: the Participant's own broadband internet connection, provided by an Internet Service Provider ("ISP").

Connection Equipment: the Participant's broadband router or cable modem, used to provide the Participant's Connection.

Intellectual Property Rights: all patents, rights to inventions, utility models, copyright and related rights, trademarks, service marks, trade, business and domain names, rights in trade dress or get-up, rights in goodwill or to sue for passing off, unfair competition rights, rights in designs, rights in computer software, database right, moral rights, rights in confidential information (including know-how and trade secrets) and any other intellectual property rights, in each case whether registered or unregistered and including all applications for and renewals or extensions of such rights, and all similar or equivalent rights or forms of protection in any part of the world.

ISP: the company providing broadband internet connection to the Participant during the term of this Program.

Participant/You/Your: the person who volunteers to participate in the Program, under these terms and conditions. The Participant must be the named account holder on the Internet service account with the ISP.

Open Source Software: the software in the Whitebox device that is licensed under an open source license (including the GPL).

Participant's Equipment: any equipment, systems, cabling or facilities provided by the Participant and used directly or indirectly in support of the Services, excluding the Connection Equipment.

Parties: both the Participant and SamKnows.

Party: one of either the Participant or SamKnows.

Requirements: the requirements specified by SamKnows as part of the sign-up process that the Participant must fulfil in order to be selected to receive the Services.

SamKnows/We/Our: the organization providing the Services and conducting the Program, namely:

SamKnows Limited (Co. No. 6510477) of 25 Harley Street, London W1G 9BR

Services / Program: the performance and measurement of certain broadband and Internet services and research program (Broadband Community Panel), as sponsored by the Federal Communications Committee (FCC), in respect of measuring broadband Internet Connections.

Software: the software that has been installed and/or remotely uploaded onto the Whitebox, by SamKnows as updated by SamKnows, from time to time, but not including any Open Source Software.

Test Results: Information concerning the Participant's ISP service results.

Whitebox: the hardware supplied to the Participant by SamKnows with the Software.

1.2. Headings in these terms and conditions shall not affect their interpretation.

1.3. A person includes a natural person, corporate or unincorporated body (whether or not having separate legal personality).

1.4. The schedules form part of these terms and conditions.

1.5. A reference to writing or written includes faxes and e-mails.

1.6.Any obligation in these terms and conditions on a person not to do something includes, without limitation, an obligation not to agree, allow, permit or acquiesce in that thing being done.

2. SamKnows' Commitment to You

2.1 Subject to the Participant complying fully with these terms and conditions, SamKnows shall use reasonable care to:

(a) provide the Participant with the Measurement Services under these terms and conditions;

(b) supply the Participant with the Whitebox and instructions detailing how it should be connected to the Participant's Connection Equipment; and

(c) if requested, SamKnows will provide a pre-paid postage label for the Whitebox to be returned.

(d) comply with all applicable United States, European Union, and United Kingdom privacy laws and directives, and will access, collect, process and distribute the information according to the following principles:

Fairness: We will process data fairly and lawfully;

Specific purpose: We will access, collect, process, store and distribute data for the purposes and reasons specified in this agreement and not in ways incompatible with those purposes;

Restricted: We will restrict our data collection and use practices to those adequate and relevant, and not excessive in relation to the purposes for which we collect the information;

Accurate: We will work to ensure that the data we collect is accurate and up-to-date, working with Participant and his/her ISP;

Destroyed when obsolete: We will not maintain personal data longer than is necessary for the purposes for which we collect and process the information;

Security: We will collect and process the information associated with this trial with adequate security through technical and organizational measures to protect personal data against destruction or loss, alteration, unauthorized disclosure or access, in particular where the processing involves the transmission of data over a network.

2.2 In addition, SamKnows shall:

(a) provide Participant with access to a Program-specific customer services email address, which the Participant may use for questions and to give feedback and comments;

(b) provide Participant with a unique login and password in order to access to an online reporting system for access to Participant's broadband performance statistics.

(c) provide Participant with a monthly email with their specific data from the Program or notifying Participant that their individual data is ready for viewing;

(d) provide Participant with support and troubleshooting services in case of problems or issues with their Whitebox;

(e) notify Participant of the end of the FCC-sponsored Program and provide a mechanism for Participant to opt out of any further performance/measuring services and research before collecting any data after termination of the Program;

(f) use only data generated by SamKnows through the Whitebox, and not use any Participant data for measuring performance without Participant's prior written consent; and

(g) not monitor/track Participant's Internet activity without Participant's prior written consent.

2.3 While SamKnows will make all reasonable efforts to ensure that the Services cause no disruption to the performance of the Participant's broadband Connection, including only running tests when there is no concurrent network activity generated by users at the Participant's location. The Participant acknowledges that the Services may occasionally impact the performance of the Connection and agrees to hold SamKnows and their ISP harmless for any impact the Services may have on the performance of their Connection.

3. Participant's Obligations

3.1 The Participant is not required to pay any fee for the provision of the Services by SamKnows or to participate in the Program.

3.2 The Participant agrees to use reasonable endeavors to:

(a) connect the Whitebox to their Connection Equipment within 14 days of receiving it;

(b) not to unplug or disconnect the Whitebox unless (i) they will be absent from the property in which it is connected for more than 3 days and/or (ii) it is reasonably necessary for maintenance of the Participant's Equipment and the Participant agrees that they shall use reasonable endeavors to minimize the length of time the Whitebox is unplugged or disconnected;

(c) in no way reverse engineer, tamper with, dispose of or damage the Whitebox, or attempt to do so;

(d) notify SamKnows within 7 days in the event that they change their ISP or their Connection tier or package (for example, downgrading/upgrading to a different broadband package), to the email address provided by SamKnows;

(e) inform SamKnows of a change of postal or email address by email; within 7 days of the change, to the email address provided by SamKnows;

(f) agrees that the Whitebox may be upgraded to incorporate changes to the Software and/or additional tests at the discretion of SamKnows, whether by remote uploads or otherwise;

(g) on completion or termination of the Services, return the Whitebox to SamKnows by mail, if requested by SamKnows. SamKnows will provide a pre-paid postage label for the Whitebox to be returned;

(h) be an active part of the Program and as such will use all reasonable endeavors to complete the market research surveys received within a reasonable period of time;

(i) not publish data, give press or other interviews regarding the Program without the prior written permission of SamKnows; and

(k) contact SamKnows directly, and not your ISP, in the event of any issues or problems with the Whitebox, by using the email address provided by SamKnows.

3.3 You will not give the Whitebox or the Software to any third party, including (without limitation) to any ISP. You may give the Open Source Software to any person in accordance with the terms of the relevant open source licence.

3.4 The Participant acknowledges that he/she is not an employee or agent of, or relative of, an employee or agent of an ISP or any affiliate of any ISP. In the event that they become one, they will inform SamKnows, who at its complete discretion may ask for the immediate return of the Whitebox.

3.5 THE PARTICIPANT'S ATTENTION IS PARTICULARLY DRAWN TO THIS CONDITION. The Participant expressly consents to having their ISP provide to SamKnows and the Federal Communications (FCC) information about the Participant's broadband service, for example: service address, speed tier, local loop length (for DSL customers), equipment identifiers and other similar information, and hereby waives any claim that its ISPs disclosure of such information to SamKnows or the FCC constitutes a violation of any right or any other right or privilege that the Participant may have under any federal, state or local statute, law, ordinance, court order, administrative rule, order or regulation, or other applicable law, including, without limitation, under 47 U.S.C**.** §§ 222 and 631 (each a "Privacy Law"). If notwithstanding Participant's consent under this Section 3.5, Participant, the FCC or any other party brings any claim or action against any ISP under a Privacy Law, upon the applicable ISPs request SamKnows promptly shall cease collecting data from such Participant and remove from its records all data collected with respect to such Participant prior to the date of such request, and shall not provide such data in any form to the FCC. The Participant further consents to transmission of information from this Program Internationally, including the information provided by the Participant's ISP, specifically the transfer of this information to SamKnows in the United Kingdom, SamKnows' processing of it there and return to the United States.

4. Intellectual Property Rights

4.1 All Intellectual Property Rights relating to the Whitebox are the property of its manufacturer. The Participant shall use the Whitebox only to allow SamKnows to provide the Services.

4.2 As between SamKnows and the Participant, SamKnows owns all Intellectual Property Rights in the Software. The Participant shall not translate, copy, adapt, vary or alter the Software. The Participant shall use the Software only for the purposes of SamKnows providing the Services and shall not disclose or otherwise use the Software.

4.3 Participation in the Broadband Community Panel gives the participant no Intellectual Property Rights in the Test Results. Ownership of all such rights is governed by Federal Acquisition Regulation Section 52.227-17, which has been incorporated by reference in the relevant contract between SamKnows and the FCC. The Participant hereby acknowledges and agrees that SamKnows may make such use of the Test Results as is required for the Program.

4.4 Certain core testing technology and aspects of the architectures, products and services are developed and maintained directly by SamKnows. SamKnows also implements various technical features of the measurement services using particular technical components from a variety of vendor partners including: NetGear, Measurement Lab, TP-Link.

5. SamKnows' Property

The Whitebox and Software will remain the property of SamKnows. SamKnows may at any time ask the Participant to return the Whitebox, which they must do within 28 days of such a request being sent. Once SamKnows has safely received the Whitebox, SamKnows will reimburse the Participant's reasonable postage costs for doing so.

6. Limitations of Liability - THE PARTICIPANT'S ATTENTION IS PARTICULARLY DRAWN TO THIS CONDITION

6.1 This condition 6 sets out the entire financial liability of SamKnows (including any liability for the acts or omissions of its employees, agents, consultants, and subcontractors) to the Participant, including and without limitation, in respect of:

(a) any use made by the Participant of the Services, the Whitebox and the Software or any part of them; and

(b) any representation, statement or tortious act or omission (including negligence) arising under or in connection with these terms and conditions.

6.2 All implied warranties, conditions and other terms implied by statute or other law are, to the fullest extent permitted by law, waived and excluded from these terms and conditions.

6.3 Notwithstanding the foregoing, nothing in these terms and conditions limits or excludes the liability of SamKnows:

(a) for death or personal injury resulting from its negligence or willful misconduct;

(b) for any damage or liability incurred by the Participant as a result of fraud or fraudulent misrepresentation by SamKnows;

(c) for any violations of U.S. consumer protection laws;

(d) in relation to any other liabilities which may not be excluded or limited by applicable law.

6.4 Subject to condition 6.2 and condition 6.3, SamKnows' total liability in contract, tort (including negligence or breach of statutory duty), misrepresentation, restitution or otherwise arising in connection with the performance, or contemplated performance, of these terms and conditions shall be limited to $100.

6.5 In the event of any defect or modification in the Whitebox, the Participant's sole remedy shall be the repair or replacement of the Whitebox at SamKnows' reasonable cost, provided that the defective Whitebox is safely returned to SamKnows, in which case SamKnows shall pay the Participant's reasonable postage costs.

6.6 The Participant acknowledges and agrees that these limitations of liability are reasonable in all the circumstances, particularly given that no fee is being charged by SamKnows for the Services or participation in the Program.

6.7 It is the Participant's responsibility to pay all service and other charges owed to its ISP in a timely manner and to comply with all other ISP applicable terms. The Participant shall ensure that their broadband traffic, including the data pushed by SamKnows during the Program, does not exceed the data allowance included in the Participant's broadband package. If usage allowances are accidentally exceeded and the Participant is billed additional charges from the ISP as a result, SamKnows is not under any obligation to cover these charges although it may choose to do so at its discretion.

7. Data protection - the participation's attention is particularly drawn to this condition.

7.1 The Participant acknowledges and agrees that his/her personal data, such as service tier, address and line performance, will be processed by SamKnows in connection with the program.

7.2 Except as required by law or regulation, SamKnows will not provide the Participant's personal data to any third party without obtaining Participant's prior consent. However, for the avoidance of doubt, the Participant acknowledges and agrees that subject to the privacy polices discussed below, the specific technical characteristics of tests and other technical features associated with the Internet Protocol environment of architecture, including the client's IP address, may be shared with third parties as necessary to conduct the Program and all aggregate statistical data produced as a result of the Services (including the Test Results) may be provided to third parties.

7.3 You acknowledge and agree that SamKnows may share some of Your information with Your ISP, and request information about You from Your ISP so that they may confirm Your service tiers and other information relevant to the Program. Accordingly You hereby expressly waive claim that any disclosure by Your ISP to SamKnows constitutes a violation of any right or privilege that you may have under any law, wherever it might apply.

8. Term and Termination

8.1 This Agreement shall continue until terminated in accordance with this clause.

8.2 Each party may terminate the Services immediately by written notice to the other party at any time. Notice of termination may be given by email. Notices sent by email shall be deemed to be served on the day of transmission if transmitted before 5.00 pm Eastern Time on a working day, but otherwise on the next following working day.

8.3 On termination of the Services for any reason:

(a) SamKnows shall have no further obligation to provide the Services; and

(b) the Participant shall safely return the Whitebox to SamKnows, if requested by SamKnows, in which case SamKnows shall pay the Participant's reasonable postage costs.

8.4 Notwithstanding termination of the Services and/or these terms and conditions, clauses 1, 3.3 and 4 to 14 (inclusive) shall continue to apply.

9. Severance

If any provision of these terms and conditions, or part of any provision, is found by any court or other authority of competent jurisdiction to be invalid, illegal or unenforceable, that provision or part-provision shall, to the extent required, be deemed not to form part of these terms and conditions, and the validity and enforceability of the other provisions these terms and conditions shall not be affected.

10. Entire agreement

10.1 These terms and conditions constitute the whole agreement between the parties and replace and supersede any previous agreements or undertakings between the parties.

10.2 Each party acknowledges that, in entering into these terms and conditions, it has not relied on, and shall have no right or remedy in respect of, any statement, representation, assurance or warranty.

11. Assignment

11.1 The Participant shall not, without the prior written consent of SamKnows, assign, transfer, charge, mortgage, subcontract all or any of its rights or obligations under these terms and conditions.

11.2 Each party that has rights under these terms and conditions acknowledges that they are acting on their own behalf and not for the benefit of another person.

12. No Partnership or Agency

Nothing in these terms and conditions is intended to, or shall be deemed to, constitute a partnership or joint venture of any kind between any of the parties, nor make any party the agent of another party for any purpose. No party shall have authority to act as agent for, or to bind, the other party in any way.

13. Rights of third parties

Except for the rights and protections conferred on ISPs under these Terms and Conditions which they may defend, a person who is not a party to these terms and conditions shall not have any rights under or in connection with these Terms and Conditions.

14. Privacy and Paperwork Reduction Acts

14.1 For the avoidance of doubt, the release of IP protocol addresses of client's Whiteboxes are not PII for the purposes of this program and the client expressly consents to the release of IP address and other technical IP protocol characteristics that may be gathered within the context of the testing architecture. SamKnows, on behalf of the FCC, is collecting and storing broadband performance information, including various personally identifiable information (PII) such as the street addresses, email addresses, sum of data transferred, and broadband performance information, from those individuals who are participating voluntarily in this test. PII not necessary to conduct this study will not be collected. Certain information provided by or collected from you will be confirmed with a third party, including your ISP, to ensure a representative study and otherwise shared with third parties as necessary to conduct the program. SamKnows will not release, disclose to the public, or share any PII with any outside entities, including the FCC, except as is consistent with the SamKnows privacy policy or these Terms and Conditions. See https://www.measuringbroadbandamerica.com/privacy/. The broadband performance

information that is made available to the public and the FCC, will be in an aggregated form and with all PII removed. For more information, see the Privacy Act of 1974, as amended (5 U.S.C. § 552a), and the SamKnows privacy policy.

14.2 The FCC is soliciting and collecting this information authorized by OMB Control No. 3060-1139 in accordance with the requirements and authority of the Paperwork Reduction Act, Pub. L. No. 96-511, 94 Stat. 2812 (Dec. 11, 1980); the Broadband Data Improvement Act of 2008, Pub. L. No. 110-385, Stat 4096 § 103(c)(1); American Reinvestment and Recovery Act of 2009 (ARRA), Pub. L. No. 111-5, 123 Stat 115 (2009); and Section 154(i) of the Communications Act of 1934, as amended.

14.3 *Paperwork Reduction Act of 1995 Notice*. We have estimated that each Participant of this study will assume a one hour time burden over the course of the Program. Our estimate includes the time to sign-up online, connect the Whitebox in the home, and periodic validation of the hardware. If you have any comments on this estimate, or on how we can improve the collection and reduce the burden it causes you, please write the Federal Communications Commission, Office of Managing Director, AMD-PERM, Washington, DC 20554, Paperwork Reduction Act Project (3060-1139). We will also accept your comments via the Internet if you send an e-mail to PRA@fcc.gov. Please DO NOT SEND COMPLETED APPLICATION FORMS TO THIS ADDRESS. You are not required to respond to a collection of information sponsored by the Federal government, and the government may not conduct or sponsor this collection, unless it displays a currently valid OMB control number and provides you with this notice. This collection has been assigned an OMB control number of 3060-1139. THIS NOTICE IS REQUIRED BY THE PAPERWORK REDUCTION ACT OF 1995, PUBLIC LAW 104-13, OCTOBER 1, 1995, 44 U.S.C. SECTION 3507. This notice may also be found at https://www.measuringbroadbandamerica.com/paperwork-reduction-act/.

15. Jurisdiction

These terms and conditions shall be governed by the laws of the state of New York.

SCHEDULE

THE SERVICES

Subject to the Participant complying with its obligations under these terms and conditions, SamKnows shall use reasonable endeavors to test the Connection so that the following information is recorded:

1. Web browsing
2. Video streaming
3. Voice over IP
4. Download speed
5. Upload speed
6. UDP latency
7. UDP packet loss
8. Consumption

9. Availability
10. DNS resolution
11. ICMP latency
12. ICMP packet loss

In performing these tests, the Whitebox will require a variable download capacity and upload capacity per month, which will be available to the Participant in motion 2.3. The Participant acknowledges that this may impact on the performance of the Connection.

1. SamKnows will perform tests on the Participant's Connection by using SamKnows' own data and will not monitor the Participant's content or internet activity. The purpose of this study is to measure the Connection and compare this data with other consumers to create a representative index of US broadband performance.

5.2 - CODE OF CONDUCT

The following Code of Conduct, available at http://data.fcc.gov/download/measuring-broadband-america/2020/Code-of-Conduct-fixed.pdf*, was signed by ISPs and other entities participating in the study:*

FCC MEASURING BROADBAND AMERICA PROGRAM

FIXED TESTING AND MEASUREMENT

STAKEHOLDERS CODE OF CONDUCT

WHEREAS the Federal Communications Commission of the United States of America (FCC) is conducting a Broadband Testing and Measurement Program, with support from its contractor SamKnows, the purpose of which is to establish a technical platform for the Measuring Broadband America Program Fixed Broadband Testing and Measurement and further to use that platform to collect data;

WHEREAS volunteer panelists have been recruited, and in so doing have agreed to provide broadband performance information measured on their Whiteboxes to support the collection of broadband performance data; and steps have been taken to protect the privacy of panelists to the program's effort to measure broadband performance. WE, THE UNDERSIGNED, as participants and stakeholders in that Fixed Broadband Testing and Measurement, do hereby agree to be bound by and conduct ourselves in accordance with the following principles and shall:

1. At all times act in good faith;
2. Not act, nor fail to act, if the intended consequence of such act or omission is inconsistent with the privacy policies of the program;
3. Not act, nor fail to act, if the intended consequence of such act or omission is to enhance, degrade, or tamper with the results of any test for any individual panelist or broadband provider, except that:

3.1. It shall not be a violation of this principle for broadband providers to:

3.1.1. Operate and manage their business, including modifying or improving services delivered to any class of subscribers that may or may not include panelists among them, provided that such actions are consistent with normal business practices, and

3.1.2. Address service issues for individual panelists at the request of the panelist or based on information not derived from the trial;

3.2. It shall not be a violation of this principle for academic and research purposes to simulate or observe tests and components of the testing architecture, provided that no impact to MBA data or the Internet Service of the subscriber volunteer panelist occurs; and

4. Not publish any data generated by the tests, nor make any public statement based on such data, until such time as the FCC releases data, or except where expressly permitted by the FCC; and

5. Not publish or make use of any test data or testing infrastructure in a manner that would significantly reduce the anonymity of collected data, compromise panelists privacy, or compromise the MBA privacy policy governing collection and analysis of data except that:
   5.1. It shall not be a violation of this principle for stakeholder signatories under the direction of the FCC to:
      5.1.1. Make use of test data or testing infrastructure to support the writing of FCC fixed Measuring Broadband America Reports;
      5.1.2. Make use of test data or testing infrastructure to support various aspects of the testing and architecture for the program including to facilitate data processing or analysis;
      5.1.3. Make use of test data or testing infrastructure to support the analysis of collected data or testing infrastructure for privacy risks or concerns, and plan for future measurement efforts;

6. Ensure that their employees, agents, and representatives, as appropriate, act in accordance with this Code of Conduct.

Signatories: _____

Printed: _____

Date: _____

5.3 - TEST NODE BRIEFING

Test Node Briefing

DOCUMENT REFERENCE:

SQ302-002-EN

TEST NODE BRIEFING

Technical information relating to

the SamKnows test nodes

August 2013

Important Notice

1 - SamKnows Test Nodes

In order to gauge an Internet Service Provider's broadband performance at a User's access point, the SamKnows Whiteboxes need to measure the service performance (*e.g.*, upload/download speeds, latency, *etc.*) from the Whitebox to a specific test node. SamKnows supports a number of "test nodes" for this purpose.

The test nodes run special software designed specifically for measuring the network performance when communicating with the Whiteboxes.

It is critical that these test nodes be deployed near to the customer (and their Whitebox). The further the test node is from the customer, the higher the latency and the greater the possibility that third-party networks may need to be traversed, making it difficult to isolate the individual ISP's performance. This is why SamKnows operates so many test nodes all around the world—locality to the customer is critical.

1.1 Test node definition

When referring to "test nodes," we are specifically referring to either the dedicated servers that are under SamKnows' control, or the virtual machines that may be provided to us. In the case of virtual machines provided by Measurement-Lab, Level3, Stackpath and others, the host operating system is under the control of and maintained by these entities and not by SamKnows.

1.2 Test node selection

The SamKnows Whiteboxes select the nearest node by running round-trip latency checks to all test nodes before measurement begins. Note that when we use the term "nearest" we are referring to the test node nearest to the Whitebox from the point of view of network delay, which may not necessarily always be the one nearest geographically.

Alternatively, it is possible to override test node selection based on latency and implement a static configuration so that the Whitebox will only test against the test node chosen by the Administrator. This is so that the Administrator can choose to test any particular test node that is of interest to the specific project and also to maintain configuration consistency. Similarly, test node selection may be done on a scheduled basis, alternating between servers, to collect test data from multiple test nodes for comparison purposes.

1.3 Test node positioning—on-net versus off-net

It is important that measurements collected by the test architecture support the comparison of ISP performance in an unbiased manner. Measurements taken from using the standardized set of "off-net" measurement test nodes (off-net here refers to a test node located outside a specific ISP's network) ensure that the performance of all ISPs can be measured under the same conditions and would avoid artificially biasing results for any one ISP over another. Test nodes located on a particular ISP's network ("on-net" test nodes), might introduce bias with respect to the ISP's own network performance. Thus data to be used to compare ISP performance are collected using "off-net" test nodes, because they reside outside the ISP network.

However, it is also very useful to have test nodes inside the ISP network ("on-net" test nodes). This allows us to:

- Determine what degradation in performance occurs when traffic leaves the ISP network; and
- Check that the off-net test nodes are performing properly (and vice versa).
- By having both on-net and off-net measurement data for each Whitebox, we can have a great deal of confidence in the quality of the data.

2.3 Data that is stored on test nodes

No measurement data collected by SamKnows is stored on test nodes.[34] The test nodes provide a "dumb" endpoint for the Whiteboxes to test against. All measurement performance results are recorded by the Whiteboxes, which are then transmitted from the Whitebox to data collection servers managed by SamKnows..

2 - Test Node Hosting and Locations

SamKnows test nodes reside in major peering locations around the world. Test nodes are carefully sited to ensure optimal connectivity on a market-by-market basis. SamKnows' test

[34] Note that Measurement-Lab runs sidestream measurements for all TCP connections against their test nodes and publishes these data in accordance with their data embargo policy.

infrastructure utilizes nodes made available by Level3, Measurement-Lab, Stackpath and various network operators, as well as under contract with select hosting providers.

2.1 Global Test Nodes

Level3 has provided SamKnows with 11 test nodes to use for the FCC's Measuring Broadband America Program. These test nodes are virtual servers meeting SamKnows specifications. Similarly, Measurement-Lab has also provided SamKnows with test nodes in various cities and countries for use with the Program's fixed measurement efforts. Measurement-Lab provides location hosting for at least three test nodes per site. SamKnows has also contracted with StackPath, a major CDN, to host virtual servers at its 10 US locations. Each location has one node with up to 100Gbps capacity.

Furthermore, SamKnows maintains its own test nodes, which are separate from the test nodes provided by Measurement-Lab and Level3 and Stackpath.

Table 1 below shows the locations of the SamKnows test node architecture supporting the Measuring Broadband America Program.[35] All of these listed test nodes reside outside individual ISP networks and therefore are designated as off-net test nodes. Note, that in many locations there are multiple test nodes installed which may be connected to different providers.

Location	SamKnows	Level3	Measurement-Lab	Stackpath
Atlanta, Georgia			✓	✓
Chicago, Illinois		✓	✓	✓
Dallas, Texas		✓	✓	✓
Los Angeles, California	✓	✓	✓	✓
Miami, Florida			✓	✓
Mountain View, California			✓	

[35] In addition to the test nodes used to support the Measuring Broadband America Program, SamKnows utilizes a diverse fleet of nodes in locations around the globe for other international programs.

New York City, New York	✓	✓	✓	✓
San Jose, California		✓		✓
Seattle, Washington			✓	✓
Washington D.C	✓	✓		
Washington, Virginia			✓	✓
Denver, Colorado			✓	✓

Table 1: Test Node Locations

SamKnows also has access to many test nodes donated by ISPs around the world. These particular test nodes reside within individual ISP networks and are therefore considered on-net test nodes.

ISPs have the advantage of measuring to both on-net and off-net test nodes, which allows them to segment end-to-end network performance and determine the performance of their own network versus third party networks. For example, an ISP can see what impact third party networks have on their end-users Quality of Experience ('QoE') by placing test nodes within their own network and at major National and International peering locations.

Diagram 1 below shows this set-up.

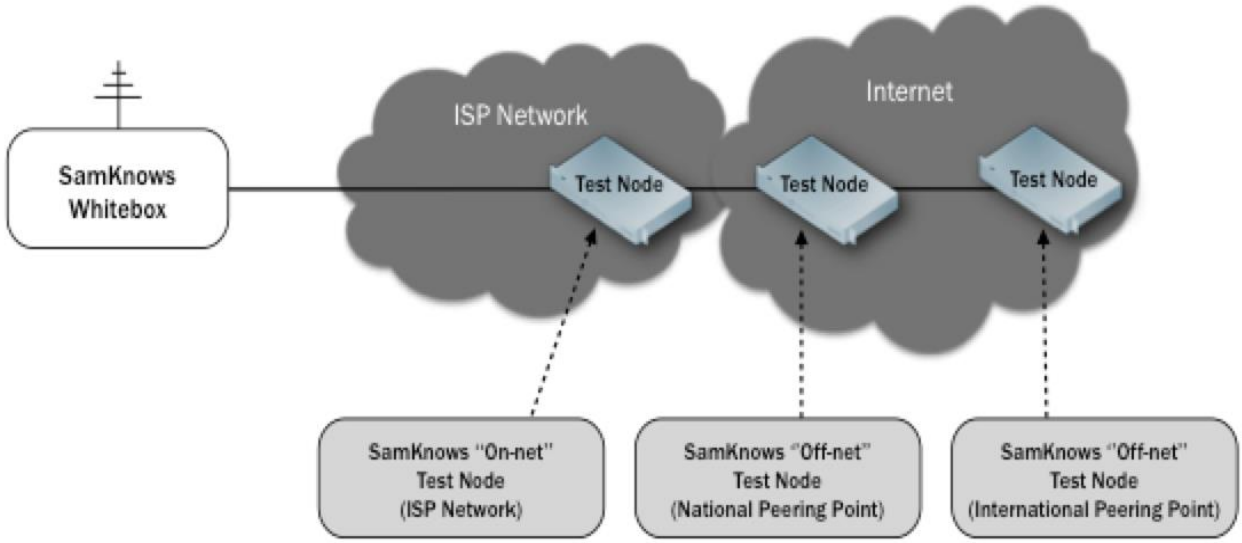

Diagram 1: On-net and Off-net Testing

Both the on-net and off-net test nodes are monitored by SamKnows as part of the global test node fleet. Test node management is explained in more detail within the next section of this document.

3 - Test Node Management

SamKnows test node infrastructure is a critical element of the SamKnows global measurement platform and includes extensive monitoring in place. SamKnows uses a management tool to control and configure the test nodes, while the platform is closely scrutinized using the Nagios monitoring application. System alerts are also in place to ensure the test node infrastructure is always available and operating well within expected threshold bounds.

The SamKnows Operations team continuously checks all test nodes to monitor capacity and overall health. Also included is data analysis to safeguard data accuracy and integrity. This level of oversight not only helps to maintain a healthy, robust platform but also allows us to spot and flag actual network issues and events as they happen. Diagnostic information also supports the Program managers' decision-making process for managing the impact of data accuracy and integrity incidents. This monitoring and administration is fully separate from any monitoring and administration of operating systems and platforms that may be necessary by hosting entities with which SamKnows may be engaged.

3.1 Seamless Test Node Management

SamKnows controls its network of test nodes via a popular open-source management tool called Puppet (**https://puppetlabs.com**). Puppet allows the SamKnows Operations team to easily

manage hundreds of test nodes and ensure that each group of test nodes is configured properly as per each project requirement. Coded in Python, Puppet uses a low-overhead agent installed on each test node that regularly communicates with the controlling SamKnows server to check for updates and ensure the integrity of the configuration.

This method of managing our test nodes allows us to deal with the large number of test nodes without affecting the user's performance in any way. We are also able to quickly and safely make changes to large parts of our test node fleet while ensuring that only the relevant test nodes are updated. This also allows us to keep a record of changes and rapidly troubleshoot any potential problems.

3.2 Proactive Test Node Monitoring

While Puppet handles the configuration and management of the test nodes, Nagios (the most popular online monitoring application) is used by SamKnows to monitor the test nodes. Each test node is configured to send Nagios regular status updates on core metrics such as CPU usage, disk space, free memory, and SamKnows-specific applications. Nagios will also perform active checks of each test nodes where possible, providing us with connectivity information—both via "ping" and connections to any webserver that may be running on the target host.

4 - Test Node Specification and Connectivity

SamKnows maintains a standard specification for all test nodes to ensure consistency and accuracy across the fleet.

4.1 SamKnows test node specifications

All dedicated test nodes must meet the following minimum specifications:

- CPU: Dual core Xeon (2 GHz+)
- RAM: 4 GB
- Disk: 80 GB
- Operating System: CentOS/RHEL 6.x
- Connectivity: Gigabit Ethernet connectivity, with gigabit upstream link.

4.2 Level3 test node specifications

All test nodes provided by level3 meet the following minimum specifications:

- CPU: 2.2 GHz Dual Core
- RAM: 4GB
- Disk: 10 GB

- Operating System: CentOS 6 (64bit)
- Connectivity: 4x1 Gigabit Ethernet (LAG protocol)

4.3 Measurement-Lab Test Node Specifications

All test nodes provided by Measurement-Lab meet the following minimum specifications:

- CPU: 2 GHz 8-core CPU
- RAM: 8 GB
- Disk: 2x100 GB
- OS: CentOS 6.4
- Connectivity: some locations 1 Gbps, some locations 10 Gbps

4.4 Stackpath test node notifications

- CPU Dual Core Xeon (2 GHz+)
- RAM: 8 GB
- Disk: 25 GB root disk
- OS: CentOS 7
- Connectivity: 10 Gbps

4.5 Test Node Connectivity

Measurement test nodes must be connected to a Tier-1 or equivalently neutral peering point. Each test node must be able to sustain 1 Gbps throughput.

At minimum, one publicly routable IPv4 address must be provisioned per-test node. The test node must not be presented with a NAT'd address. It is highly preferable for any new test nodes to also be provisioned with an IPv6 address at installation time.

It is preferred that the test nodes do not sit behind a firewall. If a firewall is used, then care must be taken to ensure that it can sustain the throughput required above.

4.6 Test Node Security

Each of the SamKnows test nodes is firewalled using the IPTables linux firewall. We close any ports that are not required, restrict remote administration to SSH only, and ensure access is only granted from a limited number of specified IP addresses. Only ports that require access from the outside world—for example TCP Port 80 on a webserver—would have that port fully open.

SamKnows regularly checks its rulesets to ensure that there are no outdated rules and that the access restriction is up to date.

SamKnows accounts on each test node are restricted to the systems administration team by default. When required for further work, an authorized SamKnows employee will have an account added.

5 - Test Node Provisioning

SamKnows also has a policy of accepting test nodes provided by network operators providing that

- The test node meets the specifications outlined earlier
- Minimum of 1 Gbps upstream is provided and downstream connectivity to national peering locations

Please note that donated test nodes may also be subject to additional local requirements.

5.1 Installation and Qualification

ISPs are requested to complete an information form for each test node they wish to provision. This will be used by SamKnows to configure the test node on the management system.

SamKnows will then provide an installation script and an associated installation guide. This will require minimal effort from the ISPs involved and will take a very similar form to the package used on existing test nodes.

Once the ISP has completed installation, SamKnows will verify the test node meets performance requirements by running server-to-server tests from known-good servers. These server-to-server measurements will be periodically repeated to verify performance levels.

5.2 Test Node Access and Maintenance

ISPs donating test nodes are free to maintain and monitor the test nodes using their existing toolsets, providing that these do not interfere with the SamKnows measurement applications or system monitoring tools. ISPs must not run resource intensive processes on the test nodes (*e.g.*, packet captures), as this may affect measurements.

ISPs donating test nodes must ensure that these test nodes are only accessed by maintenance staff when absolutely necessary.

SamKnows requests SSH access to the test nodes, with sudo abilities. sudo is a system administration tool that allows elevated privileges in a controlled granular manner. This has greatly helped diagnosis of performance issues with ISP-provided test nodes historically and would enable SamKnows to be far more responsive in investigating issues.

[DOCUMENT ENDS]