


2015

Technical Appendix

Measuring Broadband America

Fixed Broadband

**A Report on Consumer Fixed Broadband Performance
in the U.S.**



**FCC's Office of Engineering and Technology and Consumer
and Governmental Affairs Bureau**

Table of Contents

1 - INTRODUCTION/SUMMARY	4
2 - PANEL CONSTRUCTION	5
2.1 - USE OF AN ALL VOLUNTEER PANEL.....	5
2.2 - SAMPLE SIZE AND VOLUNTEER SELECTION	6
2.3 - PANELIST RECRUITMENT PROTOCOL	11
2.4 - VALIDATION OF VOLUNTEERS' SERVICE TIER.....	13
2.5 - PROTECTION OF VOLUNTEERS' PRIVACY	15
3 - BROADBAND PERFORMANCE TESTING METHODOLOGY	16
3.1 - SELECTION OF HARDWARE APPROACH.....	16
3.2 - DESIGN OBJECTIVES AND TECHNICAL APPROACH	17
3.3 - TESTING ARCHITECTURE.....	20
Testing Architecture Overview.....	20
Approach to Testing and Measurement	21
Home Deployment of the NETGEAR Based Whitebox	22
Home Deployment of the TP-Link Based Whitebox.....	22
Test Nodes (Off-Net and On-Net)	23
Test Node Breakout.....	24
Test Node Selection.....	26
3.4 - SAMKNOWS METHODOLOGY.....	27
3.5 - TEST DESCRIPTIONS.....	29
Download speed and upload speed	29
Web Browsing	29
UDP Latency and Packet Loss.....	30
Video Streaming.....	31
Voice over IP.....	31
ICMP Latency and Packet Loss	32
Latency Under Load.....	32
Cross-Talk Testing and Threshold Manager Service.....	34
4 - DATA PROCESSING AND ANALYSIS OF TEST RESULTS	35
4.1 - BACKGROUND	35

Time of Day	35
ISP and Service Tier	35
4.2 - DATA COLLECTION AND ANALYSIS METHODOLOGY	38
Data Integrity	38
Legacy Equipment	38
Collation of Results and Outlier Control	40
Peak Hours Adjusted to Local Time	41
Congestion in the Home Not Measured	41
Traffic Shaping Not Studied	41
Analysis of PowerBoost and Other 'Enhancing' Services	41
Latencies Attributable to Propagation Delay	42
Limiting Factors	42
4.3 Data Processing of Raw and Validated Data	43
5 - REFERENCE DOCUMENTS	53
5.1 - USER TERMS AND CONDITIONS	54
5.2 - Code of Conduct	64
5.3 - TEST NODE BRIEFING	65

List of Tables and Figures

Table 1: ISPs, Sample Sizes and Percentages of Total Volunteers	8
Table 2: Distribution of White-boxes by State	9
Table 3: Distribution of Whiteboxes by Census Region	11
Table 4: Panelists States Associated with Census Regions	11
Table 5: Design Objectives and Methods	17
Table 6: Overall Number of Testing Servers	23
Table 7: List of tests performed by SamKnows	27
Table 8: Estimated Total Traffic Volume Generated by Test	33
Table 9: Test to Data File Cross-Reference List	44
Table 10: Validated Data Files - Dictionary	45
Figure 1: Panelist Recruitment Protocol	13
Figure 2: Testing Architecture	20

1 - INTRODUCTION/SUMMARY

This Appendix to the 2015 Measuring Broadband America Report, a Report on Consumer Wireline Broadband Performance in the U.S., provides detailed technical background information on the methodology that produced the Report.

Specifically, this Appendix covers the process by which the panel of consumer participants was originally recruited and selected for the August 2011 Report, and then expanded over the last five years. This Appendix also discusses the testing methodology used for the report; and describes the analysis of the actual test result data.

2 - PANEL CONSTRUCTION

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

The basic objective of this study was to measure fixed broadband service performance in the United States as delivered by an 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. Although there are several ways to measure broadband performance, 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 design of the methodology allows it to be integrated with other technical measurement approaches that, in the future, could focus on other aspects of broadband performance.

2.1 - USE OF AN ALL VOLUNTEER PANEL

In 2008, SamKnows¹ conducted a test of residential broadband speed and performance in the United Kingdom² and during the course of that test determined that attrition rates for such a test were lower when an all-volunteer panel was used, rather than attempting to maintain a panel through an incentive scheme of monthly payments. Consequently, in designing the methodology for this broadband performance study, the Commission relied entirely on volunteer consumer broadband subscribers. The volunteers were selected from a large pool of prospective participants according to a plan designed to generate a representative sample of

¹ 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/>

² See http://www.samknows.com/broadband/pm/PM_Summer_08.pdf (last accessed March 27, 2015).

desired consumer demographics, including geographical location, ISP, and speed tier. As an incentive for participation, volunteers were given access to a personal reporting suite which allowed them to monitor the performance of their broadband service. They were also provided with a measurement device referred to in the study as a “Whitebox,” configured to run custom SamKnows software.³

2.2 - SAMPLE SIZE AND VOLUNTEER SELECTION

The study allowed for a target deployment of up to 10,000 Whiteboxes to volunteer panelists across the United States. The number of volunteers from each participating broadband provider was selected to ensure that the data collected would support statistically valid inferences based on a first order analysis of gathered data. Other methodological factors and considerations that influenced the selection of the sample size and makeup included:

- The panel of U.S. broadband subscribers was drawn from a pool of over 175,000 volunteers following an ongoing recruitment campaign that ran from May 2010 through September 2014.
- The volunteer sample was 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.⁴
- Target numbers for volunteers were also set across the four Census Regions—Northeast, Midwest, South and West—to help ensure geographic diversity in the volunteer panel and compensate for network variances across the U.S.⁵

³ The Whiteboxes remain in consumer homes and continue to run the tests described below. Participants may remain in the trial as long as it continues, and may retain their Whitebox when they end their participation.

⁴ The final results included volunteers from 49 states. Results collected from consumers’ Whiteboxes using some satellite and fixed terrestrial wireless technologies with low number of volunteers were not included in the report. However, data collected from those satellite and fixed terrestrial wireless subscribers are included in the detailed data files released to the public in the Raw Bulk Data Set.

⁵ 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.

- Each of the four Census Regions was split into three speed ranges: <3 Megabits per second (Mbps), 3 to 10 Mbps, ≥ 10 Mbps,⁶ with each speed tier forming an individual sample ‘cell’ against which a target number of volunteers would be selected.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 477⁸ data filed by participating ISPs for June 2011.
- An initial set of prospective participants was selected from volunteers who had responded directly to SamKnows as a result of media solicitations. Where gaps existed in the statistical sample plan, SamKnows worked with participating ISPs via email solicitations targeted at underrepresented cells. A miscellaneous cell was created across fiber-to-the-home, DSL, cable and satellite technologies, and across all regions and service tiers, to allow additional units to be allocated to accommodate volunteers who did not fit into other cells or who changed ISPs or service tiers during the trial.

Statistical experts from both the FCC and the ISPs reviewed and agreed to the plan.

Prior to the September / October 2014 testing period, 4,740 panelists from the September 2013 sample continued to supply data via their measurement devices. In addition, 1,583 subscribers were recruited after the September 2013 testing period, which brought the total subscribers reporting data in September / October 2014 to 6,323. After the data were processed, as discussed in more detail below, test results from a total of 5,583 panelists were used in the 2015 Report.

The recruitment campaign resulted in the coverage needed to ensure balanced representation of users across the U.S. Table 1 presents the number of volunteers for the month of September / October 2014 listed by ISP, as well as the percent of total volunteers accounted for by each ISP.

⁶ These speed ranges were chosen to provide alignment with broadband tiers as categorized in the “Form 477” reports that the Commission uses as its primary tool for collecting data about broadband networks and services. See *Modernizing the FCC Form 477 Data Program*, Notice of Proposed Rulemaking, 26 FCC Rcd 1508, 1512 n.27 (2011), citing *Development of Nationwide Broadband Data to Evaluate Reasonable and Timely Deployment of Advanced Services to All Americans, Improvement of Wireless Broadband Subscribership Data, and Development of Data on Interconnected Voice over Internet Protocol (VoIP) Subscribership*, Report and Order and Further Notice of Proposed Rulemaking, 23 FCC Rcd 9691, 9700-01 (2008).

⁷ The term cell is used to describe a specific number associated with a set of volunteer attributes (ISP, technology, region, speed tier) that provided a specific sample set of volunteers for the population.

⁸ 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 <http://transition.fcc.gov/form477/inst.htm# PURPOSE> for further information.

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

ISP	Sample Size	% of total volunteers
AT&T	450	8.13%
Cablevision	332	6.00%
CenturyLink	499	9.02%
Charter	525	9.49%
Comcast	901	16.29%
Cox	469	8.48%
Frontier DSL	213	3.85%
Frontier Fiber	133	2.40%
Hughes	105	1.90%
Mediacom	168	3.04%
Time Warner Cable	784	14.17%
Verizon DSL	150	2.71%
Verizon Fiber	475	8.59%
ViaSat/Exede	67	1.21%
Windstream	261	4.72%
Total	5532	

The distribution of Whiteboxes by state is found in following Table⁹ below.

⁹ Subscriber data in the 2015 Report is based on the FCC's Internet Access Services Report with data current to 31 December, 2013. See Internet Access Services: Status as of December 31, 2013, Wireline Competition Bureau, Industry Analysis and Technology Division (rel. June 2012), *available at* http://transition.fcc.gov/Daily_Releases/Daily_Business/2014/db1016/DOC-329973A1.pdf

Table 2: Distribution of White-boxes by State

State	Total Boxes	% of total Boxes	% of Total US Broadband
AL	53	0.99%	1.3%
AR	40	0.75%	0.7%
AZ	220	4.11%	2.1%
CA	512	9.56%	11.5%
CO	120	2.24%	1.8%
CT	75	1.40%	1.3%
DC	15	0.28%	0.2%
DE	8	0.15%	0.3%
FL	207	3.87%	6.9%
GA	162	3.03%	2.9%
HI	15	0.28%	0.4%
IA	173	3.23%	1.0%
ID	29	0.54%	0.5%
IL	191	3.57%	3.9%
IN	94	1.76%	1.9%
KS	33	0.62%	0.9%
KY	107	2.00%	1.3%
LA	41	0.77%	1.3%
MA	106	1.98%	2.4%
MD	86	1.61%	1.9%
ME	13	0.24%	0.5%
MI	160	2.99%	3.1%
MN	137	2.56%	1.8%
MO	124	2.32%	1.8%
MS	16	0.30%	0.7%
MT	5	0.09%	0.3%

2015 Technical Appendix

NC	194	3.62%	3.2%
ND	3	0.06%	0.2%
NE	49	0.92%	0.6%
NH	21	0.39%	0.5%
NJ	222	4.15%	3.1%
NM	52	0.97%	0.6%
NV	48	0.90%	0.9%
NY	387	7.23%	6.8%
OH	226	4.22%	3.9%
OK	63	1.18%	1.1%
OR	163	3.04%	1.3%
PA	182	3.40%	4.3%
RI	15	0.28%	0.4%
SC	74	1.38%	1.5%
SD	3	0.06%	0.3%
TN	71	1.33%	1.8%
TX	210	3.92%	7.4%
UT	41	0.77%	0.8%
VA	216	4.03%	2.6%
VT	6	0.11%	0.2%
WA	184	3.44%	2.4%
WI	166	3.10%	1.8%
WV	15	0.28%	0.6%
WY	1	0.02%	0.2%
	5354		

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
Northeast	1027	19.18%	19.47%
Midwest	1359	25.38%	21.43%
South	1578	29.47%	35.55%
West	1390	25.96%	23.55%

The distribution of States associated with the four Census Regions used statistically 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 were recruited in the 2011, 2012, 2013 and 2014 panels using the following method:

- A significant proportion of volunteers were recruited via an initial public relations and social media campaign led by the FCC. This included discussion on the FCC website and on technology blogs, as well as articles in the press regarding the study. The demographics of this initial panel were reviewed to identify any deficiencies with regard to the sample plan described above. These goals were set to produce statistically valid sets of volunteers for demographics based on ISP, speed tier, technology type, and region. This initial pool of volunteers was then supplemented by the participating ISPs,

who sent out an email to customers in desired demographics that were under-represented in the pool of publicly-solicited volunteers. Emails directed interested volunteers to contact SamKnows regarding participation in the trial. At no time during this recruitment process did the ISPs have any knowledge regarding which of their customers might be participating in the trial. In almost all cases, ISP engagement in soliciting volunteers enabled us to meet desired demographic targets.

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

A multi-mode strategy was used to qualify volunteers for this trial. 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, which was used to track subsequent submissions by these volunteers.
2. Those volunteers who were determined to be representative of the target broadband user population were sent a follow-up email, which invited participation in a web-based speed test that was developed by SamKnows in collaboration with Measurement Lab ("M-Lab") and PlanetLab.¹⁰
3. Volunteers were selected from respondents to this follow-up email based on the statistical requirements of the panel. Selected volunteers were then asked to complete an acknowledgment of User Terms and Conditions that outlined the permissions to be granted by the volunteer in key areas such as privacy.¹¹
4. Of those volunteers that completed the User Terms and Conditions, SamKnows selected the first panel of 13,000 participants,¹² each of whom received a Whitebox for self-installation. SamKnows provided full support during the Whitebox installation phase.


¹⁰ M-Lab is a consortium supporting research on broadband networks. PlanetLab is a global research network supporting the development of new network services. More information on M-Lab and PlanetLab can be found at <http://www.measurementlab.net> <http://planet-lab.org>, respectively.

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

¹² Over 15,000 Whiteboxes have been shipped to targeted volunteers since 2011, of which 6,323 were online and reporting data used in the 2015 Report from the months of September / October 2014.

Figure 1: Panelist Recruitment Protocol

The graphic below illustrates the study recruitment methodology:



2.4 - VALIDATION OF VOLUNTEERS' SERVICE TIER

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

The broadband service tier reported by each panelist was authenticated in the following way:

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 advertised 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 eighty percent of broadband consumers did not know what speed they had purchased).

- At the time of recruitment, each panelist was required to complete a speed test using an M-Lab server. This test provided a rough approximation of the panelist's service tier which served to identify panelists with targeted demographics, and highlighted anomalies in the panelist's survey response to measured speed.
- At the time the panelist installed the Whitebox, the device automatically ran an IP 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 in question.
- 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.

Of the more than 15,000 Whiteboxes that were shipped to panelists since 2011, 6,332¹⁵ units were reporting data in September / October 2014. The participating ISPs validated 80 percent of these panelists, of which 7 percent were reallocated to a different tier following the steps listed above. The remaining 20 percent of panelists were validated based on comparing the

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 2015 Report, but will be included in the Raw Bulk Data Set.

¹⁵ This figure represents the total number of boxes reporting during September / October 2014, the month chosen for the 2015 Report. Shipment of boxes continued in succeeding months and these results will be included in Raw Bulk Data Set.

performance data and line performance characteristics with the available service tiers from the appropriate ISP. Eliminating panelists who either changed ISPs during the period of September / October 2014 or did not produce data for this trial during that month produced the final data set of the approximately 5,583 volunteers included in the 2015 Report.

2.5 - PROTECTION OF VOLUNTEERS' PRIVACY

A major concern during this trial was to ensure that panelists' privacy was protected. The panel was comprised entirely of volunteers who knowingly and explicitly opted-in to the testing program. Correspondence with panelists documenting their opt-in is preserved in confidence for audit purposes.

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.

To fulfill these privacy requirements a range 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.

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 - SELECTION OF HARDWARE APPROACH

A fundamental choice when developing a solution to measure broadband performance is whether to use a hardware or software approach.

Software approaches are by far the most common and allow a very large sample to be reached relatively easily. Web-based speed tests fall into this category. These typically use Flash or Java applets, which execute within the context of the user's web browser. When initiated, these clients download content from remote web servers and measure the throughput of the transfer. Some web-based speed tests also perform upload tests, while others perform basic latency checks.

Other less common software-based approaches to performance measurement involve installing applications on the user's workstation which periodically run tests while the computer is switched on.

All software solutions implemented on a consumer's computer, smart phone, or other Internet access device suffer from the following disadvantages for the purposes of this study:

- The software may itself affect broadband performance;
- The software typically does not account for multiple machines on the same network;
- The software may be affected by the quality and build of machine;

- Potential bottlenecks (such as wireless equipment, misconfigured networks, and older computers) are generally not accounted for and result in unreliable data;
- A consumer may move the computer or laptop to a different location which can affect performance;
- The tests may only run when the computer is actually on, limiting the ability to provide a 24-hour profile;
- For manually-performed software tests, panelists may introduce a bias by when they choose to run the tests (e.g., may only run tests when they are encountering problems with their service).

In contrast, hardware approaches involve placing a device inside the user's home that is physically connected to the consumer's Internet connection, and periodically running tests to remote targets on the Internet. These hardware devices are not reliant on the user's workstation being switched on, and so allow results to be gathered throughout the day and night. The primary disadvantages of a hardware approach are that this solution is much more expensive than a software approach and requires installation of the hardware by the consumer or a third party.

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 seventeen technical objectives:

Table 5: Design Objectives and Methods

#	Technical Objectives	Methodological Accommodations
1	Must not change during the monitoring period.	The Whitebox measurement process is designed to provide automatic 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.
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 research program.	Whiteboxes can be disconnected at any time from the home network. As soon as the route 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 an in-line or bridging mode.
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.	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 from the remote control center if it contains any software or firmware components.	The Whitebox can be completely controlled remotely for updates without involvement of the consumer PC, providing the Whitebox is switched on and connected.
12	Must identify when a user changes broadband provider or package (e.g.,	Ensures regular data pool monitoring for changes in speed, ISP, IP address or performance, and flags when a panelist

	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).	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 and 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.

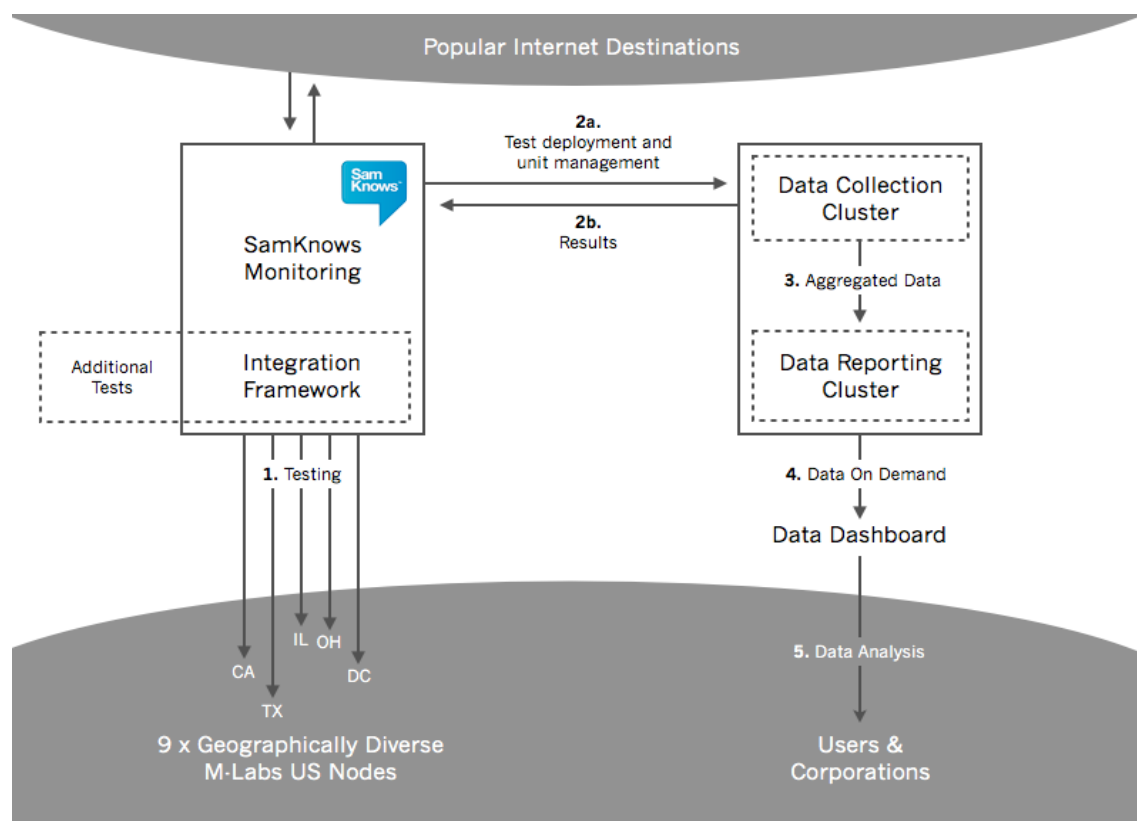
¹⁶ Signatories to the Code of Conduct are: Adtran, AT&T, Cablevision, CenturyLink, Charter, Comcast, Corning, Cox, Fiber to the Home Council, Frontier, Georgia Tech, Genband, Hughes, Insight (by TWC), Internet Society, Level3, Measurement Lab, Mediacom, Motorola, Netflix, Qwest, Time Warner Cable, US Telecom, 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

Testing Architecture Overview

As illustrated below, the performance monitoring system comprised a distributed network of Whiteboxes in the homes of members of the volunteer consumer panel, and was used to accurately measure the performance of fixed broadband connections based on real-world usage. The Whiteboxes were controlled by a cluster of servers, which hosted the test scheduler and the reporting database. The data was collated on the reporting platform and accessed via a reporting interface¹⁷ and secure FTP site. The system also included a series of speed-test servers, which the nodes 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 statistics associated with their household's broadband performance as reported by the Whitebox.

Approach to Testing and Measurement

Any network monitoring system needs to be capable of monitoring and executing tests 24 hours a day, 7 days a week. Similar to the method used by the television audience measurement industry, each panelist was equipped with a Whitebox, which was self-installed by each panelist and conducted the performance measurements.

During the initial testing period in 2011, the Whitebox provided was based upon hardware manufactured by NETGEAR, Inc. (NETGEAR) and operated as a broadband router. It was intended to replace the panelist's existing router and be the first consumer premise equipment connected to the provider's network after the modem. The installation of the Whitebox directly after the consumer's home Internet connection ensured 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.

Beginning with the 2012 testing period, a second Whitebox model was introduced. This version was based upon hardware produced by TP-Link and operated as a bridge, rather than as a router. It was designed to be installed after the customer's existing router – it did not replace the study participant's router. The TP-Link Whitebox was designed to have all home devices connected via the LAN ports on the rear of the Whitebox. The TP-Link Whitebox passively monitored wireless network activity in order to determine when the network was active (in order to defer measurements), and ran a modified version of OpenWrt, an open source router platform based upon Linux.

The 2015 Report incorporates data from both types of Whitebox. Testing has found that the results they produce are indistinguishable, and the term "Whitebox" is used throughout the 2015 Report to describe both types of devices.

The software that came pre-installed on each of the Whiteboxes was programmed to execute a series of tests designed to measure the key performance indicators (KPIs) of a broadband connection by simulating real world usage. The tests were a suite of applications, written by SamKnows in the programming language C, which were rigorously tested by the ISPs and other stakeholders. Whiteboxes have been shown to provide accurate information about lines with throughput rates in excess of 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 as per the regular NETGEAR instructions.

This approach was the default approach used across all ISPs and service tiers. However, this approach 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. Two such scenarios were:

1. Some Verizon FiOS gateways utilizing a MoCA (Multimedia over Cable) interface; and
2. AT&T U-Verse gateways providing U-Verse specific features, such as IPTV.

In such scenarios the Whitebox was connected to the existing router/gateway and all devices connected through it. In order to prevent a “double-NAT” configuration issue, in which multiple routers on the same network perform Network Address Translation (NAT) making access to the interior (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

Because the Whitebox could detect all wired and wireless traffic, no tests were performed when there was any Internet activity beyond a defined threshold value. This ensured both that testing did not interfere with consumer use of his or her Internet service and that any such use did not interfere with or invalidate testing.

Home Deployment of the TP-Link Based Whitebox

The TP-Link based Whitebox, which operated as a bridge, rather than a router, was introduced in response to the increased deployment of integrated modem/gateway devices (for example, Verizon FiOS and AT&T U-Verse). The NETGEAR based Whitebox was unable to replicate all of the features provided by these highly customized routers supplied by the ISPs, and therefore had to operate behind the ISP’s router as a bridge.

Given the increasing trend of merging modem and router into a single device, and offering additional over-the-top services via this device, SamKnows determined that it was necessary to develop a new Whitebox which operated natively as a bridge.

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.

To use the TP-Link based Whitebox, panelists were 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 was required so that the Whitebox could measure traffic volumes from wired devices in the home and defer tests accordingly. The Whitebox would not manipulate NAT traffic or provide any DHCP services – it operated as a pure Ethernet bridge.

Panelists were not asked to change their wireless network configurations. The TP-Link based Whitebox passively monitored the strongest wireless network for traffic and, like the NETGEAR Whitebox, did not conduct measurements when activity was detected.¹⁹

Test Nodes (Off-Net and On-Net)

For the tests in this study, SamKnows employed nine core measurement or reference points that were distributed geographically across nine locations. These off-net measurement points were supplemented by additional measurement points located on the networks of some of the ISPs participating in this study. Together, the core measurement points were used to measure consumers' broadband performance between the gateway device and an available reference point that was closest in transit time to the consumer's 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 scientific checks and insight into broadband service performance within an ISP's network. In total, the following 133 measurement servers were deployed in conjunction with the program:

Table 6: Overall Number of Testing Servers

Owner	Count
AT&T	11
Cablevision	2
CenturyLink	13
Charter	5
Comcast	33
Cox	2
Frontier	5
Level 3	10
M-Lab	32

¹⁹ Note that encryption of the wireless network was irrelevant for these checks as the TP-Link based Whitebox only measured traffic volume, without any analysis of the contents of packets.

Mediacom	1
Qwest	4
Time Warner Cable	6
Verizon	5
Windstream	4

Test Node Breakout

Off-Net Test Nodes

The M-Lab infrastructure served as destinations for the remote tests during this study. Nodes were located in the following major U.S. Internet peering locations:

- New York City, New York (2 locations)
- Chicago, Illinois
- Atlanta, Georgia
- Miami, Florida
- Washington, DC
- Mountain View, California (2 locations)
- Seattle, Washington
- Los Angeles, California
- Dallas, Texas

The Level 3 infrastructure served as destinations for the remote tests during this study. Nodes were located in the following major U.S. Internet peering locations:

- Chicago, Illinois (2 locations)
- Dallas, Texas
- New York City, New York
- San Jose, California (2 locations)
- Washington D.C. (2 locations)
- Los Angeles, California (2 locations)

On-Net Test Nodes

In addition to off-net nodes, some ISPs implemented their own on-net servers as an audit to the 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
- Cablevision
- CenturyLink²⁰
- Charter
- Comcast
- Cox
- Frontier
- Mediacom
- Qwest
- Time Warner Cable
- 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.

It is important to note that 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 this study are based only on tests performed using off-net nodes. The use of both on-net and off-net nodes provides a measure of confidence in the test results.

Results from both on-net and off-net nodes will be included in the Raw Bulk Data Set that will be released to the public. Test nodes were continually monitored for load and congestion; this end-to-end control of both the test node and Whitebox provided a high level of integrity in testing. The results contained in the 2015 Report are based on the measurements obtained from the M-Lab and Level 3 peering reference points only, while the Raw Bulk Data Set contains all results, including those from the ISP-only reference points.

²⁰ QWest was reported separately from Centurylink in past reports. The entities have completed merging their test infrastructure and are reported together from this report.

Test Node Selection

Having a geographically diverse set of test nodes would be of little use if the Whiteboxes running the test did not have a suitable mechanism to determine which node offered the lowest round trip time between itself and the panelist's Whitebox.

The node actually selected might not always be the geographically closest test node to the panelist; the network route between the panelist's home and the test node will often travel via an indirect route that may take it through one or more cities. This might make another test node that is physically farther away preferable.

To identify nodes with the lowest round trip time, each Whitebox fetched a complete list of test nodes from the SamKnows infrastructure upon first execution of the test batch and performed a simple round trip time measurement to each. It then selected the test node with the lowest round trip time to test against from that point forward.

Technical details of the test node's minimum requirements for hardware and software, connectivity, and systems and network management are available in the [5.2 - Test Node Briefing](#) provided in the Reference Document section of this Appendix.

3.4 - SAMKNOWS METHODOLOGY

SamKnows Methodology²¹ - Each deployed Whitebox performs the following tests:

Table 7: List of tests performed by SamKnows

Test	Primary measure(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 time to fetch a page and all of its 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
Video streaming	Initial time to buffer, number of buffer under-runs and total time for buffer delays ²²
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 regularly spaced ICMP packets
ICMP packet loss	Percentage of packets lost in the ICMP latency test

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

²² Only the total buffer delay is presented in the Tabular Test Results spreadsheet. Results of all tests will be included in the Raw Bulk Data files.

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

Latency under load	Average round trip time for a series of regularly spaced UDP packets sent during downstream/upstream sustained tests
Availability ²⁴	Total time the connection was deemed unavailable for any purpose, which could include a network fault or unavailability of a measurement point
Consumption ²⁵	A simple record of the total bytes downloaded and uploaded by the router

²⁴ The measurement of availability provided a check on how often tests could not be run and was used as a quality metric overall, but was not used in analysis of broadband performance.

²⁵ While this metric is not an active test it is included in this description as a passive test.

3.5 - TEST DESCRIPTIONS

The following sub-sections detail the methodology used in each of the individual tests.

Download speed and upload speed

These tests measured 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—was hosted on a web server on the target test node. The test operated for a fixed duration of 30 seconds. It also recorded average throughput at 5 second intervals during the test. The client attempted to download as much of the payload as possible for the duration of the test.

The test used three concurrent TCP connections (and therefore three concurrent HTTP requests) to ensure that the line was saturated. Each connection used in the test counted the numbers of bytes transferred and was 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 was taken as the total throughput of the line.

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

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

Web Browsing

The test recorded 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 was the total time taken to download the HTML front page for each web site and all associated images, JavaScript, and stylesheet resources. This test did not test against the centralized testing nodes; instead it tested against real websites, ensuring that the effects of content distribution networks and other performance enhancing factors could be taken into account.

Each Whitebox tested against the following 9 websites:²⁶

- <http://www.cnn.com>
- <http://www.youtube.com>
- <http://www.msn.com>
- <http://www.amazon.com>
- <http://www.yahoo.com>
- <http://www.ebay.com>
- <http://www.wikipedia.org>
- <http://www.facebook.com>
- <http://www.google.com>

The results include the time taken for DNS resolution. The test used up to eight concurrent TCP connections to fetch resources from targets. The test pooled TCP connections and utilized persistent connections where the remote HTTP server supports them.

The client advertised the user agent as Microsoft Internet Explorer 8. Each website was tested in sequence and the results summed and reported across all sites.

UDP Latency and Packet Loss

These tests measured 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 packet was not received back within three seconds of sending, it was treated as lost. The test recorded the number of packets sent each hour, the average round trip time and the total number of packets lost. The test computed the summarized minimum, maximum, and mean from the lowest 99% of results, effectively trimming the top (i.e., slowest) 1% of outliers.

The test operated continuously in the background. It was configured to randomly distribute the sending of the echo requests over a fixed interval of one hour, reporting the summarized results once the interval had elapsed. Approximately 600 packets were sent within a one hour period, with fewer packets sent if the line was not idle.

This test was started when the Whitebox booted and ran permanently as a background test.

²⁶ These websites were chosen based on a list by Alexa, <http://www.alexa.com/>, of the top twenty websites in October 2010.

Video Streaming

For the purpose of the video streaming test, the intent was to simulate an end user viewing a streaming video online. This test emulated live video streaming rather than a service such as YouTube that employs a ‘progressive download’ approach.

The test operated over TCP and used proprietary client and server side components. The client and server negotiated the test parameters at the start of each test.

A three-second playout buffer was configured and the client attempted to download data from the server at the maximum rate necessary to ensure that this buffer was never empty. A separate client-side thread consumed data from this buffer at a fixed rate, looking for buffer under-runs (which would manifest themselves to users as a pause in video). The Whitebox recorded the time to initial buffer, the total number of buffer under-runs and the total delay in microseconds due to these under-runs. The test operated at four bit rates: 768 kilobits per second (kbps), 1.25 Mbps, 2.25 Mbps, and 3.75 Mbps.

The test is executed from a subscriber’s Whitebox to measurement servers, along reengineered paths intended to minimize potential impacts of general routing and switching performance beyond the first tier-1 peering link of the broadband provider. Measurement servers are engineered for high-performance to minimize potential impacts of performance that might affect the measurement of network elements within the scope of control of a subscriber's broadband provider. Because the test focuses on the consumer’s actual download speed, latency, and packet loss under these controlled conditions, other real world considerations such as a subscriber's home computing, Wi-Fi and other networking environments, specific peering and transit relationships of Internet network providers, video streaming provider network and technology architectures, and other considerations that could potentially impact video performance, are not measured by this test.

Voice over IP

The Voice over IP (VoIP) test operated over UDP and, unlike the video streaming test, utilized bi-directional traffic, as is typical for voice calls.

The Whitebox would handshake with the server, and each would initiate a UDP stream with the other. The test used a 64 kbps stream with the same characteristics and properties (i.e., packet sizes, delays, bitrate) as the G.711 codec. The test measured jitter, delay and loss. These metrics were measured by subdividing the stream into blocks, and measuring the time taken to receive each block (as well as the difference between consecutive times).

Jitter was calculated using the PDV approach described in section 4.2 of RFC5481. The 99th percentile was recorded and used in all calculations when deriving the PDV.

As with the video streaming test, this test is executed under controlled conditions and does not capture other real world considerations that could potentially impact VoIP performance.

DNS Resolutions and DNS Failures

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

ICMP Latency and Packet Loss

These tests measured the round trip time (RTT) of ICMP echo requests in microseconds from the Whitebox to a target test node. The client sent 5 ICMP echo requests of 56 bytes to the target test node, waiting up to three seconds for a response to each. Packets that were not received in response were treated as lost. The mean, minimum, maximum, and standard deviation of the successful results were recorded.

Latency Under Load

The latency under load test operated for the duration of the 30 second downstream and upstream speed tests, with results for upstream and downstream recorded separately. While the speed tests ran, the latency under load test sent UDP datagrams to the target server and measured the round trip time and number of packets lost. Packets were spaced 500 milliseconds (ms) apart, and a 3 second timeout was used. The test recorded the mean, minimum and maximum round trip times in microseconds. The number of lost UDP packets was also recorded.

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

Table 8: Estimated Total Traffic Volume Generated by Test

Test Name	Test Target(s)	Test Frequency	Test Duration	Total Est. Daily Volume
Web browsing	10 popular US websites	Hourly, 24 x 7	Est. 30 seconds	80 MB
Video streaming*	1 off-net test node	Hourly, 24 x 7	Fixed 10 seconds at 768kbps, 1.25Mbps, 2.25Mbps, 3.75Mbps	60 MB
	1 on-net test node	Hourly, 24 x 7	Fixed 10 seconds at 768kbps, 1.25Mbps, 2.25Mbps, 3.75Mbps	60 MB
Voice over IP	1 on-net test node	Hourly, 24 x 7	Fixed 30 seconds at 64k	1 MB
	1 off-net test node	Hourly, 24 x 7	Fixed 30 seconds at 64k	1 MB
Download speed**	1 off-net test node	Every other hour, 24 x 7	Fixed 30 seconds***	2.1 GB at 50Mbps
				860 MB at 20Mbps
				430 MB at 10Mbps
				130 MB at 3Mbps
	1 on-net test node	Every other hour, 24 x 7	Fixed 30 seconds***	65 MB at 1.5Mbps
				2.1 GB at 50Mbps
				860 MB at 20Mbps
				430 MB at 10Mbps
Upload speed**	1 off-net test node	Every other hour, 24 x 7	Fixed 30 seconds***	130 MB at 3Mbps
				65 MB at 1.5Mbps
				174 MB at 2Mbps
	1 on-net test node	Every other hour, 24 x 7	Fixed 30 seconds***	87 MB at 1Mbps
				44 MB at 0.5Mbps
				174 MB at 2Mbps
UDP latency	1 off-net test node	Hourly, 24 x 7	Permanent	87 MB at 1Mbps
	1 on-net test node	Hourly, 24 x 7	Permanent	44 MB at 0.5Mbps
UDP packet loss	1 off-net test node	Hourly, 24 x 7	Permanent	1 MB
	1 on-net test node	Hourly, 24 x 7	Permanent	1 MB
Consumption	N/A	24 x 7	N/A	N/A (uses above)
DNS resolution	10 popular US websites	Hourly, 24 x 7	Est. 3 seconds	N/A (uses above)
ICMP latency	1 off-net test node	Hourly, 24 x 7	Est. 5 seconds	N/A (uses above)
	1 on-net test node			N/A (uses above)
ICMP packet loss	1 off-net test node	Hourly, 24 x 7	N/A (As ICMP latency)	N/A (uses above)
	1 on-net test node			N/A (uses above)

* Video streaming rates: Lines will only stream the rates they are capable of, according to the latest speed test results. If a rate is deemed unreachable (e.g. a 3.75Mbps rate on a 1Mbps line), then it will be skipped.

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

*** The speed tests run for a fixed duration of 30 seconds, but output the cumulative average transfer speed every 5 seconds (i.e. the first average is for seconds 0-5, the second is for 0-10, etc.). This allows us to measure performance over multiple durations, while running only one test.

With the introduction of Level3 measurement points, tests to the 'off-net' destinations alternate randomly between Level3 and M-Lab. The one exception to this rule is the latency and loss tests, which test continuously to both Level3 and M-Lab off-net servers, as well as the on-net server (if one exists).

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 monitored the inbound and outbound traffic across the WAN interface to calculate if a panelist was actively using the Internet connection. The threshold for traffic was set to 64 kbps downstream and 32 kbps upstream. Statistics were sampled and computed every 10 seconds. If either of these thresholds was breached, the test was delayed for a minute and the process repeated. If the connection was being actively used for an extended period of time, this pause and retry process would continue for up to 5 times before the entire test cycle was abandoned.

4 - DATA PROCESSING AND ANALYSIS OF TEST RESULTS

This section describes the background for the categorization of data gathered for the 2015 Report, and the methods employed to collect and analyze the test results.

4.1 - BACKGROUND

Time of Day

One of the key factors that affects broadband performance is usage-based congestion. At peak hours, defined for this study as the period on weekdays between 7:00 pm and 11:00 pm local time, there are more people attempting to use the Internet simultaneously, giving rise to the potential for congestion. When congestion occurs, users' performance will suffer.

ISP and Service Tier

A sufficient sample size is necessary to allow meaningful statistical analysis and the ability to robustly compare the performance of specific ISP packages. The study achieved statistically meaningful sample sizes for the following download and upload speeds²⁷ (listed in alphabetical order):

Download Speeds:

- AT&T DSL's 3 Mbps, and 6 Mbps tiers;
- AT&T U-verse's 6 Mbps, 12 Mbps, 18 Mbps, and 24 Mbps tiers;
- Cablevision's 15 Mbps, 50 Mbps and 101 Mbps tiers;
- CenturyLink's 1.5 Mbps, 3 Mbps, 7 Mbps, 10 Mbps, 12 Mbps, 20 Mbps and 40 Mbps tiers;

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 to produce enough samples to provide statistically valid data.

Charter's 15 Mbps, 30 Mbps, 60 Mbps and 100 Mbps tiers;
 Comcast's 3 Mbps, 15 Mbps, 50Mbps and 105 Mbps tiers;
 Cox's 5 Mbps, 25 Mbps, 50 Mbps and 100 Mbps tiers;
 Frontier DSL's 1 Mbps, 3 Mbps and 6 Mbps tiers;
 Frontier Fiber's 25 Mbps tier;
 Hughes' 5 Mbps and 10 Mbps tier;
 Mediacom's 15 Mbps and 50 Mbps tiers;
 Time Warner Cable's 15 Mbps, 20 Mbps, 30 Mbps, 50 Mbps and 100 Mbps tiers;
 Verizon DSL's [0.5 - 1.0] Mbps and [1.1 - 3.0] Mbps tiers;
 Verizon Fiber's 15 Mbps, 25 Mbps, 35 Mbps, 50 Mbps and 75 Mbps tiers;
 Viasat/Excede's 12 Mbps tier; and
 Windstream's 1.5 Mbps, 3 Mbps, 6 Mbps, and 12 Mbps tiers.

Upload Speeds:

AT&T DSL's 384 kbps, and 512 kbps tiers;
 AT&T U-verse's 1 Mbps, 1.5 Mbps, and 3 Mbps tiers;
 Cablevision's 5 Mbps, 25 Mbps and 35 Mbps tiers;
 CenturyLink's 256 kbps, 512 kbps, 640 kbps, 768 kbps, 896 kbps and 5 Mbps tiers;
 Charter's 3 Mbps and 4 Mbps tiers;
 Comcast's 768 kbps, 5 Mbps, 10 Mbps and 20 Mbps tiers;
 Cox's 1 Mbps, 5 Mbps, and 10 Mbps tiers;
 Frontier DSL's 384 kbps, 768 kbps tiers;
 Frontier Fiber's 5 Mbps, 10 Mbps, and 25 Mbps tiers;
 Hughes' 1 Mbps tier;
 Mediacom's 1 Mbps, 5 Mbps, and 10 Mbps tiers;
 Time Warner Cable's 1 Mbps, 2 Mbps, and 5 Mbps tiers;
 Verizon DSL's 384 kbps, and 768 kbps tiers;
 Verizon Fiber's 15 Mbps, 25 Mbps, 35 Mbps and 50 Mbps tiers;
 Viasat/Excede's 3 Mbps tier; and
 Windstream's 384 kbps and 768 kbps tiers.

Statistical averages for the validated September / October 2014 data are found on FCC's Measuring Broadband America site.²⁸ 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 See: <http://data.fcc.gov/download/measuring-broadband-america/2015/statistical-averages-Sept-2014.xlsx>

Recent advances in satellite broadband technology have allowed the inclusion of new services from ViaSat and Hughes to be incorporated into this study without a change in technical methodology. Whiteboxes were connected to ViaSat and Hughes connections in the same fashion as DSL, cable and fiber services, and ran the same suite and schedule of tests. Their results are included in the statistical averages spreadsheet.

4.2 - DATA COLLECTION AND ANALYSIS METHODOLOGY

Data Integrity

As the Whiteboxes ran tests consistently from homes across the U.S., it was important to check the data to ensure that any anomalies were removed. To ensure the integrity of the large amount of data collected, the following protocols were developed:

1. Change of ISP intra-month: found units that changed ISP intra-month (determined by performing daily WHOIS query using the panelist's IP address), and removed data for the ISP on which they spent less time over the course of that month.
2. Change of service tier intra-month: found units that changed service tier intra-month by isolating the difference between the average sustained throughput observed for the first three days in the reporting period from the average sustained throughput observed for the final three days in the reporting period. If a unit was not online at the start or end of that period, then the first/final three days that they were actually online were taken. If this difference was over 50%, 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: removed any failed or irrelevant tests by removing measurements against any non-M-Lab or Level 3 servers (to catch tests to ISP test nodes). Removed measurements against any M-Lab or Level 3 server outside of the U.S. Removed measurements against any M-Lab or Level 3 server that exhibited greater than or equal to 10% failures in a specific one hour period (the purpose was to remove periods where M-Lab or Level 3 servers were unavailable).
4. Removal of any problem units: removed measurements for any unit that exhibited greater than or equal to 10% failures in a particular one hour period (the purpose was to remove periods where units were 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.²⁹ As an example, in the 2014 report we reviewed the capabilities of cable modems within the homes of consumer volunteers participating in this study. This information is important because measured end-to-end service performance of cable broadband service is a function both of the capabilities of the service provider's network and of the capabilities of the cable modem which terminates the service within a subscriber's home. In other words, a consumer's ability to actually receive the provider's advertised speeds depends upon the capabilities of the cable modems within their home.

Cable ISPs have described two scenarios that may affect the results of this study. First, some consumers own their modems and have not yet upgraded to take advantage of the higher speeds enabled by DOCSIS 3, the latest standard cable technology. Second, some consumers who lease cable modems and have been provided with free upgrades nonetheless may have failed to install these new modems at the time of data collection. Cable ISPs requested that we review our sample of volunteers and identify panelists using legacy equipment that would not achieve the provisioned capacity available to the consumer and thus would introduce a possible inaccuracy in measured performance of the network under actual control of the ISP.

In response, we developed a policy 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 their service.


While the issue of DOCSIS 3 modems and network upgrades affect the cable industry today, this is a general issue concerning network investment and evolution and the impact on equipment that the provider places within the consumer's household and is under their direct control.

In accordance with the above stated policy we checked for the effect of inclusion of legacy cable modem on the download speed as a percentage of the advertised speed. The problems for legacy modems was observed this year for only Comcast's 105 Mbps tier and the differences in the download and upload speeds as a percentage of the advertised speed values with and without the legacy equipment was less than 2%. The results are shown in Figure 3 below:

²⁹ 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>

Figure 3: Download and Upload Speeds- legacy modem analysis



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 pm to 11:00 pm local time Monday through Friday, 12:00 am to 12:00 am local time Saturday and Sunday). Only participants who provided a minimum of one week (seven days) of valid measurements and had valid data in each of the three time intervals were included in the September 2013 test results. In addition, the top and bottom 1% of measurements were dropped to control for outliers that may have

been anomalous or otherwise not representative of actual broadband performance. All statistics were computed on the trimmed data.³⁰

Data was only charted when results from at least 25 separate Whiteboxes was available. Instances of 30 or fewer Whiteboxes were noted for possible future augmentation.

The resulting final sample of data for September / October 2014 was collected from 6,323 participants.

Peak Hours Adjusted to Local Time

Peak hours were defined as weekdays between 7:00 pm to 11:00 pm (inclusive) for the purposes of the study. All times were adjusted to the panelist's local time zone. Due to some tests that only took place once every two hours on an individual Whitebox, the period used for aggregating peak performance had to be a multiple of two.

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 U.S.-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 nine highly trafficked websites configured to serve high traffic loads may have mitigated the effects of temporary congestion.

Traffic Shaping Not Studied

The effect of traffic shaping is not studied in the 2015 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 advertised service tiers, may be most readily seen in those charts in the 2015 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 (the size of this duration may vary by ISP, service tier, and potentially other factors). For example, this could mean that a user with a contracted 6 Mbps service tier may receive 18 Mbps for the first 10MB

³⁰ These methods were reviewed with statistical experts within the FCC and by participating ISPs.

of a 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 get 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 2 seconds on a connection employing burst speed techniques would likely record the burst speed rate, whereas a speed test running for 2 hours will reduce the effect of burst speed techniques to a negligible level.

The speed test 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 2015 report we note that the use of this technology by providers is on the decline.

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 5 ms 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 17,400,823,950 measurements were taken across 147,520,083 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 is made available as OpenData for review and use by the public. Raw and processed data sets, 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 develop FAQ for database configuration and data importing instructions for mysql and postgresql are available.

<http://www.fcc.gov/measuring-broadband-america/2012/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 below flowchart and supporting files that provide detail of the process.

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. 2014 Raw Data Links
Clean Data:	Data is cleaned. This includes removing measurements when a user changed ISP/product during the period. Anomalies and significant outliers are also removed at this point. A data cleansing document describes the process in detail. 2014 Data Cleansing Document Link
Per-Unit Results(CSV):	Per-unit results are generated for each metric. Time of day averages are computed and a trimmed mean is calculated for each metric. The SQL scripts used here are contained in SQL processing scripts available with the release of each report. 2014 SQL Processing Links
SPSS Processing:	The per-unit CSV data is processed by SPSS scripts coupled with the unit profile data. This process removes ISPs/products with low sample sizes and computes statistical averages for the remainder that can be used in the report. 2014 SPSS Scripts Links
Excel Tables & Charts:	Summary data tables and charts in Excel are produced from the statistical averages. These are used directly in the report. 2014 Charts 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_udpjitter.csv
Video streaming	curr_videostream.csv
Voice over IP	curr_udpjitter.csv
DNS resolution	curr_dns.csv
DNS failures	curr_dns.csv
ICMP latency	curr_dlping.csv — Download ICMP tests curr_ulping.csv — Upload ICMP tests
ICMP packet loss	curr_dlping.csv — Download ICMP tests curr_ulping.csv — Upload ICMP tests
Latency under load	curr_udplatency.csv — IPv4 Tests curr_udplatency6.csv — IPv6 Tests
Consumption ³¹	curr_netusage.csv

³¹ While this metric is not an active test it is included in this description as a passive test.

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

<u>curr_dlping.csv</u>	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	Target hostname or IP address
rtt_avg	Average RTT in microseconds
rtt_min	Minimum RTT in microseconds
rtt_max	Maximum RTT in microseconds
rtt_std	Standard Deviation in Measured RTT in microseconds
successes	Number of successes
failiures	Number of failures
location_id	Please ignore (this is an internal key mapping to unit profile data)
<u>curr_dns.csv</u>	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
nameserver	Nameserver used to handle the DNS request
lookup_host	Hostname to be resolved
response_ip	Field unused at present
rtt	DNS resolution time in microseconds
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)
<u>curr_httpgetmt.csv</u>	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC

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

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 in microseconds
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 (Units: microseconds)
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	Please ignore (this is an internal key mapping to unit profile data)
<u>curr_httpostmt.csv</u>	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
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 in microseconds
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 (Units: microseconds)

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	Please ignore (this is an internal key mapping to unit profile data)
<u>curr_ping.csv</u>	ICMP based
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	Target hostname or IP address
rtt_avg	Average RTT in microseconds
rtt_min	Minimum RTT in microseconds
rtt_max	Maximum RTT in microseconds
rtt_std	Standard Deviation in Measured RTT in microseconds
successes	Number of successes
failures	Number of failures
location_id	Please ignore (this is an internal key mapping to unit profile data)
<u>curr_udpjitter.csv</u>	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	Target hostname or IP address
packet_size	Size of each UDP Datagram (Units: Bytes)
stream_rate	Rate at which the UDP stream is generated (Units: bits/sec)
duration	Total duration of test (Units: microseconds)
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 (Units: microseconds)
jitter_down	Downstream Jitter measured (Units: microseconds)
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	Please ignore (this is an internal key mapping to unit profile data)
<u>curr_udpl latency.csv</u>	UDP based
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	Target hostname or IP address
rtt_avg	Average RTT in microseconds
rtt_min	Minimum RTT in microseconds
rtt_max	Maximum RTT in microseconds
rtt_std	Standard Deviation in Measured RTT in microseconds
successes	Number of successes (note: use failures/(successes+failures)) for packet loss
failures	Number of failures (packets lost)
location_id	Please ignore (this is an internal key mapping to unit profile data)
<u>curr_ulping.csv</u>	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	Target hostname or IP address
rtt_avg	Average RTT in microseconds
rtt_min	Minimum RTT in microseconds
rtt_max	Maximum RTT in microseconds
rtt_std	Standard Deviation in Measured RTT in microseconds
successes	Number of successes

failures	Number of failures
location_id	Please ignore (this is an internal key mapping to unit profile data)
<u>curr_videostream.csv</u>	
unit_id	Unique identifier for an individual unit
mtime	Time test finished in UTC
target	Hostname of server
downthrpt	This metric records the average downstream throughput for the entire duration of the test. This average is calculated by taking the mean average of the speed of completing each of the blocks downloaded (as specified with the -css client parameter) in bytes/sec
downjitter	A measure of the standard deviation of the speed each block was downloaded at in microseconds
latency	The mean of all of the ping round-trip-times sent from client to server in microseconds
jitter	The standard deviation of the ping round-trip-times sent from client to server in microseconds
buffer_underruns	this metric records how many times the buffer was completely drained (i.e. the client could not receive data at the desired downstream rate, so the buffer emptied). A good connection will have zero buffer underruns.
buffer_delay	The total time in microseconds that the client spent waiting for the buffer to reach its minimum size
buffer_filltime	The total time in microseconds to fill the buffer
duration	Duration of the test in microseconds
bitrate	Stream rate in bytes/sec
buffer_size	Buffer size used in bytes
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)
<u>curr_webget.csv</u>	
unit_id	Unique identifier for an individual unit
mtime	Time test finished in UTC
target	URL to fetch

address	IP address connected to to fetch content from initial URL
fetch_time	Sum of time consumed to download Html content and then concurrently download all resources (Units: micorseconds)
bytes_total	Sum of HTML content size and all resources size (Units : Bytes)
bytes_sec	Average speed of downloading HTML content and then concurrently downloading all resources (Units: bytes/sec)
objects	Number of Resources (images, css etc) 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	Please ignore (this is an internal key mapping to unit profile data)
<u>curr_netusage.csv</u>	

unit_id	Unique identifier for an individual unit
dtype	Time test finished in UTC
wan_rx_bytes	Total bytes received via the WAN interface on the unit (incl. Ethernet/IP headers)
wan_tx_bytes	Total bytes transmitted via the WAN interface on the unit (incl. Ethernet/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	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:

User Terms and Conditions

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, trade marks, 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 licence (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 policies 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 <http://www.samknows.com/broadband/privacy.php>. 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 <http://www.testmyisp.com/paperwork.html>.

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://www.samknows.com/broadband/code-conduct>, was signed by ISPs and other entities participating in the study:



FIXED TESTING AND MEASUREMENT STAKEHOLDERS CODE OF CONDUCT

August 18, 2014

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 2014 Measuring Broadband America Program Fixed Broadband Testing and Measurement and further to use that platform to collect data;

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 to enhance, degrade, or tamper with the results of any test for any individual panelist or broadband provider, except that:
 - 2.1. It shall not be a violation of this principle for broadband providers to:
 - 2.1.1. 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
 - 2.1.2. Address service issues for individual panelists at the request of the panelist or based on information not derived from the trial;

- 2.2. 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 FCC data occurs; and
3. 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 makes a public statement regarding any results of the tests, or except where expressly permitted by the FCC; and
 4. Ensure that their employees, agents, and representatives, as appropriate, act in accordance with this Code of Conduct.

Signatories:

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

Limitation of Liability

The information contained in this document is provided for general information purposes only. While care has been taken in compiling the information herein, SamKnows does not warrant or represent that this information is free from errors or omissions. To the maximum extent permitted by law, SamKnows accepts no responsibility in respect of this document and any loss or damage suffered or incurred by a person for any reason relying on the any of the information provided in this document and for acting, or failing to act, on any information contained on or referred to in this document.

Copyright

The material in this document is protected by Copyright. No part of the materials in this document may be reproduced for any purpose whatsoever without the written permission of SamKnows.

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. 1.1 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. 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. 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

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.

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.

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.

1. 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	✓	✓	✓

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.

Miami, Florida			✓
Mountain View, California			✓
New York City, New York	✓	✓	✓
San Jose, California		✓	
Seattle, Washington			✓
Washington D.C	✓	✓	
Washington, Virginia			✓

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.

1. 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.

1. 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.

1. 4.1 SamKnows test node specifications

All dedicated test nodes must meet the following minimum specifications:

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

1. 4.2 Level3 test node specifications

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

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

- - Disk: 10GB
- - Operating System: CentOS 6 (64bit)
- - Connectivity: 4x1 Gigabit Ethernet (LAG protocol)

1. 4.3 Measurement-Lab test node specifications

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

- - CPU: 2Ghz 8-core CPU
- - RAM: 8GB
- - Disk: 2x100GB
- - OS: CentOS 6.4
- - Connectivity: minimum 1Gbps dedicated upstream

1. 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 1Gbps 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.

1. 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 1Gbps upstream is provided and downstream connectivity to national peering locations

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

1. 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.

1. 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]