

Part B: Collections of Information Employing Statistical Methods:

1. Describe (including a numerical estimate) the potential respondent universe and any sampling or other respondent selection methods to be used. Data on the number of entities (e.g., establishments, State and local government units, households, or persons) in the universe covered by the collection and in the corresponding sample are to be provided in tabular form for the universe as a whole and for each of the strata in the proposed sample. Indicate expected response rates for the collection as a whole. If the collection had been conducted previously, include the actual response rate achieved during the last collection.

The target population for this study is the American household with broadband. The survey will seek reliable national estimates of broadband speed performance for the main US ISPs (listed below).

In response to RFQ-10-000013, the research will provide a reliable national estimate of actual broadband speed performance per ISP for 34 services tiers, and a reliable estimation of performance by ISP packages within geographical region.

Based on recent experience in the UK, SamKnows will develop a panel of 10,000 broadband consumers' panelists covering 15 ISPs, and 34 packages within 4 geographical regions. Broadband speed will be collected through hardware devices placed in consumers' panelists homes.

Internet Service Provider	Type	-	Estimated Subscribers [†]
Comcast	Cable		15,930,000
AT&T	DSL		15,789,000
Time Warner	Cable		9,289,000
Verizon	DSL		9,220,000
Cox	Cable		4,200,000
Charter	Cable		3,062,300
Qwest	DSL		2,974,000
Cablevision	Cable		2,568,000
CenturyLink	DSL		2,236,000

Internet Service Provider	Type	-	Estimated Subscribers [†]
Windstream	DSL		1,132,100
Mediacom	Cable		778,000
Frontier	DSL		635,947
Insight	Cable		501,500
Cable ONE	Cable		392,832
RCN	Cable		312,000
FairPoint	DSL		295,000
Cincinnati Bell	DSL		244,000

[†] Estimated at Q4 2009.

2. Describe the procedures for the collection of information including:

- *Statistical methodology for stratification and sample selection,*
- *Estimation procedure,*
- *Degree of accuracy needed for the purpose described in the justification,*
- *Unusual problems requiring specialized sampling procedures, and*
- *Any use of periodic (less frequent than annual) data collection cycles to reduce burden.*

Building on its experience of completing a similar project in the UK, SamKnows will recruit a US-based panel and deploy its technology throughout all fifty states to cover all regions.

Sample quotas will be set for 4 regions, West, Midwest, Northeast and South. The quotas, defined by geography, technology and service level, have been named buckets.

Within each buckets, soft quota will be set on density: urban and rural population populations consistent with Census Bureau definitions.

SamKnows will use best practice in panel recruitment, which will be open to audit by 3rd party at anytime during the project. The panel recruitment will consist of two steps, using a multi mode recruitment effort to build a large pool of volunteers. Volunteers will be encouraged to sign-up to a project specific website (www.testmyisp.com), providing an initial panel of volunteers. This will be supplemented by more targeted participants who will be recruited with the support of the ISPs, who will target specific customer groups with mailings. This will provide flexibility in building the panel to ensure geographical representation, excluding outliers and, as much as possible,

maximizing the multi mode recruitment selection to avoid pitfalls such as selection bias, coverage error, and panel attrition.

SamKnows will report 'estimates' when accurate enough to be useful, accuracy is reflected by sample size and variance. In order to reflect the limitation of the sample SamKnows will show the speed results as an interval using the 95% interval around the mean. With a sample of up to 500 national wide panelists, the associated sample error is 4.4% and from experience the confidence interval for broadband speed interval is expected to be less than 0.5 Mbit/s.

It is to note that estimates for sub groups – including buckets – will be subject to higher sampling error margins based strictly on the smaller number of panellists in a specific subgroup e.g. for each bucket - 125 panelists the sample error is 8.8% and the expected confidence interval is around 0.8 Mbit/s.

Weighting

To insure representativeness of the US broadband population, and to obtain national average speed, the results will be weighted by density: urban and rural populations consistent with standard Census Bureau definitions. Please note that we need to find the penetration data by region from third parties.

To compare ISPs' performances for DSL, SamKnows would normalise the data by copper loop length (sometimes called "last mile").

Normalisation by loop length is critical for DSL, because with this technology, speed degrades as the length of the "last mile" copper line to the premises increases. Therefore operators that have a higher proportion of customer in rural areas where loop length is typically higher, may be expected to deliver lower speed than those who focus on towns and cities because they have different customer profile. In order to normalise the distance, we will use the "last mile" loop length .

A weight adjustment will be calculated to the contribution to the average speed made by each respondent based on their "last mile" copper loop length by matching the percentage observed in distance bands to the percentage in the distribution of the total sample distribution.

To reduce burden on the panelist, information technology is used extensively, all data collection will be automated after the initial installation of a hardware device. The speed and performance will be monitored through the hardware devices in consumers' homes, to accurately measure the performance of fixed line broadband connections based on real-world usage. These hardware devices are controlled by a cluster of servers, which host the test scheduler and reporting database. The data is collated on the reporting platform and accessed via a reporting interface and secure FTP.

3. Describe methods to maximize response rates and to deal with issues of non-response. The accuracy and reliability of information collected must be shown to be adequate for intended uses. For collections based on sampling, a special justification must be provided for any collection that will not yield "reliable" data that can be generalized to the universe studied.

Once the volunteer panel is recruited and the SamKnows 'Whiteboxes' deployed, the tests will run on a pre-configured schedule, subject only to changes in the schedule and the volunteers own use of their broadband connection. Therefore, 'response rates' will always be at a 'maximum'. Details on the proprietary framework are provided below:

Test Scheduling

Tests will be run every hour on a randomized timing in the hour - 24 hours a day, 7 days a week. Within a month period, the software within the white box unit will perform over 700 separate speed tests. The unit is installed immediately in front of their home internet connection. This ensures that tests can be run at any time, even if all home computers are switched off.

Testing locking

No two tests may operate simultaneously. The first test to begin will have priority, and the other will block until the first finishes (or a timeout is reached).

Threshold Manager

Both before and during testing, the Threshold Manager analyzes the UDP & TCP packet data going over the WAN interface on the unit to check whether the Internet connection is in active use. The Threshold Manager is configurable by design, but is set by default at 400Kbps. If this threshold is breached before or during the tests, activity is suspended and the threshold test is then repeated for one minute, for a maximum of 5 times until either the amount of traffic returns to a level below the threshold or the tests are suspended and that time period is marked as having a busy line.

Order of tests

The tests are run in the following order:

- ping - for all target hosts serially
- dns - for all target hosts serially
- www - for all target hosts serially
- single-threaded http get
- multiple-threaded http get

- single-threaded http post

Required number of data points

For reporting purpose, the data will be aggregated first per hour for the reporting period and then overall to minimize the effect of missing data. If less than 5 data points are recorded for one time slot within a month, the data will be discarded.

Tests

Web browsing

Measures the time taken to fetch the HTML and referenced resources from a page of a popular website. This test does not test against centralised testing nodes; instead it tests against real websites, ensuring that content distribution networks and other performance enhancing factors may be taken into account.

By default, each Whitebox will test three common websites on every test run. The time taken to download the resources, the number of bytes transferred and the calculated rate per second will be recorded. The primary measure for this test is the total time taken to download the HTML page and associated resources.

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

The test may optionally run with or without HTTP headers advertising cache support (through the inclusion or exclusion of the "Cache-Control: no-cache" request header).

Video streaming

This generic streaming test can be configured to model the characteristics of a variety of voice and video protocols. For the purpose of the video streaming test, the intention is to simulate an end user viewing a streaming video over one of the many popular websites that provide this service (e.g. YouTube).

The test operates over TCP and uses a proprietary client and server side component. The client and server negotiate the test parameters at the start of each test.

A playout buffer is configured and the client will attempt to download data from the server at the maximum rate necessary to ensure that this buffer is never empty. A separate thread is reading data from this buffer at a fixed rate, looking for buffer underruns (which would manifest themselves to users as a pause in video). The client will record the time to initial buffer, the total number of buffer underruns and the total delay in milliseconds due to these underruns.

It is expected that the bitrate of the streaming will vary according to the access line speed being tested.

Voice over IP

This test utilises the same generic streaming test as the video test, albeit with different configuration. The test operates UDP and, unlike the video streaming test, utilises bi-directional traffic.

The client initiates a UDP stream to the server and a fixed-rate stream is tested bidirectionally. A de-jitter buffer of 25ms is used to reduce the impact of jitter. The test measures this disruption by monitoring throughput, jitter, delay and loss. These metrics are measured by subdividing the stream into blocks, and measuring the time taken to receive each block (as well as the difference between consecutive times).

By default, the test uses a 64kbps stream with the same characteristics and properties (i.e. packet sizes, delays, bitrate) as the G.711 codec.

Jitter is calculated using the PDV approach described in section 4.2 of RFC5481.

Availability Test

Measures the availability of the network connection from the Whitebox to multiple target test nodes by sending and receiving TCP segments to a receiving server located on each test node.

The client establishes long lived TCP connections to the server on each test node, periodically sending TCP packets containing:

- A Magic number to enable the server to differentiate multiple clients
- Send timestamp (microseconds)

The server echoes back the same data to the client and if it fails to respond or the connection is reset via TCP RST or FIN then the client will attempt to re-establish the connection. If the client is unable to re-establish the connection to all 3 servers simultaneously, it is inferred that Internet connectivity is at fault, the test records a failure locally, along with a timestamp to record the time of failure.

To aid in diagnosing at which point in the route to the target test nodes the connectivity failed, a traceroute is launched to all target test nodes, the results of which are stored locally until connectivity is resumed and the results can be submitted.

This test is executed when the Whitebox boots and runs permanently as a background test.

UDP Latency and Packet Loss

Measures the round trip time of small UDP packets between the Whitebox and a target test node. Each packet contains consists of an 8-byte sequence number and an 8-byte timestamp. If a packet is not received back within three seconds of sending, it is treated as lost. The test records the number of packets sent each hour, the average round trip time of these and the total number of packets lost.

As with the availability test, the test operates continuously in the background and will perform multiple randomly distributed tests within a one hour period.

Speed Tests

Measures the download and upload speed of the given connection in bits per second by performing single and multi-connection GET and POST HTTP requests to a target test node.

Binary non-zero content herein referred to as the payload is hosted on a web server on the target test node. The test operates for a fixed duration (5 seconds by default). The client will attempt to download as much of the payload as possible for the duration of the test. The payload and all other testing parameters are configurable and may be subject to change in the future.

Four separate variations on the test are supported:

- Single connection GET
- Multi connection GET
- Single connection POST
- Multi connection POST

Note that only the multi connection tests are currently intended for use in the FCC project.

Each connection used in the test counts the numbers of bytes of the target payload, transferred between two points in time and calculates the speed of each thread as Bytes transferred / Time (seconds).

Factors such as TCP slow start and congestion are taken into account by repeatedly downloading small chunks (default 256KB) of the target payload before the real testing begins. This “warm up” period is said to have been completed when three consecutive chunks were downloaded at the same speed (or within a small tolerance (default 10%) of one another). In a multi connection test, three individual connections are established (each on its own thread) and are confirmed as all having completed the warm up period before timing begins.

Content downloaded is output to /dev/null or equivalent (i.e. it is discarded), whilst content uploaded is generated and streamed on the fly from /dev/urandom.

4. Describe any tests of procedures or methods to be undertaken. Testing is encouraged as an effective means of refining collections of information to minimize burden and improve utility. Tests must be approved if they call for answers to identical questions from 10 or more respondents. A proposed test or set of test may be submitted for approval separately or in combination with the main collection of information.

The SamKnows tests and methodology have been evaluated and approved by Government Regulators, Internet Service Providers, and leading Academics and researchers.

The stated service level name/identifier, headline speed, price and download limit will be used to check and assign panelist to the correct broadband packages. Different tests and checks will be held to confirm ISP and flag incorrect packages in the following manner:

At the recruitment stage

1. The ISP allocation will be validated using customer's IP against those held in ISP IP libraries. If inconsistent the record is flagged
2. Actual speed will be validated in relation to package and distance from premises to exchange distribution. Outliers will be excluded
3. Exclusion of household when distance to exchange is untypical to maximize the effective sample size after normalization

Test on continuous basis

4. Panelists have the ability to update their information through SamKnows web interface
5. Automatic flag of change of ISP and geographical location through the reporting system
6. Automatic flag of untypical maximum download speed

Independent Review: As there are a large number of stakeholders in this project, SamKnows has dedicated resource to liaising with interested 3rd parties who wish to participate in validating/enhancing the SamKnows methodology. These discussions are open to anyone who wishes to participate.

5. Provide the name and telephone number of individuals consulted on statistical aspects of the design and the name of the agency unit, contractor(s), grantee(s), or other person(s) who will actually collect and/or analyze the information for the agency.

Walter Johnston
Federal Communications Commission
Email: walter.johston@fcc.gov
Phone: 202-418-0807

Nathalie Sot
Statistician, SamKnows
Email: nathalie@samknows.com

Cell: +44 7715 484 803