

B. Data Collection and Analysis Methodology

i. 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 768 kbps 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 servers (to catch tests to ISP test nodes). Removed measurements against any M-Lab server outside of the U.S. Removed measurements against any M-Lab 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 servers were unavailable).
4. Remove 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).
5. Made any other necessary adjustments such as the removal of the Netflix web site load time measurements, which was necessary as Netflix changed their home page to default to SSL in late March 2011.

ii. 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 March 2011 test results. In addition, we dropped the top and bottom 1% of measurements to control for outliers that may have been anomalous or otherwise misrepresentative of actual broadband performance. All statistics were computed on the trimmed data.²⁷

We charted data only when at least 25 data points were available and noted instances of 30 or fewer data points.

The resulting final sample of data for March 2011 was 6,851 participants.

iii. 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.

iv. 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. Web browsing measurements were taken between the panelist's home gateway and ten 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 ten large websites configured to serve high traffic loads may have mitigated the effects of temporary congestion.

v. Traffic Shaping Not Studied

The effects of traffic shaping is not studied in this 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 main report that show performance over 24-hour periods, where tested rates for some ISPs and service tiers flatten for periods at a time.

vi. Analysis of PowerBoost and Other 'Enhancing' Services

The use of transient speed enhancing services 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 of a transfer. Once the "PowerBoost window" is

²⁷ These methods were reviewed with statistical experts within the FCC and by participating ISPs.

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 services such as “PowerBoost,” 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. PowerBoost will have a dominant effect on short speed tests: a speed test running for 2 seconds on a connection employing PowerBoost would likely record the PowerBoost rate, whereas a speed test running for 2 hours will reduce the effect of PowerBoost to a negligible level.

The speed test employed in this study isolated the effects of transient performance enhancing services such as PowerBoost 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 *sustained* speed. Testing was conducted prior to the start of trial to estimate the length of time during which PowerBoost effects might be seen. Even though the precise parameters used for PowerBoost-style services are not known, their effects were no longer observable in testing after 20 seconds of data transfer.

vii. 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 context of the other technical factors addressed by the testing methodology, a delay of 5 ms per 1000 km of distance traveled can be attributed solely to the speed of light. 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.

ix. Limiting Factors

A total of 4,281,635,408 measurements were taken across 179,913,691 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.

