



THE UNIVERSITY OF CHICAGO

DATA SCIENCE
INSTITUTE

Analysis of Internet Performance Data from Household Devices

By the Internet Equity Initiative (IEI):

Marc Richardson (mtrichardson@uchicago.edu)

Guilherme Martins (gmartins@uchicago.edu)

Taveesh Sharma (taveesh@uchicago.edu)

Alexis Schrubbe (schrubbe@uchicago.edu)

Introduction

This report presents the analysis of data collected from devices installed in **19** households to measure internet performance. The devices were deployed for a period of three weeks, during which approximately **586.636** measurements were taken per device. In total, the devices were running for over **500** hours, providing a comprehensive dataset for analysis.

Please note that there are missing measurements for some of the 19 devices. This is because some measurements, namely Ookla Speedtest and RPM, were disabled for some of the devices due to consistent appearance of errors in the logs due to extremely low bandwidth and high-latency.

Main Findings

Bandwidth Performance

- **Bandwidth Benchmarks:** Devices ran bandwidth measurements with Ookla Speedtest and Measurement Lab's NDT7 tool.
 - Out of the 19 households, approximately 6 reported bandwidth speeds below the threshold of 25/3 Mbps. This accounts for approximately 32% of the total households sampled.
 - Out of the 19 households, approximately 10 reported bandwidth speeds below the threshold of 100/20 Mbps. This accounts for approximately 53% of the total households sampled.

- o See **Figures 1, 2, 3 and 4** in the Appendix for box plots showing the distribution of the bandwidth tests for devices with speeds close or below the federal minimums for unserved and underserved speeds.
- **Bandwidth Variability:** By taking measurements over time, we can see fluctuations in bandwidth of a household's connection. Stable connections where bandwidth does not fluctuate may see better end-user performance than unstable connections where bandwidth is constantly fluctuating.
 - o Many of the households we assessed show significant variability in bandwidth measurements. Table 1 in the Appendix shows a ranking of each household by the variability in bandwidth measurements over the collection period. **Higher values in this table are associated with more unstable connections for both upload and download bandwidth as measured by Ookla and NDT7 tests.**
 - o The boxplot Figures 1–4 in the Appendix also highlight there is a wide distribution of measurements among households in the sample.

Latency Metrics

- **Poor Last-Mile Latency:** Three devices exhibited poor latency (above 100ms) in the last-mile connection (Figure 7). Poor latency impacts the user experience for real-time applications like Zoom.
- **Latency to Major Destinations:** The same three devices had poor overall latency to other major internet destinations such as Google and Facebook.
- These three devices have satellite connections to either HughesNet or Viasat.

Connection Stability

- **Unstable Connections:** The three devices connected via satellite also reported unstable connections, characterized by data outages or inconsistent speed measurements.

Round-trips Per Minute (RPM Score)

- **RPM** is measured by the *goresponsiveness* tool and it generates a score representative of how responsive the user's internet connection is under working conditions. The greater the value is, the better is the responsiveness of the connection under working conditions (1):
- **The results can be found at the Appendix, Figure 8.**

- **Analysis:** Some devices were simply unable to produce any results, in particular devices with a satellite connection. That is due to the fact that *goresponsiveness* tends to be aggressive in its way to assess the performance (high packet-per-second rate) as compared with other standard measurements as satellite connections are fragile. The results don't seem to correlate with other measurements and we do not have a subjective assessment of internet user's quality of experience that could be used to establish a relationship between RPM and QoE. However, it should be noted that the RPMs observed for the households that ran the measurement are low (a typical good measurement lands on a 150-250 range for p90), which could indicate poor user experience. More analysis is needed to conclude anything about QoE from these measurements. IEI is currently working to have a reference table for RPM Score results that can later be used to compare these results.

Number of Connected Devices

- Netrics also measures the number of devices that are using the Internet connection of a household's network. This measurement can give us a sense of usage on the network. The more users that are connecting through a single connection, the more likely it is that there will be issues with the network due to overuse. We want to look particularly for households that have a high number of connected devices but little bandwidth.
- Table 2 shows the measured download bandwidth from the Ookla tool, divided by the average number of devices connected. Devices at the top of that list have numbers below 1 Mbps / device. For a reference, a typical zoom call at 720p resolution can take 1.2 Mbps of bandwidth just for download.

Comparative Analysis with Chicago Data

When compared to data from other devices deployed in Chicago, the 19 devices in this study showed:

- The probability for a Marion device registering an upload speed of ≥ 10 Mbps or a download speed of ≥ 200 Mbps is lower than that for a Chicago device.
- Hughesnet devices consistently show a lower normalized upload throughput reduction than the majority of devices (out of 51) subscribed to Comcast in Chicago. Protek shows a higher reduction than most Comcast devices.
- Hughesnet devices consistently show a lower normalized download throughput reduction than the majority of devices (out of 51) subscribed to Comcast in Chicago. Protek shows a higher reduction than most Comcast devices.
- The normalized reduction in upload speed for a Marion device is higher than a Chicago Comcast device for 76% of the total device pairs.

- o In terms of Normalized reduction, Protek shows a lower value than only < 20% of AT&T devices in Logan Square, Chicago. Hughesnet again shows a very stable speed profile.
- o The normalized reduction in download speed for a Marion device is higher than a Chicago Comcast device for 85% of the total device pairs.

Conclusion

The data collected from the 19 household devices provides valuable insights into the state of internet performance in the targeted areas. Notably, there is a significant percentage of devices with below-average bandwidth and poor latency, which warrants further investigation and potential policy intervention. The comparative analysis with Chicago data also offers additional context, highlighting specific areas where targeted efforts could yield improvements.

(1) RPM (Round-trips Per Minute) resources:

- (a) <https://www.ietf.org/staging/draft-cpaasch-ippm-responsiveness-rpm-01.html>
- (b) <https://github.com/network-quality/goresponsiveness>

Appendix: Data Visualizations and Data Tables

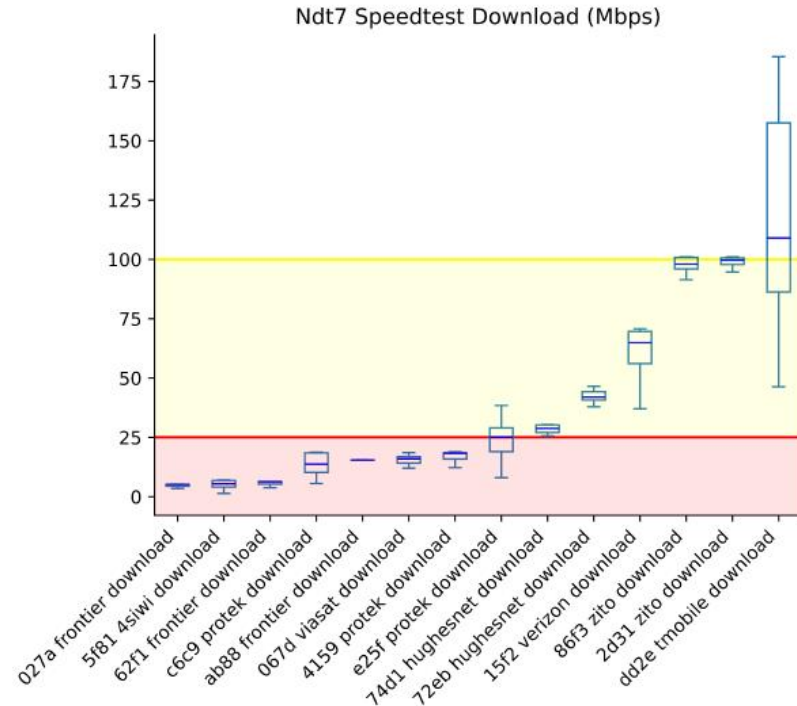
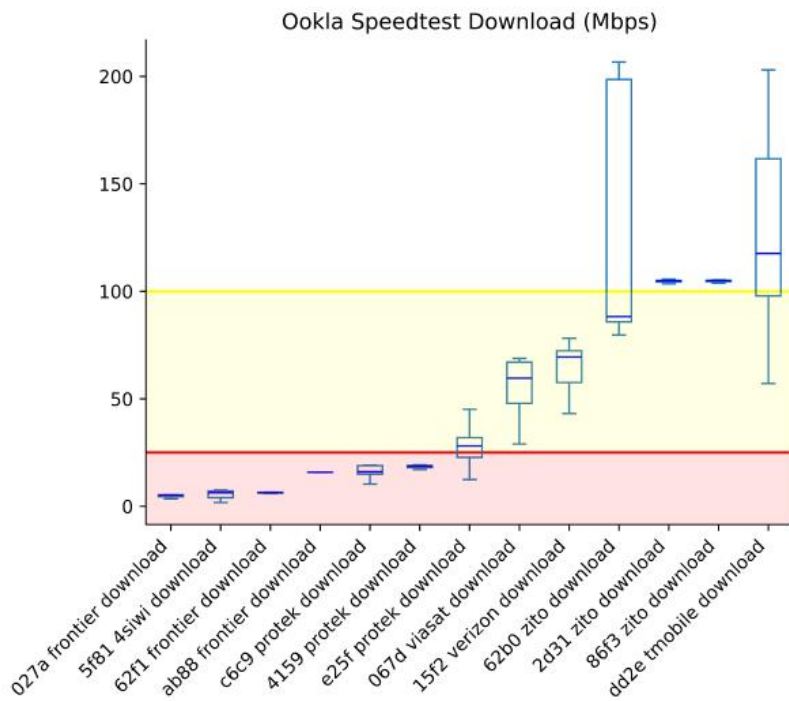


Figure 1 and 2. Ookla and Ndt7 Download Distributions Per Household

Notes: Households are ordered by median.

The 100Mbps and 25Mbps are the federal government definitions of “underserved” and “unserved” respectively. Data is filtered to only show devices median less than 120 Mbps for better visualization.

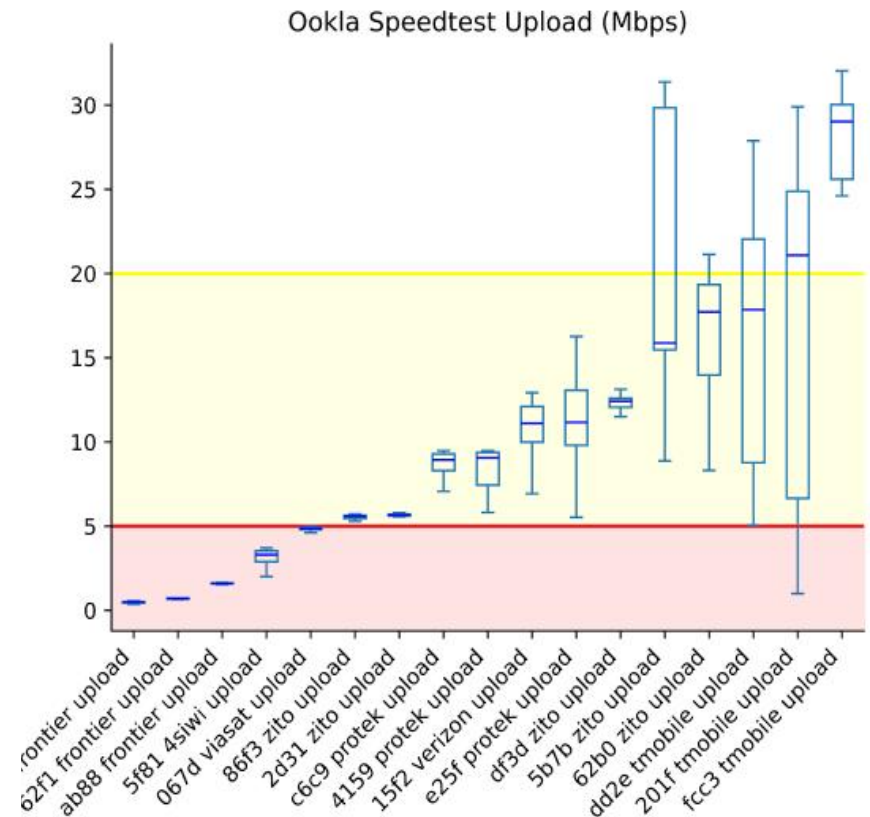
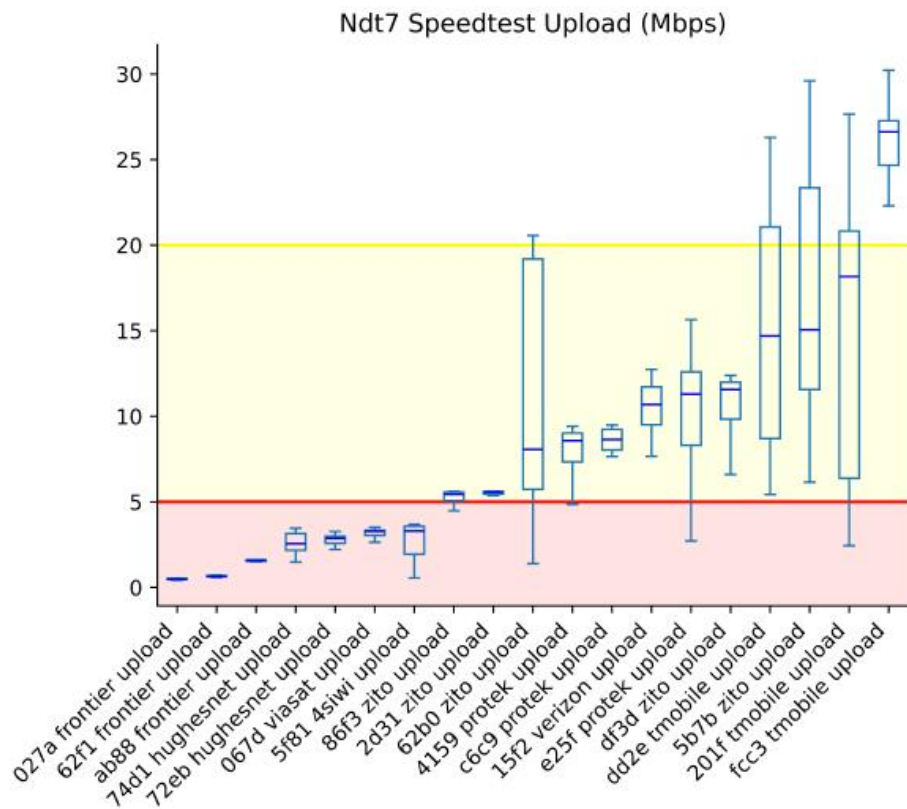


Figure 3 and 4. Ndt7 and Ookla Upload distributions per device, ordered by Median (All devices).

| deviceid | isp | Ookla NR Download | Ookla NR Upload | Ndt7 NR Download | Ndt7 NR Upload | NR Index |
|---------------------------|----------|-------------------|-----------------|------------------|----------------|----------|
| nm-mngd-20231005-a4f4ab88 | frontier | 0.044022 | 0.046335 | 0.016387 | 0.058360 | 0.041276 |
| nm-mngd-20230920-50172d31 | zito | 0.015027 | 0.036816 | 0.060372 | 0.312509 | 0.106181 |
| nm-mngd-20230922-e97686f3 | zito | 0.011780 | 0.106079 | 0.064207 | 0.432877 | 0.153736 |
| nm-mngd-20230922-fbd4df3d | zito | 0.029915 | 0.126064 | 0.256445 | 0.436244 | 0.212167 |
| nm-mngd-20230920-2bebfcc3 | tmobile | 0.384889 | 0.213743 | 0.376879 | 0.221723 | 0.299308 |
| nm-mngd-20230922-e875067d | viasat | 0.500365 | 0.075894 | 0.394831 | 0.343695 | 0.328696 |
| nm-mngd-20230920-944115f2 | verizon | 0.321545 | 0.395418 | 0.405054 | 0.321191 | 0.360802 |
| nm-mngd-20230920-c5a662f1 | frontier | 0.470624 | 0.344552 | 0.467931 | 0.164808 | 0.361979 |
| nm-mngd-20230920-b243027a | frontier | 0.604410 | 0.302444 | 0.464222 | 0.141754 | 0.378208 |
| nm-mngd-20230925-52194159 | protek | 0.320948 | 0.370789 | 0.393359 | 0.455420 | 0.385129 |
| nm-mngd-20230925-f5e4c6c9 | protek | 0.316524 | 0.266252 | 0.614585 | 0.415428 | 0.403197 |
| nm-mngd-20230925-0945e25f | protek | 0.445206 | 0.568115 | 0.547833 | 0.573326 | 0.533620 |
| nm-mngd-20230922-e66f5b7b | zito | 0.515349 | 0.599151 | 0.467201 | 0.637879 | 0.554895 |
| nm-mngd-20230920-8aee5f81 | 4siwi | 0.637918 | 0.426360 | 0.562835 | 0.655905 | 0.570754 |
| nm-mngd-20230920-d59f62b0 | zito | 0.589458 | 0.450708 | 0.524478 | 0.843942 | 0.602146 |
| nm-mngd-20230920-db41dd2e | tmobile | 0.597359 | 0.707804 | 0.571461 | 0.685776 | 0.640600 |
| nm-mngd-20230925-fdef201f | tmobile | 0.724499 | 0.822334 | 0.770984 | 0.808295 | 0.781528 |

Table 1. Devices list sorted by Normalized Reduction Index : lower numbers (lighter shade) indicate more consistent performance (less variability). NR Index is the average result of Ookla/Ndt7 Up/Down. More on Normalized reduction can be found at: <https://internetequity.uchicago.edu/wp-content/uploads/2022/08/benchmarks-or-equity-AUG2022.pdf>

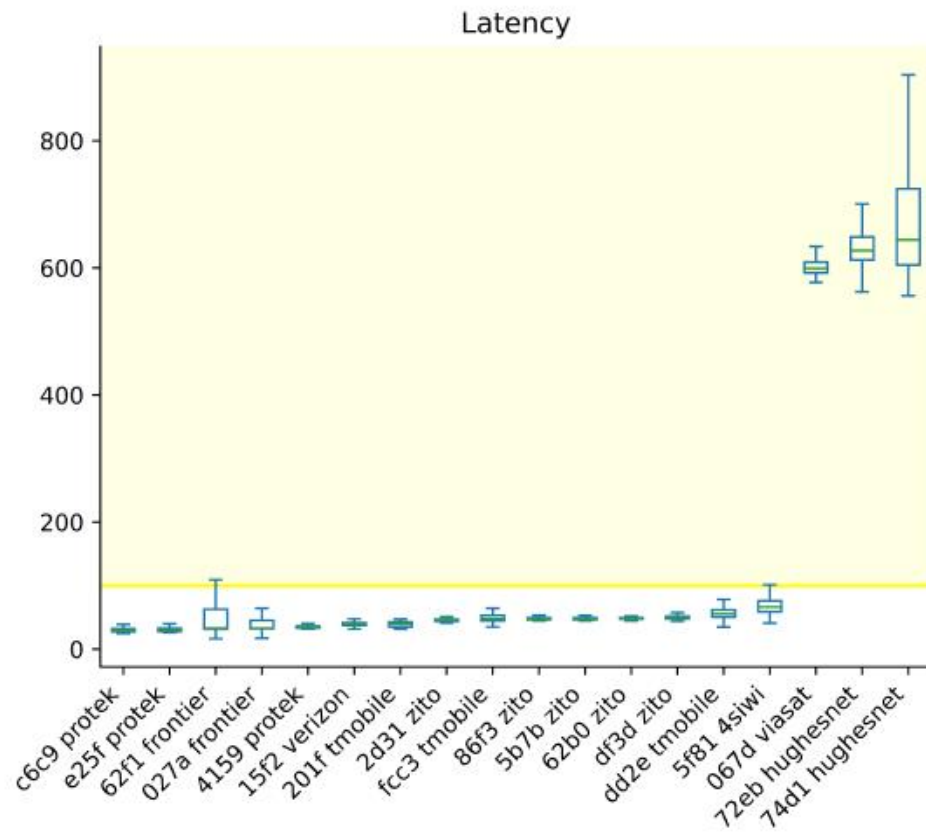


Figure 7. Latency to Google distribution for all devices in 62922. Satellite links have high latency.

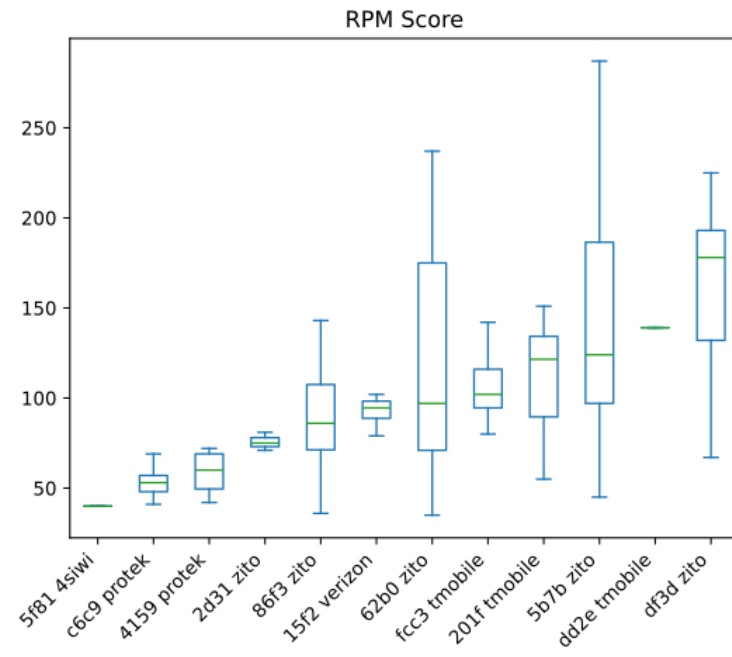


Figure 8. RPM Score for the deployed devices in 62922

| | deviceid | isp | bw_conn_dev |
|----|---------------------------|----------|-------------|
| 0 | nm-mngd-20230920-c5a662f1 | frontier | 0.484349 |
| 1 | nm-mngd-20230920-8aee5f81 | 4siwi | 0.612385 |
| 2 | nm-mngd-20230920-b243027a | frontier | 1.234655 |
| 3 | nm-mngd-20230925-f5e4c6c9 | protek | 1.644013 |
| 4 | nm-mngd-20231005-a4f4ab88 | frontier | 3.289533 |
| 5 | nm-mngd-20230925-0945e25f | protek | 4.112798 |
| 6 | nm-mngd-20230925-52194159 | protek | 4.366550 |
| 7 | nm-mngd-20230922-e97686f3 | zito | 8.616774 |
| 8 | nm-mngd-20230920-db41dd2e | tmobile | 12.954978 |
| 9 | nm-mngd-20230920-944115f2 | verizon | 12.987632 |
| 10 | nm-mngd-20230920-d59f62b0 | zito | 13.410697 |
| 11 | nm-mngd-20230920-50172d31 | zito | 16.454899 |
| 12 | nm-mngd-20230925-fdef201f | tmobile | 20.479887 |
| 13 | nm-mngd-20230922-fbd4df3d | zito | 22.225851 |
| 14 | nm-mngd-20230922-e875067d | viasat | 28.630834 |
| 15 | nm-mngd-20230920-2bebfcc3 | tmobile | 67.890812 |
| 16 | nm-mngd-20230922-e66f5b7b | zito | 152.007018 |

Table 2. Bandwidth available for device, $bw_conn_dev = \text{Ookla Bandwidth} / \text{Average Number of Devices}$