Appendix A:

Methodology for Pre-Challenge Speed Test Modifications

TEAM KENTUCKY®

OFFICE OF
BROADBAND DEVELOPMENT

Introduction

This document identifies sources and requirements for defining acceptable evidence of underserved locations when using crowdsourced test results to conduct a pre-challenge modification using crowdsourced speed test data. The Kentucky Office of Broadband Development will use Ookla Speedtest data to perform a pre-challenge modification.

This document is divided into two parts:

- Part One provides a description of the processes that will be undertaken to identify areas of need using crowdsourced measurements.

- Part Two provides an overview of the methodological rigor behind crowdsourced data, including rebuttals of common misconceptions regarding crowdsourced data.

Part One: Identifying Underserved Locations

Step one: Define criteria for underserved locations

- Areas with speeds lower than 100/20 Mbps are categorized as underserved and are eligible for funding once unserved areas have been addressed.

Note that when classifying an area as underserved, both download and upload speed thresholds are considered. That is, an area will be considered underserved if either the relevant download or upload threshold is not met.

Step two: Filter to exclude non-relevant speed tests

Measurement accuracy is critical to painting a clear picture of performance and to the equitable distribution of funding. The Kentucky Office of Broadband Development will use one year of historical Ookla Speedtest data to conduct this analysis. A full year of data is required because weather conditions and foliage that change throughout the course of a year can impact performance of some technologies. As discussed in steps three and four below, the 80th percentile of best speeds will be evaluated for comparison against reported speeds.

- Remove any speed tests that do not have a GPS-defined location. Most browser-based tests use GeoIP resolved to the centroid of a zip code (or similarly defined area) and do not provide adequate location precision.
- Exclude measurements identified as originating from corporate campuses and other business-only areas.
- Exclude tests that show a GPS location precision of no better than 300 meters in rural areas or 100 meters in urban areas (suggested) as most of these boundaries will allow for these tolerances or higher. In very remote areas with difficult terrain, you may choose to accept location definitions with lower precision.
- Exclude records that show speeds below minimum broadband thresholds and show poor WiFi connections (first traceroute hop latency exceeding 10 milliseconds)
- Exclude records where the testing server was chosen manually.

Step three: Use census blocks as initial evaluation areas

The 2020 census block polygons will be used as the basis for aggregating and evaluating the speed test data, providing the state important census data to best understand and prioritize areas due to economic need, equitable distribution of resources, and other key goals. This approach also offers states the flexibility to correlate service levels with other associated data as they see fit (e.g., demographic insights, household counts, etc.).

- Overlay the speed test points on the map containing the BSL data and the 2020 US census blocks.
- Calculate the 80th percentile speed as well as maximum and median speeds for comparison that are captured within each census block.
- Communities and areas that do not meet the broadband minimum standards will stand out, often in clusters on the map.
- Although individual census blocks will often include enough BSLs and test measurements to stand on their own, many census blocks have very low numbers of both BSLs and speed test measurements.
- If a geographically large census block includes disparate and unrelated areas, a custom polygon can be used to more selectively define the eligible area
- Sparsely populated areas contain fewer data points for most human activity, including speed test measurements — so-called "doughnut holes" are common and contiguous census blocks should be used to judge the area as a whole.

Note: The FCC has used hexagons (H3, level 8) on their national map to visualize broadband availability above *the magnification level where individual BSLs are shown. Aggregating data to these same hexagons can provide an additional step to assist eligible entities who may be trying to compare against FCC-collected data.*

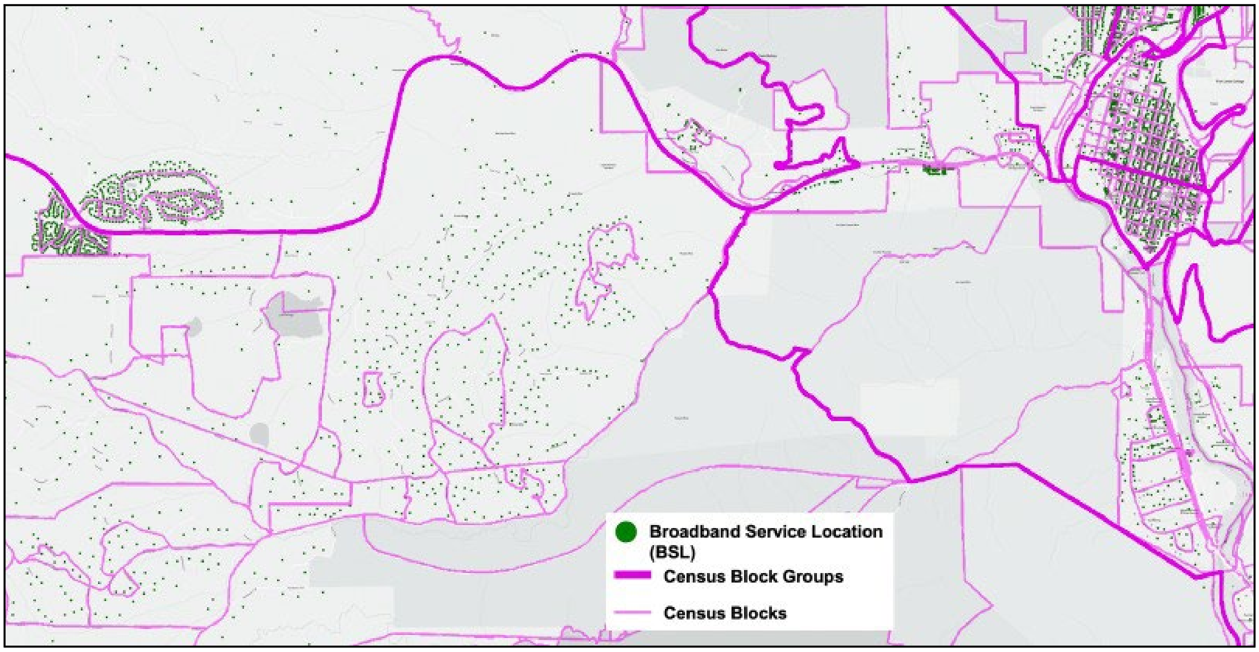

The image above shows BSLs and census boundaries for the sample study area between Durango and Durango West, Colorado.

Step four: Evaluate based on best speed results

Speedtest results include average, median, and best speeds. Each of these has its advantages. For identifying areas of need, however, Ookla recommends using best speeds.

- Best speeds act as a particularly strong indicator of need: if the highest speeds measured are below the 100/20 Mbps threshold, the defined area simply cannot meet the minimum requirements.
    - Even if a small number of tests measure slightly above the minimums, the tested network may not be providing "reliable" service as required in the IIJA legislation.
- The term "best speeds," for this purpose, is defined as the 80th percentile, respectively, for download and upload speeds.
    - An area with speeds above the numbers listed in the definition of "underserved" listed above will be considered ineligible for their respective categories unless qualifying under a separate module.
- This approach also ensures that outlier data points do not exaggerate the performance of the available network(s).
    - Although it is expected that peak internet usage hours can create temporarily slower throughput speeds, hours of lowest usage, such as early morning hours, can also provide an overly optimistic assessment of network performance.

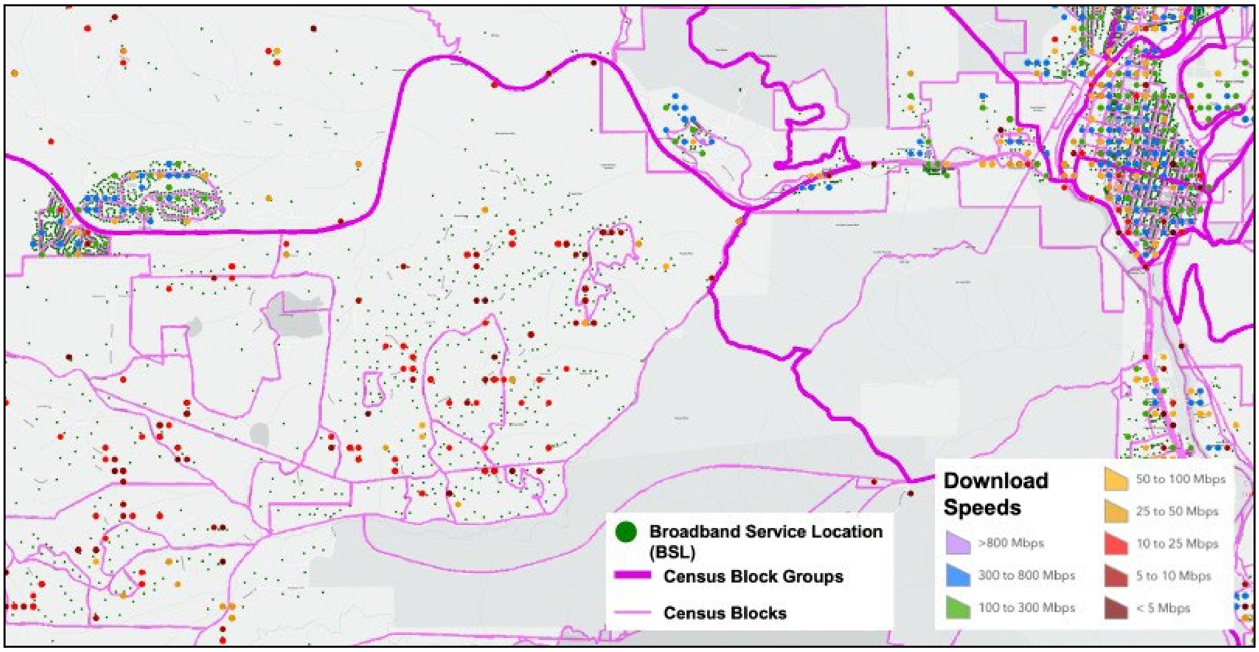

The image above shows Speedtest® results overlaying BSLs and census boundaries–the consistently red and yellow dots indicate areas that are not experiencing sufficient broadband speeds.

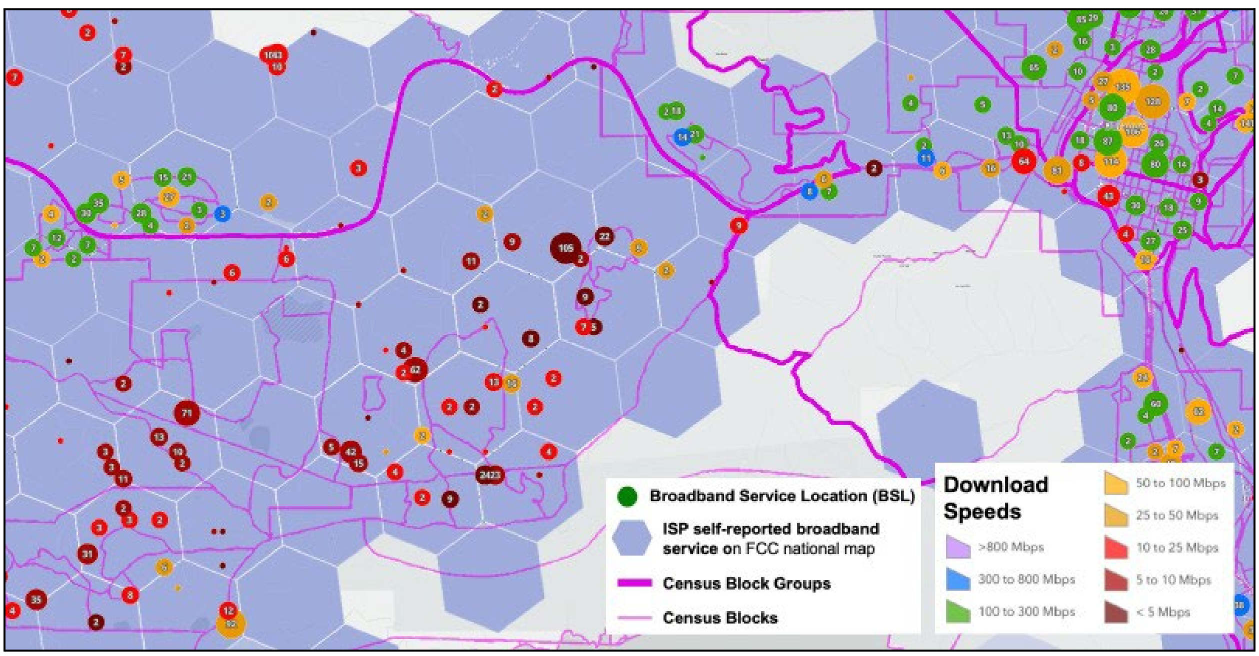

The image above shows clustered Speedtest® results, indicating the number of tests taken within close proximity, overlaying hexagons of reported broadband service availability as depicted in the FCC map.

Step five: Compare with reported speeds and known funding decisions

Once areas have been identified where best speeds at the census block level show service lagging behind the 100/20 Mbps thresholds, compare results with those reported on the FCC national map. The goal is to identify areas where broadband service is reported but evidence indicates it is not available and no other funding has been made available

- If an area has already been funded through another program, it will not be eligible unless separate proof is presented that the responsible party does not plan to build out the area or other evidence can prove they will not be able to complete the build.
- Using best speeds as described above, areas identified with service lagging behind the 100/20 Mbps thresholds will be considered eligible.

Step six: Choose the census blocks that best defines the area of concern

- Demarcate the area of concern by drawing a polygon around the areas identified as eligible for funding.
    - This can include neighboring locations that are not immediately contiguous but can be considered part of the same area.
- NOTE: Recognizing that the national map uses hexagons to define broadband availability, overlaying and comparing the FCC hexagons against the results shown in the census block will likely help facilitate additional evaluation.
    - Offering an H3-R8 hex-based view in addition to a census-block view can help local constituents, municipalities, and nonprofits make apples-to-apples comparisons between state eligibility maps and the FCC visualizations.
    - Note that BSLs and location IDs provide a critical common denominator to help correlate results across the census block to the H3-R8 levels.

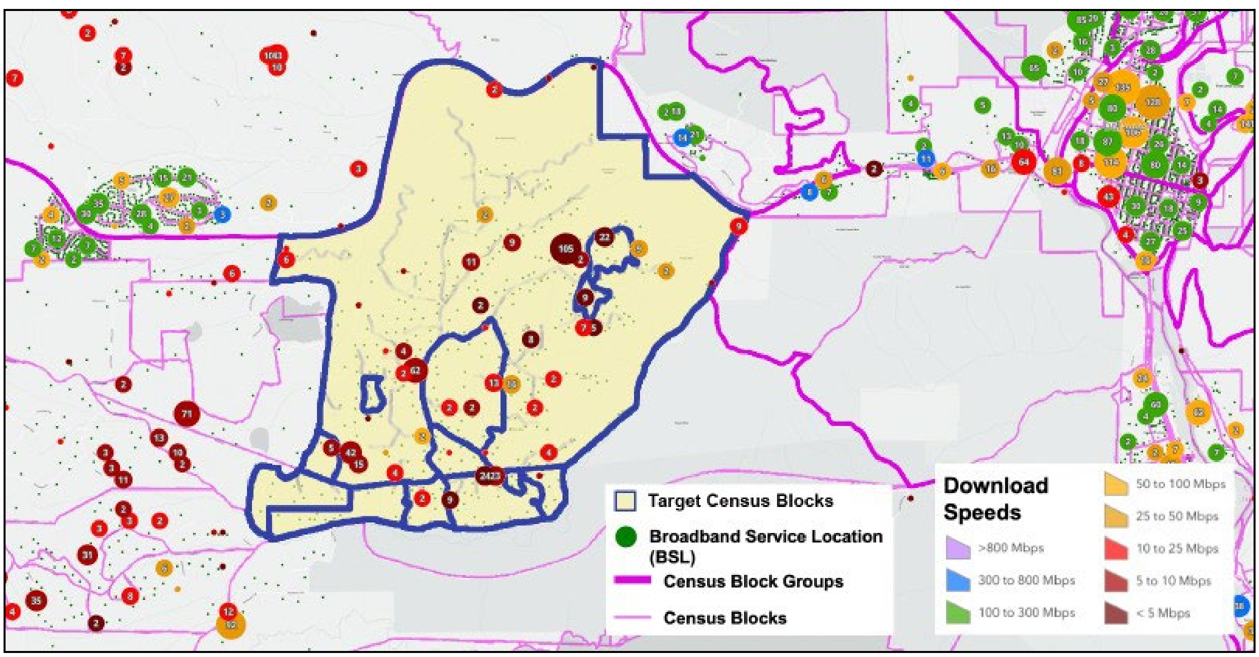

The image above shows clustered Speedtest® results, indicating the number of tests taken within close proximity, overlaying selected census blocks being reclassified as eligible for funding.

Step seven: Define BSLs by area

The challenge process is built around the Location ID for each BSL.

- Once a polygon has been created around the area of concern, the BSLs that lie within that newly defined area will be identified.
- Every BSL Location ID targeted should be captured within the census block that will be considered eligible.

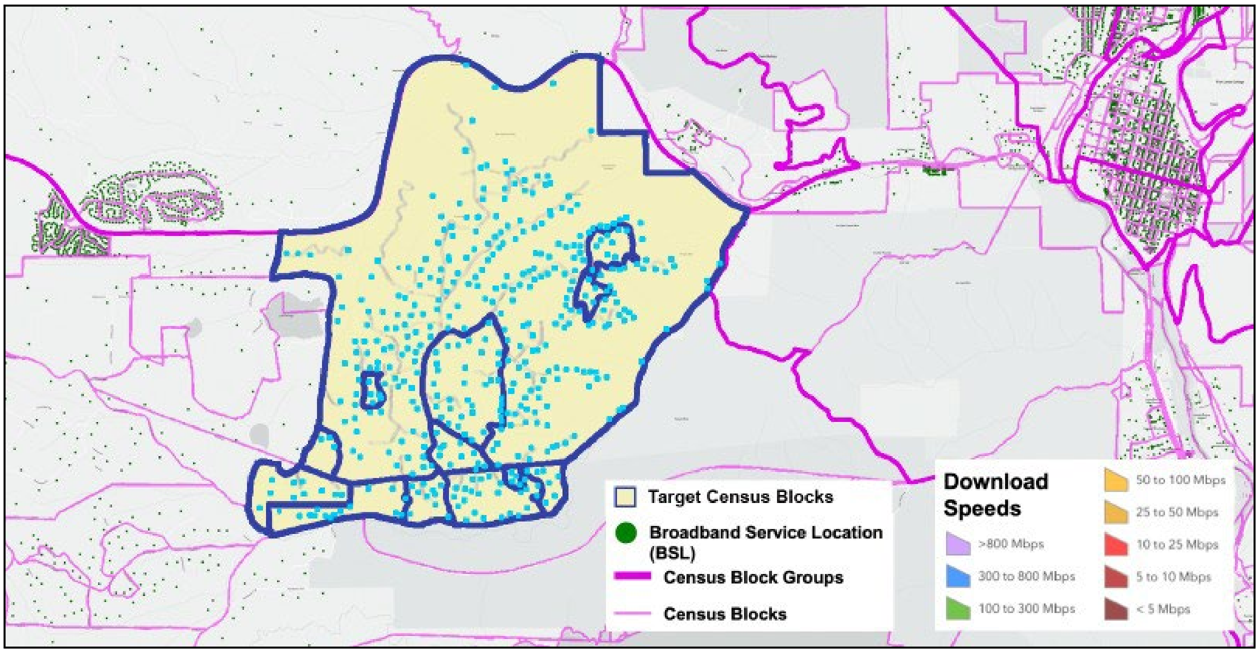

The image above shows the selected BSLs that lie within the selected census blocks being reclassified as eligible for funding.

FAQs (for the step-by-step instructions)

What are the requirements for using crowdsourced data as acceptable evidence of BEAD funding eligibility?

The document "BEAD Model Challenge Process _Final," released by NTIA on June 28, 2023, outlines several accepted methods of challenging broadband availability on the National Broadband Map, maintained by the Federal Communications Commission (FCC).

In addition to the pre-approved methods, section 1.4.7 of the document states, "If the Eligible Entity is not using the NTIA BEAD Model Challenge Process, outline the proposed sources and requirements that will be considered acceptable evidence."

The use of crowdsource data, such as throughput and latency measurements from speed tests, is already included as an acceptable form of evidence in the above-mentioned document. The additional evidentiary requirements to be submitted along with speed tests include service address, personal contact information, date of requested service and more associated with a challenge by, or on behalf of, an individual Broadband Service Location (BSL).

Why is third-party crowdsourced is needed during the challenge process?

Many state and local governments have made significant efforts to drive public engagement of reporting throughput speeds and latency. Some have set up their own speed testing websites using popular tools such as those provided by Ookla and M-Labs. Others have built their own. These state-sponsored collection efforts provide valuable information about the availability of broadband services.

Participation in state-sponsored efforts, however, is often uneven, with an initial spike of interest tied to promotional efforts followed by a sharp decline in citizen engagement. The reasons for this pattern of declining usage vary, but continued promotion efforts from a state to ensure participation is difficult to maintain as other important issues eventually supplant public attention. Moreover, even when individuals do visit these state-sponsored sites, participation is often limited by the fact that individuals are sometimes reluctant to include the personally identifiable information required by the challenge process.

Although valuable, there is simply not enough evidence resulting from these efforts. Measurements gathered from public participation represent a very small fraction of data available compared to that which is offered by leveraging existing crowdsourced datasets. Offering measurement and evidence at scale, crowdsourced data provides key insights into broadband availability and performance not available through other methods of collection.

Why should I use hexagons in addition to census blocks to evaluate service and shape polygons around areas of concern?

Used by both the FCC and NTIA, hexagons have become the de facto system of measurement for the federal entities involved in mapping broadband coverage, challenging reported service levels, and for distributing funds to improve connectivity. Using hexagons along with the associated BSLs will facilitate area challenges and allow for easier ingestion of data into existing systems.

How were Census blocks chosen as the evaluation level?

Identifying underserved areas is a balance between the granularity of zoom level used and the insights that are revealed. In general, a greater density of samples provides increased evidence, offers more indicators of available service levels, and results in greater precision regarding the estimate of cost to service each location. However, sample density varies by area, with fewer results typically available in rural areas than urban areas.

For most areas, the census block level is the sweet spot for Volume 1 of the challenge process, providing the best combination of data density and BSL density: it helps easily identify areas of need and can later be rolled-up as needed into associated hex-8 bins to help the NTIA efficiently review results that can be correlated with the existing National Broadband Map and platform.

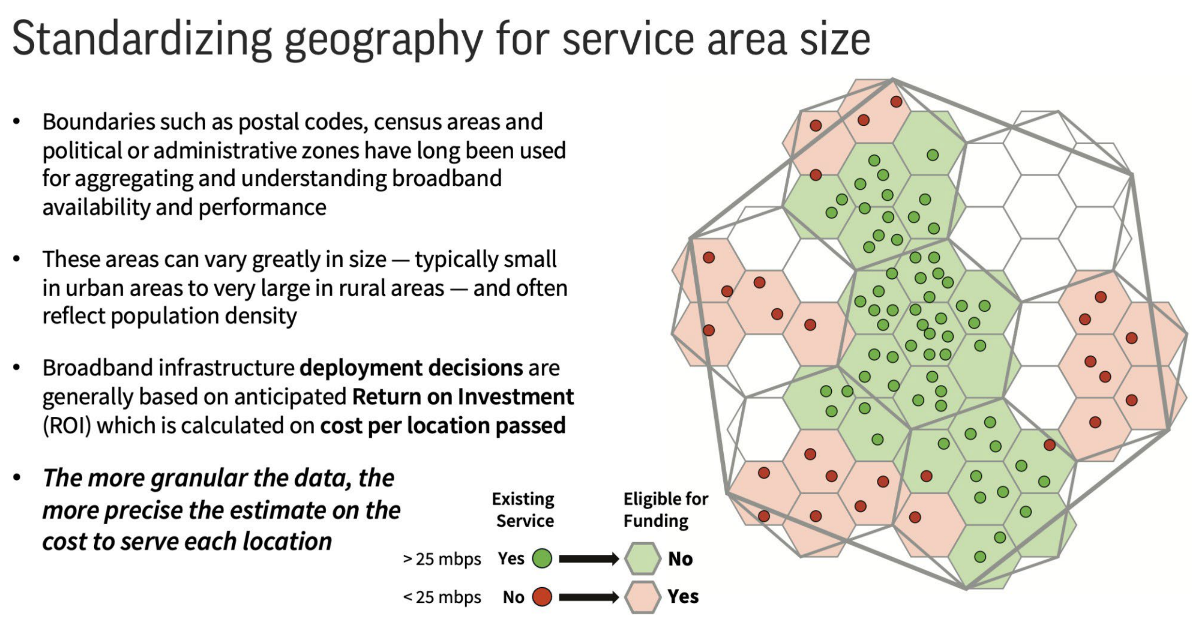

Why is location accuracy important to the challenge process?

While all crowdsourced data, regardless of source, can provide valuable information on the state of broadband across America, sources that include geographically-precise location data allow for greater defensibility when identifying areas of need. While crowdsourced data in general can provide important views of broadband availability and performance, Ookla Speedtest® data also includes GPS-accurate location data that is a key component for effectively challenging the National Broadband Map.

Browser-based tests typically do not include GPS precision and instead rely on GeoIP data that is resolved to the centroid of the nearby zip code or other similar boundary set. Because this process results in measurements that lack the required location precision, Ookla recommends filtering to only use tests captured by GPS-enabled iOS or Android devices.

Ookla further recommends filtering to include results with location accuracy set at 300 meters or better. The level-8 hex bins that the FCC and NTIA are familiar with have a diameter roughly 1 kilometer across (a bit more or less between minimum and maximum distances). Filtering to include tests with an accuracy of 300 meters or better should therefore offer acceptable location accuracy when performing an area challenge and offers a unit of measurement tied to the same level-8 hexagons relied on by the funding authorities.

However, in rural areas, where sample density often trails what is found in more populated areas, location accuracy can potentially be expanded to filter for all results at 500 meters or better. This 500 meter limit is still within the average Hex-8 edge length. This is particularly valid if the area being investigated is constituted of multiple hexagons that may represent several square kilometers. This should be particularly true if the area includes multiple neighboring hexagons.

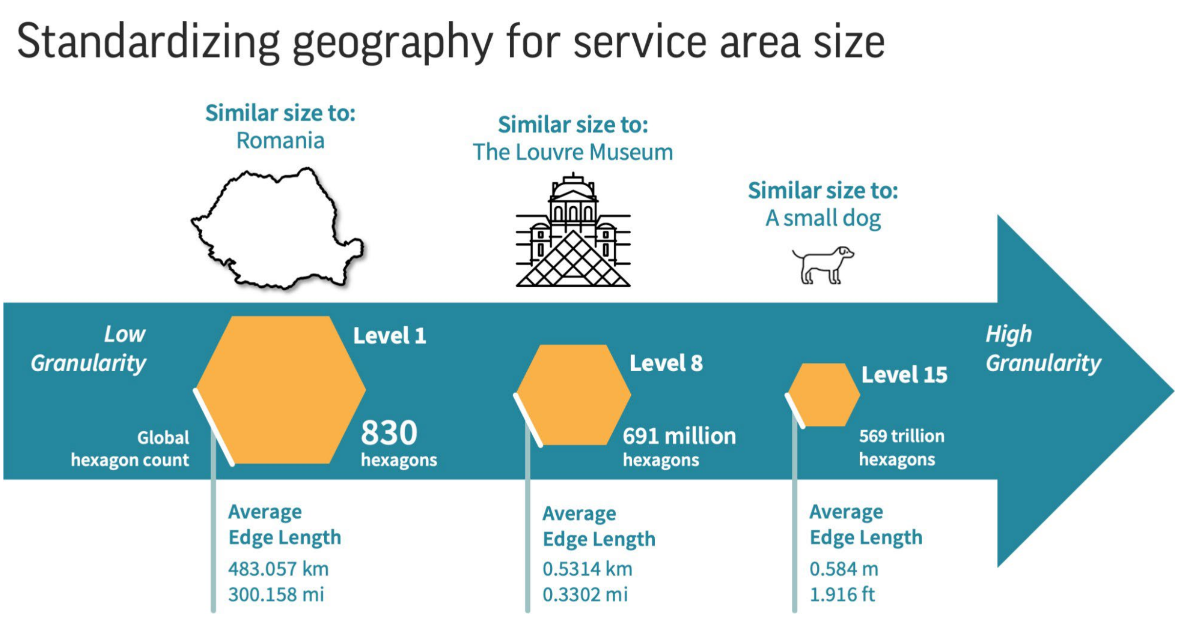

How does the Kentucky Office of Broadband Development determine an acceptable ratio of speed test measurements to Broadband Serviceable Locations (BSLs)?

At the census block level, a minimum of ten total tests from at least five unique users are required.

How and why should the 75th percentile be calculated?

This section is yet to be written. Will include explanation of outliers and anomalies.

Can speed testing prove that no service exists in a particular area?

Identifying areas completely devoid of service can be particularly difficult. After all, testing cannot be completed if service is unavailable. In this scenario, the challenge process is asking for evidence to prove a negative, and this creates a catch-22: how can you show a crowdsourced test that proves service is nonexistent if individuals cannot complete a testing precisely because service is nonexistent?

In these cases, Ookla recommends looking for areas in which clearly poor crowdsourced results (i.e., those in which max speeds are below the 25/3 or 100/20 Mbps thresholds) create a ring or rough perimeter around locations in which zero additional test results are found. This is likely a situation in which poor service at the edges has degraded into complete lack of service as you move farther along. Though not conclusive, using crowdsourced results in this way is similar to finding evidence of a black hole by looking for where light is expected but no longer exists.

Note also that the FCC National Broadband Map commonly shows isolated hexagons where only a single BSL may exist with no others nearby. This is expected where population density is very low. Examples can include locations with difficult terrain or more arid farmlands where properties can be measured in thousands of acres. These types of BSLs will often be categorized as high-cost locations.

Why are propagation models of service availability and performance inadequate for understanding real-world performance?

Providers often use propagation models that use mathematical models to estimate broadband coverage. While helpful as a first-step for planning, propagation models often do not paint a true picture of coverage and performance because they do not adequately take into consideration elements that impact signal strength and signal travel: for instance, dense foliage or changes in terrain can interrupt signals and lead to an experience that is worse than a propagation model estimates.

Part two: Crowdsourced data methodology overview

Controlling for variables and common misconceptions regarding crowdsourced data

Distributing $42 billion in funding is not a trivial task. The challenge process should offer a fair and equitable avenue to ensure funding decisions are backed by data and objective in nature. Making data-backed decisions is of course ultimately reliant on the quality of data used. Crowdsourced data offers a readily-available, peer-reviewed, and statistically-valid data source at scale.

Despite its widespread utility and well-established methodological rigor, crowdsourced data is often erroneously associated with characteristics that can lead to its preemptive dismissal. The following entries address typical misconceptions regarding crowdsourced data.

Objection: "People only take a speed test when something is wrong."

Certainly, a perceived network or performance issue can drive an individual to take a speed test as a quick, high-level diagnostic tool. If speed or latency measurements are lower than normal, this can often signal a temporary drop in service levels from the ISP due to spike in demand causing capacity issues or to temporary network technical problems. A widely distributed network of test servers is more likely to measure the performance of the local network, whereas a limited number of server locations may be measuring the performance of either the local network or the backbone serving the broader internet.

User experience can also be impacted by a third-party dependency such as Netflix, Amazon, Google, gaming platforms or others. Speed tests often show that service from the ISP is providing the expected throughput speeds and latency performance but the service the user is attempting to access is providing a sluggish response or even suffering an outage. In these instances, a speed test validates the ISP's ability to deliver the services promised.

However, diagnosing network issues or outages is only one among many reasons users undertake a speed test. Validation and curiosity, for instance, are two additional common motivations. Examples of validation would include purchasing new equipment (such as a new mobile phone, wifi router or laptop) or changing service tiers and running a speed test to confirm that the investment made has resulted in an improved connectivity experience. An example of curiosity might be making sure that throughput speed and latency are adequate for an upcoming video call or gaming session.

Moreover, crowdsourced speed tests at times show performance that exceeds what is measured by controlled drive and walk testing. This highlights another reason why individuals might undertake a speed test: rather than only testing when performance lags, users can also test to see speeds when service is particularly fast or responsive.

The key point to keep in mind is that large speed test platforms include results across a full variety of connectivity experiences available within a given geographic area. The power of crowdsourced comes from this breadth and density: testing at scale (e.g., Ookla and M-Lab each generate tens of millions of tests each day) helps eliminate outliers and results in a statistically-valid, objective view of performance that is trusted by the industry, governments, press, and public alike.

Objection: “Tests over Wi-Fi cannot be trusted to show full performance”

This objection originates from the recognition that many Wi-Fi routers cannot measure the gigabit speeds some providers are now offering. This objection, however, is mistaking the goal of measurement associated with the challenge process. The requirement is not for ISPs to deliver gigabit speeds but rather to confirm that citizens have access to the current thresholds of 25/3 Mbps and 100/20 Mbps.

Reframed in this way, the question should not be whether a router can achieve gigabit speeds but instead if it is capable of delivering speeds of at least 100 Mbps download and 20 Mbps upload. Virtually every modern Wi-Fi router can measure throughputs at those speeds and higher. As a result, when the WiFi connection is good, limitations that contribute to speeds lower than the 100/20 Mbps thresholds can be primarily attributed to the service itself, not to the router used to deliver that service within the home.

Likewise, while it is true that network speeds can be intentionally throttled for users in hotels, coffee shops and other facilities offering public WiFi, these locations are not considered for funding. Residential mesh systems, in contrast, nearly always support distributed service that exceeds the 100/20 Mbps requirement. In short, residential Wi-Fi router and mesh systems are not the causal factor for test results below the 100/20 Mbps thresholds.

What can occur with WiFi connections is that the device performing a speed test is too far away or something between the device and the router is causing interference – for example, someone may initiate a test from their backyard, or they have placed their router inside of a steel cabinet. There is no doubt that these types of scenarios can sometimes cause the network performance readings to fall well below the actual service being delivered to the building.

To control for these occurrences, filtering for tests that do not meet minimum criteria for latency can minimize impacts from unusual user behavior. When latency can be measured between each hop along the traceroute from the device to the testing server and back, the first hop represents the one from the device to the WiFi router. If the latency measurement for this first hop is unacceptably high (e.g., greater than 10 ms), that is a strong indicator that the device is having a difficult time connecting and the test should not be used as an indicator of insufficient service levels.

But even if the test fails to meet that standard, that doesn't mean that the test has no value. First and foremost, tests with high latency still serve as "proof of life" that connectivity exists in that location. More importantly, if the tests show speeds higher than 25/3 Mbps or 100/20 Mbps while contending with significant interference (as, again, represented by latency greater than 10ms), it can be assumed that the actual speeds being delivered would represent an even higher rate had the connection quality been improved. As a result, Ookla recommends considering *any and all* samples that show speeds above 25/3 Mbps, regardless of the latency measured. This can assist in avoiding overbuilding an area already receiving the target service levels.

Why Ookla data can be trusted as part of this process

What is Ookla's mission?

Ookla's mission is to measure, understand, and help improve connected experiences. Every day, over 18 million people use Ookla Speedtest® to better understand and troubleshoot the performance of their internet connections. Additionally, the Speedtest® app automatically runs 300+ million daily background tests to measure mobile network coverage. The Speedtest® application is available on numerous platforms, including the web, mobile phones, tablets, desktop computers, and TVs. Speedtest® is also embedded in routers, gateways, IoT and other connected devices to improve networking software and hardware. To date, consumers have actively initiated tens of billions of tests. In short, consumers, governments, regulators, and press rely on our data to show an accurate, unbiased picture of connectivity.

What else is Ookla doing to help facilitate the challenge process?

Unbiased and trusted, Ookla helps create a bridge between the industry, government, and consumers. This role is especially critical when considering the challenge process, which can feel complicated and burdensome. Our goal is to help streamline the submission of evidence and facilitate better and more efficient communication between states and the NTIA/FCC. To that end, your recently updated contract now supports sharing individual Speedtest® records as evidence directly to the NTIA and FCC to dispute the National Broadband Map.

Ookla industry leadership and partnership

Ookla mobile and fixed network data is used by the U.S. Federal Communications Commission (FCC) for internal analysis, reports to Congress, and public documents on the status of the telecommunications marketplace.

Ookla is the exclusive provider of global network performance data to GSMA Intelligence (GSMAi), a trade body that represents the interests of mobile operators worldwide, uniting more than 750 operators with almost 400 companies in the broader mobile ecosystem.

As an official member of the ITU-T (Study Group 12), Ookla partners with leading global operators, test and measurement companies, infrastructure and hardware providers, network analytics providers, and regulators to help develop and define quality of service (QoS) and experience (QoE) standards.

Ookla is committed to helping improve people's lives through better internet access. To that end, we license certain data sets and grant publishing rights to academic institutions and NGOs that are focused on education, public safety, public health and other areas where better internet quality provides a higher quality of life.