

MLAB

@ **CS&S** Code for
Science &
Society

Measurement Lab

Supporting Open Internet Research

Lai Yi Ohlsen laiyi@measurementlab.net

@measurementlab @laiyiohlsen

Supporting Partners

MLAB



國立暨南國際大學
National Chi Nan University

moz://a



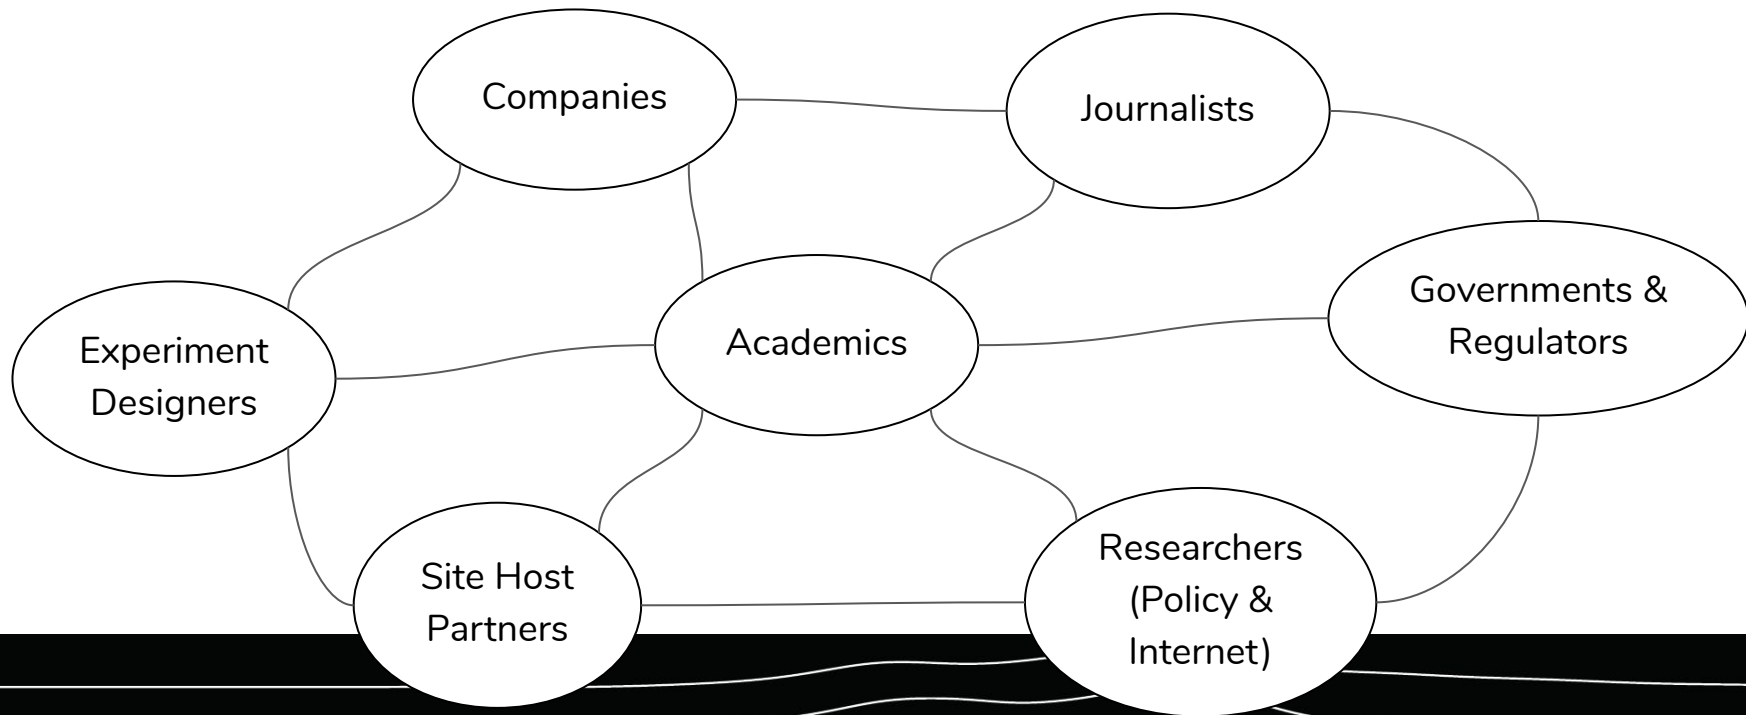
Our Community

MLAB

@

CS&S

Code for
Science &
Society



History

In 2008, it was hard to measure the Internet.

M-Lab's Mission



Measure the internet.

Save the data.


Make it universally accessible and useful.

History

In 2008, it was hard to measure the Internet.

In 2019, it still is, but slightly less so.

Measure the Internet



Measure the Internet: The Platform




We run high capacity hardware in well connected data centers, where ISPs interconnect with one another.

Measure the Internet: The Platform

MLAB

@ **CS&S** Code for
Science &
Society

M-lab measures
user experience of
the **full route** from
user to content



MLAB

@

CS&S

Code for
Science &
Society



Today — **500+** Servers in **130+** locations

Platform Upgrade:

<https://www.measurementlab.net/blog/modernizing-mlab/>

Measure the Internet: Host Experiments

MLAB

@

CS&S

Code for
Science &
Society



Glasnost
Max Planck Institute for
Software Systems



MobiPerf
University of Michigan



**Network Diagnostic
Tool**
Internet2



Neubot
Nexa Center for
Internet and Society,
Politecnico di Torino



NPAD
Pittsburgh
Supercomputing Center



Paris Traceroute
University Pierre et
Marie Curie



Project Bismark
Princeton University



Reverse Traceroute
University of
Washington



Sharperprobe
Georgia Tech College of
Computing



Windrider
Northwestern
University

Measure the Internet: NDT

NDT is a **single stream performance** measurement of a connection's capacity for “**bulk transport**” (as defined in IETF's RFC 3148). It reports **upload** and **download** speeds and **latency** metrics, and attempts to determine what problems limit speeds.

Measure the Internet: NDT

To run your own ndt-server, i.e. host your own speed test,
run:

```
docker run --net=host measurementlab/ndt
```

on any Linux machine.

Measure the Internet: NDT7

NDT7 supports BBR (compatible with IETF RFC 8837), runs over TLS and uses Websockets.

Measure the Internet: NDT

Then point your browser to

`http://localhost:3001/static/widget.html,`

you can run tests using the original version of the protocol. If you pass a TLS cert and run on port 443, you can run on NDT7 using TLS/SSL.

Measure the Internet: Experiments

Other tests include DASH, reverse-traceroute, sidestream, paris-traceroute, BISMark, WeHe/diff-detect, SamKnows, and more.

Measure the Internet: Experiments



Researchers designing server/client measurement experiments can apply to host with M-Lab and are accepted through our Experiment Review Committee.

Measure the Internet: Run Tests

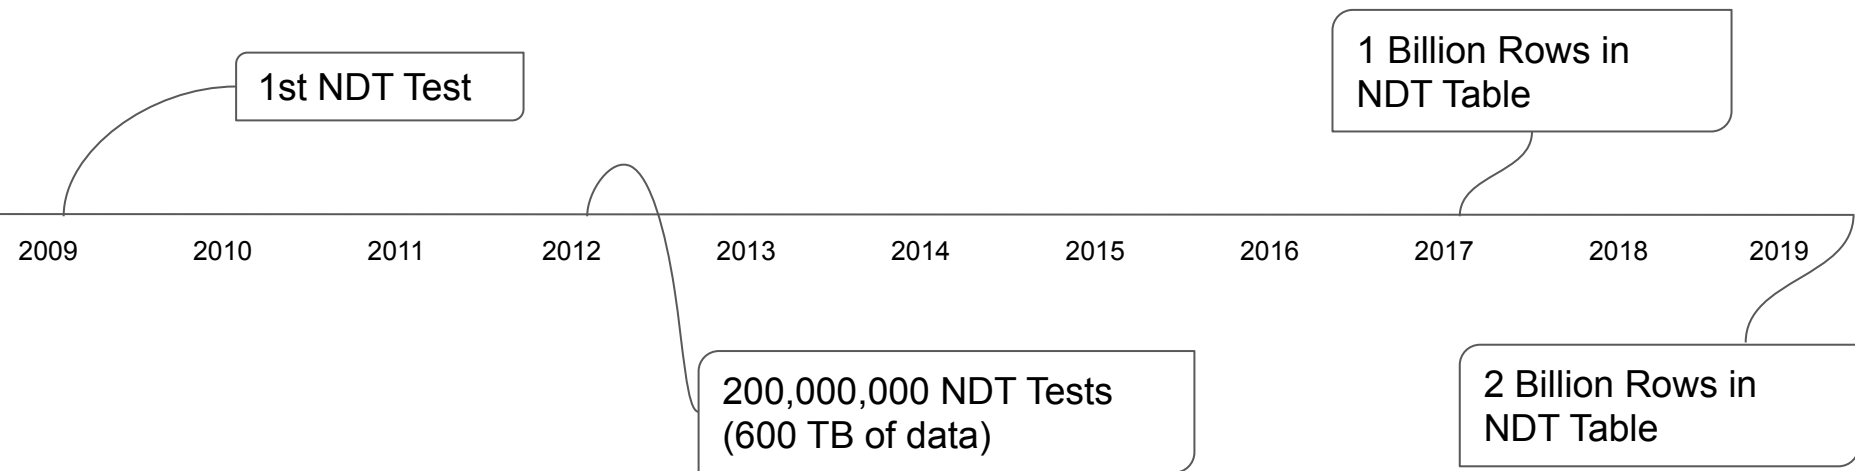
Users run active measurements through client integrations.

Clients:

<https://www.measurementlab.net/data/tools/>

Measure the Internet: Run Tests

- Current Daily volume ~3,000,000 new NDT measurements per day
- Currently, 2 billion rows in NDT Table



2billion NDT:

<https://www.measurementlab.net/blog/celebrating-2billion-ndt-tests/>

Measure the Internet: Run Tests

An easy way to run a speed test:

`speed.measurementlab.net`

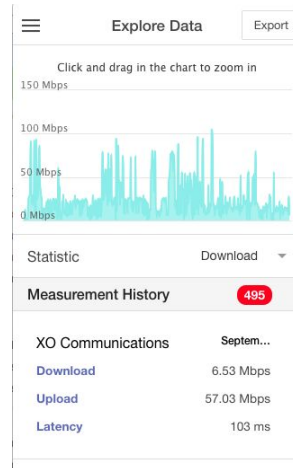
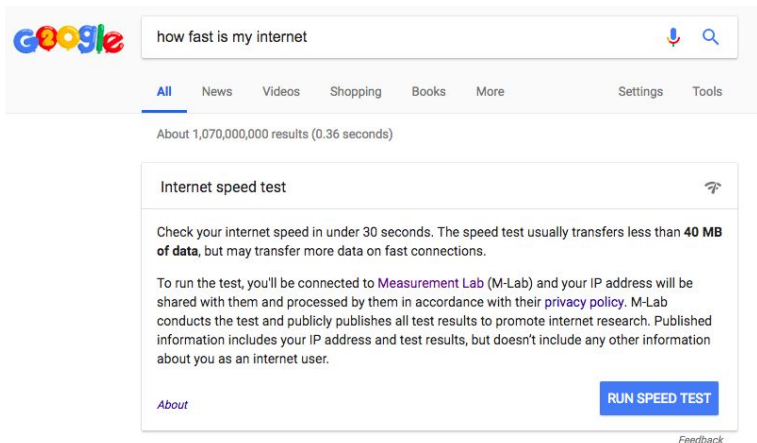
Measure the Internet: Run Tests

MLAB

@

CS&S

Code for
Science &
Society



Google Search, Software Integrations, Router Integrations, [Fingbox](#), [Chrome Extension](#)

Clients:

<https://www.measurementlab.net/data/tools/>

Save the Data



Save the Data

Store everything in Google Cloud Storage. All of it. Including raw packet traces and metadata.

Make it Universally Accessible and Useful



Make it Universally Accessible and Useful: BigQuery



For free and open access to the data, all you have to do is sign up for the M-Lab Discuss List.

Quick Start - www.measurementlab.net/quickstart/
Data Docs - www.measurementlab.net/data/docs/
Get help - support@measurementlab.net

Make it Universally Accessible and Useful: BigQuery

MLAB

@

CS&S

Code for
Science &
Society

The screenshot displays the Google Cloud Platform BigQuery interface. The top navigation bar shows 'Google Cloud Platform' and 'MLab project - public BigQuery'. The left sidebar contains a 'Query history' section and a 'Resources' section with a search bar and a list of datasets under the 'measurement-lab' project. The main area shows a query editor with a SQL query titled 'median mbps in Africa'. The query is as follows:

```
1 SELECT
2   connection_spec.client_geolocation.country_name as country,
3   COUNT(*) as count,
4   APPROX_QUANTILES(8 * (web100_log_entry.snap.HCThruOctetsAcked /
5     (web100_log_entry.snap.SndLimTimeRwin +
6     web100_log_entry.snap.SndLimTimeCwnd +
7     web100_log_entry.snap.SndLimTimeSnd)), 100)[safe_ordinal(51)] as median_mbs
8 FROM measurement-lab.ndt.downloads
9 WHERE connection_spec.client_geolocation.continent_code = 'AF'
10    AND EXTRACT(YEAR from log_time) = 2019
11 GROUP BY country
12 ORDER BY country
```

Below the query editor, there are buttons for 'Run', 'Save query', 'Save view', 'Schedule query', and 'More'. A status message indicates 'This query will process 252.6 GB when run.' with a green checkmark. The bottom section shows the 'measurement-lab' dataset list with a message 'Datasets and tables available' and a note to 'Use the Resources tree to view your data, or create a new dataset using the controls above'.

Make it Universally Accessible and Useful: BigQuery



We're always available for help with BigQuery or anything else at `support@measurementlab.net`.

Make it Universally Accessible and Useful: Visualization Site



The easiest and most popular way to access the data is using the M-Lab Visualization Site.

`viz.measurementlab.net`

Make it Universally Accessible and Useful: Visualization Site

MLAB

@

CS&S

Code for
Science &
Society

MLABviz

HomeAboutVisualizationsDataTestsPublicationsBlogLearnContribute

LocationsCompareDataAbout M-Lab Viz

Compare

The compare tool provides a way to compare speed tests along different aggregations. The **Facet By** selector on the left allows you to split the data by either Location, Client ISP or Transit ISP. Once a facet is selected, fill in the search bars on the right with the specific locations and ISPs you want to look at.

Selecting values from these lists, or from the suggestions provided, will automatically update the charts below.

01/08/2019 - 07/08/2019

Reset

Facet By ⓘ

Location

Client ISP

Transit ISP

Locations

Select one or more locations to explore measurements in. Each location will get its own chart.

Search for a location

Clear

Baltimore, Maryland ✕

Westminster, Maryland ✕

Towson, Maryland ✕

Columbia, Maryland ✕

Filter by Client ISP

Select one or more Client ISPs to filter the measurements by.

Search for a client ISP

Clear

Comcast ✕

Ting Fiber ✕

Port Networks ✕

Filter by Transit ISP ⓘ

Select one or more Transit ISPs to filter the measurements by.

Search for a transit ISP

Suggestions

Verizon

Sprint

Cogent Communications

T-Mobile USA

Johns Hopkins University

networkMaryland

Suggestions

GTT

Level 3

Cogent Communications

TELIANET

TATA COMMUNICATIONS (AMERICA)

COMMUNICATIONS (AMERICA)

Zayo Bandwidth

Internap Network Services Corporation

Corelogic, Texas

Tedco's Piedmont Exchange Point

Project

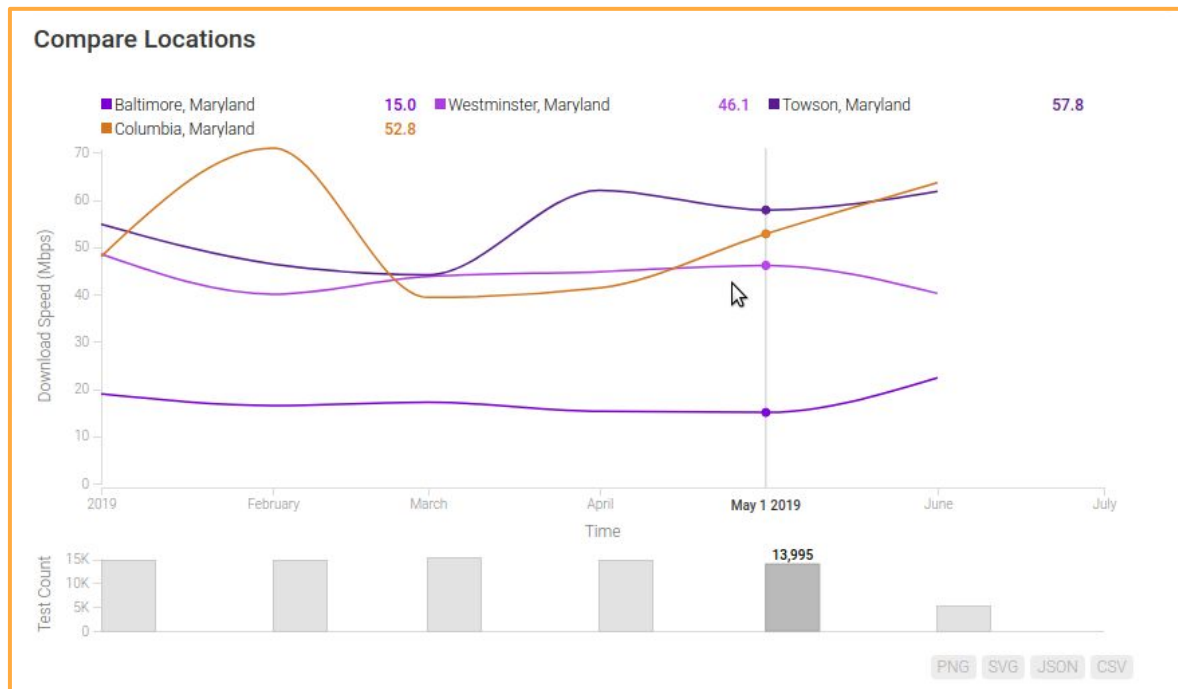
Make it Universally Accessible and Useful: Visualization Site

MLAB

@

CS&S

Code for
Science &
Society



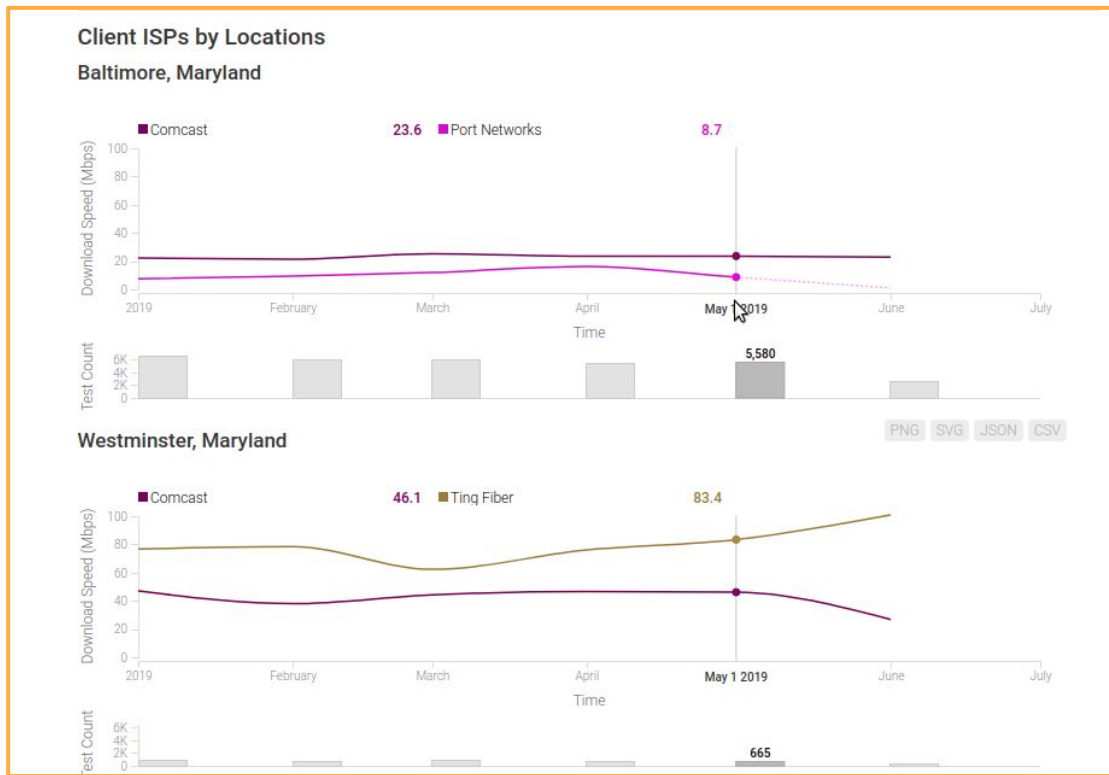
Make it Universally Accessible and Useful: Visualization Site

MLAB

@

CS&S

Code for
Science &
Society



Make it Universally Accessible and Useful


Running a speed test *once* gives you information about that one moment in time.

Collecting longitudinal, open data provides meaningful information about the behavior of the Internet over time.

MLAB

@ CS&S Code for
Science &
Society

How Researchers Use M-Lab



How Researchers Use M-Lab

Academic Publications

<https://www.measurementlab.net/publications/>

- Estimating Residential Broadband Capacity using Big Data from M-Lab
 - University of New South Wales, Sydney, Australia
- Access is more than cost: Measuring the quality of mobile broadband service
 - Web Foundation, Alliance for Affordable Internet
- Can Competition-Enhancing Regulation Bridge the Quality Divide in Internet Provision?
 - National University of Singapore
- NETPerfTrace – Predicting Internet Path Dynamics and Performance with Machine Learning
 - Université de Liège - ULiège

MLAB

@ CS&S Code for
Science &
Society

How Researchers Use M-Lab



How Researchers Use M-Lab



We support community-based research initiatives that provide an accurate depiction of Internet performance.

Community Based Research

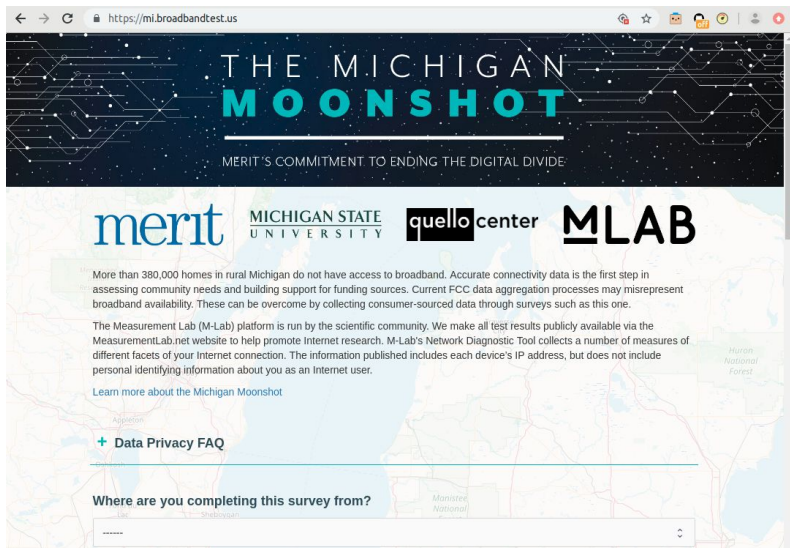
Survey & Mapping Initiatives

MLAB

@

CS&S

Code for
Science &
Society



- Michigan / MERIT - broadband testing & mapping with a homework gap hook
 - R&E network, K12 statewide systems
 - Homework assignment to run a test
- <https://mi.broadbandtest.us/>

Community Based Research

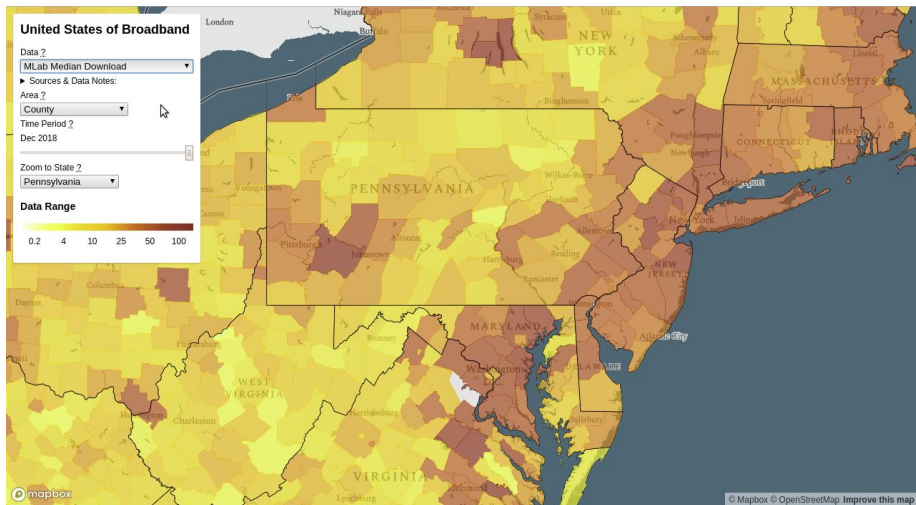
Survey & Mapping Initiatives

MLAB

@

CS&S

Code for
Science &
Society



- M-Lab prototyped community tool used to conduct broadband survey, speed test, & aggregation: *Piecewise*
- City of Seattle, WA
- Stevens County & Ferry County, WA
- Clearwater County, ID

Community Based Research

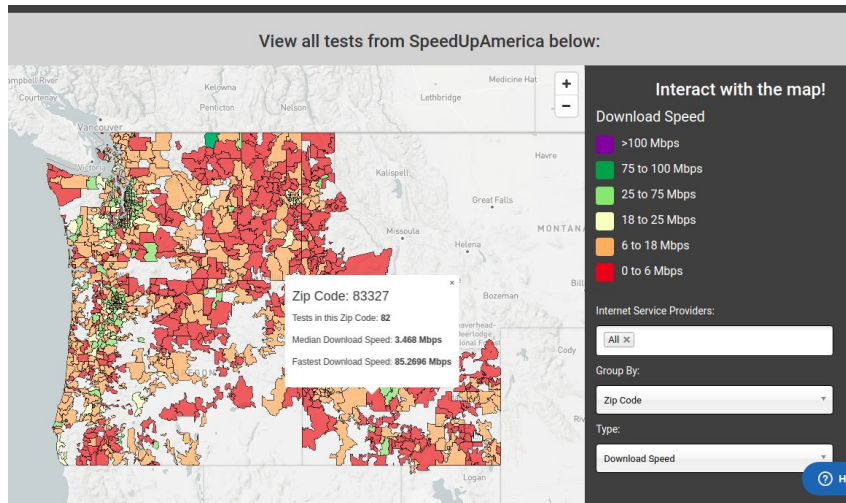
Survey & Mapping Initiatives

MLAB

@

CS&S

Code for
Science &
Society



- Speedup Louisville
 - Developed by Louisville, public/private partnership
 - Supported by M-Lab
- Speedup America
 - Further development at Technology Assoc. of Oregon

Community Based Research

App / Client for Data Collection

- National Association of Counties (NACo) using a branded app integration of NDT to collect data for rural areas
- Uses the M-Lab platform as a measurement service
- NDT test results are sent to M-Lab
- **Enhanced location data** is saved in NACo's database, with a copy of the standard test results

M**L****LAB**

@

CS&S

Code for
Science &
Society



TestIT

NACo Tools

Everyone

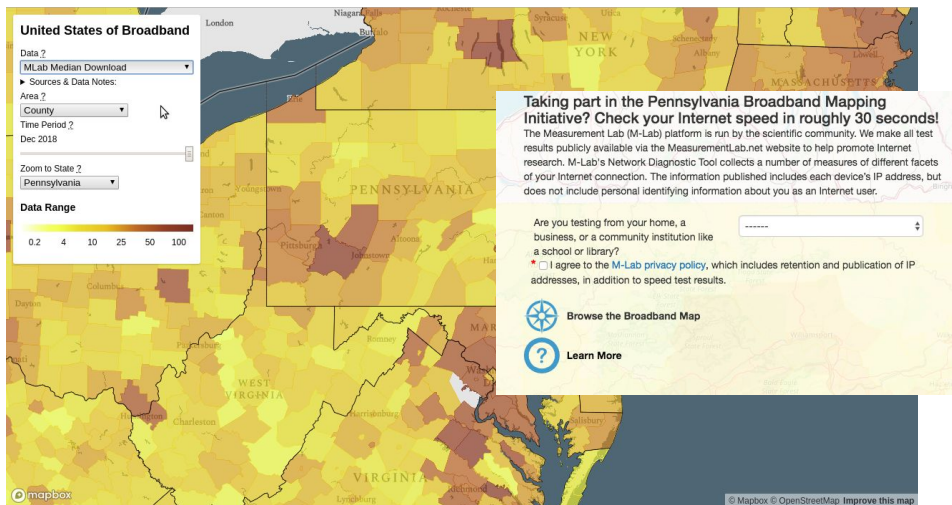
 This app is compatible

Community Based Research

Survey & Mapping Initiatives

MLAB

@ CS&S Code for
Science &
Society



- Study conducted by Penn State, compares measured speeds with each FCC 477 data release
- Select M-Lab speeds & metrics, FCC, or +/- difference layers
- Aggregate by county, state house/senate, zipcode, census tract
- <https://pa.broadbandtest.us>

How Researchers Use M-Lab **MLAB**

Measuring Library Broadband Networks



- Partner in research program to build structured data collection tools
- [IMLS Measuring Library Broadband Networks](#)
 - Partners: Internet2, Simmons University
 - Year 1: 10 Libraries / Year 2: +50 Libraries
 - NDT, speedtest.net, and more
 - <https://slis.simmons.edu/blogs/mlbn/>
- Goal: Open source code, automated testing via on-premise devices, data visualization, local/regional data comparison

Potential Research Questions

- Paris-Traceroute data offers lots of opportunity for research on Routing.
- Key Performance Indicators for cities.
- Metrics for municipally run networks.
- Enabling the detection of throttling and Internet shutdowns.

M-Lab Principles

- All measurements are active measurements
 - All the data is synthetic data, we take user privacy seriously.
 - Client initiated tests only, Servers do not start tests on their own.
- Clients are built by and for the community.
 - Anyone can develop them.
- Openness
 - All of the data is released CC0.
 - All of the code is open source.

Learn more about M-Lab: 10th Anniversary Convening blog post, videos:
<https://www.measurementlab.net/blog/mlab-10year-wrapup/>

How to Get Involved

- Propose an experiment.
- Use the data in your research and planning
- Integrate NDT into various clients (apps, hardware, etc.)
- Host M-Lab in your data center or Internet Exchange

MLAB

@ **CS&S** Code for
Science &
Society

Measurement Lab

Supporting Open Internet Research

Lai Yi Ohlsen laiyi@measurementlab.net

@measurementlab @laiyiohlsen