

MLAB

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

Measurement Lab

IFF Measurement Village

Lai Yi Ohlsen laiyi@measurementlab.net

@measurementlab @laiyiohlsen

~

Agenda

- What is M-Lab?
- How do we measure the Internet?
- What makes the data valuable?
- How can you use M-Lab?
- How does M-Lab support community research?
- How can M-Lab support the Internet Freedom community?

MLAB

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

What is M-Lab?

Mission

Measure the Internet

Save the data

Make it universally accessible and useful

Mission

Measure the Internet ←

There are many
ways to do this!

Save the data

Make it universally accessible and useful

Origin Story

- A solution to the lack of widely deployed, professionally maintained servers with ample connectivity to support Internet measurement experiments.
- Researchers also reported an inability to share large data sets with one another and other experts easily.

Fast forward 12 years...

- Platform
- Pipeline
- Data
- Tools
- Community
- Team

Team



Core Team - Code for Science & Society

Measurement Lab is a fiscally sponsored project of CS&S

Staff

- Project Director - Lai Yi Ohlsen
- Program Management & Community Lead - Chris R.
- Platform Engineers - Nathan K., Robert D.

Team



Contributors

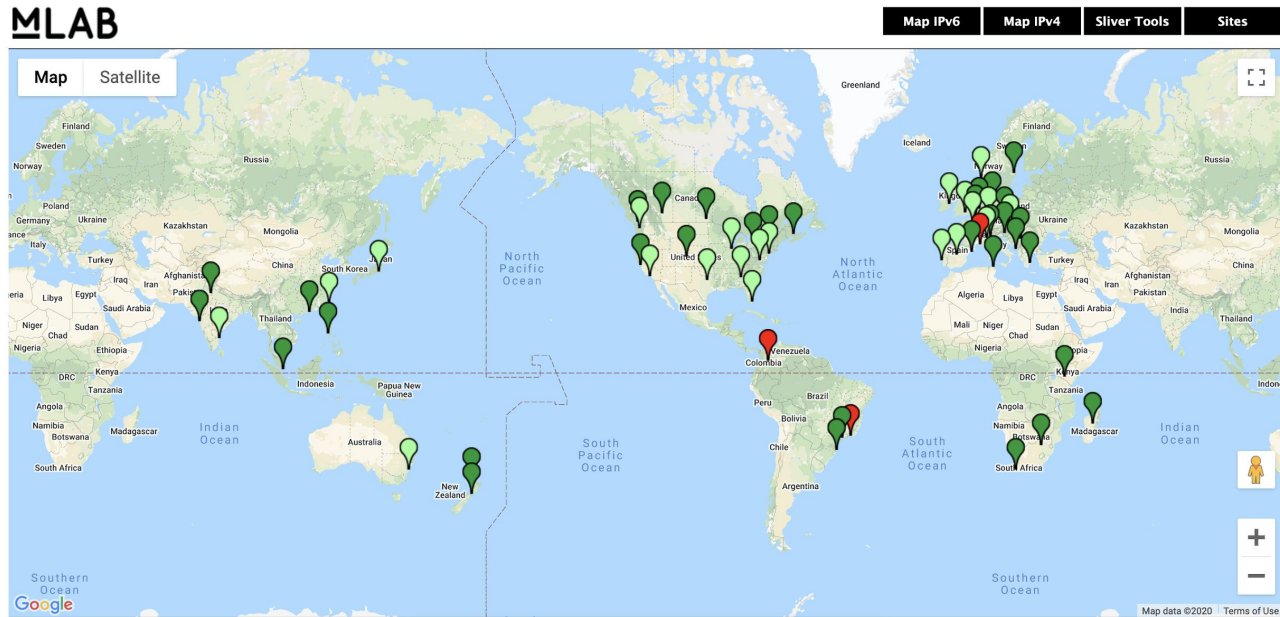
Over the years, Princeton's PlanetLab, New America's Open Technology Institute, Google, Open Technology Fund, Mozilla, Media Democracy Fund, Internet Society and more

As a core contributor, Google supports the project by contributing Internet performance research, infrastructure support, and by assigning a small team of Software Engineers to write open source code for the M-Lab platform and pipeline

How do we measure the Internet?

Off-net platform

We host about 500+ servers in about 60+ metro areas.



Off-net platform

All of M-Lab's servers are hosted in "off-net" data centers or data centers where ISPs peer with one another, outside of access networks.

Our goal is to measure the full path from user to content.

Off-net measurements measure the "Inter" part of the Internet.

Off-net platform



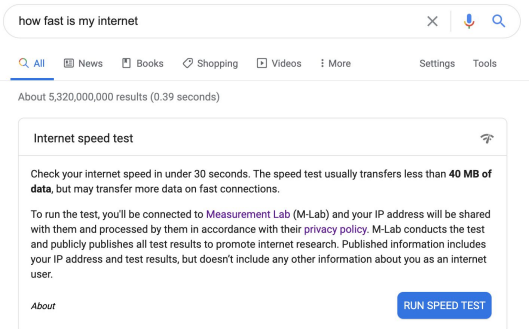
The servers host “measurement services”, proposed by tool builders (academic computer scientists, network engineers, etc.) and approved by our Review Committee.

Test clients

Anyone can develop a test client (no approval necessary). Some are community developed, some we write and maintain. Test clients run tests against the servers.

The data is then stored in public archive and can be parsed into BigQuery.

Examples of test clients:



Google Search "How fast is my Internet"

What is the speed and performance of my network?

DASH STREAMING TEST

Measures video streaming performance

NDT SPEED TEST

NDT (Network Diagnostic Test) provides a sophisticated speed and diagnostic test for understanding the performance of your network

“M-Lab data”

“M-Lab data” could be referring to the data generated by any one of the measurement services that the M-Lab platform hosts.

NDT is our most frequently run test. When people refer to “M-Lab data”, as of today they are often referring to NDT data.

Bulk transport capacity

NDT measures the single-stream performance of bulk transport capacity.

Bulk transport capacity refers to the rate that a link can deliver data with TCP --
i.e. the reliability of that link.

Link capacity refers to the maximum bitrate of the link.

Both are conflated with Internet “speed.”

Single stream

NDT measures the single-stream performance of bulk transport capacity.

Modern web browsers will use multiple streams of data, but testing for multiple streams can compensate for data packet loss over a single stream.

A multi-stream test can return measurements closer to link capacity but it would not represent packet loss.

By testing for single-stream performance, NDT is an effective baseline for measuring a user's Internet performance.

Why is my M-Lab test result different than _____ ?

1. **NDT** vs. other measurement services
2. **Off-net** vs. on-net
3. **Bulk transport capacity** vs link capacity
4. **Single stream** vs. multi-stream

More info:

[How fast is my Internet? Speed Tests, Accuracy, NDT & M-Lab](#)

M-Lab's other measurement services



DASH (Dynamic Adaptive Streaming over HTTP) measures the quality of tested networks by emulating a video streaming player. It is maintained by Simone Basso of the OONI team.

WeHe measures differential treatment of applications by ISPs. It was developed and is maintained by Dave Choffnes team at Northeastern University.

More info: <https://www.measurementlab.net/tests/>

Sidecar services



For every connection to an M-Lab server, the **Traceroute** core service collects network path information from our server back to the client IP that initiated the connection.

The M-Lab packet-headers service provides a binary which collects packet headers for all incoming TCP flows.

Sidecar services

M-Lab uses TCP INFO to collect statistics about every TCP connection used by each hosted measurement service running on the M-Lab platform. TCP measures the network as part of its normal operation. All transport protocols, including TCP, measure the network to determine how much data to send and when to optimally fill the network. Sending too much data or sending it too fast results in congestion, network queue overflows and discarded packets; sending data too slowly results in under-filled networks and wasted idle capacity. TCP INFO exposes these built in measurements for diagnostics and other applications.

What makes the data valuable?


Individual tests vs. aggregate data

By design, the value of NDT data is in the aggregation of many connection test results from around the world. Any single test is limited as an indicator for individual Internet connections due to the multiple factors that could influence the results.

However the aggregate test data provides useful views into trends in Internet performance. Patterns in the dataset enable us to ask better questions about Internet performance at scale and the factors affecting it.

Large & longitudinal

- Current Daily volume ~3,000,000 new NDT measurements per day
- As of 2020, close to 2 billion rows in NDT Table



Open, free, public

All of the code for each measurement service is open source.

All reference clients are open source.

All of the code that runs M-Lab's platform and pipeline is open source.

All of the data is publicly archived.

All of the data parsed into BigQuery is free to access.

User-contributed, global, representative



All tests are active, users opt into them.

All measurement services inherit the off-net platform methodology.

NDT tests are run globally (two thirds are run from outside of the US).

Privacy

M-Lab is aware that privacy is a concern for users running any kind of test.

All measurement services only collect the IP address assigned by a user's Internet Service Provider. This is the only piece of personal data collected by our tests. No other data about your computer or network is collected.

Users that want their IP address removed from our data are able to do so by following the process outlined in our [Privacy Policy](#).

MLAB

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

Access to M-Lab Data

Accessing M-Lab Data

MLAB

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

- There are many ways to explore and visualize M-Lab Data. We support audiences with a wide range of backgrounds, expertise, training, and needs, and therefore try to present a range of options.
- M-Lab Visualization Website - <https://viz.measurementlab.net/>
 - First stop for beginners - search by city, region, or country
 - Data presented stops in Nov. 2019, but in the process of being upgraded
- BigQuery - <https://www.measurementlab.net/data/docs/#querying-bigquery-basic>
 - Intermediate/advanced option for people or orgs with data science or database expertise
 - Most flexible, but also potentially high onboarding curve
- Third party tools that integrate with BigQuery
 - Tableau
 - R Studio
 - APIs for popular programming languages


Accessing M-Lab Data

- We've recently started publishing interactive reports using Google's Datastudio product
- BigQuery-driven reports that you can interact with to see aggregate NDT data
 - Blog post - Regional test rates & metrics re: COVID-19's Impact
<https://www.measurementlab.net/blog/datastudio-covid19-test-rates-increase/>
 - United States Dashboard - <https://datastudio.google.com/s/r3P020V1Qbw>
 - Global Dashboard - <https://datastudio.google.com/s/tUdGdBojNkM>
- Datastudio reports are an approachable way to go from a BigQuery query to charts, tabular data, maps, etc.

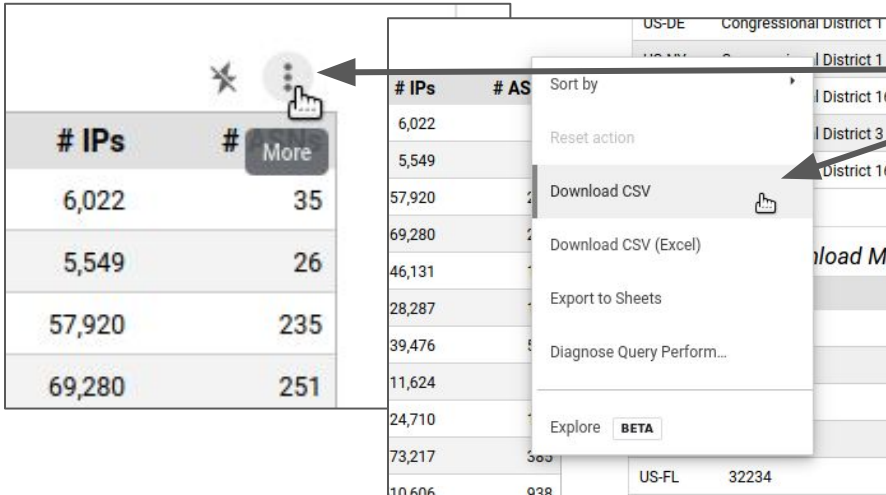
Features of Datastudio Reports



Page navigation at top left.



Filter controls like Date Range let you control aggregate output.



Selected data in tables can be exported.

United States Dashboard:

<https://datastudio.google.com/s/r3P020V1Qbw>

How does M-Lab support community research?

Piecewise

Piecewise is an open-source public engagement portal that collects both user-volunteered survey responses and speed test data using the Measurement Lab platform.

Data collected by Piecewise is visually aggregated on the web and mapped on top of M-Lab's public dataset.

Measure-app

The M-Lab Measure Chrome browser extension (measure-app) provides an extension for the Chrome web browser to run NDT tests.

Features include scheduling tests, user annotation of results, selection of M-Lab server, saving test results in an SQL Lite database in the user's browser profile, and exporting results to a local CSV file.

It also includes language localization, with translations contributed by the Open Technology Fund's Localization Lab community.

The extension is currently being upgraded in partnership with UNICEF's GIGA project.

Murakami

Murakami is a container-based service that enables automatic, recurring measurements.

Thanks to our partnership with Simmons University and support from IMLS, we are now able to run Murakami on on-premise measurement devices. We're using Odroid-xu4's but Murakami can be run on any device that can run Docker.

A fleet of devices can be managed using Balena Cloud or the Mozilla WebThings Framework.

How can M-Lab be a resource
for the IF community?

New measurement services



After our platform upgrade in 2019, the M-Lab platform is more ready than ever to have new open-source measurement services proposed.

In the next few years, we'd like to prioritize global engagements with tool developers and researchers outside of the US.

Internet Freedom advocates have a unique perspective on the meaning of a healthy Internet and we welcome your perspective on how to best use our platform to measure it.

Test clients

The IF community can write new test clients for NDT, DASH, WeHe, or any of our future measurement services or integrate them into your existing application to provide your users with more information about their Internet performance.

Use/ingest M-Lab data

M-Lab data can be ingested into an application or dashboard to provide context to the Internet performance in a specific location or time period.

For example, Psiphon has integrated NDT data into their new [dashboard](#) alongside OONI data.

Run your own NDT server

NDT-server can be run any machine that can run Docker. This means you use NDT to test the performance of a segment of a network.

To run your own ndt-server, i.e. host your own speed test, run: `docker run --net=host measurementlab/ndt` on any Linux machine.

NDT7, written by Simone Basso of OONI, supports BBR (compatible with IETF RFC 8337), runs over TLS and uses modern Websockets.

More info on NDT:

- <http://www.es.net/science-engagement/ci-engineering-brownbag-series>
- <https://youtu.be/mf65RLIPYmE>

Potential areas of research

How can each of our datasets complement one another? For example, what can data generated by RIPE Atlas probes tell us that data generated by M-Lab's platform can't?

Our traceroute data is largely unexplored. What can it tell us about hops across borders?

Furthermore, how can our datasets create meaningful metrics and indicators to help monitor and even predict events that put Internet Freedom at risk. One highly interesting example is, how can M-Lab data be used to define instances of state sponsored throttling?

Measurement Lab

csv,conf,5

Lai Yi Ohlsen laiyi@measurementlab.net

@measurementlab @laiyiohlsen