



AIM: Aggregated Internet Measurement


David Tuber, Cloudflare

Lai Yi Ohlsen, Measurement Lab









 Cloudflare
Zero Trust Services







 Cloudflare
Network Services


 Cloudflare
Application Services

1 Cloudflare One


-  Zero Trust Network Access
-  Secure Web Gateway
-  Cloud Access Security Broker
-  Cloud Email Security
-  Remote Browser Isolation
-  Data Loss Prevention




-  WAN-as-a-Service
-  Firewall-as-a-Service
-  L3 & L4 DDoS Protection
-  Network Interconnect
-  Smart Routing
-  IDS/IPS ¹

-  WAF with API Protection
-  Rate Limiting
-  Load Balancing
-  Bot Management
-  L7 DDoS Protection
-  CDN and DNS

 Cloudflare Edge
Developer Platform

-  Workers
-  Workers KV
-  Pages
-  Durable Objects
-  Video Streaming

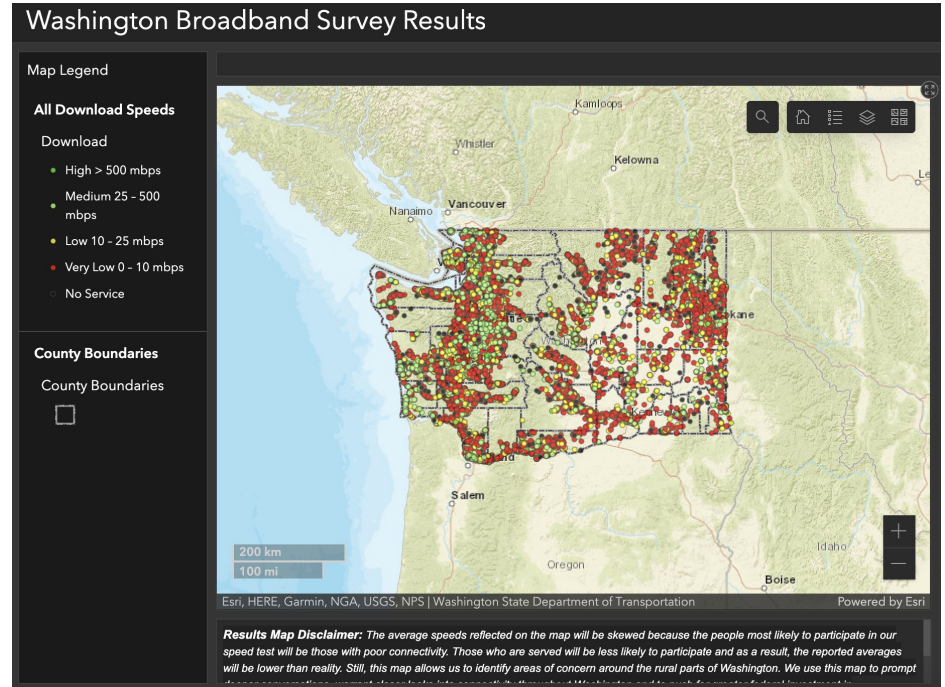
 Cloudflare
Global Network

-  **Global Edge:** 285+ cities, 95% of population within 50ms, 11,500+ interconnects, 192 Tbps capacity, China Network
-  **Building Blocks:** SSL/TLS, mTLS, Authoritative/Recursive DNS, DNSSEC, DNS over HTTP, L4-7 over Wireguard
-  **Compliance/Privacy:** FedRAMP, ISO, SOC, PCI, GDPR compliant, Logs & Analytics, Data Localization Suite

¹Contact us to request access

About Measurement Lab

- Started by Vint Cerf and multidisciplinary group of researchers in 2009
- We publish **open Internet performance data** for use by academic researchers, industry, policymakers and consumers
- We believe public policy decisions should be made with **open source methodologies**



About Measurement Lab


- We currently provide **~4 billion rows of data**, which comes from integrations of our open source clients, including the one Google Search
- All of our measurements measure to servers in interconnection points
- Though we collect and publish “speed test” data, we’re interested in **expanding the notion of Internet performance beyond bandwidth**

Internet speed test 🔗

Check your internet speed in under 30 seconds. The speed test usually transfers less than **40 MB of data**, but may transfer more data on fast connections.

To run the test, you'll be connected to [Measurement Lab \(M-Lab\)](#) and your IP address will be shared with them and processed by them in accordance with their [privacy policy](#). M-Lab conducts the test and publicly publishes all test results to promote internet research. Published information includes your IP address and test results, but doesn't include any other information about you as an internet user.

[About](#) [RUN SPEED TEST](#)



How do end-users define good Internet?

Or more importantly, what do end-users say when they have a bad Internet experience?

Meet Andrea

- Therapist, works in an office (more on that later)
- Photographer
- Internet technological experience: none



Common customer complaints

- “My Netflix won’t play”
- “I’m seeing lag in my video game”
- “My Zoom is cutting out”

How we measure the Internet doesn't map to how customers view it

- This leads to misalignment of customer expectations and reality
 - “I bought gigabit Internet but my game is still lagging”
- Existing customer-facing tools further this misalignment
- This also makes debugging network issues difficult

Speed tests should
answer: “what is my
Internet good for”

Goals of AIM

- End-users should not have to translate between connection quality data and user experience
- Application owners with a firm understanding of optimal network conditions define user experience metrics using new and existing data
- Network operators should be able to see how that experience is calculated and measure impact using forward-looking dataset
- Ensure that public policy decisions that affect all of us come from a unified set of data - we all need to speak the same language to understand each other!

User scenario

- I am a network operator and someone is complaining about poor service
- How can AIM help me?

Scenario: Debugging issues in an office

- Andrea works in an office with 1 WiFi access point
 - Her zoom keeps crashing
 - She can't upload documents
 - She can stream video
- What's going on and how can we fix it?
 - Before: "run a speed test"
 - Speed tests for 3 different offices in suite:
 - 1, close to wifi: 25mbps down, 7 up
 - 2, far away from wifi: 21mbps down, 7 up
 - 3, far away from wifi: 39mbps down, 7 up
 - What is actually the problem here and how can Andrea figure it out?

Solution: get close to your wifi router

- As you may have guessed, the problem was proximity to wifi router:
 - Far away offices had 7% packet loss and 61ms of jitter
 - That data was not present in the speed test throughput characteristics
 - Andrea doesn't know what packet loss and jitter are so she can't use them to judge her Internet
- With AIM, instead we can say:
 - Your Internet is fine for streaming
 - Your Internet is poor for video calls
 - Your Internet is poor for gaming
 - Try getting closer to your router
- A 2-week debugging scenario now takes 5 minutes of speed tests

Analysis

- Researchers can now use AIM data to make more specific claims about a region
 - Median bandwidth is x can only tell you so much
 - AIM can provide more detail by saying
 - x% of connections are good for streaming
 - y% of connections are good for video
 - z% of connections are good for gaming
- Users can also compare ISPs against each other in a region
 - I just moved to Atlanta, what Internet should I get?
 - I want the one best for streaming/video/gaming

How do I get access to the data?

You can access this data today

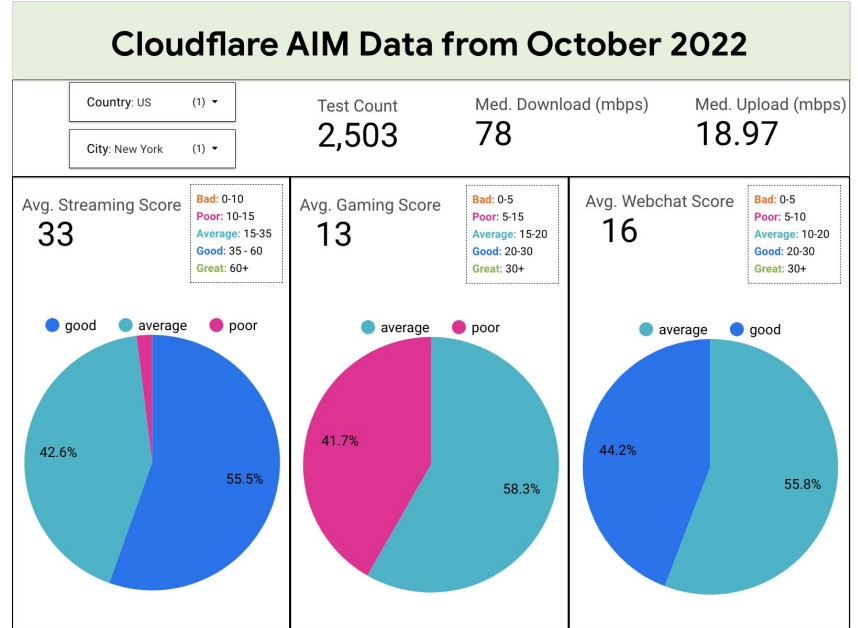
- You can access the data [via BigQuery](#)
 - All you have to do is sign up for our mailing list
 - All data is in the `measurement-lab.cloudflare.speedtest_speed1` table

```
SELECT
  *
FROM `measurement-lab.cloudflare.speedtest_speed1`
WHERE
  date >= "2022-10-01" AND date <= "2022-10-31"
  AND clientCity = "New York"
  AND clientCountry = "US"
```

[Get started here!](#)

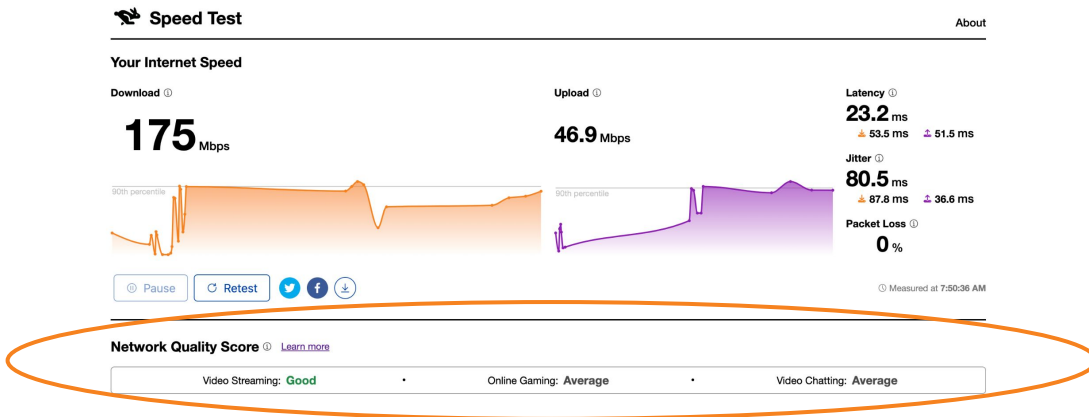
You can access this data today

- You can pull in the data into your custom dashboards for you and your customers
 - e.g. dashboard made in Google's Looker
 - **Question: what tools do you use? Grafana? Tableau? Power BI? Other?**
- Coming soon: AIM data in [Cloudflare Radar](#)



You can access this data today

- You can review our open-source repos to understand how data is calculated and collected:
 - [NDT7](#)
 - Cloudflare Speed Test: soon!
 - AIM Score calculation repository: soon!
- Check out your AIM scores on the Cloudflare Speed Test



What comes next

- Making data as easy to use as possible by network operators
 - Creating tutorials, dashboards, etc.
- Calculate AIM with M-Lab's NDT data
 - What does M-Lab's NDT say your Internet is good for? How does it compare to what Cloudflare says?
- Aligned schema for speed tests?
 - Better explanations of differences between them
 - Who uses cubic vs reno vs bbr
 - Multi-threaded responses vs single-threaded? Etc.
- Open-sourcing AIM scores to application owners
 - Getting more eyes on data means better data
 - Open sourcing of speed tests for wider adoption of standardized testing

We want feedback

- Contact us at laiyi@measurementlab.net, tubes@cloudflare.com
- We would love to know:
 - How well does our dataset help you troubleshoot customer issues?
 - What other datasets do you use?
 - How easy (or not easy) is it for you to use our data?
 - What are your preferred vis tools (e.g. Grafana, Tableau, etc.)
 - What other metrics do you want to see?
 - Suggestions, questions, comments, concerns?