



Broadband Transparency & Measurement Lab

Vint Cerf
Google

I²MTC, 3-6 May 2010



Broadband Transparency

- When an Internet application doesn't work as expected, how can you tell whether the problem is caused by your broadband ISP, the application, your PC, or something else?
- If your ISP is interfering with certain applications or websites, how would you know?
- What connection performance do you receive from your ISP? Is it what is specified in your contract?

Why Transparency?

- The Internet is designed to **put end-users in control** of the purpose for which the network is used.
- ISPs' network management practices can dramatically **impact the online services' performance** in ways that are **difficult for users to detect or understand**.
- Transparency is necessary (though not sufficient) to **keep the Internet open**.

What can be done?

Empower end-users, policy makers and the public at large with information about the actual performance of broadband connections.

But ...



Data sharing problem

- Industry, academia are not communicating well.
 - Research topics don't necessarily map to issues where industry needs academic help.
- Data are the problem.
 - Industry has data, doesn't really use it.
 - Academia can't function without it.
- Getting data and researchers together is harder than it looks.
 - Privacy policies.
 - Volume of data (both too much and too little)

Measurement Lab (M-Lab)



- *Open*, distributed server **platform** for **researchers**, to deploy Internet measurement tools.
- *Open* **tools** for **users**, to test their broadband connection.
- *Open*, better **data** for **everyone**, to build on a common pool of network measurement data.

M-Lab believes that open tools and open data access facilitate peer-review, independent analysis, and better network research.

Founding partners:

- Open Technology Institute (at New America Foundation)
- PlanetLab Consortium
- Google
- Academic researchers
 - Georgia Institute of Technology, Internet2, Max Planck Institute, Northwestern University, Pittsburgh Supercomputing Center, Princeton University



Supporting partnerships

- Servers hosting & connectivity
- Client support
- DNS hosting & server selection
- Data hosting



Google's role in M-Lab

- Deploy a "seed round" servers around the world:
 - Servers, space, power, connectivity.
- Collaborate with researchers to build measurement tools.
- Explore ways to make all data from all tools open and available to the public.
- Google's investment in platform encourages others to invest:
 - Servers, space, power, connectivity from partners
 - Software development on both servers and clients
 - Modern, large-scale data hosting and analysis solutions
 - Academic focus on data analysis, visualization, etc.



M-Lab for researchers

What?

- Purpose-built active measurement platform
 - Standards for machines, connectivity
 - Control server-side congestion, CPU oversubscription
- Connect researchers to real users, real data
 - Collect user-focused tools together

Why?

- Research tools have a hard time serving “flash crowds” (a.k.a. “Slashdot effect”)
 - Internet2 != Internet
 - University connections to Internet congested by student use
 - Variations in machine provisioning can affect results
- Research tools need help getting visibility
 - In a word: marketing
 - Two: collective marketing

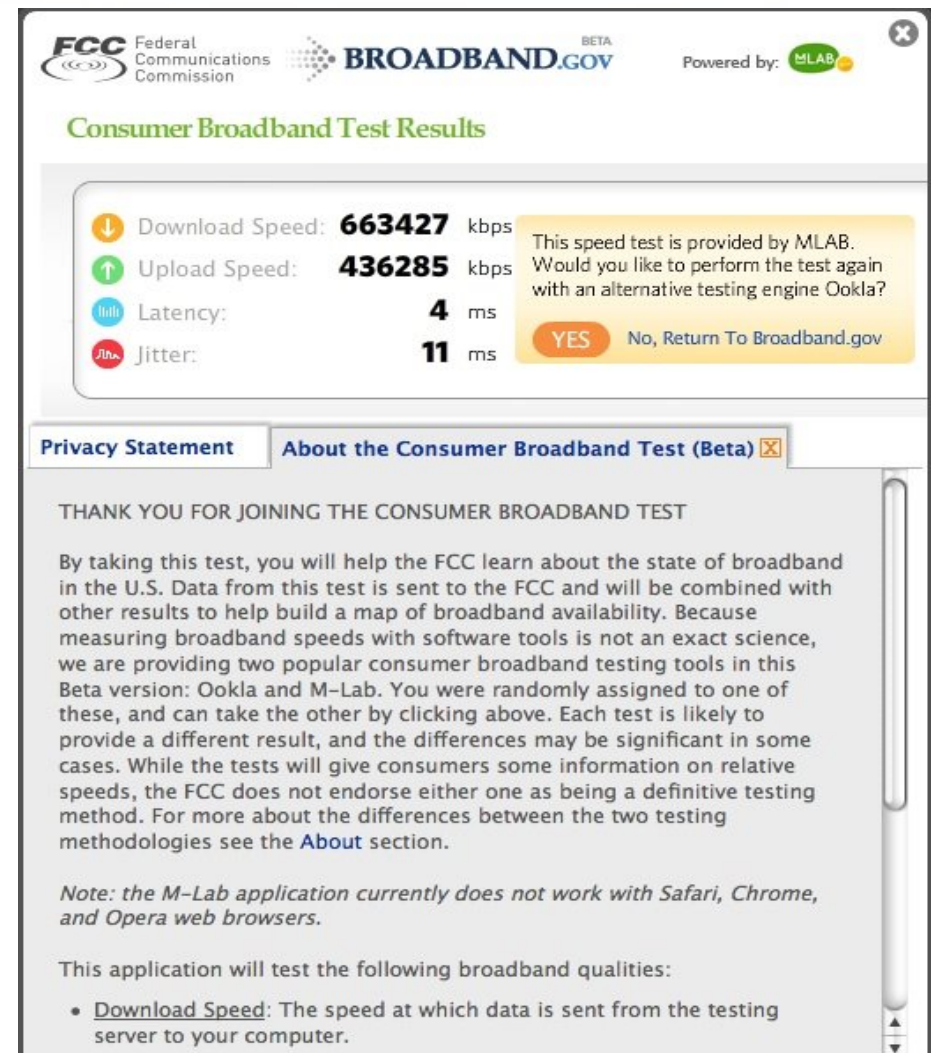


FCC's National Broadband Plan

On March 11 2010, FCC releases
Consumer Broadband Test
powered by NDT on M-Lab, and Ookla

FCC publishes initial results
at <http://blog.broadband.gov>

*National Broadband Plan recommends
expanding measurement and disclosure!*



FCC Federal Communications Commission **BROADBAND.GOV** BETA Powered by: **MLAB**

Consumer Broadband Test Results

Download Speed:	663427 kbps	This speed test is provided by MLAB. Would you like to perform the test again with an alternative testing engine Ookla?
Upload Speed:	436285 kbps	
Latency:	4 ms	<input type="button" value="YES"/> <input type="button" value="No, Return To Broadband.gov"/>
Jitter:	11 ms	

[Privacy Statement](#) [About the Consumer Broadband Test \(Beta\)](#)

THANK YOU FOR JOINING THE CONSUMER BROADBAND TEST

By taking this test, you will help the FCC learn about the state of broadband in the U.S. Data from this test is sent to the FCC and will be combined with other results to help build a map of broadband availability. Because measuring broadband speeds with software tools is not an exact science, we are providing two popular consumer broadband testing tools in this Beta version: Ookla and M-Lab. You were randomly assigned to one of these, and can take the other by clicking above. Each test is likely to provide a different result, and the differences may be significant in some cases. While the tests will give consumers some information on relative speeds, the FCC does not endorse either one as being a definitive testing method. For more about the differences between the two testing methodologies see the [About](#) section.

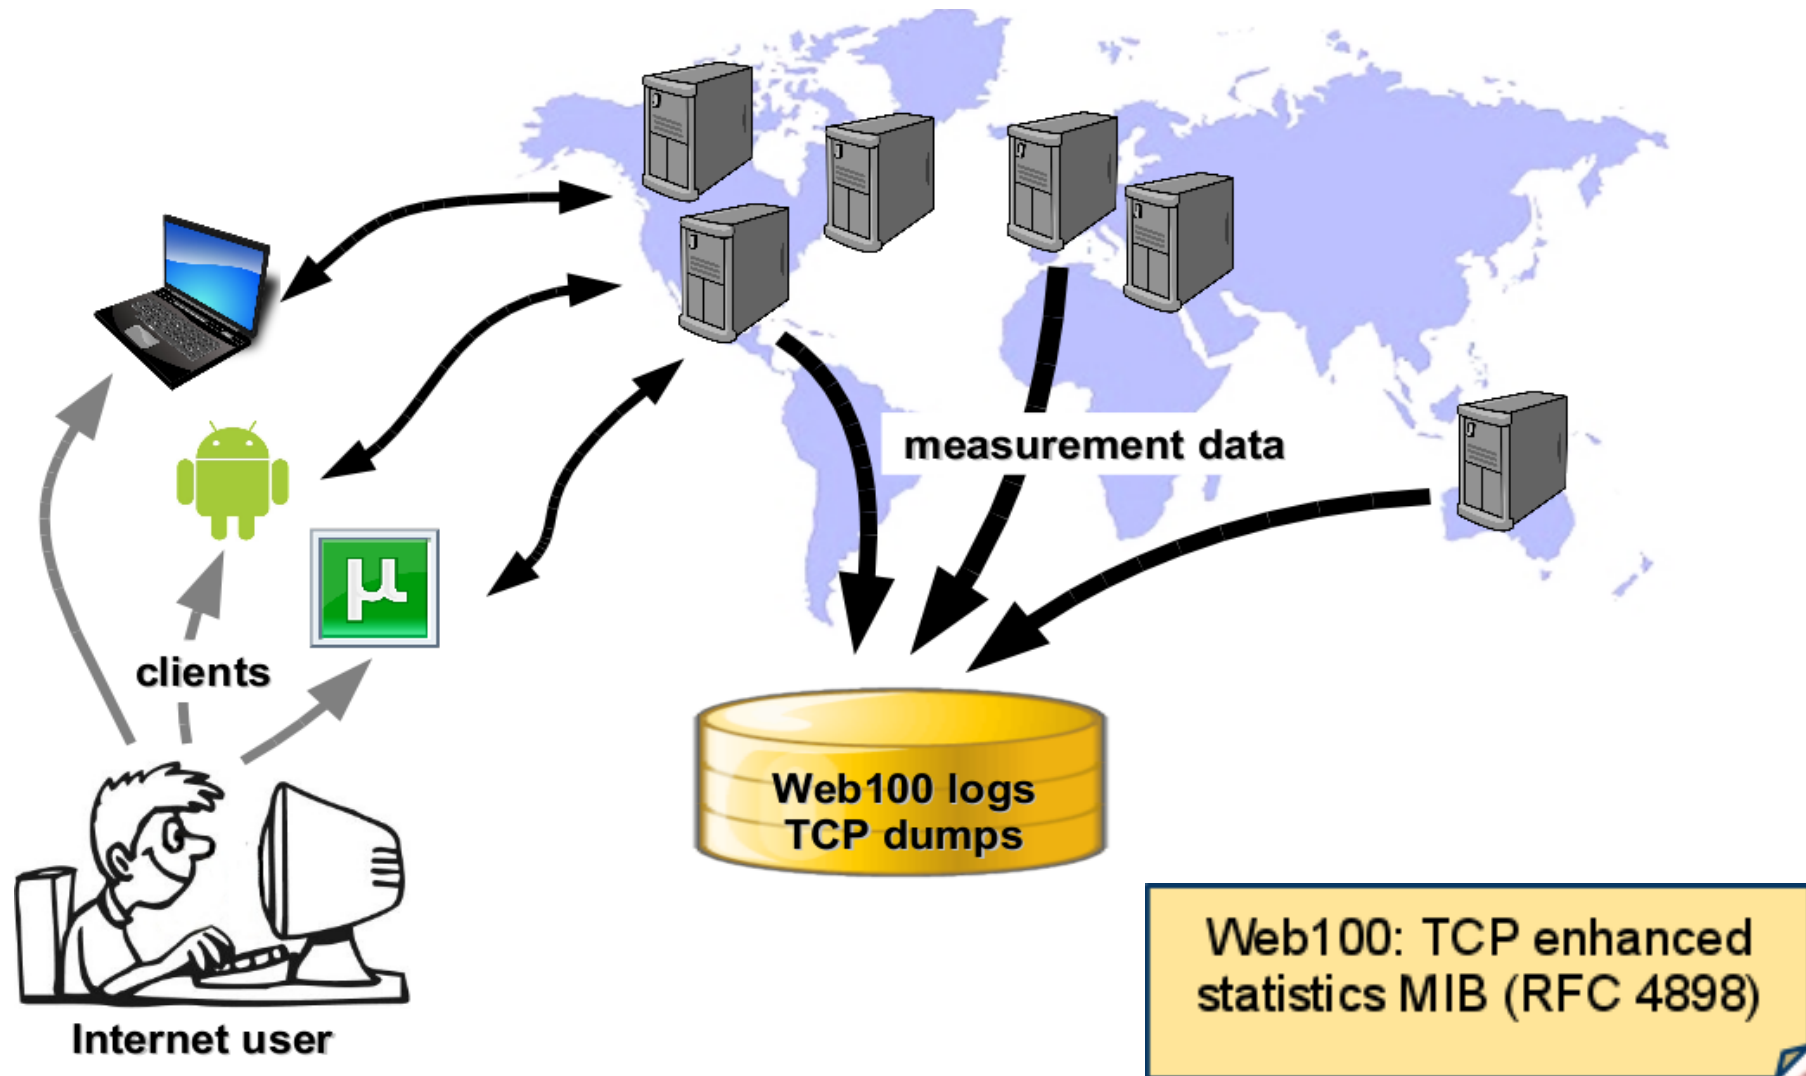
Note: the M-Lab application currently does not work with Safari, Chrome, and Opera web browsers.

This application will test the following broadband qualities:

- **Download Speed:** The speed at which data is sent from the testing server to your computer.



M-Lab platform & sample data flow



M-Lab platform

M-Lab is intended to:

- Provide **server-side resources** for **client-initiated active network measurement** of users' Internet connections.
- Provide **resource allocation** on each server that will ensure sufficient bandwidth and machine resources
- The M-Lab **platform** currently consists of **45 servers** in **15 locations**.
- The M-Lab **servers** run **tools** that instrument the server side of Internet application protocols.

The data flow

1. Whenever a user runs a test using a **tool**, the **client** on user machines communicates with the **server** in some predetermined way
2. Each **test** generates and sends some **data** back-and-forth between client and server.
3. The tools collect **data** related to the particular communication "flows" generated by the tool. *These do not monitor users' normal Internet traffic.*
4. All data collected through M-Lab are **publicly accessible**.
 - Data currently available through Amazon's EC2.
 - Currently NDT and NPAD only. Other tools' datasets are coming.



Currently available tools

- **Network Diagnostic Tool (NDT)**

Test user's connection speed and receive sophisticated diagnosis of problems limiting speed.

- Used by the **FCC Consumer Broadband Test**
- <http://www.broadband.gov/>
- Integrated with **µTorrent client**

- **Network Path and Application Diagnosis (NPAD)**

Diagnose common problems that impact last-mile broadband networks.

- **Glasnost**

Test whether traffic from user's applications is being throttled) or blocked.

- **Pathload2**

Test user's available bandwidth.

- **ShaperProbe**

Determine whether an ISP is performing traffic shaping.



NDT: What a user sees

TCP/Web100 Network Diagnostic Tool v3.6.1

click START to begin

** Starting test 1 of 1 **

Connecting to 'ndt.iupui.donar.measurement-lab.org' [ndt.iupui.donar.measurement-lab.org/64.9.225.154] to run test

Connected to: ndt.iupui.donar.measurement-lab.org -- Using IPv4 address

Checking for Middleboxes Done

checking for firewalls Done

running 10s outbound test (client-to-server [C2S]) 15.08Mb/s

running 10s inbound test (server-to-client [S2C]) 10.33Mb/s

The slowest link in the end-to-end path is a 45 Mbps T3/DS3 subnet

click START to re-test

START

Statistics

More Details...

Report Problem

----- Client System Details -----

OS data: Name = Linux, Architecture = i386, Version = 2.6.24-27-generic

Java data: Vendor = Sun Microsystems Inc., Version = 1.6.0_20

----- Web100 Detailed Analysis -----

45 Mbps T3/DS3 link found.

Link set to Full Duplex mode

No network congestion discovered.

Good network cable(s) found

Normal duplex operation found.

Web100 reports the Round trip time = 559.29 msec; the Packet size = 1360 Bytes; and No packet loss was observed.

S2C throughput test: Packet queuing detected: 1.94%

This connection is receiver limited 85.09% of the time.

Increasing the the client's receive buffer (770.0 KB) will improve performance

This connection is network limited 14.35% of the time.

Web100 reports TCP negotiated the optional Performance Settings to:

RFC 2018 Selective Acknowledgment: ON

RFC 896 Nagle Algorithm: ON

RFC 3168 Explicit Congestion Notification: OFF

RFC 1323 Time Stamping: ON

RFC 1323 Window Scaling: ON; Scaling Factors - Server=7, Client=7

Server 'ndt.iupui.donar.measurement-lab.org' is not behind a firewall. [Connection to the Client is probably behind a firewall. [Connection to the ephemeral port failed]

Information: Network Middlebox is modifying MSS variable

Server IP addresses are preserved End-to-End

Information: Network Address Translation (NAT) box is modifying the Client's IP address

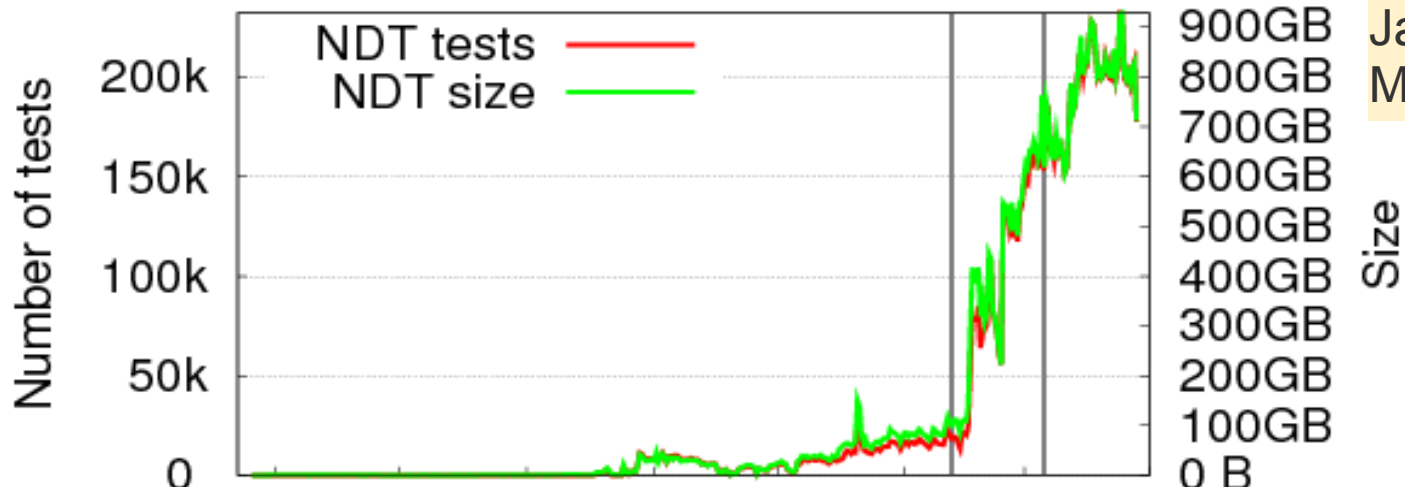
NDT: What a researcher sees



How much data (per day)?

How many tests (per day)?

µTorrent FCC

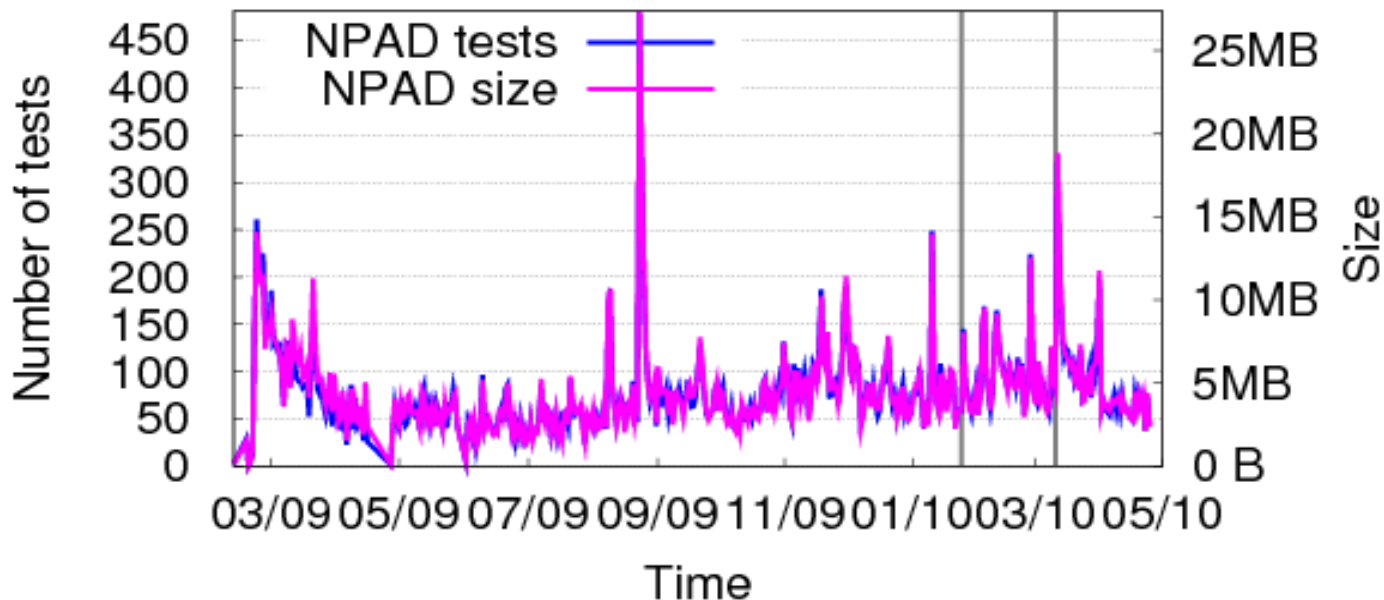


Jan 25 2010 uTorrent launch
Mar 11 2010 FCC launch

NDT

Tot tests: 15M

Tot size: 60TB

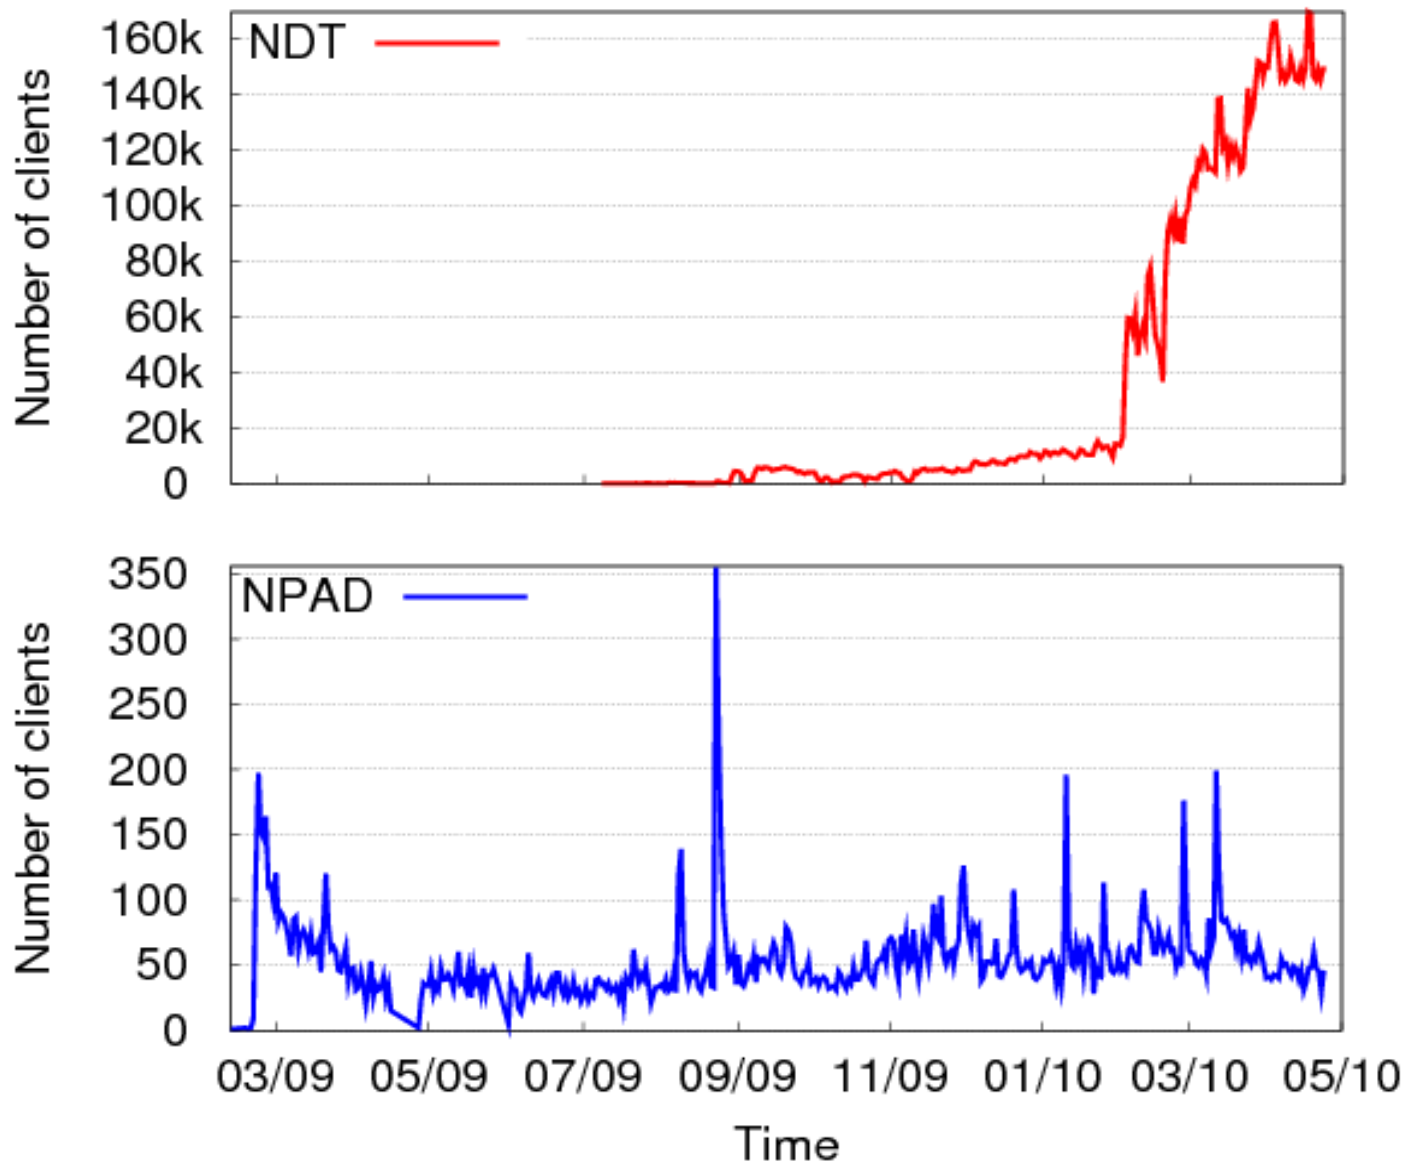


NPAD

Tot tests: 33k

Tot size: 2GB

How many clients (per day)?



Number of clients =
number of unique
client IP addresses

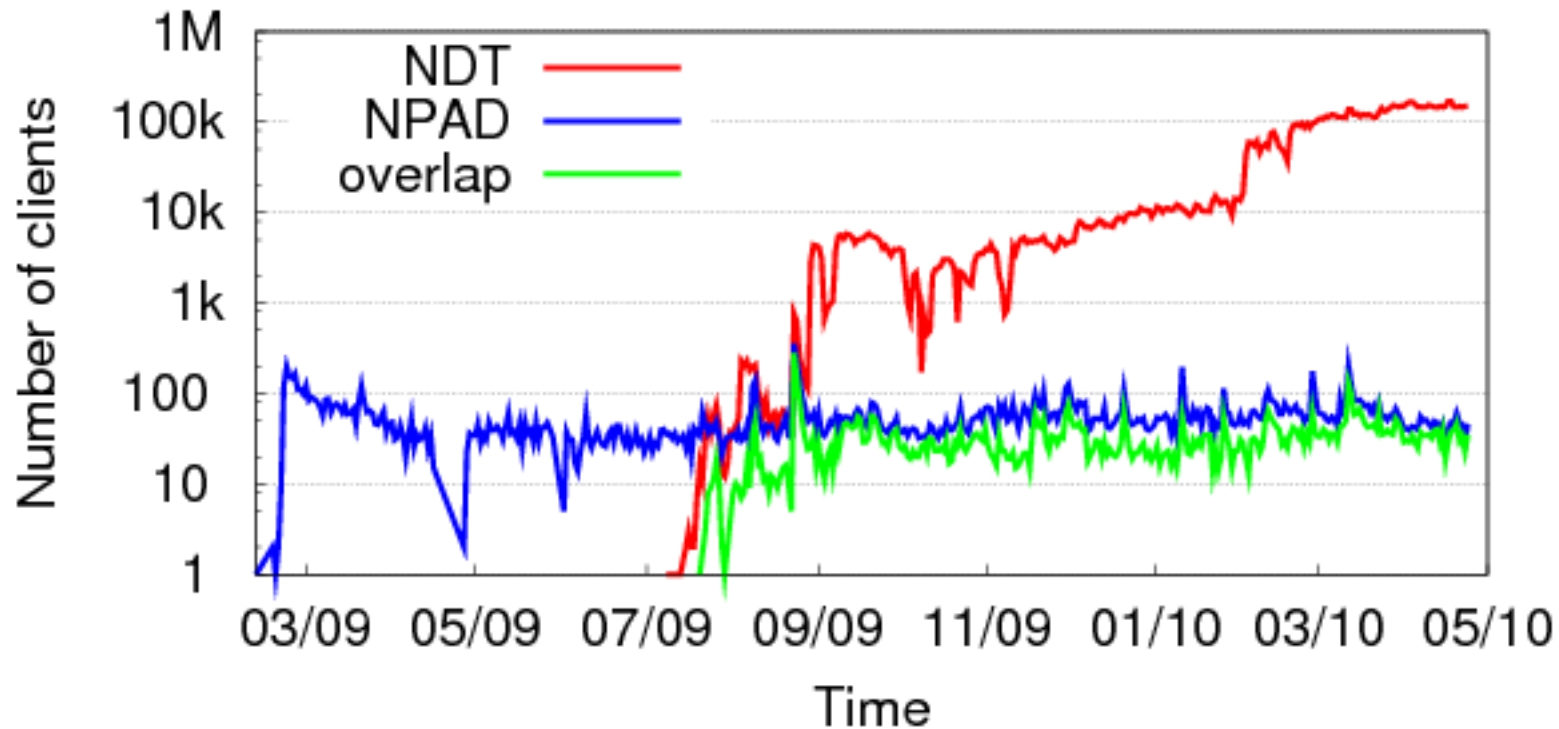
NDT

Tot clients: 8M

NPAD

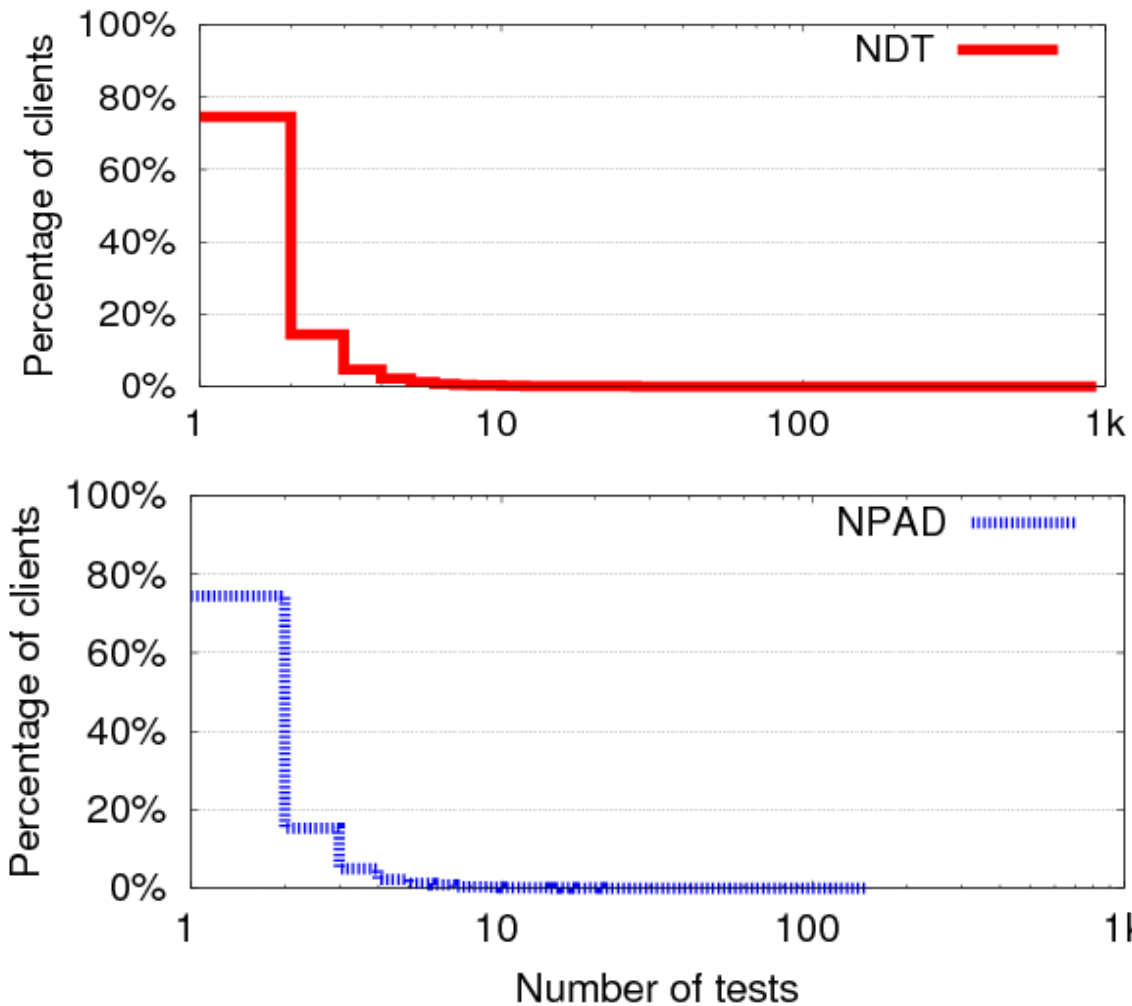
Tot clients: 21k

How many clients using both NDT & NPAD (within a day)?



Tot clients: 8M
Tot clients: 21k
Tot clients: 10k

How many tests per client since Feb09?



GET INVOLVED!

- **Analyze the collected data** and share your results with the Internet community.
- Contribute **tools** (client side, server side)
- Provide **servers** and **network connectivity** for the platform.
- Provide resources for **data hosting, aggregation and publication**.
- Provide **funding** to support the above.

A decorative header featuring a green sphere on the left and three overlapping semi-spheres in blue, red, and yellow on the right, all positioned above a horizontal line.

More information at
www.measurementlab.net

