

Measurement Lab: Overview and an Invitation to the Research Community

Constantine Dovrolis¹ Krishna Gummadi² Aleksandar Kuzmanovic³ Sascha D. Meinrath⁴

¹Georgia Tech ²MPI-SWS ³Northwestern University ⁴New America Foundation
{dovrolis@cc.gatech.edu gummadi@mpi-sws.mpg.de akuzma@cs.northwestern.edu meinrath@newamerica.net}

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.

The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

Measurement Lab (M-Lab) is an open, distributed server platform for researchers to deploy active Internet measurement tools. The goal of M-Lab is to advance network research and empower the public with useful information about their broadband connections. By enhancing Internet transparency, M-Lab helps sustain a healthy, innovative Internet. This article describes M-Lab's objectives, administrative organization, software and hardware infrastructure. It also provides an overview of the currently available measurement tools and datasets, and invites the broader networking research community to participate in the project.

Categories and Subject Descriptors: C.2.6 [Computer Communication Networks]: Internet

General Terms: Experimentation, Performance

Keywords: Network Measurement, Broadband Internet Access, Network Neutrality

1. INTRODUCTION

In 2008, Vint Cerf and others at Google initiated a broad discussion with networking researchers about major challenges in the effective study of broadband networks. Following this meeting, the New America Foundation's Open Technology Initiative (OTI) hosted a strategy convening where researchers identified several problems in this area, such as a lack of well-provisioned and well-connected measurement servers in geographically distributed areas. A second concern involved the difficulties in sharing large Internet measurement datasets between different research projects [1, 2, 3, 4]. Thirdly, decision-makers in Washington, DC were identifying the lack of broadband measurement data as a key limitation in their efforts to craft public policy. Through these discussions, Measurement Lab (M-Lab) was founded with an explicit objective to address the challenges identified by researchers and policy-makers.

M-Lab is an open, distributed server platform for researchers to deploy active Internet measurement tools, advance networking research, and empower the public with useful information about their broadband connections. By providing open access to M-Lab's data repositories to all interested parties, M-Lab helps sustain a healthy, innovative Internet. Since all data collected via M-Lab are made available to the broader research community, M-Lab is helping build a common pool of network measurement data, removing the need for every research project to collect its own data and facilitating cross-sample analyses that would not otherwise be

possible.

M-Lab also provides a platform to develop, test, and deploy new active measurement tools. The M-Lab platform uses a large number of purpose-built and well-connected measurement servers in strategic locations around the globe. Each tool is allocated dedicated resources on the M-Lab platform to facilitate accurate measurements. Server-side tools are openly licensed to allow third-parties to develop their own client-side measurement software.

Since its launch in January 2009, researchers have used M-Lab to study a wide range of broadband connection characteristics. Network scientists have already developed and deployed multiple new tools that allow users to test their broadband connections by briefly communicating with M-Lab measurement servers. For example, one increasingly important aspect of a broadband user's experience is how network management practices deployed by Internet Service Providers (ISPs) differentiate among types of traffic (i.e., by prioritizing certain services or discriminating against specific applications). In addition, researchers use M-Lab to run in-depth diagnostic testing that helps identify whether a performance problem is caused by the network, the user's end-host, or an application. Results from these tests are both provided to the users who run these tests and aggregated within the M-Lab data repository to help identify trends.

Today, M-Lab is only at the beginning of its development. Six tools are currently available, running on 45 servers in 15 locations around the globe. As a collaboration, M-Lab depends on the active support of partners from the research community, companies that are willing to host M-Lab servers or data, and application providers that can integrate M-Lab measurements in their software. By the end of 2009, M-Lab partners included scientists from Internet2, the Max Planck Institute, the Georgia Institute of Technology, and many additional organizations; data and server hosting by Amazon, Voxel, and the Hellenic Telecommunications and Post Commission; and client software development by SamKnows, BitTorrent, Virginia Tech's Enterprise GIS Research and Development Administration, among others.

Founding partners Google Inc., the PlanetLab Consortium, and OTI collectively make up the steering committee that leads the development of the organizational policies and structure of the platform going forward. As part of the "proof of concept" launch of M-Lab, Google provides servers and network connectivity for the platform. The PlanetLab Consortium provides M-Lab's software environment and acts as a point of contact for network scientists.

OTI coordinates M-Lab logistics, manages day-to-day operations, and works in Washington, DC to help formulate policy and regulatory reforms that utilize data and findings from the M-Lab project.

2. M-LAB INFRASTRUCTURE

M-Lab is an open, distributed server platform where researchers can deploy their active measurement tools. M-Lab has been developed in conjunction with the PlanetLab Consortium and so it has much in common with the well-known PlanetLab platform. M-Lab differs from PlanetLab in two key ways, however: first, the scope of tools and experiments permitted to run on M-Lab is limited to active broadband measurements, open source tools, and open data mandates; and, second, M-Lab resource allocation policies are designed to avoid resource contention between different experiments that could detrimentally impact their accuracy.

Hardware and network connectivity.

Like PlanetLab, M-Lab provides a set of geographically distributed servers on which researchers can deploy tools and run experiments. M-Lab servers are well provisioned for conducting high-bandwidth measurements at large scale. All M-Lab servers are “server-class” computers with an eight-core Intel Xeon processor running at 2 or more gigahertz and three or more gigabytes of main memory. Further, each server is connected to one or more ISPs with 1 gigabit per second upstream capacity.

Software.

M-Lab shares the same server software base with PlanetLab with some key differences. A researcher is assigned a “slice” of the server resources for her experiments and an experiment designed to run over PlanetLab servers could be run over M-Lab servers with little, if any, modification. Rather than repeat the description of the architecture and APIs, we refer interested readers to the PlanetLab design documents at <http://planet-lab.org/doc/pdn>. In addition to M-Lab’s openness and transparency requirements, there is another key difference between PlanetLab and M-Lab. Unlike PlanetLab, M-Lab provides a separate publicly addressable IP address for virtual network access to each slice. This helps avoid port contention between different experiments since, for example, by using different IPs, it becomes possible to run two experiments that both attempt to analyze the performance of BitTorrent on port 6881 on the same M-Lab server.

Administration, maintenance, and access.

M-Lab uses a similar administration, maintenance, and operations framework to PlanetLab. In fact, M-Lab can be thought of as a specialized version of PlanetLab. Researchers and network scientists that are interested in running their tools on the M-Lab platform should contact M-Lab’s steering committee, which coordinates research on the M-Lab platform. Once granted access, researchers can login and run their experiments on M-Lab servers just as they would on PlanetLab.

Resource allocation.

A key feature that distinguishes M-Lab from PlanetLab is its resource allocation policy. It is well known that experi-

ments over PlanetLab are susceptible to resource contention and server overload, largely because PlanetLab does not restrict the number of experiments that can simultaneously run on any given server. M-Lab, on the other hand, limits the number of tools that can run on each processor core to 1.5 to ensure that artifacts arising from M-Lab servers themselves are minimized. M-Lab currently implements the limit by manual admission control.

Deployment status.

As of April 2010, a total of 45 servers are operational across 15 geographically distributed sites in the United States and Europe; additional servers in Australia are expected to be online soon. More information on server operations and their locations can be found at: <http://measurementlab.net>. M-Lab servers are provided by various corporate, academic, and governmental entities, and deployment is expected to grow in the near future as new M-Lab partners join the collaboration.

3. CURRENT M-LAB TOOLS AND DATA

M-Lab tools span a variety of different applications, however they all share certain commonalities. The M-Lab platform is currently only running active network measurement tests (i.e., no passive monitoring of participant data is done). In addition, M-Lab conducts measurements between the user/client and the M-Lab measurement servers to examine the end-to-end performance characteristics along this entire path. As of April 2010, M-Lab hosts the following six active measurement tools.

Network Diagnostic Tool (NDT).

This tool measures TCP throughput between a client running at the user’s host and an M-Lab server. Test data is sent in both directions. While the test is running, the server also collects tcpdump and web100 snaplog files. Metadata containing information about the client, including OS type and version, is also collected. A test provides detailed packet level information along with kernel-level statistics on how the TCP connection performed in the given path. NDT is currently also being used by the FCC as its official broadband measurement software. A limited version of NDT is currently one of two broadband tests used by the FCC.

Network Path and Application Diagnosis (NPAD).

This tool uses TCP to measure end-to-end throughput and information about the switch/router queues along the path. web100 is used to limit the extent to which the measurement flow observes congestion caused by the measurement itself. The tool reports specific events and creates Web100 snaplog files, along with metadata, to the archive.

Glasnost.

This tool emulates a BitTorrent client and measures the performance back to a central server. It also detects if these application flows are being limited or disrupted by the ISP.

Pathload2.

This tool uses specially crafted UDP packet trains to measure the “available bandwidth” (also known as “residual capacity”) of the end-to-end path between the user’s client and an M-Lab server. It works in both directions and stores

the resulting measurements, together with statistics about packet delays and losses, in the archive.

ShaperProbe.

This tool uses specially crafted UDP packet streams to detect whether the user’s traffic is subject to traffic shaping. This network management practice means that the downstream or upstream access capacity drops significantly after a certain amount of transferred bytes and it is typically implemented with token buckets. ShaperProbe also measures the maximum downstream and upstream capacity in the measured path.

SideStream.

This tool measures TCP performance, and collects Web100 statistics, while the user browses the M-Lab web site. It runs on the same M-Lab slice with NPAD.

3.1 M-Lab data

The data collected on M-Lab provide a unique view of last-mile user experience. Instead of focusing on the Internet core, M-Lab focuses on end-to-end performance and on the characteristics of broadband access links. Measurements capture basic operational characteristics (e.g., TCP throughput, available bandwidth), advanced host diagnostics (e.g., misconfigurations, small socket buffer sizes), and ISP traffic management practices (e.g., BitTorrent blocking, traffic shaping). Analysis tools and techniques that can then be used to compare network performance across different users and ISPs are also in development. An additional goal is to provide the public, as well as Internet policy-makers, with empirically-based findings to help them make informed decisions about broadband and network management transparency.

All data collected through M-Lab is made publicly available and placed in the public domain. Currently, data from NDT and NPAD are available, and other tools’ datasets will become available in the near future. You can access these datasets through M-Lab’s public data sets and M-Lab is actively looking for additional options for making its data more accessible.

M-Lab’s steering committee is working to ensure that data are made available as promptly as possible. Occasionally, the M-Lab steering committee allows researchers to embargo their data briefly to enable them to have the first opportunity to analyze the information their tools collect. After a maximum of one year or when the researcher publishes a paper with findings from the corresponding data (whichever comes first), that data must be released publicly. However, the vast majority of M-Lab’s data is made publicly available as soon as possible after its collection. For example, NDT data is released under a *Creative Commons Zero Waiver* (<http://creativecommons.org/about/cc0>) and archived data is stored online as part of the *Amazon public dataset project* (<http://aws.amazon.com/publicdatasets>). The reader can visit this web site and click through the “Public Data Sets Catalog” link to browse through the data. Currently, NDT data is uploaded to the public repository in batches; however, M-Lab administrators hope to automate this process to update the repository daily.

Already, a host of interesting questions can be investigated with the available datasets include: (1) Has TCP autotuning, as found in Windows 7, MacOS 10.x, and Linux

2.6.x kernels, improved end-to-end network performance? (2) What traffic management policies are actually deployed by ISPs in broadband access networks? and (3) Are new protocols or services required to ensure that all users get fair Internet access? Interested researchers are encouraged to access M-Lab’s public data repositories, develop additional hypotheses, and conduct their own analysis on the available data.

4. M-LAB AND NETWORK RESEARCH

M-Lab brings benefits to both end-users and to the network research community. Specifically, researchers can use M-Lab to:

- (i) expose their measurement tools and systems to a large number of users,
 - (ii) validate their analytical and simulation models with data from real-world Internet paths,
 - (iii) share and analyze datasets collected on M-Lab,
 - (iv) avoid the administrative and operational overhead involved in managing a large-scale distributed server platform.
- We discuss these four issues in more depth below.

First, M-Lab helps researchers to make their active measurement tools available to a large number of users worldwide. The currently available tools are run by hundreds of thousands of users every month. In the past, it was quite difficult for a researcher to deploy a measurement tool to many highly available, well-connected servers around the world. PlanetLab was one of the only options, but those servers are often too heavily loaded, affecting the accuracy and precision of active measurements.

Second, the analysis of M-Lab data provides reliable information about the characteristics of Internet access networks and end-to-end performance. Instead of making assumptions about the available capacity, delays, losses, presence of traffic shapers, buffer sizes, etc., an analytical or simulation model can be grounded on measurements derived from the available M-Lab data.

Third, M-Lab can have a significant impact on how Internet measurement research is done. Currently, a researcher needs to collect her own data, and perhaps even develop her own measurement tool before investigating a given question. With M-Lab, we expect that it will be sufficient for many research questions to simply use the available M-Lab data, without facing the need to collect new measurements. In other words, M-Lab provides a researcher with the option to focus on data analysis, rather than data collection. Further, building data mining, aggregation and correlation tools that can effectively process data from multiple vantage points can provide insights far beyond what any single ISP or individual research project can do independently.

Finally, M-Lab addresses many of the challenges related to managing a large-scale distributed server platform and collecting a substantial amount of network measurements. In addition to providing operational support and access to sufficient resources, M-Lab resolves many of the administrative and legal problems that have hindered effective data sharing. The M-Lab data sharing method in which users donate their network traces to the platform in exchange for information about their connection’s performance, effectively resolves this problem. Further, the M-Lab tools do not collect sensitive data (i.e., M-Lab uses pre-determined “active tests” rather than “passive tests” which may reveal information about the participants’ online activities, applica-

tion use, etc.), helping alleviate concerns about privacy and confidentiality that other client-based measurement projects may face.

5. REFERENCES

- [1] M. Allman, E. Blanton, and W. Eddy. A Scalable System for Sharing Internet Measurements. In *Proceedings of Passive and Active Measurement (PAM) Conference*, 2002.
- [2] M. Allman and V. Paxson. Issues and Etiquette Concerning Use of Shared Measurement Data. In *Proceedings of ACM Internet Measurement Conference (IMC)*, 2007.
- [3] k. Claffy, M. Crovella, T. Friedman, C. Shannon, and N. Spring. Community-Oriented Network Measurement Infrastructure (CONMI) Workshop Report. *ACM SIGCOMM Computer Communications Review (CCR)*, pages 41–48, 2006.
- [4] V. Paxson. Strategies for Sound Internet Measurement. In *Proceedings of ACM Internet Measurement Conference (IMC)*, 2004.