

Measurement Lab (M-Lab) is an open, distributed server platform on which researchers can deploy Internet measurement tools. The goal of M-Lab is to advance research and empower the public with useful information about the characteristics of their broadband connections. By enhancing Internet transparency, we aim to help sustain a healthy, innovative Internet.

This working document outlines the general motivations, organization, and technical specification for Measurement Lab (M-Lab).

Document history

Version 1.0, 28-01-2009

Contents:

- 1 - Measurement Lab Defined
 - 1.1 - Background and Motivation
 - 1.2 - Goals
 - 1.3 - Scope
 - 1.4 - Roles
 - 1.5 - Pilot Project and Future Development

- 2 - Organization, Policies and Processes
 - 2.1 - Organization
 - 2.2 - Processes
 - 2.3 - Open Access
 - 2.3.1 - Data Access Policy
 - 2.3.2 - Software/Protocol Access Policy

- 3 - Technical Specifications of M-Lab Server Platform
 - 3.1 - Initial Node Configuration
 - 3.1.1 - Bill of Materials
 - 3.2 - Site Requirements
 - 3.2.1 - Physical Requirements
 - 3.2.1.1 - Rack Space
 - 3.2.1.2 - Power
 - 3.2.1.3 - Cooling
 - 3.2.1.4 - Site-supplied Hardware
 - 3.2.2 - Logical Requirements
 - 3.2.2.1 - IP Address Space and Bandwidth
 - 3.2.2.2 - DNS
 - 3.2.2.3 - Firewall and ACL Considerations
 - 3.3 - Installation
 - 3.3.1 - Initial Installation
 - 3.3.1.1 - ...
...(enumerating TBD installation info per location)...
 - 3.3.2 - Adding Nodes to an Existing Site

1 Measurement Lab Defined

M-Lab is an open, distributed server platform on which participating researchers can deploy Internet measurement tools.

1.1 Background and Motivation

In order to study broadband networks, researchers are currently developing a variety of tools, which allow end-users to measure varying aspects of their connection. These tools generate a pre-determined communication between a client and a server instrumented with software that emulates Internet application protocols. As purpose-built experiments, the tools do not monitor a user's other, normal Internet traffic.

In 2008, Vint Cerf and others at Google initiated conversations with network researchers to learn more about challenges to effective research in this area and discuss ways to address them. Researchers identified several problems, including the a lack of widely-deployed servers and connectivity that can be used for active network measurement tools. Deploying their own servers in a variety of locations is expensive and impractical. Due to limited server resources, only a limited number of users can run the tests at a given time.

In addition, researchers lack the means to easily share large data sets with one another. Thus, it can be difficult for researchers to build new experiments on previous network measurement experiments, compare or independently verify results, or publicize data.

M-Lab was founded by New America Foundation's Open Technology Institute (OTI), PlanetLab Consortium, Google Inc., and academic researchers in order to help address these barriers to research.

1.2 Goals

The goal of M-Lab is to advance network research and empower the public with useful information about the characteristics of their broadband connections. By enhancing Internet transparency, we aim to help sustain a healthy, innovative Internet.

M-Lab will deploy server infrastructure and connectivity for active network measurement tools, and dedicate sufficient resources on these servers to particular tools so that they can conduct accurate measurements.

In addition, M-Lab aims to develop parallel, related technical resources, such as means for data publication and sharing.

M-Lab is intended to serve as an open platform, defined in terms of the following objectives:

- Support as wide an array of measurement tools as possible.
- All data collected through M-Lab should be made publicly available and placed in the public domain.
- All researchers will be required to publish their source code.
- All researchers' server-side tools will be openly licensed and operated in a way that allows third-parties to develop client-side software for measurements.

M-Lab is intended to operate as a community-based effort. In order for M-Lab to achieve its objectives, it will need the support and participation of additional companies, institutions and researchers. M-Lab welcomes the support of other companies and researchers who would like to help expand the platform and ensure its growth and success.

1.3 Scope

Measurement Lab is developed in conjunction with PlanetLab and seeks to build on the operations, administration, and maintenance (OA&M) framework that the PlanetLab Consortium (PLC) has already usefully developed.

However, M-Lab's servers are separate and distinct from PlanetLab's existing overlay network, and the two platforms differ in key ways. PlanetLab allows slices to run a variety of network experiments and cannot guarantee that a server is sufficiently provisioned to support a given measurement tool. In contrast, M-Lab's platform is designed to:

- Provide passive server-side resources for client-initiated active network measurement of Internet users' broadband connections.
- Provide resource allocation on each server that will ensure sufficient bandwidth and machine resources.

At the outset, we limit oversubscription to 1.5 tools per every one core on the machines. Each site will maintain three servers, which will each have 8 cores per machine.

As stated in 1.2, M-Lab will *only* be used for active network measurements -- purpose-built experiments -- and will not collect or store data from passive monitoring of users' other Internet traffic.

1.4 Roles

Participants in M-Lab serve a number of roles.

- **Hosting sites:** Companies, academic institutions, and other organizations provide and host servers ("sites"), and provide them with connectivity.
- **Software environment, OA&M:** M-Lab will essentially "outsource" OA&M of M-Lab nodes to PLC, including:
 - Management and maintenance of the OS/VM software on the nodes.
 - Packaging and distribution of tools to nodes (per M-Lab's specification of which tools run on the nodes).
 - Responding to patches, need for reboots, etc.
 - Acting as point of first contact for parties having questions about the activities of the M-Lab.
 This means, for example, that RADB or SWIP lookups for M-Lab IP address space should refer to PlanetLab; likewise with DNS domain name registration.
 M-Lab will be its own management authority (MA), and does not have a separate slice authority.
- **Measurement Tool Development:** Individual researchers will create active testing tools that run on the servers (instrumented versions of application protocols). Researchers whose tools are approved and meet M-Lab's open access goals (described below, section 2) are then granted dedicated resources on M-Lab in which to deploy their tool.
- **Data Hosting, Analysis, and Other Resources:** Organizations and institutions can contribute to M-Lab by providing data aggregation and publishing resources and by collecting and archiving data from M-Lab's servers ("repository"). They might also provide other technical resources, such as computing resources for data analysis.
- **Management & Organization:** M-Lab will require some organization, whether formal or informal, to define and implement the platform's policies and processes in light of the project's objectives. For the purposes of this document, we will refer to this organizing role as "Measurement Lab Consortium" (MLC). The MLC is currently comprised of M-Lab's charter members, and will grow and expand as M-Lab matures.
- **User and researcher outreach:** M-Lab will direct users to the various tools running on the platform, through a public facing website, mailing list, and other means. M-Lab will also seek to actively engage the research community in order to extend the platform.

1.5 Pilot Project and Future Development

The initial founders of M-Lab are currently developing a "prototype" of the platform. Rather than trying to define and achieve all aspects of this project at once, our goal in this stage is to provide some immediate value to researchers who are currently building active testing tools, develop a proof of concept, and learn from that process.

The prototype phase is comprised of the following:

- **Hosting sites:** Google is providing 36 servers across 12 locations, to be deployed January through June of 2009, in line with the Technical Specification [section 3]. Google is also purchasing connectivity for these machines.
- **Software environment, OA&M:** PLC will fill this role.
- **Researchers:** The initial subset of tools are listed here: <http://measurementlab.net/measurement-lab-tools>. Researchers will be expected to be existing PlanetLab users (or, rather, part of research organizations that are PlanetLab members).
- **Data Hosting, Analysis, and Other Resource Providers:** No repository will exist at the outset. Researchers will be able to regularly download their data from the testing servers. As an interim solution, Google may collect and store an archive of the data, solely for the purpose of making it publicly available once a repository in line with M-Lab's objectives is created.
- **Management & Organization:** The founding members of M-Lab (see: <http://measurementlab.net/who>) will operate an ad-hoc steering committee comprised of relevant researchers for the purpose of developing the platform to a functional state and creating policies, processes and organizational models for the MLC.
- **Outreach:** The Open Technology Institute will manage a public-facing website as well as a public mailing list for discussion

M-Lab's future growth depends on the contribution and support of a wide community. The founding members of M-Lab welcome the participation, input and support of other companies, institutions, and researchers. Researchers or other entities that wish to support the platform, deploy a measurement tool, or otherwise provide input, should contact the steering committee: <http://measurementlab.net/contact>

2. Organization, Policies and Processes

2.1 Organization

The MLC aims to have the only the minimum amount of structure required to address processes and policies; the working assumption is that "less is more." As M-Lab grows, make-up of the MLC will evolve concurrently. Future possibilities include:

- An informal collaboration among researchers, companies, institutions.
- A formal foundation or nonprofit organization.
- An entity along the lines of the PlanetLab Consortium.

2.2 Processes

Some processes need to be defined in order to determine how server resources will be allocated to ensure data collected server-side is high-quality (for example, accurate timestamps, availability of RFC4898 TCP Extended Statistics MIB).

The following process still need to be formed and finalized:

- Process for determining what runs on the nodes, balancing the needs of “production” applications with the need to test new versions of existing applications as well as new applications.
- Process for adding to the set of tools, and changing the members of the set to produce the desired data.
- Process for scheduling which application runs where. As soon as there are more than 1.5 tools per core, and as soon as it's clear which tools generate the most traffic, resource allocation decisions will need to be made.

Additionally, M-Lab will also need to define requirements and processes for adding new members in various roles:

- Providers of tools to run on M-Lab nodes.
- Providers of new M-Lab nodes.
- Providers of data repository services.

At the outset, M-Lab users will be expected to be a subset of PlanetLab users. In other words, deploying a tool on M-Lab will require that a researcher (or her/his institution) be a member of PlanetLab. However, the MLC will consider policies to incorporate non-PlanetLab users as well.

Anyone wishing to deploy a tool on M-Lab should contact the M-Lab steering committee: <http://measurementlab.net/contact>

2.3 Open Access

The MLC will work to ensure that access to M-Lab data and source code is open, in accordance with the strategic goals documented above.

It will do so in adherence to policies that balance the interests of primary research investigators to publish their results first, as well as the broader research community's interest in independently verifying published results and accessing data.

2.3.1 Data Access Policy

Data obtained using M-Lab are required to be placed in the public domain as soon as is reasonably feasible. Individual tools will have the option of choosing from one of two policies to be enforced by the MLC and any repositories:

- Data are released as collected.
- Data are made available for the exclusive use of the original investigators until a paper is published using the data, or until twelve months after collection, whichever comes first. Afterward, the data are released to the public domain.

2.3.2 Software/Protocol Access Policy

In line with M-Lab's objectives, tools using M-Lab are required to publish their source code for client and server software such that it can be reviewed by third parties. Ideally, this would be implemented as read-only access to source code management systems (CVS, SVN, etc.) so that changes over time can be tracked. During M-Lab's initial prototype phase, source code for all tools may not be available publicly at the outset.

In order to encourage independent innovation of data collection on the client side (and broaden the base of clients generating measurement data on the M-Lab server side), it is a goal that third-party developers be able to use M-Lab for their clients and measurements, even if they themselves did not create the server-side code. For example, independently-developed client software may want to incorporate M-Lab client code for interacting with M-Lab servers and integrate the resulting client-side measurements with other measurement data.

Toward that end, the communication protocols between clients and servers should be documented and license for IP granted in a manner that allows client code to be developed and communicate with M-Lab servers, independent of the researchers who created the server-side tool.

3 Technical Specification of M-Lab Server Platform

3.1 Initial Node Configuration

Figure 1 shows a schematic level diagram for a typical 3-node site. This initial configuration will require 6 1U rack positions as shown in Figure 2.

3.1.1 Bill of Materials

Table 1 gives the bill of materials required to build a site.

Item	Qty.	Rack Space	Description
1	3	3 x 1U	Server hardware (with rack-mount kit), dual-processor, quad-core
2	1	1U	Rack switch
3	1	1U	Power strip
4	6		Ethernet cables, Cat5e or better (servers to rack switch)
5	3		Power cords (server to power strip)

Item	Qty.	Rack Space	Description
6	1		Power cord (switch to power strip)

Table 1

3.2 Site Requirements

3.2.1 Physical Requirements

3.2.1.1 Rack Space

The installation requires 6 contiguous 1U rack positions on a standard, 19-inch, 2-post rack. Servers will be required to have center-mount hardware. Lighter components (rack switch, power strip) may have front-mount hardware. A gap of 1U between the power strip and rack switch will be used for front-to-back cabling.

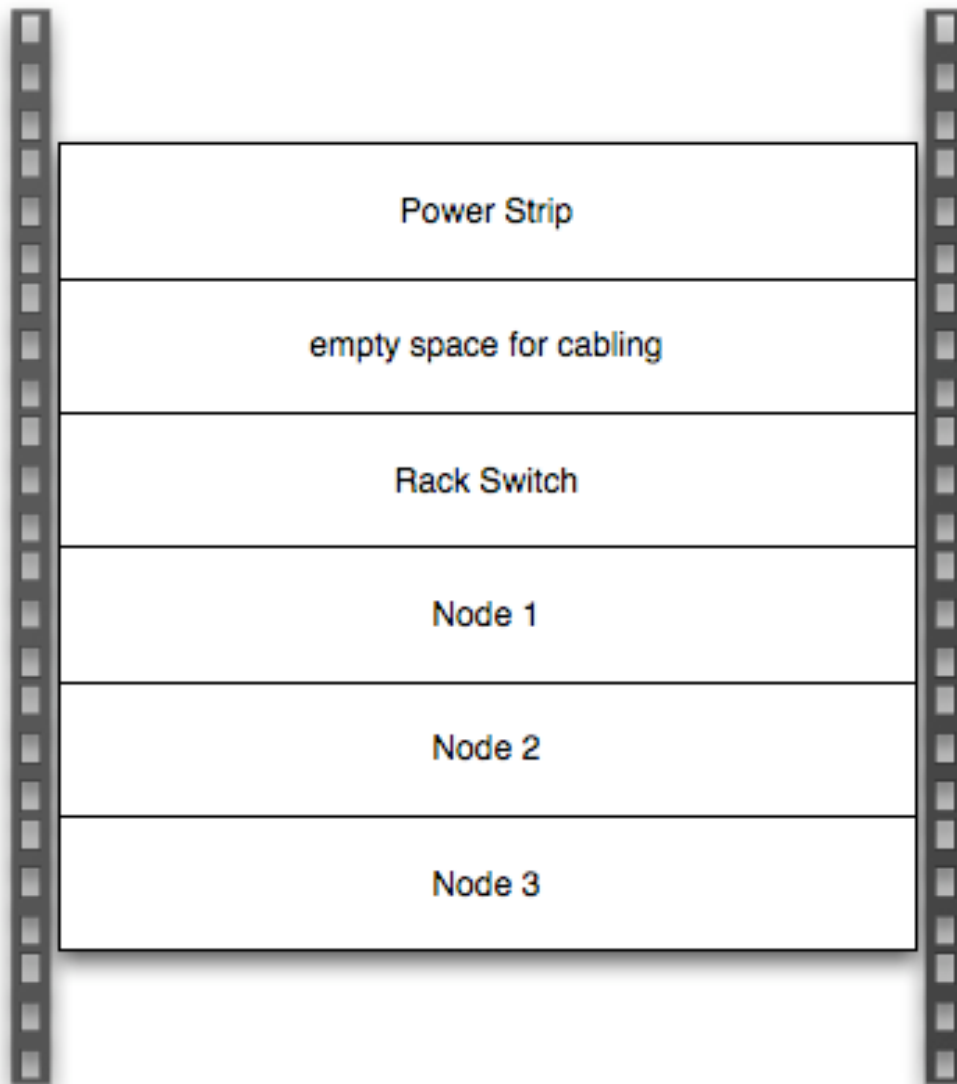


Figure 2: Rack Elevation. The gap between the power strip and the rack switch will be used for front-to-back cabling.

3.2.1.2 Power

One AC outlet is required. The recommended server option has a 670W power supply, requiring a (standard) 20A circuit to support 3 nodes and a rack switch.

3.2.1.3 Cooling

Required.

3.2.1.4 Site-supplied Hardware

An Ethernet cable of sufficient length to connect the upstream ISP to the rack switch must be provided on-site.

3.2.2 Logical Requirements

3.2.2.1 IP Address Space and Bandwidth

Each VM running on a node is required to have its own IP address for security purposes. With 8-core nodes and a 1.5:1 oversubscription of VMs to cores, we budget for 12 IP addresses for VMs, plus 2 for overhead. Three such nodes require a /26 CIDR block, which leaves room for an optional 4th node. An IPv6 /64 block is highly desirable.

Each site requires a 1G connection to the upstream ISP.

Note that the nodes may send to the broadcast address and/or collect packet traces. The subnet and associated L2 broadcast domain should be dedicated to the site's nodes and switch.

3.2.2.2 DNS

Yes.

3.2.2.3 Firewall and ACL Considerations

The M-Lab subnet should not be firewalled or otherwise constrained. This specifically includes TCP, UDP, and ICMP, but any Internet protocol packet should be allowed in and out.

3.3 Installation

This section documents the initial M-Lab sites, and processes for changing sites.

3.3.1 Initial Installation

3.3.1.1 Mountain View, CA, US

3.3.1.2 Atlanta, GA, US

3.3.1.3 Los Angeles, CA, US

3.3.1.4 Miami, FL, US

3.3.1.5 Chicago, IL, US

3.3.1.6 Portland, OR, US

3.3.1.7 Dallas, TX, US

3.3.1.8 New York, NY, US

3.3.1.9 Europe TBD

3.3.1.10 Europe TBD

3.3.1.11 Europe TBD

3.3.1.12 Europe TBD

3.3.2 Adding Nodes To An Existing Site

TBD.

3.3.3 Moving An Entire Site

TBD.