



measurementlab.net

M-Lab is an open, distributed platform of servers for researchers to deploy open source Internet measurement tools. By enhancing Internet transparency, M-Lab helps sustain a healthy, innovative Internet.

How to Contribute to M-Lab's Server Infrastructure

M-Lab supports scientific research by providing a widely-distributed server platform with ample resources and connectivity to run broadband and mobile measurement tools.

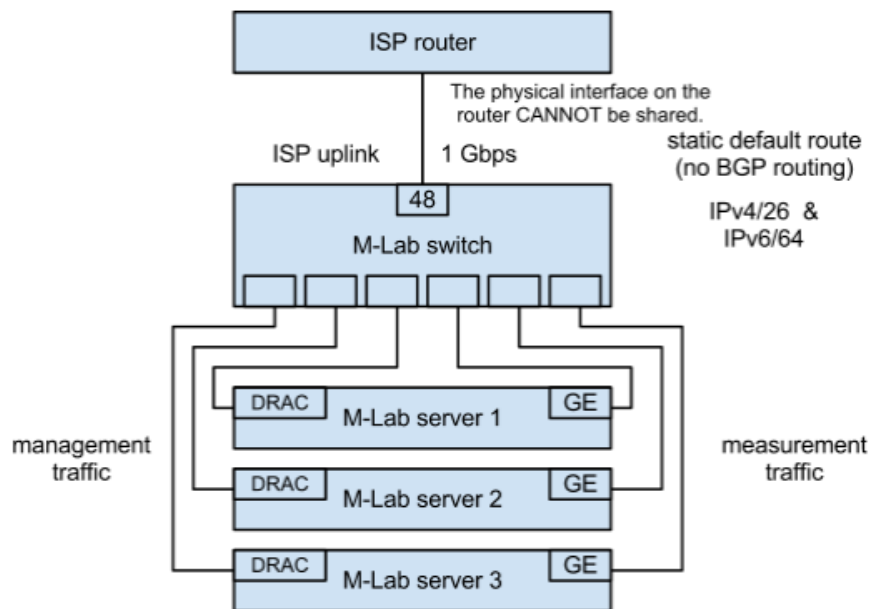
Some measurement tools are latency-sensitive. Therefore, they benefit from having servers close to their users. As a consequence, by providing local servers for M-Lab, you can help facilitate more accurate measurements in your area.

This document describes the hardware and the software of an M-Lab node as well as the connectivity requirements to host an M-Lab node. If you are interested in contributing servers and/or connectivity to M-Lab's server infrastructure, please [contact us](#), even if you are not sure whether you can meet the requirements.

Architecture of an M-Lab site

A single M-Lab site comprises 4 devices: 3 servers and 1 switch.

Every server is managed remotely via a remote console (such as a DRAC, Dell Remote Access Card) which is independently cabled to the switch. The switch is used to connect the servers to the Internet.



Physical requirements

- **Rack space:**
 - 4 Rack Units (RUs). It's preferable that the RUs are contiguous.
 - 2 extra RUs are also recommended, but not required, 1 for the power strip and 1 between the power strip and the switch, for back-to-front cabling.
- **Power:**
 - One AC outlet is required.
 - The amount of power required depends on the model of the switch and the servers.
- **Cooling:**
 - Required. The amount of heat generated depends on the model of switch and servers.

Connectivity requirements

- **IP addresses:**
 - A /26 (or shorter) IPv4 prefix, globally routed and dedicated to the site.
 - A /64 (or shorter) IPv6 prefix, globally routed and dedicated to the site.
- **Routing:**
 - A static route must be provided to the M-Lab site as the M-Lab switch does not support BGP routing. The upstream network is responsible for making all routing decisions and handling all packets, in and out of the M-Lab site network.
 - All the servers of the same M-Lab site must be located in the same LAN segment and must have the same upstream provider.
 - An IXP can host an M-Lab site only if the IXP has its own AS and upstream provider.
- **Upstream bandwidth:**
 - *Preferable* a minimum of 1Gbps dedicated upstream connectivity, to both local and international traffic.
 - When 1Gbps is not available, an M-Lab site still needs dedicated bandwidth that is at least significantly higher than what is available to common Internet users in the hosting country, to be able to run several simultaneous tests and make sure that the tests are not limited by the server upstream.
 - The traffic to the switch must not go through local firewalls or other network traffic filters.
 - Research and Education Networks (RENs) are usually not a good fit due to possibly limited connectivity to non-REN ISPs.
- **Upstream port:**
 - The port of the upstream router that is connected to the M-Lab switch cannot be shared and must be 1Gbps, even when the upstream bandwidth is less than 1Gbps.
- Within a specific geographic area (e.g., a country), we prefer an ISP/organization that is well connected to all (or most of) the ISPs that have customers in the same area.
- An M-Lab site must be hosted by an ISP or by another organization that manages its own network.
- M-Lab sites are typically co-located in the same building as major network provider's routers, typically within exchange points or in a carrier or provider's Point of Presence (POP).

On-site support requirements

- The M-Lab team monitors and maintains the site remotely.
- Facility support is only required when faults cannot be addressed remotely, like power cycling a server and replacing hardware components.

Hardware specifications

- **3 enterprise-class rack servers.**
 - We prefer 3 servers for each site in order to have redundancy and flexibility in distributing tools across multiple servers within the same site and across different sites.
 - Each server must have 8 CPU cores (at least 2GHz), 8 GB RAM, and dual 100+ GB hard drives.
 - Each server must support standard remote management functionalities, including remote console and remote power reset/reboot. Supported models are:
 - Dell Remote Access Controller (DRAC),
 - HP Integrated Lights Out (iLO),
 - Open Intelligent Platform Management Interface (IPMI), or
 - IBM Integrated Management Module (IMM).
 - We currently support the following models (in order of preference):
 - Dell PowerEdge R620, R610, R420, R410, 1950
 - ProLiant DL360 G7
 - We maybe be able to support other models. However
 - Using a model that is not yet supported requires to first test the compatibility of the new model with M-Lab's software.
 - Until the compatibility test is complete, we cannot ensure that we are going to be able to support a new model
 - Testing a new model could slow down significantly the deployment of a site.
- **1 high-performance, dedicated, rack switch.**
 - The switch will be managed by M-Lab and must provide SNMPv2 and an SSHv2-reachable CLI to support monitoring and remote configuration.
 - We currently only support the following models (in order of preference):
 - HP ProCurve 2848
 - Cisco C2960, C3560E
 - We maybe be able to support other models. However
 - Using a model that is not yet supported requires to first test the compatibility of the new model with M-Lab's software.
 - Until the compatibility test is complete, we cannot ensure that we are going to be able to support a new model.
 - Testing a new model could slow down significantly the deployment of a site.

Software specifications

The M-Lab platform is a centrally managed, globally distributed system based on the [PlanetLab Operating System](#). Each server runs the latest [CentOS](#) 6.4 with a custom kernel patched for [Web100](#) instrumentation and [Linux-VServer](#) virtualization. Web100 reveals additional TCP characteristics as described by RFC 4898. Linux-VServer slices each server into twelve container-based VMs running experiment services with dedicated public IP addresses. These experiments record measurements in response to user-initiated tests

Operation

- **Remote Control**
 - M-Lab uses a read-only boot image in all servers to enable whole-system remote control.
 - Before equipment is on the ground, new site and server information (DNS records, IP addresses, geo-location, etc) are registered in the M-Lab configuration database.
 - For each server, M-Lab creates a boot image with network configuration and a unique key.
 - The boot image is written to a read-only medium, such as CD-ROM, and permanently installed in the server.
 - On each restart, the server boots from the read-only image, which guarantees that the server is in a known state.
 - Using the key, the server authenticates with the M-Lab central management service before entering a production state.
- **Security**
 - Switch management is restricted to SSH access only from the M-Lab management network.
 - SNMP v2c community string (read-only) is random and known only to M-Lab operators.
 - The M-Lab server operating system is based on [CentOS](#) 6.4. Platform and security updates are taken from upstream mirrors daily.
 - All traffic to and from the root context travels over SSH, https, or OpenVPN tunnel.
 - The only externally facing service running in the root context is SSH on a non-standard port. SSH logins are restricted to users with registered ssh-keys.
 - Linux-VServer isolates the root context from experiment services. While the system uses a shared kernel, each experiment service runs inside a separate VM and is isolated in most other respects: IP address, resource allocation, file system.
- **Monitoring**
 - M-Lab continuously monitors site, server, and experiment activity.
 - From the switch, we record aggregate traffic across each switch port using SNMP v2c and monitor configuration changes using [RANCID](#). When unscheduled restarts or changes occur, we are notified immediately.
 - [smokeping](#) monitors losses and sudden increases in latency to all site equipment.
 - We also use [Nagios](#) to check correct operation and system integrity on servers.
 - [Zabbix](#) records resource usage of all experiment services. Resource usage allows more advanced analysis of measurement data and highlights trends over time, alerting us to potentially unusual activity.