



Procedure for approval of new experiments and allocation of slices

MeasurementLab (M-Lab) is an open platform for research, and welcomes the involvement of new participants. However, in the interests of managing the use of resources and steering the mission of the project, any experiment running on M-Lab must adhere to the following underlying principles, unless waived by the Steering Committee due to unusual circumstances:

- **The tools being used must be open-source.** The software running the experiment must be released under a permissive license approved by the Free Software Foundation (FSF), Open Source Initiative (OSI), or Creative Commons (CC). This ensures that its source code and methodology can be examined by interested parties, and that the client-side tools can be integrated into a variety of applications. Each experiment must include client as well as server code, all of which must be licensed as open source. Only the original experiment, when deployed on the platform initially, must openly license client-side code as a reference implementation. The client-side reference implementation can be separate from the version actually deployed to users, as long as it provides a reference for the application-layer protocol deployed. The purpose of open sourcing the client implementation is to allow third parties to more easily integrate the code into their own custom clients. The source code to the tools must be released either immediately upon the tool's launch or within the shorter of the following two periods: [1] a research paper is published utilizing data from the tool as run on the M-Lab platform or [2] one year from the time of the tool's launch.
- **The data must be open-sourced.** All data produced by the experiment running on the M-Lab platform must be released into the public domain (using a Creative Commons Zero (CC0) license) and made available for free in a portable format so that it may be examined and freely used for research purposes for the foreseeable future. The data must be released either immediately upon its collection or within the shorter of the two periods: [1] a research paper is published utilizing [1] data from the tool as run on the M-Lab platform or [2] one year from the time of the tool's launch. Any data collected on the M-Lab platform after that period must also be released as its collected for the duration of the tool's persistence on M-Lab.
- **The tests must be actively initiated by the client.** M-Lab employs active measurement tests only, in order to enable public release of the raw data without the need to anonymize or obfuscate personally identifying information. Tests running on M-Lab must not utilize M-Lab's infrastructure to conduct passive monitoring or conduct any measurement not purposefully initiated by the user. Accordingly, the test should not transmit any data that does not directly pertain to operation of the test, and all transmitted data should be synthetic.
- **Experiments must not collect personally identifying information.** M-Lab is committed to preserving user privacy, and tools running on the platform must not

collect personal information. Accordingly, the only identifying information that can be collected on M-Lab are IP addresses. While a given experiment may wish to collect personal information, M-Lab infrastructure cannot be used for collection, or storage of this information; non-M-Lab infrastructure must be employed for this purpose. In addition, any tool running on M-Lab that intends to begin collecting personal information outside of M-Lab must notify the Steering Committee, and be explicitly opt-in for users. For a successful example of this kind of collection, the U.S. Federal Communications Commission has separately collected address information from those participating in a bandwidth test that uses M-Lab's infrastructure.

An experiment which fits these criteria may apply by submitting the following information to the Steering Committee:

1. A description of the organization building and sponsoring the tool, if applicable, and an introduction to the team that will be maintaining the experimental tools.
2. A description of their research goals. What do they hope to study by conducting this experiment? Over what period of time do they intend that this experiment run?
3. A description of their experimental methodology. This should include an in-depth description of the data the experiment gathers and how it is gathered, as well what information is transmitted in each direction (client to server, server to client). It should make clear how their experiment intends to fit within the principles described above and should additionally describe how their experiment differs from existing M-Lab tools. In order to most efficiently use the resources of the platform, tools should be novel both in the data they collect and the experience they offer to the client. If the tool proposed already exists and is open-source then links to access the source code should be included. If research papers or articles already exist describing their experimental tools links to that documentation should be included.
4. A statement that those submitting the tools have read and understood the [M-Lab Policies and Procedures](#) document, and will operate the experiment in accordance with their responsibilities as laid out in that document.
5. An example of the data files generated by the experimental tool.
6. A timeline for deploying the tool, including what milestones exist and what effort is being put into meeting those milestones.

Once that information has been furnished to the Steering Committee, they will deliberate on the question at their next official meeting or via email, and decide whether or not to support the new experiment. The steering committee will consider applications based on their adherence to the principals outlined in this document, as well as the tools' novelty, and the availability of resources on the M-Lab platform. The reasoning behind the Steering Committee's decision will be made transparent to the those submitting the tool.