Network Working Group
Internet-Draft
Intended status: Informational
Expires: June 13, 2019

I. Learmonth
Tor Project

December 10, 2018

# Guidelines for Performing Safe Measurement on the Internet
## draft-learmonth-pearg-safe-internet-measurement-00

Abstract

Researchers from industry and academia will often use Internet measurements as a part of their work. While these measurements can give insight into the functioning and usage of the Internet, they can come at the cost of user privacy. This document describes guidelines for ensuring that such measurements can be carried out safely.

Note

Comments are solicited and should be addressed to the research group's mailing list at pearg@irtf.org and/or the author(s).

The sources for this draft are at:

https://github.com/irl/draft-safe-internet-measurement

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 13, 2019.

# 1.

When performing research using the Internet, as opposed to an isolated testbed or simulation platform, means that you research co-exists in a space with other users. This document outlines guidelines for academic and industry researchers that might use the Internet as part of scientific experiementation.

Following the guidelines contained within this document is not a substitute for any institutional ethics review process you may have, although these guidelines could help to inform that process. Similarly, these guidelines are not legal advice and local laws should be considered before starting any experiment that could have adverse impacts on user privacy.

Considerations are grouped into two categories: those that primarily apply to active measurements and those that primarily apply to passive measurements. Some of these considerations may be applicable to both depending on the experiment design.

# 2.

Active measurements generate traffic. Performance measurements such as TCP throughput testing [RFC6349] or functional measurements such as the feature-dependent connectivity failure tests performed by [PATHspider] both fall into this category.

## 2.1.

Wherever possible, use a testbed. An isolated network means that there are no other users sharing the infrastructure you are using for your experiments.

When measuring performance, competing traffic can have negative effects on the performance of your test traffic and so the testbed approach can also produce more accurate and repeatable results than experiments using the public Internet.

WAN link conditions can be emulated through artificial delays and/or packet loss using a tool like [netem]. Competing traffic can also be emulated using traffic generators.

2.2.

When performing measurements be sure to only capture traffic that you have generated. Traffic may be identified by IP ranges or by some token that is unlikely to be used by other users.

Again, this can help to improve the accuracy and repeatability of your experiment. [RFC2544], for performance benchmarking, requires that any frames received that were not part of the test traffic are discarded and not counted in the results.

2.3.

If your experiment is designed to trigger a response from infrastructure that is not your own, consider what the negative consequences of that may be. At the very least your experiment will consume bandwidth that may have to be paid for.

In more extreme circumstances, you could cause traffic to be generated that causes legal trouble for the owner of that infrastructure. The Internet is a global network crossing many legal jurisdictions and so what may be legal for you is not necessarily legal for everyone.

If you are sending a lot of traffic quickly, or otherwise generally deviate from typical client behaviour, a network may identify this as an attack which means that you will not be collecting results that are representative of what a typical client would see.

3.

Performing passive measurements requires existing traffic. Passive measurements can help to inform new developments in Internet protocols but can also carry risk.

3.1.

If you are in a position to perform passive measurements of live network traffic, you are also in a position of responsibility. Users of a network will have certain expectations of privacy and those expectations may not align with the privacy guarantees offered by the technologies they are using. As a thought experiment, consider how users might respond if you asked for their informed consent for the measurements you'd like to perform.

3.2.

When deciding on the data to collect, assume that any data collected might become public. There are many ways that this could happen, through operation security mistakes or compulsion by a judicial system.

3.3.

For all data collected, consider whether or not it is really needed.

3.4.

When collecting data, consider if the granularity can be limited by using bins or adding noise. XXX: Differential privacy.

3.5.

Do this at the source, definitely do it before you write to disk.

[Tor.2017-04-001] presents a case-study on the in-memory statistics in the software used by the Tor network, as an example.

4.

The benefits should outweigh the risks. Consider auxiliary data (e.g. third-party data sets) when assessing the risks.

5.

Take reasonable security precautions, e.g. about who has access to your data sets or experimental systems.

6.

This document has no actions for IANA.

7.

Many of these considerations are based on those from the [TorSafetyBoard] adapted and generalised to be applied to Internet research.

## 8. References

[netem] Stephen, H., "Network emulation with NetEm", April 2005.

[PATHspider]
Learmonth, I., Trammell, B., Kuehlewind, M., and G. Fairhurst, "PATHspider: A tool for active measurement of path transparency", DOI 10.1145/2959424.2959441, July 2016, <https://dl.acm.org/citation.cfm?doid=2959424.2959441>.

[RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, DOI 10.17487/RFC2544, March 1999, <https://www.rfc-editor.org/info/rfc2544>.

[RFC6349] Constantine, B., Forget, G., Geib, R., and R. Schrage, "Framework for TCP Throughput Testing", RFC 6349, DOI 10.17487/RFC6349, August 2011, <https://www.rfc-editor.org/info/rfc6349>.

[Tor.2017-04-001]
Herm, K., "Privacy analysis of Tor's in-memory statistics", Tor Tech Report 2017-04-001, <https://research.torproject.org/techreports/privacy-in-memory-2017-04-28.pdf>.

[TorSafetyBoard]
Tor Project, "Tor Research Safety Board", <https://research.torproject.org/safetyboard.html>.

## Author's Address

Iain R. Learmonth
Tor Project

Email: irl@torproject.org