Test Traffic Measurement Service

Data Disclosure Policy

*Henk Uijterwaal**

RIPE NCC
February 19, 2004

Document: RIPE-300.ps

1 Introduction

This document describes the Data Disclosure Policy (also referred to as the acceptable use policy or AUP) for the the data collected by the RIPE NCC Test Traffic Measurement Service (TTM) [1]. It describes who can access the data from this service, what one can do with the data and conditions that must be fulfilled before data can be published outside the RIPE Test Traffic Working Group (TT-WG) and the organisations hosting the test boxes.

This policy is based on discussion in the RIPE TT-WG. It replaces the old policy as described in [2]. Existing users of the service will be asked to confirm that they agree with this change in policy. New test boxes will only be installed if the hosts agree with this version of the policy.

2 General principles

This policy is based on 2 principles:

- Avoid abuse while giving as much freedom as possible to use the data.
- A simple procedure that describes the basic idea and is easy to follow rather than a heavy document that attempts to describe every possible case.

Collecting data with the test-boxes means collecting data about organisations and the performance of their networks. This is a delicate matter as

*Email: henk@ripe.net

nobody wants to see an analysis that puts the performance of his networks in a bad light, particularly if the scientific merits of the analysis cannot be proven. On the other hand, the results of TTM can be a valuable tool for day-to-day operations, long-term planning and scientific research. One does not want to be too restrictive about what can be done with the data.

To implement these principles, 3 basic rules are set:

- The RIPE NCC controls the distribution of the data.
- All data analysis should be peer reviewed before publication.
- All reports should include a pointer to the official description of the data as well as appropriate credit and copyright statements.

These basic rules are discussed in the next sections for two cases: section 3 discusses the situation for the owners of a test-box and their customers, section 4 for all others. The latter category generally consists of researchers doing scientific analysis of the data in order to better understand the behaviour of the Internet.

3 Participating Organisations

This section deals with organisations that own a Test Box and participate in the service, and their customers. Customers are loosely defined as any organization that uses the network of the test box host in order to connect to the rest of the Internet. There may or may not be a formal business relation. Sites owning a Test Box are called "test box hosts" in the remainder of this document,

3.1 Access to the data

There are several ways in which a test box host can access the data:

1. *Through an interface on the test box.* This shows the data as it is being collected, with a delay of only a few minutes. Access to this interface is restricted to IP addresses specified by the test box host and allocated to the host or its customers. The host will make the customer aware of this policy.

2. *Through the RIPE NCC website.* Processed data (plots, summaries) will be put on the RIPE NCC web site without password restrictions (but see section 4 below). Plots can be copied to another, internal web site if necessary, provided that proper credits are given and the test box host makes the users of that site aware of this policy.

3. *Through the RIPE NCC ftp-site.* Unprocessed data will be put on an ftp-server on request. The URL will be made available to the test box host and/or its customers.

Data can be downloaded and the host can share the data with its customers. The host is responsible for making the customer aware of this policy. Redistribution of the data to any other party in any form is not allowed.

3.2 Analysis and publication of the data

The data can be used freely for any analysis. The host can freely discuss the results with its customers or present the results to the RIPE NCC TT-WG.

Before an analysis is presented to the outside world, the analysis will have to be verified. This means that the organisation (including the RIPE NCC) that did the analysis will have to provide a write-up of the analysis that includes enough detail for anybody to independently verify its conclusions. This write-up will be circulated amongst all hosts for a peer review. This review period will last 2 weeks and a mailing list for comments will be made available.

If there are objections to an analysis, they will be discussed with the authors and the authors will be invited to submit a new version. When a new version is published, another 2 week review period will start. If a host still disagrees, they can ask that data related to their site is removed from the analysis. However, no single host can veto the publication of an analysis by another host or the RIPE NCC.

When data is published, it should include as few references to names of other sites, IP-addresses of test-boxes and routers etc. while keeping the document readable (for example, by calling the hosting sites ISP-A, ISP-B, etc.). Note that it will never be possible to make the data completely anonymous.

The RIPE NCC will take care of the administrative aspects of setting up a peer review, such as the distribution of draft papers and providing a mailing list for comments. It should be noted that the review process takes time. Anybody planning to present data at a conference should keep this in mind.

Once a document is published, the authors of the document should provide the RIPE NCC with a copy of the final paper. The RIPE NCC will maintain a list of published papers on its website.

3.3 Credits.

In any publication, appropriate credit to the source of the data should be given. The "official" reference for the TTM service is:

> F.Georgatos *et al*,
> "Providing Active Measurements as a Regular Service for ISP's".
> In: Proceedings of the Passive and Active Measurements Workshop
> PAM2001, Amsterdam, April 2001.
> `http://www.ripe.net/ttm`

For the data, the reference is:

> Data from the RIPE NCC TTM service, ©RIPE NCC, 2003, all rights reserved. Data may not be redistributed without prior ap-

proval of the RIPE NCC. Please see `http://www.ripe.net/ripe/docs/ripe-300.html` for details.

4 Others

This section deals with all others requesting access to the data.

4.1 Access to the data

There are two ways in which one can access the data.

1. *Processed data on the RIPE NCC website.* Plots will be visible on the RIPE NCC website without restriction. The first time the website is accessed, this policy will be explicitly shown on the screen, asking the user to confirm that he has read it before proceeding. All pages will include a reference to this document.

2. *Raw data on the RIPE NCC ftp-site.* Raw data will be made available on the ftp site. The URL will only be made available *after* the organisation has signed a copy of the policy (see appendix A). Redistribution of the data to any other party in any form is not allowed. Data should be removed when the studies have finished.

4.2 Analysis and publication of the data

The data can be used freely for any analysis. One is free to discuss the analysis inside the organization that did the analysis or the RIPE TT-WG.

Before an analysis is presented to the outside world, the analysis will have to be verified. This means that the organisation that did the analysis, will have to provide a write-up of the analysis that includes enough detail for anybody to independently verify its conclusions. This write-up will be circulated amongst the hosts for a peer review according to the procedure described in section 3.2.

When data is published, it should include as few references to names of other ISPs, IP-addresses of test-boxes and routers etc. while keeping the document readable (for example, by calling the hosting sites ISP-A, ISP-B, etc.). Note that it will never be possible to make the data completely anonymous.

The RIPE NCC will take care of the administrative aspects of setting up a peer review (see section 3.2).

Once a document is published, the authors of the document should provide the RIPE NCC with a copy of the final paper. The RIPE NCC will maintain a list of published papers on its website.

4.3 Credits.

In any publication, appropriate credit should be given (see section 3.3).

5 Changes in the policy

Changes in the data-disclosure policy can be suggested by the participating ISPs or the RIPE NCC. All changes will be discussed with test box hosts participating at that point. If all parties agree on the change, then this document will be revised. It is our goal that the new policy will be acceptable to everybody participating in the project.

6 Concluding remarks.

The RIPE NCC will take care of the administrative details related to this policy. It will make decisions on small issues not covered in this policy. Major issues with this policy will be brought forward on the tt-wg@ripe.net mailing list.

References

[1] H. Uijterwaal, O. Kolkman, "Internet Delay Measurements using Test Traffic, Design Note", RIPE-158.

[2] H. Uijterwaal, "Internet Delay Measurements using Test Traffic, Data disclosure policy", RIPE-180.

A Acceptance of the Data Disclosure Policy

Name(s): .....

Email: .....

Legally representing

Name of Organisation: .....

Address: .....

.....

.....

.....

has/have asked the RIPE NCC for access to data collected by the Test Traffic Measurement Service (TTM).

I/We have read and understood the data disclosure policy for this service ("AUP") as described in document RIPE-300.ps (`http://www.ripe.net/ripe/docs/ripe-300.html`). I/We agree to abide by these conditions.

I/We will ensure that all other members of our organisation who will get access to the data, will be made aware of this policy.

Name: .....

Date: .....

Signature(s): .....

Please print a copy of this form on your organisation's stationary, sign and send a copy of this form to the RIPE NCC by post or fax.