

## **The Neubot Project: A Collaborative Approach To Measuring Internet Neutrality**

Juan Carlos De Martin and Andrea Glorioso  
NEXA Center for Internet & Society – DAUIN – Politecnico di Torino, Italy  
<http://nexa.polito.it>  
[demartin@polito.it](mailto:demartin@polito.it), [andrea.glorioso@polito.it](mailto:andrea.glorioso@polito.it)

### **Abstract**

*The Internet was designed to be neutral with respect to kinds of applications, senders and destinations. Such design choice made very fast packet switching possible, while preserving, at the same time, strong openness towards unforeseen uses of the Internet Protocol. The result has been an extraordinary outburst of innovation, as well as a level-playing field for citizens, associations and companies worldwide. With the advent of “deep packet inspection” technology, however, fine-grained discrimination of Internet flows is now possible, be that for economical or other reasons. Collecting quantitative data on the behavior of telecommunications providers with respect to traffic discrimination thus becomes crucial, particularly at a time when policy changes are widely discussed. The “Network Neutrality Bot” (Neubot) project is based on a lightweight, open source computer program, the Neubot, that, downloaded and installed by Internet users, performs distributed measurements of the traffic characteristics of segments of the global Internet. The collected data will allow constant monitoring of the actual state of the Internet, enabling both a deeper understanding of such crucial infrastructure and a more reliable basis for discussing network neutrality policies.*

### **1. High-level description**

The Network Neutrality Bot (Neubot) is an open source application that measures the characteristics – including latency, bandwidth, jitter, packet loss rate, filtering of specific ports/applications, both instantaneously and over time – of the transmissions taking place across the section of the Internet in which the Neubot is placed.

Although it would be possible in theory to perform such measurements via probes installed in Internet Exchange Points, in practice this option is problematic

from several points of view. Therefore, the Neubot is based on a distributed model.

The Neubot, in fact, consists of a small application that end users would download and voluntarily install on their computers. Such application would run in the background, routinely performing a set of transmission tests between the Neubot computer and one (or more) servers (client-server mode), and between the Neubot computer and other Neubots (peer-to-peer mode).

The Neubot would then periodically report the test results to a central database server, therefore building a comprehensive set of data regarding locations of the Internet under the control of different Internet Service Providers / telecommunication operators. Such data set would then be analyzed to provide reliable, data-based snapshots of actual behaviors of at least portions of the Internet.

The Neubot project faces several technical challenges. Arguably the main ones are how to obtain reliable measures for the protocols, some of them proprietary, which has been chosen as relevant for this project, and how to ensure data integrity and user anonymity, particularly in the peer-to-peer mode of operation.

### **2. Rationale**

The "network neutrality" topic has lately been one of the predominant elements in the worldwide debate on Internet policies. The basic question is whether network operators should be allowed to differentiate the Internet traffic that goes through their infrastructures or, on the contrary, whether network neutrality should be explicitly safeguarded by the law, therefore enshrining what has been a characteristic of the Internet since its birth.

Proper tools and methodologies to assess operators' policies and, even more relevantly, actual practices in this area are key elements of an informed debate, both *ex ante*, i.e., assessing the current scenario, and *ex post*, i.e., verifying operators' behavior in light of regulatory decisions in this area.

In this regard, legal analysis is, of course, an important element. Among other aspects, legal analysis should consider mandatory law. In this regard, it is interesting to read, among other things, the recent proposal of the European Commission on the review of 2002 “Electronic Communications” regulatory package, mandating national regulatory authorities to apply to principle that “end-users should be able to access and distribute any lawful content and use any lawful applications and/or services of their choice” [1]. Moreover, contractual practices of Internet operators *vis-à-vis* their relationships with Internet users also deserve legal scrutiny.

However, one might arguably question the relevance of the legal framework in light of the power of Internet operators to directly control the basic infrastructures (including edge routers) and “control points” [2] [4] of relevance to end-users' Internet traffic, as well as the high information asymmetry [5] [6] that characterizes the “market” under consideration.

The combination of this power of control and of the difficulty of obtaining proper information makes it essential that tools for assessing the behavior of Internet operators are readily available.

The Neubot project tries to provide all interested stakeholders, including end-users and policy makers, with a reliable – meaning both analytically correct and resistant against attempts to unduly influence data collection – way to quantitatively assess the *ex ante* and *ex post* scenarios described above.

### 3. Perspectives on Network Neutrality

The topic of network neutrality can be tackled using a variety of disciplinary perspectives.

Taking as an anecdotal example two social disciplines whose practitioners have been particularly vocal in the network neutrality debate, a jurist might probably analyze network neutrality scenarios using the lenses of – among others – *competition law*, e.g., questioning whether network operators have any dominant power in the relevant market or whether they are possibly going to significantly influence a complementary market and/or constitute an “essential facility”, access to which is absolutely mandatory in order for other players to operate competitively on the relevant market; *consumer protection law*, e.g., understanding whether certain practices by network operators can lead to harm for consumers and end-users of digital goods and services; and *human rights law*, e.g., framing the whole discourse around network neutrality in terms of provision of, and access to, free information, as prescribed by many

international instruments for the protection and advancement of human and fundamental rights.

Conversely, an economist might approach network neutrality by focusing, depending on his/her particular expertise and preferences, on the analysis of: *innovation processes*, e.g., putting to test, whether empirically or through model building, the debate between centralized vs distributed innovation scenarios and checking whether the choice of a particular network neutrality policy influences such innovation dynamics; *macroeconomic factors* promoting the growth of more efficient network infrastructures, e.g., verifying whether certain claims by incumbent network operators, that in order for them to invest in the construction and deployment of so-called Next Generation Networks they must be given total or almost total control over the traffic passing through those networks; and *microeconomic price and service discrimination*, e.g., understanding whether and at what conditions differentiations in service levels might prove to be dynamically efficient for all the players on the relevant market.

Of particular interest for the Neubot project is the practice by Internet operators controlling a particular subset of the Internet to willfully penalize certain types of traffic, depending on the protocol used (e.g., all traffic using the Bittorrent protocol) or on the particular resource being requested (e.g., slowing down access to a particular website). The Neubot project is not directly interested in non-discriminatory differentiation of traffic when this is functionally necessary in order to guarantee the effectiveness of the related communication stream, as could be the case when data packets belonging to a real-time stream are given priority to bulk data transfer.

Whatever perspective is taken to analyze and discuss network neutrality and related topics, the non-exhaustive list above and the central importance that the Internet has acquired in all facets of social life suggest that any policy decision in this field – including a *laissez-faire*, “no policy necessary” approach – will have long-lasting societal effects.

An informed debate, based on objective data, is therefore highly advisable, hence the Neubot project, which was launched in June 2007 by the NEXA Center for Internet and Society at the Politecnico di Torino, Italy.

### 4. Architecture

The architecture of the Neubot consists of an open-source client application that volunteer end-users install on their computers and on a set of Neubot servers. The client application runs in the background and

automatically performs a set of measures, periodically sending results back to a central server (or a set of distributed servers).

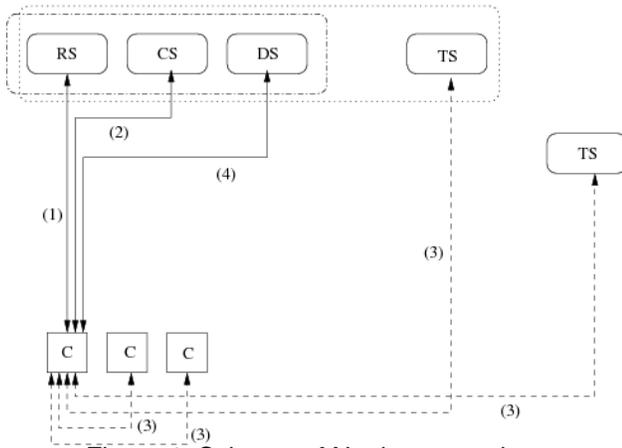


Figure 1: Scheme of Neubot operations

The diagram in Figure 1 shows the basic operations of the Neubot architecture. Each Neubot client (C) first registers (1) with a Registration Server (RS), which identifies the client. The RS also informs the Neubot client about the other available clients.

The client then connects (2) to a Configuration Server (CS) to collect instructions on which tests should be performed, i.e., which target URL's, which protocols, at what time, for how long. This allows the tests to be flexible: different “targets” could be chosen depending on the circumstances.

The Neubot client then performs the relevant tests connecting (3) to the so-called Test Servers (TS) if the Neubot is running in client-server mode, or to other Neubot clients if in peer-to-peer modality. Once the tests are completed, each Neubot client reports the results back to a Database Server (DS).

The Neubot client measures several indicators, including throughput, delay, jitter, as well as the performance of specific widely-used protocols, including the IETF Real-Time Transport Protocol (used by most Internet-based multimedia applications), the peer-to-peer BitTorrent protocol and the proprietary peer-to-peer Skype Voice over IP protocol.

The Test Servers might be under the control of the project coordinator or of a third party. The latter case will be used to test common protocols such as HTTP or SMTP, while the former will be needed to assess the performance of other, more specialized protocols.

Also, the separation between Registration Server, Configuration Server, Database Server and the Test Servers is purely logical – these services might run on the same host for performance or maintainability reasons.

Considering the uncontrolled nature of the environment, some precautions must be taken. First, the Neubot client should have minimal impact: memory footprint, computational load, and network usage should therefore be kept to a minimum, particularly when the hosting computer is in use. In addition, the Neubot client should ensure a satisfying degree of anonymity.

Another important security point is authentication; all the parts, with the exception of the Test Servers, must be authenticated to avoid server masquerading, implying that the client would get or provide information to a fake server. Authentication of the client is also necessary to shield the Database Server from receiving false measurements from untrusted sources. To implement the two-way authentication a Secure Socket Layer connection can be established, using digital certificates for both the Neubot clients and the Registration, Configuration and Database servers in order to authenticate both sides and establishing a session key to encrypt the data passing through this stateful connection. Data confidentiality is guaranteed because all the data is encrypted with standard algorithms like AES, DES or 3DES. With respect to data integrity, it is reasonable to use digital signatures on the data sent between the parties, making it possible to verify that the data has not been modified during the transit over the network.

From a high-level perspective, the architecture of the Neubot project might look similar to that of other projects, such as the OpenNet Initiative [3], whose goal is to test – using a distributed model – Internet filtering practices of countries around the world. However, the Neubot project has a wider scope, insofar as it aims to test not only the reachability, but also the behavior of particular subsets of the Internet; besides, unlike the OpenNet Initiative which focuses on HTTP traffic, the Neubot project aims to analyze a larger number of protocols as discussed above.

## 5. Conclusion

The Neubot project strives to provide all interested stakeholders with a reliable way to assess Internet behavior with respect to potential discriminations of certain classes of users, destinations or applications. The Neubot architecture uses both a client-server and a peer-to-peer approach to measure and collect data for an informed debate on the status and possible effects of

policy making regarding “network neutrality”. The value of Neubot lies both in its novel distributed architecture and in its capability to provide such data using a bottom-up methodology that, by countering the power of control by Internet operators over infrastructures and the resulting information asymmetry, is arguably conducive to a more democratic decisional process about a crucial infrastructure for our societies.

## 6. Acknowledgments

The authors wish to thank Mr Gianluigi Pignatari who, as part of his thesis work at the Politecnico di Torino, is greatly contributing to develop the first release of Neubot.

## 7. References

- [1] Commission of the European Communities, *COM(2007) 697 final – Proposal for a Directive of the European Parliament and of the Council Amending Directive 2002/21/EC on a Common Regulatory Framework for Electronic Communications Networks and Services, 2002/19/EC on Access to, and Interconnection of, Electronic Communications Networks and Services, and 2002/20/EC on the Authorisation of Electronic Communications Networks and Services*, Bruxelles, 13 November 2007.
- [2] Dame A., J.H. Guettler, K. Leeson, M. Schultz, and T.B. Jensen, *Regulatory Implications of the Introduction of Next Generation Networks and Other New Developments in Electronic Communications (Study Report for the European Commission)*, 2003.
- [3] Deibert, R., J. Palfrey, R. Rohozinski, and J. Zittrain, *Access Denied: The Practice and Policy of Global Internet Filtering*, MIT Press, Cambridge (MA), 2008
- [4] Lessig, L., *Code and Other Laws of Cyberspace*, Basic Books, 1999.
- [5] Pindyck, R.S., and D.L. Rubinfeld, *Microeconomics, 6<sup>th</sup> edition*, Prentice Hall, 2005.
- [6] Shapiro, C., and H.R. Varian, *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press, Cambridge (MA), 1998.