Assessing Application-Level Network Neutrality using Digital Twins

Péter Orosz, Tamás Marosits, Tamás Skopkó
Dept. of Telecommunications and Artificial Intelligence
Budapest University of Technology and Economics
Budapest, Hungary
orosz@tmit.bme.hu, marosits@tmit.bme.hu, skopko@tmit.bme.hu

***Abstract*—Our paper addresses the challenges of assessing network neutrality in the context of modern cloud-based services. It highlights the growing complexity for Internet service providers (ISPs) in managing network resources, which may conflict with the traditional best-effort Internet model, potentially compromising network neutrality principles. The authors propose a novel proof-of-concept system that utilizes digital twins to measure and analyze application-specific network neutrality in real-time on the public Internet. This system is designed to be horizontally scalable, adaptable to changing networks, and capable of producing low false rates in neutrality assessments. We detail the architecture of the measurement system, which includes a real client, a measurement server, and the client's digital twin, all of which work together to monitor and compare service quality. The proposed system aims to offer a reliable, real-time method for assessing network neutrality, specifically in terms of how ISPs manage traffic for different applications.**

I. Introduction

Cloud-based services have dominated global Internet communication over the last decade, driven by advanced datacenter technologies. For Internet providers, managing network resources to optimize user experience for these popular services is challenging, as it may conflict with the traditional best-effort Internet model. Prioritizing services based on popularity could hinder new technologies and services. A key question is how resource management impacts users' access to quality services, which ties into network neutrality principles supported by EU Regulation 2015/2120.

While this regulation provides a legal framework, a comprehensive technological solution to verify Internet access neutrality is lacking. This gap exists because the scientific foundations for such measurements are not yet fully developed. Analyzing network neutrality attributes, particularly for cloud services, is a new research area without established standards.

Network neutrality demands equal treatment of all user traffic, but recent years have seen cases of traffic differentiation that violate this principle. Technologies like AI-based traffic detection allow operators to identify and prioritize certain services, potentially ensuring better quality for some while relegating others to best-effort transmission.

To assess service quality, objective models are categorized as full-reference or no-reference. Full-reference models, though more accurate, require a reference source and have privacy concerns. This paper proposes a real-time full-reference model for identifying service quality restrictions imposed by Internet providers.

Network neutrality assessment methods can be categorized based on traffic patterns: real, replayed, or generated. The main questions are how these methods reflect real user scenarios and how to minimize false detection. Given the challenges of repeatability and objectivity in public Internet measurements, a comprehensive neutrality analysis must be reliable across live networks.

This paper aims to demonstrate a proof-of-concept network neutrality assessment system to support and extend the authors' related paper [1] establishing a new assessment paradigm: real-time, application- and platform-specific analysis of network neutrality on the public Internet. Leveraging public cloud services, this approach adapts to changing networks and provides low false rates in neutrality assessments.

The paper is structured as follows: Section II introduces the methodology and the architecture of the novel analytics system, Section III overviews the demo setup and the presented scenatios. Finally, Section IV concludes the paper.

II. Architecture and Assessment Method

Our proposal introduces three key innovations: i) it is horizontally scalable to assess any cloud-based service, ii) it measures real traffic from a public service provider, and iii) it creates a measurement reference in real time using the client's digital twin. While demonstrated with cloud-based services, it can also be applied to web-based applications on the client side. The primary aim is to conduct reliable application-specific network neutrality measurements on the public Internet.

A. Measurement architecture

The architecture includes three main components: the real client, the measurement server running the client's digital twin, and the content provider service, as shown on Fig. 1. Neutrality measurements are performed in parallel, with the real client and its emulated digital twin (which mirrors the real client's Internet access parameters) simultaneously measuring the same service using the same application.

The measurement server is designed as a flexible platform that runs the software components necessary for the measurement process. Its functions include assessing the Internet

978-3-903176-66-9

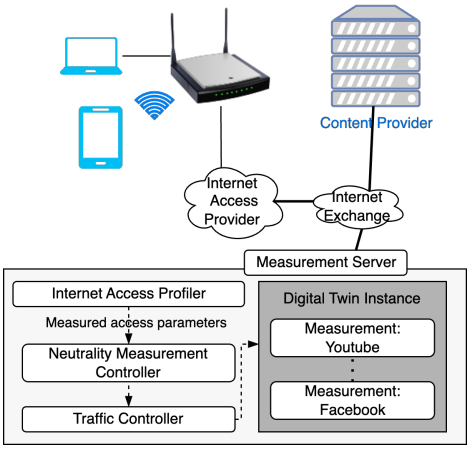

Fig. 1. Measurement architecture

access parameters of real clients, creating network containers to emulate these parameters, executing service measurements for emulated clients, and controlling the real client during testing. The server requires a high-capacity uplink (e.g., at least 10 Gbps) to accommodate multiple clients with different access profiles. After the initial bandwidth measurement, the real clients only exchange control messages with the server.

B. Real-Time Reference: Dual Measurement with Digital Twin

A key feature of the architecture is the emulated digital twin of the client, which mimics the network access characteristics of the real client, such as download/upload speeds, latency, and optionally packet loss. The measurement server should be located near an Internet Exchange Point, independent of the user's ISP, to ensure unbiased results. A preliminary QoS measurement phase identifies the real client's network access properties, which are then replicated by the emulated client for accurate comparison.

C. Measurement phases

The Neutrality Measurement Controller manages the measurement process by coordinating the real and emulated clients. The measurement session follows these steps:

1) Conduct a preliminary QoS measurement on the real client to determine key network access parameters.
2) The digital twin replicates the real client's access profile using software-based emulation within a network container.
3) Both clients request the same content from the provider to initiate the measurement.
4) During content streaming, both clients measure key performance indicators (KPIs) specific to the service.
5) Each client independently evaluates service quality based on the measured KPIs.
6) The measurement controller processes these evaluations, determining overall service quality and presenting it on a predefined quality scale (e.g., ITU P.800 MOS scale).

The preliminary QoS measurement uses a validated method from the SCL Broadband Measurement System, recognized by the Hungarian National Media and Infocommunications Authority. Additionally, the real-time reference helps identify bottlenecks on the content provider side, reducing false detection.

III. Demonstration

A. Legal and measurement background in the EU

Since the initiative for the design of the previously mentioned network neutrality assessment tool is based on the BEREC documents which explain the TSM regulation, it is worth briefly reviewing the items to be examined, which are listed in section 3.1.2 of BoR (17) 179 and sections 4.1 and 4.2 of BEREC's RAM documents[1] [4]. In general, it can be concluded that the sources of problems related to average user behavior are currently adequately described. At the same time, it is striking that BEREC tries to find a universal solution for each type of application, even though users still use unique, nameable applications whose traffic patterns and quality requirements for the network can be very different. In addition, the frequency of changes in users' content consumption and Internet usage habits justifies that the procedures examining generalized application types only provide an intermediate layer of a complex network neutrality measurement system.

Investigations of net neutrality problem sources can be divided into two groups according to sections 4.1 and 4.2 of RAM. The ones in the first group are called connection measurements, while the procedures affecting individual application quality are in the second group. Note that in the current issue of RAM, Chapter 4 is shorter than in the previous document. This is manifested in the fact that while in section 4.1, the description of the connection tests is roughly the same length, and the text also matches verbatim in most places, until in section 4.2, the possibilities of examining the manipulations affecting the traffic of some types of applications have been completely abandoned. These two groups can be divided further if, for example, the test output types or the test result classification are considered criteria. By test output, we mean the type of result generated in the measurement system after the test is performed. The qualification of the test result means what kind of information is given to the user who initiated the measurement.

In this regard, tests aimed at checking the blocking of a service or an application are the simplest since here, the fact of blocking or its absence can be clearly established, and the result of the measurement can be displayed to the user who initiates the measurement practically without any changes, possibly supplemented with some kind of interpretation. The same applies to the measurements carried out in the application cases, which refer to the availability of a specific service or, based on the specification issued by BEREC, a generalized type of service. In practical cases, these

[1]BEREC used to have documents with the same title and different number [2] [3]. The last two issues Net Neutrality Regulatory Assessment Methodology - often referred to as RAM - have the numbers BoR (17) 178 and BoR (22) 72, respectively.

The name of the test	The output of the test, the result type	Classification and nature of the test result
Technology tests		
Blocked IP discovery	IP address blocking detection Blocked/unblocked	Pass/Fail
Port scan	TCP/UDP port blocking detection Blocked/unblocked	Pass/Fail
Detection of manipulation of DNA resolution	Manipulated/not manipulated	Pass/Fail
HTTP proxy detection	Manipulated/not manipulated	Pass/Fail
Detecting website slowdowns	Numerical, measurable characteristic(s)	Pass/Fail
Application tests		
Video streaming	Detection of influence on TCP video stream (YouTube). Numerical, measurable characteristic(s) and values derived from them	Multi-valued scale
UDP Video signal stream influence detection	Numerical, measurable characteristic(s) and values derived from them	Multi-valued scale
Internet-based audio	UDP audio stream influence detection. Numerical, measurable characteristic(s) and values derived from them	Multi-valued scale
Social media	Availability and usability of the social media site (Facebook). Blocked/unblocked + Numerical, measurable characteristic(s) and values derived from them	Multi-valued scale
Peer-2-peer	Detection of blocking and traffic restriction of Peer to Peer communications (Bittorrent). Blocked/unblocked + Numerical, measurable characteristic(s) and values derived from them	Multi-valued scale

TABLE I
TESTS OF THE SMARTCOM LAB'S NET NEUTRALITY MEASURMENT SYSTEM

can be traced back to checking whether a particular port or IP address is blocked. The other type of result is when there is a multi-grade transition between the complete failure of the service and its availability in an intact form; therefore, the measurement results are projected on a multi-valued scale. These are objective metrics with which we try to approximate the subjective quality of experience user perceives.

As mentioned, BEREC specified, tendered and then ordered a measurement system that can be used for both quality of service and net neutrality tests [5]. The original goal was to provide a pan-European measurement system for national regulatory organizations that do not have their own measurement system, which would not only serve as a reference for the others, but would also enable the inclusion of all EU measurement results by fixing the database interface. The development has now been completed, but it has not yet been widely distributed[2].

[2]Some informations can be found on the web page of the project (https://net-neutrality.tools/), but the mentioned address of the measurement system (https://nntool.eu) is unreachable.

B. Demo scenario

At the same time, the SmartCom Lab's Network Neutrality Measurement System is already tested by the National Media and Infocommunications Authority of Hungary. During the demonstration, our assessment system will perform a complete network neutrality evaluation of the client's internet access. As listed in Table I, the measurement session incorporates 5 technical tests and 5 application-specific assessments. The measurement server component is located in a network directly connected to BIX (Budapest Internet Exchange), while multiple clients are deployed in different Internet access environments. One of the clients will be restricted by degraded QoS performance of the access network by emulating an administrative restriction of the access provider. The overall results will be displayed in the clients' web browser.

IV. CONCLUSIONS

The paper introduces a novel, real-time network neutrality assessment system that utilizes digital twins to accurately measure and compare the quality of service provided by Internet service providers. By replicating the real client's network conditions in a controlled environment with digital twins, the system can identify and evaluate potential violations of network neutrality with minimal false positives. This approach is horizontally scalable and adaptable to the dynamic nature of modern networks, particularly in the context of cloud-based services. The paper's findings suggest that this innovative methodology could significantly enhance the reliability of network neutrality assessments, highlighting the need for further development and research in this emerging field.

REFERENCES

[1] P. Orosz, T. Skopkó, and T. Marosits, "Application-Aware Analysis of Network Neutrality: A Scalable Real-Time Method," *Infocommunications journal*, vol. 15, no. 1, pp. 77–86, 2023.

[2] BEREC, "Net neutrality measurement tool specification, BoR (17) 179," Oct. 2017. [Online]. Available: https://www.berec.europa.eu/en/document-categories/berec/reports/net-neutrality-measurement-tool-specification

[3] ——, "Net Neutrality Regulatory Assessment Methodology, BoR (17) 178, p. 25," Oct. 2017. [Online]. Available: https://www.berec.europa.eu/en/document-categories/berec/regulatory-best-practices/methodologies/berec-net-neutrality-regulatory-assessment-methodology

[4] ——, "Net Neutrality Regulatory Assessment Methodology, BoR (22) 72, p. 38," Jun. 2022. [Online]. Available: https://www.berec.europa.eu/en/document-categories/berec/regulatory-best-practices/methodologies/berec-net-neutrality-regulatory-assessment-methodology-0

[5] ——, "Tender Specifications No BEREC/2018/01/OT – Net Neutrality Measurement Tool, BoR (18) 32," Mar. 2018. [Online]. Available: https://etendering.ted.europa.eu/document/document-file-download.html?docFileId=41876