Understanding the Impact of Network Infrastructure Changes Using Large-Scale Measurement Platforms

Vaibhav Bajpai, Jürgen Schönwälder

Understanding the Impact of Network Infrastructure Changes using Large-Scale Measurement Platforms

Vaibhav Bajpai and Jürgen Schönwälder[*]

Computer Science, Jacobs University Bremen, Germany
{v.bajpai,j.schoenwaelder}@jacobs-university.de

Abstract. A number of large-scale measurement platforms have emerged in the last few years. These platforms have deployed thousands of probes within the access, backbone networks and at residential gateways. Their primary goal is to measure the performance of broadband access networks and to help regulators sketch better policy decisions. We want to expand the goal further by using large-scale measurement platforms to understand the impact of network infrastructure changes.

1 Research Statement

The curiosity to understand the performance of the Internet from the user's vantage point led to the development of techniques to remotely probe broadband access networks. Dischinger *et al.* in [3], for instance, inject packet trains and use the responses received from residential gateways to infer broadband link characteristics. This led to the development of a number of software-based solutions such as `netalyzr` [8], that require explicit interactions with the broadband customer. Recently, the requirement for accurate measurements, coupled with efforts initiated by regulators to define data-driven standards, has led to the deployment of a number of large-scale measurement platforms that perform measurements using dedicated hardware probes not only from within ISP networks but also directly from home gateways.

In a recent study, sponsored by the FCC, Sundaresan *et al.* [10] have used measurement data from a swarm of deployed SamKnows probes to investigate the throughput and latency of access network links across multiple ISPs in the United States. They have analyzed this data together with data from their own Bismark platform to investigate different traffic shaping policies enforced by ISPs and to understand the bufferbloat phenomenon. The empirical findings of this study have recently been repraised by Canadi *et al.* in [2] where they use crowdsourced data from `speedtest.net` to compare both results. The primary aim of all these activities is to measure the performance and reliability of broadband access networks and facilitate the regulators with research findings to help them

---

make policy decisions [9]. Using a large-scale measurement platform we want to take this further and study the impact of network infrastructure changes. We want to define metrics, implement measurement tests and data analysis tools that help us answer questions of the form:

- *How does the performance of IPv6 compare to that of IPv4 in the real world?*
- *Can we identify a Carrier-Grade NAT (CGN) from a home gateway?*
- *Can we identify multiple layers of NAT from a home gateway?*
- *How much do web services centralize on Content Delivery Network (CDN)s?*
- *To what extend does the network experience depend on regionalization?*

In the past, we have performed an evaluation of IPv6 transitioning technologies to identify how well current applications and protocols interoperate with them [1]. We are now participating in the Leone[1] project, whose primary goal is to define metrics and implement tests that can asses the end-user's Quality of Experience (QoE) from measurements running on SamKnows probes.

2 Proposed Approach

SamKnows specializes in the deployment of hardware-based probes that perform measurements to assess the performance of broadband access networks. The probes function by performing active measurements when the user is not aggressively using the network. RIPE Atlas is another independent measurement infrastructure deployed by the RIPE NCC. It consists of hardware probes distributed around the globe that perform RTT and traceroute measurements to preconfigured destinations alongside DNS queries to DNS root servers.

Measurement Lab (M-Lab) [5] is an open, distributed platform to deploy Internet measurement tools. The measurement results are stored on Google's infrastructure. The tools vary from measuring TCP throughput and available bandwidth to emulating clients to identify end-user traffic differentiation policies [4,6] to performing reverse traceroute lookups from arbitrary destinations [7].

It will only be possible to answer the aforementioned research questions with access to a large-scale measurement platform. As partners of the Leone consortium, we will leverage the infrastructure of our partners. We will define metrics targeted to our research questions and complement them by implementing measurement tests. The developed tests will be deployed in our partner's networks, but may also become part of SamKnows global infrastructure, which has several thousand deployed probes and will continue to grow during the project's lifetime. The collected data will be conglomerated from multiple Measurement Agent (MA)s and analyzed to uncover information needed to help us answer these questions. This requires to develop data analysis algorithms that can integrate data from different data sources such as address block allocations from Regional Internet Registry (RIR)s or prefix and path information from BGP route views. In this pursuit, we have started with a study to assess how the user experience is effected by the deployment of IPv6.
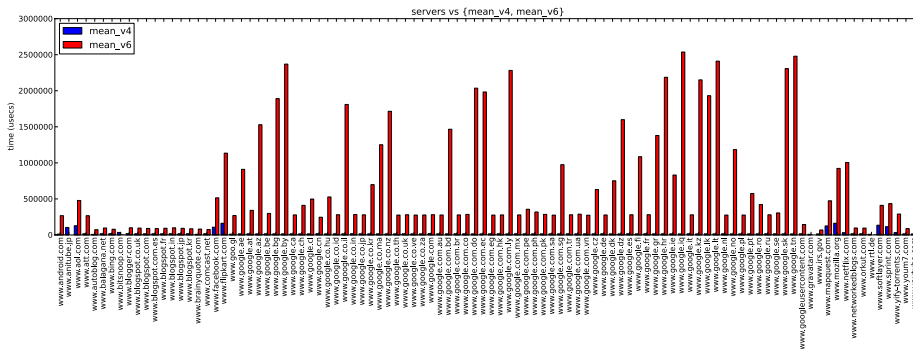
---

[1] `http://leone-project.eu`

Fig. 1. Mean time to establish TCP connections to a list of web services. The MA is a virtual machine hosted at `greatnet.de`. It has IPv4 connectivity via LamdaNet Communications [AS13237] and IPv6 connectivity via Teredo.

3 Preliminary Results

The function `getaddrinfo(...)` resolves a service name to a list of endpoints in an order that prioritizes an IPv6-upgrade path [11]. The order can dramatically reduce the application's responsiveness where IPv6 connectivity is broken, because the attempt to connect over an IPv4 endpoint will take place only when the IPv6 connection attempt has timed out, which can be in the order of seconds.

This degraded user experience can be subverted by implementing the happy eyeballs algorithm [12]. The algorithm recommends that a host, after resolving the service name, tries a TCP `connect(...)` to the first endpoint (usually IPv6). However, instead of waiting for a timeout, it only waits for 300ms, after which it must initiate another TCP `connect(...)` to an endpoint with a different address family and start a competition to pick the one that completes first. We have developed `happy`, a simple TCP happy eyeballs probing tool that uses TCP connection establishment time as a parameter to measure the algorithm's effectiveness. It uses non-blocking `connect(...)` calls to concurrently establish connections to all endpoints of a service. In order to develop data-analysis tools, we have prepared an internal test-bed of multiple MAs. The MAs have different flavors of IPv4 and IPv6 connectivity ranging from native IPv4, native IPv6, IPv6 tunnel broker endpoints, Teredo and tunnelled IPv4. We used the top 100 DNS names compiled by `he.net`[2] and ran `happy` on them.

A preliminary result comparing the mean time to establish a TCP connection to each of the services from one of the MA is shown in Fig. 1. The initial results show higher connection times over IPv6. Furthermore, on a Teredo MA, an application will never use IPv6 except when IPv4 connectivity is broken, because the Teredo IPv6 prefix has a low priority in the address selection algorithm [11]. It also appears that several services show very similar performance. These services

---

[2] `http://bgp.he.net/ipv6-progress-report.cgi`

resolve to a set of endpoints that belong to the same allocated address blocks. Digging through the `whois` information for each of the endpoints from their RIR demonstrates that major portions of the services map to address blocks owned by organizations such as Google and Akamai Technologies.

4 Conclusion

We have performed a preliminary study on how IPv6 deployment may affect the QoE of Internet users. Using a large-scale measurement platform we want to take this further, and define new metrics, measurement tests and data analysis tools that help us understand the impact of network infrastructure changes.

References

1. Bajpai, V., Melnikov, N., Sehgal, A., Schönwälder, J.: Flow-Based Identification of Failures Caused by IPv6 Transition Mechanisms. In: Proceedings of the 6th Conference on Autonomous Infrastructure, Management, and Security. AIMS'12
2. Canadi, I., Barford, P., Sommers, J.: Revisiting Broadband Performance. In: Proceedings of the ACM Internet Measurement Conference. pp. 273–286. IMC '12
3. Dischinger, M., Haeberlen, A., Gummadi, K.P., Saroiu, S.: Characterizing Residential Broadband Networks. In: Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement. pp. 43–56. IMC '07, ACM, New York, NY, USA (2007)
4. Dischinger, M., Marcon, M., Guha, S., Gummadi, K.P., Mahajan, R., Saroiu, S.: Glasnost: Enabling End Users to Detect Traffic Differentiation. In: Proceedings of Networked Systems Design and Implementation Conference. NSDI'10
5. Dovrolis, C., Gummadi, K., Kuzmanovic, A., Meinrath, S.D.: Measurement Lab: Overview and an Invitation to the Research Community. SIGCOMM Computer Communications Review 40(3), 53–56 (Jun 2010)
6. Kanuparthy, P., Dovrolis, C.: ShaperProbe: End-to-End Detection of ISP Traffic Shaping using Active Methods. In: Proceedings of the ACM Internet Measurement Conference. pp. 473–482. IMC '11, ACM, New York, NY, USA (2011)
7. Katz-Bassett, E., Madhyastha, H.V., Adhikari, V.K., Scott, C., Sherry, J., Van Wesep, P., Anderson, T., Krishnamurthy, A.: Reverse Traceroute. In: Proceedings of Networked Systems Design and Implementation Conference. NSDI'10
8. Kreibich, C., Weaver, N., Nechaev, B., Paxson, V.: Netalyzr: Illuminating the Edge Network. In: Proceedings of the ACM Internet Measurement Conference. IMC '10
9. Schulzrinne, H., Johnston, W., Miller, J.: Large-Scale Measurement of Broadband Performance: Use Cases, Architecture and Protocol Requirements, `http://tools.ietf.org/html/draft-schulzrinne-lmap-requirements-00`
10. Sundaresan, S., de Donato, W., Feamster, N., Teixeira, R., Crawford, S., Pescapè, A.: Broadband Internet Performance: A View from the Gateway. In: Proceedings of the ACM SIGCOMM 2011 Conference. pp. 134–145. SIGCOMM '11, ACM (2011)
11. Thaler, D., Draves, R., Matsumoto, A., Chown, T.: Default Address Selection for Internet Protocol Version 6 (IPv6). RFC 6724 (Proposed Standard) (Sep 2012), `http://www.ietf.org/rfc/rfc6724.txt`
12. Wing, D., Yourtchenko, A.: Happy Eyeballs: Success with Dual-Stack Hosts. RFC 6555 (Proposed Standard) (Apr 2012), `http://www.ietf.org/rfc/rfc6555.txt`