Measuring Quality and Penetration of IPv6 Services

Matěj Grégr, Miroslav Švéda
Brno University of Technology
Faculty of Information Technology
Božetěchova 2, 612 66 Brno, Czech republic
Email: igregr,sveda@fit.vutbr.cz

Tomáš Podermański
Brno University of Technology
Center of Computer and Information Services
Antonínská 1, 601 90 Brno, Czech republic
Email: tpoder@cis.vutbr.cz

Abstract—**Much has been written about the depletion of IANA's pool of IPv4 addresses and the needs of urgent transition to IPv6. Several of the biggest content and service providers have enabled IPv6, however, there is still the lack of information about worldwide IPv6 adoption and quality of services that are measured in large scale. The aim of this paper is to present our methodology to measure penetration and quality of IPv6 adoption amongst web, mail and DNS service providers. The system is built to provide an open online access to IPv6 adoption overview for the whole community. The paper discusses our methodology and measurement system and compares them with other known solutions. The analysis of collected data is presented to help to understand the IPv6 penetration.**

Keywords–ipv6,dns,mail,speed,measurement

I. Introduction

Two crucial events concerning IPv6 happened in last two years. The IPv6 Day on June 8th 2011 and IPv6 Launch on June 6th 2012. These events should have motivated the activity of service and content providers to enable IPv6. The IPv6 Day was considered as a testing day on which mainly the content providers enabled IPv6 for 24 hours. The event should have tested how many clients would connect to their dual stack web servers, how big would be the bandwidth shift from IPv4 to IPv6 and how many clients would have problems with their IPv6 connectivity. Observation confirmed an increase of IPv6 traffic and changes in application mix. Google kept IPv6 enabled for several YouTube servers, thus the main contributor for IPv6 traffic since then is YouTube as Sarrar et al. [1] showed in their report and our observation confirms their results. Thanks to the results from IPv6 day, big content providers decided to turn on IPv6 permanently one year later on IPv6 Launch day.

However, big content providers such as Google, YouTube, Akamai etc. do not represent the whole Internet. Millions of other websites still remain without IPv6 addresses. What is the ratio of enabled IPv6 websites and services (e.g., mail servers, name servers) and is the ratio increasing or stagnant? Even more, if IPv6 is enabled, is the quality of service (e.g., response time) better, worse or the same as in IPv4? These are the questions the paper tries to answer.

The paper is organized as follows. Section II describes related work. Methodology, system's architecture and implementation are described in section III. Data analysis and comparison of projects measuring IPv6 penetration are discussed in section IV. Section V evaluates validity of results. Comparison of results with other projects is presented as well. Conclusion and future work are discussed in section VI.

II. Related Work

There are several approaches to measure IPv6 adoption. Measurement can be performed on the content provider's side. The methodology is typically based on inserting a small invisible image [2] or JavaScript fragment [3] into the content of content provider's web page. Client's browser executes the code or tries to download the image using IPv4, IPv6 or other transition technology (6to4, Teredo). The analysis of requests can show the number of clients that can or cannot connect to dual-stack Web servers and their latency. This methodology measures clients. It shows, if IPv6 is supported by application (web browser), operating system and client's ISP (Internet Service Provider). The number of "consumers" is essential for content providers, because without IPv6 active customers they will not invest their time and money to IPv6 transition. Currently, the numbers are around 2.5 %, as it is shown in Google's IPv6 statistics, with some exceptions like Romania and France as shown in Geoff Huston's statistics [4].

Another method is based on measuring the number of autonomous systems announcing an IPv6 prefix. The statistics inform how ISPs and transition networks are prepared to provide IPv6 connectivity for their customers. The good analysis was presented by Karpilovsky et al. [5]. Their study has shown, that almost half of prefixes are not used at all and the rest of them is announced long after the allocation. The drawback of the methodology is that the presence of an ISP's IPv6 prefix in the global BGP (Border Gateway Protocol) table does not mean availability of IPv6 for ISP's customers. Despite of that, the number of IPv6 prefixes in global routing table is increasing and can be seen as IPv6 transition progress.

One way of measuring the quality of IPv6 service is to measure the one-way delay. Zhou et al. [6], [7] published a study comparing IPv4 and IPv6 one-way delay between several measurement points. They analyzed the RIPE IPv6 data, which include traceroute, delay, and loss measurements among a list of IPv6 sites since 2003. Their conclusion is that native IPv6 paths have small 2.5 percentile and median end-to-end delay, and comparable delay to their IPv4 counterparts. The study presented by Arthur Berger [8] found that the latency is less over v4 than v6. For example, for destinations in the North America, the mean latency is 55 ms over v4 but substantially higher, 101 ms, over v6. The difference between the IPv4 - IPv6 performance is more likely correlated with a different forward AS-level path as was reported by Amogh Dhamdhere et al. [9]. The measurement by Mehdi Nikkhah et al. [10] compares performance of IPv6 and IPv4 protocol by measuring the web page download time. They found that control plane

 ISBN: 978-1-61208-330-8

(routing) is responsible for differences between IPv4 and IPv6, because the data plane performs comparably.

The methods described above provide information about IPv6 prefix allocation, measures of clients, or test the path delay. The paper presented by Jakub Czyz et al. [11] tries to analysis the IPv6 adoption from several perspectives - allocation of IPv6 prefixes, clients readiness etc. However, the ISPs are more interested in the number of content providers that enabled IPv6 for their services. The number can indicate how much IPv6 traffic service providers will expect in case they decide to deploy IPv6. The next important information is the quantity of sites reachable from IPv6 only networks. Several measurements have been published to describe these information [12]–[17]. These papers will be analyzed in more detail in the Section V.

This paper describes the measuring platform for gathering long-term statistics about IPv6 penetration amongst content providers. The measuring is focused on the penetration of web services, mail services and name services available over IPv6. The paper compares also IPv4 and IPv6 one-way delay for web services to measure the quality of connection from the users perspective. The results obtained during the information gathering are discussed together with comparison to other services using similar methodology.

III. Methodology, Architecture and Implementation

IPv6 penetration amongst content providers can be measured by checking the appropriate resource records (RR) in DNS. Table I briefly describes the RRs tested in our methodology. If any record contains `CNAME` record it is followed until a valid record for IPv4 (`A`) or IPv6 (`AAAA`) is found. The web services are also checked if an alternative record for IPv6 is available such as `www6, ipv6, www.v6`. The alternative records are sometimes used by network administrators for testing purposes. Availability of records in DNS is tested periodically and is described in more detail in the section III-A.

TABLE I
Resource records checked

Service	Record type	Test
Web	A	www.<domain>
	AAAA	www.<domain>
	AAAA	www6|ipv6|www.v6.<domain>
Mail	A for MX	<domain>
	AAAA for MX	<domain>
DNS	A for NS	<domain>
	AAAA for NS	<domain>

The response time of web services is measured for sites announcing web services over IPv4 and IPv6 protocol as well. Using IPv4 and IPv6, the system tries to connect to the remote web server. The time is measured between the first packet initiating relation (SYN) and the answer from the server (SYN, ACK). Comparing the results obtained via IPv4 and IPv6, it can be observed, which protocol has a better response. We are aware of the fact, that the IPv6/IPv4 response time can vary from location to location (due to asymmetric routing, missing IPv4/IPv6 peering, different number of hops, link quality, etc.), however, the goal is to measure the quality from the perspective of our users.

A. Architecture and Implementation

The architecture of the system is divided into several blocks as is depicted in the Figure 1. The core of the system is SQL database containing list of domains and statistical data. There are two subsystems connected to the database. The first one is responsible for querying DNS system. It takes the list of domains from the database and periodically updates data with the information gathered from DNS. The history of changes for each record is stored as well, allowing us to provide current and historic data for each domain in the database. The next subsystem performs the check of IPv6 quality by measuring the one path delay as was described in the previous section. It also updates information into domains database and stores historical information of each measurement.

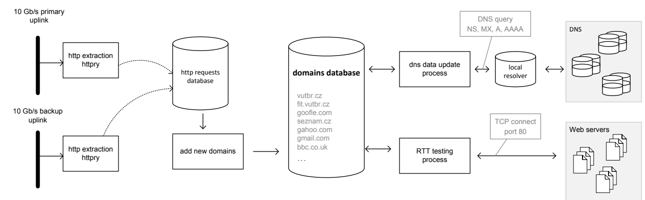

Figure 1. Architecture of the system.

B. Frequency of Updates

The DNS update process is executed every day, running in parallel using 100 threads to get the maximum performance. To update records and gather all data related to one domain, approximately 10 queries are needed. The system is able to obtain data for more than 100 domains per second. That means, the system must perform approximately 1000 DNS queries per second. Using database containing almost 9 million records, the whole update process takes approximately 2 days. Domains are queried in a random order to avoid the overloading of DNS servers. As was stated in the previous section, web domains supporting both protocols are checked to measure the quality of service. The update process is executed every day for domains that enabled IPv6 for a web service in past 24 hours. All dual stacked domains in the database are checked periodically once a week. The whole update process takes close to 6 hours for the database containing currently approximately 415 000 of dual stacked domains. However, the DNS system contains errors causing data inaccuracies. To get rid of invalid data, the system uses following rules:

- Queries returning addresses that belong to the reserved address space (private, link local) are ignored e.g., `192.168.0.0/16, 10.0.0.0/8, fe80::/10`.
- Address of loopback is ignored (`127.0.0.1/8, ::1/128`)
- Records with endless loop are ignored - a `CNAME` record points to another `CNAME` which points back to the original `CNAME`.
- The `CNAME` chain is followed up to 10 levels
- Manual patterns for domain names, e.g., domains containing random names like

 ISBN: 978-1-61208-330-8

www.335297538.acaoradical.com, www.335325653.acaoradical.com etc. are ignored. The list of patterns is managed manually.

C. Data Sources

The total number of tested domains will determine the precision of the system as is shown in the section IV. The results will be more accurate with the growing number of tested domains. A popular source of domains is the list of domains provided by Alexa The Web Information Company. It is possible to download daily updated list of the top 1 million sites for free in csv format from the Alexa's webpage. However, we believe that only the top list is not enough. For example, there are about 4 500 domains within Czech TLD in the Alexa's top list. The total number of domains in Czech TLD is approximately 1000000 thus the list contains less than 0.5 % registered CZ TLD domains. We believe that this extensively small number of domains is not enough to provide accurate results of IPv6 penetration. Another drawback of using Alexa data set only is that all sub domains are aggregated to appropriate TLD. It is quite common, that subdomains are used for different services, e.g., scholar.google.com and maps.google.com, however, subdomains are included neither in top 1 million sites nor in the list of top 500 sites per country. Using only top 1 million sites, it is impossible to check if subdomains are accessible via IPv6 or not. In the previous example, scholar.google.com is accessible only over IPv4 and maps.google.com is dual stacked website.

Other sources populating the database can be logs from mail server, DNS cache and reverse lookup of addresses accessing a network. The amount of domains that can be collected using these methods depends mainly on the network design and the number of users. In our case, we were able to collect initially approximately 100 000 of valid domains. Unfortunately, the growth of domains in the database was slow. The main reason for the slow growth is that we are a campus network without control of DNS and mail servers used by faculties of the university. To extend the number of domains and overcome the drawbacks of using only Alexa top sites, and DNS cache, the following solution was developed. Several probes are used in the Brno University of Technology's campus network listening to all HTTP requests performed by users. The requests are inspected using the httpry HTTP analyzer [18]. The advantage of this solution is that the probes are able to intercept all HTTP traffic from the whole university. The output from the analyzer is sent to a collector where the requests are stored in the *HTTP requests database* as depicted in the Figure 1. Once per day, the update process adds the new unique domains into the main database. Another source of new domains is a web interface that we provide on our site. Anybody can use the interface to check data of any domain and if the requested domain is not found, it will be added into the database. Lastly, the system is open to add any other source - e.g., import of whole zone file.

It is also important to remove domains that do not exist or have disappeared from DNS. If the domain does not contain any valid record (A, AAAA, CNAME, MX, NS), the domain is removed from the database including all related historical records. If the domain is added later using one of the approaches described above, it is treated as a new record in the database.

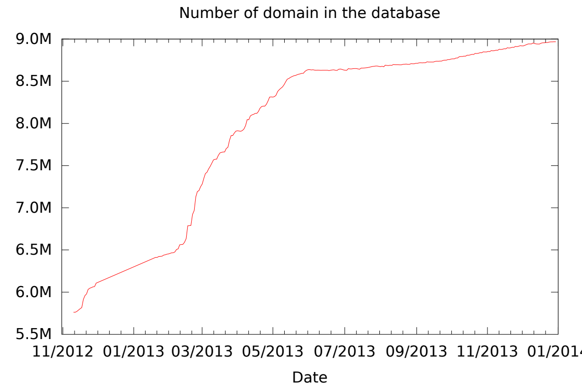

Figure 2. The number of domains in the database.

As of December 2013, the database contains approximately 9 million of domains [19]. The number is depicted in the Figure 2 and it is still growing as connected users are actively building the database with their HTTP requests. To summarize, the benefits of this approach is following:

- Independence on third party data sets (Alexa list).
- Visibility of IPv6 enabled sites, which are interesting for our users.
- Visibility of sub domains in a TLDdomain. This is useful because Alexa list does not provide information, about visited subdomains, only aggregated data.
- Long term solution with minimum maintenance.

IV. Data Analysis

Data collection is still ongoing. This paper presents data collected up to August 2013. The total number of web domains is defined as NWT, the number of web domains supporting web services over IPv6 as NWV6, the number of web domains supporting dualstack as NWDS and the number of web domains supporting IPv6 web through alternative name as NWA6.

- NWT - domains having at least one IPv4 or IPv6 record announced for www.<domain>.
- NWV6 - domains having at least one IPv6 (AAAA) record and do not have IPv4 record (A) announced for www.<domain>.
- NWV4 - domains having at least one IPv4 (A) record and do not have IPv6 record (AAAA) announced for www.<domain> and do not have alternative IPv6 record www6|ipv6|www.v6.<domain>
- NWDS - domains having both IPv6 and IPv4 records announced for www.<domain>.
- NWA6 - domains announcing any of www6|ipv6|www.v6.<domain> via IPv6 (AAAA) and do not announce IPv6 for a record www.<domain>.

TABLE II
IPV6 PENETRATION AMONGST WEB, MAIL AND DNS SERVICES - RATIO ON 15TH OF AUGUST 2013

	IPv4 only	IPv6 only	Dual stack	IPv6 alt. name
Web service	94.77 %	0.00150 %	4.78 %	0.45 %
Mail service	87.29 %	0.00031 %	12.71 %	
DNS service	74.58 %	0.00032 %	25.42 %	

The penetration ratio of IPv4 only sites is computed as shown in equation 1. Other ratios are computed using the same formula but the numerator is changed accordingly to NWV6 for IPv6 only sites ratio, NWDS for dualstack ratio etc.

$$NWV4ratio = \frac{NWV4}{NWT}100 \tag{1}$$

Based on these rules, we can analyze the data in our database to obtain the IPv6 penetration amongst web service as shown in Table II. As we can see, the majority of web pages is accessible using the IPv4 protocols or more precisely 99.9985 %. Despite the two big IPv6 events in last two years and the fact, that IPv4 addresses are depleted in APNIC and RIPE regions, the number of web sites accessible over IPv6 still stays on a very low level. The IPv6 only domains are usually without any meaningful content for end users. These are test websites used for testing user's IPv4/ IPv6 connectivity [4], sites where a www.<domain> has only AAAA record, but there is also record for <domain> which is accessible via IPv4 etc. DNS and mail services accessibility using IPv4 protocol is 99.99969 % and 99.99968 % but the availability using IPv6 protocol is much higher in comparison to web sites. The higher penetration of these services corresponds to deployment strategy for a new service, where an administrator goes from the key services to the less important ones.

V. VALIDITY OF THE RESULTS

Every measuring system must prove that data provided by the system are trustworthy. To provide the most accurate penetration of IPv6 services, all DNS data would have to be collected. This approach is however impossible thus we compared our data with similar projects. All projects with no exception use list of domains provided by Alexa and usually only a subset of the top 1 million is used.

The projects are compared in the Table III. The **Tests** column describes which tests projects run. The web test obtains the evidence of an AAAA record for selected domain. alt test checks an existence of alternative names for the domain (e.g., v6.<domain>). MX and NS tests are testing presence of AAAA record for mail and DNS services, DNSSEC and SPF verify the support for these services and avail test measures the quality of connection using both IPv4 and IPv6.

Results for global penetration and selected TLDs that in most statistics indicate the largest IPv6 penetration are compared in the Table IV. The comparison is based on the data from 17th August 2013 or latest available. As we can see in the Table IV the obtained results are very different. The Figure 3 shows one of the reasons for such distinction. The chart can be interpreted as follows: IPv6 penetration (y axis) is calculated for every number of domains (x axis). For example, IPv6 penetration for first 3 domains according to Alexa order

TABLE III
COMPARISON OF PROJECTS MEASURING IPV6 PENETRATION

Project	Data	Records	Tests	Freq.
IPv6matrix [12]	Alexa	top 1 million	web, alt, MX, NS	month
IPv6observatory [13]	Alexa	top 500/TLD	web, alt, MX, NS	daily
Eric Vyncke [14]	Alexa	top 50/TLD	web, alt, MX, NS	daily
Hurican Electric [15]	Alexa, zones	165 million	web, MX, NS	daily
Ari Keranen [16]	Alexa	10000	web, alt, avail	twice
6lab.cisco.com [17]	Alexa	top 500/TLD	web, alt	daily
cz.nic [20]	.cz TLD	1 million	web, MX, NS, DNSSEC	month
6lab.cz [19]	Alexa, Users	8.7 million	web, alt, MX, NS, avail, DNSSEC, SPF	daily

TABLE IV
RESULTS OF IPV6 PENETRATION PROVIDED BY DIFFERENT PROJECTS

Project	Global	.cz	.de	.fr	.ch
ipv6matrix	18.26 %	58.55 %	49.44 %	46.28 %	34.32 %
ipv6observatory	–	13.6 %	7.1 %	5.7 %	–
Eric Vyncke	7 %	30 %	10 %	8 %	46 %
Hurican Eletric	3.06 %	–	16.9 %	–	–
Ari Keranen	3 %	–	–	–	–
6lab.cisco.com	7.96 %	60.63 %	46.83 %	47.75 %	50.63 %
cz.nic	–	18.4 %	–	–	–
6lab.cz	4.79 %	14.45 %	12.09 %	3.08 %	2.34 %

is 100 % in .com TLD. If the ratio is evaluated for top 5 domains, it drops to 60 %. If we use the same data source and top 500 domains, the penetration decreases to 4.86 %. The results show that it strongly depends on the number of involved records used for calculation of IPv6 ratio. This also means, that IPv6 penetration calculated only on several top websites will tend to over-estimate the overall IPv6 penetration in web domains.

It can be seen that the IPv6 ratio becomes more stable for more than 1000 records. The increasing IPv6 penetration for some TLDs after 5 000 domains is caused by the fact, that the rest of the data set does not have any ranking, thus is sorted randomly. The number of requests must be increased at least to 10 000 records in the case of global IPv6 penetration where all domains and countries are involved. Using a smaller number of domains will again tent to over-estimate the IPv6 penetration. Similar observation was published by IPv6 Observatory [13]. A minimum of 10 000 domains is necessary to estimate the ratio of domains having an AAAA record as calculated from the top 1 million domains with an error in the range [-0.5;0.5]. The IPv6 Observatory, however, runs analysis using top 500 sites per TLD as was described in the Table IV.

Another reason for such a big difference between the projects' IPv6 penetration is that the projects use different methodologies. For example, the IPv6matrix project shows 58.55 % of IPv6 penetration for .cz TLD. This is due to the fact that IPv6matrix counts a domain as IPv6 enabled if the domain has an AAAA record for NS or MX or web or NTP server. As shown in the Table II, the penetration of AAAA records for NS and MX records is high, thus IPv6matrix shows much higher penetration than others. The 6lab.cisco.com project uses methodology which counts a weight of a domain. The weight of a domain is an approximation based on pageviews from Alexa statistics. The consequence of the weight is that more visited domains with IPv6 increase the IPv6 ratio for the country. This is based on an idea, that users are more likely to connect, spend time and access content on a very small number of sites as stated in their report [17]. To analyze the assumption, we used all HTTP requests made by all users in the

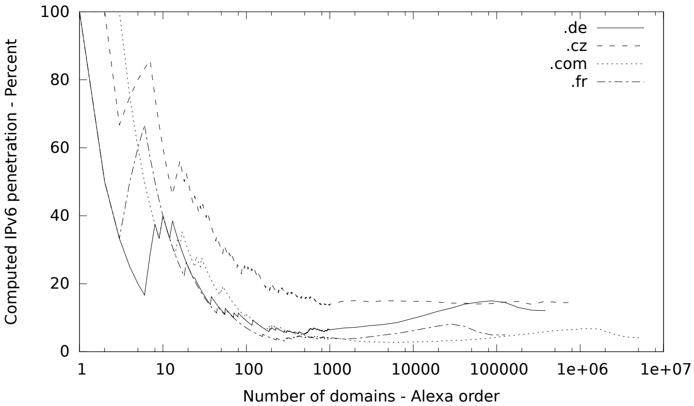

Figure 3. Dependence of the IPv6 ratio on the number of domains.

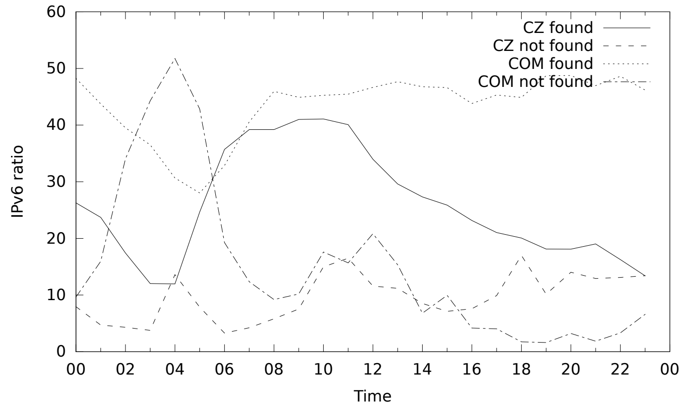

Figure 4. The IPv6 ratio of requests (not) matching Alexa list in a day.

TABLE V
HTTP REQUESTS FOUND/NOT FOUND IN THE ALEXA LIST

	Found	Not found	Ratio
All	79 216 045	15 555 245	83.58 %
.com	37 273 163	2 946 005	92.1 %
.cz	24 973 190	7 082 944	71.6 %

BUT network in a day. The requested domains were aggregated to the first subdomain after TLD - e.g., `maps.google.com` to `google.com`. This is due the fact, that Alexa list does not contain subdomains. The Table V describes the requested domains which were found or were not found in the Alexa list. The column **Ratio** stand for the chance, that a domain is found in the list. The results confirm that a majority of requests is found in the Alexa top list, but there is still a significant number of requests that is not found - almost 30 % of requests for .cz TLD are not found in the Alexa list. Relying only on Alexa list, the IPv6 penetration could be overestimated, because the domains which are not found in the Alexa list show much lower IPv6 penetration.

The `HTTP` requests for TLDs `.com` and `.cz` were further analyzed and the Figure 4 shows the difference between IPv6 penetration within the `HTTP` requests which were/were not found in the Alexa list. For example, the line `COM found` shows that IPv6 ratio is almost `50 %` amongst the `HTTP` requests which were found in the Alexa list; however, it is bellow `10 %` for the `HTTP` requests which were not found as shows the line `COM not found`. The higher IPv6 penetration during the night is mainly because of automatic updates, tests of IPv6 connection etc. - e.g., connections to `ds.download.windowsupdate.com`. These requests are not triggered by users. Although, it seems significant that approximately half of `HTTP` requests in `.com` TLD made by our users is IPv6 enabled, it is necessary to understand, that the reason lies in extensive usage of services like Google ads, Google analytics and social plugins (Facebook's 'like' button, etc.). These services are IPv6 enabled therefore, regarding a user visiting a domain which uses these services, there would be one request to obtain the domain and several other subrequests connecting to Facebook, Google and Twitter services.

The Figures 5 and 6 depict the performance between IPv4 and IPv6 connectivity using data from last two years. The round trip time is measured as described in the section III. The difference between these two protocols is counted as IPv4 - IPv6, thus negative values represents measurements where IPv6 is faster then IPv4. The graph in Figurer 5 plots the distribution of the difference using cumulative distribution function. The Figure 5 shows the frequency of measurements using histogram to provide better insight into the numbers in the <-5, 5>interval. Further analysis of the data shows, that IPv6 is more than 1 ms faster in 16% of measurements in 2012 and 22 % in 2013. IPv4 is faster more than 1 ms in 40 % of measurements in 2012 and 18 % in 2013. It can be seen, that the performance is usually similar - majority of measurements fits in the <-20, 20>interval. During the last year, the performance of IPv6 improved and usually is almost the same as IPv4. This can be a sign, that IPv6 is starting to be used for production traffic, however there is a substantial number of sites (3.51 %) we were not able to connect to. These non-working domains or slow IPv6 connectivity are nowadays usually handled by implementation of Happy Eyeballs in modern web browsers, however there is still a large fraction of clients without Happy Eyeballs implementation.

VI. Conclusion and Future Work

The paper describes a system and a methodology used for measuring IPv6 penetration amongst content providers. The contribution of the paper is following. The system gathers data from various sources and the latest data were analyzed. The system does not rely on the Alexa list only but the database is built actively from users' HTTP requests in a fully automatic way. The methodology for computing the IPv6 penetration has been compared with other projects measuring the IPv6 adoption with finding that IPv6 penetration amongst domains can vary from 13 to 60 percent. These differences are caused by distinct methodologies and strongly depend on the number of measured records. Unfortunately, only a subset of available data set is used, thus the IPv6 adoption presented by different projects is usually overestimated.

Based on the analyzed data, we are recommending to use at least 1000 domains to compute the penetration for a TLD. Global IPv6 ratio should be computed using as many domains as possible. The all domains should be treat equally without using any weight or only top x sites. The current IPv6

 ISBN: 978-1-61208-330-8

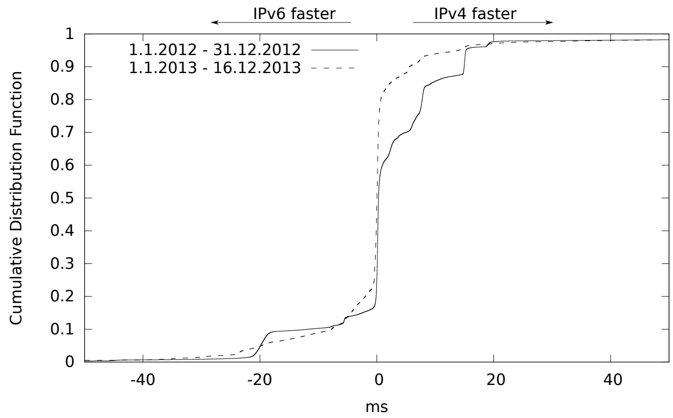

Figure 5. IPv4,IPv6 performance - CDF function

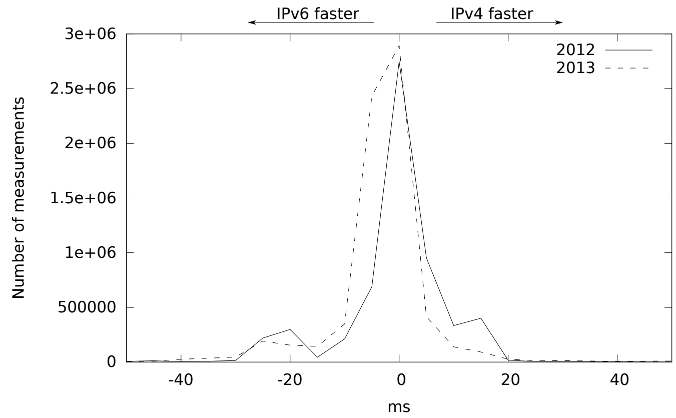

Figure 6. IPv4,IPv6 performance - number of measurements

methodologies measuring IPv6 penetration per TLD use top 50 or top 500 which overestimates the IPv6 penetration.

The presented system measures also quality of IPv6 connectivity and the results are analyzed with finding that overall IPv6 performance has been improved, although there are still number of sites (3.51 %) that are not accessible over IPv6 even though the website announce `AAAA` record.

The access to measured data together with all historical data is available and open to everybody on the project webpage [19]. As far as we know, our framework and methodology is the only one, which save all historical data of every page, thus we are able to see, when the `AAAA` record is added, removed etc. Other methodologies measuring the IPv6 adoption update the global IPv6 penetration only and does not keep the historical records. The future work could check sites' unavailability from different measuring points. We are also collecting the information, if a domain is signed with DNSSEC [21] or if a mail server provides SPF record [22]. This could also be a part of a future analysis. We plan to keep collecting and processing data, together with checking IPv4 and IPv6 performance.

ACKNOWLEDGMENT

This work is a part of the project VG20102015022 supported by Ministry of the Interior of the Czech Republic. This work was also supported by the project FIT-S-14-2299 Research and application of advanced methods in ICT.

REFERENCES

[1] N. Sarrar, G. Maier, B. Ager, R. Sommer, and S. Uhlig, "Investigating IPv6 Traffic," in Passive and Active Measurement. Springer Berlin / Heidelberg, 2012.

[2] G. Huston, "Testing IPv6 for World IPv6 Day," May 2011, [cit. 27.2.2014]. [Online]. Available: http://www.potaroo.net/ispcol/2011-05/ip6test.html

[3] L. Colitti, S. Gunderson, E. Kline, and T. Refice, "Evaluating IPv6 Adoption in the Internet," in Passive and Active Measurement. Springer Berlin / Heidelberg, 2010.

[4] G. Huston, "Measuring IPv6 - Country by Country," [online], July 2012. [Online]. Available: http://www.potaroo.net/ispcol/2012-07/v6report.html

[5] E. Karpilovski, A. Gerber, D. Pei, J. Rexford, and A. Shaikh, "Quantifying the Extent of IPv6 Deployment," in Network Measurement. Springer Berlin, 2009.

[6] X. Zhou and P. Van Mieghem, "Hopcount and E2E Delay: IPv6 Versus IPv4," in Passive and Active Measurement. Springer Berlin / Heidelberg, 2005.

[7] X. Zhou, M. Jacobsson, H. Uijterwaal, and P. Van Mieghem, "IPv6 delay and loss performance evolution," International Journal of Communication Systems, vol. 21, no. 6, 2008, pp. 643–663.

[8] A. Berger, "Comparison of performance over ipv6 versus ipv4," Akamai Technologies, 2011.

[9] A. Dhamdhere, M. Luckie, B. Huffaker, k. claffy, A. Elmokashfi, and E. Aben, "Measuring the deployment of IPv6: topology, routing and performance," in Proceedings of the 2012 ACM conference on Internet measurement conference. ACM, 2012, pp. 537–550.

[10] M. Nikkhah, R. Guérin, Y. Lee, and R. Woundy, "Assessing ipv6 through web access a measurement study and its findings," in Proceedings of the Seventh COnference on emerging Networking EXperiments and Technologies. ACM, 2011, p. 26.

[11] J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, and M. Bailey, "Measuring IPv6 Adoption," Technical Report TR-13-004, ICSI, Tech. Rep., 2013.

[12] O. M. Crpin-Leblond, "IPv6 Matrix Project," [cit. 27.2.2014]. [Online]. Available: http://www.ipv6matrix.org

[13] IPv6 Observatory, "Top-500 websites with AAAA records," [cit. 27.2.2014]. [Online]. Available: http://www.ipv6observatory.eu/indicator/

[14] E. Vyncke, "IPv6 Deployment Aggregated Status," [cit. 5.3.2014]. [Online]. Available: http://www.vyncke.org/ipv6status

[15] M. Leber, "Global IPv6 Deployment Progress Report," [cit. 16.12.2013]. [Online]. Available: http://bgp.he.net/ipv6-progress-report.cgi

[16] A. Keranen and J. Arkko, "Some Measurements on World IPv6 Day from an End-User Perspective," RFC 6948, 2013. [Online]. Available: http://www.ietf.org/rfc/rfc6948.txt

[17] H. Kaczmarek, "Internet IPv6 Adoption: Methodology, Measurement and Tools," [online], July 2012. [Online]. Available: http://6lab.cisco.com/stats/information

[18] J. Bitte, "HTTP logging and information retrieval tool," [cit. 27.2.2014]. [Online]. Available: http://dumpsterventures.com/jason/httpry/

[19] T. Podermaski and M. Grgr, "Worldwide online IPv6 penetration," August 2013, [cit. 27.2.2014]. [Online]. Available: http://6lab.cz/live-statistics/data-source

[20] CZ.NIC, ".cz statistics," [cit. 12.2.2014]. [Online]. Available: http://stats.nic.cz

[21] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "Resource Records for the DNS Security Extensions," RFC 4034, Mar. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4034.txt

[22] M. Wong and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1," RFC 4408, Apr. 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4408.txt