

Independent Submission
Request for Comments: 6948
Category: Informational
ISSN: 2070-1721

A. Keranen
J. Arkko
Ericsson
July 2013

Some Measurements on World IPv6 Day from an End-User Perspective

Abstract

During World IPv6 Day on June 8, 2011, several key content providers enabled their networks to offer both IPv4 and IPv6 services. Hundreds of organizations participated in this effort, and in the months and weeks leading up to the event worked hard on preparing their networks to support this event. The event was largely unnoticed by the general public, which is a good thing since it means that no major problems were detected. For the Internet, however, there was a major change on a short timescale. This memo discusses measurements that the authors made from the perspective of an end user with good IPv4 and IPv6 connectivity. Our measurements include the number of most popular networks providing AAAA records for their service, as well as delay and connection failure statistics.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6948>.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

- 1. Introduction 2
- 2. Motivation and Goals 3
- 3. Measurement Methodology 4
- 4. Measurement Results 5
 - 4.1. DNS AAAA Records 5
 - 4.2. TCP Connection Setup 8
 - 4.3. TCP Connection Delays 9
- 5. Conclusions 11
- 6. Security Considerations 12
- 7. Informative References 12
- Appendix A. Acknowledgments 14

1. Introduction

Many large content providers participated in World IPv6 Day on June 8, 2011. On that day, IPv6 [RFC2460] was enabled by default for 24 hours on numerous networks and sites that previously supported only IPv4. The aim was to identify any remaining issues with widespread IPv6 usage in these networks. Most of the potential problems associated with using IPv6 are, after all, of a practical nature, such as ensuring that the necessary components have IPv6 turned on, that configurations are correct, and that any implementation bugs have been removed.

Some content providers have been reluctant to enable IPv6. The reasons for this include delays for applications attempting to connect over broken IPv6 links before falling back to IPv4 [RFC6555] and unreliable IPv6 connectivity. Bad IPv6 routing has been behind many of the problems. Among the causes are broken 6to4 tunneling protocol [RFC3056] connectivity, experimental IPv6 setups that are untested and unmonitored, and configuration problems with firewalls. The situation is improving as more users and operators put IPv6 to use and fix the problems that emerge.

The World IPv6 Day event was largely unnoticed by the general public, which is a good thing since it means that no major problems were detected. Existing IPv4 connectivity was not damaged by IPv6, and also new IPv6 connectivity worked as expected in vast majority of cases. For the Internet, however, there was a major change on a short timescale. This memo discusses measurements that the authors made from the perspective of an end user with well-working IPv4 and IPv6 connectivity. Our measurements include the number of the most popular networks providing AAAA records for their service, as well as delay and connection failure statistics.

The rest of this memo is structured as follows. Section 2 discusses the goals of our measurements, Section 3 describes our measurement methodology, Section 4 gives our preliminary results, and Section 5 draws some conclusions.

2. Motivation and Goals

Practical IPv6 deployment plans benefit from accurate information about the extent to which IPv6 can be used for communication and how its characteristics differ from those of IPv4. For instance, operators planning to deploy dual-stack networking may wish to understand what fraction of their traffic would move to IPv6. This information is useful for estimating the capacity necessary to deal with the IPv6 traffic and the impacts to the operator's IPv4 infrastructure or carrier-grade NAT devices as their traffic is reduced. Network owners also wish to understand the extent to which they can expect different delay characteristics or problems with IPv6 connectivity. The goals of our measurements were to help with these topics by answering the following questions:

- o What fraction of the most popular Internet sites offer AAAA records? How did World IPv6 Day change the situation?
- o How do the traffic characteristics differ between IPv4 and IPv6 on sites offering AAAA records? Are the connection failure rates similar? How are round-trip times (RTTs) impacted?

There have been many measurements about some of these aspects from a service provider perspective, such as Google studies about broken connectivity between Google and its end users. Our measurements start from a different angle, by assuming good dual-stack connectivity at the measurement end, and then probing the rest of the Internet to understand, for instance, how likely there are to be IPv6 connectivity problems or what the delay differences are between IPv4 and IPv6. Similar studies have been performed by the University of Pennsylvania and Comcast [IPv6Monitor] and RIPE NCC [RIPEv6Day].

3. Measurement Methodology

We used the top 10,000 sites of the Alexa 1 million most popular sites list [Alexa] from June 1, 2011. For each domain name in the list, we performed DNS queries with different host names. For IPv4 addresses (A records), we used host name "www" and also performed a query with just the domain name. For IPv6 addresses (AAAA records), we used different combinations of host names that have been used for IPv6 sites, namely, "www6", "ipv6", "v6", "ipv6.www", "www.ipv6", "v6.www", and "www.v6".

All DNS queries were initiated in the order listed above (first "www" and just the domain name for A records, then "www", domain name, and different IPv6-host names for AAAA records), but the queries were done in parallel (i.e., without waiting for the previous query to finish). The first response for A and AAAA records and the corresponding host names were recorded. The queries had a 3-second retransmission timeout, and if there was no response for 10 seconds, all remaining queries for the site were canceled. We used a custom Perl script and the Net::DNS [net-dns] module for the DNS queries.

The measurement script used a bind9 DNS server running on the same host as was performing the measurement. The DNS cache of the server was flushed before each measurement run in order to detect the changes in the DNS records in real time. The host, and thus the DNS server, was not part of DNS IPv6 whitelisting agreements. (See Section 4.3 of [RFC6589] for information on DNS resolver whitelisting.)

The local network where the host performing the measurements was had native IPv6 (dual-stack) connectivity. The IPv6 connectivity to the local network was provided by an IPv6-over-IPv4 tunnel from the network's default router to the ISP's IPv6 peering point.

After obtaining IP addresses for the site, if a site had both A and AAAA records, a simple C program was used to create TCP connections to port 80 (HTTP) simultaneously using both IPv4 and IPv6 to the (first) IP addresses discovered from the DNS. The connection setup was repeated up to 10 times, giving up after the first failed attempt (but only after normal TCP retransmissions). The connection setup delay was measured by recording the time immediately before and after the connect system call. The host used for measurements was a regular Linux PC with a 2.6.32 version kernel and a dual-stack Internet connection via Ethernet.

The measurements were started one week before World IPv6 Day (on Wednesday, June 1, 17:30 UTC) and ran once every three hours until July 11. One test run took from two to two-and-a-half hours to complete.

The accuracy and generality of the measurement results are limited by several factors. While we ran the tests at three different sites, most of the results discussed in this document present snapshots of the situation from just one measurement point, the Ericsson Research Finland premises, near Helsinki. Also, since one measurement run took quite a long time, the network characteristics and DNS records might have changed even during a single run. The first DNS response was used for the TCP connectivity tests, and this selection might have resulted in selection of a non-optimal host; yet, a slight preference was given to the "www" and only-domain-name records since their queries were started before the others. While the host performing the measurements was otherwise idle, the local network was in regular office use during the measurements. The connectivity setup delay was collected in user space, with a regular, non-real-time kernel implementation, resulting in small inaccuracies in the timing information.

4. Measurement Results

4.1. DNS AAAA Records

The number of top 10,000 sites with AAAA DNS records before, during, and after World IPv6 Day is shown in Figure 1. The measurements performed during World IPv6 Day are shown on the light gray background.

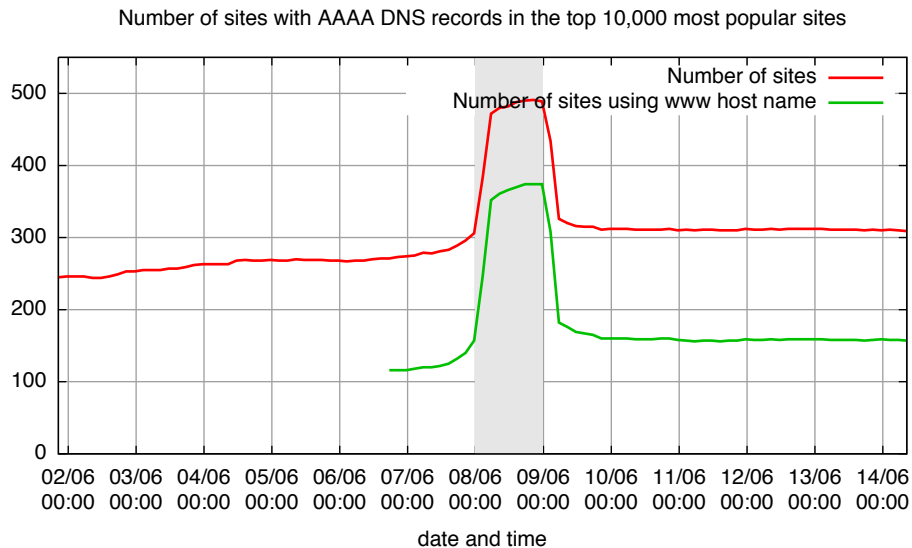


Figure 1: Number of sites with AAAA DNS records in the top 10,000 most popular sites

When the measurements began on June 1, 245 sites (2.45%) of the top 10,000 sites had both A and AAAA records. During the following days, the number of such sites slowly increased, reaching 306 sites in the measurement that was started at 22:30 UTC on June 7, the evening before World IPv6 Day. When World IPv6 Day officially started, the following measurement (at 01:30 UTC) recorded 383 sites, and the next one 472 sites. During the day, the number of sites with AAAA records peaked at 491 (4.91% of the measured 10,000 sites), at 19:30 UTC.

When World IPv6 Day was over, the number of AAAA records dropped nearly as fast as it had increased just 24 hours earlier. However, the number of sites stabilized at around 310 and did not drop below 300 after that, resulting in over 3% of the top 10,000 sites still having AAAA records at the end of our measurements, on July 11.

While 274 sites had IPv6 enabled in their DNS for some of the tested host names one day before World IPv6 Day, only 116 had it for the "www" host name that is commonly used when accessing a web site. The number of "www" host names with AAAA records more than tripled during World IPv6 Day, reaching 374 sites for 3 consecutive measurement runs (i.e., for at least 6 hours). Also, the number of AAAA records for the "www" host name dropped steeply after the day and remained at around 160 sites after that.

Similar but more pronounced trends can be seen if only the top 100 of the most popular sites are taken into considerations, as shown in Figure 2.

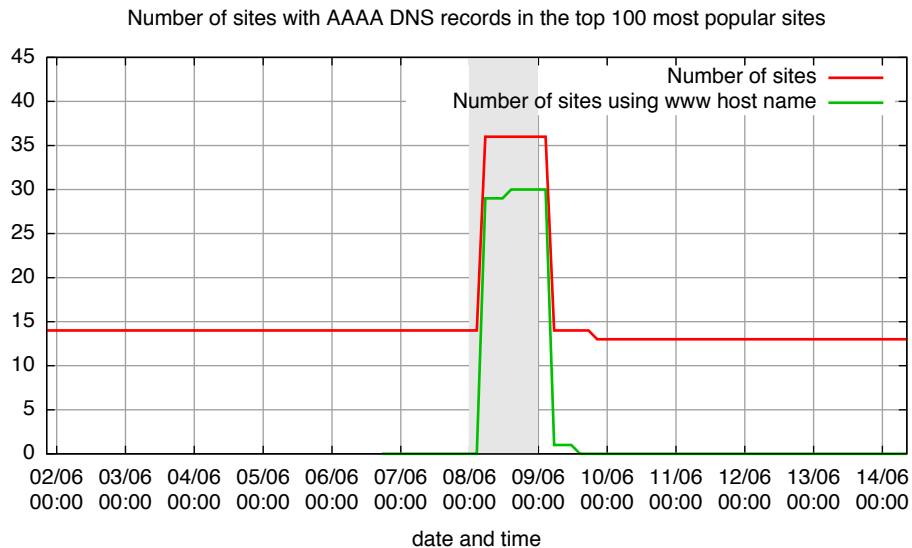


Figure 2: Number of sites with AAAA DNS records in the top 100 most popular sites

Here, the number of sites with some of the tested host names having a AAAA record was initially 14; then, it jumped to 36 during the day and eventually dropped to 13. Also, while none of the top 100 sites apparently had a AAAA record for their "www" host name before and after World IPv6 day, during the day the number peaked at 30. Thus, roughly one third of the 100 most popular sites had IPv6 enabled for World IPv6 Day.

Two other test sites in Sweden and Canada experienced similar trends with the DNS records. However, one of the sites used an external DNS server that was part of whitelisting agreements. There, the number of sites with AAAA records before World IPv6 Day was already higher (more than 400), and hence the impact of the day was smaller, because the amount of sites increased to the same numbers as seen by the test site in Finland. With the whitelisted DNS server, the number of sites remained above 450 after the day.

4.2. TCP Connection Setup

To test whether the IP addresses given by the DNS actually provide connectivity to the web site and whether there is any difference in the connection setup delay and failure rates with IPv4 and IPv6, we attempted to create TCP connections for all domains that contained both A and AAAA DNS records. The fraction of sites for which the first DNS response gave addresses that were not accessible with TCP to port 80 over IPv4 or IPv6 is shown in Figure 3.




Figure 3: TCP connection setup failure ratio (for the first DNS response)

There was a baseline failure rate with IPv4 of around 1-3% that was fairly static throughout the test period. For hosts with AAAA records, the fraction of inaccessible sites was much higher: in the beginning, up to one fourth of the tested hosts did not respond to TCP connection attempts. Much of this was likely due to the various test sites with different "IPv6 prefixes" (as discussed in Section 3); in the first run, more than half of the tested sites with AAAA records used them for the first DNS response. Also, some of the hosts were not even supposed to be accessed with HTTP but provided AAAA records for other purposes, while some sites had clear configuration errors, such as localhost or link-local IPv6 addresses.

As World IPv6 Day came closer, the number of inaccessible IPv6 sites decreased slowly and the number of sites with AAAA records increased at the same time, resulting in the failure ratio dropping to roughly 20% before the day. During the day, the number of IPv6 sites increased rapidly, but also the number of failures decreased, and hence, at the end of the day, the failure ratio dropped to just above 10%. After World IPv6 Day, when many of the participating IPv6 hosts were taken off-line, the fraction of failed sites for IPv6 increased. However, since there was no increase in the absolute number of failed sites, the fraction of inaccessible sites remained at a lower level, between 15% and 20%, than before the day.

4.3. TCP Connection Delays

For sites that were accessible with both IPv4 and IPv6, we measured the time difference between establishing a TCP connection with IPv4 and with IPv6. We took the median (as defined in Section 11.3 of [RFC2330]) of the time differences of all 10 connections, and then the median and mean (of the median) over all sites. The results are shown in Figure 4.




Figure 4: TCP connection setup delay differences (IPv4 - IPv6)

In general, the delay differences were small: the median of medians remained less than 3 ms off from zero (i.e., IPv4 and IPv6 delays were equal), and even the mean, which is more sensitive to outliers, remained within ± 5 ms most of the time, with the greatest spikes reaching to roughly -15 ms (i.e., the mean of median IPv6 delays was 15 ms larger than for IPv4 delays). Closer inspection of the results shows that the spikes were often caused by only one site or a handful of sites with bad connectivity and multiple retransmissions of TCP SYN and ACK packets, resulting in connection setup delays an order of magnitude larger than those for the other sites.

Surprisingly, the median delay for IPv6 connections was, in most cases, equal to or smaller than the IPv4 delay, but during World IPv6 Day, the IPv6 delays increased slightly and became (as a median) slower than their IPv4 counterparts. One reason for such an effect was that some of the sites that enabled IPv6 for World IPv6 Day had an extremely low IPv4 delay, less than 10 ms (e.g., due to the Content Delivery Network (CDN) provider hosting the IPv4 site), but a "regular" delay (over 100 ms) for the IPv6 host.

More detailed analysis of the TCP connection setup delay differences, and the reasons for them, is left for future work.

5. Conclusions

World IPv6 Day had a very visible impact on the availability of content over IPv6, particularly when considering the top 100 content providers. It is difficult to find other examples of bigger one-day swings in some characteristics of the Internet. However, the impact on end users was small, given that when dual-stack works correctly, it should not be visible at the user level, and given that IPv6 availability for end users themselves is small.

The key conclusions are as follows:

- o On that day, there was a large jump in the number of content providers providing AAAA DNS records.
- o On that day, there was a smaller but apparently permanent increase in the number of content providers supporting AAAA.
- o Large and sudden swings in the relative amount of IPv4 vs. IPv6 traffic are possible merely by supporting a dual-stack access network and having a few large content providers offer their services either globally or to a particular network over IPv6.
- o A large fraction of sites that published AAAA records for a name under their domain (be it "www", "www6", or something else) were actually not responding to TCP SYN requests on IPv6. This fraction was far higher than that which we've seen in our previous measurements, and we are still determining why that was the case. Measurement errors or problems on our side of the network cannot be ruled out at this stage. In any case, it is also clear that as new sites joined, incomplete or in-progress configurations create more connectivity problems in the IPv6 Internet than we've seen before. Other measurements are needed to verify what the general level of IPv6 connectivity is to addresses publicly listed in AAAA records.
- o Even if the overall level of connection failures was high, activities on and around the IPv6 day appear to have caused a significant permanent drop in the number of these failures.

- o When IPv6 and IPv4 connectivity were both available, their delay characteristics appeared very similar. In other words, most of the providers that made IPv6 connectivity available appear to have provided a production-quality network. TCP connection setup delay differences due to RTT differences between IPv4 and IPv6 connections were, in general, low. In the remaining differences in our measurements, random packet loss played a major role. However, some sites could experience considerable differences simply because of different content distribution mechanisms used for IPv4 and IPv6 content.

It is promising that the amount of the most popular Internet content on IPv6 was surprisingly high, roughly one third of top 100 sites (during World IPv6 Day or with whitelisting enabled). However, other content on the Internet forms a long tail that is harder to move to IPv6. For instance, only 3% of the 10,000 most popular web sites provided their content over IPv6 before World IPv6 Day. On a positive note, the top 100 sites form a very large part of overall Internet traffic [Labovitz], and thus even the top sites moving to IPv6 could represent a significant fraction of Internet traffic on IPv6. However, this requires that users be enabled to use IPv6 in their access networks. We believe that this should be the goal of future global IPv6 efforts.

6. Security Considerations

Security issues have not been discussed in this memo.

7. Informative References

- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, April 2012.
- [RFC6589] Livingood, J., "Considerations for Transitioning Content to IPv6", RFC 6589, April 2012.
- [net-dns] Fuhr, M., "Net::DNS", <<http://www.net-dns.org/>>.

[IPv6Monitor]

University of Pennsylvania and Comcast, "IPv6 Monitoring @ Penn", 2012, <<http://mnlab-ipv6.seas.upenn.edu/>>.

[RIPEv6Day]

RIPE NCC, "World IPv6 Day Measurements", <<http://v6day.ripe.net/>>.

[Alexa]

Alexa the Web Information Company, "Alexa Top 1,000,000 Sites", <<http://s3.amazonaws.com/alexa-static/top-1m.csv.zip>>.

[Labovitz]

Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., and F. Jahanian, "Internet Inter-Domain Traffic", Proceedings of ACM SIGCOMM 2010, August 2010.

Appendix A. Acknowledgments

The authors would like to thank Suresh Krishnan, Fredrik Garneij, Lorenzo Colitti, Jason Livingood, Alain Durand, Emile Aben, Jan Melen, and Tero Kauppinen for interesting discussions in this problem space. Thanks also to Tom Petch and Bob Hinden for thorough reviews and many helpful comments.

Authors' Addresses

Ari Keranen
Ericsson
Jorvas 02420
Finland

EMail: ari.keranen@ericsson.com

Jari Arkko
Ericsson
Jorvas 02420
Finland

EMail: jari.arkko@piuha.net