

Understanding Current IPv6 Performance: A Measurement Study

Yi Wang¹, Shaozhi Ye², Xing Li³

Department of Electronic Engineering, Tsinghua University,
Beijing 100084, P. R. China

¹wangyi@ns.6test.edu.cn

²ys@compass.net.edu.cn

³xing@cernet.edu.cn

Abstract

Much work has been done on IPv6 standards and testbeds deployment. However, little is known about the performance of the real IPv6 Internet, especially from the perspective of end users. In this paper, we present a measurement study of current IPv6 performance conducted from CERNET¹. We study 585,680 packet-level traces with 133,340 million packets collected from 936 IPv6/IPv4 dual-stack Web servers located in 44 countries. We present a comprehensive performance comparison of IPv6 and IPv4, including connectivity, packet loss rate, round-trip time, etc. Our measurement results show that IPv6 connections tend to have smaller RTTs than their IPv4 counterparts, but suffer higher packet loss rate at the same time. We also notice that tunneled paths do not show a notable degraded performance compared with native paths. To our knowledge, this paper is the first performance study based on both large scale TCP and ICMP traffic measurement in real IPv6 Internet.

1. Introduction

IPv6 is proposed by IETF to provide the Internet with larger address space and better performance [1]. In the past decade, a lot of work has been done on the protocol design [4], connection [5], [6], routing mechanism [7], and transition mechanisms [8], [9] of IPv6.

After its first decade of protocol design and testing, the IPv6 Internet is now in a transition phase from experimental research networks to global operational networks. Performance is always an important issue in any public service network. As to IPv6, it is an even

more complex question mark. On the one hand, IPv6 Internet is not constructed on a clear map. It faces transition problem from the very beginning. Tunneling plays an important role in IPv6 testbeds and the migration from IPv4 to IPv6. IPv6-in-IPv4 tunnels are widely used where no native IPv6 connectivity available. However, it is now believed that tunnels degrade the network performance and reliability [2], [3]. On the other hand, little is known about the actual performance of the growing IPv6 Internet, especially from the perspective of an average IPv6 end user. We believe that end users are responsible for a large fraction of Internet traffic, and their experiences of network performance are helpful to the design and deployment of IPv6 networks.

Recently, more attention is paid to the performance and operational issues of IPv6 networks [2], [3]. In [2], the authors discuss the IPv6-in-IPv4 tunnel discovery issue, and propose a set of techniques to infer tunnels by combinations of basic methods including Path MTU discovery, DNS lookups, IP spoofing, hop limit manipulation and IPv6 header modifying. Their experimental results show that even “native” networks have to go through tunnels to reach more than 60% of the entire IPv6 Internet. The authors of [3] argue that poorly managed experimental IPv6 sites are one of the major hurdles to the perceived quality of the IPv6 Internet. With focuses on troubleshooting, they select a group of IPv6/IPv4 dual-stack nodes by DNS lookups. And then they study the IPv6:IPv4 RTT ratios by dual-stack *ping* and do path analysis using traceroute with Path MTU discovery from three different locations in Japan and Spain.

The novelty of our work in this paper is emphasized by the fact that no previous work attempts to characterize the performance of IPv6 Internet by data traffic measurement (e.g., TCP) other than ICMP *ping*.

¹ China Education and Research Network

Besides, by doing path MTU discovery and studying the joint distribution of IPv6/IPv4 RTTs, we study the path IPv6 path characteristic. Our measurement results indicate that current IPv6 Internet suffers from lower connectivity and higher packet loss rate than its IPv4 counterpart, but experiences smaller average RTTs at the same time. We do not find evidence that tunneling notably degrades IPv6 performance.

The reminder of this paper is organized as follows. In Section 2, we propose a methodology of IPv6/IPv4 dual-stack Web server measurement. In Section 3, we present the measurement results and performance analysis comparing IPv6/IPv4 counterparts. In Section 4, we study the IPv6 path characteristic of being native or tunneled, and its effect on IPv6 performance. In Section 5, we conclude the whole paper.

2. Methodology

As an active measurement study, we design our methodology as follows. First, we obtain a list of dual-stack Web servers from our IPv6 Web search database which has been collecting information of IPv6 Web sites since 2001. Then we perform large scale data gathering by crawling web pages and logging the connection traces. We also run *ping* and *ping6* tests to study the performance differences between TCP packets and ICMP packets.

2.1. Dual-Stack Web Server List

WWW (more precisely, the HTTP on port 80) is one of the most widely used service in current IPv4-based Internet [12]. For the purpose of smooth transition from IPv4 to IPv6, more and more Web servers are implemented and configured with both IPv4 and IPv6 protocol stacks. We choose these dual-stack Web sties as our data source since we can gain better understanding of IPv6 performance as well as its distinctive problems by comparison with its IPv4 counterpart.

We have been monitoring the evolution of Web sites on the IPv6 Internet since May 2001 and found that the number of IPv6 accessible Web sites keeps growing steadily [13]. For the experiment in this paper, we use an initial hostname list of 1,306 IPv6 Web sites. These Web sites has been observed on service at least for one month since May 2001. On Sept. 1st, 2004, the beginning of our experiment, 1,235 of them can still be resolved to either IPv6 or IPv4 addresses (or both). In this paper, we are interested in those dual-stack Web servers with both IPv6 and IPv4 global routable addresses. After removing duplicates with identical IPv6/IPv4 address pairs but different hostnames, we

finally obtain a list of 936 dual-stack Web servers, which is used for our data collection.

2.2. Design of Testbed and Experiment

In order to imitate the experience of average IPv6 end users, our testbed consists of a PC running FreeBSD 4.10 operating system with IPv4 and IPv6 dual stacks, as well as 100Mbps Ethernet links to both IPv4 and native IPv6 Internet through CERNET.

During the experiment, we run two programs, *wget* [14] and *tcpdump* [15]. We use *wget* and its IPv6 patched version *wget6* to download the same files from Web servers through IPv4 and IPv6 links respectively. The first level files of Web sites are often the index.htm pages, which are usually no more than several Kbytes. To avoid being biased by short connections, we crawl two or three levels deeper for larger files. We use a script to start multiple processes of *wget* and *wget6* together and make the IPv4 and IPv6 crawling jobs of the same sites working at almost the same time. Meanwhile, we make two *tcpdump* listen to the IPv4 and IPv6 interfaces respectively and log traces. We use a cronjob to collect data and dump traces periodically.

3. Measurement Results

3.1. Experiment Summary

We performed Web crawling and trace logging from Sept. 1st to Sept. 10th, 2004. Altogether, we collected over 585,680 connection traces of about 133.34 million packets from the 936 dual-stack web servers. In terms of bytes, about 63.8G bytes data are from IPv6 connections and the other 109.2G bytes data are from IPv4 connections.

Table 1. Distribution of dual-stack web servers by country codes

JP: 298	CH: 16	MY: 10	AT: 6	BR: 2	IN: 1
NL: 133	EE: 16	PL: 10	KR: 6	ID: 2	LU: 1
US: 83	FI: 16	NO: 9	BE: 5	IS: 2	UK: 1
DE: 77	CA: 13	SE: 8	GR: 5	PH: 2	YU: 1
IT: 39	DK: 13	TW: 8	PT: 4	RU: 2	
GB: 32	AU: 12	ZA: 8	SK: 4	BG: 1	
IE: 28	CZ: 12	CN: 7	MX: 3	HK: 1	
FR: 20	ES: 10	HU: 7	TH: 3	HR: 1	

Table 1 shows the distribution of the dual-stack nodes by their country codes, representing 44 countries. (It is possible that the real location of a node is different from the registered country.). We obtain the country codes for 936 nodes from whois databases of APNIC, RIPE and ARIN.

Figure 1 shows the hop counts distribution of the connections in the experiment. In our experiment, the

average hop count is 8.7 in IPv6 Internet, and 17.5 in IPv4 Internet. Moreover, the hop count in IPv4 tends to be diversified in a wider range than in IPv6. The simplicity of topology and the existence of tunnels may account for these discrepancies.




Figure 1. Distribution of hop counts of IPv4 and IPv6 connections in the experiment

3.2. Connectivity

Connectivity is a fundamental requirement to a network. Most of the time in current IPv4 Internet, connectivity is not such a big problem as it was decades ago. However, it might be an issue in the IPv6 Internet.

To examine the connectivity of the 936 dual-stack Web servers, we use both the traditional *ping/ping6* and *wget* to test their connectivity. There is possibility that some server is up and responses to ICMP *ping* packets, but has nothing on service at the 80 port, or vice versa (probably the ping service is disabled).

Table 2. Average accessibility of dual-stack Web servers by *ping/ping6* in different regions

IPv6		✓	✓	×	×
IPv4		✓	×	✓	×
Total	936 (100%)	642 (68.6%)	89 (9.5%)	132 (14.1%)	73 (7.8%)
RIPE	478 (100%)	312 (65.3%)	60 (12.5%)	76 (15.9%)	30 (6.3%)
JP	298 (100%)	235 (78.9%)	23 (7.7%)	24 (8.1%)	16 (5.3%)
ARIN	108 (100%)	68 (63.0%)	2 (1.8%)	24 (22.2%)	14 (13.0%)
APNIC	52 (100%)	27 (41.9%)	4 (7.7%)	8 (15.4%)	13 (25.0%)

✓: accessible; ×: inaccessible

Table 2 summarizes the accessibility of the 936 servers by partitioning them into four general regions. Table 3 shows the results of *wget/wget6* experiment. Since Japan contributes a large fraction of servers with

relatively close locations, we separate them from the non-JP APNIC servers as [3] does. We also merge the small number of LACNIC servers into the ARIN ones. The Japanese servers perform best connectivity, since majority IPv6 links from CERNET to RIPE, ARIN and other APNIC nodes first go through Japan. In Table 2, APNIC servers have exceptional low connectivity by *ping* for both IPv4 and IPv6. Comparing with Table 3, it is probably because many APNIC Web servers have disabled the *ping* service.

Table 3. Average accessibility of dual-stack Web servers by *wget/wget6* in different regions

IPv6		✓	✓	×	×
IPv4		✓	×	✓	×
Total	936 (100%)	584 (62.4%)	73 (7.8%)	191 (20.4%)	88 (9.4%)
RIPE	478 (100%)	289 (60.4%)	53 (11.1%)	94 (19.7%)	42 (8.8%)
JP	298 (100%)	202 (67.8%)	17 (5.7%)	52 (17.4%)	27 (9.1%)
ARIN	108 (100%)	63 (58.3%)	2 (1.9%)	30 (27.8%)	13 (12.0%)
APNIC	52 (100%)	30 (57.7%)	1 (1.9%)	15 (28.8%)	6 (11.5%)

✓: accessible; ×: inaccessible

3.3. Packet Loss

Due to the development and enormous diversity of the Internet, average packet loss rate in different studies is reported in a large range. It is between 0.36% and 3.54% by Boralla et al. [16] based on the study of speech data transmission, between 1.38% and 11% Yajnik et al. [17] based on the measurement at the Mbone receiver, between 2.7% and 5.2% by Paxson [18] based on his long-term experiment of bulk data transference. The most similar data source as our paper may be the one used by Balakrishnan et al. [19]. They analyzed the dynamics of a large number of TCP web sessions at a busy Web server, and reported an average packet loss rate of 0.49%.

In our measurement, the IPv6 and the IPv4 connections have an average packet loss rate of 3.09% and 0.76% respectively. Table 4 shows the detailed distribution. We notice that the IPv6 links tend to experience more packet loss than the IPv4 ones as a whole.

Figure 2 shows the packet loss rate of the four general regions and some representative countries. We notice that the IPv6 packet loss rates of Germany (DE) and Ireland (IE) nodes are extraordinarily high. After examining the corresponding traces, we find that there are one or more “special links” from each of these countries to our testbed. These links are of large volume of data with high packet loss rate (exceed 5%, sometimes even 10%), which increase the average loss

rate of the country as a whole. However, some other links using the same BGP routing table do not experience such high loss rate, so we argue problems of these “special links” probably occur in the access links near the Web servers, not in the IPv6 backbone.

Table 4. Packet loss distribution by wget/wget6 measurement

IPv4	No loss	No loss	Loss	Loss
IPv6	No loss	Loss	No loss	Loss
Percent	49.3%	27.8%	8.8%	14.1%





Figure 2. Distribution of packet loss rate

3.4. Round-trip Time

Round-trip time (RTT) is an important parameter to indicate the quality-of-service (QoS) of a network. Moreover, in the scenario of this paper we find its distribution also an insightful tool to reveal the network topology and deployment information.

Figure 3 shows the distribution of the observed RTTs in scatter plot. It plots about 584 RTT value pairs of the Web servers accessible by both IPv6 and IPv4 web crawlers in the 10-day measurements. For each value pair, the IPv4 RTT is plotted on the X-axis and the IPv6 RTT is plotted on the Y-axis. We also

draw the unity line, $y = x$ to help the comparison. Accordingly, figure 4 shows the probability distribution function (PDF) of the IPv4/IPv6 RTTs.




Figure 3. Scatter plot of IPv4/IPv6 RTT distribution




Figure 4. PDF of IPv4/IPv6 RTT distribution

Figure 3 shows that different regions tend to have different typical RTT value ranges, which accords with our intuition. The RIPE nodes have the most concentrated RTT distribution, considering its largest total number. Interestingly, they cluster into two narrow bands with IPv6 RTT range approximately from 320 ms to 420 ms and from 450 ms to 550 ms respectively. The Japanese nodes also exhibit this characteristic, which indicates the simplicity of the current IPv6 Internet topology. The ARIN nodes do not have such a notable clustering characteristic as the RIPE and Japanese nodes. Most of the ARIN nodes are around the unity line. In spite of their small total number, the APNIC nodes have large variance of RTT values due to their topology diversity. It is clear in Figure 4 that the three major peaks of the IPv6 PDF

curve correspond to the three major IPv6 RTT clusters in Figure 3. Figure 3 and Figure 4 also indicate the majority of the dual-stack nodes have smaller IPv6 RTTs than IPv4 RTTs.

4 Native or Tunneled?

So far, we have got a rough picture of current IPv6 performance. In this section, we take a detailed look at the IPv6 path characteristic, namely native or tunneled.

4.1. What's the role of tunnels?

We first get the native/tunneled paths proportion by path MTU (Maximum Transmission Unit) discovery [10], [11]. The result is summarized in Table 5. A path MTU of 1500 indicates a native IPv6 path, and certain MTU values can suggest the persistence of a tunnel. We note that the most common MTU is 1280, which is minimum IPv6 MTU defined by [4]. The second most common MTU value other than 1500 is 1480, a typical IPv6-in-IPv4 tunnel MTU in popular tunnel implementations. The paths with a MTU of 1476 are probably due to GRE tunnels.

Table 5. IPv6 Path MTU Summary

Path MTU	Total	RIPE	JP	ARIN	APNIC
1500	234	103	104	19	8
1480	130	70	31	18	11
1476	28	11	-	17	-
1472	1	-	-	-	1
1460	10	10	-	-	-
1280	415	205	144	44	22
Other	10	8	2	-	-
unreach	108	71	17	10	10
Total	936	478	298	108	52




Figure 5. Percentage of tunneled paths

Figure 5 shows the percentage of tunneled paths by the four regions and major countries. We note that Japanese nodes have the highest native path percentage among the four regions, probably because of their efforts to promote native IPv6 deployment and the good native IPv6 connection with CERNET. However, among the major countries, Ireland (IE) nodes stand out with the lowest tunneled paths percentage.

4.2. How do tunnels perform?

It is recently believed that tunnels degrade the network performance and reliability [2], [3]. Given the large percentage tunneled paths take in currently IPv6 Internet, it is important to know their actual performance.





Figure 6. Scatter plot of IPv4/IPv6 RTT distribution of major tunneled nodes: (a)JP (b)RIPE

Figure 6 shows the IPv4/IPv6 RTT scatter plots of dual-stack nodes with IPv6 connections through major tunnel end points to our test host. Here we define an intermediate node as a tunnel end point if the path

MTU from that node to our test host is 1500 and the MTU value drops to a typical tunnel value after that hop. And we define a tunnel end point as a “major” one if there are more than four paths go through it to different nodes. It is obvious that tunneling does not seem to introduce notable extra delays. Most tunneled IPv6 paths have smaller RTT values than their IPv4 counterparts, and the reductions are often more than 100 ms. Compare with Figure 3, we notice that more native IPv6 paths are above the $y = x$ line than tunneled IPv6 paths.

We also study the correlation between tunneling and reliability in terms of packet loss rate. Recall Figure 2(b), Germany (DE) and Ireland (IE) nodes have the highest IPv6 packet loss rate. The high loss rate of Germany nodes are all from tunneled paths, however, the high loss rate of Ireland nodes are all from native paths. Our measurement result does not imply any generic correlation between tunneling and packet loss rate.

5. Conclusions

Today, IPv6 is in an important transition state from testbeds to operational networks. Future IPv6 development and engineering require clear understanding of current IPv6 performance. In this paper, we design an active measurement methodology to compare IPv6 and IPv4 performance from the perspective of a common end user. Our methodology includes obtaining and screening of IPv4/IPv6 dual-stack Web servers as well as Web file crawling and trace logging techniques. We perform our measurement from native IPv6 network in CERNET to 936 dual-stack web sites worldwide and present the first experimental results of IPv6 performance based on both TCP traffic and ICMP *ping* in real Internet. Our measurement results indicate that there is still space left to improve the IPv6 performance, which includes increasing its connectivity and reducing its packet loss rate. However, IPv6 paths experience smaller average RTTs than their IPv4 counterparts. We also notice that tunneled paths do not show a notable degraded performance compared with native paths.

Our effort to understand the IPv6 performance and operational issues is not completed. We are interested in the performance differences of native and tunnel links of IPv6 in more details. A systematic comparative study using tunnel discovery techniques is on-going. We also plan to follow its performance evolution in larger scope, and study its interaction with IPv4 Internet by measurement and modeling.

6. References

1. S. Bradner and A. Mankin, “IP: Next Generation (IPng) White Paper Solicitation,” *RFC 1550*, Dec. 1993.
2. L. Colitti, G. D. Battista, and M. Patrignani, “IPv6-in-IPv4 tunnel discovery: methods and experimental results,” *IEEE Transactions on Network and Service Management*, vol. 1, no.1, April. 2004.
3. K. Cho, M. Luckie, and B. Huffaker, “Identifying IPv6 Network Problems in the Dual-Stack World,” *SIGCOMM’04 Network Troubleshooting Workshop*, Portland, USA, 2004.
4. S. Deering and R. Hinden, “Internet Protocol, version 6 (IPv6) Specification,” *RFC 2460*, Dec. 1998.
5. B. Carpenter and K. Moore, “Connection of IPv6 Domains via IPv4 Clouds,” *RFC 3056*, Feb. 2001.
6. A. Durand, P. Fasano, I. Guardini, and D. Lento, “IPv6 Tunnel Broker,” *RFC 3053*, Jan. 2001.
7. R. Rockell and R. Fink, “6Bone Backbone Routing Guidelines,” *RFC2772*, Feb. 2000.
8. R. Gilligan and E. Nordmark, “Transition Mechanisms for IPv6 Hosts and Routers,” *RFC 2893*, Aug. 2000.
9. Ioan Raicu and Sherali Zeadally, “Evaluating IPv4 to IPv6 Transition Mechanisms,” *IEEE International Conference on Telecommunications 2003*, Feb. 2003.
10. J. Mogul and S. Deering, “Path MTU Discovery,” *RFC 1191*, Nov. 1990.
11. J. McCann, S. Deering and J. Mogul, “Path MTU Discovery for IP version 6,” *RFC 1981*, Aug. 1996.
12. S. McCreary and K. Claffy, “Trends in Wide Area IP Traffic Patterns: A View from Ames Internet Exchange,” <http://www.caida.org/outreach/papers/2000/AIX0005/>
13. Yue Li, Hui Liu, Gang Zhu, Shaozhi Ye, and Xing Li, “Analysis of IPv6 over Search Engine,” *The Fifth Joint AEARU Workshop on Web Technology and Computer Science*, Oct 2003.
14. GNU wget. <http://www.gnu.org/software/wget/wget.html>
15. TCPDUMP. <http://www.tcpdump.org/>
16. M.S. Borella, D. Swider, S. Uludag, and G.B. Brewster, “Internet Packet Loss: Measurement and Implications for End-to-End QoS,” *International Conference on Parallel Processing*, Aug. 1998.
17. M. Yajnik, S. Moon, J. Kurose, and D. Townsley, “Measurement and Modeling of the Temporal Dependence in Packet Loss,” *IEEE INFOCOM’99*, March 1999.
18. V. Paxson, “Measurements and Analysis of End-to-End Internet Dynamics,” *Ph.D. dissertation*, University of California at Berkeley, 1997.
19. H. Balakrishnan, V.N. Padmanablah, S. Seshan, M. Stemm, and R.H. Katz, “TCP Behavior of a Busy Internet Server: Analysis and Improvements,” *IEEE INFOCOM’98*, March 1998.