

Measuring the IPv4-IPv6 IVI Translation Overhead

Anthony K. Tsetse Alexander L. Wijesinha Ramesh K. Karne Alae Loukili
Towson University Towson University Towson University Towson University
Dept. of Comp. & Info. Sci. Dept. of Comp. & Info. Sci. Dept. of Comp. & Info. Sci. Dept. of Comp. & Info. Sci.
Towson, MD 21252, U.S.A. Towson, MD 21252, U.S.A. Towson, MD 21252, U.S.A. Towson, MD 21252, U.S.A.

atsetse@towson.edu awijesinha@towson.edu rkarne@towson.edu aloukili@towson.edu

ABSTRACT

While IPv6 deployment in the Internet continues to grow slowly at present, the imminent exhaustion of IPv4 addresses will encourage its increased use over the next several years. However, due to the predominance of IPv4 in the Internet, the transition to IPv6 is likely to take a long time. During the transition period, translation mechanisms will enable IPv6 hosts and IPv4 hosts to communicate with each other. For example, translation can be used when a server or application works with IPv4 but not with IPv6, and the effort or cost to modify the code is large. Stateless and stateful translation is the subject of several recent IETF RFCs. We measure the overhead due to the new IVI translator, which is viewed as a design for stateless translation, by measuring the internal timing for a freely available Linux implementation of IVI. To study the impact of operating system overhead on IVI translation, we also implement the IVI translator on a bare PC that runs applications without an operating system or kernel. Our results show that translating IPv4 packets into IPv6 packets is more expensive than the reverse, and that address mapping is the most expensive IVI operation. While both implementations of IVI have low overhead, a modest performance gain is obtained due to using a bare PC.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols – *Applications, Protocol architecture.*

C.2.6 [Computer-Communication Networks]: Internetworking – *Routers, Standards.*

C.4 [Performance of Systems]: *Measurement techniques.*

D.4.4 [Operating Systems]: Communications Management – *Network communication.*

General Terms

Measurement and Performance.

Keywords

IPv6; IVI Translation; IPv4-IPv6 transition; bare PC; Linux.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

RACS'12, October 23-26, 2012, San Antonio, TX, USA.
Copyright 2012 ACM 978-1-4503-1492-3/12/10...\$10.00.

1. INTRODUCTION

Internet Protocol version 6 (IPv6) [1] is the next generation IP protocol for the Internet. IPv6 was introduced to replace the existing Internet Protocol version 4 (IPv4). Although IPv6 provides the benefit of increased address space (16 byte addresses), its deployment in the Internet has been slow primarily due to the established and familiar base of IPv4. Interim measures such as NAT and CIDR have also served to conserve the limited IPv4 address space. However, it is expected that IPv6 usage will increase as the number of networks and devices (such as low power sensors) that require IP addresses continues to grow. IPv6 and IPv4 are not compatible due to using different header sizes and formats.

To enable IPv4 and IPv6 to co-exist during the transition period, several mechanisms have been proposed. Dual stack and tunneling are described in RFC 2893 [2], with an update in RFC 4213 [3] to reflect their usage in practice. Dual stack devices have implementations of both IPv4 and IPv6 protocol stacks running independently. This makes it possible for such devices to process both IPv4 and IPv6 packets. Tunneling allows IPv6 packets to be carried on IPv4 networks and vice versa. Address translation is a transition mechanism that allows communication between IPv6 devices and IPv4 devices by converting between packets of each type. Several RFCs dealing with stateless and stateful translation have been recently published by the IETF. The new IVI translator [4] is an implementation of stateless IPv4-IPv6 translation that can also support stateful packet translation (the abbreviation IVI is derived from the Roman numerals IV for 4 and VI for 6).

In this paper, we measure the overhead due to IVI translation by measuring internal timings for the existing Linux implementation. We also implement the IVI translator on a bare PC without an operating system or kernel. This enables us to compare the overhead due to only the IVI translation process with the overhead measured using Linux (Fedora) IVI translator.

The remainder of this paper is organized as follows. In Section 2, we briefly discuss related work. In Section 3, we provide an overview of IVI translation. In Sections 4 and 5, we describe the Linux and bare PC implementations of the IVI translator. In Section 6, we present the results of the performance study to determine IVI overhead. In Section 7, we give the conclusion.

2. RELATED WORK

Several techniques have been proposed for translating between IPv4 and IPv6 packets to IPv6 packets. One such technique, Network Address Translation-Protocol Translation (NAT-PT and

NAPT-PT) [5], was defined in RFC 2766. It allowed a set of IPv6 hosts to share a single IPv4 address for IPv6 packets destined for IPv4 hosts, and also allowed mapping of transport identifiers of the IPv6 hosts. RFC 2766 has since been recommended for historical status (in RFC 4966) [6] because of several issues. In RFC 3142 [7], the operation of IPv6-to-IPv4 transport relay translators (TRT) is discussed. TRT enables IPv6-only hosts to exchange two-way TCP or UDP traffic.

Recent work on IPv4-IPv6 translation is detailed in a group of several related RFCs. A mechanism enabling an IPv6 client to communicate with an IPv4 server (stateful NAT64) is described in RFC 6146 [8]. Translation of IP/ICMP headers in NAT64 is done based on the Stateless IP/ICMP Translation Algorithm (SIIT) of RFC 6145 [9]. NAT64 uses an address mapping algorithm discussed in RFC 6052 [10] for IPv4-IPv6 address mapping. RFC 6147 [11] specifies the DNS64 mechanism that can be used with NAT64 to enable IPv6 clients to communicate with IPv4 servers using the fully qualified domain name of the server. Translation as a tool for IPv6/IPv4 coexistence along with eight translation scenarios are discussed in RFC 6144 [12]. These RFCs provide the background for the IVI translator and translation approach given in [4], whose performance is the focus of this study. Details concerning an IVI implementation for stateless and stateful packet translation, and its deployment, are provided in [13].

3. OVERVIEW OF IVI TRANSLATION

IVI translation employs a prefix-specific and stateless address mapping scheme that enables IPv4 hosts to communicate with IPv6 hosts as described in RFC 6219 [4]. In this scheme, which is based on that given in [10], subsets of an ISP's IPv4 addresses are embedded in the ISP's IPv6 addresses. During translation, an address of one type (v4 or v6) is converted to the other type. To represent IPv4 addresses in IPv6, the ISP inserts its IPv4 address prefix after a unique IPv6 prefix consisting of PL bits, where PL=32 (4), 40 (5), 48 (6), 56 (7), 64 (8), or 96 (12) (the number in parenthesis is the number of bytes $p=PL/8$). The representation for $p<12$ is shown in Fig. 1, where $V4_p$ and $V4_s$ respectively denote the prefix and suffix of the IPv4 address and u is a byte of all 0s (the length of each field in bytes appears above it). For example, if $p=5$ bytes (i.e. PL=40 bits), the PREFIX, $V4_p$, u , $V4_s$, and SUFFIX fields consist of 5, 3, 1, 1, and 6 bytes respectively. If $p=12$, the PREFIX field is followed immediately by the entire IPv4 address (u is not present)..

p	8-p	1	p-4	11-p
PREFIX	$V4_p$	u	$V4_s$	SUFFIX

Figure 1. Representing IPv4 ISP addresses in IPv6.

Thus, an ISP with an IPv6 /32 address block can have a prefix of /40, where bits 32 through 39 are set to all ones. Similarly, /24 IPv4 address blocks are translated into /64 IPv6 addresses. The suffixes of these IPv6 addresses are normally set to all zeros. So with this mapping scheme [4], an ISP with a /32 IPv6 address block 2001:dba:: will translate the IPv4 address 202.38.97.205/24 into the equivalent IPv6 address 2001:dba:ffca:2661:cd::.

Translation of IPv4 and IPv6 headers is done as prescribed in [9] with the source and destination address translated according to the address mapping scheme described above. After translating the

transport layer (TCP or UDP) and ICMP headers, their checksums are recomputed to reflect the changes in the IP header. The IP header mapping scheme specified in [4] enables an IPv6 header to be constructed from an IPv4 header, and conversely, in the following manner (in addition to converting source and destination IPv4 or IPv6 addresses into their IVI mapped equivalent addresses). To translate an IPv4 header into a IPv6 header, the IHL, identification, flags, offset, header checksum, and options are discarded; version (0x4), TOS, protocol, and TTL fields are mapped into version (0x6), traffic class, next header, and hop limit respectively; and the value of total length is decremented by 20 and assigned to payload length. Conversely, to translate an IPv6 header into an IPv4 header, flow label is discarded; version (0x6), traffic class, next header, and hop limit fields are mapped into version (0x4), TOS, protocol, and TTL respectively; the value of payload length is incremented by 20 and assigned to total length; IHL is set to 5; and the remaining fields in the IPv4 header are assigned in the usual way (such as generating a value for the identification field and computing the IPv4 checksum). For convenience, we include Tables 1 and 2 taken from [4], which summarize the key elements of header translation from IPv4 to IPv6 and from IPv6 to IPv4 respectively that were discussed above.

Table 1. IPv4-to-IPv6 Header Translation (from [4])

IPv4 Field	IPv6 Equivalent
Version (0x4)	Version (0x6)
IHL	Discarded
Type of Service	Traffic Class
Total Length	Payload Length = Total Length - 20
Identification	Discarded
Flags	Discarded
Offset	Discarded
TTL	Hop Limit
Protocol	Next Header
Header Checksum	Discarded
Source Address	IVI address mapping
Destination Address	IVI address mapping
Options	Discarded

Table 2. IPv6-to-IPv4 Header Translation (from [4])

IPv6 Field	IPv4 Equivalent
Version (0x6)	Version (0x4)
Traffic Class	Type of Service
Flow Label	Discarded
Payload Length	Total Length = Payload Length + 20
Next Header	Protocol
Hop Limit	TTL
Source Address	IVI address mapping
Destination Address	IVI address mapping
	IHL=5
	Header Checksum Recalculated

4. IVI LINUX IMPLEMENTATION

The Linux implementation of the IVI translator [14] has two main components: a configuration utility and a translation utility. The

configuration utility is used to setup the routes using *mroute* or *mroute6* commands. The translation utility is responsible for translating the packets. To run the IVI translator on Linux, a patch has to be applied to the Linux kernel source. Once the patch is applied, the Linux kernel can be configured for IVI translation after recompiling the kernel. Patching the kernel results in some kernel source files being changed, and 4 other files (*mapping.c*, *proto_trans.c*, *mroute.c* and *mroute6.c*) being added for configuration and translation.

The functions required for translation referred to in the discussion below are all in the file *proto_trans.c*. When an IPv4 packet is received, the function *mapping_address4to6* is invoked to map the IPv4 address to an IPv6 address. Based on the value of the protocol field in the received IPv4 header, *icmp4to6_trans*, *tcp4to6_trans* or *udp4to6_trans* is invoked to translate the appropriate protocol. Each of these methods calls a function that computes the appropriate higher layer protocol (ICMP, TCP, or UDP) checksum as well. Once the protocol translation is done, *iphdr6to4_trans* is called to translate the IP header. For translating IPv6 to IPv4 packets, a similar process is used.

5. IVI BARE PC IMPLEMENTATION

The IVI translator application that runs on a bare PC with no operating system/kernel support was built by adding IVI translation capability on top of lean bare PC versions of the IPv6 and IPv4 protocols. The bare PC translator application included the implementation of two Ethernet objects ETH0 and ETH1 that directly communicate with its two network interface cards (NICs) each connecting to one of the two IP networks. A key difference between the Linux and bare PC implementations is that the IP protocols, address mappings, and protocol mappings in the latter are implemented as part of a single object called the XLATEOBJ. The architecture of the bare PC IVI translator is shown in Fig. 2.

The MAIN and RECEIVE (RCV) tasks are the only tasks running in the bare PC IVI translator application. The MAIN task runs when the system is started and whenever the RCV task is not running. The RCV task is invoked from the MAIN task when a packet (Ethernet frame) arrives on either interface. The capability to send out router advertisements on the IPv6 interface of the bare PC IVI Translator is implemented in the GW object as part of the MAIN task. In XLATEOBJ, the *xto4handler* and *xto6handler* methods respectively translate between IPv6 packets and IPv4 packets. To map the source and destination IPv6 addresses to the corresponding IPv4 addresses, the *xto4handler* invokes the *addr6to4* method.

The processing logic of the bare PC IVI translator is shown in Fig. 3. The translator initially checks the arriving packet to determine if it is an IPv6 or IPv4 packet. If the packet is an IPv6 packet, the payload protocol type is determined from the next header field, and the appropriate protocol checksum is computed. If the checksum is valid, the IPv6 packet is translated into an IPv4 packet and forwarded. If the packet is an IPv4 packet, the IP checksum is computed. If the checksum is valid, the packet is translated into an IPv6 packet based on the protocol type in the packet and forwarded after validating the higher layer protocol checksum as needed. If any checksum fails, the packet is dropped.

6. INTERNAL TIMINGS

6.1 Experimental Setup

Fig. 4 shows the test LAN used to generate TCP data. A similar set up was used to generate UDP and ICMP data. In the figure, S1, S2, S3 and S4 are 100 Mbps Ethernet switches, and R1 and R2 are IPv4 and IPv6 routers respectively. Router R1 serves the IPv4 routing domain, which consists of an IPv4 client C4 and an IPv4 server SV4. On the IPv6 side, C6 is an IPv6 client and SV6 is an IPv6 server. C6 and SV6 obtain the necessary routing information from the IPv6 router R2. The IVI translator is represented by XT. It runs on a Dell Optiplex GX270 desktop with an Intel Pentium IV 2.4 GHz processor, 512 MB RAM, and Intel PRO 10/100 and 3Com 10/100 NICs. The translator either runs with Fedora Linux (Fedora 12, Linux kernel 2.6.31) as the operating system, or as a bare PC application with no operating system/kernel. The client and server systems run on Dell Optiplex GX520s with a 2.6 GHz processor, 1 GB RAM, and an Intel PRO/1000 NIC, and Fedora Linux as the operating system. Apache HTTP Server 2.2.16 and Mozilla Firefox v3.6.7 were used respectively as the Web server and Web browser.

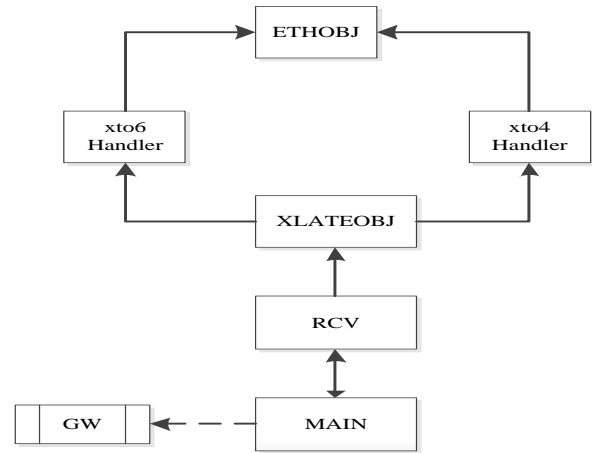


Figure 2. Bare PC IVI architecture.

To capture timings during translation with HTTP/TCP data, connections were made as follows. For IPv4 to IPv6 translation, the IPv4 client (C4) initiated a connection via the browser to the IPv6 Web server (SV6) in the IPv6 network. Router R1 forwarded the IPv4 request to the IVI translator XT, which translated the received packets to IPv6 packets, and sent them to the IPv6 Web server (SV6) via the IPv6 router R2. For IPv6 to IPv4 translation, the IPv6 client (C6) similarly initiated a connection via the browser to the IPv4 Web server (SV4). Router R2 forwarded this request to the IVI translator XT. The packet was translated to an IPv6 packet and forwarded to the IPv6 router R1, which delivered the packet to the IPv4 server (SV4).

6.2 Results

The internal timing data presented in this section was captured by inserting timing points in the Linux (Fedora) and the bare PC source code. The TCP traffic was generated by sending HTTP requests from clients to servers based on the network setup in Fig. 4 as described earlier. Using a similar setup, Mgen [15] was used to generate UDP traffic and the ping/ping6 utility was used to generate ICMP packets. Each experiment was run three times and

the average data is reported (we omit the deviation since the differences in the values for each experiment were small).

Fig. 5 shows the time it takes internally to translate packets using the Linux implementation of the IVI translator. It is seen that the IP header translation is the most expensive operation for both IPv6-IPv4 and IPv4-IPv6 translation, while UDP, TCP or

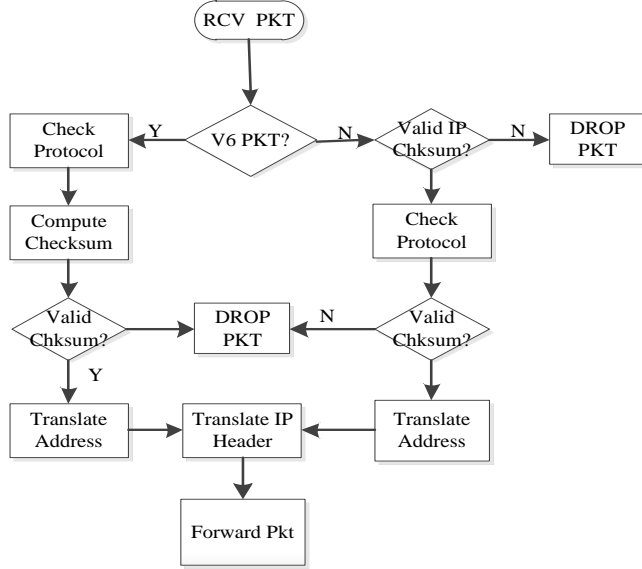


Figure 3. Bare PC IVI packet processing logic

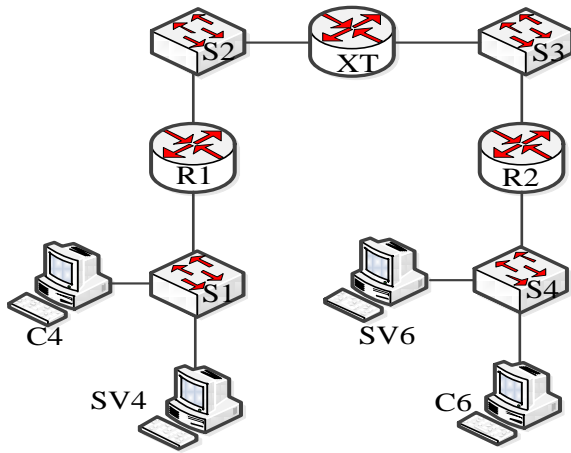


Figure 4. Test LAN for TCP connections.

ICMP translation has almost the same translation time. The header translation time is larger since it includes both address and protocol (TCP, UDP or ICMP) translation.

It is also evident that the IP address translation time is larger than the UDP, TCP, and ICMP header translation time since the higher layer translations involve very little processing. Comparing translations between the IP versions in each direction, the IP address translation time is 3.9 μ s and 0.5 μ s higher than TCP, UDP, and ICMP translation time for the IPv4 to IPv6 and IPv6 to IPv4 translations respectively. The larger time for the IPv4 to

IPv6 address translation compared to the reverse is because the IPv6 address needs to be constructed by extending the IPv4 address, whereas the IPv4 address is simply extracted from the IPv6 address.

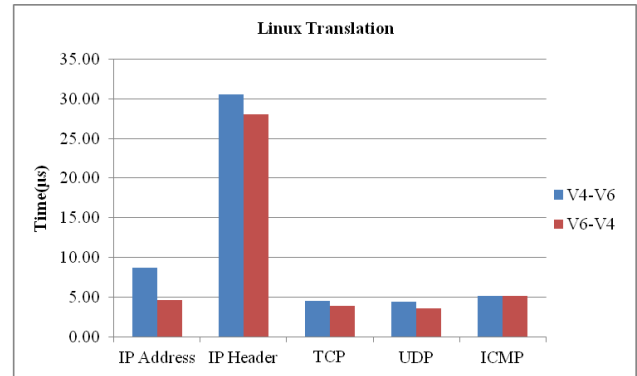


Figure 5. Linux translation overhead.

Fig. 6 shows the corresponding IVI translation times on a bare PC. The times for IP header translation from IPv4 to IPv6 is 15.8 μ s compared to 12 μ s for translation from IPv6 to IPv4. TCP and ICMP translation take almost the same time: about 2.5 μ s for IPv4 to IPv6 translation and 2.0 μ s for IPv6 to IPv4 translation. For IP address mapping, the processing time is 5 μ s from IPv4 to IPv6, and 2 μ s from IPv6 to IPv4. For UDP, translating from IPv4 to IPv6 and from IPv6 to IPv4 takes 2 μ s and 1.5 μ s respectively.

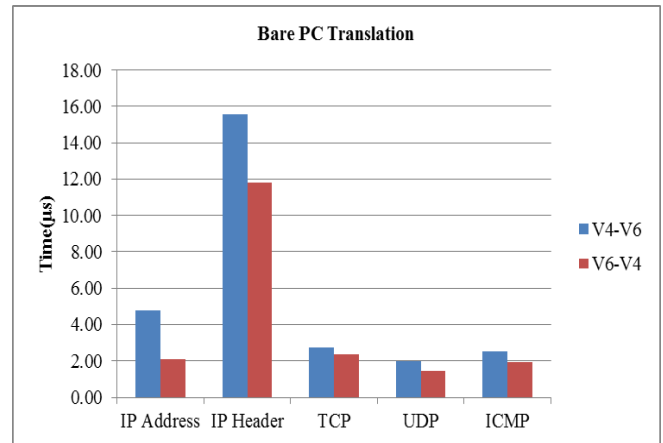


Figure 6. Bare PC translation overhead.

Comparing the times for IVI translation from IPv4 to IPv6 for the Linux and bare PC implementations (using Figs. 5 and 6), it can be seen that 1) for TCP, UDP and ICMP translations, the time for the bare PC is about 2.3 μ s less than for the Linux implementation; 2) for IP address translation, the processing time for the bare PC is about 3.9 μ s less than for the Linux implementation; and 3) for IP header translation, the processing time for the bare PC is 15 μ s less than for the Linux implementation.

Similarly, comparing the times for IVI translation from IPv6 to IPv4 translation for the Linux and bare PC implementations, it can also be seen that 4) for TCP and UDP traffic, the improvement in processing time due to using a bare PC is 2.2 μ s;

5) for ICMP, the improvement in processing time is 3.2 μ s; and 6) for IP address mapping, it is 2.6 μ s.

To compare packets per second (pps) processed by the IVI translator for IPv4 to IPv6 translation and IPv6 to IPv4 translation, the Mgen generator was used to make TCP connections between source and sink. The results when transferring 1024-byte packets when the throughput is at a maximum are shown in Fig. 7. As expected, pps for IPv4 to IPv6 translation is less than for IPv6 to IPv4 translation which reflects the decrease in processing time for the latter; the drop in pps for IPv4 to IPv6 translation compared to IPv6 to IPv4 translation is about 18% for Linux and about 13% for the bare PC. The improvement due to using a bare PC instead of Linux is about 17% for IPv4 to IPv6 translation and about 9% for IPv6 to IPv4 translation.

We also compared the round trip times for 1024 byte ICMP Ping packets. Request packets undergo IPv4 to IPv6 translation and the reply packets undergo IPv6 to IPv4 translation. The results are shown in Fig. 8 and indicate that round trip time is reduced by about 17% on a bare PC.

These results show that the overhead for IVI translation for both the Linux and bare PC implementations is small. However, there is a small improvement in processing time (and hence, packets processed per second) when using a bare PC.

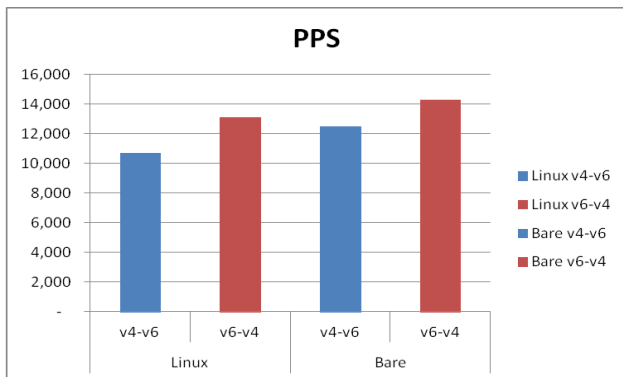


Figure 7. Packets per second (pps).

7. CONCLUSION

We determined the overhead due to IVI translation by measuring the internal timings for both IPv6 to IPv4 and IPv4 to IPv6 translation on an ordinary desktop running Linux Fedora. We also implemented an IVI translator as a bare PC application on the same machine with no operating system or kernel, and compared the corresponding timings. The results show that in general it is more expensive to translate packets from IPv4 to IPv6 than from IPv6 to IPv4 with a 4.1 μ s difference on Linux and a 3.8 μ s difference on a bare PC. It was also found that address mapping is the most expensive operation in IVI translation regardless of the system on which the translator is implemented. IVI translation has little overhead for both the Linux and bare PC implementations regardless of the higher layer protocol carried in the payload. However, eliminating the operating system overhead would

enable more efficient translation. The left-over CPU cycles could be used to process more packets, or for enhanced security-related processing during packet translation.

8. REFERENCES

- [1] Deering, S., and Hinden, R. *Internet Protocol, Version 6 (IPv6) Specification*. RFC 2460, Dec. 1998.
- [2] Gilligan, R., and Nordmark, E. *Transition Mechanisms for IPv6 Hosts and Routers*. RFC 2893, Aug. 2000.
- [3] Nordmark, E., and Gilligan, R. *Basic Transition Mechanisms for IPv6 Hosts and Routers*. RFC 4213, Oct. 2005.
- [4] Li X., et al. *The China Education and Research Network (CERNET) IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition*. RFC 6219, May 2011.
- [5] Tsirtsis, G., and Srisuresh, P. *Network Address Translation – Protocol Translation (NAT-PT)*. RFC 2766, Feb. 2000 .
- [6] Aoun, C., and Davies, E. *Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status*. RFC 4966, Jul. 2007.
- [7] Hagino, J., and Yamamoto, K. *An IPv6-to-IPv4 Transport Relay Translator*. RFC 3142, Jun. 2001.
- [8] Bagnulo, M., Matthews, P., and van Beijnum, I. *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*. RFC 6146, Apr. 2011.
- [9] Li, X., Bao, C., and Baker, F. *IP/ICMP Translation Algorithm*. RFC 6145, Mar. 2011.
- [10] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and Li, X. *IPv6 Addressing of IPv4/IPv6 translators*. RFC 6052, Oct. 2010.
- [11] Bagnulo, M., Sullivan, A., Mathews, P., and van Beijnum, I. *DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers*. RFC 6147, Apr. 2011.
- [12] Baker, F., Li, X., Bao, C., and Yin, K. *Framework for IPv4/IPv6 Translation*. RFC 6144, Apr. 2011.
- [13] Zhai, Y., Bao, C., and Li, X. *Transition from IPv4 to IPv6: A Translation Approach*. 2011 Sixth IEEE International Conference on Networking, Architecture and Storage. 30-39.
- [14] Source Code of the IVI implementation for Linux: <http://linux.ivi2.org/impl/>. accessed 05/10/2012.
- [15] Mgen, <http://cs.itd.nrl.navy.mil>, accessed 05/10/2012.

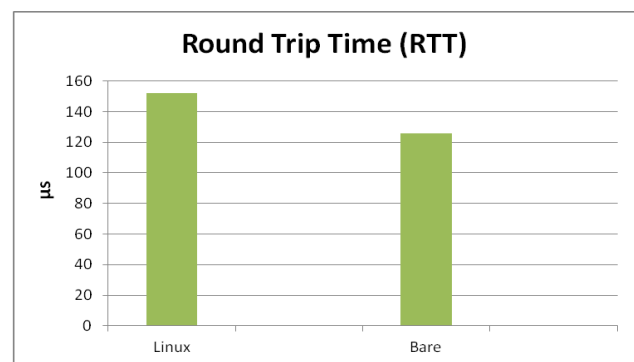


Figure 8. Ping round trip time (RTT).