

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: June 10, 2016

W. Chimiak
S. Patton
J. Brown
Laboratory for Telecommunication Sciences
J. Bezerra
Florida International University
H. Galiza
Rede Nacional de Ensino e Pesquisa (RNP) - NEG AmLight
J. Smith
University of Pennsylvania
December 08, 2015

IPv4 with 64 bit Address Space
draft-chimiak-enhanced-ipv4-02

Abstract

This document describes a solution to the Internet address depletion problem through use of a clever IPv4 options mechanism as a solution. This IPv4 protocol extension is called enhanced IP (EnIP). Because it is IPv4, it maximizes backward compatibility while increasing address space by a factor of 17.9 million. Unlike other similar proposals, care was taken to avoid costly changes and requirements to the core network and border routers, with the exception that options be passed in that equipment as described below. Because it is backward compatible, current IPv4 software, network equipment, firewalls, intrusion detection/protection, and layer 5 firewalls can be maintained until IPv6 system information security reaches acceptable maturity and availability.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 10, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Protocol Transitions	3
1.2.	EnIP's Use of IP Options	4
1.3.	Contents of this Paper	4
2.	Introduction to EnIP (EnIP)	5
2.1.	EnIP Addressing Example	5
2.2.	IPv4 Header with EnIP Option Header	5
2.2.1.	IPv4 Header Fields used for EnIP	6
3.	EnIP Operation	8
3.1.	IPv4 NAT Explanation using Figure 2	8
3.1.1.	Typical Private IPv4 Host to a Typical Public IPv4 Host using NAT	8
3.1.2.	Typical Private IPv4 Host to a Typical Private IPv4 Host using NAT	9
3.1.3.	Private EnIP Host to a Typical Public IPv4 Host NAT	9
3.1.4.	Private EnIP Host to a Typical Private IPv4 Host using NAT	10
3.2.	Enhanced IPv4 Explanation using an Example Figure 7	11
3.2.1.	EIP1 constructs the Packet	11
3.2.2.	EIP1 Transmits the Packet to N1	12
3.2.3.	Packet Arrives at N2 (192.0.2.2)	12
3.2.4.	EIP2 Receives the Packet	12
3.2.5.	EIP2 Transmits a Packet to EIP1	13
3.2.6.	Packet Arrives at N2	13
3.2.7.	Packet Arrives at N1	13
3.3.	Upgrading Servers to Support EnIP	14
3.4.	DNS Operation	14
3.5.	Information Security	14
4.	Tests and Comparisons of EnIP and Typical IPv4	15
5.	Conclusion	15

6.	IANA Considerations	15
7.	Informative References	16
	Authors' Addresses	17

[1.](#) Introduction

For various reasons, there is a large demand for IP addresses. It would be useful to have a unique address for each Internet device to allow, if desired, that any device could call another [[Lee-2012](#)]. The Internet of Things would also be able to make use of more routable addresses if they were available. Currently, this is not possible with IPv4. IPv6 has presented a potential solution to this problem but has faced problems with global adoption. Carrier Grade Network Address Translation and NAT (NAPT) have provided the solution thus far.

Enhanced IP (EnIP) was designed to minimize impact on core and border routers. DHCP, ARP, and existing IPv4 routing protocols are compatible with EnIP. By leveraging private address space in a similar manner of IP4+4 [[Turanyi-2002](#)], EnIP increases available address space by a factor of 17.9 million [[Chimiak-2014](#)] or 17.9 million new addresses for each current routable IP address.

This could allow the reassignment of small segments of unused address blocks in /8 networks to registries with chronic shortages of IP addresses.

EnIP packets carry additional address bits and state in an IP option, eliminating routing table updates like IPv6. EnIP supports end-to-end connectivity, a shortcoming of NAT, making it easier to implement mobile networks. Host renumbering is also not required in EnIP as has been the case with other 64-bit protocol proposals [[Turanyi-2002](#)]. This is a major topic of [section 4](#).

Because current NATs are resource intensive, requiring that state be maintained, this paper will call these NATs NATs, The stateless NAT used in EnIP will be called an EnIP NAT.

[1.1.](#) Protocol Transitions

Several transitions in the Internet Protocol suite are discussed in the IEEE paper [[Chimiak-2014](#)]. These transitions include Network Control Program (NCP) to TCP/IP, the evolution of IPv4, and IPv6. It also refers to network address port translation (NAPT), usually called NAT [[Paul-2014](#)], Nat444 [[Shirasaki-2012](#)], Dual-Stack Lite [[RFC6333](#)], NAT64 [[RFC6146](#)] and [RFC 6144](#) [[RFC6144](#)].

1.2. EnIP's Use of IP Options

EnIP uses IP options to increase address space. Over time, this fixed-length IP option would need to be added to the fast path implementations available on core routers [[Aweya-1999](#)]. Fast path implementations are discussed by Hidell et. al. [[Hidell-2014](#)]. Fast path allows core routers to optimize data paths to line cards based on table lookups. The fast path does not usually process packets containing IP options [[Aweya-1999](#)]. As a result, packets containing IP options are forwarded to the slow path CPU for processing and forwarding to the correct line card.

However, a very simple upgrade could be made to the core and border routers to process EnIP packets in the fast path with the following simple pseudo code [[Chimiak-2014](#)]:

```
if (IgnoreOptions)
    FastPathOperations();
else if (IP_Option_Present AND Option_Byte1 == 0X9A)
    FastPathOperations() ;
else
    SlowPathOperations() ;
```

EnIP and IPv6 both require upgrades to the fast path implementations of routers. EnIP's fast path upgrade has four advantages:

- (1) The EnIP upgrade is the pseudo code above, instead of a large code base insertion.
- (2) There is no equipment reconfiguration.
- (3) Internet providers can independently upgrade their fast paths.
- (4) Finally, there are no requirements to update the IPv4 routing tables.

1.3. Contents of this Paper

[Section 2](#) describes EnIP. [Section 3](#) is the heart of the paper, describing the mechanics of EnIP. [Section 4](#) describes the University of Maryland and University of Delaware deployment. It also describes tests between two nodes using EnIP and using normal IPv4 and gives the results of the tests. The final section contains some concluding remarks.

2. Introduction to EnIP (EnIP)

EnIP increases IP address space while maintaining compatibility with existing IPv4 implementations. Current EnIP implementations use IP option 26 to create a twelve byte extension to the IP header.

2.1. EnIP Addressing Example

The EnIP extension contains four bytes of overhead, and two four byte fields used as additional storage for the EnIP source and destination addresses. EnIP addresses are written as two IPv4 addresses concatenated together. IPv6 and EnIP addressing schemes are shown below:

IPv6 address	2001:0101:c000:0202:0a01:0102::0
EnIP address	192.0.2.2.10.1.1.2

In the above EnIP address, 192.0.2.2 is a public IPv4 address called the site address; the 10.1.1.2 address is called the host address. The host address is always a private IPv4 address assigned to a host behind a lightweight, stateless EnIP-enabled NAT. It is the private IPv4 address that identifies the host in the network behind that NAT and is compatible with other normal IPv4 hosts. The site address allows the packet to traverse public networks because its IPv4 Source Host Number (or source address) is 192.0.2.2, a fully routable IPv4 address.

2.2. IPv4 Header with EnIP Option Header

The IPv4 header supporting EnIP is shown in Figure 1. It is an IPv4 header with additional option space used as follows:

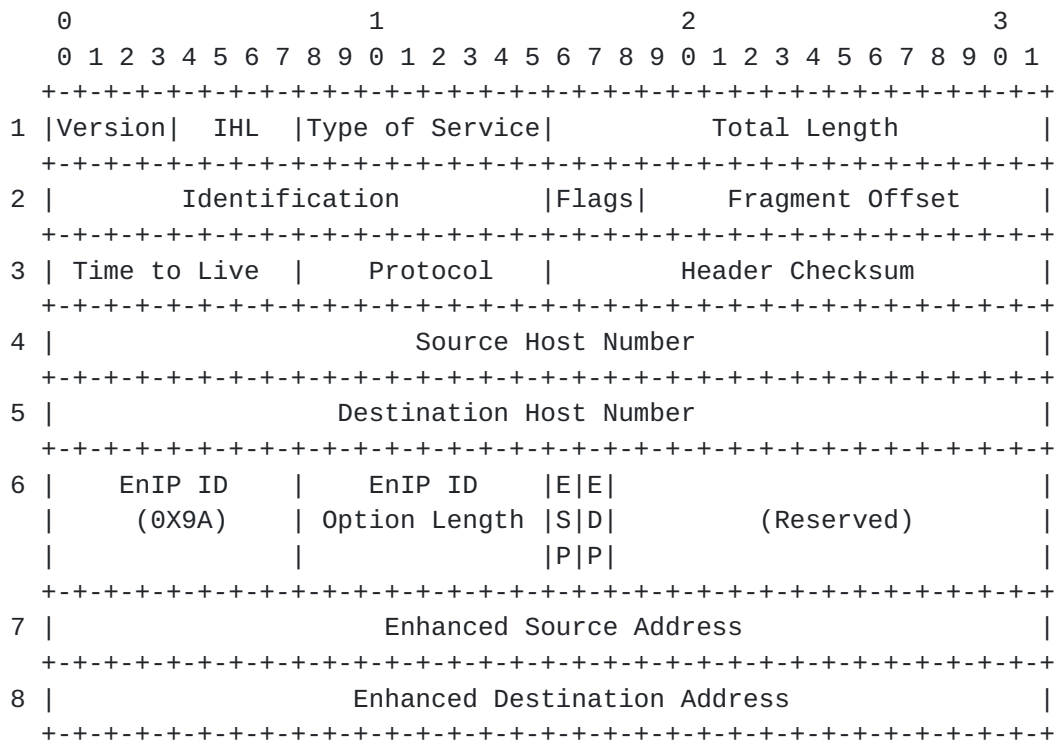


Figure 1: EnIP Header

2.2.1. IPv4 Header Fields used for EnIP

The fields

Field	Bits	Description
Version	4	The Version field is identical to IPv4's
IHL	4	Internet Header Length is identical to IPv4's and is set to the length of the EnIP header
Type of Service	8	The ToS field is identical to IPv4's
Total Length	16	The Total Length field is identical to that of IPv4 and is 32 octets
Identification	16	The Identification field is identical to IPv4's
Flags	3	The Flags field is identical to IPv4's
Fragment Offset	13	The Fragment Offset field is identical to IPv4's
Time to Live	8	The Time to Live field is

Protocol	8	identical to IPv4's The Protocol field is identical to IPv4's
Header Checksum	16	The Header Checksum field is to IPv4's

EnIP specific		

Source Host Number	32	The Source Host Number field identifies the source host. The EnIP use is described in section 3
Destination Host Number	32	The Destination Host Number field identifies the destination host. The EnIP is described in section 3
EnIP ID	8	The EnIP ID field equals to 0x9A or 1 00 11010. It's meaning is given at the end of this section
EnIP Option Length	8	The EnIP Option Length field is 12 octets
ESP	1	This selector determines if an Enhanced Source Address is present
EDP	1	This selector determines if an EnIP Destination Address is present
Reserved	14	Reserved for future use
Enhanced Source Address	32	This is the host address used by EnIP. It is described in section 3
Enhanced Destination Address	32	This is the destination host address used by EnIP. It is described in section 3

NOTE: The protocol has 12 bytes of overhead per packet.

The meaning of the EnIP ID field,

1 00 11010 (0x9A hexadecimal)

is as follows:

If an EnIP packet traverses a router and must be fragmented because of a link with a smaller MTU, the copy bit ensures that fragments

include the 12-byte IP option header in all fragmented packets. The rest of the bits and their meaning are as follows.

```

+-----+
|           Meaning of 1 00 11010 the EnIP ID           |
+-----+-----+-----+-----+
|      1      |      00      | 11010 (or 26 base 10) |
+-----+-----+-----+-----+
|copy bit set | Class is Control | Option is a new value |
+-----+-----+-----+-----+

```

EnIP Id Field Interpretation

Note that 26 is the IP option value used in the EnIP experiments.

3. EnIP Operation

3.1. IPv4 NAT Explanation using Figure 2

```

+-----+                                     +-----+
| IP1 |                                     | IP2 |
+-----+                                     +-----+
10.1.1.1                                     10.3.3.1

          192.0.2.1          192.0.2.2
          +-----+          +-----+
+-----+          | N1 |          | N2 |          +-----+
| EIP1 |          +-----+          +-----+          | EIP2 |
+-----+          10.1.1.254          10.3.3.254          +-----+
10.1.1.2                                     10.3.3.2

```

EIP1 and EIP2: machines supporting IPv4
and Enhanced IPv4
IP1 and IP2: machines supporting only IPv4
N1 and N2: NAT and EnIPNAT/translator devices

Figure 2 - Enhanced IP and Unenhanced IPv4

This section shows how two nodes, without EnIP, can communicate. More complete explanations are available in [[Chimiak-2014](#)].

3.1.1. Typical Private IPv4 Host to a Typical Public IPv4 Host using NAT

IP1 (10.1.1.1), a typical IPv4 host, wants to reach N2 (192.0.2.2) (See Figure 3). It uses NAT (as on Linux iptables) to masquerade as the public IP address of N1 (192.0.2.1). N1 sets up a port forwarding rule for IP1 for return packets from N2.




Figure 3 - Typical Private IPv4 Host to a Typical Public IPv4 Host NAT

3.1.2. Typical Private IPv4 Host to a Typical Private IPv4 Host using NAT

The same IP1 (10.1.1.1) wants to reach tcp port 80 on IP2 (10.3.3.1), another typical IPv4 host (see Figure 4). The packets from 10.1.1.1 use NAT on N1 to masquerade as source IP address 192.0.2.1. At N2 (192.0.2.2), the packets arrive. There they have a NAT port forwarding rule setup on N2 to map tcp port 80 on 192.0.2.2 to forward packets to the internal host IP2 (10.3.3.1).

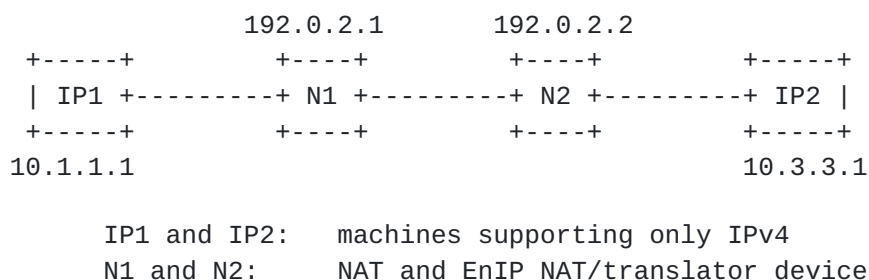


Figure 4 - Typical Private IPv4 Host to a Typical Private IPv4 Host using NAT

So packets move from IP1 to N1, taking N1's address as the source address and move to N2. N2 has a mapping to IP2 and sends the packet to IP2. When IP2 sends replies, the NATs used their table entries to send the packet back to IP1, adjusting the addresses as necessary.

3.1.3. Private EnIP Host to a Typical Public IPv4 Host NAT

Suppose EIP1 (10.1.1.2) wants to reach N2 (192.0.2.2) as in Figure 5. It is the same as an IPv4 node wanting to reach N2. So the destination IP address EIP1 used to communicate with N2 is 192.0.2.2. The NAT device (N1) uses IPv4 NAT to translate the source of the packets to come from 192.0.2.1. EIP1 behaves as though it is an IPv4 host. It is fully compatible with normal IPv4 communications.

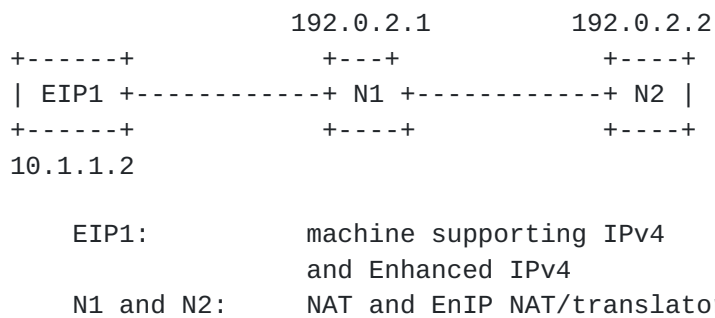


Figure 5 - Private EnIP Host to a Typical Public IPv4 Host using NAT

3.1.4. Private EnIP Host to a Typical Private IPv4 Host using NAT

EIP1 (10.1.1.2) initiates a tcp port 80 connection to IP2 (10.3.3.1) (see Figure 6). Packets from EIP1 reach N1. As above, they are masqueraded as IPv4 address 192.0.2.1. When the packets arrive at N2 (192.0.2.2) from N1 (192.0.2.1), there is a port forwarding entry for the port 80 setup to send the packets to IP2 (10.3.3.1).

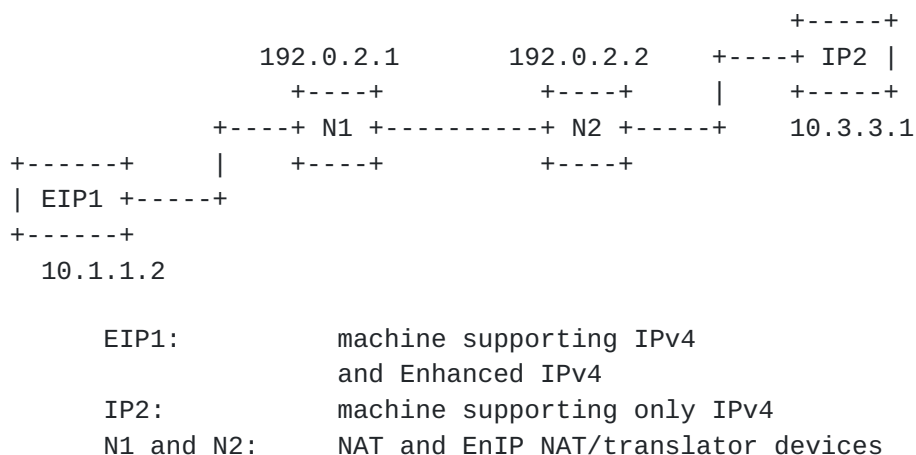



Figure 6 - Private EnIP Host to a Typical Private IPv4 Host using NAT

In the reverse direction, IP2 looks at the return address. It is the address of N1 (192.0.2.1). It sends the packet to N2. N2 sees the connection in its state table and realizes it is the N1 to IP2 connection. It strips IP2's 10.3.3.1 and places its 192.0.2.2 address in the source field and places the appropriate port value and sends the packet to N1. N1 sees that this is the EIP1 to IP2 connection in its state table. It uses the state table entry to rewrite the proper port translation, sets the destination address as 10.1.1.2 and sends the packet back to EIP1 as a normal IPv4 packet.

It is important to note that thus far this paper
has not demonstrated any usage of EnIP features,
just compatibility with current IPv4 implementations

3.2. Enhanced IPv4 Explanation using an Example Figure 7

Suppose EIP1 wants to send packets to EIP2. Here both hosts are running Enhanced IPv4 stacks and N1 and N2 support EnIP. Suppose EIP1 knows the address of EIP2 is 192.0.2.2.10.3.3.2. EIP1 knows its internal IP address of 10.1.1.2 but is not aware of the external address of N1 (192.0.2.1). The following is done (see Figure 7):



EIP1 and EIP2: machines supporting IPv4
and Enhanced IPv4

N1 and N2: NAT and EnIP NAT/translator devices

Figure 7 - Enhanced IP

3.2.1. EIP1 constructs the Packet

- (1) EIP1 sets the Source Host Number field (the source IPv4 address) to 10.1.1.2;
- (2) The EnIP ID field is set to 0X9A. the ESP bit is zeroed in the EnIP header.
- (3) The Enhanced Source Address in the EnIP header is set to all ones since an EnIP source address is not currently present.
- (4) The most significant 32 bits of the EnIP address is set by storing 192.0.2.2 in the IPv4 Destination Host Number.
- (5) The least significant 32 bits of the EnIP address is set by storing 10.3.3.2 in the EnIP Destination Address field.
- (6) Finally, EIP1 sets the EDP bit to 1.

3.2.2. EIP1 Transmits the Packet to N1

The packet is routed to N1 using normal IPv4. N1 does the following:

- (1) It notes the EnIP option (0X9A) is present.
- (2) N1 reads 10.1.1.2 from the IPv4 Source Host Number field and places that in the Enhanced Source Address field. This field no longer contains 255.255.255.255.
- (3) N1 sets the ESP bit to 1
- (4) N1 makes 192.0.2.1 the IPv4 Source Host Number.
- (5) N1 recomputes the IP checksum of the new packet. If the packet carries TCP or UDP, it recomputes these checksums as they have also changed.

3.2.3. Packet Arrives at N2 (192.0.2.2)

When the packet Arrives at N2, N2 does the following:

- (1) N2 identifies the packet as an EnIP packet (0X9A).
- (2) It reads the Enhanced Destination Address, 10.3.3.2, and places this value as the IP header's Destination Host Number.
- (3) N2 zeroes the EDP bit and the Enhanced Destination Address to zero.
- (4) It recomputes the IP checksum. If the packet carries TCP or UDP, recomputes these checksums as they have changed as a result of a change to the IP destination address.
- (5) N2 sends the packet to EIP2.

3.2.4. EIP2 Receives the Packet

When the packet arrives at EIP2, EIP2 does the following:

- (1) It computes the source address of the packet. The IPv4 Source Host Number (192.0.2.1) is concatenated with the Enhanced source address (10.1.1.2) The result is 192.0.2.1.10.1.1.2.
- (2) The IPv4 Destination Host Number (address) is 10.3.3.2.

3.2.5. EIP2 Transmits a Packet to EIP1

To construct a packet from EIP2 to EIP1, EIP2 does the following:

- (1) EIP2 sets the Option ID field to 0X9A.
- (2) It takes the IPv4 Source Host Number from the incoming packet, 192.0.2.1, and sets it as the IPv4 Destination Host Number.
- (3) It places the Enhanced Source Address (10.1.1.2) in the Enhanced Destination Address field.
- (4) EIP2 sets the EDP field to 1 and the IPv4 Source Host Number to 10.3.3.2.
- (5) It sets the Enhanced Source Address to all ones.
- (6) It sets the ESP bit to 0.

3.2.6. Packet Arrives at N2

When the packet Arrives at N2, N2 does the following:

- (1) It writes 10.3.3.2 in the Enhanced Source Address field.
- (2) The ESP bit is set to 1.
- (3) N2 writes 192.0.2.2 in the IPv4 Source Host Number field and recomputes the IP checksum. If the packet carries TCP or UDP, it recomputes these checksums as well.
- (4) N2 sends the packet to N1 (192.0.2.1).

3.2.7. Packet Arrives at N1

When the packet arrives at N1, it does the following:

- (1) N1 reads 10.1.1.2 from the Enhanced Destination Address.
- (2) It writes this value in the IPv4 Destination Host Number field.
- (3) It zeroes the EDP bit and the Enhanced Destination Address field.
- (4) N1 recomputes the IP checksum and if the protocol is TCP or UDP, recomputes these checksums as well.

(5) N1 sends the packet to EIP1.

3.3. Upgrading Servers to Support EnIP

In order to maintain backward compatibility with old IPv4 systems, it is important that the EnIP Source Address is an allowed private address. Otherwise, that address becomes a routable address outside of an autonomous system and there is a potential for routing ambiguity.

With a kernel modification to support EnIP, an existing server deployed using an IPv4 address can be upgraded. The example of a TCP echo server [[RFC862](#)] is shown in detail elsewhere [[Chimiak-2014](#)]. However, highlights are now presented here.

An echo server logs the source IP address of each data packet with the `getpeername` function. However, the length of the IPv4 address structure returned is currently 16 bytes for IPv4 (the size of a `struct sockaddr_in`). For IPv6 the size of `struct sockaddr_in6` is 28 bytes. As an example of the Alpha implementation mentioned previously, EnIP returns a new structure called `struct sockaddr_ein` that is 26 bytes. It is necessary to use the length 26 to detect a `struct sockaddr_ein`. This new struct has two values: `sin_addr1` and `sin_addr2`. These represent the IPv4 source address and the EnIP Source Address. An implementation of `getpeername` could provide a compatibility mode to treat EnIP addresses as IPv4 addresses. Once the echo client software is upgraded, the `getpeername` implementation could return `struct sockaddr_ein`.

3.4. DNS Operation

[RFC 2928](#) provides 2001:0101 as the experimental IPv6 prefix[RFC2928]. EnIP lookups can use already standardized AAAA records with this prefix. This prefix uses 32 bits of the 128 bit AAAA record. EnIP uses only 64 bits of the remaining 96 bits. If 192.0.2.2.10.1.1.2 is the EnIP address, the EnIP AAAA record is 2001:0101:c000:0202:0a01:0102::0. We call this an AA record. A new record type is not needed and eliminates the need for DNS server software upgrades on a massive scale.

3.5. Information Security

The same article shows the care taken to minimize changes in IDS/Firewalls Operation. It also addresses the prevention of EnIP NAT or hosts from forwarding illegal packets, and its operation with most of IPSEC, except in the full tunnel mode AH+ESP scenario.

4. Tests and Comparisons of EnIP and Typical IPv4

The IEEE Computer Magazine article describes the two-node EnIP test deployment over the Internet2 network with the EnIP NATs. The tests, along with their results, are given that compare typical IPv4 connections with EnIP connections [[Chimiak-2014](#)]. EnIP was used across 7 routers on Internet2 to demonstrate viability in environments that need more addressing.

In brief, the test had a node at the University of Maryland running the Linux 2.6.38 kernel with patches to include the EnIP alpha code. Another similar node was set up at the University of Delaware. Each had an EnIP NAT. http, samba, and ssh were used without modification with Enhanced IP DNS calls used.

The results, in [section 5.3](#) of the paper[Chimiak-2014], show EnIP performance gains over traditional NAT, as there are no hash table lookups, as in NAT. It provided roughly a factor of 2 improvement.

EnIP has also had basic testing with a 4G Long Term Evolution (LTE) simulator along with an Android phone. However, these results have not been compiled.

5. Conclusion

This RFC describes the operation of IPv4 using IP options, called Enhanced IP or EnIP. EnIP provides a solution to IPv4 address scarcity. Because the EnIP design criteria is to extend IPv4 instead of replacing it, it reduces impact on routers and information security mechanisms already in place. The operation of two EnIP nodes at the University of Maryland and the University of Delaware, running http, samba, and ssh without modification of the software or routers in the path demonstrates the feasibility of EnIP on a small but realistic scale that spanned seven routers.

Tests were conducted and documented, that show expected performance improvement over the old NAT currently used, and the new EnIP NAT. It provided roughly a factor of 2 improvement [[Chimiak-2014](#)].

6. IANA Considerations

This document does not create a new registry nor does it register any values in existing registries; no IANA action is required.

7. Informative References

[Chimiak-2014]

Chimiak, W., Patton, S., and S. Janansky, "Enhanced IP: IPv4 with 64-Bit Addresses", IEEE Computer Magazine Vol. 47, Issue 2, pp 62-69, February 2014.

[Lee-2012]

Lee, G., "The Internet of Things - Concept and Problem Statement", Internet Draft (work in progress) [draft-lee-iot-problem-statement-05](#), July 2012.

[Turanyi-2002]

Turanyi, Z. and A. Valko, "IPv4+4", Proceedings of the 10th IEEE International Conference on Network Protocols, 2002.

[Paul-2014]

Paul, T. and F. Tsuchiya, "Extending the ip Internet through address reuse", ACM SIGCOMM Computer Communication Review vol. 23, Issue 1, pp. 16-33, January 1993.

[Shirasaki-2012]

Yamaguchi, J., Shirasaki, Y., Miyakawa, S., Nakagawa, A., and H. Ashida, "Enhanced IP: IPv4 with 64-Bit Addresses", Internet Draft (work in progress) [draft-shirasaki-nat444-isp-shared-addr-08](#), July 2012.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", Proposed Standard [RFC 6333](#), August 2011.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", Proposed Standard [RFC 6146](#), Apr 2011.

[RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation.", Internet Draft (Informational) [RFC 6144](#), April 2011.

[Aweya-1999]

Aweya, J., "IP router architecture: An overview", Technical Report University of Virginia, 1999.

[Hidell-2014]

Hidell, M., Sjodin, P., and O. Hagsand, "Router architectures: Tutorial at Networking 2004", Networking 2004 Athens, Greece, May 2004.

[RFC862] Postel, J., "Echo Protocol", May 1983.

[RFC2928] Postel, J., "Initial IPv6 Sub-TLA ID Assignments", September 2000.

Authors' Addresses

William Chimiak
Laboratory for Telecommunication Sciences
8080 Greenmead Drive
College Park, MD 20740
US

Phone: +1 301 422 5217
EMail: w.chimiak@ieee.org

Samuel Patton
Laboratory for Telecommunication Sciences
8080 Greenmead Drive
College Park, MD 20740
US

Phone: +1 301 422 5217
EMail: sam@enhancedip.org

James F. Brown
Laboratory for Telecommunication Sciences
8080 Greenmead Drive
College Park, MD 20740
US

Phone: +1 301 422 5217
EMail: brown@ltsnet.net

Jeronimo A Bezerra
Florida International University
11200 S.W. 8th St PC building Suite 312
Miami, FL 33199
US

Phone: +1 786 616 3081
EMail: jbezerra@fiu.edu

Humberto Silva Galiza de Freitas
Rede Nacional de Ensino e Pesquisa (RNP) - NEG AmLight
Av. Dr. Andre Tosello, 209
Cidade Universitaria Zeferino Vaz Campinas, SP 13083-886
Brazil

Phone: +55 19 3787-3300
EMail: galiza@amlight.net

Jonathan Smith
University of Pennsylvania
Levine Hall 3330 Walnut Street
Philadelphia, PA 19104
US

Phone: 215.898.9509
EMail: jms@cis.upenn.edu

