Softwire Working Group
Internet-Draft
Intended status: Informational
Expires: January 31, 2011

Y. Lee
Comcast
P. Kapoor
Xavient
July 30, 2010

UDP Encapsulation of 6rd
draft-lee-softwire-6rd-udp-02

Abstract

This memo specifies the UDP encapsulation to IPv6 Rapid Deployment (6rd) protocol which enables hosts behind unmodified Home Gateway device to access 6rd service.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

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 January 31, 2011.

Copyright Notice

Table of Contents

1. Introduction

6rd protocol [I-D.ietf-softwire-ipv6-6rd] enables service providers to rapidly deploy IPv6 over IPv4 network. In [RFC5569], it describes the 6rd architecture to enable a service provider to deploy IPv6 connectivity on an IPv4 network for 1.5 million residential customers. The architecture involves two major components: (1) the 6rd relays (6rd Border Router) in the ISP network and (2) 6rd CPE (6rd CE) in the customer premise.

The 6rd CE in the customer home is a NAT [RFC3022] device which shares a single external IPv4 address to multiple internal IPv4 hosts. Note that the external IPv4 address can be either public or private. If private address is used, the ISP will provide NAT function in the provider network. Details is described in [RFC5569] Section 4. The 6rd CE is also an IPv6 router which advertises an IPv6 prefix (6rd Delegated Prefix) to the internal IPv6 hosts. The 6rd Delegated Prefix consists of two part: the 6rd Prefix and the external IPv4 address. The 6rd CE learns the 6rd Prefix, IPv4MaskLen, and 6rdPrefixLen from DHCP [I-D.ietf-softwire-ipv6-6rd] and calculate the 6rd Delegated Prefix. Then, the 6rd CE advertises the 6rd Delegated Prefix to the customer's home network via Router Advertisement [RFC4861] or DHCPv6 [RFC3315].

The 6rd specification fits well to those ISPs who manage the customers' home gateway (HGW). When the ISP is ready to deploy the 6rd, a new firmware which contains the 6rd CE implementation will be pushed to the managed HGW. For the ISPs who do not manage the customers' HWG, they cannot upgrade the customer's HWGs to support 6rd. This memo specific a UDP encapsulation to encapsulate 6rd over UDP which enables ISP to deploy 6rd to hosts behind unmodified HWG. There are two scenarios in which the ISP can use this specification.

1. First deployment scenario is the O/S in the host implements this specification. The host is connect to the unmodified HGW's LAN interface. One the host establishes IPv4 connectivity, it initiates the 6rd discovery procedure Section 5 and construct the 6rd Delegated Prefix. When the host sends an IPv6 datagram, it encapsulates the IPv4 datagram with a standard UDP header [RFC0768], then it encapsulates the IPv6 datagram in an IPv4 header following the procedure defined in [RFC5569]. Details of the UDP encapsulation is described in Section 4.

2. Second deployment scenario is a dedicated server implements this specification. The server is connected to the unmodified HGW's LAN interface. Once the server establishes IPv4 connectivity, it initiates the 6rd discovery procedure Section 5 and construct the 6rd Delegated Prefix. Then, it advertises the 6rd Prefix via

Router Advertisement or DHCPv6 to the HGW LAN so that IPv6 datagram will use the server as a gateway to establish IPv6 sessions. This enables the hosts connecting to the HGW's LAN to access 6rd service without additional configuration.

This specification defines the udp encapsulation that allows to deploy 6rd behind a 6rd unaware NAT-ed gateway. The general mechanism is still stateless and requires little change to the IPv4 networking.

2. Terminology

This documents users terms defined in [I-D.ietf-softwire-ipv6-6rd]. In addition, we defines the following new terms:

- HGW: is referred to the edge device installed in customer home. It typically provides NAT function to the home equipments. The HGW in this context is unaware of 6rd.
- External IPv4 address: is referred to the external IPv4 address assigned by the ISP to the HGW. This address can be either a public IPv4 address of a private [RFC1918] address.
- Internal IPv4 address: is referred to the internal IPv4 address assigned by the HGW to the 6rd host. This address is a [RFC1918] address.
- 6rd UDP Host: is the host implemented this specification.
- 6rd UDP Delegated Prefix: is referred to the 6rd Delegated Prefix [I-D.ietf-softwire-ipv6-6rd] generated by the 6rd BR to identify the host implemented this specification. The 6rd UDP Delegated Prefix is different from the 6rd Delegated Prefix in which the 16-bit udp source port information is part of the 6rd delegated prefix calculation.
- 6rd UDP Host Delegated Prefix: is referred to the 6rd Delegated Prefix used by the host. This is different from the 6rd UDP Delegated Prefix in which the 16-bit udp source port information in the IPv6 address will be replaced by zeros.

3. Overview

3.1. 6rd UDP Host

Before the host can use 6rd, it needs to discover five pieces of information: the 6rd prefix, the external IPv4 address, the IPv4MaskLen, the 6rdPreflixLen, and the 6rd BR IPv4 address. Section 5 discusses the process to discover the information. The host calculates the 6rd UDP Host Delegated Prefix following the procedure described in Section 3.5.

The 6rd udp host implemented this specification must be able to encapsulate and de-capsulate udp and IPv4 headers when sending and receiving IPv6 datagrams. It must also allocate an available udp port during startup. This port must be used in the 6rd encapsulation during the life time of the process.

When the host wants to initiate an IPv6 session to an outside IPv6 host, the first encapsulates is to append the IPv6 datagram with an udp header. The udp source port is the udp port allocated at startup; the destination port is an IANA defined port. The second encapsulate is to append the IPv4 header. The source address will be the internal IPv4 address assigned by the HGW; the destination address will be the 6rd BR's IPv4 address.

Since 6rd udp host is behind the HGW, it must send a keepalive message to the 6rd BR to maintain the udp NAT binding alive. The keepalive is a simple udp packet which has special IPv6 destination address so that the 6rd BR will rsecognize this is a keepalive packet. When the 6rd receives this special udp datagram, it must discard the datagram and sends a response to the 6rd udp host. In the response, the IPv6 address contains only the 6rd udp delegated prefix and zero-pad the rest of the bits. When the 6rd udp host receives this datagram. It must discard it. The 6rd udp host must send the keepalive frequent enough to keep the binding.

3.2. Home Gateway

The HGW in this specification is a typical HGW found in retailed stores. In the WAN side, it connects to the ISP and runs DHCP client. The ISP offers an IPv4 address, a default gateway, and list of DNS servers to the HGW. In the LAN side, it runs DHCP server and offers [RFC1918] address to the hosts on the LAN. The HGW provides standard NAT [RFC3022] functions to allow multiple hosts to share a single external IPv4 address. When a host connects to the HGW, the host requests the Internal IPv4 address and the Internal IPv4 Gateway via DHCP. The HGW is not aware of the 6rd service.

3.3. 6rd Border Router

The 6rd BR implemented this specification must be configured to receive UDP packet on port XXXX (TBD) in its IPv4 interface facing the 6rd udp hosts. When it receives a udp packet, it de-capsulates the udp header and extract the udp source port information. Then it replaces the 16-bit zero-padded bit to the udp source port information and forms the 6rd delegated prefix. The 6rd BR must be pre-configured with the IPv4MaxLen and 6rdPrefixLen.

When the 6rd BR receives an IPv6 packet in its IPv6 interface, it must extract the 16-bit udp source port from the IPv6 destination address and use it to construct the udp header for encapsulation. The 6rd BR must also zero-pad the 16-bit udp source port information in the IPv6 destination address before sending to the 6rd udp host.

The 6rd BR performs the stateless header replacement function to embed the NAT-ed udp port information into the 6rd prefix.

3.4. 6rd UDP Delegated Prefix

In order to keep 6rd BR operation stateless and to make 6rd implementations co-exist with the HGW NAT at the same time, this specification defines a new 6rd delegated prefix [I-D.ietf-softwire-ipv6-6rd] calculation procedure which is slightly different from the classic 6rd delegated prefix calculation. This specification proposes to include the 16-bit UDP source port information in the 6rd delegated prefix.

The 6rd udp delegated prefix is calculated by concatenating the 6rd prefix, a consecutive set of bits from the CE IPv4 address, and the 16-bit udp source port. The sum of the number of bits must be less than or equal to 64.

For example, the 6rd prefix is 2001:DB8::/32; the IPv4 address is 192.0.1.100/24; the IPv4MaskLen is 16; the 6rdPrefixLen is 32; and the source UDP is 12345. The 6rd udp delegated prefix is 2001:DB8:0164:3039::/64

The 6rd udp delegated prefix is generated and used by 6rd BR.

3.5. 6rd UDP Host Delegated Prefix

Since the 6rd Host calculating the delegrated prefix is behind the NAT, it does not know the NAT-ed udp source port. When the 6rd constructs the 6rd UDP Delegated Prefix, it will not populate the 16-bit UDP source port information. Instead, it will zero-pad the 16-bit field. This will also ensure that the 6rd udp host delegated

prefix will not change if the old NAT-ed binding in the HGW expires and a new binding is formed.

Following the previous example, the 6rd host udp delegated prefix behind the HGW is 2001:DB8:0164::/64

The 6rd udp host delegated prefix is used by the 6rd udp host.

3.6. Procedures

3.6.1. Outbound Flow from the 6rd UDP Host

3.6.1.1. 6rd UDP Host

When the 6rd udp operation on the host starts, it must allocate an available udp port. This port is used in the source udp port of the udp encapsulation.

When the 6rd host wants to send an IPv6 datagram to an IPv6 destination, it puts the 6rd udp host delegated prefix in the source IPv6 address field. Then the host encapsulates the IPv6 datagram with a udp header. The source port is the udp port allocated during the startup process and the destination port is the well-known port assigned by IANA. Then, the host encapsulates the udp datagram with an IPv4 header. The IPv4 header contains the 6rd BR in the destination address field and the internal IPv4 address in the source address field. The datagram is passed to the relevant routing mechanism.

3.6.1.2. Home Gateway

When the HGW receives the first udp datagram, it NATs the source port and source IPv4 address to create a NAT binding in its table. Then, the HGW forwards the IPv4 udp datagram to the 6rd BR. Any subsequent udp datagram designating the same udp port and IPv4 address from the same udp port and IPv4 address will use the same NAT binding. This specification contains no new requirement to the HGW.

3.6.1.3. 6rd Border Router

When the 6rd BR receives the udp datagram, it de-capsulates the IPv4 and udp headers, and extracts the 16-bit udp source port information in the udp header. Then, it calculates a new IPv6 source address by replacing the 6rd udp host delegated prefix with the 6rd udp delegated prefix. Then, the 6rd BR forwards the IPv6 datagram to the IPv6 destination.

3.6.2. Inbound Flow from the Subscriber IPv6 Host

3.6.2.1. 6rd Border Router

When the 6rd BR receives the IPv6 datagram, it extracts the 16-bit udp destination port information from the IPv6 destination address. It calculates the new IPv6 destination address by replacing the 6rd udp delegated address with the 6rd udp host delegated address. It also constructs the udp header by putting the extracted 16-bit into the destination port and XXXX (TBD) in the source port. Finally, the 6rd BR extracts the IPv4 address from the IPv6 destination address, encapsulate the IPv6 datagram with the udp header and IPv4 header, and forward it to the provider network.

3.6.2.2. Home Gateway

When the HGW receives the IPv4 datagram, it performs the standard NAT function by replacing the destination IPv4 address and udp destination port to the address and port stored in the NAT binding table. Then, it forwards the datagram to the host in the LAN. This specification contains no new requirement to the HGW.

3.6.2.3. 6rd UDP Host

When the 6rd host receives the datagram. It de-capsulates the udp and IPv4 headers and forwards the IPv6 datagram to its IPv6 interface or to connected destination.

3.7. Examples

3.7.1. Host Model

This example explains how the Host Model works. Both 6rd Host A and 6rd Host B learn the 6rd prefix via static or dynamic configuration. They also learn the external IPv4 address by some external mechanism. Host A and Host B create the 6rd delegated prefix by concatenating 6rd prefix, the external IPv4 address, and pads 16-bit at the end of the external IPv4 address. The resulting 6rd udp host delegated prefix must be less than or equal to 64.

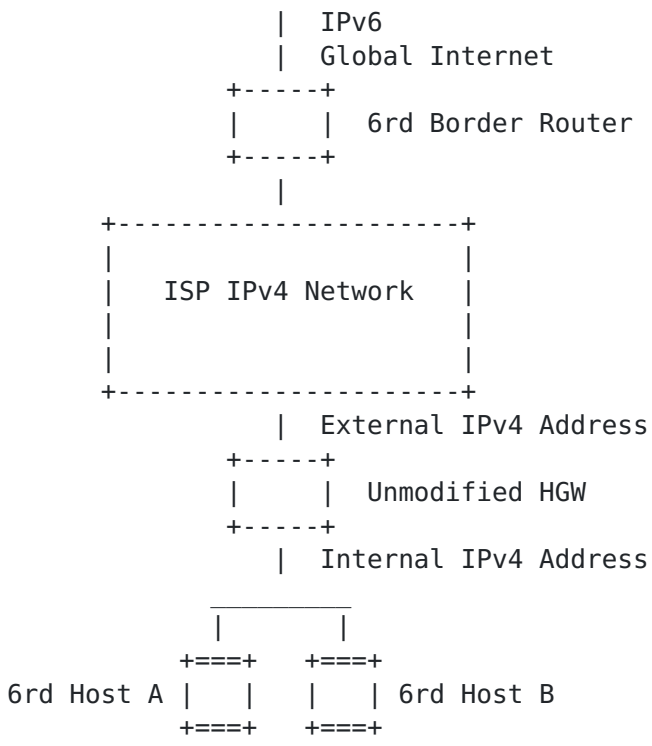

Figure 1

3.7.2. Server Model

This example is very similar to the Host Model. Instead of the host implementing this specification, only a server device is required to implement this specification. The server uses the same mechanism described in the Host Model to construct the 6rd udp host delegated prefix. When the server is ready to provide the 6rd service, it announces the 6rd udp host delegated prefix in the Router Advertisement. Client A and Client B receive the 6rd udp host delegated Prefix and auto-config the IPv6 interface.

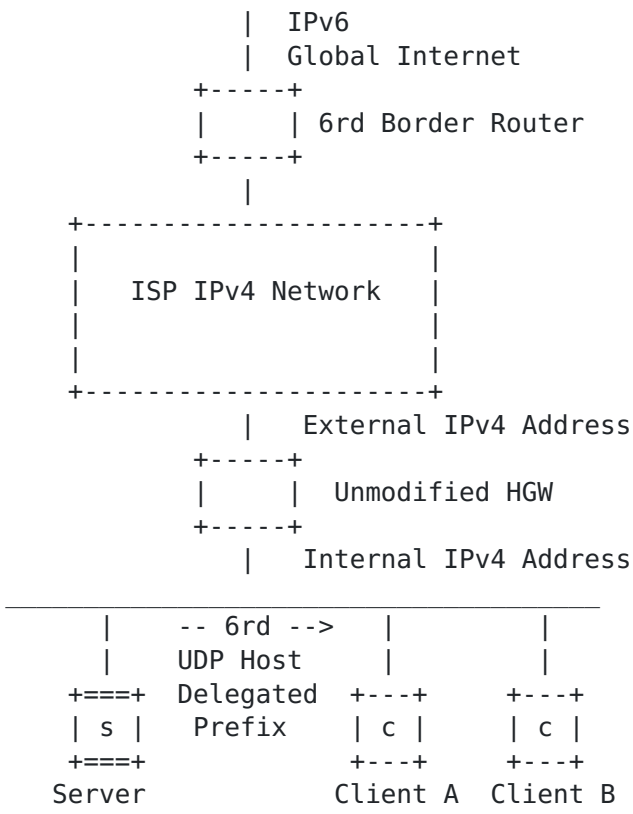

Figure 2

In this model, Client A and B are ordinary IPv6 hosts unaware of any 6rd specific knowledge. They learn the prefix and default router (which is the Server) via standard RA.

4. 6rd Prefix UDP Encapsulation

4.1. UDP-Encapsulated 6rd Header

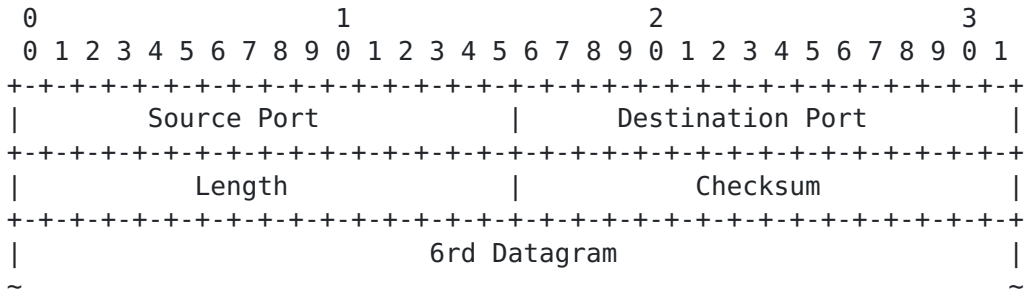

Where

- Source Port is any available port.
- Destination Port must be the well-known port assigned by IANA.
- this specification does not introduce new requirement to the IPv4 UDP Length and Checksum.

5. 6rd Border Router Discovery

In the classic 6rd framework, the 6rd CE connects directly to the ISP. The ISP uses DHCP to pass the 6rd Prefix, IPv4MaskLen, 6rdPrefixLen, and the 6rd BR address to the CPE CE. In this specification, the 6rd host is not directly connected to the ISP. Between the ISP and the 6rd host, there is a HGW which is unaware of 6rd. The ISP can't use DHCP to pass the necessary information to the 6rd hosts. In this specification, we suggest two discovery mechanisms.

5.1. Manual Discovery

The ISP gives the customers the 6rd Prefix, IPv4MaskLen, 6rdPreflixLen, and 6rd BR information. If the customers want to start the 6rd service, they must enter the information manually. This method is only feasible in very small scale deployment and not recommended for any large scale deployment.

5.2. Automatic Discovery

The ISP uses RADIUS [RFC2865] or DIAMETER [RFC3588] to distribute the 6rd Prefix, IPv4MaskLen, 6rdPrefixlen, and 6rd BR address information. When the customer starts the 6rd service, he/she must authenticate him/herself to the ISP. Upon a successful authentication, he/she will be given the necessary parameters through either RADIUS or DIAMETER response.

6. MTU

Similar to other tunnel encapsulations, this specification reduces the effect MTU size. The encapsulation overhead is 20-byte for IPv4 header and 8-byte for UDP. The host and 6rd BR must account for this

overhead.

7. Comparison to the Classic 6rd

This specification is considered more restrictive than the classic 6rd. There are three major areas:

7.1. 6rd Prefix Length

In the classic 6rd model, the maximum 6rd Prefix length can be as long as 32 + IPv4MaskLen. In this specification, the maximum 6rd Prefix length is 16 + IPv4MaskLen due to encapsulating the udp source port information in the 6rd udp delegated prefix.

7.2. Additional 6rd BR Operation

In the classic 6rd architecture, the 6rd BR does a simple de-capsulation and encapsulation. When the 6rd BR receives a udp datagram from the ISP network, the 6rd BR must calculate a new IPv6 source address after the de-capsulation. The new IPv6 source address contains the 6rd udp delegated prefix in which the udp source port information is embedded in it. When the 6rd BR receives an IPv6 datagram from the IPv6 Internet, it reverses the process by replacing the 6rd udp delegated prefix to the 6rd host udp delegated prefix.

7.3. Life Time of 6rd Delegated Prefix

In classic 6rd model, the life time of the 6rd delegate prefix is bounced by the life time of the external IPv4 address. The life time of the external IPv4 address could be hours or even days. In this specification, the life time of the 6rd udp delegated prefix is bounced by the life time of the NAT binding in the HGW. The HGW can expire the udp binding if there is no traffic passing for few seconds. This specification requires the 6rd udp host to send keepalive to the 6rd BR to refresh the binding frequently. However, the 6rd host udp delegated prefix used in the LAN will not suffer from the same short-lived interval. Since the 6rd host udp delegated prefix does not contain the udp port information, its life time is equal to 6rd delegated prefix which is bounced by the life time of the IPv4 external address.

8. Comparison to Softwire Hub-and-Spoke and Teredo

Softwire Hub-and-Spoke [RFC5571] and Teredo [RFC4380] are two protocols that provide IPv6 connectivity to hosts behind typical HGW. Softwire Hub-and-Spoke uses L2TPv2 over UDP [RFC2661] for the tunnel

protocol; Teredo defines its own tunneling protocol for UDP encapsulation. This specification provides similar functionality. The benefit of this specification is that the operation is stateless and requires no control protocol. However, this specification require external bootstrapping process to pass provisioning information to the 6rd udp host.

9. Deployment Considerations

The same 6rd BR can support both classic 6rd and 6rd UDP encapsulation. To achieve this, the classic 6rd and 6rd udp encapsulation must use different 6rd prefixes. In a deployment scenario where customers have mixed 6rd CE and typical HGW, this specification potentially saves operation cost by deploying only one type of network equipment. This specification also useful for operator to speedup the 6rd deployment process by offering users a softwire download of 6rd udp host which works behind their home gateways, and for users who do not want to change their home gateways.

10. IANA Considerations

This specification requests IANA to assign a UDP port for the 6rd UDP encapsulation.

11. Security Considerations

TBD

12. Acknowledgements

TBD

13. References

13.1. Normative References

[I-D.ietf-softwire-ipv6-6rd]
Townsley, M. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)", draft-ietf-softwire-ipv6-6rd-10 (work in progress), May 2010.

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768,

August 1980.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.

[RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.

[RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

[RFC5569] Despres, R., "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)", RFC 5569, January 2010.

[RFC5571] Storer, B., Pignataro, C., Dos Santos, M., Stevant, B., Toutain, L., and J. Tremblay, "Softwire Hub and Spoke Deployment Framework with Layer Two Tunneling Protocol Version 2 (L2TPv2)", RFC 5571, June 2009.

13.2. Informative References

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.

[RFC2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3588] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

Authors' Addresses

Yiu L. Lee
Comcast

Email: yiu_lee@cable.comcast.com
URI: http://www.comcast.com

Prashant Kapoor
Xavient

Email: pkapoor@xavient.com
URI: http://www.xavient.com