

Operations Area
Internet-Draft
Intended status: Informational
Expires: April 15, 2011

Y. Lee
Comcast
V. Kuarsingh
Rogers Communications
October 12, 2010

IPv6 Transition Cable Access Network Use Cases
draft-lee-v6ops-tran-cable-usecase-00

Abstract

This memo describes some use cases to transition to IPv6 in cable access network. This memo discusses enabling dual-stack to users over various types of network infrastructures. It also describes impacts to network, operation, CPE, and applications.

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 April 15, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Internet-Draft

IPv6 Transition Cable Use Cases

October 2010

Table of Contents

1.	Introduction	3
2.	Offer Dual-Stack on Top of Existing Access Network	3
2.1.	IPv4-only Access Network	4
2.1.1.	6rd	4
2.1.1.1.	Deployment Requirements	4
2.1.1.2.	Network Impact	5
2.1.1.3.	Operation Impact	5
2.1.1.4.	CPE Impact	6
2.1.1.5.	Application Impact	6
2.1.2.	MPLS	6
2.2.	Native Dual-Stack Use Case	6
2.2.1.	IPv6 Address Design	7
2.2.2.	Provisioning	7
2.2.3.	Advertising Customer's Prefixes to the Access Network	7
2.2.4.	Benefits of Native Dual Stack	7
2.3.	Native Dual-Stack with Shared IPv4 Addresses Use Case	8
3.	Offer Dual-Stack on IPv6-only Access Network	8
3.1.	Shared IPv4 Address Use Case	8
3.1.1.	DS-lite	8
3.1.1.1.	Deployment Requirements	8
3.1.1.2.	Network Impact	9
3.1.1.3.	Operation Impact	9
3.1.1.4.	CPE Impact	10
3.1.1.5.	Application Impact	10
3.2.	Public IPv4 Address Use Case	11
3.2.1.	IPv4 Over IPv6	11
3.2.1.1.	Deployment Requirements	11
3.2.1.2.	Network Impact	12
3.2.1.3.	Operation Impact	12
3.2.1.4.	CPE Impact	12
3.2.1.5.	Application Impact	12
4.	Security Considerations	12
5.	Acknowledgements	13
6.	IANA Considerations	13
7.	References	13
7.1.	Normative References	13
7.2.	Informative References	13
	Authors' Addresses	14

1. Introduction

The Cable access network primarily uses DOCSIS technology defined by CableLabs to deliver IP services to users. DOCSIS provides an abstraction to deliver IP packets over coaxial cable. DOCSIS is a shared media technology and use Ethernet for Layer-2, it doesn't use PPP or ATM for encapsulation.

A Cable Modem which is a DOCSIS enabled modem is the device to transmit the user's Ethernet frames over DOCSIS to the Cable Modem Termination System (CMTS) in the cable operator's network. DOCSIS has gone through few generations. The most current version is DOCSIS 3.0. By specifications, DOCSIS 2.0 and DOCSIS 3.0 both support IPv6 for cable modem management and user's traffic. However, DOCSIS 1.x specification and some older DOCSIS 2.0's implementations do not. Cable operators will take time to retire all the legacy cable modems and replace them to the newer version of cable modems. So there will be a transition period to upgrade all the equipments to support IPv6.

The complexity of upgrading the regional and core network to dual-stack is relatively low compared to upgrading the access network to support IPv6 for thousands of CMTSes and millions of cable modems and CPEs. So this memo focuses on use cases to enable IPv6 in the cable access network. The transition methodology is to provide dual-stack to the users regardless the underneath technology inside a cable operator. When IPv6 services become majority and IPv4 services gradually diminish, the operator may consider to provide only IPv6 to users and provide IPv4-IPv6 translation in the network when users access IPv4 services. This memo describes use cases to provide dual-stack to users because we have more experience.

We divide the use cases into two primary categories. The first category describes dual-stack deployment to the users using the existing access network. The access network could be IPv4-only or dual-stack. The second category describes dual-stack deployment to the users using IPv6-only access network. The goal of these use

cases is providing service continuity during the transition.

[2.](#) Offer Dual-Stack on Top of Existing Access Network

We discuss three use cases that offer dual-stack to users. The first use case describes the scenario where the access network is IPv4-only and operators utilize tunneling technologies to give dual-stack access to users. The second use case describes the standard native dual-stack deployment model. The third use case describes native dual-stack where the IPv4 connection may be provided using shared public IPv4 addresses (NAT444).

[2.1.](#) IPv4-only Access Network

According to [[I-D.arkko-ipv6-transition-guidelines](#)], native dual-stack is the simplest model for transition. However, this requires the entire network to be dual-stack. Moreover, the provisioning system and other support systems must be upgraded to support IPv6. Most operators will need to upgrade the network in phases along with the provisioning system(s) and supporting systems. During the transition period, there will be IPv4-only islands. In order to offer dual-stack access to users over IPv4 islands, operators may consider the use of tunneling technologies such as 6rd and MPLS.

There are incentives to offer IPv6 to users before completing the upgrade. For example: early IPv6 adopters can start experiencing IPv6 services and have connectivity to IPv6-only content should it be available. Operational groups can also begin to familiarize themselves with IPv6 and begin troubleshooting IPv6. Application developers and content providers can start providing services over IPv6. In the end, this may help to speedup the overall IPv6 adoption.

[2.1.1.](#) 6rd

[2.1.1.1.](#) Deployment Requirements

6rd [[RFC5969](#)] is a technology that provides IPv6 connectivity over the existing IPv4 access network. The idea is simple, it leverages the 6to4 model [[RFC3056](#)] and uses the provider's specific prefix instead of the IANA assigned well-known prefix. This will give the

operator's control of both ingress and egress flows. This technology has been proven to be successful in real operator deployments [\[RFC5569\]](#).

6rd is comprised of two elements: 6rd-CE and 6rd-BR. 6rd-CE initiates an IPv6-in-IP tunnel to the 6rd-BR. 6rd-BR terminates the tunnel and forward the IPv6 packets to the IPv6 Internet. Similar to 6to4, 6rd uses the IPv4 address provisioned to the user to construct the IPv6 address. Since the IPv4 address is stored in the IPv6 prefix, the address translation is stateless.

6rd works when a user was provisioned with a public IPv4 address. It also works with [\[RFC1918\]](#) address when it is combined with a provider NAT44 function in the network. In this use case, we discuss only the public IPv4 address model.

[2.1.1.2](#). Network Impact

This describes the egress connection from the 6rd-CE to the IPv6 Internet. After the IPv6 packet was encapsulated in an IPv4 packet by 6rd-CE, the network will forward the packet similar to any other IPv4 packet. 6rd model is transparent to the IPv4 network. The packet will eventually arrive in the closest 6rd-BR for decapsulation, then it will be forwarded to IPv6 destination. The "closest" 6rd-BR is defined by the IP address used in combination with network routing conditions.

This describes the ingress connection from the IPv6 Internet to 6rd-CE. IPv6 packet with 6rd prefix in the destination address will be forwarded normally and arrive to the closest 6rd-BR. The 6rd-BR extracts the IPv4 information from the IPv6 address and encapsulates the IPv6 packet in an IPv4 packet. Then, it will forward the encapsulated packet to the IPv4 network.

The 6rd prefix is advertised by the 6rd-BR or by an upstream router on it's behalf. The operator will advertise this prefix within their network and towards the Internet and other neighboring peers. The operator also needs to assign an anycast address to the 6rd-BR. This

anycast address will be shared by all the 6rd-BR and will be advertised in the operator's IPv4 serving IGP. The 6rd-CE will send the encapsulated packets to this anycast address.

IPv6 packets are delivered on the IPv6-in-IP tunnel. MTU is a common consideration for any tunnel technology. Since 6rd is a stateless technology, the tunnel endpoints cannot perform fragmentation. The simplest solution is to increase default MTU size larger than 1500 bytes in the access network. More discussion can be found in [\[RFC5969\]](#).

Hosts behind the 6rd-CE may not be able to dynamically learn any DNS server via SLAAC, so they may query DNS from a DNS server in the IPv4 network. The DNS server in the IPv4 network should be configured process AAAA records.

[2.1.1.3.](#) Operation Impact

6rd is a stateless technology. It greatly simplifies the network design for scalability and high availability. Traffic engineering of the tunnels is not explicitly required since the 6rd-BRs are known via an IGP (or IGP assisted path). Operators can add or remove 6rd-BR in the network without transferring service states from one 6rd-BR to another 6rd-BR. Operators also need not assign any particular 6rd-BR to a 6rd-CE. 6rd-CE will rely on routing to find the closest 6rd-BR.

6rd is similar to VPN technology. 6rd packets are encapsulated and transparent to the network. Operator can operate, monitor and troubleshoot the 6rd network independently.

Considerations for 6rd include any in-line service or network device that monitors, controls or assists with traffic flows. Since 6rd sends IPv6 packets insider an IPv4 tunnel, all such systems must be 6rd aware to continue to supply the same functions for this new traffic type. Additionally, if an operator has enabled dynamic QoS within their access network, the overall detection, policy and enforcement infrastructure will need to be able to manage the control of IPv6 flows within an IPv4 tunnel.

[2.1.1.4.](#) CPE Impact

CPE is required to implement the 6rd-CE specification. 6rd-CE must be the first device connecting to the cable modem and is responsible for learning the 6rd prefix and construct the 6rd delegated prefix. The CPE is also responsible to advertise the 6rd delegated prefix to hosts behind the CPE. If the CPE implements SLAAC, the hosts behind the CPE learns the prefix and default gateway via Router Advertisement. As with the network portion, any service information, including QoS, will need to be carefully managed to support the IPv6-in-IP function.

[2.1.1.5.](#) Application Impact

Applications will have dual-stack and should behave identically as of running on a native dual-stack host. Applications which are served via IPv6 will add additional load to BRs within the network. The operator may want to take this under consideration if they are planning to deploy high bandwidth services over IPv6. The operator may choose to offer some services over IPv4 in this case to lower the load on the BRs and allow for more efficient traffic delivery inside the network (since the BR and application systems may not share network locations).

[2.1.2.](#) MPLS

TBD

[2.2.](#) Native Dual-Stack Use Case

Providing native dual-stack to user may be the simplest for transition to IPv6, but it requires operators to upgrade the network, provisioning systems, and supporting systems to give production grade service to users. In this memo, native dual-stack means to provision a public IPv4 address, a global IPv6 address, and a global IPv6

prefix to a user.

[2.2.1.](#) IPv6 Address Design

In general, most of the IPv4 address architecture rules still apply to the IPv6 address architecture. For example: each service (e.g. VoIP vs. IPTV) should use different prefixes. Also, operators should use two separate prefixes for network infrastructure and customer

services.

Due to the high utilization and the allocation policies of IPv4 prefixes, the result is each organization got many discontinuous blocks of prefixes rather than a large aggregate. The drawback is a fairly large Internet routing table. The overall IPv6 address pool is 128-bit long. Operators are normally given a prefix that contains an enormous number of addresses. If an operator carefully plans for address allocation and aggregation, it should only advertise the provider's prefix to the IPv6 Internet routing table. For example: each regional network should be a suffix of the overall provider's prefix. The result should be a smaller and more organized Internet routing table. In contrast, bad IPv6 address design may result a divided routing table and unnecessarily bubble its size.

[2.2.2.](#) Provisioning

TBD

[2.2.3.](#) Advertising Customer's Prefixes to the Access Network

Apart from an IPv6 address assignment to the CPE, the network will also delegate a prefix to the CPE for the hosts behind the CPE. This prefix is normally assigned by a DHCP server. The access network will need to learn the prefix and the associated cable modem and CPE. [[I-D.droms-dhc-dhcpv6-agentopt-delegate](#)] suggests that the DHCP Relay Agent which is the CMTS can query the DHCP server and learn the prefix. Then, it installs the prefix into its routing table. Another way is the DHCP Relay Agent inspects the DHCP IA_PD reply from the DHCP server and installs the prefix to the routing table. This topic remains open and more development is coming.

[2.2.4.](#) Benefits of Native Dual Stack

Utilizing a native dual stack option for IPv4 and IPv6 connectivity includes the overall integration ease for the provider. Although this option requires the deployment of IPv6, it is the more understood and support option. Other than standard IPv6 functionality within the network providers space and in the CPE, no new options are necessarily needed. Many inline services will need

to support IPv6, but are likely to support IPv6 native before newer

connectivity options which includes DS-lite, 6rd and other such tunneling modes.

[2.3.](#) Native Dual-Stack with Shared IPv4 Addresses Use Case

This use case is an extension of the previous native dual stack option. In this particular case, all the IPv6 deployment considerations are made with an added complexity of shared IPv4 access. Shared IPv4 connectivity with a provider controlled NAT44 function may be required for dual stack deployments after IPv4 exhaustion. This option provides many of the same advantages as the native dual stack option which includes in the clear IPv4 and IPv6 flows. The provider can still utilize existing systems that support native IPv4 and IPv6 flows, but will need to add in network functionally related to the NAT44 function.

[3.](#) Offer Dual-Stack on IPv6-only Access Network

When the access network is IPv6-only, IPv6 traffic can be delivered natively over IPv6. So, there is no new requirement to enable IPv6. However, the access network will not be able to deliver IPv4 services. We provide two use cases to give dual-stack to users in an IPv6-only access network.

[3.1.](#) Shared IPv4 Address Use Case

When IPv4 addresses are limited, operators may consider multiplexing IPv4 addresses among internal users. Users will not be provisioned with a public IPv4 address. Instead, users will share a pool of public IPv4 addresses in the network.

DS-lite [[I-D.ietf-softwire-dual-stack-lite](#)] is a technology that provides IPv4 access over an IPv6-only access network. This also provides NAT44 functionality in the operator's network to multiplex a pool of public IPv4 addresses amongst users.

[3.1.1.](#) DS-lite

[3.1.1.1.](#) Deployment Requirements

DS-lite is composed of two elements: B4 element and AFTR element. B4 element initiates an IP-in-IPv6 tunnel to the AFTR. AFTR terminates the tunnel and performs NAT44. B4 element can be implemented in a CPE or in a host. For this use case, we only discuss the CPE B4 element model.

An operator is required to deploy B4 to user premises. B4 will replace the existing CPE and must be the first network device in front of the cable modem. The operator will provision an IPv6 address to the B4 element. It will not provision any IPv4 address to the B4. Operator will also provision an IPv6 Prefix to the B4 and B4 will advertise this IPv6 prefix to the hosts behind it so that IPv6-capable hosts will have native IPv6 services.

B4 will run as DHCP server to the hosts behind it. It also acts as IPv4 default gateway and DNS proxy to the hosts. IPv4 packets will be delivered over the IP-in-IPv6 tunnel between the B4 and AFTR. From the host perspective, it will be provisioned with dual-stack and the applications running on the host can decide to use IPv4 or IPv6.

An operator is required to deploy a set of AFTR elements in the network. The AFTR should be dual-stack to terminate the IP-in-IPv6 tunnel from B4 elements and deliver NAT-ed packets to IPv4 Internet.

[3.1.1.2.](#) Network Impact

DS-lite requires the access network to support IPv6. This requires the CMTS and cable modem to be IPv6 enabled. It also requires to deploy a set of AFTR elements in the operator network. AFTR is a stateful network device, it inherits the cost to manage a stateful network device inside the network.

IPv4 packets are delivered on the IP-in-IPv6 tunnel. This reduces the effective MTU size. Neither hosts behind the B4 element nor services in front of the AFTR are aware of the tunnel. The operator can increase the MTU size in the access network. However, many cable modem implementations do not support MTU larger than default 1500 bytes, so the B4 and AFTR elements must handle fragmentation caused by the tunnel overhead.

The AFTR owns the NAT pool, it will be the aggregation point of the IPv4 addresses defined in the NAT pool. AFTR must advertise the NAT pool prefix to the IPv4 Internet. In contrast, the IPv6 tunnel interface should stay only inside the operator's IGP and should not be advertised to the IPv6 Internet.

[3.1.1.3.](#) Operation Impact

DS-lite identifies a user by IPv6 address. Operators should be trained to understand how to map a user from an IPv6 address in the AFTR. AFTR is a NAT device, operator should maintain the NAT binding information to satisfy the government regulations. This is standard

procedure for operating any NAT44 device.

DS-Lite introduces the operational mode where historical IPv4 connectivity (as experienced) is now totally dependent on IPv6. This significant change in operating conditions must be well understood by the operator. If DS-lite is introduced during deployment infancy in the operators IPv6 network, it will require careful attention to operational practices and capabilities to maintain the IPv6 network.

AFTR is critical to continuously offer IPv4 access in IPv6-only access network. Operator should scale AFTR to provide non-interruptive access to users. Operators should closely monitor two AFTR's resources: (1) Network Capacity and (2) Port Utilization. When network capacity is reached, the operator should decide to upgrade the AFTR to higher network capacity or to deploy a new AFTR to balance the workload. When port utilization is high, the operator should increase the NAT pool size.

AFTR is stateful, it will complicate the high-availability (HA) design. Operators should apply the standard HA design (e.g. cold or hot) which best fits to their network operations.

[3.1.1.4.](#) CPE Impact

CPE is required to implement the B4 element specification. Also, port-forwarding and UPnP IGD protocol will no longer function. IETF PCP Working Group was formed to address the port-forwarding and UPnP IGD issues.

CPE must know the IPv6 address of the AFTR tunnel interface. This information can be obtained from DHCP. Since there is only IPv6 access to the B4 element. Any IPv4 network service learned from DHCP must be proxy by the B4 element.

If the operator cannot increase the access network MTU size, the B4 element must handle fragmentation to ensure IPv4 service using maximum MTU size won't be affected by the tunnel overhead.

[3.1.1.5.](#) Application Impact

[3.1.1.5.1.](#) Egress Connection

Since hosts behind B4 are provisioned with dual-stack, the application can decide to use IPv4 or IPv6. If the external service is also dual-stack, the host will automatically prefer IPv6 over IPv4 if the host O/S has implemented [\[RFC3484\]](#). If the host prefers IPv4 due to application logic, it will use the private IPv4 address provisioned by the B4 element. For applications expecting to use specific source port will be impacted because the AFTR inside the network won't be able to allocate a specific source port.

Applications use random source port will continue to function without modification.

[3.1.1.5.2.](#) Ingress Connection

Similar to traditional NAT, ingress connection will be blocked by default. The current techniques such as port-forwarding and UPnP IGD are required modification. Technically this could be done. But this will requires some changes in user's procedure to enable the service. It also adds cost to operators to offer port-forwarding service.

[3.2.](#) Public IPv4 Address Use Case

Some applications requires specific source port and some applications requires ingress connection. Users using those applications may want to be provisioned with a public IPv4 address to ease the potential challenges caused by NAT in the network. IPv4-over-IPv6 (4over6) [\[I-D.cui-softwire-host-4over6\]](#) is a simple technology to provision a public IPv4 address to a user and provide IPv4 access over an IPv6-only network.

[3.2.1.](#) IPv4 Over IPv6

[3.2.1.1.](#) Deployment Requirements

4over6 consists of two elements: 4over6 Initiator and 4over6 Tunnel Concentrator (TC). 4over6 is similar to DS-lite except two features: (1) Unlike B4 element, 4over6 Initiator will be provisioned with a public IPv4 address. (1) 4over6 TC only terminates the IP-in-IPv6 tunnel and won't perform any NAT44 function.

4over6 supports both host and CPE models. We will only discuss the

6over6 CPE model.

An operator is required to deploy 4over6 Initiator in premises. The 4over6 initiator will replace the existing CPE and must be the first network device in front of the cable modem. The operator will provide an IPv6 address and an IPv6 prefix to the CPE. The procedure is similar to Native IPv6 use case and DS-lite use case.

4over6 Initiator is very similar to the B4 element. It serves as DHCP server, IPv4 default gateway and DNS server to hosts behind it. The only difference is 4over6 will be provisioned with a public IPv4 address while B4 element will not. Once 4over6 Initiator discovers the 4over6 TC, it will issue standard DHCP request over the tunnel to the 4over6 TC. The 4over6 TC either relays the DHCP request to a centralized DHCP server or replies to the request if it is the authoritative DHCP server for the 4over6 service. Once the CPE

acquires the public IPv4 address, the user can run all his legacy IPv4 applications similar to what he is doing with a regular IPv4 home gateway.

[3.2.1.2](#). Network Impact

Similar to DS-lite, the access network must support IPv6. This requires the CMTS and cable modem must be IPv6 enabled. It also requires the operator to deploy a set of 4over6 TC in the network.

Despite no NAT in the 6over4 TC, 6over4 TC is required to maintain the 4over6 Initiate IPv6 address (tunnel-id) and IPv4 address binding. Also, the 4over6 TC must advertise the IPv4 prefix to the Internet. It is the aggregation point of the IPv4 address prefix.

4over6 suffers the same MTU limitation which is common to any tunnel protocols. Please refer to [Section 3.1.1.2](#) for details.

[3.2.1.3](#). Operation Impact

Since each user will be assigned a public IPv4 address, it doesn't require operator to log any binding. Operator should be able to identify a user by either IPv4 or IPv6 address.

Similar to AFTR, network capacity and IPv4 address utilization are

critical resources to 4over6 TC. Operator must closely monitor the resources to ensure continuous IPv4 access.

Operators also need to coordinate the IPv4 address space in the DHCP server and the 4over6 Initiator which manages the space. This requires careful coordination and management.

[3.2.1.4.](#) CPE Impact

CPE is required to implement the 4over6 TC specification. Unlike B4 element, port-forwarding and the UPnP IGD will work without modification.

[3.2.1.5.](#) Application Impact

Applications will have dual-stack and should behave identically as of running on a native dual-stack host.

[4.](#) Security Considerations

TBD

[5.](#) Acknowledgements

TBD

[6.](#) IANA Considerations

This memo includes no request to IANA.

[7.](#) References

[7.1.](#) Normative References

[I-D.arkko-ipv6-transition-guidelines]

Arkko, J. and F. Baker, "Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment", [draft-arkko-ipv6-transition-guidelines-06](#) (work in

progress), August 2010.

[I-D.cui-software-host-4over6]

Cui, Y., Wu, J., and P. Wu, "Host 4over6 for IPv6 host connecting IPv4 Internet", [draft-cui-software-host-4over6-01](#) (work in progress), July 2010.

[I-D.ietf-software-dual-stack-lite]

Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", [draft-ietf-software-dual-stack-lite-06](#) (work in progress), August 2010.

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

[RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", [RFC 5969](#), August 2010.

[7.2](#). Informative References

[I-D.droms-dhc-dhcpv6-agentopt-delegate]

Droms, R., "DHCP Relay Agent Assignment Notification Option", [draft-droms-dhc-dhcpv6-agentopt-delegate-00](#) (work in progress), November 2005.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets",

[BCP 5](#), [RFC 1918](#), February 1996.

[RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", [RFC 3056](#), February 2001.

[RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", [RFC 3484](#), February 2003.

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

Authors' Addresses

Yiu L. Lee
Comcast
One Comcast Center
Philadelphia, PA 19103
U.S.A.

Email: yiul_lee@cable.comcast.com
URI: <http://www.comcast.com>

Victor Kuarsingh
Rogers Communications
8200 Dixie Road
Brampton, Ontario L6T 0C1
Canada

Email: victor.kuarsingh@rci.rogers.com
URI: <http://www.rogers.com>