

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 28, 2011

Y. Cui
M. Xu
P. Wu
S. Wang
J. Wu
X. Li
Tsinghua University
C. Metz
Cisco Systems, Inc.
October 25, 2010

Translation Spot Negotiation in IPv4/IPv6-Coexist Mesh
draft-cui-software-pet-03

Abstract

IPv4 and IPv6 are expected to coexist for a long period. Currently, there are many IPv4/IPv6 transition/coexistence techniques, roughly divided into the categories of tunneling and translation. Tunneling and translation have respective application scopes, and translation has some technical limitations, including scalability issue, application layer translation, operation complexity, etc. To improve the availability of translation, this draft proposes the method of selecting appropriate translation spot to execute translation. When the translation spot is not on IPv4-IPv6 border, tunnel is used to achieve the traversing between translation spot and IP border. This method applies well in mesh scenario where both IPv4 and IPv6 client network exists, and BGP can be extended to achieve a translation spot signaling.

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 28, 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	4
2. Translation Spot Selection	6
3. Translation Spot Selection in IPv4/IPv6-coexist Mesh	8
3.1. Scenario description	8
3.2. Translation between IPvX and IPvY networks	9
3.3. Translation between IPvX network and IPvY Internet	9
3.4. Translation between IPvY network and IPvX Internet	9
4. Translation Spot Signaling	10
4.1. Signaling content	10
4.2. Extensions in MP-BGP	10
5. Further discussion	13
5.1. Achievement of translation spot selection	13
5.2. Cooperate with softwire	13
5.3. Using NAT64 or IVI as translation mechanism	13
6. IANA considerations	14
7. Acknowledgements	15
8. References	16
8.1. Normative References	16
8.2. Informative References	17
Authors' Addresses	18

1. Introduction

Recently more and more IPv6 networks have been deployed, especially IPv6 backbone networks. However the existing IPv4 networks still carry the major network traffic and hold the major network services and applications. It has been widely believed that IPv4 and IPv6 networks will coexist for a long term. This leads to the demand for IPv4-IPv6 coexistence technology.

Till now there are two types of IPv4-IPv6 coexistence techniques: tunneling and translation. Tunneling can achieve IPv4-over-IPv6/IPv6-over-IPv4 traversing, by means of encapsulation and decapsulation. Examples of tunneling methods include IP-in-IP tunnel [RFC2893][RFC4213], GRE tunnel [RFC1702], 6to4 tunnel [RFC3056], 6over4 tunnel [RFC2529], softwire mesh technique [RFC5565], etc. Tunneling is transparent and light-weighted. It can be implemented fully by hardware.

On the other hand, translation is used to achieve IPv4-IPv6 inter-communication, by means of converting the semantic between IPv4 and IPv6. Examples of translation methods include SIIT [RFC2765], NAT-PT [RFC2766], BIS [RFC2767], BIA [RFC3338], IVI [I-D.xli-behave-ivil], NAT64 [I-D.ietf-behave-v6v4-xlate-stateful] and so on. Translation can achieve IPv4-IPv6 interworking which tunneling cannot do, but it has several technical limitations:

- o Scalability. In stateful translation, the dynamic mapping of (address, port) tuple should be maintained on the translation device. The total number of mapping entries is up to the order of flow number. As to stateless translation, it has to consume IPv4 addresses to satisfy IPv6 hosts. This is also not scalable since IPv6 address space is much larger than IPv4 address.
- o Application layer translation. Since translation will modify the address of an IP packet, or we say an end host, an application protocol that contains IP addresses in its payload won't work if we don't convert the addresses. However, due to the variety of applications protocols, it's unrealistic for the translation device to support all of them.
- o Operation complexity. To accomplish correct translation, the following operations are required: address or (address, port) tuple conversion, IP and ICMP fields translation, TCP/UDP checksum re-computing, application layer detection and translation, fragmentation when necessary. It's rather complex for a per-packet process and probably unacceptable when the volume is high.

- o Lack of efficient NAT46 translation mechanism. No efficient IPv4 to IPv6 communication mechanism has been proposed since NAT-PT. A fundamental difficulty here is that IPv6 address space is much larger than IPv4 so the translation mechanism has to make DNS or other addressing method stateful. Obviously this is not scalable.

Though facing all these issues, translation is irreplaceable in its application scope, so it's necessary to find a way to improve its availability. To solve this problem, this draft proposes the method of finding the appropriate translation spot to execute translation. The method adopts tunnel when necessary, to achieve traversing between translation spot and IP border. As an attempt, this draft applies the method in IPv4/IPv6-coexist mesh scenario, and extends BGP to achieve translation spot signaling in the scenario.

2. Translation Spot Selection

The issues of translation listed in section 1 are inherent disadvantages due to the principle of translation. Hence it's difficult to solve these problems by improving the mechanism. However, by choosing the appropriate location to perform translation, these problems can be solved or lightened, and translation can be more available. This draft calls the location to perform translation as "translation spot".

The basic idea of translation spot selection is to choose the place where the scalability and complexity is not a concern, i.e., the place where the translator is capable for its own translation traffic. Following this thought, a straightforward principle is to push translation down to edge networks. Since the volume of translation traffic in edge networks is relatively low, it's possible to achieve a real-time per-flow mapping and per-packet modification there. On the contrary, traffic in backbone is aggregated and hence much higher in volume. So routers in backbone would rather only support routing and forwarding than take charge of high-speed translation. However, when the total translation volume is low, it's easier to perform a unified translation in backbone than to distribute the job to many edge networks.

To achieve flexible translation spot selection, there's still a difficulty in packet forwarding: in a given topology, the IPv4-IPv6 border spot is fixed; If the translation spot isn't identical to the IP border spot, the packets can't be forwarded between the two spot due to IP diversity. See the example in Figure 1. The IP border is on spot 2 between IPvY backbone and IPvX Internet while the translation spot can be spot 1 or spot 2. If spot 1 is chosen, then packets from IPvY edge network are translated into IPvX on spot 1; they have to traverse to IPvY backbone to reach IPvX Internet. , and packets from IPvX Internet have to traverse the IPvY backbone to reach spot 1 for translation. Similar thing happens when spot 2 is chosen in Figure 2.




Figure 1 Translation Spot selection




Figure 2 Translation Spot selection

This is actually a traversing problem and the typical solution is tunneling. By building a tunnel to connect IP border and the translation spot, the forwarding path can be achieved. In the example of Figure 1, an IPvX-over-IPvY tunnel between spot 1 and spot 2 can be used to forward translated-to-IPvX packets from spot 1 to spot 2, and to-be-translated IPvX packets from spot 2 to spot 1. In Figure 2, an IPvY-over-IPvX tunnel between spot 1 and spot 2 can be used to forward to-be-translated IPvY packets from spot 1 to spot 2, and translated-to-IPvY packet from spot 2 to spot 1. Although the flexible translation spot selection may require an extra tunnel, its cost is much lower than translation, and hence acceptable.

3. Translation Spot Selection in IPv4/IPv6-coexist Mesh

3.1. Scenario description

Translation spot selection can be used in many scenarios. As a demonstration this draft applies it to the mesh scenario described in Figure 3. In this scenario, an IPvX-only backbone is connected to both IPvX networks and IPvY networks. The backbone may also have entrance to IPvX and IPvY Internet. Besides native traffic and IPvY-over-IPvX software traffic described in [RFC4925], there're also traffics between IPvX and IPvY networks, between IPvX network and IPvY Internet, and between IPvY network and IPvX Internet. All these three types of traffics require translation, which should be performed on AFBRs (Address Family Border Router) or BRs (Border Router) on the border of the backbone.

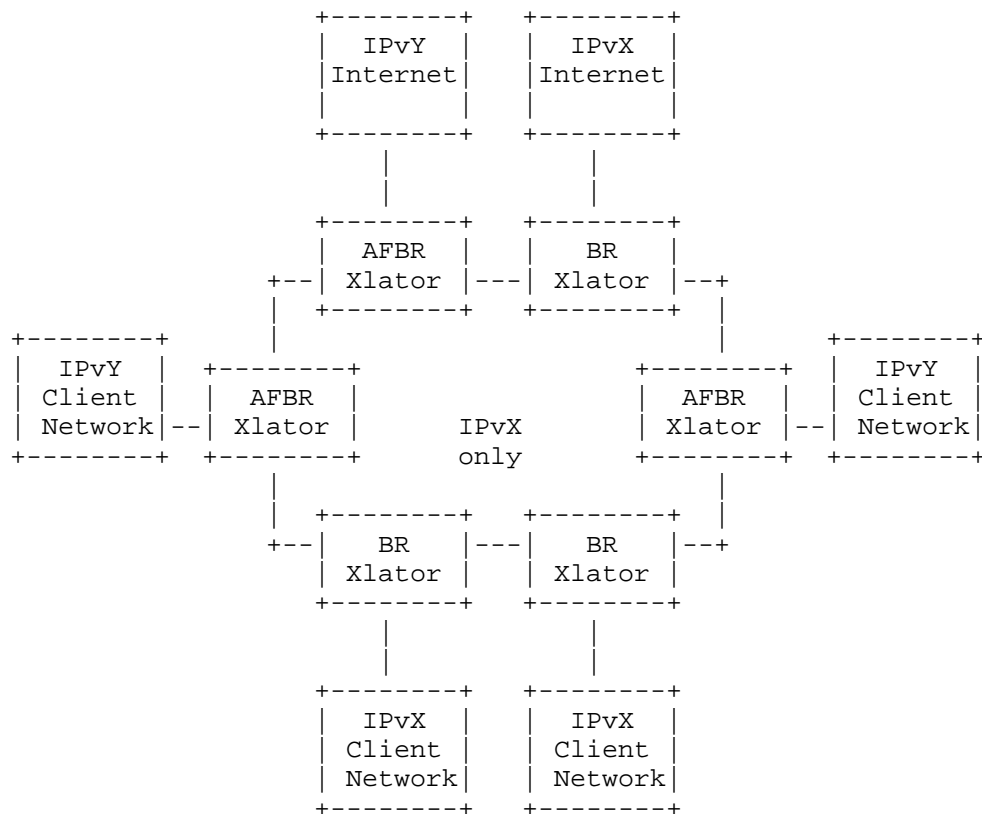


Figure 3 Translation Spot Selection in IPv4/IPv6-coexist Mesh

3.2. Translation between IPvX and IPvY networks

The communication between an IPvX network and an IPvY network follows the path "IPvX network - BR - IPvX backbone - AFBR - IPvY network". The translation can be performed either on the BR between IPvX network and backbone, or on the AFBR between IPvX backbone and IPvY network.

If the BR is chosen to be translation spot, a tunnel should be established for packet forwarding between the BR and the AFBR. Naturally it could be a softwire tunnel since it's a mesh scenario. Besides, to perform correct translation, BR needs the translation context delivered from the AFBR. This will be discussed in the next section.

3.3. Translation between IPvX network and IPvY Internet

The communication between an IPvX network and IPvY Internet follows the path "IPvX network - BR - IPvX backbone - AFBR - IPvY Internet". The translation spot can be either the BR between IPvX network and backbone, or the AFBR between IPvX backbone and IPvY Internet. BR can be chosen to avoid scalability and operation complexity issues, and AFBR can be chosen for unified translation purpose.

If the BR is chosen to be translation spot, a softwire tunnel should be established between the BR and the AFBR. Also BR needs the translation context delivered from the AFBR.

3.4. Translation between IPvY network and IPvX Internet

The communication between an IPvY network and IPvX Internet follows the path "IPvY network - AFBR - IPvX backbone - BR - IPvX Internet". The translation spot can be either the AFBR between IPvY network and IPvX backbone, or the BR between IPvX backbone and IPvX Internet. Usually the AFBR is preferred in this case, since it's the IP border and traffic is not so aggregated as in BR. However, BR can be chosen for unified translation purpose.

If the BR is chosen to be translation spot, a softwire tunnel should be established between the BR and the AFBR. Also BR needs the translation context delivered from the AFBR.

In all three types of translation-involved communication, translation spot selection is feasible. Yet an auto negotiation method is required to make the translation spot selection and translation context advertisement process more practical in the mesh scenario. This will be discussed in the next section.

4. Translation Spot Signaling

In the IPv4/IPv6-coexist mesh, the total number of client networks, and hence the total number of AFBRs and BRs could be quite high, so an auto negotiation method is required to select the translation spot for all translation-involved communications, rather than manual configuration on every AFBR and BR. This negotiation method is called translation spot signaling.

4.1. Signaling content

It's clear that translation should be performed on an appropriate translator, or as in this scenario, an AFBR or BR device. Here the concept of Translation Preference (TP) is defined to represent the appropriateness of a device to perform translation. TP is a quantified value set by the administrator of the corresponding AFBR or BR device. By exchanging and comparing TP values, two translators can decide which one to be the translation spot.

The TP value should be decided by the administrator. The general criterion here is, the translator whose performance is better, whose traffic volume is lower, and the size of network behind which is smaller (thus the translation traffic is less aggregated), is preferred to do translation and should have a high value. TP can also be configured based on administrator's policy, such as unified translation.

Tps for stateless and stateful translation are separated because they have different foundations (stateless translation requires IPv6 host to possess IPv4 address). In a mixed scenario, some translators can't perform stateless translation like others because IPv6 hosts in its network don't own IPv4 addresses.


Besides TP, translation context should also be advertised through signaling. The translation context is the necessary knowledge to perform a translation. For stateless translation it's the mapping prefix, and for stateful translation it's the address pool used for address mapping. For example, in the type of "IPv6 network - BR - IPv6 Backbone - AFBR - IPv4 Internet" communication, if stateless translation is adopted, then AFBR should tell BR the prefix for IPv4-IPv6 address mapping when BR performs the translation; if stateful translation is adopted, then AFBR should tell BR the IPv4 addresses BR can use for address mapping when BR performs the translation.

4.2. Extensions in MP-BGP

MP-BGP is adopted to carry the translation spot signaling process since it fits the mesh scenario and is already used in software

mesh[RFC5565].

We define a new a new BGP Attribute, "Translation Information Attribute" to carry the TP and translation context information. It's an optional transitive attribute, and the attribute type code is TBD by IANA. The value field of this attribute is composed of a set of Type-Length-Value (TLV) encodings. The TLV is structured as follows. The Length field stands for the total number of octets in the Value field.



We define 4 TLVs here: Stateless_TP TLV, Stateful_TP TLV, IPv6_Prefix TLV and IPv4_pool TLV. More TLVs may be defined in the future when necessary.

- o Stateless_TP TLV has the type field assigned to 1 and length field assigned to 2. The value field is filled with the 16bit TP value for stateless translation. High the TP value means high preference to perform translation.
- o Stateful_TP TLV has the type field 2 and length field 2. The value field is filled with the 16bit TP value for stateful translation. High the TP value means high preference to perform translation.
- o IPv6_Prefix TLV has the type field assigned to 3. The length field is variable. The value field is filled with the IPv6 prefix for address mapping in stateless translation, encoding in NLRI format[RFC4760].
- o IPv4_pool TLV has the type field assigned to 4. The length field is variable. The value field is filled with the IPv4 pool for address mapping in stateful translation, encoding in NLRI format.

The AFBRs and BRs in the mesh should run MP-BGP process and peer with each other. When a new BGP session is established, AFBR and BR send a update containing the Translation Information Attribute to each other, which contains the Stateless_TP TLV or Stateful_TP TLV. Each router independently decides translation spot based on received TP

value. When the selected translation spot isn't the AFBR, then the AFBR should send another update with the Translation Information Attribute containing the IPv6_Prefix TLV or the IPv4_pool TLV to the BR. The tunnel-related routing should be triggered too, if there's any.

5. Further discussion

5.1. Achievement of translation spot selection

To be precise, through translation spot selection, we can solve the scalability problem of stateful translation and the operation complexity problem for both stateless and stateful translation. Also we make it more possible to perform application layer translation and adopt NAT46 mechanisms (NAT-PT) by pushing the translation spot down to the edge.

5.2. Cooperate with software

In the mesh scenario, software[RFC5565] is usually adopted as the tunnel mechanism. If it's used to support forwarding between the BR and the AFBR, then after translation spot signaling, BR and AFBR should trigger the software routing process, in which AFBR should advertise the actual IPv4 prefixes, while BR should advertise to AFBR either the address pool assigned from the AFBR (stateful case), or the IPv4 address prefix containing the IPv4 address possessed by the IPv6 hosts (stateless case).

5.3. Using NAT64 or IVI as translation mechanism

NAT64[I-D.ietf-behave-v6v4-xlate-stateful] is a typical stateful translation mechanism. It can be used in the IPv4/IPv6-coexist mesh for translation-involved communications across the backbone. If AFBR is chosen to be the translation spot, then the traffic will follow a traditional NAT64 process; else BR is chosen to be the translation spot, then AFBR should divided its public IPv4 address pool and assigned one block to the BR through translation spot signaling. BR will perform the NAT64 translation using the assigned IPv4 address block. In software routing, BR should advertise this block to AFBR.

IVI[I-D.xli-behave-ivi] is a typical stateless translation mechanism. It can be used in the IPv4/IPv6-coexist mesh for translation-involved communications across the backbone. If AFBR is chosen to be the translation spot, then the traffic will follow a traditional IVI process; else BR is chosen to be the translation spot, then AFBR should inform BR the IVI prefix, then BR can learn the address mapping role and the IPv4 prefix possessed by its network. In software routing, BR should advertise this IPv4 prefix to AFBR.

6. IANA considerations

IANA is requested to assign a value from the "BGP Path Attributes" Registry, to be called "Translation Information Attribute", with this document as the reference.

7. Acknowledgements

The authors would like to thank Lixia Zhang, Eric Nordmark, Jari Arkko, Alain Durand and David Ward for their valuable comments on this draft.

8. References

8.1. Normative References

- [RFC1702] Hanks, S., Li, T., Farinacci, D., and P. Traina, "Generic Routing Encapsulation over IPv4 networks", RFC 1702, October 1994.
- [RFC2529] Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC 2529, March 1999.
- [RFC2765] Nordmark, E., "Stateless IP/ICMP Translation Algorithm (SIIT)", RFC 2765, February 2000.
- [RFC2766] Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.
- [RFC2767] Tsuchiya, K., HIGUCHI, H., and Y. Atarashi, "Dual Stack Hosts using the "Bump-In-the-Stack" Technique (BIS)", RFC 2767, February 2000.
- [RFC2893] Gilligan, R. and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", RFC 2893, August 2000.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3338] Lee, S., Shin, M-K., Kim, Y-J., Nordmark, E., and A. Durand, "Dual Stack Hosts Using "Bump-in-the-API" (BIA)", RFC 3338, October 2002.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.
- [RFC4925] Li, X., Dawkins, S., Ward, D., and A. Durand, "Softwire Problem Statement", RFC 4925, July 2007.
- [RFC5565] Wu, J., Cui, Y., Metz, C., and E. Rosen, "Softwire Mesh Framework", RFC 5565, June 2009.

8.2. Informative References

[I-D.ietf-behave-v6v4-xlate-stateful]

Bagnulo, M., Matthews, P., and I. Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", draft-ietf-behave-v6v4-xlate-stateful-12 (work in progress), July 2010.

[I-D.xli-behave-ivi]

Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, "The CERNET IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition", draft-xli-behave-ivi-07 (work in progress), January 2010.

Authors' Addresses

Yong Cui
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R.China

Phone: +86-10-6278-5822
Email: cy@csnet1.cs.tsinghua.edu.cn

Mingwei Xu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P. R. China

Phone: +86-10-6278-5822
Email: xmw@csnet1.cs.tsinghua.edu.cn

Peng Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P. R. China

Phone: +86-10-6278-5822
Email: weapon@csnet1.cs.tsinghua.edu.cn

Shengling Wang
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P. R. China

Phone: +86-10-6278-5822
Email: slwang@csnet1.cs.tsinghua.edu.cn

Jianping Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P. R. China

Phone: +86-10-6278-5983
Email: jianping@cernet.edu.cn

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

Phone: +86-10-6278-5983
Email: xing@cernet.edu.cn

Chris Metz
Cisco Systems, Inc.
3700 Cisco Way
San Jose, Ca. 95134
USA

Email: chmetz@cisco.com

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: June 23, 2011

H. Singh
W. Beebe
Cisco Systems, Inc.
C. Donley
CableLabs
B. Stark
AT&T
O. Troan, Ed.
Cisco Systems, Inc.
December 20, 2010

Basic Requirements for IPv6 Customer Edge Routers
draft-ietf-v6ops-ipv6-cpe-router-09

Abstract

This document specifies requirements for an IPv6 Customer Edge (CE) router. Specifically, the current version of this document focuses on the basic provisioning of an IPv6 CE router and the provisioning of IPv6 hosts attached to it.

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 23, 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.

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
2. Terminology	3
3. Architecture	4
3.1. Current IPv4 End-user Network Architecture	4
3.2. IPv6 End-user Network Architecture	5
3.2.1. Local communication	6
4. Requirements	7
4.1. General Requirements	7
4.2. WAN Side Configuration	7
4.3. LAN Side Configuration	10
4.4. Security Considerations	13
5. Acknowledgements	13
6. Contributors	14
7. IANA Considerations	14
8. References	14
8.1. Normative References	14
8.2. Informative References	16
Authors' Addresses	16

1. Introduction

This document defines basic IPv6 features for a residential or small office router referred to as an IPv6 CE router. Typically these routers also support IPv4.

Mixed environments of dual-stack hosts and IPv6-only hosts (behind the CE router) can be more complex if the IPv6-only devices are using a translator to access IPv4 servers [I-D.ietf-behave-v6v4-framework]. Support for such mixed environments is not in scope of this document.

This document specifies how an IPv6 CE router automatically provisions its WAN interface, acquires address space for provisioning of its LAN interfaces and fetches other configuration information from the service provider network. Automatic provisioning of more complex topology than a single router with multiple LAN interfaces is out of scope for this document.

See [RFC4779] for a discussion of options available for deploying IPv6 in Service Provider access networks.

1.1. 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].

2. Terminology

End-user Network	one or more links attached to the IPv6 CE router that connect IPv6 hosts.
IPv6 Customer Edge router	a node intended for home or small office use which forwards IPv6 packets not explicitly addressed to itself. The IPv6 CE router connects the end-user network to a service provider network.
IPv6 host	any device implementing an IPv6 stack receiving IPv6 connectivity through the IPv6 CE router
LAN interface	an IPv6 CE router's attachment to a link in the end-user network. Examples are Ethernets (simple or bridged), 802.11 wireless or other LAN technologies. An IPv6 CE router may have one or more network

layer LAN Interfaces.

Service Provider	an entity that provides access to the Internet. In this document, a Service Provider specifically offers Internet access using IPv6, and may also offer IPv4 Internet access. The Service Provider can provide such access over a variety of different transport methods such as DSL, cable, wireless, and others.
WAN interface	an IPv6 CE router's attachment to a link used to provide connectivity to the Service Provider network; example link technologies include Ethernets (simple or bridged), PPP links, Frame Relay, or ATM networks as well as Internet-layer (or higher-layer) "tunnels", such as tunnels over IPv4 or IPv6 itself.

3. Architecture

3.1. Current IPv4 End-user Network Architecture

An end-user network will likely support both IPv4 and IPv6. It is not expected that an end-user will change their existing network topology with the introduction of IPv6. There are some differences in how IPv6 works and is provisioned which has implications for the network architecture. A typical IPv4 end-user network consist of a "plug and play" router with NAT functionality and a single link behind it, connected to the Service Provider network.

A typical IPv4 NAT deployment by default blocks all incoming connections. Opening of ports is typically allowed using UPnP IGD [UPnP-IGD] or some other firewall control protocol.

Another consequence of using private address space in the end-user network is that it provides stable addressing, i.e. it never changes even when you change Service Providers, and the addresses are always there even when the WAN interface is down or the customer edge router has not yet been provisioned.

Rewriting addresses on the edge of the network also allows for some rudimentary multi-homing; even though using NATs for multi-homing does not preserve connections during a fail-over event [RFC4864].

Many existing routers support dynamic routing, and advanced end users

can build arbitrary, complex networks using manual configuration of address prefixes combined with a dynamic routing protocol.

3.2. IPv6 End-user Network Architecture

The end-user network architecture for IPv6 should provide equivalent or better capabilities and functionality than the current IPv4 architecture.

The end-user network is a stub network. Figure 1 illustrates the model topology for the end-user network.

An example of a typical end-user network.

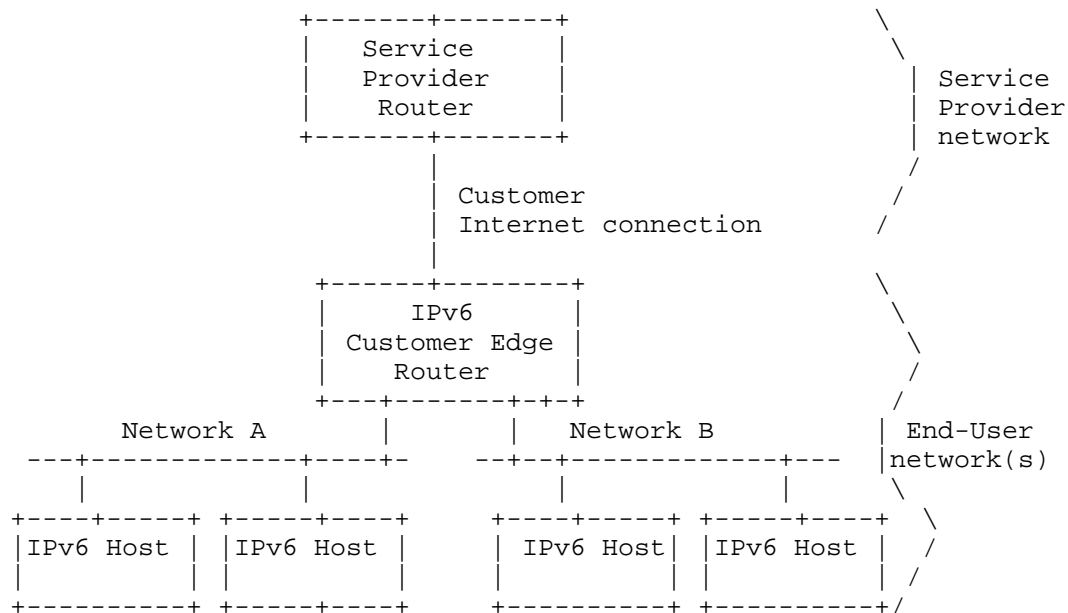


Figure 1

This architecture describes the:

- o Basic capabilities of an IPv6 CE router
- o Provisioning of the WAN interface connecting to the Service Provider

- o Provisioning of the LAN interfaces

For IPv6 multicast traffic the IPv6 CE router may act as an Multicast Listener Discovery (MLD) proxy [RFC4605] and may support a dynamic multicast routing protocol.

The IPv6 CE router may be manually configured in an arbitrary topology with a dynamic routing protocol. Automatic provisioning and configuration is described for a single IPv6 CE router only.

3.2.1. Local communication

Link-local IPv6 addresses are used by hosts communicating on a single link. Unique Local IPv6 Unicast Addresses (ULA) [RFC4193] are used by hosts communicating within the End-user Network across multiple links, but without requiring the application to use a globally routable address. The IPv6 CE router defaults to acting as the demarcation point between two networks by providing a ULA boundary, a multicast zone boundary and ingress and egress traffic filters.

A dual-stacked host is multi-homed to IPv4 and IPv6 networks. The IPv4 and IPv6 topologies may not be congruent and different addresses may have different reachability, e.g. ULA addresses. A host stack has to be able to quickly failover and try a different source address and destination address pair if communication fails as outlined in [I-D.wing-v6ops-happy-eyeballs-ipv6].

At the time of writing, several hosts implementations do not handle the case where they have an IPv6 address configured and no IPv6 connectivity. Either because the address itself has a limited topological reachability (e.g. ULA) or because the IPv6 CE router is not connected to the IPv6 network on its WAN interface. To support host implementations that do not handle multi-homing in a multi-prefix environment [I-D.ietf-v6ops-multihoming-without-nat66], the IPv6 CE router should, as detailed in the below requirements, not advertise itself as a default router on the LAN interface(s) when it does not have IPv6 connectivity on the WAN interface or when it is not provisioned with IPv6 addresses. For local IPv6 communication the mechanisms specified in [RFC4191] are used.

ULA addressing is useful where the IPv6 CE router has multiple LAN interfaces with hosts that need to communicate with each other. If the IPv6 CE router has only a single LAN interface (IPv6 link) then link-local addressing can be used instead.

In the event more than one IPv6 CE router is present on the LAN, then coexistence with IPv4 requires all of them to conform to these recommendations, especially requirements ULA-5 and L-4.

4. Requirements

4.1. General Requirements

The IPv6 CE router is responsible for implementing IPv6 routing; that is, the IPv6 CE router must look up the IPv6 Destination address in its routing table to decide to which interface it should send the packet.

In this role, the IPv6 CE router is responsible for ensuring that traffic using its ULA addressing does not go out the WAN interface, and does not originate from the WAN interface.

- G-1: An IPv6 CE router is an IPv6 node according to the IPv6 Node Requirements [RFC4294] specification.
- G-2: The IPv6 CE router MUST implement ICMP according to [RFC4443]. In particular point to point links MUST be handled as described in section 3.1 of [RFC4443].
- G-3: The IPv6 CE router MUST NOT forward any IPv6 traffic between its LAN Interface(s) and its WAN Interface until the router has successfully completed the IPv6 address acquisition process.
- G-4: By default an IPv6 CE router that has no default router(s) on its WAN interface MUST NOT advertise itself as an IPv6 default router on its LAN interfaces. That is, the "Router Lifetime" field is set to zero in all Router Advertisement messages it originates [RFC4861].
- G-5: By default if the IPv6 CE router is an advertising router and loses its IPv6 default router(s) on the WAN interface, it MUST explicitly invalidate itself as an IPv6 default router on each of its advertising interfaces by immediately transmitting one or more Router Advertisement messages with the "Router Lifetime" field set to zero [RFC4861].

4.2. WAN Side Configuration

The IPv6 CE router will need to support connectivity to one or more access network architectures. This document describes an IPv6 CE router that is not specific to any particular architecture or Service Provider, and supports all commonly used architectures.

IPv6 Neighbor Discovery and DHCPv6 protocols operate over any type of IPv6 supported link-layer and there is no need for a link-layer specific configuration protocol for IPv6 network layer configuration options as in e.g. PPP IPCP for IPv4. This section makes the

assumption that the same mechanism will work for any link-layer, be it Ethernet, DOCSIS, PPP or others.

WAN side requirements:

- W-1: When the router is attached to the WAN interface link it MUST act as an IPv6 host for the purposes of stateless or stateful interface address assignment ([RFC4862] / [RFC3315]).
- W-2: The IPv6 CE router MUST generate a link-local address and finish Duplicate Address Detection according to [RFC4862] prior to sending any Router Solicitations on the interface. The source address used in the subsequent Router Solicitation MUST be the link-local address on the WAN interface.
- W-3: Absent of other routing information the IPv6 CE router MUST use Router Discovery as specified in [RFC4861] to discover a default router(s) and install default route(s) in its routing table with the discovered router's address as the next-hop.
- W-4: The router MUST act as a requesting router for the purposes of DHCPv6 prefix delegation ([RFC3633]).
- W-5: DHCPv6 address assignment (IA_NA) and DHCPv6 prefix delegation (IA_PD) SHOULD be done as a single DHCPv6 session.
- W-6: The IPv6 CE router MUST use a persistent DUID for DHCPv6 messages. The DUID MUST NOT change between network interface resets or IPv6 CE router reboot.

Link-layer requirements:

- WLL-1: If the WAN interface supports Ethernet encapsulation, then the IPv6 CE router MUST support IPv6 over Ethernet [RFC2464].
- WLL-2: If the WAN interface supports PPP encapsulation the IPv6 CE router MUST support IPv6 over PPP [RFC5072].
- WLL-3: If the WAN interface supports PPP encapsulation, in a dual-stack environment with IPCP and IPV6CP running over one PPP logical channel, the NCPs MUST be treated as independent of each other and start and terminate independently.

Address assignment requirements:

- WAA-1: The IPv6 CE router MUST support SLAAC [RFC4862].
- WAA-2: The IPv6 CE router MUST follow the recommendation in [RFC5942]. and in particular the handling of the L-flag in the Router Advertisement Prefix Information Option.
- WAA-3: The IPv6 CE router MUST support DHCPv6 [RFC3315] client behavior.
- WAA-4: The IPv6 CE router MUST be able to support the following DHCPv6 options: IA_NA, Reconfigure Accept [RFC3315], DNS_SERVERS [RFC3646].
- WAA-5: The IPv6 CE router SHOULD support the DHCPv6 SNTP option [RFC4075] and the Information Refresh Time Option [RFC4242].
- WAA-6: If the IPv6 CE router receives an RA message (described in [RFC4861]) with the M-flag set to 1, the IPv6 CE router MUST do DHCPv6 address assignment (request an IA_NA option).
- WAA-7: If the IPv6 CE router is unable to assign address(es) through SLAAC it MAY do DHCPv6 address assignment (request an IA_NA) even if the M-flag is set to 0.
- WAA-8: If the IPv6 CE router does not acquire global IPv6 address(es) from either SLAAC or DHCPv6, then it MUST create global IPv6 address(es) from its delegated prefix(es) and configure those on one of its internal virtual network interfaces.
- WAA-9: As a router the IPv6 CE router MUST follow the weak host model [RFC1122]. When originating packets out an interface it will use a source address from another of its interfaces if the outgoing interface does not have an address of suitable scope.

Prefix Delegation requirements:

- WPD-1: The IPv6 CE router MUST support DHCPv6 prefix delegation requesting router behavior as specified in [RFC3633] (IA_PD option).
- WPD-2: The IPv6 CE router MAY indicate as a hint to the delegating router the size of the prefix it requires. If so, it MUST ask for a prefix large enough to assign one /64 for each of its interfaces rounded up to the nearest nibble and MUST be configurable to ask for more.

- WPD-3: The IPv6 CE router MUST be prepared to accept a delegated prefix size different from what is given in the hint. If the delegated prefix is too small to address all of its interfaces, the IPv6 CE router SHOULD log a system management error.
- WPD-4: The IPv6 CE router MUST always initiate DHCPv6 prefix delegation, regardless of the M and O-flags in a received Router Advertisement message.
- WPD-5: If the IPv6 CE Router initiates DHCPv6 before receiving a Router Advertisement it MUST also request an IA_NA option in DHCPv6.
- WPD-6: If the delegated prefix(es) are aggregate route(s) of multiple, more-specific routes, the IPv6 CE router MUST discard packets that match the aggregate route(s), but not any of the more-specific routes. In other words, the next-hop for the aggregate route(s) should be the null destination. This is necessary to prevent forwarding loops when some addresses covered by the aggregate are not reachable [RFC4632].
- (a) The IPv6 CE router SHOULD send an ICMPv6 Destination Unreachable according to section 3.1 [RFC4443] back to the source of the packet, if the packet is to be dropped due to this rule.
- WPD-7: If the IPv6 CE router requests both an IA_NA and an IA_PD in DHCPv6, it MUST accept an IA_PD in DHCPv6 Advertise/Reply messages, even if the message does not contain any addresses.
- WPD-8: By default an IPv6 CE router MUST NOT initiate any dynamic routing protocol on its WAN interface.

4.3. LAN Side Configuration

The IPv6 CE router distributes configuration information obtained during WAN interface provisioning to IPv6 hosts and assists IPv6 hosts in obtaining IPv6 addresses. It also supports connectivity of these devices in the absence of any working WAN interface.

An IPv6 CE router is expected to support an IPv6 end-user network and IPv6 hosts that exhibit the following characteristics:

1. Link-local addresses may be insufficient for allowing IPv6 applications to communicate with each other in the end-user network. The IPv6 CE router will need to enable this

communication by providing globally-scoped unicast addresses or ULAs [RFC4193] whether or not WAN connectivity exists.

2. IPv6 hosts should be capable of using SLAAC and may be capable of using DHCPv6 for acquiring their addresses.
3. IPv6 hosts may use DHCPv6 for other configuration information, such as the DNS_SERVERS option for acquiring DNS information.

Unless otherwise specified, the following requirements apply to the IPv6 CE router's LAN interfaces only.

ULA requirements:

- ULA-1: The IPv6 CE router SHOULD be capable of generating a ULA prefix [RFC4193].
- ULA-2: A IPv6 CE router with a ULA prefix, MUST maintain this consistently across reboots.
- ULA-3: The value of the ULA prefix SHOULD be user configurable.
- ULA-4: By default the IPv6 CE router MUST act as a site border router according to section 4.3 of [RFC4193] and filter packets with Local IPv6 source or destination addresses accordingly.
- ULA-5: An IPv6 CE router MUST NOT advertise itself as a default router with Router Lifetime greater than zero whenever all of its configured and delegated prefixes are ULA prefixes.

LAN requirements:

- L-1: The IPv6 CE router MUST support router behavior according to Neighbor Discovery for IPv6 [RFC4861].
- L-2: The IPv6 CE router MUST assign a separate /64 from its delegated prefix(es) (and ULA prefix if configured to provide ULA addressing) for each of its LAN interfaces.
- L-3: An IPv6 CE router MUST advertise itself as a router for the delegated prefix(es) (and ULA prefix if configured to provide ULA addressing) using the "Route Information Option" specified in section 2.3 of [RFC4191]. This advertisement is independent of having IPv6 connectivity on the WAN interface or not.

- L-4: An IPv6 CE router MUST NOT advertise itself as a default router with a Router Lifetime [RFC4861] greater than zero if it has no prefixes configured or delegated to it.
- L-5: The IPv6 CE router MUST make each LAN interface an advertising interface according to [RFC4861].
- L-6: In Router Advertisements messages, the Prefix Information Option's A and L-flags MUST be set to 1 by default.
- L-7: The A and L-flags setting SHOULD be user configurable.
- L-8: The IPv6 CE router MUST support a DHCPv6 server capable of IPv6 address assignment according to [RFC3315] OR a stateless DHCPv6 server according to [RFC3736] on its LAN interfaces.
- L-9: Unless the IPv6 CE router is configured to support the DHCPv6 IA_NA option, it SHOULD set M=0 and O=1 in its Router Advertisement messages [RFC4861].
- L-10: The IPv6 CE router MUST support providing DNS information in the DHCPv6 DNS_SERVERS and DOMAIN_LIST options [RFC3646].
- L-11: The IPv6 CE router SHOULD support providing DNS information in Router Advertisement RDNSS and DNSSL options as specified in [RFC6106].
- L-12: The IPv6 CE router SHOULD make available a subset of DHCPv6 options (as listed in section 5.3 of [RFC3736]) received from the DHCPv6 client on its WAN interface to its LAN side DHCPv6 server.
- L-13: If the delegated prefix changes, i.e. the current prefix is replaced with a new prefix without any overlapping time period, then the IPv6 CE router MUST immediately advertise the old prefix with a preferred lifetime of 0 and a valid lifetime of 2 hours (which must be decremented in real time) in a Router Advertisement message.
- L-14: The IPv6 CE router MUST send an ICMP Destination Unreachable Message, code 5 (Source address failed ingress/egress policy) for packets forwarded to it using an address from a prefix which has been deprecated.

4.4. Security Considerations

It is considered a best practice to filter obviously malicious traffic (e.g. spoofed packets, "martian" addresses, etc.). Thus, the IPv6 CE router ought to support basic stateless egress and ingress filters. The CE router is also expected to offer mechanisms to filter traffic entering the customer network; however, the method by which vendors implement configurable packet filtering is beyond the scope of this document.

Security requirements:

- S-1: The IPv6 CE router SHOULD support [I-D.ietf-v6ops-cpe-simple-security]. In particular, the IPv6 CE router SHOULD support functionality sufficient for implementing the set of recommendations in [I-D.ietf-v6ops-cpe-simple-security] section 4. This document takes no position on whether such functionality is enabled by default or mechanisms by which users would configure it.
- S-2: The IPv6 CE router MUST support ingress filtering in accordance with [RFC2827] (BCP 38)

5. Acknowledgements

Thanks to the following people (in alphabetical order) for their guidance and feedback:

Mikael Abrahamsson, Tore Anderson, Merete Asak, Scott Beuker, Mohamed Boucadair, Rex Bullinger, Brian Carpenter, Lorenzo Colitti, Remi Denis-Courmont, Gert Doering, Alain Durand, Katsunori Fukuoka, Tony Hain, Thomas Herbst, Kevin Johns, Erik Kline, Stephen Kramer, Victor Kuarsingh, Francois-Xavier Le Bail, David Miles, Arifumi Matsumoto, Shin Miyakawa, Jean-Francois Mule, Michael Newbery, Carlos Pignataro, John Pomeroy, Antonio Querubin, Teemu Savolainen, Matt Schmitt, Hiroki Sato, David Thaler, Mark Townsley, Bernie Volz, James Woodyatt, Dan Wing and Cor Zwart

This draft is based in part on CableLabs' eRouter specification. The authors wish to acknowledge the additional contributors from the eRouter team:

Ben Bekele, Amol Bhagwat, Ralph Brown, Eduardo Cardona, Margo Dolas, Toerless Eckert, Doc Evans, Roger Fish, Michelle Kuska, Diego Mazzola, John McQueen, Harsh Parandekar, Michael Patrick, Saifur Rahman, Lakshmi Raman, Ryan Ross, Ron da Silva, Madhu Sudan, Dan Torbet and Greg White

6. Contributors

The following people have participated as co-authors or provided substantial contributions to this document: Ralph Droms, Kirk Erichsen, Fred Baker, Jason Weil, Lee Howard, Jean-Francois Tremblay, Yiu Lee, John Jason Brzozowski and Heather Kirksey.

7. IANA Considerations

This memo includes no request to IANA.

8. References

8.1. Normative References

- [I-D.ietf-v6ops-cpe-simple-security]
Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment for Providing Residential IPv6 Internet Service", draft-ietf-v6ops-cpe-simple-security-16 (work in progress), October 2010.
- [RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 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.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC3646] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, December 2003.

- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.
- [RFC4075] Kalusivalingam, V., "Simple Network Time Protocol (SNTP) Configuration Option for DHCPv6", RFC 4075, May 2005.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, November 2005.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC4242] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, November 2005.
- [RFC4294] Loughney, J., "IPv6 Node Requirements", RFC 4294, April 2006.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, August 2006.
- [RFC4779] Asadullah, S., Ahmed, A., Popoviciu, C., Savola, P., and J. Palet, "ISP IPv6 Deployment Scenarios in Broadband Access Networks", RFC 4779, January 2007.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.
- [RFC5072] S.Varada, Haskins, D., and E. Allen, "IP Version 6 over

PPP", RFC 5072, September 2007.

[RFC5942] Singh, H., Beebee, W., and E. Nordmark, "IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes", RFC 5942, July 2010.

[RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, November 2010.

8.2. Informative References

[I-D.ietf-behave-v6v4-framework]
Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation",
draft-ietf-behave-v6v4-framework-10 (work in progress),
August 2010.

[I-D.ietf-v6ops-multihoming-without-nat66]
Troan, O., Miles, D., Matsushima, S., Okimoto, T., and D. Wing, "IPv6 Multihoming without Network Address Translation",
draft-ietf-v6ops-multihoming-without-nat66-00 (work in progress), December 2010.

[I-D.wing-v6ops-happy-eyeballs-ipv6]
Wing, D. and A. Yourtchenko, "Happy Eyeballs: Trending Towards Success with Dual-Stack Hosts",
draft-wing-v6ops-happy-eyeballs-ipv6-01 (work in progress), October 2010.

[UPnP-IGD]
UPnP Forum, "Universal Plug and Play (UPnP) Internet Gateway Device (IGD)", November 2001,
<<http://www.upnp.org/standardizeddcps/igd.asp>>.

Authors' Addresses

Hemant Singh
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719
USA

Phone: +1 978 936 1622
Email: shemant@cisco.com
URI: <http://www.cisco.com/>

Wes Beebee
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719
USA

Phone: +1 978 936 2030
Email: wbeebee@cisco.com
URI: <http://www.cisco.com/>

Chris Donley
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
USA

Email: c.donley@cablelabs.com

Barbara Stark
AT&T
725 W Peachtree St
Atlanta, GA 30308
USA

Email: barbara.stark@att.com

Ole Troan (editor)
Cisco Systems, Inc.
Veversmauet 8
N-5017 BERGEN,
Norway

Email: ot@cisco.com

Homenet
Internet-Draft
Intended status: Informational
Expires: May 3, 2012

E. Vyncke
A. Yourtchenko
M. Townsley
Cisco Systems
October 31, 2011

Advanced Security for IPv6 CPE
draft-vyncke-advanced-ipv6-security-03.txt

Abstract

This document describes how an IPv6 residential Customer Premise Equipment (CPE) can leverage modern security techniques to have strong security, while retaining as much of the end-to-end reachability of IPv6 as possible.

It is a re-submission in the framework of the HOMENET working group. The reputation part of this document should leverage the work done in the REPUTE working group of the Application are.

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 May 3, 2012.

Copyright Notice

Copyright (c) 2011 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.

Table of Contents

1. Introduction	3
2. Threats	3
3. Overview	4
3.1. Rules for Security Policy	5
3.2. Security Analysis	6
4. IANA Considerations	7
5. Security Considerations	7
6. Acknowledgements	8
7. References	8
7.1. Normative References	8
7.2. Informative References	8
Authors' Addresses	8

1. Introduction

Internet access in residential IPv4 deployments generally consist of a single IPv4 address provided by the service provider for each home. Residential CPE then translates the single address into multiple private addresses allowing more than one device in the home, but at the cost of losing end-to-end reachability. IPv6 allows all devices to have a unique, global, IP address, restoring end-to-end reachability directly between any device. Such reachability is very powerful for ubiquitous global connectivity, and is often heralded as one of the significant advantages to IPv6 over IPv4. Despite this, concern about exposure to inbound packets from the IPv6 Internet (which would otherwise be dropped by the address translation function if they had been sent from the IPv4 Internet) remain. This document describes firewall functionality for an IPv6 CPE which departs from the "simple security" model described in [RFC6092]. The intention is to provide an example of a security model which allows most traffic, including incoming unsolicited packets and connections, to traverse the CPE unless the CPE identifies the traffic as potentially harmful based on a set of signatures (and other correlation data and heuristics) that are kept up to date on a regular basis. The computational resources necessary to support some, not all, functionalities of this model are likely more intensive than those described in [RFC6092], but are easily within the realm of what is commonly available in 2011 on medium to high-end network based firewall systems for small and medium businesses, or host-based commercial firewalls that run on laptop and desktop PCs. This set of techniques is also known as Universal Threat Mitigation (UTM).

2. Threats

For a typical residential network connected to the Internet over a broadband connection, the threats can be classified into:

- o denial of service by packet flooding: overwhelming either the access bandwidth or the bandwidth of a slower link in the residential network (like a slow home automation network) or the CPU power of a slow IPv6 host (like networked thermostat or any other sensor type nodes)
- o denial of service by service requests: like sending print jobs from the Internet to an ink jet printer until the ink cartridge is empty or like filing some file server with junk data
- o unauthorized use of services: like accessing a webcam or a file server which are open to anonymous access within the residential network but should not be accessed freely and anonymously from

outside of the home network

- o exploiting a vulnerability in the host in order to get access to data or to execute some arbitrary code in the attacked host. Exploitation can be further divided in two classes:
 1. day-0 attack when this attack has never been seen before (hence nothing can really detect it) and
 2. day+n attack where this attack is known and can be detected by the use of an attack signature
- o trojanized host (belonging to a Botnet) can communicate via a covert channel to its master and launch attacks to Internet targets.

3. Overview

The basic goal is to provide an adaptive security policy which aims to block known harmful traffic and allow the rest, restoring as much of end-to-end communication as possible. In addition, new protocols may evolve and be deployed over time; only if they become a threat vector does the CPE firewall receive a signature update (including dynamic correlation data) to classify and block them. This is in direct contrast to [RFC6092], which requires built-in knowledge of a number of protocols, or requires Internet communication to be limited to a handful of protocols that the CPE understands how to process.

- o Intrusion Prevention System (IPS) is a signature-based technology which inspects a pre-defined set of protocols at all layers (from layer-3 to layer-7) and uses a vast set of heuristics to detect attacks within one or several flow. Upon detection, the flow is terminated and an event is logged for further optional auditing. As exploits are added every day, the signature database must be updated daily and is usually quite large (more than 100 MB). This requires both large local storage (large flash or even a hard disk) and a subscription to an update service.
- o Reputation database is a centralized database which gives a reputation score to any IPv6 address (or prefix). The score varies from untrusted to trusted. Untrusted IPv6 addresses are typically addresses of a well-known attacker or from a Botnet member or from an ISP with a poor track of security... Protocols exist to dynamically request a reputation (based on DNS or HTTP). This usually requires a subscription. Note: in IPv6 the reputation database concept is still in its infancy, for example, little experience exists on the scope of the reputation: a host

/128, a LAN prefix /64 or a delegated prefix size of /56 or /48...

- o Local correlation uses another set of heuristics (like TCP distribution of Initial Sequence Number or used TCP ports or protocol handshake banners) to assert the variety of local hosts (namely operating system (OS) version and set of application) and raise or decrease the importance of a specific attack signature. For example, if the OS of host A is OS-A, then there is no point to inspect traffic to or from host A for attacks which are only relevant to OS-B.
- o Global correlation leverage all IPS distributed on the Internet to build the reputation database as well as changing the relevance of an IPS signature (for example, a propagating worm will trigger a lot of identical signatures on several IPS, this should raise the relevance of a specific signature up to the point of blocking all inbound/outbound connections on a specific layer-4 port).

The above techniques are common in the large network where budget is enough to buy firewalls, IPS and subscribe to signature or reputation source. The authors of this document believes that competition and Moore's law will make the set of those techniques (commonly referred to as 'Universal Threat Mitigation') affordable for consumer space.

3.1. Rules for Security Policy

These are an example set of rules to be applied. Each would normally be configurable, either by the user directly or on behalf of the user by a subscription service. The default preferred state hasn't been listed, though it is expected that all rules would be on by default.

If we named all hosts on the residential side of the CPE as 'inside' and all hosts on the Internet as 'outside', then the behavior of the CPE is described by a small set of rules:

1. Rule RejectBogon: apply unicast reverse path forwarding (RPF) checks (anti-spoofing) for all inbound and outbound traffic (implicitly blocking link-local and ULA in the same shot)
2. Rule BlockBadReputation: block all inbound and outbound packets whose outside IPv6 address has a bad reputation score
3. Rule AllowReturn: inspect all outbound traffic and allow the return traffic matching the states (5-tuple + TCP sequence number or any layer-4 state), apply IPS on the outbound (to block Botnet) and inbound (to block malicious/cracked servers which could inject malware) with IPS. If the protocol is not supported/recognized by the IPS, accept it anyway.

4. Rule AllowToPublicDnsHost: allow all inbound traffic to any inside address which is listed in the public DNS with a AAAA record (this requires that the CPE/RG can do a zone transfer, i.e., that the CPE/RG appears like a secondary name server), all inbound traffic is also inspected with IPS. If the protocol is not supported/recognized by the IPS, accept it anyway.
5. Rule ProtectLocalOnly: block all inbound traffic to any inside address as long as the inside address has never sent a packet to the outside. The intent is to protect local-only devices like thermostat or printers. Most (if not all) hosts expecting inbound connections have to send a couple of outbound packets to the outside (registration, DNS request, ...). This is the usual IPv4 firewall behavior augmented with IPS and reputation
6. Rule CryptoIntercept: at the exception of IPsec, all inbound connections that are encrypted (notably TLS [RFC5246]) must be intercepted (this is terminated by the CPE that will present its own self-signed certificate to the remote party which should have installed the CPE self-signed certificate in a secure way in its trust anchors store) in order to allow for further inspection. The decrypted flow is then passed again through those rules and encrypted again before being forwarded to the local host. This is actually a Man-in-the-Middle attack done for a good reason: protect the naive residential user. Of course, documentation and GUI MUST be provided to educate the user and help him/her to understand how to do it in a secure way. Note: this technique is also used nowadays by large enterprise web proxies with the self-signed certificate being securely distributed to all clients.
7. Rule ParanoidOpeness: allow all unsolicited inbound connections rate limited to protect against port and address scanning attacks or overloading devices or slow links within the home. The connection MUST be inspected by the IPS engine. If the connection is anonymous or using a default password (like connecting to a webcam as a guest), then the flow SHOULD be dropped. If the IPS detects an attack, then the flow MUST be closed. If the protocol is not recognized as supported by the IPS, the flow MAY be allowed.

3.2. Security Analysis

This proposal of 'paranoid openness' stops the following attacks:

- o unauthorized use of services/denial of service: because all anonymous access to inside servers are blocked.

- o Denial of services on low bandwidth or low CPU inside hosts IFF those hosts never access the Internet
- o Exploiting of a day+1 attack, those attacks are blocked with the IPS signature and address reputation database

The CryptoIntercept part can also be leveraged as a small Certification Authority (CA) that could generate RSA key pairs and X.509 certificates at the CPE/RG owner's request. Those key pairs and certificates can then be given to trusted devices or users (like the owner's laptop so that he/she could easily and safely connect from the outside).

This proposal cannot help with the following attacks:

- o flooding the access link to the Internet, this is exactly the same as with the old layers-3/4 firewall approach as only the ISP can effectively stop the flooding of the CE-PE link;
- o weak password on inside services, of course the IPS component will detect multiple failed attempts (dictionary attack) and report the offender to the Global Correlation system;
- o exploiting of day-0 attack: until now, these day-0 attacks are caused either by rapidly propagating worms (then the global correlation of unusual traffic pattern will raise an alert and block the traffic after a couple of hundred's of successful attacks) or by targeted attacks against high-profile targets (like Government or banks or ;..) which should be protected by conventional less open security policies;
- o exploiting a vulnerability in a rare or new protocol (not yet supported by the IPS), this case will probably never occur on a wide scale in a residential use of Internet.

4. IANA Considerations

There are no extra IANA consideration for this document.

5. Security Considerations

All security considerations have been done in the Security Analysis Section 3.2.

It is also advisable that the inbound rate limiter system could be added to the [RFC6092] as it is light and does not depend on a

centralized policy server.

6. Acknowledgements

Many thanks to Ole Troan, Stuart Cheshire, Dave Oran and Eliot Lear for the review of the -00 version and to Ron Bonica, Sam Hartmans, Lee Howard, Greg Lebovitz, Jordi Palet, Tina Tsou and others for their comments during and after the first presentation at the Hiroshima IETF meeting in November 2009.

A previous IETF work has similar ideas
[I-D.palet-v6ops-ipv6security].

7. References

7.1. Normative References

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

7.2. Informative References

- [I-D.palet-v6ops-ipv6security]
Palet, J., Vives, A., Martinez, G., and A. Gomez, "IPv6 distributed security requirements",
draft-palet-v6ops-ipv6security-02 (work in progress),
February 2005.
- [RFC2993] Hain, T., "Architectural Implications of NAT", RFC 2993,
November 2000.
- [RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092,
January 2011.

Authors' Addresses

Eric Vyncke
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2 778 4677
Email: evyncke@cisco.com

Andrew Yourtchenko
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2 704 5494
Email: ayourtch@cisco.com

Mark Townsley
Cisco Systems
11, Rue Camille Desmoulins
Issy Les Moulineaux 92782
France

Phone: +33 15 804 3483
Email: townsley@cisco.com

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 28, 2011

H. Singh
W. Beebee
Cisco Systems, Inc.
C. Donley
CableLabs
B. Stark
AT&T
O. Troan, Ed.
Cisco Systems, Inc.
October 25, 2010

Advanced Requirements for IPv6 Customer Edge Routers
draft-wbeebee-v6ops-ipv6-cpe-router-bis-04

Abstract

This document continues the work undertaken by the IPv6 CE Router Phase I work in the IETF v6ops Working Group. Advanced requirements or Phase II work is covered in this document.

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 28, 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.

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
2. Terminology	3
3. Conceptual Configuration Variables	4
4. Architecture	4
5. Advanced Features and Feature Requirements	6
5.1. DNS	6
5.2. Multicast Behavior	6
5.3. ND Proxy	7
5.4. Prefix Delegation on LAN interface(s) (More details are TBD)	8
5.5. Routed network behavior(General Cases TBD)	8
5.6. Transition Technologies Support	9
5.6.1. Dual-Stack(DS)-Lite	9
5.6.2. 6rd	10
5.6.3. Transition Technologies Coexistence	10
5.7. Quality Of Service	11
5.8. Unicast Data Forwarding	11
5.9. ZeroConf	11
6. Security Considerations	11
7. Acknowledgements	11
8. Contributors	12
9. IANA Considerations	12
10. References	12
10.1. Normative References	12
10.2. Informative References	15
Authors' Addresses	15

1. Introduction

This document defines Advanced IPv6 features for a residential or small office router referred to as an IPv6 CE router. Typically these routers also support IPv4. The IPv6 End-user Network Architecture for such a router is described in [I-D.ietf-v6ops-ipv6-cpe-router]. This version of the document includes the requirements for Advanced features.

1.1. 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].

2. Terminology

End-user Network	one or more links attached to the IPv6 CE router that connect IPv6 hosts.
IPv6 Customer Edge router	a node intended for home or small office use which forwards IPv6 packets not explicitly addressed to itself. The IPv6 CE router connects the end-user network to a service provider network.
IPv6 host	any device implementing an IPv6 stack receiving IPv6 connectivity through the IPv6 CE router
LAN interface	an IPv6 CE router's attachment to a link in the end-user network. Examples are Ethernet (simple or bridged), 802.11 wireless or other LAN technologies. An IPv6 CE router may have one or more network layer LAN Interfaces.
Service Provider	an entity that provides access to the Internet. In this document, a Service Provider specifically offers Internet access using IPv6, and may also offer IPv4 Internet access. The Service Provider can provide such access over a variety of different transport methods such as DSL, cable, wireless, and others.

WAN interface an IPv6 CE router's attachment to a link used to provide connectivity to the Service Provider network; example link technologies include Ethernets (simple or bridged), PPP links, Frame Relay, or ATM networks as well as Internet-layer (or higher-layer) "tunnels", such as tunnels over IPv4 or IPv6 itself.

3. Conceptual Configuration Variables

The CE Router maintains such a list of conceptual optional configuration variables.

1. Enable an IGP on the LAN.

4. Architecture

This document extends the architecture described in [I-D.ietf-v6ops-ipv6-cpe-router] to cover a strictly larger set of operational scenarios. In particular, QoS, multicast, DNS, routed network in the home, transition technologies, and conceptual configuration variables. This document also extends the model described in [I-D.ietf-v6ops-ipv6-cpe-router] to a two router topology where the two routers are connected back-to-back (the LAN of one router is connected to the WAN of the other router). This topology is depicted below:

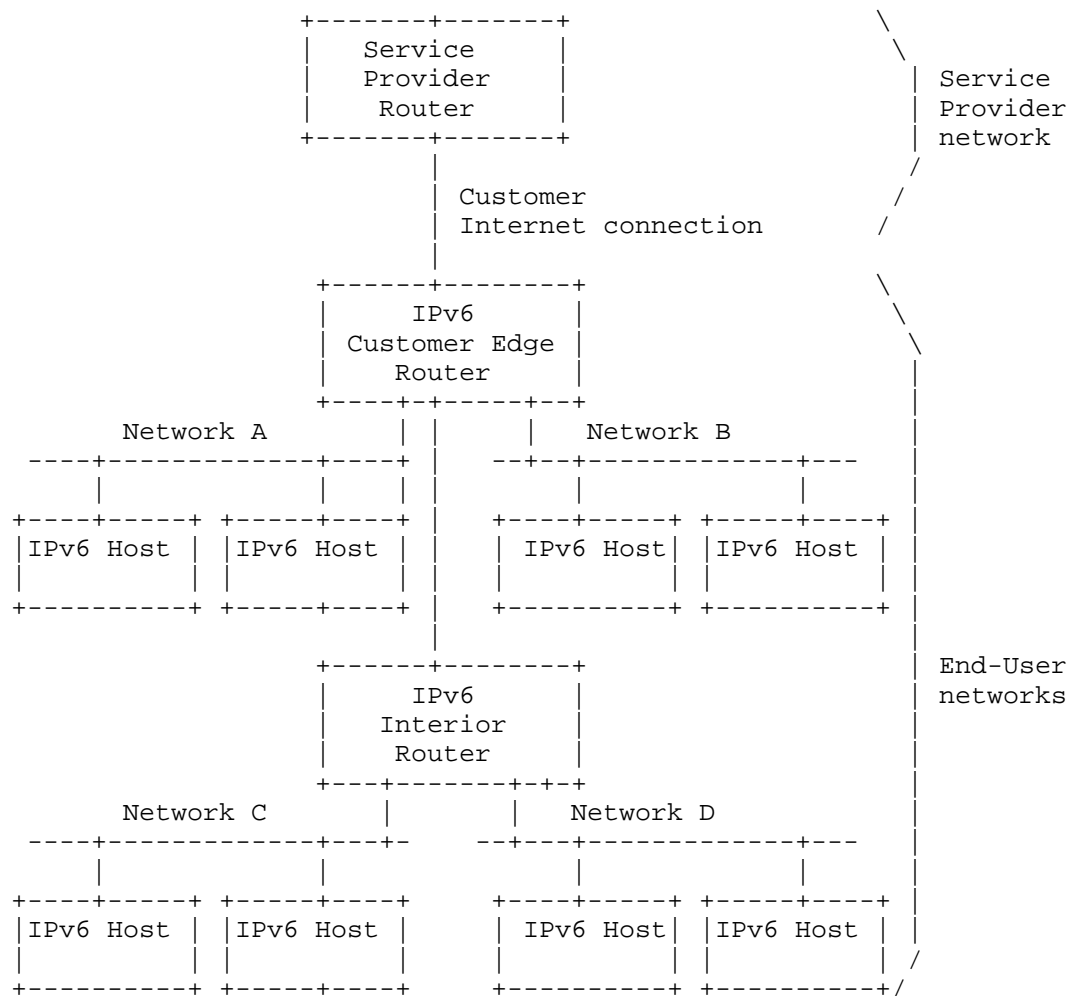


Figure 1.

For DNS, the operational expectation is that the end-user would be able to access home hosts from the home using DNS names instead of more cumbersome IPv6 addresses. Note that this is distinct from the requirement to access home hosts from outside the home.

End-users are expected to be able to receive multicast video in the home without requiring the CE router to include the cost of supporting full multicast routing protocols.

5. Advanced Features and Feature Requirements

The IPv6 CE router will need to support connectivity to one or more access network architectures. This document describes an IPv6 CE router that is not specific to any particular architecture or Service Provider, and supports all commonly used architectures.

5.1. DNS

- D-1: For local DNS queries for configuration, the CE Router may include a DNS server to handle local queries. Non-local queries can be forwarded unchanged to a DNS server specified in the DNS server DHCPv6 option. The CE Router may also include DNS64 functionality which is specified in [I-D.bagnulo-behave-dns64].
- D-2: The local DNS server MAY also handle renumbering from the Service Provider provided prefix for local names used exclusively inside the home (the local AAAA and PTR records are updated). This capability provides connectivity using local DNS names in the home after a Service Provider renumbering. A CE Router MAY add local DNS entries based on dynamic requests from the LAN segment(s). The protocol to carry such requests from hosts to the CE Router is yet to be described.

5.2. Multicast Behavior

This section is only applicable to a CE Router with at least one LAN interface. A host in the home is expected to receive multicast video. Note the CE Router resides at edge of the home and the Service Provider, and the CE Router has at least one WAN connection for multiple LAN connections. In such a multiple LAN to a WAN topology at the CE Router edge, it is not necessary to run a multicast routing protocol and thus MLD Proxy as specified in [RFC4605] can be used. The CE Router discovers the hosts via a MLDv2 Router implementation on a LAN interface. A WAN interface of the CE Router interacts with the Service Provider router by sending MLD Reports and replying to MLD queries for multicast Group memberships for hosts in the home.

The CE router SHOULD implement MLD Proxy as specified in [RFC4605]. For the routed topology shown in Figure 1, each router implements a MLD Proxy. If the CE router implements MLD Proxy, the requirements on the CE Router for MLD Proxy are listed below.

WAN requirements, MLD Proxy:

WMLD-1: Consistent with [RFC4605], the CE router MUST NOT implement the router portion of MLDv2 for the WAN interface.

LAN requirements, MLD Proxy:

LMMLD-1: The CPE Router MUST follow the model described for MLD Proxy in [RFC4605] to implement multicast.

LMMLD-2: Consistent with [RFC4605], the LAN interfaces on the CPE router MUST NOT implement an MLDv2 Multicast Listener.

LAN requirements:

LM-1: If the CE Router has bridging configured between the LAN interfaces, then the LAN interfaces MUST support snooping of MLD [RFC3810] messages.

5.3. ND Proxy

LAN requirements:

LNDP-1: If the CE Router has only one /64 prefix to be used across multiple LAN interfaces and the CE Router supports any two LAN interfaces that cannot bridge data between them because the two interfaces have disparate MAC layers, then the CE Router MUST support Proxying Neighbor Advertisements as specified in Section 7.2.8 of [RFC4861]. If any two LAN interfaces support bridging between the interfaces, then Proxying Neighbor Advertisements is not necessary between the two interfaces. Legacy 3GPP networks have the following requirements:

1. No DHCPv6 prefix is delegated to the CE Router.
2. Only one /64 is available on the WAN link.
3. The link types between the WAN interface and LAN interface(s) are disparate and, therefore, can't be bridged.
4. No NAT66 is to be used.
5. Each LAN interface needs global connectivity.
6. Uses SLAAC to configure LAN interface addresses.

For these legacy 3GPP networks, the CPE Router MUST support ND Proxy between the WAN and LAN interface(s). If a CE

Router will never be deployed in an environment with these characteristics, then ND Proxy is not necessary.


5.4. Prefix Delegation on LAN interface(s) (More details are TBD)

This section is only applicable to a CE Router with at least one LAN interface. The LAN interface(s) are delegated prefixes subnetted from the delegated prefix acquired by the WAN interface and the ULA prefix. After the CE router has assigned prefixes for all of its internally defined needs (its interfaces and any other purposes defined in its internal logic), any leftover prefixes are available for delegation. Any automated prefix delegation mechanism is TBD.

5.5. Routed network behavior (General Cases TBD)

CPE Router Behavior in a routed network:

R-1: One example of the CPE Router use in the home is shown below. The home has a broadband modem combined with a CPE Router, all in one device. The LAN interface of the device is connected to another standalone CPE Router that supports a wireless access point. To support such a network, this document recommends using prefix delegation of the prefix obtained either via IA_PD from WAN interface or a ULA from the LAN interface. The network interface of the downstream router may obtain an IA_PD via stateful DHCPv6. If the CPE router supports the routed network through automatic prefix delegation, the CPE router MUST support a DHCPv6 server or DHCPv6 relay agent. Further, if an IA_PD is used, the Service Provider or user MUST allocate an IA_PD or ULA prefix short enough to be delegated and subsequently used for SLAAC. Therefore, a prefix length shorter than /64 is needed. The CPE Router MAY support and IGP in the home network.



WAP = Wireless Access Point

Figure 2.

5.6. Transition Technologies Support

5.6.1. Dual-Stack(DS)-Lite

Even as users migrate from IPv4 to IPv6 addressing, a significant percentage of Internet resources and content will remain accessible only through IPv4. Also, many end-user devices will only support IPv4. As a consequence, Service Providers require mechanisms to allow customers to continue to access content and resources using IPv4 even after the last IPv4 allocations have been fully depleted. One technology that can be used for IPv4 address extension is DS-Lite.

DS-Lite enables a Service Provider to share IPv4 addresses among multiple customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) tunneling and Carrier Grade NAT. More specifically, Dual-Stack-Lite encapsulates IPv4 traffic inside an IPv6 tunnel at the IPv6 CE Router and sends it to a Service Provider Address Family Translation Router (AFTR). Configuration of the IPv6 CE Router to support IPv4 LAN traffic is outside the scope of this document.

The IPv6 CE Router SHOULD implement DS-Lite functionality as specified in [I-D.ietf-softwire-dual-stack-lite].

WAN requirements:

- DLW-1: To facilitate IPv4 extension over an IPv6 network, if the CE Router supports DS-Lite functionality, the CE Router WAN interface MUST implement a B4 Interface as specified in [I-D.ietf-softwire-dual-stack-lite].
- DLW-2: If the IPv6 CE Router implements DS-Lite functionality, the CE Router MUST support using a DS-Lite DHCPv6 option [I-D.ietf-softwire-ds-lite-tunnel-option] to configure the DS-Lite tunnel. The IPv6 CE Router MAY use other mechanisms to configure DS-Lite parameters. Such mechanisms are outside the scope of this document.
- DLW-3: IPv6 CE Router MUST NOT perform IPv4 Network Address Translation (NAT) on IPv4 traffic encapsulated using DS-Lite.
- DLW-4: If the IPv6 CE Router is configured with a non-RFC1918 IPv4 address on its WAN interface, the IPv6 CE Router MUST disable the DS-Lite B4 element.

DLW-5: If DS-Lite is operational on the IPv6 CE Router, multicast data MUST NOT be sent on any DS-Lite tunnel.

5.6.2. 6rd

The IPv6 CE Router can be used to offer IPv6 service to a LAN, even when the WAN access network only supports IPv4. One technology that supports IPv6 service over an IPv4 network is IPv6 Rapid Deployment (6rd). 6rd encapsulates IPv6 traffic from the end user LAN inside IPv4 at the IPv6 CE Router and sends it to a Service Provider Border Relay (BR). The IPv6 CE Router calculates a 6rd delegated IPv6 prefix during 6rd configuration, and sub-delegates the 6rd delegated prefix to devices in the LAN.

The IPv6 CE Router SHOULD implement 6rd functionality as specified in [RFC5969].

6rd requirements:

6RD-1: If the IPv6 CE Router implements 6rd functionality, the CE Router WAN interface MUST support at least one 6rd Virtual Interface and 6rd CE functionality as specified in [RFC5969].

6RD-2: If the IPv6 CE Router implements 6rd CE functionality, it MUST support using the 6rd DHCPv4 Option (212) for 6rd configuration. The IPv6 CE Router MAY use other mechanisms to configure 6rd parameters. Such mechanisms are outside the scope of this document.

6RD-3: If 6rd is operational on the IPv6 CE Router, multicast data MUST NOT be sent on any 6rd tunnel.

5.6.3. Transition Technologies Coexistence

Run the following four in parallel to provision CPE router connectivity to the Service Provider:

1. Initiate IPv4 address acquisition.
2. Initiate IPv6 address acquisition as specified by [I-D.ietf-v6ops-ipv6-cpe-router].
3. If 6rd is provisioned, initiate 6rd.
4. If DS-Lite is provisioned, initiate DS-Lite.

The default route for IPv6 through the native physical interface should have preference over the 6rd tunnel interface. The default

route for IPv4 through the native physical interface should have preference over the DS-Lite tunnel interface.

5.7. Quality Of Service

Q-1: The CPE router MAY support differentiated services [RFC2474].

5.8. Unicast Data Forwarding

The null route introduced by the WPD-6 requirement in [I-D.ietf-v6ops-ipv6-cpe-router] has lower precedence than other routes except for the default route.

5.9. ZeroConf

The CE Router MAY support manual configuration via the web using a URL string like `http://router.local` as per multicast DNS (mDNS). Zero-configuration is vendor-dependent.

6. Security Considerations

None.

7. Acknowledgements

Thanks to the following people (in alphabetical order) for their guidance and feedback:

Mikael Abrahamsson, Merete Asak, Scott Beuker, Mohamed Boucadair, Rex Bullinger, Brian Carpenter, Remi Denis-Courmont, Gert Doering, Alain Durand, Katsunori Fukuoka, Tony Hain, Thomas Herbst, Kevin Johns, Stephen Kramer, Victor Kuarsingh, Francois-Xavier Le Bail, David Miles, Shin Miyakawa, Jean-Francois Mule, Michael Newbery, Carlos Pignataro, John Pomeroy, Antonio Querubin, Teemu Savolainen, Matt Schmitt, Hiroki Sato, Mark Townsley, Bernie Volz, James Woodyatt, Dan Wing and Cor Zwart

This draft is based in part on CableLabs' eRouter specification. The authors wish to acknowledge the additional contributors from the eRouter team:

Ben Bekele, Amol Bhagwat, Ralph Brown, Eduardo Cardona, Margo Dolas, Toerless Eckert, Doc Evans, Roger Fish, Michelle Kuska, Diego Mazzola, John McQueen, Harsh Parandekar, Michael Patrick, Saifur Rahman, Lakshmi Raman, Ryan Ross, Ron da Silva, Madhu Sudan, Dan Torbet and Greg White.

8. Contributors

The following people have participated as co-authors or provided substantial contributions to this document: Ralph Droms, Kirk Erichsen, Fred Baker, Jason Weil, Lee Howard, Jean-Francois Tremblay, Yiu Lee, John Jason Brzozowski and Heather Kirksey.

9. IANA Considerations

This memo includes no request to IANA.

10. References

10.1. Normative References

[I-D.bagnulo-behave-dns64]

Bagnulo, M., Sullivan, A., Matthews, P., Beijnum, I., and M. Endo, "DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", draft-bagnulo-behave-dns64-02 (work in progress), March 2009.

[I-D.ietf-softwire-ds-lite-tunnel-option]

Hankins, D. and T. Mrugalski, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual- Stack Lite", draft-ietf-softwire-ds-lite-tunnel-option-05 (work in progress), September 2010.

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

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

[I-D.ietf-v6ops-ipv6-cpe-router]

Singh, H., Beebee, W., Donley, C., Stark, B., and O. Troan, "Basic Requirements for IPv6 Customer Edge Routers", draft-ietf-v6ops-ipv6-cpe-router-07 (work in progress), August 2010.

[I-D.vyncke-advanced-ipv6-security]

Vyncke, E. and M. Townsley, "Advanced Security for IPv6 CPE", draft-vyncke-advanced-ipv6-security-01 (work in progress), March 2010.

[RFC1122] Braden, R., "Requirements for Internet Hosts -

Communication Layers", STD 3, RFC 1122, October 1989.

- [RFC2080] Malkin, G. and R. Minnear, "RIPng for IPv6", RFC 2080, January 1997.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 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.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC3646] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, December 2003.
- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [RFC4075] Kalusivalingam, V., "Simple Network Time Protocol (SNTP) Configuration Option for DHCPv6", RFC 4075, May 2005.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC4242] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, November 2005.

- [RFC4294] Loughney, J., "IPv6 Node Requirements", RFC 4294, April 2006.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, May 2006.
- [RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, August 2006.
- [RFC4779] Asadullah, S., Ahmed, A., Popoviciu, C., Savola, P., and J. Palet, "ISP IPv6 Deployment Scenarios in Broadband Access Networks", RFC 4779, January 2007.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.
- [RFC5072] S.Varada, Haskins, D., and E. Allen, "IP Version 6 over PPP", RFC 5072, September 2007.
- [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.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.

10.2. Informative References

[I-D.ietf-behave-v6v4-framework]

Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", draft-ietf-behave-v6v4-framework-10 (work in progress), August 2010.

[UPnP-IGD]

UPnP Forum, "Universal Plug and Play (UPnP) Internet Gateway Device (IGD)", November 2001, <<http://www.upnp.org/standardizeddcps/igd.asp>>.

Authors' Addresses

Hemant Singh
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719
USA

Phone: +1 978 936 1622
Email: shemant@cisco.com
URI: <http://www.cisco.com/>

Wes Beebee
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719
USA

Phone: +1 978 936 2030
Email: wbeebee@cisco.com
URI: <http://www.cisco.com/>

Chris Donley
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
USA

Email: c.donley@cablelabs.com

Barbara Stark
AT&T
725 W Peachtree St
Atlanta, GA 30308
USA

Email: barbara.stark@att.com

Ole Troan (editor)
Cisco Systems, Inc.
Veversmauet 8
N-5017 BERGEN,
Norway

Email: ot@cisco.com

