

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

H. Singh
W. Beebee
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

v6ops Working Group
Internet-Draft
Intended status: Informational
Expires: March 4, 2011

G. Van de Velde
O. Troan
Cisco Systems
T. Chown
University of Southampton
August 31, 2010

Non-Managed IPv6 Tunnels considered Harmful
<draft-vandavelde-v6ops-harmful-tunnels-01.txt>

Abstract

IPv6 is ongoing and natively being deployed by a growing community and it is important that the quality perception and traffic flows are as optimal as possible. Ideally it would be as good as the IPv4 perceptive experience.

This paper looks into a set of transitional technologies where the actual user has IPv6 connectivity through the means of IPv6-in-IPv4 tunnels. A subset of the available tunnels has the property of being non-managed (i.e. 6to4 [RFC3056] and Teredo [RFC4380]).

While native IPv6 deployments will keep growing it is uncertain or even expected that non-managed IPv6 tunnels will be providing the same user experience and operational quality as managed tunnels or native IPv6 connectivity.

This paper will detail some considerations around non-managed tunnels and will document the harmful element of these for the future growth of networks and the Internet.

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 March 4, 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. Managed Tunnelling Properties	4
3. Tunnel User Experience Views	5
4. Why do non-managed tunnels exist?	5
5. Non-Managed Tunnelling Properties	6
5.1. Performance	6
5.2. Topological Considerations	7
5.3. Operational Provisioning	7
5.4. Operational Troubleshooting	7
5.5. Security	8
5.6. Content Services	8
6. Conclusion	9
7. IANA Considerations	9
8. Security Considerations	9
9. Acknowledgements	9
10. References	9
10.1. Normative References	9
10.2. Informative References	9
Authors' Addresses	10

1. Introduction

While the Internet and networks continue to grow, it is found that the deployment of IPv6 within these networks is an ongoing activity due to global IPv4 address pool depletion. An important aspect is that the quality, availability and security of the IPv6 connectivity is as good as possible, and when possible even more advanced as the IPv4 connectivity.

Historically IETF has been facilitating a variety of technologies and procedures to deploy IPv6 successfully in addition to existing IPv4 connectivity. In general and for the sake of this draft these procedures and technologies can be divided into three major groups: (1) native (dual-stack) IPv6, (2) Tunnelled IPv6 and (3) Translation. While native IPv6 deployments has been steadily growing, the value and the drawbacks of some tunnelling mechanisms can be investigated. Translational techniques provide a total different aspect of considerations and applicability and is beyond the scope of this paper. Transition techniques have been and still are in many cases important for the bootstrapping of IPv6, this paper will look into a range of property aspects of non-managed IPv6 tunnelling techniques. Areas of perverse traffic paths, security considerations, lack of business incentives to run tunnel relays/gateways, black holing and ownership of supportability will be analysed. Finally the paper will conclude that for the growth of IP connectivity, non-managed tunnelling techniques are considered harmful especially for those that want to access applications over the network through pervasive IPv6 connectivity and have no particular interest on how connectivity to the applications is established (IPv4, translation, IPv6, etc...)

2. Managed Tunnelling Properties

A managed tunnel is a tunnel has a few properties supporting the ownership and quality of the tunnel.

When using a managed service, there tends to be an administrative entity which provides quality assurance and can take action if users of the service are experiencing a degraded service. An example would be 6rd tunnels [RFC5969]

In addition there is a general trust awareness and agreement between the user of the managed tunnel service and the provider of the managed tunnel service.

3. Tunnel User Experience Views

The tunnel experience can be divided into three distinct segments: (1) the End-user view, (2) the Enterprise View and (3) the Service Provider View.

The End-user view exists mainly out of two different user profiles. The technical power user and the general user mainly trying to reach their favourite application on the network. The technical power user may have a particular interest to run IPv6 as a transport mechanism, and if his upstream service provider has no native IPv6 connectivity available, then non-managed tunneling mechanisms may provide a solution satisfying to the immediate needs of the technical power user. Alternatively, the general user trying to reach his favourite network application, may have no interest or awareness of his IPv6 usage, particularly when non-managed tunnels are utilized.

The Enterprise View is a more traffic flows and network oriented positioning. When the upstream service provider does not have an IPv6 offer, then the enterprise may start to rely upon a technology as 6to4 [RFC3056]. However this technology has the potential of creating quite perverse traffic paths when user want to reach applications on the Internet. When user would like to reach other 6to4 [RFC3056] users, then more optimized traffic paths, generally following the IPv4 traffic paths are realized

The final view is how a Internet service provider looks into non-managed tunnel usage. A service provider may decide to deploy a 6to4 relay to increase the IPv6 quality of their customers. This a service which require resources (money, maintenance, etc...). Often the 6to4 relay service is not just (always) restricted to only the service providers customers, which as result provides often results in a demotivation to provide quality tunnel relay devices. From a content service provider perspective the usage of non-managed tunnel often results in measurable differences in RTT and reliability in some cases, and hence are reluctant to bring all services to mainstream IPv6 for all users 'just yet'.

4. Why do non-managed tunnels exist?

Non-managed tunnels exist due to a variety of reasons.

Early adopters: people and organisations with a desire to use new and potentially market disrupting technologies and applications may have a desire to use the latest IP even when the upstream provider doesn't have an available service offering.

Lock-step process to implement IPv6: It is not trivial to move a system or an organisation in lock-step towards IPv6 and the aid of tunnels help in this process.

The utilisation of tunnels aid in providing a de-coupling between infrastructure readiness and application readiness, and hence contribute to the development of both elements.

5. Non-Managed Tunnelling Properties

The properties of Non-managed tunnels span many different areas. In this section the properties are analysed and segmented within different areas of impact. In each case the comparison is made between native IPv6 connectivity and connectivity through a non-managed tunnel. A common property of non-managed tunnels is that they often use well-known anycast addresses or other well known addresses and anticipate upon the goodwill of middleware (typically a relay or gateway) device to serve as a tunnel termination point. In some cases, for example a 6to4 relay can be provided by a connected responsible service provider, and hence good quality operation can be guaranteed.

Non-managed tunnels often have asymmetric behaviour. There is an outbound and an inbound connectivity behaviour from the tunnel initiator. It is possible to influence the good quality tunnel behaviour of the outbound connectivity (e.g. by explicit setting of the 6to4 relay), however, influencing good inbound connectivity is often an issue.

5.1. Performance

Deploying a tunnelling mechanism mostly results in encapsulation and de-capsulation efforts. Often this activity has a performance impact on the device, especially when the device does not use hardware acceleration for this functionality. If the performance impact is scoped into the device its lifetime through performance capacity management then the actual impact is predictive. Non-deterministic tunnels tend to have a non-predictive behaviour for capacity, and hence application and network performance is non-predictive. The key reason for this is the decoupling of the capacity management of the tunnel aggregation devices from the capacity desired by users of the aggregation devices.

During initial IPv6 deployment there have been mainly technical savvy people that have been using non-managed tunnel technologies and it has for many been working well. However, if non-managed tunnelling would be deployed in mass and especially when enabled by default by

CPE vendors or host vendors then those aggregation points could become overloaded and result in bad performance. There are a few measures that can be taken, i.e. upgrade the CPU power of the aggregation device or its bandwidth available, however this may not happen without the right motivation for the operator of the aggregation device (i.e. cash flows, reputation, competitive reasons, etc...).

5.2. Topological Considerations

Due to non-managed IPv6 tunnels the traffic flows may result in sub-optimal flows through the network topology between two communicating devices. The impact for example can cause increase of the RTT and packet loss, especially considering the availability (or better non-availability) of tunnel aggregation/de-aggregation points of certain topological areas or realms. The increase of non-managed tunnel usage would amplify the negative impact on good quality connectivity. For many operators of tunnel aggregation/de-aggregation devices there is little motivation to increase the quality and number of available devices within a topological area or logistical realm.

5.3. Operational Provisioning

Some elements regarding provisioning of both managed and non-managed tunnels can be controlled, while others are beyond control or influence of people and applications using tunnels. To make applications highly reliable and performing, all elements within the traffic path must provide an expected quality service and performance. For managed tunnels, the user or provider of the tunnel can exercise a degree of operational management and hence influence good quality behaviour upon the tunnel especially upon the aggregation and de-aggregation devices. In some cases even the traffic path between both aggregation and de-aggregation can be controlled. Non-managed tunnels however have less good quality behaviour of both tunnel aggregation and de-aggregation devices because often good quality behaviour is beyond the control or influence of the tunnel user. For non-managed tunnels the tunnel aggregator and/or tunnel de-aggregator are operated by a 3rd party which may have a conflicting interest with the user of the non-managed tunnel. An exception is where the use of the tunnel mechanism is all within one ISP, or ISPs who are 'well coupled', e.g. as happens between many NRENS.

5.4. Operational Troubleshooting

When one is using non-managed tunnels, then these tunnels may get aggregated or de-aggregated by a 3rd party or a device outside the control of a contracted service provider. Troubleshooting these

devices these devices will be pretty hard for the tunnel user or to work around the issue.

Also some tools like traceroute don't work too well on asymmetric paths. Another aspect is that tunnels show as one hop in a traceroute, not indicating where problems may be.

5.5. Security

For an aggregating or de-aggregating tunnel device it is a non-trivial issue to separate the valid traffic from non-valid traffic because it is from the aggregation device perspective almost impossible to know -from- and -towards- about the tunnel traffic. This imposes potential attacks on the available resources of the aggregating/de-aggregating router. A detailed security analysis for 6to4 tunnels can be found in [RFC3964].

For the user of the non-managed IPv6 tunnel there is an underlying trust that the aggregating/de-aggregating device is a trustworthy device. However, some of the devices used are run by anonymous 3rd parties outside the trusted infrastructure from the user perspective, which is not an ideal situation. The usage of non-managed tunnels increases the risk of rogue aggregation/de-aggregation devices and may be open to malicious packet analyses or manipulation.

From the operator perspective, managing the aggregating/de-aggregating tunnel device, there is a trust assumption that no-one abuses the service. Abuse may impact preset or assumed service quality levels, and hence the quality provided can be impacted

There is also an impact caused by ipv4 firewalling upon non-managed tunnels. Common firewall policies recommend to block tunnels, especially non-managed tunnels, because there is no trust that the traffic within the tunnel is not of malicious intent. This restricts the applicability of some non-managed tunnel mechanisms (e.g. 6to4). Other tunnel mechanisms have found manners to avoid traditional firewall filtering (e.g. Teredo) and open the local network infrastructure for malicious influence (e.g. virus, worms, infrastructure attacks, etc..).

5.6. Content Services

When providing content services a very important related aspect is that these services are accessible with high reliability, are trustworthy and have a high performance. Using non-managed tunnels makes this a much harder equation and can result in all three elements to suffer negatively, without the ability to uniquely identify and resolve the root cause. The statistical impact of non-

Managed tunnels has been measured by some Internet Content providers and is often an additional delay of O(100msec) (need to add reference here)

This reduces the interest of content providers to provide content services over IPv6 when non-managed tunnels are used.

6. Conclusion

Non-managed tunnels have properties impacting the growth of networks and the Internet in a negative way. Consequences regarding black-holing, perverse traffic paths, lack of business incentive and operational management influence and security issues are a real pragmatic concern, while universal supportability for the tunnel relay services appear to be non-trivial. Due to these elements the usage of non-managed tunnelling can be considered harmful for the growth of networks and the Internet.

7. IANA Considerations

There are no extra IANA consideration for this document.

8. Security Considerations

There are no extra Security consideration for this document.

9. Acknowledgements

10. References

10.1. Normative References

10.2. Informative References

- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3964] Savola, P. and C. Patel, "Security Considerations for 6to4", RFC 3964, December 2004.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.

- [RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798, February 2007.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.

Authors' Addresses

Gunter Van de Velde
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2704 5473
Email: gvandeve@cisco.com

Ole Troan
Cisco Systems
Folldalslia 17B
Bergen N-5239
Norway

Phone: +47 917 38519
Email: ot@cisco.com

Tim Chown
University of Southampton
Highfield
Southampton, SO17 1BJ
United Kingdom

Phone: +44 23 8059 3257
Email: tjc@ecs.soton.ac.uk

