Network Working Group
Internet-Draft
Obsoletes: 7084 (if approved)
Intended status: Best Current Practice
Expires: 3 September 2026

G. Lencse
Sz辿chenyi Istv叩n University
J. Palet Martinez
The IPv6 Company
B. Patton
UNH-IOL
T. Winters
QA Cafe

2 March 2026

# Basic Requirements for IPv6 Customer Edge Routers
## draft-ietf-v6ops-rfc7084bis-05

## 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. The document obsoletes RFC 7084.

## 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 https://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 3 September 2026.

## Copyright Notice

Table of Contents



## 1. Introduction

This document defines basic IPv6 features for a residential or small-office router, referred to as an "IPv6 CE router", in order to establish an industry baseline for features to be implemented on such a router.

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.

The document does not cover the IP transition technologies available to IPv6 CE routers. For information about IP transition technologies please refer to [RFC8585].

### 1.1. Requirements Language

Take careful note: Unlike other IETF documents, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are not used as described in RFC 2119 [RFC2119]. This document uses these keywords not strictly for the purpose of interoperability, but rather for the purpose of establishing industry-common baseline functionality. As such, the document points to several other specifications (preferable in RFC or stable form) to provide additional guidance to implementers regarding any protocol implementation required to produce a successful CE router that interoperates successfully with a particular subset of currently deploying and planned common IPv6 access networks.

## 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 that 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), IEEE 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 it 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 Ethernet (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; these differences have implications for the network architecture. A typical IPv4 end-user network consists 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 a Universal Plug and Play Internet Gateway Device (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; that is, 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.

Many existing routers support dynamic routing (which learns routes from other routers), 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.

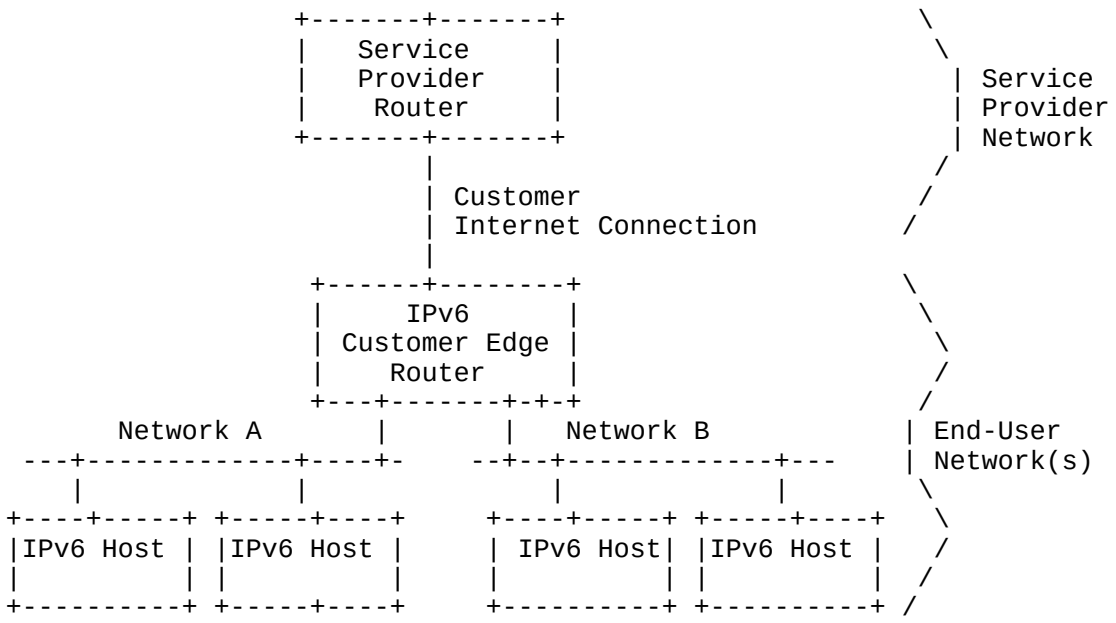


Figure 1: An Example of a Typical End-User Network

This architecture describes the:

* Basic capabilities of an IPv6 CE router

* Provisioning of the WAN interface connecting to the service provider

* Provisioning of the LAN interfaces

For IPv6 multicast traffic, the IPv6 CE router may act as a 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 are 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 (ULAs) [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.

Host implementations may 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 multihoming in a multi-prefix environment [RFC7157], the IPv6 CE router should not, as detailed in the requirements below, 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.

Coexistence with IPv4 requires any IPv6 CE router(s) on the LAN to conform to these recommendations, especially requirements ULA-5 and L-4 below.

## 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 specification [RFC8504].

G-2: The IPv6 CE router MUST implement ICMPv6 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 and the delegated prefix 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) and/or detects loss of connectivity 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].

G-6: The IPv6 CE router MUST NOT forward packets with Link-Local source or destination addresses to other links [RFC4291]], Section 2.5.6. However, there is a corner case discussed in [RFC4007], Section 9, where a router is expected to forward a packet with Link-Local source or destination addresses back to the link it was received on, if appropriate, and subject to Neighbor Discovery resolution of the destination address.

G-7: The IPv6 CE router MUST support Hop-by-Hop Option processing as described in [RFC9673].

## 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 that 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 IP Control Protocol (IPCP) for IPv4. This section makes the assumption that the same mechanism will work for any link layer, be it Ethernet, the Data Over Cable Service Interface Specification (DOCSIS), PPP, or others.

WAN-side requirements:

W-1: When the IPv6 CE router is attached to the WAN interface link, it MUST act as an IPv6 host for the purposes of stateless [RFC4862] or stateful [RFC8415] interface address assignment.

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 other routing information, the IPv6 CE router MUST use Router Discovery as specified in [RFC4861] to discover a default router(s) and install a default route(s) in its routing table with the discovered router’s address as the next hop.

W-4: The IPv6 CE router MUST act as a requesting router for the purposes of DHCPv6 prefix delegation ([RFC8415]).

W-5: The IPv6 CE router MUST use a persistent DHCP Unique Identifier (DUID) for DHCPv6 messages. The DUID MUST NOT change between network-interface resets or IPv6 CE router reboots.

W-6: The WAN interface of the CE router SHOULD support a Port Control Protocol (PCP) client as specified in [RFC6887] for use by applications on the CE router. The PCP client SHOULD follow the procedure specified in Section 8.1 of [RFC6887] to discover its PCP server. This document takes no position on whether such functionality is enabled by default or mechanisms by which users would configure the functionality. Handling PCP requests from PCP clients in the LAN side of the CE router is out of scope.

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 Network Control Protocols (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 Stateless Address Autoconfiguration (SLAAC) [RFC4862].

WAA-2: The IPv6 CE router MUST follow the recommendations in Section 4 of [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 [RFC8415] client behavior.

WAA-4: The IPv6 CE router MUST be able to support the following DHCPv6 options: Identity Association for Non-temporary Address (IA_NA), Reconfigure Accept [RFC8415], and DNS_SERVERS [RFC3646]. The IPv6 CE router SHOULD be able to support the DNS Search List (DNSSL) option as specified in [RFC3646].

WAA-5: The IPv6 CE router SHOULD implement the Network Time Protocol (NTP) as specified in [RFC5905] to provide a time reference common to the service provider for other protocols, such as DHCPv6, to use. If the IPv6 CE router implements NTP, it requests the NTP Server DHCPv6 option [RFC5908] and uses the received list of servers as primary time reference, unless explicitly configured otherwise. The IPv6 CE router SHOULD use the WAN NTP list of servers in response to DHCPv6 message requesting NTP Server DHCPv6 option on the LAN side.

WAA-6: If the IPv6 CE router receives a Router Advertisement 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 does not acquire a global IPv6 address(es) from either SLAAC or DHCPv6, then it MUST create a global IPv6 address(es) from its delegated prefix(es) and configure those on one of its internal virtual network interfaces, unless configured to require a global IPv6 address on the WAN interface.

WAA-8: The IPv6 CE router MUST support the SOL_MAX_RT option [RFC8415] and request the SOL_MAX_RT option in an Option Request Option (ORO).

WAA-9: As a router, the IPv6 CE router MUST follow the weak host (Weak End System) model [RFC1122]. When originating packets from an interface, it will use a source address from another one of its interfaces if the outgoing interface does not have an address of suitable scope.

WAA-10: The IPv6 CE router SHOULD implement the Information Refresh Time option and associated client behavior as specified in [RFC8415].

WAA-11: The IPv6 CE router MUST NOT use an EUI-64 based address as discussed in [RFC7721]. IPv6 CE router SHOULD follow [RFC8064] when generating an IPv6 address.

WAA-12: If the value of the O or M flag in Router Advertisement changes from TRUE to FALSE, the CE router MUST continue running the stateful address autoconfiguration protocol, i.e., the change in the value of Flags has no effect. In particular, a CE router MUST NOT reinvoke stateful configuration if it is already participating in the stateful protocol as a result of an earlier advertisement.

Prefix delegation requirements:

WPD-1: The IPv6 CE router MUST support DHCPv6 prefix delegation requesting router behavior as specified in [RFC8415] (Identity Association for Prefix Delegation (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 SHOULD 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. [RFC6177] covers the recommendations for service providers for prefix allocation sizes.

WPD-4: By default, the IPv6 CE router MUST initiate DHCPv6 prefix delegation when either the M or O flags are set to 1 in a received Router Advertisement (RA) message. Behavior of the CE router to use DHCPv6 prefix delegation when the CE router has not received any RA or received an RA with the M and the O bits set to zero is out of scope for this document.

WPD-5: Any packet received by the IPv6 CE router with a destination address in the prefix(es) delegated to the CE router but not in the set of prefixes assigned by the CE router to the LAN must be dropped. In other words, the next hop for the prefix(es) delegated to the CE router 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 message in accordance with Section 3.1 of [RFC4443] back to the source of the packet, if the packet is to be dropped due to this rule.

WPD-6: If the IPv6 CE router requests both an IA_NA and an IA_PD option in DHCPv6, it MUST accept an IA_PD option in DHCPv6 Advertise/Reply messages, even if the message does not contain any addresses, unless configured to only obtain its WAN IPv6 address via DHCPv6; see [RFC8415].

WPD-7: By default, an IPv6 CE router MUST NOT initiate any dynamic routing protocol on its WAN interface.

WPD-8: The IPv6 CE router SHOULD support the [RFC6603] Prefix Exclude option.

WPD-9: IPv6 CE routers SHOULD NOT automatically send a DHCPv6 message with IA_PD RELEASE messages upon restart events. See Section 3.1 [RFC9096] for further details.

WPD-10: IPv6 CE routers MUST by default use a WAN-side Identity Association Identifier (IAID) value that is stable between CE router restarts, DHCPv6 client restarts, or interface state changes (e.g., transient PPP interfaces), unless the CE router employs the IAID techniques discussed in Section 4.5 of [RFC7844]. See Section 3.2 of [RFC9096]for further details.

### 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: An IPv6 CE router with a ULA prefix MUST maintain this prefix consistently across reboots.

ULA-3: The value of the ULA prefix SHOULD be 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 a 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 or not having IPv6 connectivity on the WAN interface.

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 Advertisement messages ([RFC4861]), the Prefix Information option’s A and L flags MUST be set to 1 by default.

L-7: The A and L flags’ ([RFC4861]) settings SHOULD be user configurable.

L-8: The IPv6 CE router MUST support a DHCPv6 server capable of IPv6 address assignment according to OR a stateless DHCPv6 server according to [RFC8415] on its LAN interfaces.

L-9: Unless the IPv6 CE router is configured to support the DHCPv6 IA_NA option, it SHOULD set the M flag to zero and the O flag to 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 MUST support providing DNS information in the Router Advertisement Recursive DNS Server (RDNSS) and DNS Search List options. Both options are specified in [RFC8106].

L-12: The IPv6 CE router SHOULD make available a subset of DHCPv6 options (as listed in Section 21 of [RFC8415]) received from the DHCPv6 client on its WAN interface to its LAN-side DHCPv6 server.

L-13: The IPv6 CE router MUST signal stale configuration information as specified in Section 3.5 of [RFC9096].

L-14: The IPv6 CE router MUST send an ICMPv6 Destination Unreachable message, code 5 (Source address failed ingress/egress policy) for packets forwarded to it that use an address from a prefix that has been invalidated.

L-15: The IPv6 CE router MUST NOT advertise prefixes via SLAAC or assign addresses or delegate prefixes via DHCPv6 on the LAN side using lifetimes that exceed the remaining lifetimes of the corresponding prefixes learned on the WAN side via DHCPv6.

L-16: The IPv6 CE router SHOULD advertise capped SLAAC option lifetimes, capped DHCPv6 IA Address option lifetimes, and capped IA Prefix option lifetimes, as specified in Section 3.4 of [RFC9096].

L-17: The IPv6 CE router SHOULD implement [RFC9131] on the LAN to avoid packet loss.

L-18: The IPv6 CE router MUST be configured to respond to Router Solicitations with a unicast Router Advertisements, as specified in [RFC7772].

L-19: The IPv6 CE router SHOULD have a default MaxRtrAdvInterval value of 300 seconds and default MinRtrAdvInterval value of 100 seconds from [RFC4861]. Note, the interval at which Router Advertisements are transmited must be lower then the smallest address lifetimes advertised.

L-20: The IPv6 CE router SHOULD implement SLAAC renumbering events as documented in [I-D.ietf-6man-slaac-renum].

L-21: The IPv6 CE router MUST NOT enable IPv6 Router Advertisement Guard, [RFC6105], by default.

LAN Prefix Delegation requirements (RFC 9818):

LPD-1: Each IPv6 CE router MUST support IPv6 prefix assignment according to Section 13.3 of [RFC8415] (Identity Association for Prefix Delegation (IA_PD) option) on its LAN interface(s).

LPD-2: Each IPv6 CE router MUST assign a prefix from the delegated prefix as specified by L-2. If insufficient prefixes are available, the IPv6 CE router MUST log a system management error.

LPD-3: The prefix assigned to a link MUST NOT change in the absence of a local policy or a topology change.

LPD-4: After LAN link prefix assignments, the IPv6 CE router MUST keep the remaining IPv6 prefixes available to other routers via Prefix Delegation.

LPD-5: IPv6 CE routers MUST maintain a local routing table that is dynamically updated with leases and the associated next hops as they are delegated to clients. Packets with destination addresses in a delegated prefix MUST be routed to that prefix regardless of which interface they are received on. When a delegated prefix is released or expires, the associated route MUST be removed from the IPv6 CE router's routing table. A delegated prefix expires when the valid lifetime assigned in the IA_PD expires without being renewed. When a prefix is released or expires, it MUST be returned the pool of available prefixes.

LPD-6: By default, the IPv6 CE router's filtering rules MUST allow forwarding of packets with an outer IPv6 header containing a source address belonging to delegated prefixes, along with reciprocal packets from the same flow, following the recommendations of [RFC6092]. This updates WPD-5 to not drop packets from prefixes that have been delegated. IPv6 CE routers MUST continue to drop packets with destination addresses in prefixes that are not assigned to the LAN or delegated.

LPD-7: The IPv6 CE router MUST provision IA_PD prefixes with a prefix-length of 64 on the LAN-facing interface unless configured to use a different prefix-length by the CE router administrator. The prefix-length of 64 is used as that is the current prefix-length supported by SLAAC [RFC4862]. For hierarchical prefix delegation, a prefix-length shorter than 64 may be configured.

LPD-8: IPv6 CE routers configured to generate a ULA prefix as defined in ULA-1 MUST continue to provision available GUA IPv6 prefixes.

LPD-9: If an IPv6 CE router is provisioning both a ULA and GUA via prefix delegation, the GUA SHOULD appear first in the DHCPv6 packets.

LPD-10: IPv6 CE routers MUST NOT delegate prefixes via DHCPv6 on the LAN using lifetimes that exceed the remaining lifetimes of the corresponding prefixes learned on the WAN.

LPD-11: IPv6 CE routers SHOULD utilize the P Flag, as described in [RFC9762], if they intend for every client to obtain prefixes for the network.

### 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 [RFC6092]. In particular, the IPv6 CE router SHOULD support functionality sufficient for implementing the set of recommendations in [RFC6092], 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 BCP 38 [RFC2827]. This SHOULD be enabled by default but users MUST be able to disable this.

S-3: If the IPv6 CE router firewall is configured to filter incoming tunneled data, the firewall SHOULD provide the capability to filter decapsulated packets from a tunnel.

## 5. Acknowledgements

The following people have participated as co-authors or provided substantial contributionas to the orginal RFC 7084 document: Ralph Droms, Kirk Erichsen, Fred Baker, Jason Weil, Lee Howard, Jean-Francois Tremblay, Yiu Lee, John Jason Brzozowski, and Heather Kirksey. Thanks to Ole Troan for editorship in the original RFC 6204 document.

Thanks to the following people (in alphabetical order) for their guidance and feedback on 7084:

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

Thanks to the following people (in alphabetical order) for their guidance and feedback on 7084-bis:

Brian Carpenter, Lorenzo Colitti, David Farmer, Bob Hinden and Ted Lemon.

This document 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: Tim Carlin and Marion Dillon.

## 7. Appendix: Changes from RFC 7084

There have been many editorial clarifications as well as significant additions and updates. While this section highlights some of the changes, readers should not rely on this section for a comprehensive list of all changes.

1. Added P flag, and when it can be supported.

2. Added PRefix Delegation on LAN requirements, RFC 9818

3. Added RA Guard, which MUST NOT be enabled by default.

4. Added Hop-by-Hop Processing to the General Section.

5. Added support for SLAAC renumbering.

6. Clarified that Link-Local packets should not be forwarded.

7. Added LAN requirements for Router Advertisements Intervals.

8. Updated to a BCP from Informational.

9. Added support for RFC 9131.

10. Added text for NTP Servers relayed from WAN to LAN.

11. Updated with RFC 9096 changes for renumbering.

12. Updated to use RFC 8585 for transition technologies.

13. Removed transition technologies 6RD and DS-Lite requirements.

14. Updated to use RFC 8415 for DHCPv6

15. Updated to use RFC 7157 for mutlihoming discussion.

16. Updated to use RFC 8106 for DNS options in Router Advertisements.

17. Updated to use RFC 8405 for IPv6 node requirements.

18. Updated S-2 requirement to a MUST to prevent spoofing attacks.

19. Added a requirement to not utilize EUI-64 address.

# 8. References

## 8.1. Normative References

[I-D.ietf-6man-slaac-renum]
Gont, F., Zorz, J., Patterson, R., and J. Linkova, "Improving the Robustness of Stateless Address Autoconfiguration (SLAAC) to Flash Renumbering Events", Work in Progress, Internet-Draft, draft-ietf-6man-slaac-renum-13, 15 February 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-6man-slaac-renum-13>.

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <https://www.rfc-editor.org/info/rfc1122>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <https://www.rfc-editor.org/info/rfc2464>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, DOI 10.17487/RFC3646, December 2003, <https://www.rfc-editor.org/info/rfc3646>.

[RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <https://www.rfc-editor.org/info/rfc4191>.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <https://www.rfc-editor.org/info/rfc4193>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.

[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, DOI 10.17487/RFC4605, August 2006, <https://www.rfc-editor.org/info/rfc4605>.

[RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, DOI 10.17487/RFC4632, August 2006, <https://www.rfc-editor.org/info/rfc4632>.

[RFC4779] Asadullah, S., Ahmed, A., Popoviciu, C., Savola, P., and J. Palet, "ISP IPv6 Deployment Scenarios in Broadband Access Networks", RFC 4779, DOI 10.17487/RFC4779, January 2007, <https://www.rfc-editor.org/info/rfc4779>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <https://www.rfc-editor.org/info/rfc4862>.

[RFC5072] Varada, S., Ed., Haskins, D., and E. Allen, "IP Version 6 over PPP", RFC 5072, DOI 10.17487/RFC5072, September 2007, <https://www.rfc-editor.org/info/rfc5072>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <https://www.rfc-editor.org/info/rfc5905>.

[RFC5908] Gayraud, R. and B. Lourdelet, "Network Time Protocol (NTP) Server Option for DHCPv6", RFC 5908, DOI 10.17487/RFC5908, June 2010, <https://www.rfc-editor.org/info/rfc5908>.

[RFC5942] Singh, H., Beebee, W., and E. Nordmark, "IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes", RFC 5942, DOI 10.17487/RFC5942, July 2010, <https://www.rfc-editor.org/info/rfc5942>.

[RFC6092] Woodyatt, J., Ed., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, DOI 10.17487/RFC6092, January 2011, <https://www.rfc-editor.org/info/rfc6092>.

[RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, DOI 10.17487/RFC6105, February 2011, <https://www.rfc-editor.org/info/rfc6105>.

[RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, DOI 10.17487/RFC6177, March 2011, <https://www.rfc-editor.org/info/rfc6177>.

[RFC6603] Korhonen, J., Ed., Savolainen, T., Krishnan, S., and O. Troan, "Prefix Exclude Option for DHCPv6-based Prefix Delegation", RFC 6603, DOI 10.17487/RFC6603, May 2012, <https://www.rfc-editor.org/info/rfc6603>.

[RFC6887] Wing, D., Ed., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, DOI 10.17487/RFC6887, April 2013, <https://www.rfc-editor.org/info/rfc6887>.

[RFC7157] Troan, O., Ed., Miles, D., Matsushima, S., Okimoto, T., and D. Wing, "IPv6 Multihoming without Network Address Translation", RFC 7157, DOI 10.17487/RFC7157, March 2014, <https://www.rfc-editor.org/info/rfc7157>.

[RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", RFC 7721, DOI 10.17487/RFC7721, March 2016, <https://www.rfc-editor.org/info/rfc7721>.

[RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy Consumption of Router Advertisements", BCP 202, RFC 7772, DOI 10.17487/RFC7772, February 2016, <https://www.rfc-editor.org/info/rfc7772>.

[RFC7844] Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity Profiles for DHCP Clients", RFC 7844, DOI 10.17487/RFC7844, May 2016, <https://www.rfc-editor.org/info/rfc7844>.

[RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", RFC 8064, DOI 10.17487/RFC8064, February 2017, <https://www.rfc-editor.org/info/rfc8064>.

[RFC8106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 8106, DOI 10.17487/RFC8106, March 2017, <https://www.rfc-editor.org/info/rfc8106>.

[RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <https://www.rfc-editor.org/info/rfc8415>.

[RFC8504] Chown, T., Loughney, J., and T. Winters, "IPv6 Node Requirements", BCP 220, RFC 8504, DOI 10.17487/RFC8504, January 2019, <https://www.rfc-editor.org/info/rfc8504>.

[RFC8585] Palet Martinez, J., Liu, H. M.-H., and M. Kawashima, "Requirements for IPv6 Customer Edge Routers to Support IPv4-as-a-Service", RFC 8585, DOI 10.17487/RFC8585, May 2019, <https://www.rfc-editor.org/info/rfc8585>.

[RFC9096] Gont, F., 貼or転, J., Patterson, R., and B. Volz, "Improving the Reaction of Customer Edge Routers to IPv6 Renumbering Events", BCP 234, RFC 9096, DOI 10.17487/RFC9096, August 2021, <https://www.rfc-editor.org/info/rfc9096>.

[RFC9131] Linkova, J., "Gratuitous Neighbor Discovery: Creating Neighbor Cache Entries on First-Hop Routers", RFC 9131, DOI 10.17487/RFC9131, October 2021, <https://www.rfc-editor.org/info/rfc9131>.

[RFC9673] Hinden, R. and G. Fairhurst, "IPv6 Hop-by-Hop Options Processing Procedures", RFC 9673, DOI 10.17487/RFC9673, October 2024, <https://www.rfc-editor.org/info/rfc9673>.

[RFC9762] Colitti, L., Linkova, J., Ma, X., Ed., and D. Lamparter, "Using Router Advertisements to Signal the Availability of DHCPv6 Prefix Delegation to Clients", RFC 9762, DOI 10.17487/RFC9762, June 2025, <https://www.rfc-editor.org/info/rfc9762>.

## 8.2. Informative References

[UPnP-IGD] Forum, U., "InternetGatewayDevice:2 Device Template Version 1.01", December 2010, <http://upnp.org/specs/gw/igd2/>.

## Authors' Addresses

G叩bor Lencse
Sz辿chenyi Istv叩n University
Gyナ喪
Egyetem tテゥr 1.
H-9026
Hungary
Email: lencse@sze.hu

Jordi Palet Martinez
The IPv6 Company
Molino de la Navata, 75
28420 La Navata - Galapagar Madrid
Spain
Email: jordi.palet@theipv6company.com
URI: http://www.theipv6company.com/

Ben Patton
University of New Hampshire, Interoperability Lab (UNH-IOL)
Durham, NH
United States
Email: bpatton@iol.unh.edu

Timothy Winters
QA Cafe
100 Main Street, Suite #212
Dover, NH 03820
United States of America
Email: tim@qacafe.com