

IPv6 Operations  
Internet-Draft  
Intended status: Best Current Practice  
Expires: December 19, 2017

F. Baker  
June 17, 2017

Requirements for a Zero-Configuration IPv6 CPE  
draft-baker-v6ops-cpe-autoconfigure-00

Abstract

This note is a breif exploration of what is required for a CPE to be auto-configurable from the perspective on an ISP or other upstream network. It assumes that the CPE may also be IPv4-capable (probably using NAT), but that the requirements for that are well understood and need no further specification.

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 December 19, 2017.

Copyright Notice

Copyright (c) 2017 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 . . . . .	2
1.1. Requirements Language . . . . .	2
2. Operational Requirements . . . . .	2
3. Expected Behavior . . . . .	3
4. Prefix Delegation . . . . .	4
5. IANA Considerations . . . . .	4
6. Security Considerations . . . . .	4
7. Privacy Considerations . . . . .	4
8. Human Rights Considerations . . . . .	5
9. Acknowledgements . . . . .	5
10. References . . . . .	5
10.1. Normative References . . . . .	5
10.2. Informative References . . . . .	5
Appendix A. Change Log . . . . .	6
Author's Address . . . . .	6

## 1. Introduction

We observe that, in today's offerings, "IPv6-capable" has many different meanings. These often require specific configuration and are non-interoperable.

The objective is to enable a customer to purchase a CPE router from a mass market store, or for an ISP to purchase CPE Routers for its managed service offering, that implement IPv6 [RFC2460] and can be attached to any residential/SOHO network and any ISP or other upstream network "as is out of the box", and work correctly.

## 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 [RFC2119].

## 2. Operational Requirements

The goal stated in Section 1 requires that downstream, which is to say within the home or SOHO, the CPE must presume that there may exist systems that will autoconfigure [RFC4862] themselves using information in a Router Advertisement [RFC4861], and that there may exist systems that require address assignment using DHCPv6 [RFC3315]. It may offer a DNS service using a provider such as OpenDNS, Google Public DNS, Amazon Route 53, or some other such service, or relay the address of an ISP-provided DNS server.

Similarly, the stated goal requires that upstream, the CPE must presume that it will be required to solicit and observe a Router Advertisement [RFC4861], and

- o learn an upstream DHCPv6 server address,
- o either autoconfigure [RFC4862] its upstream address or derive one using DHCPv6 [RFC3315],
- o potentially learn an DNS server address from an RDNSS [RFC4339] or from DHCPv6,
- o and allocate IPv6 /64 prefixes for each of its interior subnets using the IPv6 Prefix Options for DHCP [RFC3633].

Given that, it is in a position to offer IPv6 services in the residential/SOHO network depending on the upstream IPv6 capabilities.

### 3. Expected Behavior

As a result, a CPE needs to perform several steps, and come out of the box configured to do so. These include:

1. Upon detecting the upstream interface as "up", emit a Router Solicitation [RFC4861] on it.
2. If it receives a Router Advertisement [RFC4861], verify its contents. These may include:
  - \* If the RA contains a valid Prefix Information Option whose prefix is available for autoconfiguration, create an address in that prefix for that interface as specified in SLAAC [RFC4862].
  - \* Failing that, use DHCPv6 [RFC3315] to request an address from the upstream network.
  - \* In that same DHCP request, it MAY request an IA\_PD [RFC3633] delegation of a set of prefixes as described in Section 4.
3. If it has not already done so, the router should request an IA\_PD [RFC3633] delegation of a set of prefixes as described in Section 4.
4. Given an upstream interface and a delegation of prefixes to use downstream, it should
  - \* subdelegate a /64 prefix to each downstream interface

- \* allocate an address to each downstream interface using the relevant prefixes
- \* start announcing a periodic RA on each downstream interface. This RA should include, in addition to usual information elements, the RDNSS [RFC4339].

#### 4. Prefix Delegation

When the CPE requests a set of prefixes from its upstream network, there are several conditions that may apply:

- o [RFC4291] and [RFC7421] presume a /64 prefix on each IPv6 subnet.
- o Each LAN to which the CPE connects may be presumed to require a subnet - if not immediately, at some point in the future.
- o There may be LANs in the residential/SOHO network that are not attached to the CPE, but require subdelegation within the network using DHCPv6 or HNCP [RFC7788].

The IA\_PD requests a prefix, and indicates its preference for a "Length for this prefix in bits". By nature, this is exponential: if a home requires 17 subnets, it will require the prefix to be no longer than 59 bits, and therefore technically requesting at least 32 /64 prefixes. In fact, some ISPs have stated privately that they actually allocate prefix lengths of 56, 60, or 64 (and therefore sets of 256, 16, or 1 /64) depending on the CPE's request.

The CPE should request as many as it thinks it might need, including interior sub-delegation if it has an idea of what that may require.

#### 5. IANA Considerations

This memo asks the IANA for no new parameters.

#### 6. Security Considerations

This note describes the use of existing features, each of which has its own Security Considerations, and as such adds no new security vulnerabilities.

#### 7. Privacy Considerations

This memo calls for no personally identifiable information. The data conveyed may, however, be correlatable with other data that is personally identifiable. Such things are beyond the scope of this document.

## 8. Human Rights Considerations

Technologies described in this memo are not necessarily associated with a human being, and as such violate no human rights.

## 9. Acknowledgements

## 10. References

### 10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.
- [RFC4339] Jeong, J., Ed., "IPv6 Host Configuration of DNS Server Information Approaches", RFC 4339, DOI 10.17487/RFC4339, February 2006, <<http://www.rfc-editor.org/info/rfc4339>>.
- [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, <<http://www.rfc-editor.org/info/rfc4861>>.

### 10.2. Informative References

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

- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<http://www.rfc-editor.org/info/rfc4941>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<http://www.rfc-editor.org/info/rfc7217>>.
- [RFC7421] Carpenter, B., Ed., Chown, T., Gont, F., Jiang, S., Petrescu, A., and A. Yourtchenko, "Analysis of the 64-bit Boundary in IPv6 Addressing", RFC 7421, DOI 10.17487/RFC7421, January 2015, <<http://www.rfc-editor.org/info/rfc7421>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<http://www.rfc-editor.org/info/rfc7788>>.

#### Appendix A. Change Log

Initial Version: Jun 13 2017

#### Author's Address

Fred Baker  
Santa Barbara, California 93117  
USA

Email: [FredBaker.IETF@gmail.com](mailto:FredBaker.IETF@gmail.com)

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: November 27, 2018

Z. Kahn, Ed.  
LinkedIn  
J. Brzozowski, Ed.  
Comcast  
R. White, Ed.  
LinkedIn  
May 26, 2018

Requirements for IPv6 Routers  
draft-ietf-v6ops-ipv6rtr-reqs-04

Abstract

The Internet is not one network, but rather a collection of networks. The interconnected nature of these networks, and the nature of the interconnected systems that make up these networks, is often more fragile than it appears. Perhaps "robust but fragile" is an overstatement, but the actions of each vendor, implementor, and operator in such an interconnected environment can have a major impact on the stability of the overall Internet (as a system). The widespread adoption of IPv6 could, particularly, disrupt network operations, in a way that impacts the entire system.

This time of transition is an opportune time to take stock of lessons learned through the operation of large-scale networks on IPv4, and consider how to apply these lessons to IPv6. This document provides an overview of the design and architectural decisions that attend IPv6 deployment, and a set of IPv6 requirements for routers, switches, and middleboxes deployed in IPv6 networks. The hope of the editors and contributors is to provide the necessary background to guide equipment manufacturers, protocol implementors, and network operators in effective IPv6 deployment.

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 November 27, 2018.

#### Copyright Notice

Copyright (c) 2018 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 (<https://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. Contributors . . . . .	3
1.2. Acknowledgments . . . . .	4
1.3. Use and Applicability . . . . .	4
2. Review of the Internet Architecture . . . . .	5
2.1. Robustness Principle . . . . .	5
2.2. Complexity . . . . .	7
2.2.1. Elegance . . . . .	7
2.2.2. Trade-offs . . . . .	8
2.3. Layered Structure . . . . .	9
2.4. Routers . . . . .	10
3. Requirements Related to Device Management and Security . . . . .	12
3.1. Programmable Device Access . . . . .	12
3.2. Human Readable Device Access . . . . .	13
3.3. Supporting Zero Touch Provisioning for Connected Devices . . . . .	13
3.4. Device Protection against Denial of Service Attacks . . . . .	15
4. Requirements Related to Telemetry . . . . .	15
4.1. Device State and Traceability . . . . .	16
4.2. Topology State and Traceability . . . . .	16
4.3. Flow State and Traceability . . . . .	17
5. Requirements Related to IPv6 Forwarding and Addressing . . . . .	17
5.1. The IPv6 Address is not a Host Identifier . . . . .	17
5.2. Router IPv6 Addresses . . . . .	18
5.3. The Maximum Transmission Unit . . . . .	19
5.4. ICMP Considerations . . . . .	20
5.5. Machine Access to the Forwarding Table . . . . .	21
5.6. Processing IPv6 Extension Headers . . . . .	22
5.7. IPv6 Operation by Default . . . . .	22
5.8. IPv6 Only Operation . . . . .	22



5.9. Prefix Length Handling in IPv6 Packet Forwarding . . . .	23
5.10. IPv6 Mobility Support . . . . .	23
6. Security Considerations . . . . .	23
6.1. Robustness and Security . . . . .	23
6.2. Programmable Device Access and Security . . . . .	24
6.3. Zero Touch Provisioning and Security . . . . .	24
6.4. Defaulting to IPv6 Forwarding and Security . . . . .	24
7. IANA Considerations . . . . .	25
8. Conclusion . . . . .	25
9. References . . . . .	25
9.1. Normative References . . . . .	25
9.2. Informative References . . . . .	25
Authors' Addresses . . . . .	32

## 1. Introduction

This memo defines and discusses requirements for devices that perform forwarding for Internet Protocol version 6 (IPv6). The "use and applicability" section below contains more information on the specific target of this draft, and the envisioned use of the draft.

Readers should recognize that while this memo applies to IPv6, routers and middleboxes IPv6 packets will often also process IPv4 packets, forward based on MPLS labels, and potentially process many other protocols. This memo will only discuss IPv4, MPLS, and other protocols as they impact the behavior of an IPv6 forwarding device; no attempt is made to specify requirements for protocols other than IPv6. The reader should, therefore, not count on this document as a "sole source of truth," but rather use this document as a guide.

For IPv4 router requirements, readers are referred to [RFC1812]. For simplicity, the term "devices" is used interchangeably with the phrase "routers and middleboxes" and the term "routers" throughout this document. These three terms represent stylistic differences, rather than substantive differences.

This document is broken into the following sections: a review of Internet architecture and principles, requirements relating to device management, requirements related to telemetry, requirements related to IPv6 forwarding and addressing, and future considerations. Following these sections, a short conclusion is provided for review.

### 1.1. Contributors

Shawn Zandi, Pete Lumbis, Fred Baker, James Woodyatt, Erik Muller, Lee Howard, and Joe Clarke contributed significant text and ideas to this draft.

## 1.2. Acknowledgments

The editors and contributors would like to thank Ron Bonica, Lorenzo Coitti, Brian E. Carpenter, Tim Chown, Peter Lothberg, and Mikael Abrahamsson for their comments, edits, and ideas on the text of this draft.

## 1.3. Use and Applicability

The conceived use of this draft is as a reference point. The first part of the draft is designed to help IPv6 implementors and network operators to understand Internet and Internetworking technologies, so they can better understand the context of IPv6. The second part of this draft outlines a common set of requirements for devices which are designed to forward IPv6 traffic. This can include (but is not limited to) the devices described below.

- o Devices which are primarily designed to forward traffic between more than two interfaces. These are normally referred to by the Internet community as routers or, in some cases, intermediate systems.
- o Devices which are designed to modify packets rather than "just" forwarding them. These are often referred to by the Internet community as middleboxes. See [RFC7663] for a fuller definition of middleboxes.

This draft is not designed to apply to consumer devices, such as smart devices (refrigerators, light bulbs, garage door openers, etc.), Internet of Things (IoT) devices, cell phones primarily used as an end user device (such as checking email, social media, games, and use as a voice device), and other devices of this class. It is up to each provider or equipment purchaser to determine how best to apply this document to their environment.

The intended use of this document is for operators to be able to point to a common set of functionality which should be available across all IPv6 implementations. Several members of the community have argued there is no common set of IPv6 features; rather each deployment of IPv6 calls for different feature sets. However, the authors of this draft believe outlining a common set of features expected of every IPv6 forwarding device is useful. Specifically:

- o If every IPv6 deployment situation is unique, and requires a different set of features, there will not be a solid definition of what an IPv6 forwarding device is, or performs. This fragments the concept of IPv6 forwarding devices in an unhelpful way, especially as IPv6 deployment is already seen as difficult.

- o It encourages developers and vendors to code a multitude of different IPv6 stacks, one for each possible set of features. This fragments the experience with these stacks, potentially preventing the development of a well designed, fully featured stacks the entire community can rely on.

Because this document is designed to be a reference point rather than a best common practice or a standard, this document does not use [RFC2119] upper case "must" and "should" throughout. Rather, it uses lower case "must" and "should" throughout, anticipating operators will find such guidance clear and useful.

## 2. Review of the Internet Architecture

The Internet relies on a number of basic concepts and considerations. These concepts are not explicitly called out in any specification, nor do they necessarily impact protocol design or packet forwarding directly. This section provides an overview of these concepts and considerations to help the reader understand the larger context of this document.

### 2.1. Robustness Principle

Every point where multiple protocols interact, is an interaction surface that can threaten the robustness of the overall system. While it may seem the global Internet has achieved a level of stability that makes it immune to such considerations, the reality is every network is a complex system, and is therefore subject to massive non repeatable unanticipated failures. Postel's Robustness Principle countered this problem with a simple statement, explicated in [RFC1122]: "Be conservative in what you do, and liberal in what you accept from others."

However, since this time, it has been noted that following this law allows errors in protocols to accumulate over time, with overall negative effects on the system as a whole. [RFC1918] describes several points in conjunction with this principle that bear updating based on further experience with large-scale protocol and network deployments within the Internet community, including:

- o Applications should deal with error states gracefully; an application should not degrade in a way that will cause the failure of adjacent systems when possible. For instance, when a routing protocol implementation fails, it should not do so in a way that will cause the spreading of or continued existence of false reachability information, nor should it fail in a way that overloads adjacent routers or interacting protocols and causing a cascading failure.

- o It is best to assume the network is filled with poor implementations and malevolent actors, both of which will find every possible failure mode over time.
- o It is best to assume every technology will be used to the limits of its technical capabilities, rather than assuming a particular protocol's scope of use will align (in any way) with the intent of the original designer(s). [RFC5218] defines a wildly successful protocol as one that "far exceeds its original goals, in terms of purpose (being used in scenarios far beyond the initial design), in terms of scale (being deployed on a scale much greater than originally envisaged), or both." Successful implementations attract more functionality, much like a few nodes in a scale free graph eventually become connectivity hubs.
- o Protocols and implementations change over time. A corollary of the assumption that protocols will be used until they reach their technical limits is that protocols will change over time as they gain new functionality. [RFC5218] points out several problems with "wild success" in a protocol: undesirable side effects, performance problems, and becoming a high value attack target. Protocol and implementation design should take into account use cases that have not yet been thought of by building flexibility into protocols. Protocols should also remained focused on a narrow range of use cases; it is often wise to invent a new protocol than to extend a single protocol into a broad set of use cases.
- o Protocols are sometimes replaced or updated to new versions in order to add new capabilities or features. Updating a protocol requires great care in providing for a transition mechanism between older and newer versions. [RFC8170] provides sound advice on protocol transition planning and mechanisms.
- o Obscure, but legal, protocol features are often ignored or left unimplemented. Protocols must handle receiving unexpected information gracefully so they do not fail because of incomplete or partial implementations. Protocols should avoid specifying contradictory states, or features that will cause interoperability issues if multiple implementations choose to implement different feature sets.
- o Monocultures are almost always bad. While multiple implementations can represent an interaction surface which increases complexity, particularly if a broad set of protocol capabilities and/or implementation features are used, using the same implementation at every point in a deployment results in a mono-culture. In a monoculture, a single event can trigger a

defect in every router, causing a network failure. Mono-cultures must be carefully balanced against interaction surfaces; often this is best accomplished by using multiple implementations and minimal, widely implemented, and well understood protocol features.

A summary of the points above might be this: It is important to work within the bounds of what is actually implemented in any given protocol, and to leave corner cases for another day. It is often easy to assume "virtual oceans" are easier to boil than physical ones, or for an ocean to appear much smaller because it is being implemented in software. This is often deceptive. It is never helpful to boil the ocean whether in a design, an implementation, or a protocol.

## 2.2. Complexity

Complexity, as articulated by Mike O'Dell (see [RFC3439]), is "the primary mechanism which impedes efficient scaling, and as a result is the primary driver of increases in both capital expenditures (CAPEX) and operational expenditures (OPEX)." At the same time, complexity cannot be "solved," but rather must be "managed." The simplest and most obvious solution to any problem is often easy to design, deploy, and manage. It's also often wrong and/or broken. As much as developers, designers, and operators might like to make things as simple as possible, hard problems require complex solutions. See Alderson and Doyle [COMPLEXHARD] for a discussion of the relationship between hard problems and complex solutions.

The following sections contain observations which apply to the management of complexity in both protocol and network design.

### 2.2.1. Elegance

Elegance should be the goal of protocol and network design. Rather than seeking out simple solutions because they are simple, seek out solutions that will solve the problem in the simplest way possible (and no simpler). Often this will require:

- o Ensuring the goal is actually the goal. Many times the goal is taken from the operational realm into the protocol design realm before enough thought has been applied to ensure the correct problem is being addressed.
- o Seeing the problem from different angles, trying to break the problem up in multiple ways; and trying, abandoning, and rebuilding ideas and implementations until a better way is found.

- o Sometimes the complexity of the solution will overwhelm the use case; sometimes it is better to leave the apparent problem unsolved, or allow the community time and space to find a simpler solution.

#### 2.2.2. Trade-offs

There are always trade-offs. For any protocol, network, or operational design decision, there will always be a trade-off between at least two competing goals. If some problem appears to have a single solution without trade-offs, this doesn't mean the trade-offs don't exist. Rather, it means the trade-offs haven't been discovered yet. In the area of protocol and network design, these trade-offs often take the form of common "choose two of three" situations, such as "quick, cheap, high quality." In network and protocol design, the trade-offs are often:

- o The amount of state carried in the system and the speed at which it changes, or simply the state. The amount of state required to operate a system as it scales tends to be nonlinear. Some instances of this are described in [RFC3439] section 2.2.1, the Amplification Principle.
- o The number of interaction surfaces between the components that make up the complete system, and the depth of those interaction surfaces. Some examples of surfaces are described in [RFC3439]section 2.2.2, the Coupling Principle. Layering is essentially a form of abstraction; all abstractions are subject to the law of leaky abstractions, [LEAKYABS] which states: "all nontrivial abstractions leak."
- o The desired optimization, including efficient use of network resources, optimal support for business objectives, and optimal support for a specific set of applications.

These three make up a "triangle problem." For instance, to increase the optimization of traffic flow through a network generally requires adding more state to the control plane, leading to problems in complexity due to amplification. To reduce amplification, the control plane (or perhaps the various functions the control plane serves) can be broken up into subsystems, or modules. Breaking the control plane up into subsystems, however, introduces interaction surfaces between the components, which is another form of complexity. [RFC7980] provides a good overview of network complexity; in particular, section 3 of that document provides some examples of complexity trade-offs.

### 2.3. Layered Structure

The Internet data plane is organized around broad top and bottom layers, and much thinner middle layer. This is illustrated in the figure below.




Figure 1

This layering emulates or mirrors many naturally occurring systems; it is a common strategy for managing complexity (see Meyer's presentation on complexity). [COMPLEXLAYER] The single protocol in the center, IPv6, serves to separate the complexity of the lower layers from the complexity of the upper layers. This center layer of the Internet ecosystem has traditionally been called the Network Layer, in reference to the Department of Defense (DoD) [DoD] and OSI models. [OSI] The Internet ecosystem includes two different protocols in this central location.

- o IPv4, an older network protocol that, it is anticipated, will be replaced over time as the Internet ecosystem standardizes on IPv6
- o IPv6, a newer network protocol that is being adopted

MPLS is often used as a "middle" sub-transport layer, and at other times as "middle" sub data link layer; hence MPLS is difficult to classify within the strictly hierarchical model depicted here. These protocols are often treated as if they exist in strict hierarchical layers with a well defined and followed Application Programming Interface (API), data models, Remote Procedure Calls (RPCs), sockets, etc. The reality, however, is there are often solid reasons for violating these layers, creating interaction surfaces that are often deeper than intended or understood without some experience. Beyond this, such layering mechanisms act as information abstractions. It is well known that all such abstractions leak (see above on the law of leaky abstractions). Because of these intentional and

unintentional leakages of information, the interactions between protocols is often subtle.

## 2.4. Routers


A router connects to two or more logical interfaces and at least one physical interface. A router processes packets by:

- o Receiving a packet through an interface
- o Stripping the data link, physical header, or tunnel encapsulation off the packet
- o Examining the packet for errors, and determining if this packet needs to be punted to another process on the router
- o Looking up the destination in a local forwarding table
- o Rewriting the data link and/or physical layer header
- o Transmitting the packet out an interface

When consulting the forwarding table, the router searches for a match based on:

- o The longest prefix containing the destination address (this is the most common matching element)
- o A label, such as a flow label or MPLS label
- o The source address or other header fields (not common)

The router then examines the information in the matching entry to determine the next hop, or rather the next logically connected device to forward the packet to. The next hop will either be another router, which will presumably carry the packet closer to the final destination, or it will be the destination host itself. The following figure provides a conceptual model of a router; not all routers actually have this set of tables and interactions, and some have many more moving parts. This model is simply used as a common reference to promote understanding.








Figure 2

The configuration datastores in this figure follow [RFC8342].

### 3. Requirements Related to Device Management and Security

Network engineering began in the era of Command Line Interfaces (CLIs), and has generally stayed with these CLIs even as the Graphical User Interface (GUI) has become the standard way of interacting with almost every other computing device. Direct human interaction with routers and middleboxes in large-scale and complex environments, however, tends to result in an unacceptably low Mean Time Between Mistakes (MTBM), directly impacting the overall availability of the network. In reaction to this, operators have increased their reliance on automation, specifically targeting machine to machine interfaces, such as Remote Procedure Calls (RPCs) and Application Programming Interface (API) solutions, to manage and configure routers and middleboxes. This section considers the various components of device management.

Across all interface types, devices should provide and use complete, idempotent, stateless configurations. Further, default settings should be accessible in some way, even if they are hidden by default for configuration readability.

#### 3.1. Programmable Device Access

Configuration primarily relates to the startup, candidate, running, intended, and operational configurations in the router model shown above. In order to deploy networks at scale, operators rely on automated management of router configuration. This effort has traditionally focused on screen scraping and other proprietary methods of "reading" and "writing" configuration information through a CLI. In the future, operators expect to move towards open source/open standards YANG models, regardless of how these are encoded and/or carried (or marshaled).

Vendors and implementors should implement machine readable interfaces with overlays to support human interaction, rather than human readable interfaces with overlays to support machine to machine interaction. Emphasis should be placed on machine to machine interaction for day to day operations, rather than on human readable interfaces, which are largely used in the process of troubleshooting. Within the realm of machine to machine interfaces, emphasis should be placed on marshaling information in YANG models.

To support automated router configuration, IPv6 routers and routers should support YANG configuration, including (but not limited to):

- o Openconfig models [OPENCONF] related to the protocols configured on the device, interface state, and device state
- o [RFC8343]: A YANG Data Model for Interface Management
- o [RFC7224]: IANA Interface Type YANG Module
- o [RFC8344]: A YANG Data Model for IP Management
- o [RFC7317]: A YANG Data Model for System Management
- o [RFC8349]: A YANG Data Model for Routing Management (NMDA Version)

### 3.2. Human Readable Device Access

To operate a network at scale, operators rely on the ability to access routers and middleboxes to troubleshoot and gather state manually through a number of different interfaces. These interfaces should provide current device configuration, current device state (such as interface state, packets drops, etc.), and current control plane contents (such as the RIB in the figure above). In other words, manual interfaces should provide information about the router (the whole device stack).

To support manual state gathering and troubleshooting, IPv6 routers and middleboxes should support:

- o TELNET ([RFC0854]): TELNET should be disabled by default, but should be available for operational purposes as required or as configured by the operator
- o SSH ([RFC4253]): SSH should be the default access for IPv6 capable routers
- o All network devices supporting IPv6 must support access through an Ethernet management port

### 3.3. Supporting Zero Touch Provisioning for Connected Devices

To operate a network at scale, operators rely on protocols and mechanisms that reduce provisioning time to a minimum. The preferred state is zero touch provisioning; plug a new router in and it just works without any manual configuration. The closer an operator can come to this ideal, the more MTBM and Operational Expenses (OPEX) can be reduced -- important goals in the real world. IPv6 routers should support several standards, including, but not limited to:

- o [I-D.ietf-dhc-rfc3315bis]: Dynamic Configuration Protocol for IPv6 must be supported.
- o [RFC4862]: IPv6 Stateless Address Autoconfiguration (SLAAC) must be supported, and must be enabled by default on all router interfaces. SLAAC must be able to be disabled by operators who prefer to use some other mechanism for address management and assignment (specifically for customer facing edge ports).
- o [RFC7217]: Semantically Opaque Interface Identifiers should be supported unless there's a need to embed MAC address.
- o [RFC7934]: Host Address Availability, the ability to assign multiple addresses to a host, should be supported.
- o [RFC7527]: Enhanced Duplicate Address Detection should be supported.
- o [RFC7527]: Enhanced Duplicate Address Detection may be disabled for manually configured interfaces.
- o [RFC8028]: First-Hop Router Selection by Hosts, specifically section 2.1, which says a router should be able to send a PIO with both the L and A bits cleared.
- o [RFC3810]: Routers supporting IPv6 must support Multicast Listener Discovery Version 2
- o [RFC7772]: Routers supporting IPv6 should support Reducing Energy Consumption of Router Advertisements
- o [RFC8273]: Routers supporting IPv6 should support Unique IPv6 Prefix per Host

The provisioning of Domain Name Systems (DNS) system information is a contentious topic, based on provider, operating system, interface, and other requirements. This document therefore addresses the mechanisms that must be included in IPv6 router implementations, but leaves the option of what to configure and deploy to the network operator. Routers supporting IPv6, and intended for user facing connections, must support:

- o [RFC3646]: DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6) if DHCPv6 is supported.
- o [RFC8106]: IPv6 Router Advertisement Options for DNS Configuration. This includes the ability to send Router Advertisements (RAs) with DNS information.

Whether these are enabled by default, or require extra configuration, is left as an exercise for providers and implementation developers to determine on a case by case basis.

### 3.4. Device Protection against Denial of Service Attacks

Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks are unfortunately common in the Internet globally; these types of attacks cost network operators a great deal in opportunity and operational costs in prevention and responses. To provide for effective counters to DoS and DDoS attacks directly on routers:

- o Manufacturers and system integrators should test and clearly report the packet/traffic load handling capabilities of devices with and without various encryption methods enabled
- o Routers should be able to police traffic destined to the control plane based on the rate of traffic received, including the ability to police individual flows, targeted services, etc., at individual rates as described in [RFC6192]
- o Ideally, devices should be able to statefully filter traffic destined to the control plane

There are other useful techniques for dealing with DDoS attacks at the network level, including: transferring sessions to a new address and abandoning the address under attack, using BGP communities to spread the attack over multiple ingress ports and "consume" it, and requiring mutual authentication before allocating larger resource pools to a connection. These techniques are not "device level," and hence are not considered further here.

## 4. Requirements Related to Telemetry

Telemetry relates to information devices push to systems used to monitor and track the state of the network. This applies to individual devices as well as the network as a system. Two major challenges face operators in the area of telemetry:

- o Information that is laid out primarily for human, rather than machine, consumption. While human consumption of telemetry is important in some situations, this information should be supplied in a form that focuses on machine readability with an overlay or interpreter that allows human consumption.
- o Software systems that require information to be queried (or polled or even pushed) on a per-item basis. This form of organization can produce a lot of information, and a lot of individual packets,

very quickly, overwhelming monitoring systems and consuming a large amount of available network resources. Instead, telemetry should be focused on bulk collection.

There are three broad categories of telemetry: device state and traceability, topology state and traceability, and flow traceability. These three roughly correspond to the management plane, the control plane, and the forwarding plane of the network. Each of the sections below considers one of these three telemetry types.

#### 4.1. Device State and Traceability

Ideally, the entire network could be monitored using a single modeling language to ease implementation of telemetry systems and increase the pace at which new software can be deployed in production environments. In real deployments, it is often impossible to reach this ideal; however, reducing the languages and methods used, while focusing on machine readability, can greatly ease the deployment and management of a large-scale network. Specifically, IPv6 routers should support:

- o [RFC6241] and [RFC8040]: NETCONF/RESTCONF transporting telemetry formatted according to YANG (see above)
- o [I-D.ietf-i2rs-yang-l2-network-topology]: An I2RS model for layer 2 topologies
- o [I-D.ietf-netconf-yang-push]: YANG Datastore Subscription
- o [RFC5424]: Syslog
- o gRPC based telemetry interfaces [GRPC]
- o Simple Network Management Protocol (SNMP) MIBs as appropriate

Syslog and SNMP access for telemetry should be considered "legacy," and should not be the focus of new telemetry access development efforts.

#### 4.2. Topology State and Traceability

IPv6 routers are part of a system of devices that, combined, make up the entire network. Viewing the network as a system is often crucial for operational purposes. For instance, being able to understand changes in the topology and utilization of a network can lead to insights about traffic flow and network growth that lead to a greater understanding of how the network is operating, where problems are developing, and how to improve the network's performance. To support

systemic monitoring of the network topology, IPv6 devices should support at least:

- o [RFC5424]: North-Bound Distribution of Link-State and Traffic Engineering (TE) Information using BGP
- o [I-D.ietf-i2rs-yang-l2-network-topology]: An I2RS model for layer 2 topologies
- o [RFC8346]: An I2RS model for layer 3 topologies
- o [RFC8345]: A Data Model for Network Topologies

#### 4.3. Flow State and Traceability

Network operators frequently need to observe and understand the types, sources, and destinations of traffic passing through devices. For example, information about traffic flows may be used to identify abuse (such as DDOS attacks) or to plan network expansions based on traffic patterns. To support insight and analysis of this traffic, IPv6 devices should support IPFIX as described in [RFC7011], PSAMP as described in [RFC5474], or some other flow state mechanism.

In-situ Operational and Management (iOAM) is a technology that being developed at the time of this writing; see [I-D.ietf-ippm-ioam-data]. Operators and vendors should consider the deployment of iOAM to provide deeper information about flow and topology information.

### 5. Requirements Related to IPv6 Forwarding and Addressing

There are a number of capabilities that a device should have to be deployed into an IPv6 network, and several forwarding plane considerations operators and vendors need to bear in mind. The sections below explain these considerations.

#### 5.1. The IPv6 Address is not a Host Identifier

The IPv6 address is commonly treated as a host identifier; it is not. Rather, it is an interface identifier that describes the topological point where a particular host connects to the Internet. Specifically:

- o The IPv6 address will change when a device changes where it connects to the network.
- o A single host can have multiple addresses. For instance, a host may have one address per interface, or multiple addresses assigned

through different mechanisms, or through multiple connection points.

- o A single IPv6 address may represent many hosts, as in the case of a group of hosts reachable through a multicast address, or a set of services reachable through an anycast address.

Because the host address may change at any time, it is generally harmful to embed IPv6 addresses inside upper layer headers to identify a particular host.

## 5.2. Router IPv6 Addresses

Internet Routing Registries may allocate a network operator a wide range of prefix lengths (see [RFC6177] for further information). Within this allocation, network operators will often suballocate address space along nibble boundaries (/48, /52, /56, /60, and /64) for ease of configuration and management. Several common practices are:

- o Each multiaccess interface is allocated a /64
- o Point-to-point links are allocated a /64, but should be addressed with a longer prefix length to prevent certain kinds of denial of service attacks ([RFC6547] originally mandated 64 bit prefix lengths on point-to-point links; [RFC6164] explains possible security issues with assigning a 64 bit prefix length to a point-to-point, and recommends a /127 instead)
- o Although aggregation is typically only performed to the nibble boundaries noted above, variances are possible
- o Loopback addresses are assigned a /128

Given these common practices, routers designed to run IPv6 should support the following addressing conventions:

- o The default prefix length on any interface other than a loopback should be a /64
- o Configuring a prefix length longer than a /64 on any multi-access interface should require additional configuration steps to prevent manual configuration errors
- o Routers must not assume IPv6 prefix lengths only on nibble boundaries



- o Routers should support any prefix length shorter or greater than /64
- o Loopback interfaces should default to a /128 prefix length unless some additional configuration is undertaken to override this default setting
- o Routers must be able to generate link local addresses on all links and/or interfaces using stateless address autoconfiguration (see [RFC6434]).

### 5.3. The Maximum Transmission Unit

The long history of the Maximum Transmission Unit (MTU) in networks is not a happy one. Specific problems with MTU sizing include:

- o Many different default sizes on different media types, from very small (576 bytes on X.25) to very large (17914 bytes on 16Mbps Token Ring)
- o Many different ways to calculate the MTU on any given link; for instance a 9000 byte MTU can be calculated as 8184 bytes on one operating system, 8972 on another, and 9000 on a third
- o The increasing use of tunnel encapsulations in the network; for instance MPLS over GRE over IP over...
- o The wide variety of default MTUs across many different end hosts and operating systems
- o The general ineffectiveness of path MTU discovery to operate correctly in the face of packet filters and rate limiters (see the section on ICMP filtering below)
- o Lower speed links at the network edge which require a lot of time to serialize a packet with a large MTU
- o Increased jitter caused by the disparity between large and small packet size across a lower bandwidth links

The final point requires some further elucidation. The time required to serialize various packets at various speeds are:

- o 64 byte packet onto a 10Mb/s link: .5ms
- o 1500 byte packet onto a 10Mb/s link: 1.2ms
- o 9000 byte packet onto a 10Mb/s link: 7.2ms

- o 64 byte packet onto a 100Mb/s link: .05ms
- o 1500 byte packet onto a 100Mb/s link: .12ms
- o 9000 byte packet onto a 100Mb/s link: .72ms

A 64 byte packet trapped behind a single 1500 byte packet on a 10Mb/s link suffers 1.2ms of serialization delay. Each additional 1500 byte packet added to the queue in front of the 64 byte packet adds an additional 1.2ms of delay. In contrast, a 64 byte packet trapped behind a single 9000 byte packet on a 10Mb/s link suffers 7.7ms of serialization delay. Each additional 9000 byte packet added to the queue adds an additional 7.2ms of serialization delay. The practical result is that larger MTU sizes on lower speed links can add a significant amount of delay and jitter into a flow. On the other hand, increasing the MTU on higher speed links appears to add negligible additional delay and jitter.

The result is that it costs less in terms of overall systemic performance to use higher MTUs on higher speed links than on lower speed links. Based on this, increasing the MTU across any particular link may not increase overall end-to-end performance, but can greatly enhance the performance of local applications (such as a local BGP peering session, or a large/long standing elephant flow used to transfer data across a local fabric), while also providing room for tunnel encapsulations to be added with less impact on lower MTU end systems.

The general rule of thumb is to assume the largest size MTU should be used on higher speed transit only links in order to support a wide array of available link sizes, default MTUs, and tunnel encapsulations. Routers designed for a network core, data center core, or use on the global Internet should support at least 9000 byte MTUs on all interfaces. MTU detection mechanisms, such as IS-IS hello padding, described in [RFC3719], should be enabled to ensure correct point-to-point MTU configuration. Devices should also support:

- o [RFC8201]: Path MTU Discovery for IP version 6
- o [RFC4821]: Packetization Layer Path MTU Discovery

#### 5.4. ICMP Considerations

Internet Control Message Protocol (ICMP) is described in [RFC0792] and [RFC4443]. ICMP is often used to perform a traceroute through a network (normally by using a TTL expired ICMP message), for Path MTU discovery, and, in IPv6, for autoconfiguration and neighbor

discovery. ICMP is often blocked by middleboxes of various kinds and/or ICMP filters configured on the ingress edge of a provider network, most often to prevent the discovery of reachable hosts and network topology. Routers implementing IPv6:

- o Should rate limit the generation of ICMP echo and echo responses by default (for instance, using a token bucket method as described in [RFC4443]). The device should support the configuration of not generating ICMP echo, echo response, and time exceeded packets to prevent topology discovery.
- o Should rate limit the generation of ICMP error messages with a token bucket method as described in [RFC4443]. Rate limits should be narrow enough to (a) protect the device's ability to generate packets and (b) reduce the usefulness of ICMP error packets as part of a distributed denial of service attack. Limits should be generous enough to allow successful path MTU discovery and traceroute. For example, in a small/mid-size device, the possible defaults could be bucket size=100, refill rate=100/s. Larger devices can afford more generous rate limits.
- o Should implement the filtering suggestions in [I-D.gont-opsec-icmp-ingress-filtering]
- o Should not filter Destination Unreachable or Packet Too Big ICMP error messages by default, as this has negative impacts on many aspects of IPv6 operation, particularly path MTU discovery.

There are implications for path MTU discovery and other useful mechanisms in filtering and rate limiting ICMP. The trade-off here is between allowing unlimited ICMP, which would allow path MTU detection to work, or limiting ICMP in a way that prevents negative side effects for individual devices, and hence the operational capabilities of the network as a whole. Operators rightly limit ICMP to reduce the attack surface against their network, as well as the opportunity for "perfect storm" events that inadvertently reduce the capability of routers and middleboxes. Hence ICMP can be treated as "quasi-reliable" in many situations; existence of an ICMP message can prove, for instance, that a particular host is unreachable. The non-existence of an ICMP message, however, does not prove a particular host exists or does not.

#### 5.5. Machine Access to the Forwarding Table

In order to support treating the "network as a whole" as a single programmable system, it is important for each router have the ability to directly program forwarding information. This programmatic interface allows controllers, which are programmed to support

specific business logic and applications, to modify and filter traffic flows without interfering with the distributed control plane. While there are several programmatic interfaces available, this document suggests that the I2RS interface to the RIB be supported in all IPv6 routers. Specifically, these drafts should be supported to enable network programmability:

- o [I-D.ietf-i2rs-fb-rib-data-model]: Filter-Based RIB Data Model
- o [I-D.ietf-i2rs-fb-rib-info-model]: Filter-Based RIB Information Model
- o [I-D.ietf-i2rs-rib-data-model]: A YANG Data Model for Routing Information Base (RIB)
- o [RFC7922]: I2RS Traceability

#### 5.6. Processing IPv6 Extension Headers

(To be added)

#### 5.7. IPv6 Operation by Default

If a device forwards and/or originates IPv4 packets by default (without explicit configuration by the operator), it should forward and/or originate IPv6 packets by default. See the security considerations section below for reflections on the automatic configuration of IPv6 forwarding in parallel with IPv4.

#### 5.8. IPv6 Only Operation

While the transition to IPv6 only networks may take years (or perhaps decades), a number of operators are moving to deploy IPv6 on internal networks supporting transport and data center fabric applications more quickly. Routers and middleboxes that support IPv6 should support IPv6 only operation, including:

- o Link Local addressing must be configurable and usable as the primary address on all interfaces on a device.
- o IPv4 and/or MPLS should not be required for proper device operation. For instance, an IPv4 address should not be required to determine the router ID for any protocol. See [RFC6540] section 2.
- o Any control plane protocol implementations must support the recommendations in [RFC7404] for operation using link local addresses only.

### 5.9. Prefix Length Handling in IPv6 Packet Forwarding

Routers must support IPv6 destination lookups in the forwarding process on a single bit prefix length increments, in accordance with [RFC7608].

### 5.10. IPv6 Mobility Support

Mobile IPv6 [RFC6275] and associated specifications, including [RFC3776] and [RFC4877] allow a node to change its point of attachment within the Internet, while maintaining (and using) a permanent address. All communication using the permanent address continues to proceed as expected even as the node moves around. At the present time, Mobile IP has seen only limited implementation. More usage and deployment experience is needed with mobility before any specific approach can be recommended for broad implementation in hosts and routers. Consequently, routers may support [RFC6275] and associated specifications (these specifications are not required for IPv6 routers).

## 6. Security Considerations

This document addresses several ways in which devices designed to support IPv6 forwarding. Some of the recommendations here are designed to increase device security; for instance, see the section on device access. Others may intersect with security, but are not specifically targeted at security, such as running IPv6 link local only on links. These are not discussed further here, as they improve the security stance of the network. Other areas discussed in this draft are more nuanced. This section gathers the intersection between operational concerns and security concerns into one place.

ICMP security is already considered in the section on ICMP; it will not be considered further here. Link local only addressing will increase security by removing transit only links within the network as a reachable destination.

### 6.1. Robustness and Security

Robustness, particularly in the area of error handling, largely improves security if designed and implemented correctly. Many attacks take advantage of mistakes in implementations and variations in protocols. In particular, any feature that is unevenly implemented among a number of implementations often offers an attack surface. Hence, reducing protocol complexity helps reduce the breadth of attack surfaces.

Another point to consider at the intersection of robustness and security is the issue of monocultures. Monocultures are in and of themselves a potential attack surface, in that finding a single failure mode can be exploited to take an entire network (or operator) down. On the other hand, reducing the number of implementations for any particular protocol will decrease the set of "random" features deployed in the network. These two goals will often be opposed to one another. Network designers and operators need to consider these two sides of this trade-off, and make an intelligent decision about how much diversity to implement versus how to control the attack surface represented by deploying a wide array of implementations.

#### 6.2. Programmable Device Access and Security

Programmable interfaces, including programmable configuration, telemetry, and machine interface to the routing table, introduce a large attack surface; operators should be careful to ensure this attack surface is properly secured. Specifically:

- o Prevent external access to any administrative access points used for device programmability
- o Use AAA systems to ensure only valid devices and/or users access devices
- o Rate limit the change rate and protect management interfaces from DoS and DDoS attacks

Such interfaces should be treated no differently than SSH, SFTP, and other interfaces available to manage routers and middleboxes.

#### 6.3. Zero Touch Provisioning and Security

Zero touch provisioning opens a new attack surface; insider attackers can simply install a new device, and assume it will be autoconfigured into the network. A "simple" solution would be to install door locks, but this will likely not be enough; defenses need to be layered to be effective. It is recommended that devices installed in the network need to contain a hardware or software identification system that allows the operator to identify devices that are installed in the network.

#### 6.4. Defaulting to IPv6 Forwarding and Security

Operators should be aware that devices which forward IPv6 by default can introduce a new attack surface or new threats without explicit configuration. Operators should verify that IPv6 policies, including filtering, match or fulfill the same intent as any existing IPv4

policies when deploying devices capable of forwarding both IPv4 and IPv6.

## 7. IANA Considerations

This document has no actions for IANA.

## 8. Conclusion

The deployment of IPv6 throughout the Internet marks a point in time where it is good to review the overall Internet architecture, and assess the impact on operations of these changes. This document provides an overview of a lot of these changes and lessons learned, as well as providing pointers to many of the relevant documents to understand each topic more deeply.

## 9. References

### 9.1. Normative References

[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>>.

### 9.2. Informative References

- [COMPLEXHARD]  
Alderson, D. and J. Doyle, "Contrasting Views of Complexity and Their Implications For Network-Centric Infrastructures", 2010, <<http://ieeexplore.ieee.org/abstract/document/5477188/?reload=true>>.
- [COMPLEXLAYER]  
Meyer, D., "Macro Trends, Architecture, and the Hidden Nature of Complexity", 2010, <<http://www.slideshare.net/dmm613/macro-trends-complexityandsdn-32951199>>.
- [DoD]  
Wikipedia, "The Internet Protocol Suite", 2016, <[https://en.wikipedia.org/wiki/Internet\\_protocol\\_suite](https://en.wikipedia.org/wiki/Internet_protocol_suite)>.
- [GRPC]  
gRPC, "gRPC", 2016, <<http://www.grpc.io>>.

- [I-D.gont-opsec-icmp-ingress-filtering]  
Gont, F., Hunter, R., Massar, J., and W. LIU, "Defeating Attacks which employ Forged ICMPv4/ICMPv6 Error Messages", draft-gont-opsec-icmp-ingress-filtering-03 (work in progress), July 2017.
- [I-D.ietf-dhc-rfc3315bis]  
Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) bis", draft-ietf-dhc-rfc3315bis-13 (work in progress), April 2018.
- [I-D.ietf-i2rs-fb-rib-data-model]  
Hares, S., Kini, S., Dunbar, L., Krishnan, R., Bogdanovic, D., and R. White, "Filter-Based RIB Data Model", draft-ietf-i2rs-fb-rib-data-model-01 (work in progress), March 2017.
- [I-D.ietf-i2rs-fb-rib-info-model]  
Kini, S., Hares, S., Dunbar, L., Ghanwani, A., Krishnan, R., Bogdanovic, D., and R. White, "Filter-Based RIB Information Model", draft-ietf-i2rs-fb-rib-info-model-00 (work in progress), June 2016.
- [I-D.ietf-i2rs-rib-data-model]  
Wang, L., Chen, M., Dass, A., Ananthakrishnan, H., Kini, S., and N. Bahadur, "A YANG Data Model for Routing Information Base (RIB)", draft-ietf-i2rs-rib-data-model-15 (work in progress), May 2018.
- [I-D.ietf-i2rs-yang-l2-network-topology]  
Dong, J. and X. Wei, "A YANG Data Model for Layer-2 Network Topologies", draft-ietf-i2rs-yang-l2-network-topology-04 (work in progress), March 2018.
- [I-D.ietf-ippm-ioam-data]  
Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., Chang, R., daniel.bernier@bell.ca, d., and J. Lemon, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-02 (work in progress), March 2018.
- [I-D.ietf-netconf-yang-push]  
Clemm, A., Voit, E., Prieto, A., Tripathy, A., Nilsen-Nygaard, E., Bierman, A., and B. Lengyel, "YANG Datastore Subscription", draft-ietf-netconf-yang-push-15 (work in progress), February 2018.



- [LEAKYABS] Spolsky, J., "The Law of Leaky Abstractions", 2002, <<https://www.joelonsoftware.com/2002/11/11/the-law-of-leaky-abstractions/>>.
- [OPENCONF] OpenConfig, "Openconfig release YANG models", 2016, <<https://github.com/openconfig/public/tree/master/release>>.
- [OSI] Wikipedia, "OSI Model", 2016, <[https://en.wikipedia.org/wiki/OSI\\_model](https://en.wikipedia.org/wiki/OSI_model)>.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <<https://www.rfc-editor.org/info/rfc792>>.
- [RFC0854] Postel, J. and J. Reynolds, "Telnet Protocol Specification", STD 8, RFC 854, DOI 10.17487/RFC0854, May 1983, <<https://www.rfc-editor.org/info/rfc854>>.
- [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>>.
- [RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", RFC 1812, DOI 10.17487/RFC1812, June 1995, <<https://www.rfc-editor.org/info/rfc1812>>.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC3439] Bush, R. and D. Meyer, "Some Internet Architectural Guidelines and Philosophy", RFC 3439, DOI 10.17487/RFC3439, December 2002, <<https://www.rfc-editor.org/info/rfc3439>>.
- [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>>.

- [RFC3719] Parker, J., Ed., "Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)", RFC 3719, DOI 10.17487/RFC3719, February 2004, <<https://www.rfc-editor.org/info/rfc3719>>.
- [RFC3776] Arkko, J., Devarapalli, V., and F. Dupont, "Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents", RFC 3776, DOI 10.17487/RFC3776, June 2004, <<https://www.rfc-editor.org/info/rfc3776>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC4253] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Transport Layer Protocol", RFC 4253, DOI 10.17487/RFC4253, January 2006, <<https://www.rfc-editor.org/info/rfc4253>>.
- [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>>.
- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, DOI 10.17487/RFC4821, March 2007, <<https://www.rfc-editor.org/info/rfc4821>>.
- [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>>.
- [RFC4877] Devarapalli, V. and F. Dupont, "Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture", RFC 4877, DOI 10.17487/RFC4877, April 2007, <<https://www.rfc-editor.org/info/rfc4877>>.
- [RFC5218] Thaler, D. and B. Aboba, "What Makes for a Successful Protocol?", RFC 5218, DOI 10.17487/RFC5218, July 2008, <<https://www.rfc-editor.org/info/rfc5218>>.
- [RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, DOI 10.17487/RFC5424, March 2009, <<https://www.rfc-editor.org/info/rfc5424>>.

- [RFC5474] Duffield, N., Ed., Chiou, D., Claise, B., Greenberg, A., Grossglauser, M., and J. Rexford, "A Framework for Packet Selection and Reporting", RFC 5474, DOI 10.17487/RFC5474, March 2009, <<https://www.rfc-editor.org/info/rfc5474>>.
- [RFC6164] Kohno, M., Nitzan, B., Bush, R., Matsuzaki, Y., Colitti, L., and T. Narten, "Using 127-Bit IPv6 Prefixes on Inter-Router Links", RFC 6164, DOI 10.17487/RFC6164, April 2011, <<https://www.rfc-editor.org/info/rfc6164>>.
- [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>>.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, DOI 10.17487/RFC6192, March 2011, <<https://www.rfc-editor.org/info/rfc6192>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.
- [RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, DOI 10.17487/RFC6434, December 2011, <<https://www.rfc-editor.org/info/rfc6434>>.
- [RFC6540] George, W., Donley, C., Liljenstolpe, C., and L. Howard, "IPv6 Support Required for All IP-Capable Nodes", BCP 177, RFC 6540, DOI 10.17487/RFC6540, April 2012, <<https://www.rfc-editor.org/info/rfc6540>>.
- [RFC6547] George, W., "RFC 3627 to Historic Status", RFC 6547, DOI 10.17487/RFC6547, February 2012, <<https://www.rfc-editor.org/info/rfc6547>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.

- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.
- [RFC7224] Bjorklund, M., "IANA Interface Type YANG Module", RFC 7224, DOI 10.17487/RFC7224, May 2014, <<https://www.rfc-editor.org/info/rfc7224>>.
- [RFC7317] Bierman, A. and M. Bjorklund, "A YANG Data Model for System Management", RFC 7317, DOI 10.17487/RFC7317, August 2014, <<https://www.rfc-editor.org/info/rfc7317>>.
- [RFC7404] Behringer, M. and E. Vyncke, "Using Only Link-Local Addressing inside an IPv6 Network", RFC 7404, DOI 10.17487/RFC7404, November 2014, <<https://www.rfc-editor.org/info/rfc7404>>.
- [RFC7527] Asati, R., Singh, H., Beebee, W., Pignataro, C., Dart, E., and W. George, "Enhanced Duplicate Address Detection", RFC 7527, DOI 10.17487/RFC7527, April 2015, <<https://www.rfc-editor.org/info/rfc7527>>.
- [RFC7608] Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", BCP 198, RFC 7608, DOI 10.17487/RFC7608, July 2015, <<https://www.rfc-editor.org/info/rfc7608>>.
- [RFC7663] Trammell, B., Ed. and M. Kuehlewind, Ed., "Report from the IAB Workshop on Stack Evolution in a Middlebox Internet (SEMI)", RFC 7663, DOI 10.17487/RFC7663, October 2015, <<https://www.rfc-editor.org/info/rfc7663>>.
- [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>>.
- [RFC7922] Clarke, J., Salgueiro, G., and C. Pignataro, "Interface to the Routing System (I2RS) Traceability: Framework and Information Model", RFC 7922, DOI 10.17487/RFC7922, June 2016, <<https://www.rfc-editor.org/info/rfc7922>>.
- [RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", BCP 204, RFC 7934, DOI 10.17487/RFC7934, July 2016, <<https://www.rfc-editor.org/info/rfc7934>>.

- [RFC7980] Behringer, M., Retana, A., White, R., and G. Huston, "A Framework for Defining Network Complexity", RFC 7980, DOI 10.17487/RFC7980, October 2016, <<https://www.rfc-editor.org/info/rfc7980>>.
- [RFC7991] Hoffman, P., "The "xml2rfc" Version 3 Vocabulary", RFC 7991, DOI 10.17487/RFC7991, December 2016, <<https://www.rfc-editor.org/info/rfc7991>>.
- [RFC8028] Baker, F. and B. Carpenter, "First-Hop Router Selection by Hosts in a Multi-Prefix Network", RFC 8028, DOI 10.17487/RFC8028, November 2016, <<https://www.rfc-editor.org/info/rfc8028>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [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>>.
- [RFC8170] Thaler, D., Ed., "Planning for Protocol Adoption and Subsequent Transitions", RFC 8170, DOI 10.17487/RFC8170, May 2017, <<https://www.rfc-editor.org/info/rfc8170>>.
- [RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <<https://www.rfc-editor.org/info/rfc8201>>.
- [RFC8273] Brzozowski, J. and G. Van de Velde, "Unique IPv6 Prefix per Host", RFC 8273, DOI 10.17487/RFC8273, December 2017, <<https://www.rfc-editor.org/info/rfc8273>>.
- [RFC8342] Bjorklund, M., Schoenwaelder, J., Shafer, P., Watsen, K., and R. Wilton, "Network Management Datastore Architecture (NMDA)", RFC 8342, DOI 10.17487/RFC8342, March 2018, <<https://www.rfc-editor.org/info/rfc8342>>.
- [RFC8343] Bjorklund, M., "A YANG Data Model for Interface Management", RFC 8343, DOI 10.17487/RFC8343, March 2018, <<https://www.rfc-editor.org/info/rfc8343>>.
- [RFC8344] Bjorklund, M., "A YANG Data Model for IP Management", RFC 8344, DOI 10.17487/RFC8344, March 2018, <<https://www.rfc-editor.org/info/rfc8344>>.

- [RFC8345] Clemm, A., Medved, J., Varga, R., Bahadur, N., Ananthakrishnan, H., and X. Liu, "A YANG Data Model for Network Topologies", RFC 8345, DOI 10.17487/RFC8345, March 2018, <<https://www.rfc-editor.org/info/rfc8345>>.
- [RFC8346] Clemm, A., Medved, J., Varga, R., Liu, X., Ananthakrishnan, H., and N. Bahadur, "A YANG Data Model for Layer 3 Topologies", RFC 8346, DOI 10.17487/RFC8346, March 2018, <<https://www.rfc-editor.org/info/rfc8346>>.
- [RFC8349] Lhotka, L., Lindem, A., and Y. Qu, "A YANG Data Model for Routing Management (NMDA Version)", RFC 8349, DOI 10.17487/RFC8349, March 2018, <<https://www.rfc-editor.org/info/rfc8349>>.

## Authors' Addresses

Zaid Ali Kahn (editor)  
LinkedIn  
CA  
USA

Email: [zaid@linkedin.com](mailto:zaid@linkedin.com)

John Brzozowski (editor)  
Comcast  
USA

Email: [John\\_Brzozowski@comcast.com](mailto:John_Brzozowski@comcast.com)

Russ White (editor)  
LinkedIn  
Oak Island, NC 28465  
USA

Email: [russ@riw.us](mailto:russ@riw.us)

Network  
Internet-Draft  
Obsoletes: 6555 (if approved)  
Intended status: Standards Track  
Expires: April 28, 2018

D. Schinazi  
T. Pauly  
Apple Inc.  
October 25, 2017

Happy Eyeballs Version 2: Better Connectivity Using Concurrency  
draft-ietf-v6ops-rfc6555bis-07

Abstract

Many communication protocols operated over the modern Internet use host names. These often resolve to multiple IP addresses, each of which may have different performance and connectivity characteristics. Since specific addresses or address families (IPv4 or IPv6) may be blocked, broken, or sub-optimal on a network, clients that attempt multiple connections in parallel have a higher chance of establishing a connection sooner. This document specifies requirements for algorithms that reduce this user-visible delay and provides an example algorithm, referred to as "Happy Eyeballs". This document obsoletes the original algorithm description in [RFC6555].

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 April 28, 2018.

Copyright Notice

Copyright (c) 2017 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 (<https://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. Overview . . . . .	3
3. Hostname Resolution Query Handling . . . . .	4
3.1. Handling Multiple DNS Server Addresses . . . . .	5
4. Sorting Addresses . . . . .	5
5. Connection Attempts . . . . .	6
6. DNS Answer Changes during Happy Eyeballs Connection Setup . .	7
7. Supporting IPv6-only Networks with NAT64 and DNS64 . . . . .	8
7.1. IPv4 Address Literals . . . . .	8
7.2. Host Names with Broken AAAA Records . . . . .	9
7.3. Virtual Private Networks . . . . .	10
8. Summary of Configurable Values . . . . .	11
9. Limitations . . . . .	11
9.1. Path Maximum Transmission Unit Discovery . . . . .	12
9.2. Application Layer . . . . .	12
9.3. Hiding Operational Issues . . . . .	12
10. Security Considerations . . . . .	12
11. IANA Considerations . . . . .	12
12. Acknowledgments . . . . .	12
13. References . . . . .	13
13.1. Normative References . . . . .	13
13.2. Informative References . . . . .	14
Appendix A. Differences from RFC6555 . . . . .	14
Authors' Addresses . . . . .	15



## 1. Introduction

Many communication protocols operated over the modern Internet use host names. These often resolve to multiple IP addresses, each of which may have different performance and connectivity characteristics. Since specific addresses or address families (IPv4 or IPv6) may be blocked, broken, or sub-optimal on a network, clients that attempt multiple connections in parallel have a higher chance of establishing a connection sooner. This document specifies requirements for algorithms that reduce this user-visible delay and provides an example algorithm.

This document defines the algorithm for "Happy Eyeballs", a technique of reducing user-visible delays on dual-stack hosts. This definition obsoletes the original description in [RFC6555]. Now that this approach has been deployed at scale and measured for several years, the algorithm specification can be refined to improve its reliability and generalization.

The Happy Eyeballs algorithm of racing resolved addresses has several stages of ordering and racing to avoid delays to the user whenever possible, while preferring the use of IPv6. This document discusses how to handle DNS queries when starting a connection on a dual-stack client, how to create an ordered list of destination addresses to which to attempt connections, and how to race the connection attempts.

### 1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 2. Overview

This document defines a method of connection establishment, named "Happy Eyeballs Connection Setup". This approach has several distinct phases:

1. Initiation of asynchronous DNS queries [Section 3]
2. Sorting of resolved destination addresses [Section 4]
3. Initiation of asynchronous connection attempts [Section 5]
4. Establishment of one connection, which cancels all other attempts

Note that this document assumes that the host destination address preference policy favors IPv6 over IPv4. IPv6 has many desirable properties designed to be improvements over IPv4 [RFC8200]. If the host is configured to have a different preference, the recommendations in this document can be easily adapted.

### 3. Hostname Resolution Query Handling

When a client has both IPv4 and IPv6 connectivity, and is trying to establish a connection with a named host, it needs to send out both AAAA and A DNS queries. Both queries SHOULD be made as soon after one another as possible, with the AAAA query made first, immediately followed by the A query.

Implementations SHOULD NOT wait for both families of answers to return before attempting connection establishment. If one query fails to return, or takes significantly longer to return, waiting for the second address family can significantly delay the connection establishment of the first one. Therefore, the client SHOULD treat DNS resolution as asynchronous. Note that if the platform does not offer an asynchronous DNS API, this behavior can be simulated by making two separate synchronous queries on different threads, one per address family.

The algorithm proceeds as follows: if a positive AAAA response (a response with at least one valid AAAA record) is received first, the first IPv6 connection attempt is immediately started. If a positive A response is received first due to reordering, the client SHOULD wait for a short time for the AAAA response to ensure preference is given to IPv6 (it is common for the AAAA response to follow the A response by a few milliseconds). This delay will be referred to as the "Resolution Delay". The recommended value for the Resolution Delay is 50 milliseconds. If a positive AAAA response is received within the Resolution Delay period, the client immediately starts the IPv6 connection attempt. If a negative AAAA response (no error, no data) is received within the Resolution Delay period or the AAAA response has not been received by the end of the Resolution Delay period, the client SHOULD proceed to Sorting Addresses [Section 4] and staggered connection attempts [Section 5] using any IPv4 addresses returned so far. If the AAAA response arrives while these connection attempts are in progress, but before any connection has been established, then the newly received IPv6 addresses are incorporated into the list of available candidate addresses [Section 6] and the process of connection attempts will continue with the IPv6 addresses added, until one connection is established.

### 3.1. Handling Multiple DNS Server Addresses

If multiple DNS server addresses are configured for the current network, the client may have the option of sending its DNS queries over IPv4 or IPv6. In keeping with the Happy Eyeballs approach, queries SHOULD be sent over IPv6 first (note that this is not referring to the sending of AAAA or A queries, but rather the address of the DNS server itself and IP version used to transport DNS messages). If DNS queries sent to the IPv6 address do not receive responses, that address may be marked as penalized, and queries can be sent to other DNS server addresses.

As native IPv6 deployments become more prevalent, and IPv4 addresses are exhausted, it is expected that IPv6 connectivity will have preferential treatment within networks. If a DNS server is configured to be accessible over IPv6, IPv6 should be assumed to be the preferred address family.

Client systems SHOULD NOT have an explicit limit to the number of DNS servers that can be configured, either manually or by the network. If such a limit is required by hardware limitations, the client SHOULD use at least one address from each address family from the available list.

## 4. Sorting Addresses

Before attempting to connect to any of the resolved destination addresses, the client should define the order in which to start the attempts. Once the order has been defined, the client can use a simple algorithm for racing each option after a short delay [Section 5]. It is important that the ordered list involves all addresses from both families that have been received by this point, as this allows the client to get the racing effect of Happy Eyeballs for the entire list, not just the first IPv4 and first IPv6 addresses.

First, the client MUST sort the addresses received up to this point using Destination Address Selection ([RFC6724], Section 6).

If the client is stateful and has history of expected round-trip times (RTT) for the routes to access each address, it SHOULD add a Destination Address Selection rule between rules 8 and 9 that prefers addresses with lower RTTs. If the client keeps track of which addresses it has used in the past, it SHOULD add another destination address selection rule between the RTT rule and rule 9, which prefers used addresses over unused ones. This helps servers that use the client's IP address during authentication, as is the case for TCP Fast Open [RFC7413] and some HTTP cookies. This historical data MUST NOT be used across different network interfaces, and SHOULD be flushed whenever a device changes the network to which it is attached.

Next, the client SHOULD modify the ordered list to interleave address families. Whichever address family is first in the list should be followed by an address of the other address family; that is, if the first address in the sorted list is IPv6, then the first IPv4 address should be moved up in the list to be second in the list. An implementation MAY want to favor one address family more by allowing multiple addresses of that family to be attempted before trying the other family. The number of contiguous addresses of the first address family will be referred to as the "First Address Family Count", and can be a configurable value. This is performed to avoid waiting through a long list of addresses from a given address family if connectivity over that address family is impaired.

Note that the address selection described in this section only applies to destination addresses; Source Address Selection ([RFC6724], Section 5) is performed once per destination address and is out of scope of this document.

## 5. Connection Attempts

Once the list of addresses received up to this point has been constructed, the client will attempt to make connections. In order to avoid unreasonable network load, connection attempts SHOULD NOT be made simultaneously. Instead, one connection attempt to a single address is started first, followed by the others in the list, one at a time. Starting a new connection attempt does not affect previous attempts, as multiple connection attempts may occur in parallel. Once one of the connection attempts succeeds (generally when the TCP handshake completes), all other connections attempts that have not yet succeeded SHOULD be cancelled. Any address that was not yet attempted as a connection SHOULD be ignored. At that time, the asynchronous DNS query MAY be cancelled as new addresses will not be used for this connection. However, the DNS client resolver SHOULD still process DNS replies from the network for a short period of time

(recommended to be 1 second), as they will populate the DNS cache and can be used for subsequent connections.

A simple implementation can have a fixed delay for how long to wait before starting the next connection attempt. This delay is referred to as the "Connection Attempt Delay". One recommended value for a default delay is 250 milliseconds. A more nuanced implementation's delay should correspond to the time when the previous attempt is sending its second TCP SYN, based on TCP's retransmission timer [RFC6298]. If the client has historical RTT data gathered from other connections to the same host or prefix, it can use this information to influence its delay. Note that this algorithm should only try to approximate the time of the first SYN retransmission, and not any further retransmissions which may be influenced by exponential timer back off.

The Connection Attempt Delay MUST have a lower bound, especially if it is computed using historical data. More specifically, a subsequent connection MUST NOT be started within 10 milliseconds of the previous attempt. The recommended minimum value is 100 milliseconds, which is referred to as the "Minimum Connection Attempt Delay". This minimum value is required to avoid congestion collapse in the presence of high packet loss rates. The Connection Attempt Delay SHOULD have an upper bound, referred to as the "Maximum Connection Attempt Delay". The current recommended value is 2 seconds.

## 6. DNS Answer Changes during Happy Eyeballs Connection Setup

If, during the course of connection establishment, the DNS answers change either by adding resolved addresses (for example, due to DNS push notifications [DNS-PUSH]), or removing previously resolved addresses (for example, due to expiry of the TTL on that DNS record), the client should react based on its current progress.

If an address is removed from the list that already had a connection attempt started, the connection attempt SHOULD NOT be cancelled, but rather be allowed to continue. If the removed address had not yet had a connection attempt started, it SHOULD be removed from the list of addresses to try.

If an address is added to the list, it should be sorted into the list of addresses not yet attempted according to the rules above (Section 4).

## 7. Supporting IPv6-only Networks with NAT64 and DNS64

While many IPv6 transition protocols have been standardized and deployed, most are transparent to client devices. The combined use of NAT64 [RFC6146] and DNS64 [RFC6147] is a popular solution that is being deployed and requires changes in client devices. One possible way to handle these networks is for the client device networking stack to implement 464XLAT [RFC6877]. 464XLAT has the advantage of not requiring changes to user space software, however it requires per-packet translation if the application is using IPv4 literals and does not encourage client application software to support native IPv6. On platforms that do not support 464XLAT, the Happy Eyeballs engine SHOULD follow the recommendations in this section to properly support IPv6-only networks with NAT64 and DNS64.

The features described in this section SHOULD only be enabled when the host detects one of these networks. A simple heuristic to achieve that is to check if the network offers routable IPv6 addressing, does not offer routable IPv4 addressing, and offers a DNS resolver address.

### 7.1. IPv4 Address Literals

If client applications or users wish to connect to IPv4 address literals, the Happy Eyeballs engine will need to perform NAT64 address synthesis for them. The solution is similar to "Bump-in-the-Host" [RFC6535] but is implemented inside the Happy Eyeballs library.

When an IPv4 address is passed in to the library instead of a host name, the device queries the network for the NAT64 prefix using "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis" [RFC7050] then synthesizes an appropriate IPv6 address (or several) using the encoding described in "IPv6 Addressing of IPv4/IPv6 Translators" [RFC6052]. The synthesized addresses are then inserted into the list of addresses as if they were results from DNS queries; connection attempts follow the algorithm described above (Section 5).

## 7.2. Host Names with Broken AAAA Records

At the time of writing, there exist a small but non negligible number of host names that resolve to valid A records and broken AAAA records, which we define as AAAA records that contain seemingly valid IPv6 addresses but those addresses never reply when contacted on the usual ports. These can be for example caused by:

- o Mistyping of the IPv6 address in the DNS zone configuration
- o Routing black holes
- o Service outages

While an algorithm complying with the other sections of this document would correctly handle such host names on a dual-stack network, they will not necessarily function correctly on IPv6-only networks with NAT64 and DNS64. Since DNS64 recursive resolvers rely on the authoritative name servers sending negative ("no error no answer") responses for AAAA records in order to synthesize, they will not synthesize records for these particular host names, and will instead pass through the broken AAAA record.

In order to support these scenarios, the client device needs to query the DNS for the A record then perform local synthesis. Since these types of host names are rare, and in order to minimize load on DNS servers, this A query should only be performed when the client has given up on the AAAA records it initially received. This can be achieved by using a longer timeout, referred to as the "Last Resort Local Synthesis Delay" and recommended to be 2 seconds. The timer is started when the last connection attempt is fired. If no connection attempt has succeeded when this timer fires, the device queries the DNS for the IPv4 address and on reception of a valid A record, treats it as if it were provided by the application (Section 7.1).

### 7.3. Virtual Private Networks

Some Virtual Private Networks (VPN) may be configured to handle DNS queries from the device. The configuration could encompass all queries, or a subset such as `*.internal.example.com`. These VPNs can also be configured to only route part of the IPv4 address space, such as `192.0.2.0/24`. However, if an internal hostname resolves to an external IPv4 address, these can cause issues if the underlying network is IPv6-only. As an example, let's assume that `www.internal.example.com` has exactly one A record, `198.51.100.42`, and no AAAA records. The client will send the DNS query to the company's recursive resolver and that resolver will reply with these records. The device now only has an IPv4 address to connect to, and no route to that address. Since the company's resolver does not know the NAT64 prefix of the underlying network, it cannot synthesize the address. Similarly, the underlying network's DNS64 recursive resolver does not know the company's internal addresses, so it cannot resolve the hostname. Because of this, the client device needs to resolve the A record using the company's resolver then locally synthesize an IPv6 address, as if the resolved IPv4 address were provided by the application (Section 7.1).



## 8. Summary of Configurable Values

The values that may be configured as defaults on a client for use in Happy Eyeballs are as follows:

- o Resolution Delay (Section 3): The time to wait for a AAAA response after receiving an A response. Recommended to be 50 milliseconds.
- o First Address Family Count (Section 4): The number of addresses belonging to the first address family (such as IPv6) that should be attempted before attempting another address family. Recommended to be 1, or 2 to more aggressively favor one address family.
- o Connection Attempt Delay (Section 5): The time to wait between connection attempts in the absence of RTT data. Recommended to be 250 milliseconds.
- o Minimum Connection Attempt Delay (Section 5): The minimum time to wait between connection attempts. Recommended to be 100 milliseconds. MUST NOT be less than 10 milliseconds.
- o Maximum Connection Attempt Delay (Section 5): The maximum time to wait between connection attempts. Recommended to be 2 seconds.
- o Last Resort Local Synthesis Delay (Section 7.2): The time to wait after starting the last IPv6 attempt and before sending the A query. Recommended to be 2 seconds.

The delay values described in this section were determined empirically by measuring the timing of connections on a very wide set of production devices. They were picked to reduce wait times noticed by users while minimizing load on the network. As time passes, it is expected that the properties of networks will evolve. For that reason, it is expected that these values will change over time. Implementors should feel welcome to use different values without changing this specification. Since IPv6 issues are expected to be less common, the delays SHOULD be increased with time as client software is updated.

## 9. Limitations

Happy Eyeballs will handle initial connection failures at the TCP/IP layer, however other failures or performance issues may still affect the chosen connection.

### 9.1. Path Maximum Transmission Unit Discovery

Since Happy Eyeballs is only active during the initial handshake and TCP does not pass the initial handshake, issues related to MTU can be masked and go unnoticed during Happy Eyeballs. Solving this issue is out of scope of this document. One solution is to use Packetization Layer Path MTU Discovery [RFC4821].

### 9.2. Application Layer

If the DNS returns multiple addresses for different application servers, the application itself may not be operational and functional on all of them. Common examples include Transport Layer Security (TLS) and the Hypertext Transport Protocol (HTTP).

### 9.3. Hiding Operational Issues

It has been observed in practice that Happy Eyeballs can hide issues in networks. For example, if a misconfiguration causes IPv6 to consistently fail on a given network while IPv4 is still functional, Happy Eyeballs may impair the operator's ability to notice the issue. It is recommended that network operators deploy external means of monitoring to ensure functionality of all address families.

## 10. Security Considerations

Note that applications should not rely upon a stable hostname-to-address mapping to ensure any security properties, since DNS results may change between queries. Happy Eyeballs may make it more likely that subsequent connections to a single hostname use different IP addresses.

## 11. IANA Considerations

This memo includes no request to IANA.

## 12. Acknowledgments

The authors thank Dan Wing, Andrew Yourtchenko, and everyone else who worked on the original Happy Eyeballs design [RFC6555], Josh Graessley, Stuart Cheshire, and the rest of team at Apple that helped implement and instrument this algorithm, and Jason Fesler and Paul Saab who helped measure and refine this algorithm. The authors would also like to thank Fred Baker, Nick Chettle, Lorenzo Colitti, Igor Gashinsky, Geoff Huston, Jen Linkova, Paul Hoffman, Philip Homburg, Warren Kumari, Erik Nygren, Jordi Palet Martinez, Rui Paulo, Stephen Strowes, Jinmei Tatuya, Dave Thaler, Joe Touch and James Woodyatt for their input and contributions.

## 13. References

### 13.1. Normative References

- [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>>.
- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, DOI 10.17487/RFC4821, March 2007, <<https://www.rfc-editor.org/info/rfc4821>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<https://www.rfc-editor.org/info/rfc6052>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <<https://www.rfc-editor.org/info/rfc6146>>.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, DOI 10.17487/RFC6147, April 2011, <<https://www.rfc-editor.org/info/rfc6147>>.
- [RFC6298] Paxson, V., Allman, M., Chu, J., and M. Sargent, "Computing TCP's Retransmission Timer", RFC 6298, DOI 10.17487/RFC6298, June 2011, <<https://www.rfc-editor.org/info/rfc6298>>.
- [RFC6535] Huang, B., Deng, H., and T. Savolainen, "Dual-Stack Hosts Using "Bump-in-the-Host" (BIH)", RFC 6535, DOI 10.17487/RFC6535, February 2012, <<https://www.rfc-editor.org/info/rfc6535>>.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, DOI 10.17487/RFC6555, April 2012, <<https://www.rfc-editor.org/info/rfc6555>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<https://www.rfc-editor.org/info/rfc6724>>.

- [RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <<https://www.rfc-editor.org/info/rfc7050>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

### 13.2. Informative References

- [DNS-PUSH] Pusateri, T. and S. Cheshire, "DNS Push Notifications", Work in Progress, draft-ietf-dnssd-push, March 2017.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <<https://www.rfc-editor.org/info/rfc6877>>.
- [RFC7413] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, DOI 10.17487/RFC7413, December 2014, <<https://www.rfc-editor.org/info/rfc7413>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

### Appendix A. Differences from RFC6555

"Happy Eyeballs: Success with Dual-Stack Hosts" [RFC6555] mostly concentrates on how to stagger connections to a hostname that has an AAAA and an A record. This document additionally discusses:

- o how to perform DNS queries to obtain these addresses
- o how to handle multiple addresses from each address family
- o how to handle DNS updates while connections are being raced
- o how to leverage historical information
- o how to support IPv6-only networks with NAT64 and DNS64

Note that a simple implementation of the algorithm described in this document is still compliant with the previous specification

[RFC6555]. Implementations should take the new considerations into account when applicable to optimize their behavior.

Authors' Addresses

David Schinazi  
Apple Inc.  
1 Infinite Loop  
Cupertino, California 95014  
US

Email: dschinazi@apple.com

Tommy Pauly  
Apple Inc.  
1 Infinite Loop  
Cupertino, California 95014  
US

Email: tpaully@apple.com

IPv6 Operations (v6ops)  
Internet-Draft  
Obsoletes: 7084 (if approved)  
Intended status: Informational  
Expires: December 13, 2017

J. Palet Martinez  
Consulintel, S.L.  
June 11, 2017

Basic Requirements for IPv6 Customer Edge Routers  
draft-ietf-v6ops-rfc7084-bis-04

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 and the support of HNCP ([RFC7788]) for automated provisioning of downstream routers. The document also covers several transition technologies, as required in a world where IPv4 addresses are no longer available, so hosts in the customer LANs with IPv4-only or IPv6-only applications or devices, requiring to communicate with IPv4-only services at the Internet, are able to do so. 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 <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 December 13, 2017.

Copyright Notice

Copyright (c) 2017 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 . . . . .	4
3. Usage Scenarios . . . . .	5
4. Architecture . . . . .	6
4.1. Current IPv4 End-User Network Architecture . . . . .	6
4.2. IPv6 End-User Network Architecture . . . . .	7
4.2.1. Local Communication . . . . .	9
5. Requirements . . . . .	9
5.1. General Requirements . . . . .	9
5.2. WAN-Side Configuration . . . . .	10
5.3. LAN-Side Configuration . . . . .	14
5.4. Transition Technologies Support . . . . .	16
5.4.1. IPv4 Service Continuity in Customer LANs . . . . .	16
5.4.1.1. 464XLAT . . . . .	16
5.4.1.2. Dual-Stack Lite (DS-Lite) . . . . .	17
5.4.1.3. Lightweight 4over6 (lw4o6) . . . . .	18
5.4.1.4. MAP-E . . . . .	18
5.4.1.5. MAP-T . . . . .	19
5.4.2. Support of IPv6 in IPv4-only WAN access . . . . .	19
5.4.2.1. 6in4 . . . . .	19
5.4.2.2. 6rd . . . . .	20
5.5. IPv4 Multicast Support . . . . .	22
5.6. Security Considerations . . . . .	22
6. Acknowledgements . . . . .	22
7. Contributors . . . . .	23
8. ANNEX A: Code Considerations . . . . .	23
9. ANNEX B: Changes from RFC7084 . . . . .	24
10. ANNEX C: Changes from RFC7084-bis-00 . . . . .	24
11. ANNEX D: Changes from RFC7084-bis-01 . . . . .	25
12. ANNEX E: Changes from RFC7084-bis-02 . . . . .	25
13. ANNEX F: Changes from RFC7084-bis-03 . . . . .	25
14. References . . . . .	26
14.1. Normative References . . . . .	26
14.2. Informative References . . . . .	31
Author's Address . . . . .	31

## 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.

These routers typically also support IPv4, at least in the LAN side.

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 may be handled by means of HNCP ([RFC7788]). In some cases, manual provisioning may be acceptable, when intended for a small number of customers.

This document doesn't cover the specific details of each possible access technology. For example, if the IPv6 CE is supporting built-in or external 3GPP/LTE interfaces, [RFC7849] is a relevant reference. See [RFC4779] for a discussion of options available for deploying IPv6 in wireline service provider access networks.

This document also covers the IP transition technologies required in a world where IPv4 addresses are no longer available, so the service providers need to provision IPv6-only WAN access, while at the same time ensuring that IPv4-only or IPv6-only devices or applications in the customer LANs can still reach IPv4-only devices or applications in Internet, which still don't have IPv6 support.

### 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 IPv6 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. In other documents, the IPv6 CE is named as CPE (Customer Premises Equipment or Customer Provided Equipment). In the context of this document, both terminologies are synonymous.
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 it may also offer IPv4 Internet access. The service provider can provide such access over a variety of different transport methods such as FTTH, DSL, cable, wireless, 3GPP/LTE, 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. Usage Scenarios

The IPv6 CE router described in this document is expected to be used typically, in any of the following scenarios:

1. Residential/household users. Common usage is any kind of Internet access (web, email, streaming, online gaming, etc.).
2. Residential with Small Office/Home Office (SOHO). Same usage as for the first scenario.
3. Small Office/Home Office (SOHO). Same usage as for the first scenario.
4. Small and Medium Enterprise (SME). Same usage as for the first scenario.
5. Residential/household with advanced requirements. Same basic usage as for the first scenario, however there may be requirements for exporting services to the WAN (IP cameras, web, DNS, email, VPN, etc.).
6. Small and Medium Enterprise (SME) with advanced requirements. Same basic usage as for the first scenario, however there may be requirements for exporting services to the WAN (IP cameras, web, DNS, email, VPN, etc.).

The above list is not intended to be comprehensive of all the possible usage scenarios, just the main ones. In fact, combinations of the above usages are also possible, for example a residential with SOHO and advanced requirements.

The mechanisms for exporting IPv6 services are commonly "naturally" available in any IPv6 router, as when using GUA, unless they are blocked by firewall rules, which may require some manual configuration by means of a GUI and/or CLI.

However, in the case of IPv4, because the usage of private addresses and NAT, it typically requires some degree of manual configuration such as setting up a DMZ, virtual servers, or port/protocol forwarding. In general, CE routers already provide GUI and/or CLI to manually configure them, or the possibility to setup the CE in bridge mode, so another CE behind it, takes care of that. It is out of the scope of this document the definition of any requirements for that.

The main difference for an IPv6 CE router to support one or several of the above indicated scenarios, is related to the packet processing capabilities, performance, even other details such as the number of

WAN/LAN interfaces, their maximum speed, memory for keeping tables or tracking connections, etc. So, it is out of the scope of this document to classify them.

For example, an SME may have just 10 employees (micro-SME), which commonly will be considered same as a SOHO, but a small SME can have up to 50 employees, or 250 for a medium one. Depending on the IPv6 CE router capabilities or even how it is being configured (for instance, using SLAAC or DHCPv6), it may support even a higher number of employees if the traffic in the LANs is low, or switched by another device(s), or the WAN bandwidth requirements are low, etc. The actual bandwidth capabilities of access with technologies such as FTTH, cable and even 3GPP/LTE, allows the support of such usages, and indeed, is a very common situation that access networks and the IPv6 CE provided by the service provider are the same for SMEs and residential users.

There is also no difference in terms of who actually provides the IPv6 CE router. In most of the cases is the service provider, and in fact is responsible, typically, of provisioning/managing at least the WAN side. However, commonly the user has access to configure the LAN interfaces, firewall, DMZ, and many other aspects. In fact, in many cases, the user must supply, or at least can replace the IPv6 CE router, which makes even more relevant that all the IPv6 CE routers, support the same requirements defined in this document.

The IPv6 CE router described in this document is not intended for usage in other scenarios such as bigger Enterprises, Data Centers, Content Providers, etc. So, even if the documented requirements meet their needs, may have additional requirements, which are out of the scope of this document.

## 4. Architecture

### 4.1. 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.

#### 4.2. 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, in the sense that is not providing transit to other external networks. However HNCP ([RFC7788]) allows support for automatic provisioning of downstream routers. Figure 1 illustrates the model topology for the end-user network.

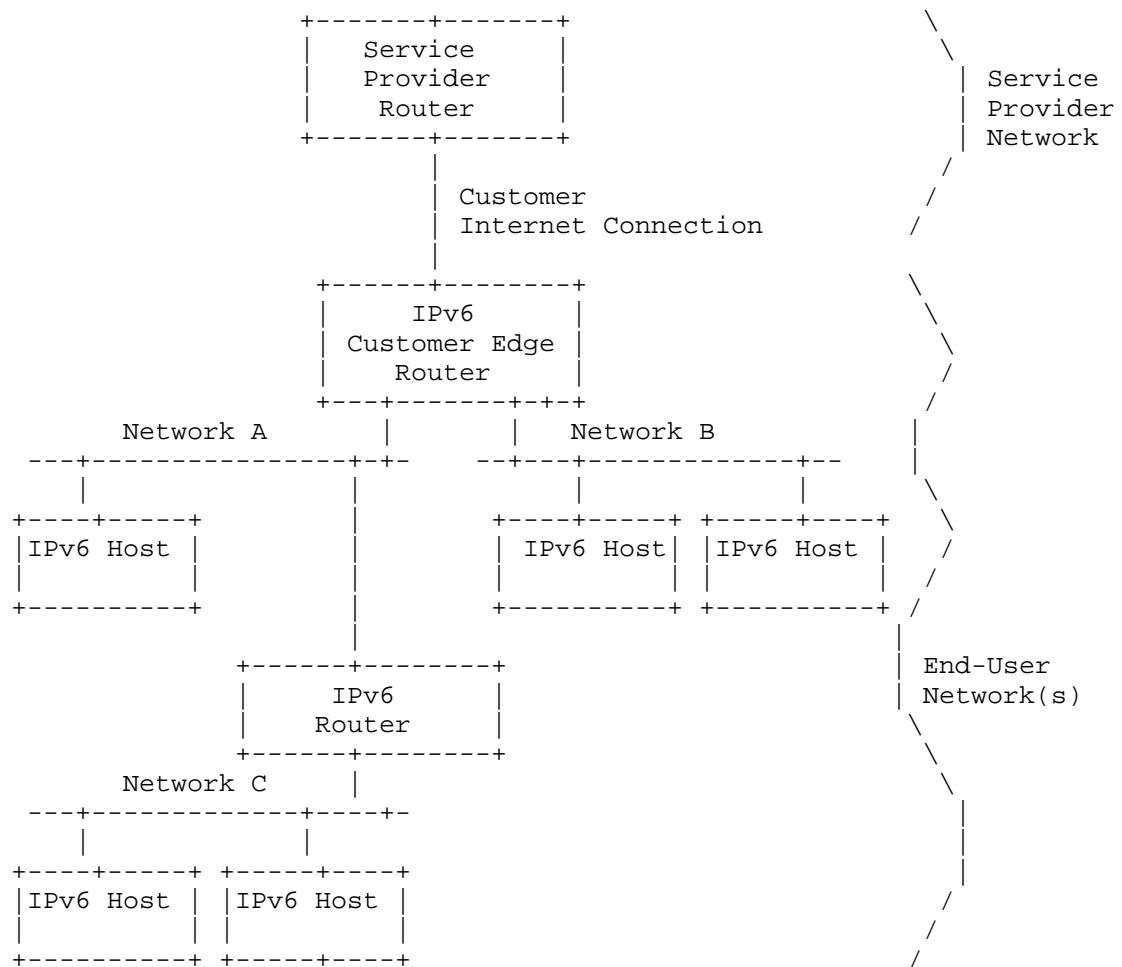


Figure 1: An Example of a Typical End-User Network

This architecture describes the:

- 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 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 or using HNCP ([RFC7788]). Automatic provisioning and configuration is described for a single IPv6 CE router only.

#### 4.2.1. 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.

At the time of this writing, several host 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 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.

### 5. Requirements

#### 5.1. 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 [RFC6434].
- 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 comply with [RFC7608].

## 5.2. 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 router is attached to the WAN interface link, it MUST

act as an IPv6 host for the purposes of stateless [RFC4862] or stateful [RFC3315] 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 router MUST act as a requesting router for the purposes of DHCPv6 prefix delegation ([RFC3633]).
- 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 IPv6 CE router SHOULD support a Port Control Protocol (PCP) client as specified in [RFC6887] for use by applications on the IPv6 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 IPv6 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 [RFC3315] 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 [RFC3315], 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. LAN side support of NTP is out of scope for this document.
- 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 [RFC7083] 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 [RFC4242].

Prefix delegation requirements:

- WPD-1: The IPv6 CE router MUST support DHCPv6 prefix delegation requesting router behavior as specified in [RFC3633] (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 IPv6 CE router to use DHCPv6 prefix delegation when the IPv6 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 IPv6 CE router but not in the set of prefixes assigned by the IPv6 CE router to the LAN must be dropped. In other words, the next hop for the prefix(es) delegated to the IPv6 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 [RFC7550].

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.

### 5.3. 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 [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 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 [RFC6106].

- L-12: The IPv6 CE router SHOULD implement a DNS proxy as described in [RFC5625].
- L-13: 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-14: 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 zero and a Valid Lifetime of either a) zero or b) the lower of the current Valid Lifetime and two hours (which must be decremented in real time) in a Router Advertisement message as described in Section 5.5.3, (e) of [RFC4862].
- L-15: 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-16: The IPv6 CE router SHOULD provide HNCP (Home Networking Control Protocol) services, as specified in [RFC7788].

#### 5.4. Transition Technologies Support

Even if the main target of this document is the support of IPv6-only WAN access, for some time, there will be a need to support IPv4-only devices and applications in the customers LANs, in one side of the picture. In the other side, some Service Providers willing to deploy IPv6, may not be able to do so in the first stage, neither as IPv6-only or dual-stack in the WAN. Consequently, transition technologies to resolve both issues should be taken in consideration.

##### 5.4.1. IPv4 Service Continuity in Customer LANs

###### 5.4.1.1. 464XLAT

464XLAT [RFC6877] is a technique to provide IPv4 access service to IPv6-only edge networks without encapsulation.

The IPv6 CE router SHOULD support CLAT functionality. If 464XLAT is supported, it MUST be implemented according to [RFC6877]. The following CE Requirements also apply:

464XLAT requirements:

- 464XLAT-1: The IPv6 CE router MUST perform IPv4 Network Address Translation (NAT) on IPv4 traffic translated using the CLAT, unless a dedicated /64 prefix has been acquired using DHCPv6-PD [RFC3633].
- 464XLAT-2: The IPv6 CE router MUST implement [RFC7050] in order to discover the PLAT-side translation IPv4 and IPv6 prefix(es)/suffix(es). In environments with PCP support, the IPv6 CE SHOULD follow [RFC7225] to learn the PLAT-side translation IPv4 and IPv6 prefix(es)/suffix(es) used by an upstream PCP-controlled NAT64 device.

#### 5.4.1.2. Dual-Stack Lite (DS-Lite)

Dual-Stack Lite [RFC6333] enables both continued support for IPv4 services and incentives for the deployment of IPv6. It also de-couples IPv6 deployment in the service provider network from the rest of the Internet, making incremental deployment easier. Dual-Stack Lite enables a broadband service provider to share IPv4 addresses among customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) and Network Address Translation (NAT). It is expected that DS-Lite traffic is forwarded over the IPv6 CE router's native IPv6 WAN interface, and not encapsulated in another tunnel.

The IPv6 CE router SHOULD implement DS-Lite functionality. If DS-Lite is supported, it MUST be implemented according to [RFC6333]. This document takes no position on simultaneous operation of Dual-Stack Lite and native IPv4. The following IPv6 CE router requirements also apply:

DS-Lite requirements:

- DSLITE-1: The IPv6 CE router MUST support configuration of DS-Lite via the DS-Lite DHCPv6 option [RFC6334]. The IPv6 CE router MAY use other mechanisms to configure DS-Lite parameters. Such mechanisms are outside the scope of this document.
- DSLITE-2: The IPv6 CE router MUST support the DHCPv6 S46 priority option described in [RFC8026].
- DSLITE-3: The IPv6 CE router MUST NOT perform IPv4 Network Address Translation (NAT) on IPv4 traffic encapsulated using DS-Lite.
- DSLITE-4: If the IPv6 CE router is configured with an IPv4 address on its WAN interface, then the IPv6 CE router SHOULD disable the DS-Lite Basic Bridging BroadBand (B4) element.

#### 5.4.1.3. Lightweight 4over6 (lw4o6)

Lw4o6 [RFC7596] specifies an extension to DS-Lite, which moves the NAPT function from the DS-Lite tunnel concentrator to the tunnel client located in the IPv6 CE router, removing the requirement for a CGN function in the tunnel concentrator and reducing the amount of centralized state.

The IPv6 CE router SHOULD implement lw4o6 functionality. If DS-Lite is implemented, lw4o6 MUST be supported as well. If lw4o6 is supported, it MUST be implemented according to [RFC7596]. This document takes no position on simultaneous operation of lw4o6 and native IPv4. The following IPv6 CE router Requirements also apply:

Lw4o6 requirements:

- LW4O6-1: The IPv6 CE router MUST support configuration of lw4o6 via the lw4o6 DHCPv6 options [RFC7598]. The IPv6 CE router MAY use other mechanisms to configure lw4o6 parameters. Such mechanisms are outside the scope of this document.
- LW4O6-2: The IPv6 CE router MUST support the DHCPv6 S46 priority option described in [RFC8026].
- LW4O6-3: The IPv6 CE router MUST support the DHCPv4-over-DHCPv6 (DHCP 4o6) transport described in [RFC7341].
- LW4O6-4: The IPv6 CE router MAY support Dynamic Allocation of Shared IPv4 Addresses as described in [RFC7618].

#### 5.4.1.4. MAP-E

MAP-E [RFC7597] is a mechanism for transporting IPv4 packets across an IPv6 network using IP encapsulation, including a generic mechanism for mapping between IPv6 addresses and IPv4 addresses as well as transport-layer ports.

The IPv6 CE router SHOULD support MAP-E functionality. If MAP-E is supported, it MUST be implemented according to [RFC7597]. The following CE Requirements also apply:

MAP-E requirements:

- MAPE-1: The IPv6 CE router MUST support configuration of MAP-E via the MAP-E DHCPv6 options [RFC7598]. The IPv6 CE router MAY use other mechanisms to configure MAP-E parameters. Such mechanisms are outside the scope of this document.

MAPE-2: The IPv6 CE router MUST support the DHCPv6 S46 priority option described in [RFC8026].

#### 5.4.1.5. MAP-T

MAP-T [RFC7599] is a mechanism similar to MAP-E, differing from it in that MAP-T uses IPv4-IPv6 translation, rather than encapsulation, as the form of IPv6 domain transport.

The IPv6 CE router SHOULD support MAP-T functionality. If MAP-T is supported, it MUST be implemented according to [RFC7599]. The following IPv6 CE Requirements also apply:

MAP-T requirements:

MAPT-1: The CE router MUST support configuration of MAP-T via the MAP-E DHCPv6 options [RFC7598]. The IPv6 CE router MAY use other mechanisms to configure MAP-E parameters. Such mechanisms are outside the scope of this document.

MAPT-2: The IPv6 CE router MUST support the DHCPv6 S46 priority option described in [RFC8026].

#### 5.4.2. Support of IPv6 in IPv4-only WAN access

##### 5.4.2.1. 6in4

6in4 [RFC4213] specifies a tunneling mechanism to allow end-users to manually configure IPv6 support via a service provider's IPv4 network infrastructure.

The IPv6 CE router MAY support 6in4 functionality. 6in4 used for a manually configured tunnel requires a subset of the 6rd parameters (delegated prefix and remote IPv4 end-point). The on-wire and forwarding plane is identical for both mechanisms, however 6in4 doesn't support mesh traffic and requires manually provisioning. Thus, if the device supports either 6rd or 6in4, it's commonly a minor UI addition to support both. If 6in4 is supported, it MUST be implemented according to [RFC4213]. The following CE Requirements also apply:

6in4 requirements:

6IN4-1: The IPv6 CE router SHOULD support 6in4 automated configuration by means of the 6rd DHCPv4 Option 212. If the IPv6 CE router has obtained an IPv4 network address through some other means such as PPP, it SHOULD use the DHCPINFORM request message [RFC2131] to request the 6rd DHCPv4 Option.



The IPv6 CE router MAY use other mechanisms to configure 6in4 parameters. Such mechanisms are outside the scope of this document.

- 6IN4-2: If the IPv6 CE router is capable of automated configuration of IPv4 through IPCP (i.e., over a PPP connection), it MUST support user-entered configuration of 6in4.
- 6IN4-3: If the IPv6 CE router supports configuration mechanisms other than the 6rd DHCPv4 Option 212 (user-entered, TR-069 [TR-069], etc.), the IPv6 CE router MUST support 6in4 in "hub and spoke" mode. 6in4 in "hub and spoke" requires all IPv6 traffic to go to the 6rd Border Relay, which in this case is the tunnel-end-point. In effect, this requirement removes the "direct connect to 6rd" route defined in Section 7.1.1 of [RFC5969].
- 6IN4-4: The IPv6 CE router MUST allow 6in4 and native IPv6 WAN interfaces to be active alone as well as simultaneously in order to support coexistence of the two technologies during an incremental transition period such as a transition from 6in4 to native IPv6.
- 6IN4-5: Each packet sent on a 6in4 or native WAN interface MUST be directed such that its source IP address is derived from the delegated prefix associated with the particular interface from which the packet is being sent (Section 4.3 of [RFC3704]).
- 6IN4-6: The IPv6 CE router MUST allow different as well as identical delegated prefixes to be configured via each (6in4 or native) WAN interface.
- 6IN4-7: In the event that forwarding rules produce a tie between 6in4 and native IPv6, by default, the IPv6 CE router MUST prefer native IPv6.

#### 5.4.2.2. 6rd

6rd [RFC5969] specifies an automatic tunneling mechanism tailored to advance deployment of IPv6 to end users via a service provider's IPv4 network infrastructure. Key aspects include automatic IPv6 prefix delegation to sites, stateless operation, simple provisioning, and service that is equivalent to native IPv6 at the sites that are served by the mechanism. It is expected that such traffic is forwarded over the IPv6 CE router's native IPv4 WAN interface and not encapsulated in another tunnel.

The IPv6 CE router MAY support 6rd functionality. If 6rd is supported, it MUST be implemented according to [RFC5969]. The following CE Requirements also apply:

6rd requirements:

- 6RD-1: The IPv6 CE router MUST support 6rd configuration via the 6rd DHCPv4 Option 212. If the IPv6 CE router has obtained an IPv4 network address through some other means such as PPP, it SHOULD use the DHCPINFORM request message [RFC2131] to request the 6rd DHCPv4 Option. The IPv6 CE router MAY use other mechanisms to configure 6rd parameters. Such mechanisms are outside the scope of this document.
- 6RD-2: If the IPv6 CE router is capable of automated configuration of IPv4 through IPCP (i.e., over a PPP connection), it MUST support user-entered configuration of 6rd.
- 6RD-3: If the IPv6 CE router supports configuration mechanisms other than the 6rd DHCPv4 Option 212 (user-entered, TR-069 [TR-069], etc.), the IPv6 CE router MUST support 6rd in "hub and spoke" mode. 6rd in "hub and spoke" requires all IPv6 traffic to go to the 6rd Border Relay. In effect, this requirement removes the "direct connect to 6rd" route defined in Section 7.1.1 of [RFC5969].
- 6RD-4: The IPv6 CE router MUST allow 6rd and native IPv6 WAN interfaces to be active alone as well as simultaneously in order to support coexistence of the two technologies during an incremental transition period such as a transition from 6rd to native IPv6.
- 6RD-5: Each packet sent on a 6rd or native WAN interface MUST be directed such that its source IP address is derived from the delegated prefix associated with the particular interface from which the packet is being sent (Section 4.3 of [RFC3704]).
- 6RD-6: The IPv6 CE router MUST allow different as well as identical delegated prefixes to be configured via each (6rd or native) WAN interface.
- 6RD-7: In the event that forwarding rules produce a tie between 6rd and native IPv6, by default, the IPv6 CE router MUST prefer native IPv6.

### 5.5. IPv4 Multicast Support

Actual deployments support IPv4 multicast for services such as IPTV. In the transition phase it is expected that multicast services will still be provided using IPv4 to the customer LANs.

In order to support the delivery of IPv4 multicast services to IPv4 clients over an IPv6 multicast network, the IPv6 CE router SHOULD support [RFC8114] and [RFC8115].

### 5.6. 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 IPv6 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 SHOULD support ingress filtering in accordance with BCP 38 [RFC2827]. Note that this requirement was downgraded from a MUST from RFC 6204 due to the difficulty of implementation in the IPv6 CE router and the feature's redundancy with upstream router ingress filtering.
- 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.

## 6. Acknowledgements

Thanks to James Woodyatt, Mohamed Boucadair, Masanobu Kawashima, Mikael Abrahamsson, Barbara Stark, Ole Troan and Brian Carpenter for their review and comments.

This document is an update of RFC7084, whose original authors were: Hemant Singh, Wes Beebe, Chris Donley and Barbara Stark. The rest of the text on this section and the Contributors section, are the

original acknowledgements and Contributors sections of the earlier version of this document.

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

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.

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.

## 7. 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. Thanks to Ole Troan for editorship in the original RFC 6204 document.

## 8. ANNEX A: Code Considerations

One of the apparent main issues for vendors to include new functionalities, such as support for new transition mechanisms, is the lack of space in the flash (or equivalent) memory. However, it has been confirmed from existing open source implementations (OpenWRT/LEDE), that adding the support for the new transitions mechanisms, requires around 10-12 Kbytes (because most of the code is shared among several transition mechanisms), which typically means about 0,15% of the existing code size in popular CEs in the market.

It is also clear that the new requirements don't have extra cost in

terms of RAM memory, neither other hardware requirements such as more powerful CPUs.

The other issue seems to be the cost of developing the code for those new functionalities. However at the time of writing this document, it has been confirmed that there are several open source versions of the required code for supporting the new transition mechanisms, so the development cost is negligent, and only integration and testing cost may become a minor issue.

#### 9. ANNEX B: Changes from RFC7084

The -bis version of this document has some minor text edits here and there. Significant updates are:

1. New section "Usage Scenarios".
2. Added support of HNCP ([RFC7788]) in LAN (L-16).
3. Added support of 464XLAT ([RFC6877]).
4. Added support of lw4o6 ([RFC7596]).
5. Added support of MAP-E ([RFC7597]) and MAP-T ([RFC7599]).
6. As the main scope of this document is the IPv6-only CE (IPv6-only in the WAN link), the support of 6rd ([RFC5969]) has been changed to MAY. 6in4 ([RFC4213]) support has been included as well in case 6rd is supported, as it doesn't require additional code.
7. New section "IPv4 Multicast Support".
8. Added support for DNS proxy [RFC5625] as general LAN requirement.
9. Split of transition in two sub-sections for the sake of clarity.

#### 10. ANNEX C: Changes from RFC7084-bis-00

Section to be removed for WGLC. Significant updates are:

1. LW4O6-5 changed to port-restricted to conform with [RFC7596].
2. MAPE-3 changed to port-restricted to conform with [RFC7597].
3. MAPT-3 changed to port-restricted to conform with [RFC7599].
4. [RFC7341] removed from 464XLAT, DS-LITE, MAP-E and MAP-T requirements.

5. [RFC5625] removed from 464XLAT, and included as general LAN requirement.
  6. [RFC7618] included as MAY for lw4o6.
  7. 6in4 text clarifications.
  8. Included non-normative reference to [RFC7849] to clarify that the details of the connectivity to 3GPP/LTE networks is out of the scope.
  9. Split of transition in two sub-sections for the sake of clarity.
11. ANNEX D: Changes from RFC7084-bis-01
- Section to be removed for WGLC. Significant updates are:
1. G-6 added in order to comply with [RFC7608].
  2. LW4O6-5 removed.
  3. MAPE-3 removed.
  4. MAPT-3 removed.
  5. Included non-normative reference to [RFC7849] to clarify that the details of the connectivity to 3GPP/LTE networks is out of the scope.
  6. Split of transition in two sub-sections for the sake of clarity.
12. ANNEX E: Changes from RFC7084-bis-02
- Section to be removed for WGLC. Significant updates are:
1. LW4O6-5 removed, was a mistake due to copy-paste from DS-LITE.
  2. Removed citation to individual I-Ds for DHCPv6 options.
13. ANNEX F: Changes from RFC7084-bis-03
- Section to be removed for WGLC. Significant updates are:
1. Clarifications on text regarding downstream routers support.

## 14. References

### 14.1. Normative References

- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<http://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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<http://www.rfc-editor.org/info/rfc2131>>.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<http://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, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.
- [RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, DOI 10.17487/RFC3646, December 2003, <<http://www.rfc-editor.org/info/rfc3646>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<http://www.rfc-editor.org/info/rfc3704>>.

- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, DOI 10.17487/RFC3736, April 2004, <<http://www.rfc-editor.org/info/rfc3736>>.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <<http://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, <<http://www.rfc-editor.org/info/rfc4193>>.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, DOI 10.17487/RFC4213, October 2005, <<http://www.rfc-editor.org/info/rfc4213>>.
- [RFC4242] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, DOI 10.17487/RFC4242, November 2005, <<http://www.rfc-editor.org/info/rfc4242>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, DOI 10.17487/RFC4443, March 2006, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc5072>>.
- [RFC5625] Bellis, R., "DNS Proxy Implementation Guidelines", BCP 152, RFC 5625, DOI 10.17487/RFC5625, August 2009, <<http://www.rfc-editor.org/info/rfc5625>>.
- [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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc5942>>.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, DOI 10.17487/RFC5969, August 2010, <<http://www.rfc-editor.org/info/rfc5969>>.
- [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, <<http://www.rfc-editor.org/info/rfc6092>>.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, DOI 10.17487/RFC6106, November 2010, <<http://www.rfc-editor.org/info/rfc6106>>.
- [RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, DOI 10.17487/RFC6177, March 2011, <<http://www.rfc-editor.org/info/rfc6177>>.

- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <<http://www.rfc-editor.org/info/rfc6333>>.
- [RFC6334] Hankins, D. and T. Mrugalski, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite", RFC 6334, DOI 10.17487/RFC6334, August 2011, <<http://www.rfc-editor.org/info/rfc6334>>.
- [RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, DOI 10.17487/RFC6434, December 2011, <<http://www.rfc-editor.org/info/rfc6434>>.
- [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, <<http://www.rfc-editor.org/info/rfc6603>>.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <<http://www.rfc-editor.org/info/rfc6877>>.
- [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, <<http://www.rfc-editor.org/info/rfc6887>>.
- [RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <<http://www.rfc-editor.org/info/rfc7050>>.
- [RFC7083] Droms, R., "Modification to Default Values of SOL\_MAX\_RT and INF\_MAX\_RT", RFC 7083, DOI 10.17487/RFC7083, November 2013, <<http://www.rfc-editor.org/info/rfc7083>>.
- [RFC7225] Boucadair, M., "Discovering NAT64 IPv6 Prefixes Using the Port Control Protocol (PCP)", RFC 7225, DOI 10.17487/RFC7225, May 2014, <<http://www.rfc-editor.org/info/rfc7225>>.
- [RFC7341] Sun, Q., Cui, Y., Siodelski, M., Krishnan, S., and I. Farrer, "DHCPv4-over-DHCPv6 (DHCP 4o6) Transport", RFC 7341, DOI 10.17487/RFC7341, August 2014, <<http://www.rfc-editor.org/info/rfc7341>>.

- [RFC7596] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture", RFC 7596, DOI 10.17487/RFC7596, July 2015, <<http://www.rfc-editor.org/info/rfc7596>>.
- [RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <<http://www.rfc-editor.org/info/rfc7597>>.
- [RFC7598] Mrugalski, T., Troan, O., Farrer, I., Perreault, S., Dec, W., Bao, C., Yeh, L., and X. Deng, "DHCPv6 Options for Configuration of Software Address and Port-Mapped Clients", RFC 7598, DOI 10.17487/RFC7598, July 2015, <<http://www.rfc-editor.org/info/rfc7598>>.
- [RFC7599] Li, X., Bao, C., Dec, W., Ed., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", RFC 7599, DOI 10.17487/RFC7599, July 2015, <<http://www.rfc-editor.org/info/rfc7599>>.
- [RFC7608] Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", BCP 198, RFC 7608, DOI 10.17487/RFC7608, July 2015, <<http://www.rfc-editor.org/info/rfc7608>>.
- [RFC7618] Cui, Y., Sun, Q., Farrer, I., Lee, Y., Sun, Q., and M. Boucadair, "Dynamic Allocation of Shared IPv4 Addresses", RFC 7618, DOI 10.17487/RFC7618, August 2015, <<http://www.rfc-editor.org/info/rfc7618>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<http://www.rfc-editor.org/info/rfc7788>>.
- [RFC8026] Boucadair, M. and I. Farrer, "Unified IPv4-in-IPv6 Software Customer Premises Equipment (CPE): A DHCPv6-Based Prioritization Mechanism", RFC 8026, DOI 10.17487/RFC8026, November 2016, <<http://www.rfc-editor.org/info/rfc8026>>.
- [RFC8114] Boucadair, M., Qin, C., Jacquenet, C., Lee, Y., and Q. Wang, "Delivery of IPv4 Multicast Services to IPv4 Clients over an IPv6 Multicast Network", RFC 8114, DOI 10.17487/RFC8114, March 2017, <<http://www.rfc-editor.org/info/rfc8114>>.

- [RFC8115] Boucadair, M., Qin, J., Tsou, T., and X. Deng, "DHCPv6 Option for IPv4-Embedded Multicast and Unicast IPv6 Prefixes", RFC 8115, DOI 10.17487/RFC8115, March 2017, <<http://www.rfc-editor.org/info/rfc8115>>.

#### 14.2. Informative References

- [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, <<http://www.rfc-editor.org/info/rfc7157>>.
- [RFC7550] Troan, O., Volz, B., and M. Siodelski, "Issues and Recommendations with Multiple Stateful DHCPv6 Options", RFC 7550, DOI 10.17487/RFC7550, May 2015, <<http://www.rfc-editor.org/info/rfc7550>>.
- [RFC7849] Binet, D., Boucadair, M., Vizdal, A., Chen, G., Heatley, N., Chandler, R., Michaud, D., Lopez, D., and W. Haeffner, "An IPv6 Profile for 3GPP Mobile Devices", RFC 7849, DOI 10.17487/RFC7849, May 2016, <<http://www.rfc-editor.org/info/rfc7849>>.
- [TR-069] Broadband Forum, "CPE WAN Management Protocol", TR-069 Amendment 4, July 2011, <<http://www.broadband-forum.org/technical/trlist.php>>.
- [UPnP-IGD] UPnP Forum, "InternetGatewayDevice:2 Device Template Version 1.01", December 2010, <<http://upnp.org/specs/gw/igd2/>>.

#### Author's Address

Jordi Palet Martinez  
Consulintel, S.L.  
Molino de la Navata, 75  
La Navata - Galapagar, Madrid 28420  
Spain

EMail: [jordi.palet@consulintel.es](mailto:jordi.palet@consulintel.es)  
URI: <http://www.consulintel.es/>

Internet Engineering Task Force  
Internet-Draft  
Intended status: Informational  
Expires: September 14, 2017

B. Liu  
S. Jiang  
Huawei Technologies  
March 13, 2017

Considerations For Using Unique Local Addresses  
draft-ietf-v6ops-ula-usage-considerations-02

Abstract

This document provides considerations for using IPv6 Unique Local Addresses (ULAs). Based on an analysis of different ULA usage scenarios, this document identifies use cases where ULA addresses are helpful as well as potential problems caused by using them.

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 September 14, 2017.

Copyright Notice

Copyright (c) 2017 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 . . . . .	2
2. Requirements Language . . . . .	3
3. General Considerations For Using ULAs . . . . .	3
3.1. Do Not Treat ULA Equal to RFC1918 . . . . .	3
3.2. Using ULAs in a Limited Scope . . . . .	4
4. Analysis and Operational Considerations for Scenarios Using ULAs . . . . .	4
4.1. ULA-only in Isolated Networks . . . . .	4
4.2. ULA+PA in Connected Networks . . . . .	5
4.3. ULA-Only in Connected Networks . . . . .	7
4.4. Some Specific Use Cases . . . . .	8
4.4.1. Special Routing . . . . .	8
4.4.2. Used as Identifier . . . . .	8
4.5. IPv4 Co-existence Considerations . . . . .	9
5. Security Considerations . . . . .	9
6. IANA Considerations . . . . .	10
7. Acknowledgements . . . . .	10
8. References . . . . .	10
8.1. Normative References . . . . .	10
8.2. Informative References . . . . .	10
Authors' Addresses . . . . .	13

## 1. Introduction

Unique Local Addresses (ULA) is defined in [RFC4193], and it is an alternative to site-local address (deprecated in [RFC3879]). ULAs have the following features:

- Automatically Generated

ULA prefixes can be automatically generated using the algorithms described in [RFC4193]. This feature allows automatic prefix allocation. Thus one can get a network working immediately without applying for prefix(es) from an RIR/LIR (Regional Internet Registry/Local Internet Registry).

- Globally Unique

ULAs are defined as a global scope address space. However, they are not intended to be used globally on the public Internet; in contrast, they are mostly used locally, for example, in isolated networks, internal networks, or VPNs.

ULAs are intended to have an extremely low probability of collision. The randomization of 40 bits in a ULA prefix is considered sufficient enough to ensure a high degree of uniqueness

(refer to [RFC4193] Section 3.2.3 for details) and simplifies merging of networks by avoiding the need to renumber overlapping IP address space.

- Provider Independent Address Space

ULAs can be used for internal communications even without Internet connectivity. They need no registration, so they can support on-demand usage and do not carry any RIR/LIR burden of documentation or fees.

- Well Known Prefix

The prefixes of ULAs are well known thus they are easily identified and filtered.

This document aims to introduce the usage of ULAs in various scenarios, provide some operational considerations, and clarify the advantages and disadvantages of the usage in each scenario. Thus, the administrators could choose to use ULAs in a certain way that considered beneficial for them.

## 2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] when they appear in ALL CAPS. When these words are not in ALL CAPS (such as "should" or "Should"), they have their usual English meanings, and are not to be interpreted as [RFC2119] key words.

## 3. General Considerations For Using ULAs

### 3.1. Do Not Treat ULA Equal to RFC1918

ULA and [RFC1918] are similar in some aspects. The most obvious one is as described in Section 3.1.3 that ULA provides an internal address independence capability in IPv6 that is similar to how [RFC1918] is commonly used. ULA allows administrators to configure the internal network of each platform the same way it is configured in IPv4. Many organizations have security policies and architectures based around the local-only routing of [RFC1918] addresses and those policies may directly map to ULA [RFC4864].

But this does not mean that ULA is equal to an IPv6 version of [RFC1918] deployment. [RFC1918] usually combines with NAT/NAPT for global connectivity. But it is not necessary to combine ULAs with

any kind of NAT. Operators can use ULA for local communications along with global addresses for global communications (see Section 4.2). This is a big advantage brought by default support of multiple-addresses-per-interface feature in IPv6. (People may still have a requirement for NAT with ULA, this is discussed in Section 4.3. But people also need to keep in mind that ULA is not intentionally designed for this kind of use case.)

Another important difference is the ability to merge two ULA networks without renumbering (because of the uniqueness), which is a big advantage over [RFC1918].

### 3.2. Using ULAs in a Limited Scope

A ULA is by definition a prefix that is never advertised outside a given domain, and is used within that domain by agreement of those networked by the domain.

So when using ULAs in a network, the administrators need to clearly set the scope of the ULAs and configure ACLs on relevant border routers to block them out of the scope. And if internal DNS is enabled, the administrators might also need to use internal-only DNS names for ULAs and might need to split the DNS so that the internal DNS server includes records that are not presented in the external DNS server.

## 4. Analysis and Operational Considerations for Scenarios Using ULAs

### 4.1. ULA-only in Isolated Networks

IP is used ubiquitously. Some networks like industrial control bus (e.g. [RS-485], [SCADA], or even non-networked digital interfaces like [MIL-STD-1397] have begun to use IP. In these kinds of networks, the system may lack the ability to communicate with the public networks.

As another example, there may be some networks in which the equipment has the technical capability to connect to the Internet, but is prohibited by administration. These networks may include data center networks, separate financial networks, lab networks. machine-to-machine (e.g. vehicle networks), sensor networks, or even normal LANs, and can include very large numbers of addresses.

ULA is a straightforward way to assign the IP addresses in the kinds of networks just described, with minimal administrative cost or burden. Also, ULAs fit in multiple subnet scenarios, in which each subnet has its own ULA prefix. For example, when assigning vehicles



with ULAs, it is then possible to separate in-vehicle embedded networks into different subnets depending on real-time situation.

However, each isolated network has the possibility to be connected in the future. Administrators need to consider the following before deciding whether to use ULAs:

- o If the network eventually connects to another isolated or private network, the potential for address collision arises. However, if the ULAs were generated in the standard way, this will not be a big problem.
- o If the network eventually connects to the global Internet, then the operator will need to add a new global prefix and ensure that the address selection policy is properly set up on all interfaces.

Operational considerations:

- o Prefix generation: randomly generated according to the algorithms defined in [RFC4193] or manually assigned. Normally, automatic generation of the prefixes is recommended, following [RFC4193]. If there are some specific reasons that call for manual assignment, administrators have to plan the prefixes carefully to avoid collision.
- o Prefix announcement: in some cases, networks might need to announce prefixes to each other. For example, in vehicle networks with infrastructure-less settings such as Vehicle-to-Vehicle (V2V) communication, prior knowledge of the respective prefixes is unlikely. Hence, a prefix announcement mechanism is needed to enable inter-vehicle communications based on IP. As one possibility, such announcements could rely on extensions to the Router Advertisement message of the Neighbor Discovery Protocol (e.g., [I-D.petrescu-autoconf-ra-based-routing] and [I-D.jhlee-mext-mnpp]).

#### 4.2. ULA+PA in Connected Networks

Two classes of network might need to use ULA with PA (Provider Aggregated) addresses:

- o Home network. Home networks are normally assigned with one or more globally routed PA prefixes to connect to the uplink of an ISP. In addition, they may need internal routed networking even when the ISP link is down. Then ULA is a proper tool to fit the requirement. [RFC7084] requires the CPE to support ULA. Note: ULAs provide more benefit for multiple-segment home networks; for

home networks containing only one segment, link-local addresses are better alternatives.

- o Enterprise network. An enterprise network is usually a managed network with one or more PA prefixes or with a PI prefix, all of which are globally routed. The ULA can be used to improve internal connectivity and make it more resilient, or to isolate certain functions like OAM for servers.

Benefits of Using ULAs in this scenario:

- o Separated local communication plane: for either home networks or enterprise networks, the main purpose of using ULAs along with PA addresses is to provide a logically local routing plane separated from the global routing plane. The benefit is to ensure stable and specific local communication regardless of the ISP uplink failure. This benefit is especially meaningful for the home network or for private OAM function in an enterprise.
- o Renumbering: in some special cases such as renumbering, enterprise administrators may want to avoid the need to renumber their internal-only, private nodes when they have to renumber the PA addresses of the rest of the network because they are changing ISPs, because the ISP has restructured its address allocations, or for some other reason. In these situations, ULA is an effective tool for addressing internal-only nodes. Even public nodes can benefit from ULA for renumbering, on their internal interfaces. When renumbering, as [RFC4192] suggests, old prefixes continue to be valid until the new prefix(es) is(are) stable. In the process of adding new prefix(es) and deprecating old prefix(es), it is not easy to keep local communication disentangled from global routing plane change. If we use ULAs for local communication, the separated local routing plane can isolate the effects of global routing change.

Drawbacks:

- o Operational Complexity: there are some arguments that in practice the use of ULA+PA creates additional operational complexity. This is not a ULA-specific problem; the multiple-addresses-per-interface is an important feature of IPv6 protocol. Nevertheless, running multiple prefixes needs more operational consideration than running a single one.

Operational considerations:

- o Default Routing: connectivity may be broken if ULAs are used as default route. When using RIO (Route Information Option) in

[RFC4191], specific routes can be added without a default route, thus avoiding bad user experience due to timeouts on ICMPv6 redirects. This behavior was well documented in [RFC7084] as rule 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." and along with rule 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.". However, it needs to be noticed that current OSes don't all support [RFC4191].

- o SLAAC/DHCPv6 co-existing: Since SLAAC and DHCPv6 might be enabled in one network simultaneously; the administrators need to carefully plan how to assign ULA and PA prefixes in accordance with the two mechanisms. The administrators need to know the current issue of the SLAAC/DHCPv6 interaction (please refer to [I-D.ietf-v6ops-dhcpv6-slaac-problem] for details).
- o Address selection: As mentioned in [RFC5220], there is a possibility that the longest matching rule will not be able to choose the correct address between ULAs and global unicast addresses for correct intra-site and extra-site communication. [RFC6724] claims that a site-specific policy entry can be used to cause ULAs within a site to be preferred over global addresses.
- o DNS relevant: if administrators choose not to do reverse DNS delegation inside of their local control of ULA prefixes, a significant amount of information about the ULA population may leak to the outside world. Because reverse queries will be made and naturally routed to the global reverse tree, so external parties will be exposed to the existence of a population of ULA addresses. [ULA-IN-WILD] provides more detailed situations on this issue. Administrators may need a split DNS to separate the queries from internal and external for ULA entries and GUA entries.

#### 4.3. ULA-Only in Connected Networks

In theory, a site numbered with ULAs only can get connected via a NPTv6[RFC6296] (which is an experimental specification that provides a stateless one-to-one mapping between internal addresses and external addresses) or application-layer proxy. This approach could get provider independent addresses or get connected from the isolated stage without applying to any RIRs/LIRs. This might make small organizations saving time and address fee.

However, this approach breaks the end-to-end transparency. People have suffered from the NAT/Proxy middle boxes so much in the IPv4 era, there is no reason to continue the suffering when IPv6 is available. This document does not consider ULA+NPTv6/Proxy as a good choice for normal cases. Rather, this document considers ULA+PA (Provider Aggregated) as a better approach to connect to the global network when ULAs are expected to be retained.

#### 4.4. Some Specific Use Cases

Along with the general scenarios, this section provides some specific use cases that could benefit from using ULA.

##### 4.4.1. Special Routing

For various reasons the administrators may want to have private routing be controlled and separated from other routing. For example, in the business-to-business case described in [I-D.baker-v6ops-b2b-private-routing], two companies might want to use direct connectivity that only connects stated machines, such as a silicon foundry with client engineers that use it. A ULA provides a simple way to assign prefixes that would be used in accordance with an agreement between the parties.

##### 4.4.2. Used as Identifier

ULAs could be self-generated and easily grabbed from the standard IPv6 stack. And ULAs don't need to be changed as the GUA prefixes do. So they are very suitable to be used as identifiers by the up layer applications. And since ULA is not intended to be globally routed, it is not harmful to the routing system.

Such kind of benefit has been utilized in real implementations. For example, in [RFC6281], the protocol BTMM (Back To My Mac) needs to assign a topology-independent identifier to each client host according to the following considerations:

- o TCP connections between two end hosts wish to survive in network changes.
- o Sometimes one needs a constant identifier to be associated with a key so that the Security Association can survive the location changes.

It needs to be noticed again that in theory ULA has the possibility of collision. However, the probability is desirably small enough and can be ignored in most cases when ULAs are used as identifiers.

#### 4.5. IPv4 Co-existence Considerations

Generally, this document does not consider IPv4 to be in scope. But regarding ULA, there is a special case needs to be recognized, which is described in Section 3.2.2 of [RFC5220]. When an enterprise has IPv4 Internet connectivity but does not yet have IPv6 Internet connectivity, and the enterprise wants to provide site-local IPv6 connectivity, a ULA is the best choice for site-local IPv6 connectivity. Each employee host will have both an IPv4 global or private address and a ULA. Here, when this host tries to connect to an outside node that has registered both A and AAAA records in the DNS, the host will choose AAAA as the destination address and the ULA for the source address according to the IPv6 preference of the default policy table defined in the old address selection standard [RFC3484]. This will clearly result in a connection failure. The new address selection standard [RFC6724] has corrected this behavior by preferring IPv4 than ULAs in the default policy table. However, there are still lots of hosts using the old standard [RFC3484], thus this could be an issue in real networks.

Happy Eyeballs [RFC6555] solves this connection failure problem, but unwanted timeouts will obviously lower the user experience. One possible approach to eliminating the timeouts is to deprecate the IPv6 default route and simply configure a scoped route on hosts (in the context of this document, only configure the ULA prefix routes). Another alternative is to configure IPv4 preference on the hosts, and not include DNS A records but only AAAA records for the internal nodes in the internal DNS server. Then outside nodes have both A and AAAA records and can be connected through IPv4 as default and internal nodes can always connect through IPv6. But since IPv6 preference is default, changing the default in all nodes is not suitable at scale.

#### 5. Security Considerations

Security considerations regarding ULAs, in general, please refer to the ULA specification [RFC4193]. Also refer to [RFC4864], which shows how ULAs help with local network protection.

As mentioned in Section 4.2, when using NPTv6, the administrators need to know where the firewall is located to set proper filtering rules.

Also as mentioned in Section 4.2, if administrators choose not to do reverse DNS delegation inside their local control of ULA prefixes, a significant amount of information about the ULA population may leak to the outside world.

## 6. IANA Considerations

This memo has no actions for IANA.

## 7. Acknowledgements

Many valuable comments were received in the IETF v6ops WG mail list, especially from Cameron Byrne, Fred Baker, Brian Carpenter, Lee Howard, Victor Kuarsingh, Alexandru Petrescu, Mikael Abrahamsson, Tim Chown, Jen Linkova, Christopher Palmer Jong-Hyouk Lee, Mark Andrews, Lorenzo Colitti, Ted Lemon, Joel Jaeggli, David Farmer, Doug Barton, Owen Delong, Gert Doering, Bill Jouris, Bill Cervený, Dave Thaler, Nick Hilliard, Jan Zorz, Randy Bush, Anders Brandt, , Sofiane Imadali and Wesley George.

Some test of using ULA in the lab was done by our research partner BNRC-BUPT (Broad Network Research Centre in Beijing University of Posts and Telecommunications). Thanks for the work of Prof. Xiangyang Gong and student Dengjia Xu.

Tom Taylor did a language review and revision throughout the whole document. The authors appreciate a lot for his help.

This document was produced using the xml2rfc tool [RFC2629] (initially prepared using 2-Word-v2.0.template.dot.).

## 8. References

### 8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, DOI 10.17487/RFC2629, June 1999, <<http://www.rfc-editor.org/info/rfc2629>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<http://www.rfc-editor.org/info/rfc4193>>.

### 8.2. Informative References

- [I-D.baker-v6ops-b2b-private-routing]  
Baker, F., "Business to Business Private Routing", draft-baker-v6ops-b2b-private-routing-00 (work in progress), July 2007.
- [I-D.ietf-v6ops-dhcpv6-slaac-problem]  
Liu, B., Jiang, S., Gong, X., Wang, W., and E. Rey, "DHCPv6/SLAAC Interaction Problems on Address and DNS Configuration", draft-ietf-v6ops-dhcpv6-slaac-problem-07 (work in progress), August 2016.
- [I-D.jhlee-mext-mnpp]  
Tsukada, M., Ernst, T., and J. Lee, "Mobile Network Prefix Provisioning", draft-jhlee-mext-mnpp-00 (work in progress), October 2009.
- [I-D.petrescu-autoconf-ra-based-routing]  
Petrescu, A., Janneteau, C., Demailly, N., and S. Imadali, "Router Advertisements for Routing between Moving Networks", draft-petrescu-autoconf-ra-based-routing-05 (work in progress), July 2014.
- [MIL-STD-1397]  
"Military Standard, Input/Output Interfaces, Standard Digital Data, Navy Systems (MIL-STD-1397B), 3 March 1989".
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<http://www.rfc-editor.org/info/rfc1918>>.
- [RFC2993] Hain, T., "Architectural Implications of NAT", RFC 2993, DOI 10.17487/RFC2993, November 2000, <<http://www.rfc-editor.org/info/rfc2993>>.
- [RFC3027] Holdrege, M. and P. Srisuresh, "Protocol Complications with the IP Network Address Translator", RFC 3027, DOI 10.17487/RFC3027, January 2001, <<http://www.rfc-editor.org/info/rfc3027>>.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, DOI 10.17487/RFC3484, February 2003, <<http://www.rfc-editor.org/info/rfc3484>>.
- [RFC3879] Huitema, C. and B. Carpenter, "Deprecating Site Local Addresses", RFC 3879, DOI 10.17487/RFC3879, September 2004, <<http://www.rfc-editor.org/info/rfc3879>>.

- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <<http://www.rfc-editor.org/info/rfc4191>>.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, DOI 10.17487/RFC4192, September 2005, <<http://www.rfc-editor.org/info/rfc4192>>.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, DOI 10.17487/RFC4864, May 2007, <<http://www.rfc-editor.org/info/rfc4864>>.
- [RFC5220] Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Problem Statement for Default Address Selection in Multi-Prefix Environments: Operational Issues of RFC 3484 Default Rules", RFC 5220, DOI 10.17487/RFC5220, July 2008, <<http://www.rfc-editor.org/info/rfc5220>>.
- [RFC5902] Thaler, D., Zhang, L., and G. Lebovitz, "IAB Thoughts on IPv6 Network Address Translation", RFC 5902, DOI 10.17487/RFC5902, July 2010, <<http://www.rfc-editor.org/info/rfc5902>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<http://www.rfc-editor.org/info/rfc6052>>.
- [RFC6281] Cheshire, S., Zhu, Z., Wakikawa, R., and L. Zhang, "Understanding Apple's Back to My Mac (BTMM) Service", RFC 6281, DOI 10.17487/RFC6281, June 2011, <<http://www.rfc-editor.org/info/rfc6281>>.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, DOI 10.17487/RFC6296, June 2011, <<http://www.rfc-editor.org/info/rfc6296>>.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, DOI 10.17487/RFC6555, April 2012, <<http://www.rfc-editor.org/info/rfc6555>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<http://www.rfc-editor.org/info/rfc6724>>.



- [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013, <<http://www.rfc-editor.org/info/rfc7084>>.
- [RS-485] "Electronic Industries Association (1983). Electrical Characteristics of Generators and Receivers for Use in Balanced Multipoint Systems. EIA Standard RS-485."
- [SCADA] "Boyer, Stuart A. (2010). SCADA Supervisory Control and Data Acquisition. USA: ISA - International Society of Automation."
- [ULA-IN-WILD] "G. Michaelson, "conference.apnic.net/data/36/apnic-36-ula\_1377495768.pdf" ".

## Authors' Addresses

Bing Liu  
Huawei Technologies  
Q14, Huawei Campus, No.156 Beiqing Road  
Hai-Dian District, Beijing, 100095  
P.R. China

Email: [leo.liubing@huawei.com](mailto:leo.liubing@huawei.com)

Sheng Jiang  
Huawei Technologies  
Q14, Huawei Campus, No.156 Beiqing Road  
Hai-Dian District, Beijing, 100095  
P.R. China

Email: [jiangsheng@huawei.com](mailto:jiangsheng@huawei.com)

v6ops  
Internet-Draft  
Intended status: Best Current Practice  
Expires: January 1, 2018

J. Brzozowski  
Comcast Cable  
D. Schinazi  
S. Cheshire  
Apple Inc.  
L. Colitti  
E. Kline  
J. Linkova  
Google  
M. Keane  
Microsoft  
P. Saab  
Facebook  
June 30, 2017

Incremental Deployment of IPv6-only Wi-Fi for IETF Meetings  
draft-jjmb-v6ops-ietf-ipv6-only-incremental-00

Abstract

The purpose of this document is to provide a blueprint and guidance for deploying IPv6-only Wi-Fi at IETF meetings. This document outlines infrastructure and operational guidance that operators should consider when deploying IPv6-only networks using NAT64 and DNS64 to support communication to legacy IPv4-only services.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 1, 2018.

## Copyright Notice

Copyright (c) 2017 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. Design Principles . . . . .	4
2.1. Network Infrastructure . . . . .	4
3. Network Services . . . . .	5
3.1. DNS64 . . . . .	5
3.2. NAT64 . . . . .	6
3.3. DHCPv6 . . . . .	6
4. User Equipment . . . . .	7
4.1. Host Address Assignment and Configuration . . . . .	7
4.2. IPv4 support . . . . .	7
5. Network Management . . . . .	8
6. Telemetry and Monitoring . . . . .	9
7. Support for User Applications and Services . . . . .	10
8. Support and Operations . . . . .	10
8.1. Reporting Issues (Ticketing) . . . . .	10
8.2. Interactive Support . . . . .	11
9. Known Client-side Issues . . . . .	11
10. IANA Considerations . . . . .	11
10.1. Security Considerations . . . . .	12
11. Future Work . . . . .	12
12. Related Industry Efforts . . . . .	12
13. References . . . . .	13
13.1. Normative References . . . . .	13
13.2. Informative References . . . . .	14
Authors' Addresses . . . . .	14

## 1. Introduction

The purpose of this document is to provide a blueprint and guidance for deploying IPv6-only Wi-Fi at IETF meetings. This document outlines infrastructure and operational guidance that operators should consider when deploying IPv6-only networks using NAT64 and DNS64 to support communication to legacy IPv4-only services.

One of the main strengths of the IETF has always been an insistence on running code. As such, IETF meetings were one of the first deployments of a dual-stack network to help test the first implementations of IPv6. Many years later, as several networks are shifting towards IPv6-only, it is the responsibility of the IETF to lead the trend and make their main network IPv6-only.

This document outlines the requirements and design principles for an IPv6-only network infrastructure that includes support for IPv4-only content. It also discusses techniques and requirements for network management, telemetry, and the operations and support for the IPv6-only network. Recommendations and best practices for operations and support will be provided, however, alternate approaches may be utilized. Disabling or removal of IPv4 stacks is out of scope for this document. This document focuses on the explicit provisioning of IPv6-only using NAT64 [RFC6146] and DNS64 [RFC6147] to access IPv4-only content and services.

### 1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in "Key words for use in RFCs to Indicate Requirement Levels" [RFC2119].

## 2. Design Principles

### 2.1. Network Infrastructure

The following are specific network design details that are minimally required to support an IPv6-only network that utilize NAT64 and DNS64. The following have been drawn from real deployment scenarios for large scale uses of IPv6-only with NAT64 and DNS64. The parameters specified here are specific to providing IPv6-only connectivity. It is assumed that IPv6-only is provisioned and that IPv4 stacks remain active on network and host interfaces. The disabling or removal of IPv4 stacks from hosts or routers is out of scope for this document. As such, it is important to note that link local IPv4 [RFC3927] will likely remain active and will appear on hosts and network infrastructure.

The following section outlines the requirement to provisioning IPv6-only. We minimally assume that SLAAC will be utilized, however, for completeness the parameters required for DHCPv6 [RFC3315] and [RFC3736] are also provided:

- o IPv6-only hosts are expected to be provisioned with IPv6-only connectivity, however, link local IPv4 is likely to be present.
- o RA interval is RECOMMENDED to be minimally set to 600 seconds per the guidance outlined in [RFC7772].
- o Support for solicited unicast router advertisements are also recommended per [RFC7772]
- o At least one prefix information option (PIO) MUST be included in router advertisements, the transmitted PIO MUST correspond to the IPv6 prefix that is valid for a given IPv6 link.
- o The use of SLAAC [RFC4862] MUST be signalled by the network, specifically for each transmitted PIO the A bit MUST be set to one.
- o DHCPv6 support SHOULD be included to support legacy operating systems that do not support DNS RA options but is not required. Whether stateless or stateful DHCPv6 is used, both the DNS Server IPv6 address and DNS Search List options [RFC8106] MUST minimally be included. The DNS server IPv6 address(es) MUST be those used for DNS64. It is RECOMMENDED that these values be identical to those used in the IPv6 router advertisements that include the DNS options [RFC8106]. If DHCPv6 support is deployed, stateless DHCPv6 MUST minimally be available.

- o IPv6 router advertisements MUST include the DNS options [RFC8106]. Both the DNS Server IPv6 address(es) and DNS Search List are REQUIRED. If DHCPv6 support is deployed the values sent here for DNS RA options are RECOMMENDED to match those sent via DHCPv6.

To ensure seamless and to support an incremental deployment of IPv6-only access to legacy dual stack infrastructure should remain available. The following are recommended approaches that may be considered to achieve the same.

The deployment of IPv6-only with NAT64 and DNS64 may very well help to identify applications, services, or use cases that are not entirely compatible with the same. It is therefore important to ensure that users of IP networks, whether wired or wireless, have access to legacy dual stack infrastructure as a fallback. For wireless network it is recommend to have a secondary SSID labelled accordingly, e.g. example-ssid-dual-stack or example-ssid-legacy. For wired network connectivity having secondary ports that are dual stack enabled is also recommended. Note that while it is recommended to ensure the presence of a fallback network, the goal remains to make the IPv6-only network the primary network.

This document assumes that dual stack connectivity is available by default and that IPv4-only connectivity is no longer supported. As such, it is out of scope for this document to outline fallback or access to legacy connectivity that is IPv4-only.

### 3. Network Services

The following network services are required for an IPv6-only where support for and access to IPv4 content, services, and applications are required.

#### 3.1. DNS64

The following recommendations apply to the use and deployment of DNS64:

- o Use of the well known DNS64 prefix per [RFC6052]
- o It is also recommended that query logging be enabled for DNS64, performance impacts of query logging must be noted but are largely out of scope for this document. Query logging is essential to determine the volume and make up of DNS queries and replies that are specific to DNS64 and IPv4-only content, services, and applications.

### 3.2. NAT64

The following recommendations apply to the use and deployment of NAT64:

- o DNS64 is a critical aspect to direct requests from IPv6-only hosts to a NAT64 service.
- o NAT64 configurations vary widely, port allocation techniques are largely out of scope for this document. One-to-one (1:1) mappings can be used to allocate an IPv4 address per connected device or alternatively blocks of IPv4 ports can also be assigned per device, each has different properties. It is generally recommended to allocate IPv4 ports per device in an effort to maximize IPv4 utilization for NAT64.

### 3.3. DHCPv6

Support for DHCPv6 may be required in some deployments. If required, parameters pertaining to IPv6 router discovery may require adjustment. The following outlines the guidance specific to the use of DHCPv6:

- o Stateless DHCPv6 SHOULD be supported to facilitate the transmission of DNS servers IPv6 address(es) and DNS search lists to legacy hosts that do not support DNS RA options.
- o Stateful DHCPv6 for address assignment MAY be supported, but is not required. If stateful DHCPv6 is used the DNS parameters mentioned above MUST be included.
- o If, at some future date, support for IPv6 prefix delegation becomes necessary, stateful DHCPv6 will likely be mandatory (Future Work (Section 11)). The details of IPv6 prefix delegation are out of scope for this document.

## 4. User Equipment

### 4.1. Host Address Assignment and Configuration

- o Hosts MUST support SLAAC.
- o Hosts SHOULD support DNS RA options [RFC8106] for the acquisition of DNS server IPv6 addresses and a DNS Search List.
- o Hosts MAY support DHCPv6 for address acquisition, the use of DHCPv6 for address acquisition is not prohibited.
- o DHCPv6 option to configure DNS server option 23 and domain search list option 24 [RFC3646] address MUST be implemented if DHCPv6 is to be utilized.

### 4.2. IPv4 support

The IPv4 stacks of hosts MAY remain enabled, which means that Link Local IPv4 [RFC3927] (169.254/16) addresses MAY continue to be present and in use. Disabling of the IPv4 stack of hosts is out of scope for this document.

Host operating systems SHOULD provide a means for applications to easily connect to IPv4-only servers by using the NAT64/DNS64. While modern applications simply need to make AAAA queries and connect to the resulting IPv6 address, operating systems SHOULD provide simple ways for applications to do so or even connect to IPv4 literals in the absence of host names. Possible solutions include 464XLAT [RFC6877], "Bump-in-the-Host" [RFC6535] and Happy Eyeballs v2 [HEv2].

Finally, it is RECOMMENDED that support for DHCPv4 be explicitly suppressed in particular to prevent the inadvertent assignment of IPv4 addresses on networks that do not have a valid IPv4 egress. DHCPv4 servers, rogue or otherwise, could adversely impact the experience of end users of the IPv6-only network.



## 5. Network Management

The focus of this document is user equipment and hosts. The network and network service requirements are oriented around providing IPv6-only connectivity that allows for the use of NAT64 and DNS64 to maintain reachability to IPv4-only content, applications, and services. Operations and management of the underlying network is technically out of scope for this document, however, given the relevance of the same to the focus of this draft some guidance is being provided.

Strictly speaking the primary requirement for the underlying network is that IPv6 is supported along with the services required to enable the use of NAT64 and DNS64. This suggests that the underlying network could in fact be dual stack for management and operations. It is required that the provisioning of IPv4 for user equipment and host connectivity not be supported. User equipment or host facing interfaces **MUST NOT** acquire non-link-local IPv4 addresses or IPv4 DNS server addresses. Additionally, the network **MUST NOT** respond to DHCPv4 requests or DNS queries sent over IPv4.

Given the above, within a given VLAN it is possible and likely that IPv4 may be observed, present, and possibly used. It is out of scope for this document to prevent the use of IPv4 entirely.

Depending on the level of readiness IPv6-only network management may or may not be possible. Network management and operations includes but is not limited to the following:

- o Remote access to network infrastructure via SSH or telnet
- o Remote SNMP communications
- o Remote NETCONF communications
- o Remote Syslog communications

While it is strongly recommended that all network management and operations be performed over IPv6-only it is not strictly required. However, it is important to note that the presence and use of IPv4 for network management and operations must not impede or impact the use of IPv6-only with NAT64 and DNS64.

## 6. Telemetry and Monitoring

At this point in time, IPv6-only networks with no IPv4 support at all are still not widespread and may expose issues in host operating systems or applications. It is therefore recommended that telemetry summarizing how hosts are being provisioned and accessing the Internet be collected and analyzed. In order to preserve the privacy of users of the network, it is paramount that connectivity information (e.g. DNS64 records) cannot be correlated with individual client nodes.

We can measure how hosts:

- o Configure IPv6 addresses (SLAAC, DHCPv6) and which ones they use
- o Configure DNS server addresses (DNS RA options vs DHCPv6)

We can measure what percentage of the traffic:

- o Uses native IPv6
- o Uses NAT64

Recording the most common hostnames that require the DNS64 would also allow operators to establish a list of the most prominent IPv4-only services.

Observing the TCP/UDP ports used by applications that still leverage IPv4 link-local on an IPv6-only network will also help prepare for the time when routers stop supporting IPv4 communications altogether.

Given that some users may have devices running legacy IPv4-only software, the network should provide a different fallback network that is dual-stack. It is worth measuring the number of users that switch to this network, and possibly use an anonymous survey asking users what software failure caused them to switch. Additionally, the fallback network SHOULD use different authentication credentials per meeting (such as SSID) to make sure a failure causing a user to switch does not mean they will stay on the fallback network forever.

## 7. Support for User Applications and Services

Following is a list of commonly used applications and services that are expected to operate, without incident, when used in an IPv6-only environment that utilizes NAT64 and DNS64. The list below is not exhaustive.

- o VPN
- o Chat
- o Email
- o SSH/Telnet
- o Git
- o Voice

## 8. Support and Operations

Most every network has customers or end users of some sort, therefore it is essential to ensure that end users or consumers of the network have means to do the following while transitions are occurring in networks and related infrastructure. One key item referenced earlier is the availability of temporary fallback networks that support legacy communications.

The following outline additional items that end users must have available to communicate with network operators. All of the items below must be available via dual stack connectivity.

### 8.1. Reporting Issues (Ticketing)

Tools and systems that can be used to report issues with applications, services, or content must be available for end-users. Network and systems operators are responsible for acknowledging and classifying issues and ultimately ensuring that the same are properly addressed. Specifically to this document "fixed" is meant to imply that proper support for IPv6 is available. In some cases network and system operators may need to implement temporary workarounds to ensure that end users can access the desired content, application, or service.

In order for users experiencing IPv6-specific issues to be able to report them, the ticketing system **MUST** also be reachable over the dual-stack fallback network. The existence of the fallback network **SHOULD** also be made clear to users ahead of time. In order to help narrow down issues, the ticketing system **SHOULD** ask the user whether the issue is specific to IPv6-only and whether they have experienced the issue or a different outcome on the fallback network.

## 8.2. Interactive Support

Interactive support is often desired in lieu or in conjunction with traditional support models like trouble ticket creation. It is recommended that interactive support be available via real time and near real time mechanisms like Slack or electronic mail (e-mail).

## 9. Known Client-side Issues

Following are known client side issues that are specific to the deployment of IPv6-only networks and/or the use of NAT64/DNS64:

- o Use of literal IPv4 addresses - the use of literal IPv4 addresses is a known issue given the approach that is documented in this I-D. Addressing the use of literal IPv4 addresses is out of scope for this document.
- o Applications that explicitly require IPv4 by only performing AAAA queries or restricting the type of underlying socket they use.
- o Unreachable but valid AAAA RR in the DNS - in some cases a valid AAAA RR is returned by the DNS, however, if the same is unreachable or is not configured the presence of the same will prevent a DNS64 query which in turn prevents the use of the NAT64 to reach the target host references by the address in the AAAA DNS RR.

## 10. IANA Considerations

This memo includes no request to IANA.

### 10.1. Security Considerations

The vastness of the IPv6 address space often makes it more difficult to scan the same unlike legacy IPv4-only or dual stack IP networks. It is conceivable that IPv6-only network represent a reduction in attack surface area which in turn could be viewed a security improvement compared to IPv4-only or dual stack IP networks.

Given the criticality of the DNS64 for reachability to the NAT64, poisoning of one or both could represent a vector for the attack of the DNS64 and NAT64 which could in turn impact the end user experience. Worse poisoning of the DNS64 and/or NAT64 could result in redirection of end use devices to malicious hosts. It is likely that this vulnerability is no greater in IPv6-only networks utilizing DNS64 and NAT64 compared to traditional IPv4-only or dual stack networks.

### 11. Future Work

The following items are out of scope for this document, however, the following are listed as future work items specific to incremental IPv6-only deployments:

- o Support for IPv6 prefix delegation
- o Disabling IPv4 stacks at some point in the future
- o Fully deprecating the fallback legacy IPv4 network

### 12. Related Industry Efforts

- o Comcast new building and IPv6-only (John Jason Brzozowski <john\_brzozowski@comcast.com>)
- o Microsoft corporate IT IPv6-only (Marcus Keane <marcus.keane@microsoft.com>)
- o Google (Jen Linkova <furry@google.com>)

## 13. References

### 13.1. Normative References

- [HEv2] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2", Work in Progress, draft-ietf-v6ops-rfc6555bis, June 2017.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, DOI 10.17487/RFC3646, December 2003, <<http://www.rfc-editor.org/info/rfc3646>>.
- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, DOI 10.17487/RFC3736, April 2004, <<http://www.rfc-editor.org/info/rfc3736>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<http://www.rfc-editor.org/info/rfc6052>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <<http://www.rfc-editor.org/info/rfc6146>>.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, DOI 10.17487/RFC6147, April 2011, <<http://www.rfc-editor.org/info/rfc6147>>.

- [RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy Consumption of Router Advertisements", BCP 202, RFC 7772, DOI 10.17487/RFC7772, February 2016, <<http://www.rfc-editor.org/info/rfc7772>>.
- [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, <<http://www.rfc-editor.org/info/rfc8106>>.

### 13.2. Informative References

- [RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, DOI 10.17487/RFC3927, May 2005, <<http://www.rfc-editor.org/info/rfc3927>>.
- [RFC6535] Huang, B., Deng, H., and T. Savolainen, "Dual-Stack Hosts Using "Bump-in-the-Host" (BIH)", RFC 6535, DOI 10.17487/RFC6535, February 2012, <<http://www.rfc-editor.org/info/rfc6535>>.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <<http://www.rfc-editor.org/info/rfc6877>>.

### Authors' Addresses

John Jason Brzozowski  
Comcast Cable  
1701 John F. Kennedy Blvd.  
Philadelphia, PA  
USA

Email: [john\\_brzozowski@cable.comcast.com](mailto:john_brzozowski@cable.comcast.com)

David Schinazi  
Apple Inc.  
1 Infinite Loop  
Cupertino, California 95014  
US

Email: [dschinazi@apple.com](mailto:dschinazi@apple.com)

Stuart Cheshire  
Apple Inc.  
1 Infinite Loop  
Cupertino, California 95014  
USA

Email: cheshire@apple.com

Lorenzo Colitti  
Google

Email: lorenzo@google.com

Erik Kline  
Google

Email: ek@google.com

Jen Linkova  
Google

Email: furry@google.com

Marcus Keane  
Microsoft

Email: marcus.keane@microsoft.com

Paul Saab  
Facebook

Email: ps@fb.com



IPv6 Operations  
Internet-Draft  
Intended status: Informational  
Expires: January 3, 2018

J. Linkova  
Google  
M. Stucchi  
July 2, 2017

Using Conditional Router Advertisements for Enterprise Multihoming  
draft-linkova-v6ops-conditional-ras-01

Abstract

This document discusses most common scenarios of connecting an enterprise network to multiple ISPs using an address space assigned by an ISP. The problem of enterprise multihoming without address translation of any form has not been solved yet as it requires both the network to select the correct egress ISP based on the packet source address and hosts to select the correct source address based on the desired egress ISP for that traffic.

[I-D.ietf-rtgwg-enterprise-pa-multihoming] proposes a solution to this problem by introducing a new routing functionality (Source Address Dependent Routing) to solve the uplink selection issue and using Router Advertisements to influence the host source address selection. While the above-mentioned document focuses on solving the general problem and on covering various complex use cases, this document describes how the solution proposed in

[I-D.ietf-rtgwg-enterprise-pa-multihoming] can be adopted for limited number of common use cases. In particular, the focus is on scenarios where an enterprise network has two Internet uplinks used either in primary/backup mode or simultaneously and hosts in that network might not yet properly support multihoming as described in [RFC8028].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 3, 2018.

## Copyright Notice

Copyright (c) 2017 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 . . . . .	2
2. Common Enterprise Multihoming Scenarios . . . . .	3
2.1. Two ISP Uplinks, Primary and Backup . . . . .	3
2.2. Two ISP Uplinks, Used for Load Balancing . . . . .	4
3. Conditional Router Advertisements . . . . .	4
3.1. Solution Overview . . . . .	4
3.1.1. Uplink Selection . . . . .	4
3.1.2. Source Address Selection and Conditional RAs . . . . .	4
3.2. Example Scenarios . . . . .	6
3.2.1. Single Router, Primary/Backup Uplinks . . . . .	6
3.2.2. Two Routers, Primary/Backup Uplinks . . . . .	7
3.2.3. Single Router, Load Balancing Between Uplinks . . . . .	9
3.2.4. Two Router, Load Balancing Between Uplinks . . . . .	10
3.2.5. Topologies with Dedicated Border Routers . . . . .	10
4. IANA Considerations . . . . .	12
5. Security Considerations . . . . .	12
5.1. Privacy Considerations . . . . .	12
6. Acknowledgements . . . . .	12
7. References . . . . .	12
7.1. Normative References . . . . .	12
7.2. Informative References . . . . .	14
Appendix A. Change Log . . . . .	15
Authors' Addresses . . . . .	15

## 1. Introduction

Multihoming is an obvious requirement for many enterprise networks to ensure the desired level of network reliability. However, using more than one ISP (and address space assigned by those ISPs) introduces the problem of assigning IP addresses to hosts. In IPv4 there is no choice but using [RFC1918] address space and NAT ([RFC3022]) at the

network edge. Using Provider Independent or PI address space is not always an option as it requires running BGP between the enterprise network and the ISPs). As IPv6 host can, by design, have multiple addresses of the global scope, multihoming using provider address looks even easier for IPv6: each ISP assigns an IPv6 block (usually /48) and hosts in the enterprise network have addresses assigned from each ISP block. However using IPv6 PA blocks in multihoming scenario introduces some challenges, including but not limited to:

- o Selecting the correct uplink based on the packet source address;
- o Signaling to hosts that some source addresses should or should not be used (e.g. an uplink to the ISP went down or became available again).

The document [I-D.ietf-rtgwg-enterprise-pa-multihoming] discusses these and other related challenges in details in relation to the general multihoming scenario for enterprise networks. Unfortunately the proposed solution heavily relies on the rule 5.5 of the default address selection algorithm ([RFC6724]) which has not been widely implemented at the moment this document was written. Therefore network administrators in enterprise networks can't yet assume that all devices in their network support the rule 5.5, especially in the quite common BYOD ("Bring Your Own Device") scenario. However, while it does not seem feasible to solve all the possible multihoming scenarios without relying on rule 5.5, it is possible to provide IPv6 multihoming using provider-assigned (PA) address space for the most common use cases. This document discusses how the general solution described in [I-D.ietf-rtgwg-enterprise-pa-multihoming] can be applied to those two specific cases.

## 2. Common Enterprise Multihoming Scenarios

### 2.1. Two ISP Uplinks, Primary and Backup

This scenario has the following key characteristics:

- o The enterprise network is using uplinks to two (or more) ISPs for Internet access;
- o Each ISP assigns IPv6 PA address space for the network;
- o Uplink(s) to one ISP is a primary (preferred) one. All other uplinks are backup and are not expected to be used while the primary one is operational;
- o If the primary uplink is operational, all Internet traffic should flow via that uplink;

- o When the primary uplink fails the Internet traffic needs to flow via the backup uplinks;
- o Recovery of the primary uplink needs to trigger the traffic switchover from the backup uplinks back to primary one.

## 2.2. Two ISP Uplinks, Used for Load Balancing

This scenario has the following key characteristics:

- o The enterprise network is using uplinks to two (or more) ISPs for Internet access;
- o Each ISP assigns an IPv6 PA address space;
- o All the uplinks may be used simultaneously, with the traffic being randomly balanced between them.

## 3. Conditional Router Advertisements

### 3.1. Solution Overview

#### 3.1.1. Uplink Selection

As discussed in [I-D.ietf-rtgwg-enterprise-pa-multihoming], one of the two main problems to be solved in the enterprise multihoming scenario is the problem of the next-hop (uplink) selection based on the packet source address. For example, if the enterprise network has two uplinks, to ISP\_A and ISP\_B, and hosts have addresses from subnet\_A and subnet\_B (belonging to ISP\_A and ISP\_B respectively) then packets sourced from subnet\_A must be sent to ISP\_A uplink while packets sourced from subnet\_B must be sent to ISP\_B uplink.

While some work is being done in the Source Address Dependent Routing (SADR) area, the simplest way to implement the desired functionality currently is to apply a policy which selects a next-hop or an egress interface based on the packet source address. Most of the SMB/Enterprise grade routers have such functionality available currently.

#### 3.1.2. Source Address Selection and Conditional RAs

Another problem to be solved in the multihoming scenario is the source address selection on hosts. In the normal situation (all uplinks are up/operational) hosts have multiple global unique addresses and can rely on the default address selection algorithm ([RFC6724]) to pick up a source address, while the network is responsible for choosing the correct uplink based on the source address selected by a host as described in Section 3.1.2. However,

some network topology changes (i.e. changing uplink status) might affect the global reachability for packets sourced from the particular prefixes and therefore such changes have to be signaled back to the hosts. For example:

- o An uplink to an ISP\_A went down. Hosts should not use addresses from ISP\_A prefix;
- o A primary uplink to ISP\_A which was not operational has come back up. Hosts should start using the source addresses from ISP\_A prefix.

[I-D.ietf-rtgwg-enterprise-pa-multihoming] provides a detailed explanation on why SLAAC and router advertisements are the most suitable mechanism for signaling network topology changes to hosts and thereby influencing the source address selection. Sending a router advertisement to change the preferred lifetime for a given prefix provides the following functionality:

- o deprecating addresses (by sending an RA with the preferred\_lifetime set to 0 in the corresponding POI) to indicate to hosts that that addresses from that prefix should not be used;
- o making a previously unused (deprecated) prefix usable again (by sending an RA containing a POI with non-zero preferred lifetime) to indicate to hosts that addresses from that prefix can be used again.

To provide the desired functionality, first-hop routers are required to

- o send RA triggered by defined event policies in response to uplink status change event; and
- o while sending periodic or solicited RAs, set the value in the given RA field (e.g. PIO preferred lifetime) based on the uplink status.

The exact definition of the 'uplink status' depends on the network topology and may include conditions like:

- o uplink interface status change;
- o presence of a particular route in the routing table;
- o presence of a particular route with a particular attribute (next-hop, tag etc) in the routing table;

- o protocol adjacency change.

etc.

In some scenarios, when two routers are providing first-hop redundancy via VRRP, the master-backup status can be considered as a condition for sending RAs and changing the preferred lifetime value. See Section 3.2.2 for more details.

If hosts are provided with ISP DNS servers IPv6 addresses via RDNSS [RFC8106] it might be desirable for the conditional RAs to update the Lifetime field of the RDNSS option as well.

### 3.2. Example Scenarios

This section illustrates how the conditional RAs solution can be applied to most common enterprise multihoming scenarios.

#### 3.2.1. Single Router, Primary/Backup Uplinks




Figure 1: Single Router, Primary/Backup Uplinks

Let's look at a simple network topology where a single router acts as a border router to terminate two ISP uplinks and as a first-hop router for hosts. Each ISP assigns a /48 to the network, and the ISP\_A uplink is a primary one, to be used for all Internet traffic, while the ISP\_B uplink is a backup, to be used only when the primary uplink is not operational.

To ensure that packets with source addresses from ISP\_A and ISP\_B are only routed to ISP\_A and ISP\_B uplinks respectively, the network administrator needs to configure a policy on R1:

```
if {
    packet_destination_address is not in 2001:db8:1::/48 or 2001:db8:2::/48
    packet_source_address is in 2001:db8:1::/48
} then {
    next-hop is ISP_A_uplink
}
if {
    packet_destination_address is not in 2001:db8:1::/48 or 2001:db8:2::/48
    packet_source_address is in 2001:db8:2::/48
}
then {
    next-hop is ISP_B_uplink
}
```

Under normal circumstances it is desirable that all traffic be sent via the ISP\_A uplink, therefore hosts (the host H1 in the example topology figure) should be using source addresses from 2001:db8:1:1::/64. When/if ISP\_A uplink fails, hosts should stop using the 2001:db8:1:1::/64 prefix and start using 2001:db8:2:1::/64 until the ISP\_A uplink comes back up. To achieve the desired behavior the router advertisement configuration on the R1 device for the interface facing H1 needs to have the following policy:

```
prefix 2001:db8:1:1::/64 {
    if ISP_A_uplink is up
        then preferred_lifetime = 604800
    else preferred_lifetime = 0
}

prefix 2001:db8:2:1::/64 {
    if ISP_A_Uplink is up
        then preferred_lifetime = 0
    else preferred_lifetime = 604800
}
```

A similar policy needs to be applied to the RDNSS Lifetime if ISP\_A and ISP\_B DNS servers are used.

### 3.2.2. Two Routers, Primary/Backup Uplinks

Let's look at a more complex scenario where two border routers are terminating two ISP uplinks (one each), acting as redundant first-hop routers for hosts. The topology is shown on Fig.2




Figure 2: Two Routers, Primary/Backup Uplinks

In this scenario R1 sends RAs with PIO for 2001:db8:1:1::/64 (ISP\_A address space) and R2 sends RAs with PIO for 2001:db8:2:1::/64 (ISP\_B address space). Each router needs to have a forwarding policy configured for packets received on its hosts-facing interface:

```

if {
    packet_destination_address is not in 2001:db8:1::/48 or 2001:db8:2::/48
    packet_source_address is in 2001:db8:1::/48
} then {
    next-hop is ISP_A_uplink
}
if {
    packet_destination_address is not in 2001:db8:1::/48 or 2001:db8:2::/48
    packet_source_address is in 2001:db8:2::/48
} then {
    next-hop is ISP_B_uplink
}

```

In this case there is more than one way to ensure that hosts are selecting the correct source address based on the uplink status. If VRRP is used to provide first-hop redundancy and the master router is the one with the active uplink, then the simplest way is to use the VRRP mastership as a condition for router advertisement. So, if ISP\_A is the primary uplink, the routers R1 and R2 need to be configured in the following way:

R1 is the VRRP master by default (when ISP\_A uplink is up). If ISP\_A uplink is down, then R1 becomes a backup. Router advertisements on R1's interface facing H1 needs to have the following policy applied:



```
prefix 2001:db8:1:1::/64 {  
    if vrrp_master then preferred_lifetime = 604800  
    else preferred_lifetime = 0  
}
```

R2 is VRRP backup by default. Router advertisement on R2 interface facing H1 needs to have the following policy applied:

```
prefix 2001:db8:2:1::/64 {  
    if vrrp_master then preferred_lifetime = 604800  
    else preferred_lifetime = 0  
}
```

If VRRP is not used or interface status tracking is not used for mastership switchover, then each router needs to be able to detect the uplink failure/recovery on the neighboring router, so that RAs with updated preferred lifetime values are triggered. Depending on the network setup various triggers like a route to the uplink interface subnet or a default route received from the uplink can be used. The obvious drawback of using the routing table to trigger the conditional RAs is that some additional configuration is required. For example, if a route to the prefix assigned to the ISP uplink is used as a trigger, then the conditional RA policy would have the following logic:

R1:

```
prefix 2001:db8:1:1::/64 {  
    if ISP_A_uplink is up then preferred_lifetime = 604800  
    else preferred_lifetime = 0  
}
```

R2:

```
prefix 2001:db8:2:1::/64 {  
    if ISP_A_uplink_route is present then preferred_lifetime = 0  
    else preferred_lifetime = 604800  
}
```

### 3.2.3. Single Router, Load Balancing Between Uplinks

Let's look at the example topology shown in Figure 1, but with both uplinks used simultaneously. In this case R1 would send RAs containing PIOs for both prefixes, 2001:db8:1:1::/64 and 2001:db8:2:1::/64, changing the preferred lifetime based on particular uplink availability. If the interface status is used as uplink availability indicator, then the policy logic would look like the following:

```
prefix 2001:db8:1:1::/64 {  
    if ISP_A_uplink is up then preferred_lifetime = 604800  
    else preferred_lifetime = 0  
}  
prefix 2001:db8:2:1::/64 {  
    if ISP_B_uplink is up then preferred_lifetime = 604800  
    else preferred_lifetime = 0  
}
```

R1 needs a forwarding policy to be applied to forward packets to the correct uplink based on the source address as described in Section 3.2.1.

#### 3.2.4. Two Router, Load Balancing Between Uplinks

In this scenario the example topology is similar to the one shown in Figure 2, but both uplinks can be used at the same time. It means that both R1 and R2 need to have the corresponding forwarding policy to forward packets based on their source addresses.

Each router would send RAs with POI for the corresponding prefix, setting preferred\_lifetime to a non-zero value when the ISP uplink is up, and deprecating the prefix by setting the preferred lifetime to 0 in case of uplink failure. The uplink recovery would trigger another RA with non-zero preferred lifetime to make the addresses from the prefix preferred again. The example RA policy on R1 and R2 would look like:

R1:

```
prefix 2001:db8:1:1::/64 {  
    if ISP_A_uplink is up then preferred_lifetime = 604800  
    else preferred_lifetime = 0  
}
```

R2:

```
prefix 2001:db8:2:1::/64 {  
    if ISP_B_uplink is up then preferred_lifetime = 604800  
    else preferred_lifetime = 0  
}
```

#### 3.2.5. Topologies with Dedicated Border Routers

For simplicity reasons all topologies below show the ISP uplinks terminated on the first-hop routers. Obviously, the proposed approach can be used in more complex topologies when dedicated devices are used for terminating ISP uplinks. In that case VRRP

mastership or interface status can not be used as a trigger for conditional RAs and route presence as described above should be used instead.

Let's look at the example topology shown on the Figure 3:

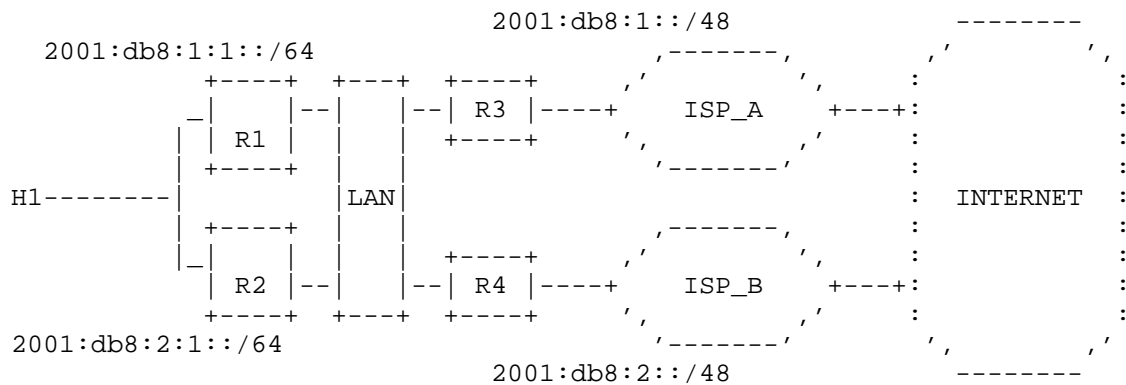


Figure 3: Dedicated Border Routers

For example, if ISP\_A is a primary uplink and ISP\_B is a backup one then the following policy might be used to achieve the desired behaviour (H1 is using ISP\_A address space, 2001:db8:1:1::/64 while ISP\_A uplink is up and only using ISP\_B 2001:db8:2:1::/64 prefix if the uplink is non-operational):

R1 and R2 policy:

```
prefix 2001:db8:1:1::/64 {
  if ISP_A_uplink_route is present then preferred_lifetime = 604800
  else preferred_lifetime = 0
}
prefix 2001:db8:2:1::/64 {
  if ISP_A_uplink_route is present then preferred_lifetime = 0
  else preferred_lifetime = 604800
}
```

For load-balancing case the policy would look slightly different: each prefix has non-zero preferred\_lifetime only if the corresponding ISP uplink route is present:

```
prefix 2001:db8:1:1::/64 {  
  if ISP_A_uplink_route is present then preferred_lifetime = 604800  
  else preferred_lifetime = 0  
}  
prefix 2001:db8:2:1::/64 {  
  if ISP_B_uplink_route is present then preferred_lifetime = 0  
  else preferred_lifetime = 604800  
}
```

#### 4. IANA Considerations

This memo asks the IANA for no new parameters.

#### 5. Security Considerations

##### 5.1. Privacy Considerations

#### 6. Acknowledgements

#### 7. References

##### 7.1. Normative References

- [I-D.ietf-rtgwg-enterprise-pa-multihoming]  
Baker, F., Bowers, C., and J. Linkova, "Enterprise Multihoming using Provider-Assigned Addresses without Network Prefix Translation: Requirements and Solution", draft-ietf-rtgwg-enterprise-pa-multihoming-00 (work in progress), March 2017.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<http://www.rfc-editor.org/info/rfc1918>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [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, <<http://www.rfc-editor.org/info/rfc2827>>.

- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, DOI 10.17487/RFC3022, January 2001, <<http://www.rfc-editor.org/info/rfc3022>>.
- [RFC3582] Abley, J., Black, B., and V. Gill, "Goals for IPv6 Site-Multihoming Architectures", RFC 3582, DOI 10.17487/RFC3582, August 2003, <<http://www.rfc-editor.org/info/rfc3582>>.
- [RFC4116] Abley, J., Lindqvist, K., Davies, E., Black, B., and V. Gill, "IPv4 Multihoming Practices and Limitations", RFC 4116, DOI 10.17487/RFC4116, July 2005, <<http://www.rfc-editor.org/info/rfc4116>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<http://www.rfc-editor.org/info/rfc4193>>.
- [RFC4218] Nordmark, E. and T. Li, "Threats Relating to IPv6 Multihoming Solutions", RFC 4218, DOI 10.17487/RFC4218, October 2005, <<http://www.rfc-editor.org/info/rfc4218>>.
- [RFC4219] Lear, E., "Things Multihoming in IPv6 (MULTI6) Developers Should Think About", RFC 4219, DOI 10.17487/RFC4219, October 2005, <<http://www.rfc-editor.org/info/rfc4219>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, DOI 10.17487/RFC6296, June 2011, <<http://www.rfc-editor.org/info/rfc6296>>.
- [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, <<http://www.rfc-editor.org/info/rfc7157>>.
- [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, <<http://www.rfc-editor.org/info/rfc8106>>.

## 7.2. Informative References

- [I-D.ietf-rtgwg-dst-src-routing]  
Lamparter, D. and A. Smirnov, "Destination/Source Routing", draft-ietf-rtgwg-dst-src-routing-04 (work in progress), May 2017.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<http://www.rfc-editor.org/info/rfc3704>>.
- [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, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<http://www.rfc-editor.org/info/rfc4941>>.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, DOI 10.17487/RFC5533, June 2009, <<http://www.rfc-editor.org/info/rfc5533>>.
- [RFC5534] Arkko, J. and I. van Beijnum, "Failure Detection and Locator Pair Exploration Protocol for IPv6 Multihoming", RFC 5534, DOI 10.17487/RFC5534, June 2009, <<http://www.rfc-editor.org/info/rfc5534>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<http://www.rfc-editor.org/info/rfc6724>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<http://www.rfc-editor.org/info/rfc7788>>.
- [RFC8028] Baker, F. and B. Carpenter, "First-Hop Router Selection by Hosts in a Multi-Prefix Network", RFC 8028, DOI 10.17487/RFC8028, November 2016, <<http://www.rfc-editor.org/info/rfc8028>>.

Appendix A. Change Log

Initial Version: July 2017

Authors' Addresses

Jen Linkova  
Google  
Mountain View, California 94043  
USA

Email: [furry@google.com](mailto:furry@google.com)

Massimiliano Stucchi

Email: [max@stucchi.ch](mailto:max@stucchi.ch)

v6ops  
Internet-Draft  
Intended status: Standards Track  
Expires: January 4, 2018

J. Palet Martinez  
Consulintel, S.L.  
July 3, 2017

Reporting of Happy Eyeballs v2 Failures  
draft-palet-ietf-v6ops-he-reporting-00

Abstract

This document describes an extension to Happy Eyeballs in order to report IPv6 failures that force the fall-back to IPv4.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 4, 2018.

Copyright Notice

Copyright (c) 2017 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 . . . . .	2
2. Using Syslog . . . . .	2
3. Discovery of the syslog collector NSP . . . . .	3
4. HE behaviour on failure detection . . . . .	3
5. Privacy Considerations . . . . .	3
6. Security Considerations . . . . .	3
7. IANA Considerations . . . . .	3
8. Acknowledgements . . . . .	3
9. Normative References . . . . .	4
Author's Address . . . . .	4

## 1. Introduction

Happy Eyeballs ([RFC6555]) provides a way for improving user-visible delay when IPv6 connectivity is performing worst than the IPv4 one.

However, this hides the possible IPv6 connectivity issues to the operator because users don't notice anything broken, so they aren't reporting it to their providers.

The goal of this document is to specify an extension of HE, in order to use existing protocols for providing a reporting to the operator, which can be used to setup alarms and trigger further investigation so to improve.

## 2. Using Syslog

In order to simplify the reporting of the HE failures, syslog ([RFC5424]) over UDP ([RFC5426]), MUST be used, by means of the default port (514) with IPv6-only.

The intend is to make this reporting very simple, so no choice of alternative ports or transport protocols is offered.

Operators willing to use this reporting MUST configure at least one syslog collector at the IPv6 prefix formed as:

Network-Specific Prefix::192.88.99.1

The Network-Specific Prefix (NSP) MUST be chosen by the operator from its RIR allocated IPv6 addressing space.

Additional collectors can be made available by using anycast at the NSP + 192.88.99.0/24 prefix

### 3. Discovery of the syslog collector NSP

The same mechanism described by RFC7050 ([RFC7050]) should be used to define the address of the syslog collector(s).

Because the collectors will be using an IPv6 address with the 32 low order bits from the reserved range 192.88.99.0/24, this will not be in conflict with any public addresses used in Internet, so this mechanism is compatible with the expected usage of the NSP for NAT64.

### 4. HE behaviour on failure detection

This section will specify the exact behaviour of HE in order to initiate the reporting and the specific format/parameters of the HE failure message to be sent to the syslog collector.

A preliminary consideration is to include, in addition to the syslog required parameters, the timeouts detected, the failed destination address and the source prefix from where the destination has failed.

TBD.

### 5. Privacy Considerations

The goal is to provide the operator information about the failures detected by HE, without requiring specific users traffic information. Towards this, it will be sufficient to provide to the syslog collector details about the failed destination address and source prefix. So privacy issues regarding identification of a specific device or user are avoided.

TBD.

### 6. Security Considerations

This document does not have any specific security considerations.

### 7. IANA Considerations

IANA is requested to reserve 192.88.99.0/24, which was previously released by ([RFC7526]) for this RFC.

### 8. Acknowledgements

The author would like to acknowledge the inputs of TBD ...

## 9. Normative References

- [RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, DOI 10.17487/RFC5424, March 2009, <<http://www.rfc-editor.org/info/rfc5424>>.
- [RFC5426] Okmianski, A., "Transmission of Syslog Messages over UDP", RFC 5426, DOI 10.17487/RFC5426, March 2009, <<http://www.rfc-editor.org/info/rfc5426>>.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, DOI 10.17487/RFC6555, April 2012, <<http://www.rfc-editor.org/info/rfc6555>>.
- [RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <<http://www.rfc-editor.org/info/rfc7050>>.
- [RFC7526] Troan, O. and B. Carpenter, Ed., "Deprecating the Anycast Prefix for 6to4 Relay Routers", BCP 196, RFC 7526, DOI 10.17487/RFC7526, May 2015, <<http://www.rfc-editor.org/info/rfc7526>>.

## Author's Address

Jordi Palet Martinez  
Consulintel, S.L.  
Molino de la Navata, 75  
La Navata - Galapagar, Madrid 28420  
Spain

Email: [jordi.palet@consulintel.es](mailto:jordi.palet@consulintel.es)  
URI: <http://www.consulintel.es/>

IPv6 Operations (v6ops)  
Internet-Draft  
Intended status: Informational  
Expires: April 20, 2018

J. Palet Martinez  
Consulintel, S.L.  
October 17, 2017

Transition Requirements for IPv6 Customer Edge Routers  
draft-palet-v6ops-rfc7084-bis-transition-01

Abstract

This document specifies the transition requirements for an IPv6 Customer Edge (CE) router. Specifically, this document extends the "Basic Requirements for IPv6-only Customer Edge Routers" ([RFC7084]) in order to allow the provisioning of IPv6 transition services for the hosts attached to it. The document covers several transition technologies, either for delivering IPv6 in IPv4-only access networks and specially for delivering IPv4 "as-a-service" as required in a world where IPv4 addresses are no longer available, so hosts in the customer LANs with IPv4-only or IPv6-only applications or devices, requiring to communicate with IPv4-only services at the Internet, are able to do so.

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 April 20, 2018.

Copyright Notice

Copyright (c) 2017 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 (<https://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 . . . . .	2
1.1. Requirements Language . . . . .	3
2. Terminology . . . . .	3
3. Usage Scenarios . . . . .	4
4. Architecture . . . . .	5
4.1. Current IPv4 End-User Network Architecture . . . . .	5
4.2. IPv6 End-User Network Architecture . . . . .	6
5. Requirements . . . . .	8
5.1. General Requirements . . . . .	8
5.2. LAN-Side Configuration . . . . .	8
5.3. Transition Technologies Support . . . . .	8
5.3.1. IPv4 Service Continuity in Customer LANs . . . . .	8
5.3.1.1. 464XLAT . . . . .	8
5.3.1.2. Dual-Stack Lite (DS-Lite) . . . . .	9
5.3.1.3. Lightweight 4over6 (lw4o6) . . . . .	10
5.3.1.4. MAP-E . . . . .	10
5.3.1.5. MAP-T . . . . .	11
5.3.2. Support of IPv6 in IPv4-only WAN access . . . . .	11
5.3.2.1. 6in4 . . . . .	11
5.3.2.2. 6rd . . . . .	12
5.4. IPv4 Multicast Support . . . . .	14
5.5. Security Considerations . . . . .	14
6. Acknowledgements . . . . .	14
7. ANNEX A: Code Considerations . . . . .	14
8. References . . . . .	15
8.1. Normative References . . . . .	15
8.2. Informative References . . . . .	17
Author's Address . . . . .	17

## 1. Introduction

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

These routers are based on "Basic Requirements for IPv6-only Customer Edge Routers" ([RFC7084]), so the scope of this documents is to

include also IPv4 support, at least in the LAN side.

This document covers the IP transition technologies required when ISPs have already and IPv4-only access network that they can't turn to dual-stack or IPv6-only, as well as the situation in a world where IPv4 addresses are no longer available, so the service providers need to provision IPv6-only WAN access, while at the same time ensuring that IPv4-only or IPv6-only devices or applications in the customer LANs can still reach IPv4-only devices or applications in Internet, which still don't have IPv6 support.

This document specifies the transition mechanisms to be supported by an IPv6 transition CE router, relevant provisioning or configuration information differences from [RFC7084]. Automatic provisioning of more complex topology than a single router with multiple LAN interfaces may be handled by means of HNCP ([RFC7788]), which is out of the scope of this document.

### 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 IPv6 Transition CE router that interoperates successfully with a particular subset of currently deploying and planned common IPv6 access networks.

## 2. Terminology

This document uses the same terminology as in [RFC7084], with two minor clarifications.

The term "IPv6 transition Customer Edge Router" is defined as an "IPv6 Customer Edge Router" that provides transition support to allow IPv4-IPv6 coexistence either in the WAN, the LAN or both.

The "WAN Interface" term used across this document, means that can also support link technologies based in Internet-layer (or higher-layers) "tunnels", such as tunnels IPv4-in-IPv6 or IPv6-in-IPv4.

### 3. Usage Scenarios

The IPv6 Transition CE router described in this document is expected to be used typically, in any of the following scenarios:

1. Residential/household users. Common usage is any kind of Internet access (web, email, streaming, online gaming, etc.).
2. Residential with Small Office/Home Office (SOHO). Same usage as for the first scenario.
3. Small Office/Home Office (SOHO). Same usage as for the first scenario.
4. Small and Medium Enterprise (SME). Same usage as for the first scenario.
5. Residential/household with advanced requirements. Same basic usage as for the first scenario, however there may be requirements for exporting services to the WAN (IP cameras, web, DNS, email, VPN, etc.).
6. Small and Medium Enterprise (SME) with advanced requirements. Same basic usage as for the first scenario, however there may be requirements for exporting services to the WAN (IP cameras, web, DNS, email, VPN, etc.).

The above list is not intended to be comprehensive of all the possible usage scenarios, just the main ones. In fact, combinations of the above usages are also possible, for example a residential with SOHO and advanced requirements.

The mechanisms for exporting IPv6 services are commonly "naturally" available in any IPv6 router, as when using GUA, unless they are blocked by firewall rules, which may require some manual configuration by means of a GUI and/or CLI.

However, in the case of IPv4, because the usage of private addresses and NAT, it typically requires some degree of manual configuration such as setting up a DMZ, virtual servers, or port/protocol forwarding. In general, CE routers already provide GUI and/or CLI to manually configure them, or the possibility to setup the CE in bridge mode, so another CE behind it, takes care of that. It is out of the scope of this document the definition of any requirements for that.

The main difference for an IPv6 Transition CE router to support one or several of the above indicated scenarios, is related to the packet processing capabilities, performance, even other details such as the

number of WAN/LAN interfaces, their maximum speed, memory for keeping tables or tracking connections, etc. So, it is out of the scope of this document to classify them.

For example, an SME may have just 10 employees (micro-SME), which commonly will be considered same as a SOHO, but a small SME can have up to 50 employees, or 250 for a medium one. Depending on the IPv6 Transition CE router capabilities or even how it is being configured (for instance, using SLAAC or DHCPv6), it may support even a higher number of employees if the traffic in the LANs is low, or switched by another device(s), or the WAN bandwidth requirements are low, etc. The actual bandwidth capabilities of access with technologies such as FTTH, cable and even 3GPP/LTE, allows the support of such usages, and indeed, is a very common situation that access networks and the IPv6 Transition CE provided by the service provider are the same for SMEs and residential users.

There is also no difference in terms of who actually provides the IPv6 Transition CE router. In most of the cases is the service provider, and in fact is responsible, typically, of provisioning/managing at least the WAN side. However, commonly the user has access to configure the LAN interfaces, firewall, DMZ, and many other aspects. In fact, in many cases, the user must supply, or at least can replace the IPv6 Transition CE router, which makes even more relevant that all the IPv6 Transition CE routers, support the same requirements defined in this document.

The IPv6 Transition CE router described in this document is not intended for usage in other scenarios such as bigger Enterprises, Data Centers, Content Providers, etc. So, even if the documented requirements meet their needs, may have additional requirements, which are out of the scope of this document.

#### 4. Architecture

##### 4.1. 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.

#### 4.2. 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, in the sense that is not providing transit to other external networks. However HNCP ([RFC7788]) allows support for automatic provisioning of downstream routers. Figure 1 illustrates the model topology for the end-user network.

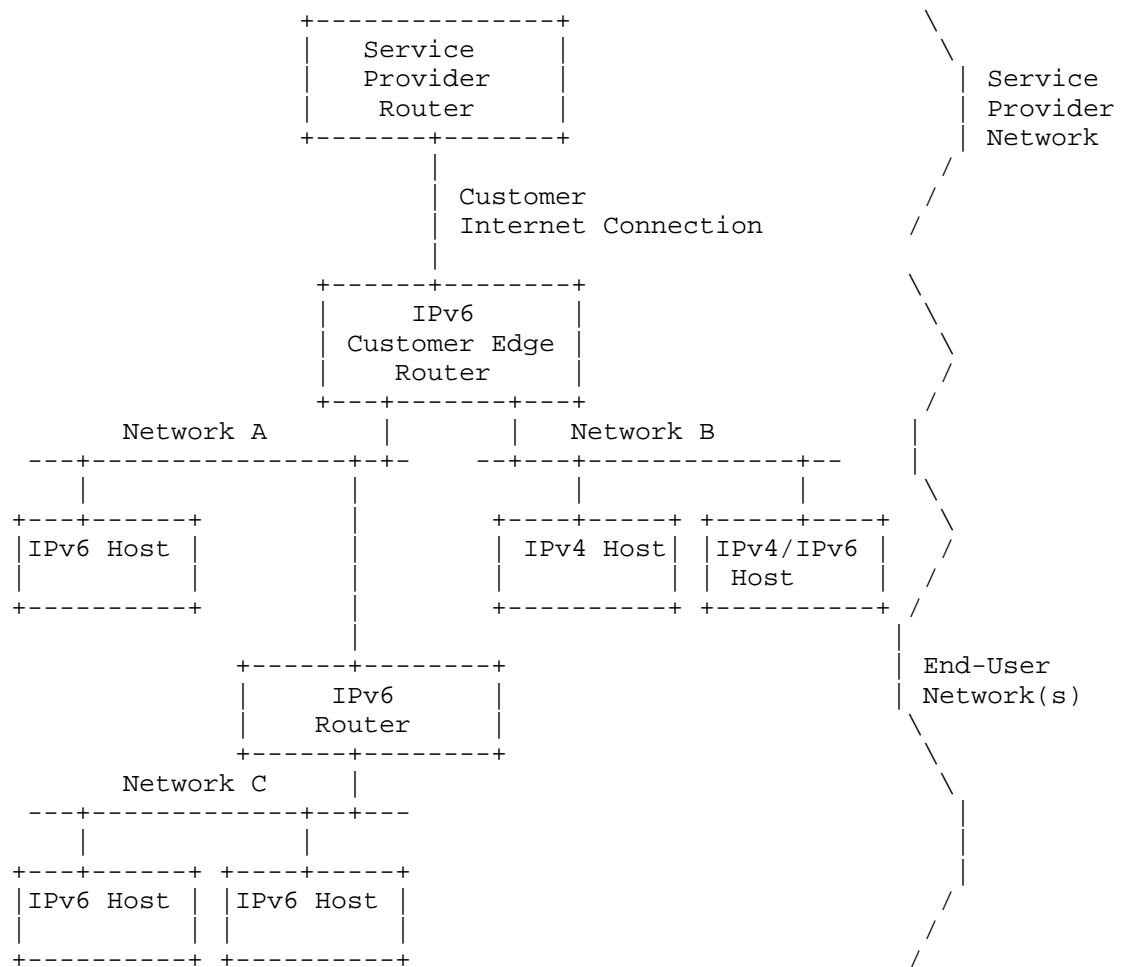


Figure 1: An Example of a Typical End-User Network

This architecture describes the:

- o Basic capabilities of an IPv6 Transition CE router
- o Provisioning of the WAN interface connecting to the service provider
- o Provisioning of the LAN interfaces

The IPv6 Transition CE router may be manually configured in an arbitrary topology with a dynamic routing protocol or using HNCP ([RFC7788]). Automatic provisioning and configuration is described

for a single IPv6 Transition CE router only.

## 5. Requirements

### 5.1. General Requirements

The IPv6 Transition CE router must comply with the general requirements stated in [RFC7084]. Furthermore, a new general requirement is added:

G-6 The IPv6-only CE router MUST comply with [RFC7608].

### 5.2. LAN-Side Configuration

The IPv6 Transition CE router must comply with LAN-Side Configuration as stated in [RFC7084].

In addition, a new LAN Requirement is:

L-15 The IPv6 CE router SHOULD implement a DNS proxy as described in [RFC5625].

### 5.3. Transition Technologies Support

Even if the main target of this document is the support of IPv6-only WAN access, for some time, there will be a need to support IPv4-only devices and applications in the customers LANs, in one side of the picture. In the other side, some Service Providers willing to deploy IPv6, may not be able to do so in the first stage, neither as IPv6-only or dual-stack in the WAN. Consequently, transition technologies to resolve both issues should be taken in consideration.

#### 5.3.1. IPv4 Service Continuity in Customer LANs

##### 5.3.1.1. 464XLAT

464XLAT [RFC6877] is a technique to provide IPv4 access service to IPv6-only edge networks without encapsulation.

The IPv6 Transition CE router SHOULD support CLAT functionality. If 464XLAT is supported, it MUST be implemented according to [RFC6877]. The following CE Requirements also apply:

464XLAT requirements:

464XLAT-1: The IPv6 Transition CE router MUST perform IPv4 Network Address Translation (NAT) on IPv4 traffic translated using the CLAT, unless a dedicated /64 prefix has been

acquired using DHCPv6-PD [RFC3633].

- 464XLAT-2: The IPv6 Transition CE router MUST implement [RFC7050] in order to discover the PLAT-side translation IPv4 and IPv6 prefix(es)/suffix(es). In environments with PCP support, the IPv6 Transition CE SHOULD follow [RFC7225] to learn the PLAT-side translation IPv4 and IPv6 prefix(es)/suffix(es) used by an upstream PCP-controlled NAT64 device.

#### 5.3.1.2. Dual-Stack Lite (DS-Lite)

Dual-Stack Lite [RFC6333] enables both continued support for IPv4 services and incentives for the deployment of IPv6. It also de-couples IPv6 deployment in the service provider network from the rest of the Internet, making incremental deployment easier. Dual-Stack Lite enables a broadband service provider to share IPv4 addresses among customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) and Network Address Translation (NAT). It is expected that DS-Lite traffic is forwarded over the IPv6 Transition CE router's native IPv6 WAN interface, and not encapsulated in another tunnel.

The IPv6 Transition CE router SHOULD implement DS-Lite functionality. If DS-Lite is supported, it MUST be implemented according to [RFC6333]. This document takes no position on simultaneous operation of Dual-Stack Lite and native IPv4. The following IPv6 Transition CE router requirements also apply:

DS-Lite requirements:

- DSLITE-1: The IPv6 Transition CE router MUST support configuration of DS-Lite via the DS-Lite DHCPv6 option [RFC6334]. The IPv6 Transition CE router MAY use other mechanisms to configure DS-Lite parameters. Such mechanisms are outside the scope of this document.
- DSLITE-2: The IPv6 Transition CE router MUST support the DHCPv6 S46 priority option described in [RFC8026].
- DSLITE-3: The IPv6 Transition CE router MUST NOT perform IPv4 Network Address Translation (NAT) on IPv4 traffic encapsulated using DS-Lite.
- DSLITE-4: If the IPv6 Transition CE router is configured with an IPv4 address on its WAN interface, then the IPv6 Transition CE router SHOULD disable the DS-Lite Basic Bridging BroadBand (B4) element.

#### 5.3.1.3. Lightweight 4over6 (lw4o6)

Lw4o6 [RFC7596] specifies an extension to DS-Lite, which moves the NAPT function from the DS-Lite tunnel concentrator to the tunnel client located in the IPv6 Transition CE router, removing the requirement for a CGN function in the tunnel concentrator and reducing the amount of centralized state.

The IPv6 Transition CE router SHOULD implement lw4o6 functionality. If DS-Lite is implemented, lw4o6 MUST be supported as well. If lw4o6 is supported, it MUST be implemented according to [RFC7596]. This document takes no position on simultaneous operation of lw4o6 and native IPv4. The following IPv6 Transition CE router Requirements also apply:

Lw4o6 requirements:

- LW4O6-1: The IPv6 Transition CE router MUST support configuration of lw4o6 via the lw4o6 DHCPv6 options [RFC7598]. The IPv6 Transition CE router MAY use other mechanisms to configure lw4o6 parameters. Such mechanisms are outside the scope of this document.
- LW4O6-2: The IPv6 Transition CE router MUST support the DHCPv6 S46 priority option described in [RFC8026].
- LW4O6-3: The IPv6 Transition CE router MUST support the DHCPv4-over-DHCPv6 (DHCP 4o6) transport described in [RFC7341].
- LW4O6-4: The IPv6 Transition CE router MAY support Dynamic Allocation of Shared IPv4 Addresses as described in [RFC7618].

#### 5.3.1.4. MAP-E

MAP-E [RFC7597] is a mechanism for transporting IPv4 packets across an IPv6 network using IP encapsulation, including a generic mechanism for mapping between IPv6 addresses and IPv4 addresses as well as transport-layer ports.

The IPv6 Transition CE router SHOULD support MAP-E functionality. If MAP-E is supported, it MUST be implemented according to [RFC7597]. The following CE Requirements also apply:

MAP-E requirements:

- MAPE-1: The IPv6 Transition CE router MUST support configuration of MAP-E via the MAP-E DHCPv6 options [RFC7598]. The IPv6

Transition CE router MAY use other mechanisms to configure MAP-E parameters. Such mechanisms are outside the scope of this document.

MAPE-2: The IPv6 Transition CE router MUST support the DHCPv6 S46 priority option described in [RFC8026].

#### 5.3.1.5. MAP-T

MAP-T [RFC7599] is a mechanism similar to MAP-E, differing from it in that MAP-T uses IPv4-IPv6 translation, rather than encapsulation, as the form of IPv6 domain transport.

The IPv6 Transition CE router SHOULD support MAP-T functionality. If MAP-T is supported, it MUST be implemented according to [RFC7599]. The following IPv6 Transition CE Requirements also apply:

MAP-T requirements:

MAPT-1: The CE router MUST support configuration of MAP-T via the MAP-E DHCPv6 options [RFC7598]. The IPv6 Transition CE router MAY use other mechanisms to configure MAP-E parameters. Such mechanisms are outside the scope of this document.

MAPT-2: The IPv6 Transition CE router MUST support the DHCPv6 S46 priority option described in [RFC8026].

#### 5.3.2. Support of IPv6 in IPv4-only WAN access

##### 5.3.2.1. 6in4

6in4 [RFC4213] specifies a tunneling mechanism to allow end-users to manually configure IPv6 support via a service provider's IPv4 network infrastructure.

The IPv6 Transition CE router MAY support 6in4 functionality. 6in4 used for a manually configured tunnel requires a subset of the 6rd parameters (delegated prefix and remote IPv4 end-point). The on-wire and forwarding plane is identical for both mechanisms, however 6in4 doesn't support mesh traffic and requires manually provisioning. Thus, if the device supports either 6rd or 6in4, it's commonly a minor UI addition to support both. If 6in4 is supported, it MUST be implemented according to [RFC4213]. The following CE Requirements also apply:

6in4 requirements:

- 6IN4-1: The IPv6 Transition CE router SHOULD support 6in4 automated configuration by means of the 6rd DHCPv4 Option 212. If the IPv6 Transition CE router has obtained an IPv4 network address through some other means such as PPP, it SHOULD use the DHCPINFORM request message [RFC2131] to request the 6rd DHCPv4 Option. The IPv6 Transition CE router MAY use other mechanisms to configure 6in4 parameters. Such mechanisms are outside the scope of this document.
- 6IN4-2: If the IPv6 Transition CE router is capable of automated configuration of IPv4 through IPCP (i.e., over a PPP connection), it MUST support user-entered configuration of 6in4.
- 6IN4-3: If the IPv6 Transition CE router supports configuration mechanisms other than the 6rd DHCPv4 Option 212 (user-entered, TR-069 [TR-069], etc.), the IPv6 Transition CE router MUST support 6in4 in "hub and spoke" mode. 6in4 in "hub and spoke" requires all IPv6 traffic to go to the 6rd Border Relay, which in this case is the tunnel-end-point. In effect, this requirement removes the "direct connect to 6rd" route defined in Section 7.1.1 of [RFC5969].
- 6IN4-4: The IPv6 Transition CE router MUST allow 6in4 and native IPv6 WAN interfaces to be active alone as well as simultaneously in order to support coexistence of the two technologies during an incremental transition period such as a transition from 6in4 to native IPv6.
- 6IN4-5: Each packet sent on a 6in4 or native WAN interface MUST be directed such that its source IP address is derived from the delegated prefix associated with the particular interface from which the packet is being sent (Section 4.3 of [RFC3704]).
- 6IN4-6: The IPv6 Transition CE router MUST allow different as well as identical delegated prefixes to be configured via each (6in4 or native) WAN interface.
- 6IN4-7: In the event that forwarding rules produce a tie between 6in4 and native IPv6, by default, the IPv6 Transition CE router MUST prefer native IPv6.

#### 5.3.2.2. 6rd

6rd [RFC5969] specifies an automatic tunneling mechanism tailored to advance deployment of IPv6 to end users via a service provider's IPv4 network infrastructure. Key aspects include automatic IPv6 prefix

delegation to sites, stateless operation, simple provisioning, and service that is equivalent to native IPv6 at the sites that are served by the mechanism. It is expected that such traffic is forwarded over the IPv6 Transition CE router's native IPv4 WAN interface and not encapsulated in another tunnel.

The IPv6 Transition CE router MAY support 6rd functionality. If 6rd is supported, it MUST be implemented according to [RFC5969]. The following CE Requirements also apply:

6rd requirements:

- 6RD-1: The IPv6 Transition CE router MUST support 6rd configuration via the 6rd DHCPv4 Option 212. If the IPv6 Transition CE router has obtained an IPv4 network address through some other means such as PPP, it SHOULD use the DHCPINFORM request message [RFC2131] to request the 6rd DHCPv4 Option. The IPv6 Transition CE router MAY use other mechanisms to configure 6rd parameters. Such mechanisms are outside the scope of this document.
- 6RD-2: If the IPv6 Transition CE router is capable of automated configuration of IPv4 through IPCP (i.e., over a PPP connection), it MUST support user-entered configuration of 6rd.
- 6RD-3: If the IPv6 Transition CE router supports configuration mechanisms other than the 6rd DHCPv4 Option 212 (user-entered, TR-069 [TR-069], etc.), the IPv6 Transition CE router MUST support 6rd in "hub and spoke" mode. 6rd in "hub and spoke" requires all IPv6 traffic to go to the 6rd Border Relay. In effect, this requirement removes the "direct connect to 6rd" route defined in Section 7.1.1 of [RFC5969].
- 6RD-4: The IPv6 Transition CE router MUST allow 6rd and native IPv6 WAN interfaces to be active alone as well as simultaneously in order to support coexistence of the two technologies during an incremental transition period such as a transition from 6rd to native IPv6.
- 6RD-5: Each packet sent on a 6rd or native WAN interface MUST be directed such that its source IP address is derived from the delegated prefix associated with the particular interface from which the packet is being sent (Section 4.3 of [RFC3704]).
- 6RD-6: The IPv6 Transition CE router MUST allow different as well as identical delegated prefixes to be configured via each (6rd



or native) WAN interface.

6RD-7: In the event that forwarding rules produce a tie between 6rd and native IPv6, by default, the IPv6 Transition CE router MUST prefer native IPv6.

#### 5.4. IPv4 Multicast Support

Actual deployments support IPv4 multicast for services such as IPTV. In the transition phase it is expected that multicast services will still be provided using IPv4 to the customer LANs.

In order to support the delivery of IPv4 multicast services to IPv4 clients over an IPv6 multicast network, the IPv6 Transition CE router SHOULD support [RFC8114] and [RFC8115].

#### 5.5. Security Considerations

The IPv6 Transition CE router must comply with the Security Considerations as stated in draft-palet-v6ops-rfc7084-bis2.

#### 6. Acknowledgements

Thanks to James Woodyatt, Mohamed Boucadair, Masanobu Kawashima, Mikael Abrahamsson, Barbara Stark, Ole Troan and Brian Carpenter for their review and comments.

#### 7. ANNEX A: Code Considerations

One of the apparent main issues for vendors to include new functionalities, such as support for new transition mechanisms, is the lack of space in the flash (or equivalent) memory. However, it has been confirmed from existing open source implementations (OpenWRT/LEDE), that adding the support for the new transitions mechanisms, requires around 10-12 Kbytes (because most of the code is shared among several transition mechanisms), which typically means about 0,15% of the existing code size in popular CEs in the market.

It is also clear that the new requirements don't have extra cost in terms of RAM memory, neither other hardware requirements such as more powerful CPUs.

The other issue seems to be the cost of developing the code for those new functionalities. However at the time of writing this document, it has been confirmed that there are several open source versions of the required code for supporting the new transition mechanisms, so the development cost is negligent, and only integration and testing cost may become a minor issue.

## 8. References

### 8.1. Normative References

- [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>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <<https://www.rfc-editor.org/info/rfc3633>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<https://www.rfc-editor.org/info/rfc3704>>.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, DOI 10.17487/RFC4213, October 2005, <<https://www.rfc-editor.org/info/rfc4213>>.
- [RFC5625] Bellis, R., "DNS Proxy Implementation Guidelines", BCP 152, RFC 5625, DOI 10.17487/RFC5625, August 2009, <<https://www.rfc-editor.org/info/rfc5625>>.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, DOI 10.17487/RFC5969, August 2010, <<https://www.rfc-editor.org/info/rfc5969>>.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <<https://www.rfc-editor.org/info/rfc6333>>.
- [RFC6334] Hankins, D. and T. Mrugalski, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite", RFC 6334, DOI 10.17487/RFC6334, August 2011, <<https://www.rfc-editor.org/info/rfc6334>>.

- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <<https://www.rfc-editor.org/info/rfc6877>>.
- [RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <<https://www.rfc-editor.org/info/rfc7050>>.
- [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013, <<https://www.rfc-editor.org/info/rfc7084>>.
- [RFC7225] Boucadair, M., "Discovering NAT64 IPv6 Prefixes Using the Port Control Protocol (PCP)", RFC 7225, DOI 10.17487/RFC7225, May 2014, <<https://www.rfc-editor.org/info/rfc7225>>.
- [RFC7341] Sun, Q., Cui, Y., Siodelski, M., Krishnan, S., and I. Farrer, "DHCPv4-over-DHCPv6 (DHCP 4o6) Transport", RFC 7341, DOI 10.17487/RFC7341, August 2014, <<https://www.rfc-editor.org/info/rfc7341>>.
- [RFC7596] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture", RFC 7596, DOI 10.17487/RFC7596, July 2015, <<https://www.rfc-editor.org/info/rfc7596>>.
- [RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <<https://www.rfc-editor.org/info/rfc7597>>.
- [RFC7598] Mrugalski, T., Troan, O., Farrer, I., Perreault, S., Dec, W., Bao, C., Yeh, L., and X. Deng, "DHCPv6 Options for Configuration of Software Address and Port-Mapped Clients", RFC 7598, DOI 10.17487/RFC7598, July 2015, <<https://www.rfc-editor.org/info/rfc7598>>.
- [RFC7599] Li, X., Bao, C., Dec, W., Ed., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", RFC 7599, DOI 10.17487/RFC7599, July 2015, <<https://www.rfc-editor.org/info/rfc7599>>.

- [RFC7608] Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", BCP 198, RFC 7608, DOI 10.17487/RFC7608, July 2015, <<https://www.rfc-editor.org/info/rfc7608>>.
- [RFC7618] Cui, Y., Sun, Q., Farrer, I., Lee, Y., Sun, Q., and M. Boucadair, "Dynamic Allocation of Shared IPv4 Addresses", RFC 7618, DOI 10.17487/RFC7618, August 2015, <<https://www.rfc-editor.org/info/rfc7618>>.
- [RFC8026] Boucadair, M. and I. Farrer, "Unified IPv4-in-IPv6 Software Customer Premises Equipment (CPE): A DHCPv6-Based Prioritization Mechanism", RFC 8026, DOI 10.17487/RFC8026, November 2016, <<https://www.rfc-editor.org/info/rfc8026>>.
- [RFC8114] Boucadair, M., Qin, C., Jacquenet, C., Lee, Y., and Q. Wang, "Delivery of IPv4 Multicast Services to IPv4 Clients over an IPv6 Multicast Network", RFC 8114, DOI 10.17487/RFC8114, March 2017, <<https://www.rfc-editor.org/info/rfc8114>>.
- [RFC8115] Boucadair, M., Qin, J., Tsou, T., and X. Deng, "DHCPv6 Option for IPv4-Embedded Multicast and Unicast IPv6 Prefixes", RFC 8115, DOI 10.17487/RFC8115, March 2017, <<https://www.rfc-editor.org/info/rfc8115>>.

## 8.2. Informative References

- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<https://www.rfc-editor.org/info/rfc7788>>.
- [TR-069] Broadband Forum, "CPE WAN Management Protocol", TR-069 Amendment 4, July 2011, <<http://www.broadband-forum.org/technical/trlist.php>>.
- [UPnP-IGD] UPnP Forum, "InternetGatewayDevice:2 Device Template Version 1.01", December 2010, <<http://upnp.org/specs/gw/igd2/>>.

Author's Address

Jordi Palet Martinez  
Consulintel, S.L.  
Molino de la Navata, 75  
La Navata - Galapagar, Madrid 28420  
Spain

EMail: [jordi.palet@consulintel.es](mailto:jordi.palet@consulintel.es)  
URI: <http://www.consulintel.es/>

IPv6 Operations (v6ops)  
Internet-Draft  
Obsoletes: 7084 (if approved)  
Intended status: Informational  
Expires: December 12, 2017

J. Palet Martinez  
Consulintel, S.L.  
June 10, 2017

Minimum Requirements for IPv6-only Customer Edge Routers  
draft-palet-v6ops-rfc7084-bis2-00

Abstract

This document specifies minimum requirements for an IPv6-only Customer Edge (CE) router. Specifically, the current version of this document focuses on the basic provisioning of an IPv6-only CE router and the provisioning of IPv6 hosts attached to it. Neither the provisioning of IPv4 nor downstream routers are in the scope of this document. 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 <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 December 12, 2017.

Copyright Notice

Copyright (c) 2017 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 . . . . .	2
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 . . . . .	4
3.2.1. Local Communication . . . . .	6
4. Requirements . . . . .	6
4.1. General Requirements . . . . .	6
4.2. WAN-Side Configuration . . . . .	7
4.3. LAN-Side Configuration . . . . .	11
4.4. Security Considerations . . . . .	13
5. Acknowledgements . . . . .	14
6. Contributors . . . . .	14
7. ANNEX A: Changes from RFC7084 . . . . .	14
8. References . . . . .	15
8.1. Normative References . . . . .	15
8.2. Informative References . . . . .	18
Author's Address . . . . .	18

## 1. Introduction

This document defines minimum IPv6 features for a very basic residential or small- office router, referred to as an "IPv6-only CE router", in order to establish an industry baseline for features to be implemented on such a router. This is, as well, the router to be used in remote locations where IPv6-only LANs/devices are used to connect sensors, displays, etc.

These routers don't support IPv4, neither functionalities for provisioning downstream routers.

This document specifies how an IPv6-only 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.

This document doesn't cover the specific details of each possible access technology. For example, if the CE is supporting built-in or external 3GPP/LTE interfaces, [RFC7849] is a relevant reference. See [RFC4779] for a discussion of options available for deploying IPv6 in

wireline service provider access networks.

### 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-only CE router that connect IPv6 hosts.
IPv6-only Customer Edge Router	a node intended for home or small-office use that forwards IPv6 packets not explicitly addressed to itself. The IPv6-only CE router connects the end-user network to a service provider network. In other documents, the CE is named as CPE (Customer Premises Equipment or Customer Provided Equipment). In the context of this document, both terminologies are synonymous.
IPv6 Host	any device implementing an IPv6 stack receiving IPv6 connectivity through the IPv6-only CE router.
LAN Interface	an IPv6-only 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-only 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-only, and it may also



offer IPv4 Internet access, but non intended to be supported by this IPv6-only CE router. The service provider can provide such access over a variety of different transport methods such as FTTH, DSL, cable, wireless, 3GPP/LTE, and others.

#### WAN Interface

an IPv6-only 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.

### 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, however very simple networks may work perfectly with IPv6-only. 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 a simple IPv6-only network 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:

- o Basic capabilities of an IPv6-only 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-only CE router may act as a Multicast Listener Discovery (MLD) proxy [RFC4605] and may support a dynamic multicast routing protocol.

The IPv6-only 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-only 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-only 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.

At the time of this writing, several host 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-only 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-only 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-only 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.

## 4. Requirements

### 4.1. General Requirements

The IPv6-only CE router is responsible for implementing IPv6 routing; that is, the IPv6-only 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-only 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-only CE router is an IPv6 node according to the IPv6 Node Requirements specification [RFC6434].

G-2: The IPv6-only 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-only 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-only 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-only 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-only CE router MUST comply with [RFC7608].

#### 4.2. WAN-Side Configuration

The IPv6-only CE router will need to support connectivity to one or more access network architectures. This document describes an IPv6-only 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-only CE router is attached to the WAN interface link, it MUST act as an IPv6 host for the purposes of stateless [RFC4862] or stateful [RFC3315] interface address assignment.
- W-2: The IPv6-only 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-only 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-only CE router MUST act as a requesting router for the purposes of DHCPv6 prefix delegation ([RFC3633]).
- W-5: The IPv6-only 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 IPv6-only 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-only CE router MUST support IPv6 over Ethernet [RFC2464].
- WLL-2: If the WAN interface supports PPP encapsulation, the IPv6-only CE router MUST support IPv6 over PPP [RFC5072].

Address assignment requirements:

- WAA-1: The IPv6-only CE router MUST support Stateless Address Autoconfiguration (SLAAC) [RFC4862].
- WAA-2: The IPv6-only 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-only CE router MUST support DHCPv6 [RFC3315] client behavior.
- WAA-4: The IPv6-only CE router MUST be able to support the following DHCPv6 options: Identity Association for Non-temporary Address (IA\_NA), Reconfigure Accept [RFC3315], and

DNS\_SERVERS [RFC3646]. The IPv6-only CE router SHOULD be able to support the DNS Search List (DNSSL) option as specified in [RFC3646].

- WAA-5: The IPv6-only 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-only 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. LAN side support of NTP is out of scope for this document.
- WAA-6: If the IPv6-only 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-only 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-only CE router MUST support the SOL\_MAX\_RT option [RFC7083] and request the SOL\_MAX\_RT option in an Option Request Option (ORO).
- WAA-9: As a router, the IPv6-only 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-only CE router SHOULD implement the Information Refresh Time option and associated client behavior as specified in [RFC4242].

Prefix delegation requirements:

- WPD-1: The IPv6-only CE router MUST support DHCPv6 prefix delegation requesting router behavior as specified in [RFC3633] (Identity Association for Prefix Delegation (IA\_PD) option).
- WPD-2: The IPv6-only 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-only 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-only 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-only 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 IPv6-only 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-only 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-only 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-only 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 [RFC7550].

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

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

#### 4.3. LAN-Side Configuration

The IPv6-only 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-only 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-only 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-only CE router's LAN interfaces only.

ULA requirements:

- ULA-1: The IPv6-only CE router SHOULD be capable of generating a ULA prefix [RFC4193].
- ULA-2: An IPv6-only 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-only 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-only 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-only CE router MUST support router behavior according



to Neighbor Discovery for IPv6 [RFC4861].

- L-2: The IPv6-only 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-only 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-only 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-only 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-only 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-only 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-only CE router MUST support providing DNS information in the DHCPv6 DNS\_SERVERS and DOMAIN\_LIST options [RFC3646].
- L-11: The IPv6-only 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 [RFC6106].
- L-12: The IPv6-only CE router SHOULD implement a DNS proxy as described in [RFC5625].
- L-13: The IPv6-only 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-14: If the delegated prefix changes, i.e., the current prefix is replaced with a new prefix without any overlapping time period, then the IPv6-only CE router MUST immediately advertise the old prefix with a Preferred Lifetime of zero and a Valid Lifetime of either a) zero or b) the lower of the current Valid Lifetime and two hours (which must be decremented in real time) in a Router Advertisement message as described in Section 5.5.3, (e) of [RFC4862].
- L-15: The IPv6-only 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.

#### 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-only CE router ought to support basic stateless egress and ingress filters. The IPv6-only 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-only CE router SHOULD support [RFC6092]. In particular, the IPv6-only 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-only CE router SHOULD support ingress filtering in accordance with BCP 38 [RFC2827]. Note that this requirement was downgraded from a MUST from RFC 6204 due to the difficulty of implementation in the CE router and the feature's redundancy with upstream router ingress filtering.
- S-3: If the IPv6-only 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

This document is an update of RFC7084, whose original authors were: Hemant Singh, Wes Beebee, Chris Donley and Barbara Stark. The rest of the text on this section and the Contributors section, are the original acknowledgements and Contributors sections of the earlier version of this document.

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

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.

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: 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.

## 7. ANNEX A: Changes from RFC7084

The -bis2 version of this document has some minor text edits here and there. Significant updates are:

1. Removed, in general requirements related to IPv4/dual-stack support.

2. References to IPv6 CE router changed to IPv6-only CE router.
3. Removed WLL-3, as it was referring to dual-stack support.
4. G-6 added in order to comply with [RFC7608].
5. L-12 added to support for DNS proxy [RFC5625] as general LAN requirement.
6. Removed transition support.

## 8. References

### 8.1. Normative References

- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<http://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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<http://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, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.
- [RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, DOI 10.17487/RFC3646, December 2003, <<http://www.rfc-editor.org/info/rfc3646>>.

- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, DOI 10.17487/RFC3736, April 2004, <<http://www.rfc-editor.org/info/rfc3736>>.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <<http://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, <<http://www.rfc-editor.org/info/rfc4193>>.
- [RFC4242] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, DOI 10.17487/RFC4242, November 2005, <<http://www.rfc-editor.org/info/rfc4242>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, DOI 10.17487/RFC4443, March 2006, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc5072>>.
- [RFC5625] Bellis, R., "DNS Proxy Implementation Guidelines", BCP 152, RFC 5625, DOI 10.17487/RFC5625, August 2009, <<http://www.rfc-editor.org/info/rfc5625>>.
- [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, <<http://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, <<http://www.rfc-editor.org/info/rfc5908>>.
- [RFC5942] Singh, H., Beebe, W., and E. Nordmark, "IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes", RFC 5942, DOI 10.17487/RFC5942, July 2010, <<http://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, <<http://www.rfc-editor.org/info/rfc6092>>.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, DOI 10.17487/RFC6106, November 2010, <<http://www.rfc-editor.org/info/rfc6106>>.
- [RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, DOI 10.17487/RFC6177, March 2011, <<http://www.rfc-editor.org/info/rfc6177>>.
- [RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, DOI 10.17487/RFC6434, December 2011, <<http://www.rfc-editor.org/info/rfc6434>>.
- [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, <<http://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, <<http://www.rfc-editor.org/info/rfc6887>>.
- [RFC7083] Droms, R., "Modification to Default Values of SOL\_MAX\_RT and INF\_MAX\_RT", RFC 7083, DOI 10.17487/RFC7083, November 2013, <<http://www.rfc-editor.org/info/rfc7083>>.
- [RFC7608] Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", BCP 198, RFC 7608, DOI 10.17487/RFC7608, July 2015, <<http://www.rfc-editor.org/info/rfc7608>>.

## 8.2. Informative References

- [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, <<http://www.rfc-editor.org/info/rfc7157>>.
- [RFC7550] Troan, O., Volz, B., and M. Siodelski, "Issues and Recommendations with Multiple Stateful DHCPv6 Options", RFC 7550, DOI 10.17487/RFC7550, May 2015, <<http://www.rfc-editor.org/info/rfc7550>>.
- [RFC7849] Binet, D., Boucadair, M., Vizdal, A., Chen, G., Heatley, N., Chandler, R., Michaud, D., Lopez, D., and W. Haeffner, "An IPv6 Profile for 3GPP Mobile Devices", RFC 7849, DOI 10.17487/RFC7849, May 2016, <<http://www.rfc-editor.org/info/rfc7849>>.
- [TR-069] Broadband Forum, "CPE WAN Management Protocol", TR-069 Amendment 4, July 2011, <<http://www.broadband-forum.org/technical/trlist.php>>.
- [UPnP-IGD] UPnP Forum, "InternetGatewayDevice:2 Device Template Version 1.01", December 2010, <<http://upnp.org/specs/gw/igd2/>>.

Author's Address

Jordi Palet Martinez  
Consulintel, S.L.  
Molino de la Navata, 75  
La Navata - Galapagar, Madrid 28420  
Spain

EMail: [jordi.palet@consulintel.es](mailto:jordi.palet@consulintel.es)  
URI: <http://www.consulintel.es/>



IPv6 Operations (v6ops)  
Internet-Draft  
Obsoletes: 7084 (if approved)  
Intended status: Informational  
Expires: December 12, 2017

J. Palet Martinez  
Consulintel, S.L.  
June 10, 2017

Basic Requirements for IPv6 Customer Edge Routers with HNCP  
draft-palet-v6ops-rfc7084-bis4-hnccp-00

Abstract

This document specifies minimum 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. Includes support of HNCP ([RFC7788]) for automated provisioning of downstream routers. 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 <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 December 12, 2017.

Copyright Notice

Copyright (c) 2017 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 . . . . .	2
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 . . . . .	4
3.2.1. Local Communication . . . . .	6
4. Requirements . . . . .	6
4.1. General Requirements . . . . .	6
4.2. WAN-Side Configuration . . . . .	7
4.3. LAN-Side Configuration . . . . .	11
4.4. Security Considerations . . . . .	13
5. Acknowledgements . . . . .	14
6. Contributors . . . . .	14
7. ANNEX A: Changes from RFC7084 . . . . .	14
8. References . . . . .	15
8.1. Normative References . . . . .	15
8.2. Informative References . . . . .	18
Author's Address . . . . .	18

## 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.

These routers typically also support IPv4, at least in the LAN side.

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 may be handled by means of HNCP ([RFC7788]).

This document doesn't cover the specific details of each possible access technology. For example, if the CE is supporting built-in or external 3GPP/LTE interfaces, [RFC7849] is a relevant reference. See [RFC4779] for a discussion of options available for deploying IPv6 in wireline service provider access networks.

### 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 IPv6 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. In other documents, the CE is named as CPE (Customer Premises Equipment or Customer Provided Equipment). In the context of this document, both terminologies are synonymous.
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-only, and it may also offer IPv4 Internet access, but non intended to be supported by this IPv6 CE

router. The service provider can provide such access over a variety of different transport methods such as FTTH, DSL, cable, wireless, 3GPP/LTE, 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.

### 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 a simple IPv6-only network should provide equivalent or better capabilities and functionality than the current IPv4 architecture.

The end-user network is a stub network, in the sense that is not providing transit to other external networks. However HNCP

([RFC7788]) allows supporting automatic provisioning of downstream routers. Figure 1 illustrates the model topology for the end-user network.

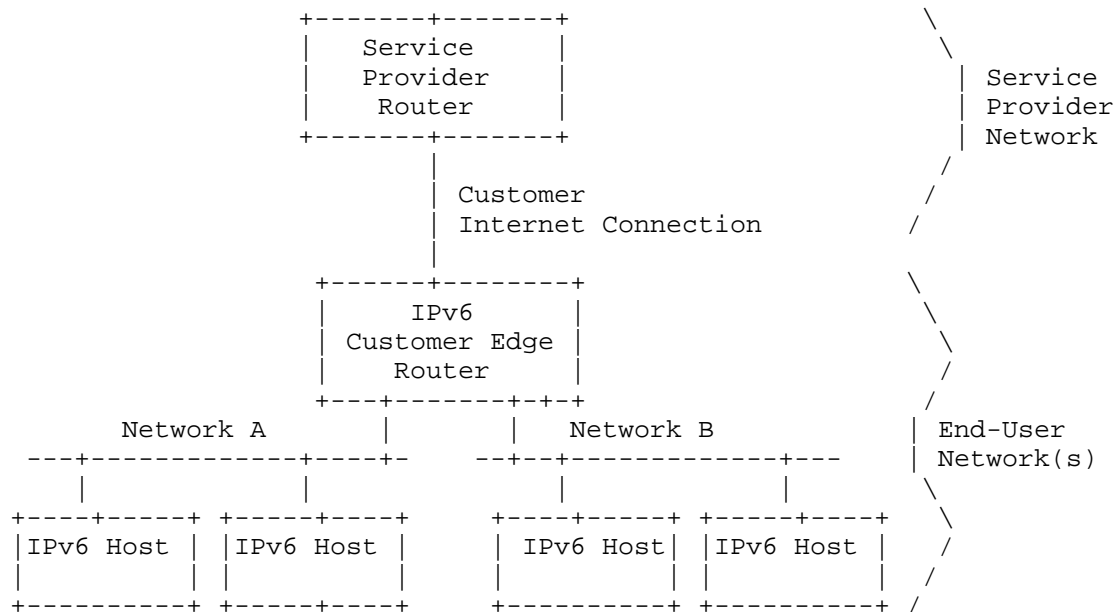


Figure 1: An Example of a Typical End-User Network

This architecture describes the:

- 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 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 or using HNCP ([RFC7788]). 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 (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.

At the time of this writing, several host 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 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 [RFC6434].

- 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 comply with [RFC7608].

#### 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 [RFC3315] 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 ([RFC3633]).
- 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 IPv6 CE router SHOULD support a Port Control Protocol (PCP) client as specified in [RFC6887] for use by applications on the IPv6 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 IPv6 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 [RFC3315] 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 [RFC3315], 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. LAN side support of NTP is out of scope for this document.
- 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 [RFC7083] 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 [RFC4242].

Prefix delegation requirements:

- WPD-1: The IPv6 CE router MUST support DHCPv6 prefix delegation requesting router behavior as specified in [RFC3633] (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 IPv6 CE router to use DHCPv6 prefix delegation when the IPv6 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 IPv6 CE router but not in the set of prefixes assigned by the IPv6 CE router to the LAN must be dropped. In other words, the next hop for the prefix(es) delegated to the IPv6 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 [RFC7550].
- 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.

#### 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 [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 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 [RFC6106].
- L-12: The IPv6 CE router SHOULD implement a DNS proxy as described in [RFC5625].
- L-13: 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-14: 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 zero and a Valid Lifetime of either a) zero or b) the lower of the current Valid Lifetime and two hours (which must be decremented in real time) in a Router Advertisement message as described in Section 5.5.3, (e) of [RFC4862].
- L-15: 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-16: The IPv6 CE router SHOULD provide HNCP (Home Networking Control Protocol) services, as specified in [RFC7788].

#### 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 IPv6 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 SHOULD support ingress filtering in accordance with BCP 38 [RFC2827]. Note that this requirement was downgraded from a MUST from RFC 6204 due to the difficulty of implementation in the IPv6 CE router and the feature's redundancy with upstream router ingress filtering.
- 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

This document is an update of RFC7084, whose original authors were: Hemant Singh, Wes Beebee, Chris Donley and Barbara Stark. The rest of the text on this section and the Contributors section, are the original acknowledgements and Contributors sections of the earlier version of this document.

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

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.

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: 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.

## 7. ANNEX A: Changes from RFC7084

The -bis-4-hncp version of this document has some minor text edits here and there. Significant updates are:

1. G-6 added in order to comply with [RFC7608].

2. L-12 added to support for DNS proxy [RFC5625] as general LAN requirement.
3. Added support of HNCP ([RFC7788]) in LAN (L-16).
4. Removed transition support.

## 8. References

### 8.1. Normative References

- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<http://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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<http://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, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.
- [RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, DOI 10.17487/RFC3646, December 2003, <<http://www.rfc-editor.org/info/rfc3646>>.
- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, DOI 10.17487/RFC3736, April 2004, <<http://www.rfc-editor.org/info/rfc3736>>.

- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <<http://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, <<http://www.rfc-editor.org/info/rfc4193>>.
- [RFC4242] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, DOI 10.17487/RFC4242, November 2005, <<http://www.rfc-editor.org/info/rfc4242>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, DOI 10.17487/RFC4443, March 2006, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc5072>>.



- [RFC5625] Bellis, R., "DNS Proxy Implementation Guidelines", BCP 152, RFC 5625, DOI 10.17487/RFC5625, August 2009, <<http://www.rfc-editor.org/info/rfc5625>>.
- [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, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc6092>>.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, DOI 10.17487/RFC6106, November 2010, <<http://www.rfc-editor.org/info/rfc6106>>.
- [RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, DOI 10.17487/RFC6177, March 2011, <<http://www.rfc-editor.org/info/rfc6177>>.
- [RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, DOI 10.17487/RFC6434, December 2011, <<http://www.rfc-editor.org/info/rfc6434>>.
- [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, <<http://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, <<http://www.rfc-editor.org/info/rfc6887>>.

- [RFC7083] Droms, R., "Modification to Default Values of SOL\_MAX\_RT and INF\_MAX\_RT", RFC 7083, DOI 10.17487/RFC7083, November 2013, <<http://www.rfc-editor.org/info/rfc7083>>.
- [RFC7608] Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", BCP 198, RFC 7608, DOI 10.17487/RFC7608, July 2015, <<http://www.rfc-editor.org/info/rfc7608>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<http://www.rfc-editor.org/info/rfc7788>>.

## 8.2. Informative References

- [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, <<http://www.rfc-editor.org/info/rfc7157>>.
- [RFC7550] Troan, O., Volz, B., and M. Siodelski, "Issues and Recommendations with Multiple Stateful DHCPv6 Options", RFC 7550, DOI 10.17487/RFC7550, May 2015, <<http://www.rfc-editor.org/info/rfc7550>>.
- [RFC7849] Binet, D., Boucadair, M., Vizdal, A., Chen, G., Heatley, N., Chandler, R., Michaud, D., Lopez, D., and W. Haeffner, "An IPv6 Profile for 3GPP Mobile Devices", RFC 7849, DOI 10.17487/RFC7849, May 2016, <<http://www.rfc-editor.org/info/rfc7849>>.
- [TR-069] Broadband Forum, "CPE WAN Management Protocol", TR-069 Amendment 4, July 2011, <<http://www.broadband-forum.org/technical/trlist.php>>.
- [UPnP-IGD] UPnP Forum, "InternetGatewayDevice:2 Device Template Version 1.01", December 2010, <<http://upnp.org/specs/gw/igd2/>>.

Author's Address

Jordi Palet Martinez  
Consulintel, S.L.  
Molino de la Navata, 75  
La Navata - Galapagar, Madrid 28420  
Spain

EMail: [jordi.palet@consulintel.es](mailto:jordi.palet@consulintel.es)  
URI: <http://www.consulintel.es/>