IPv6 Maintenance (6man) Working Group
Internet Draft
Updates: 4861, 4862, 6724 (if approved)
Intended status: Standards Track
Expires: September 2023

E. Vasilenko
P. Volpato
Huawei Technologies

March 26, 2023

Neighbor Discovery support for Multi-home Multi-prefix
draft-vv-6man-nd-support-mhmp-02

Abstract

Multi-home Multi-prefix (MHMP) IPv6 environment is the norm for businesses that need to have uplink resiliency.
Several solutions have been already discussed and proposed to address MHMP and how it can be enabled in different network contexts. This draft looks at MHMP from the perspective of Neighbour Discovery Protocols (NDP).

For any considered destination, the MHMP challenge may be split into 3 sub-challenges (important to solve in the below order):
1) the host should choose the proper source address for the packet,
2) the host should choose the best default router as the next hop,
3) site topology may be complicated and may need the source routing through the site.

This draft is concerned with the solution for the first two problems that need improvement for the ND (RFC 4861) SLAAC (RFC 4862) and Default Address Selection (RFC 6724). The last problem is considered as properly discussed by Multihoming in Enterprise (RFC 8678).

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 September 2023.

Table of Contents

1. Terminology and pre-requisite.....3
2. Introduction.....3
3. The NDP analysis in MHMP.....5
4. Solution for the case "equal prefixes".....7
5. Solution for the case "non-equal prefixes".....8
6. Resolution for a not-reachable provider.....10
7. Extensions of the existing standards.....12
7.1. Preference to choose source address before the next hop..12
7.2. Default router choice by host.....13
7.3. Prefixes become dynamic.....13
7.4. Default router announcement rules.....15
7.5. Clean prefixes for the changed default router list.....15
8. Interoperability analysis.....15
9. Security Considerations.....16
10. IANA Considerations.....16
11. References.....16
11.1. Normative References.....16
11.2. Informative References.....17
12. Acknowledgments.....18

1. Terminology and pre-requisite

[Default Address], [ND], and [SLAAC] are pre-requisite to understanding this document.
The terms are inherited from these standards.

Additional terms:

PA - Provider-Aggregatable addresses leased by the Carrier to the client or subscriber

PI - Provider-Independent addresses received from some Regional Internet Registry

MHMP - Multi-Homing Multi-Prefix. An environment with hosts connected to different PA providers (multi-homing) through different address spaces announced from different providers (multi-prefix)

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Introduction

Businesses usually have multiple network connections with different service providers to guarantee network resiliency.

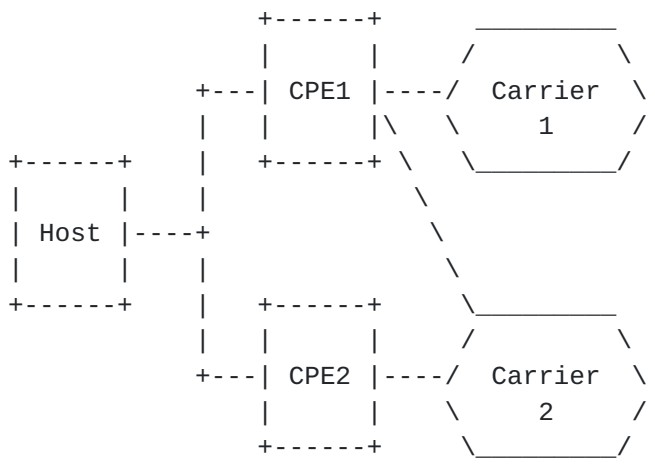

Figure 1: The basic Multi-Homing Multi-Prefix topology

Such a scenario is identified as Multi-Home Multi-Prefix (MHMP) and properly discussed in [MHMP] and [MHMP Enterprise] which discuss many possible topologies.

IPv4 is solving such a scenario by independent NAT translation on every CPE to the Carrier in combination with private address space on site. [IPv6] is capable to avoid NAT.

An MHMP site may have a complex architecture in [IPv6], potentially, with many links and routers (CPEs) connected to different carriers. CPEs may receive different Provider Aggregatable (PA) prefixes from the upstream carriers.

Hosts located in an MHMP environment may also have multiple different addresses assigned to their interfaces that come from multiple delegations (from different carriers).

This may create challenges when a host located in an MHMP environment wishes to communicate with a certain destination. Such a host typically receives the list of destination addresses from DNS sorted by [Default Address]. Knowing the Destination Address (DA), a host proceeds with the selection of the other parameters necessary for sending the packet following the algorithm assumed in [Default Address].

This process determines the Next-Hop (NH) router used for forwarding the outgoing packet and its Source Address (SA).

In MHMP scenarios though a host may experience difficulties to determine the "right" SA given the NH router selected to communicate. If the network prefix chosen for the SA does not belong to the address space advertised by the selected NH router, the result is the packet to be filtered out at the ingress of the Carrier.

As explained further in the next sections, two aspects may lead a host to select the "wrong" SA.

A first issue may depend on the lack of support for specific functions, such as Rule 5.5 of the SA selection process [Default Address] (select a prefix advertised by the NH router) or adoption of [Conditional PIO] (use a router advertisement to influence the host in the selection of SA).

A second aspect is related to the complexity of the communication service a host may wish to use. This may include destinations

located in controlled environments (e.g. a "walled garden"), or reachable via traffic policy (e.g. imposing a low-latency path). The network prefixes advertised to the host in these cases have to be handled with special care.

This draft looks at the possibilities provided by Neighbor Discovery Protocols (NDP) to solve the MHMP issues. The reason to analyze these possibilities is that NDP is key to IPv6 and may offer a general solution.

The next section introduces the rationale for analyzing NDP to solve the MHMP issues. Sections 4 and 5 dig further into the relevant cases for which NDP may be useful. Section 6 deals with how to handle a non-reachable provider. Section 7 proposes some extensions to the existing NDP. Section 8 deals with aspects related to interoperability.

3. The NDP analysis in MHMP

The desired destination address is the pre-request for the discussion of this document, as pointed out in [Default Address]. This document does not change the way how destination addresses are sorted by [Default Address] or [Happy Eyeballs] rather it analyzes the options once the destination address is selected:

1. choose next-hop first (current practice for clients that prevail in MHMP environment; servers probably use bind() to choose source address but the server side is not relevant to the discussion because it typically has Provider Independent addresses),

or

2. choose the source address first
(more optimal strategy recommended here).

Both need some corrections to the [ND] and [Default Address] selection.

In some cases, these choices may create issues, in particular when a faulty situation occurs (e.g. network prefixes invalidity due to loss of connectivity or abrupt CPE reconfiguration). During such events, a host may find it difficult to establish communication with a destination if the proper next-hop or source address is not selected, for example, due to the packet filtering policies applied by the upstream carriers (for anti-spoofing).

So, overall the association between a destination address, the next-hop, and the source address in an MHMP environment may be challenging.

The MHMP challenge may be split into 3 sub-challenges (important to solve in the below order):

1) the host should choose the proper source address for the packet by [Default Address],

2) the host should choose the best default router as the next-hop [ND] (not strictly mandatory, may be fixed later by the source routing with some loss of efficiency),

3) despite the assumed good choice for the default router selection, site topology may be complicated and as a result, may need the source routing through the site anyway - see [MHMP Enterprise].

It is important to point out that the challenge of selecting the next hop and source address exists for every desired destination.

There are two reminders before the discussion for solutions:

- [ND] section 6.3.6 recommends the next-hop choice between default routers in a round-robin style. Traffic policy or even reachability of particular resources through a particular default router is not considered at the [ND] level.

- [Default Address] section 7 (and a few other places) assumes that source address selection should happen after the next-hop (or interface) choice by [ND].

Before digging into the discussions, it is worth noticing that:

- This document has put NAT (including NPT) out of consideration. The attempt is here to get fully transparent E2E connectivity.

- This document assumes PA environment only, PI (Provider Independent) address space needs a BGP connection to Carrier that does not create a problem to solve from the technology point of view but it creates an enormous problem of scalability for the Internet with tens of millions of routes.

It is possible to introduce one additional classification to separate what it is possible to implement now from what needs additional standardization efforts:

1. Case "equal prefixes": Announced prefixes are fully equal by scope and value, all resources interested for hosts could be reachable through any announced PA prefix. Additionally, traffic distribution between carriers could be non-predictable (no traffic engineering or policy).

2. Case "non-equal prefixes": Announced prefixes are not equal because (1) some resources could be accessed only through a particular prefix (for example "walled garden" of one carrier) or (2) it is desirable to have some policy for traffic distribution between PA prefixes (cost of traffic, delay, packet loss, jitter, proportional load, etc.).

4. Solution for the case "equal prefixes"

This use case is potentially possible to operate without any changes to standardization but it would be not optimal (because currently next-hop is chosen first) and would need additional functionality on routers and hosts anyway (rule 5.5 of [Default Address], [Conditional PIO], source routing). Let's discuss the current standardization situation.

Case "equal prefixes" does not create any requirement on what prefix should be used for the source address. It is only needed that the source address would be chosen to be compatible with the next hop which should be in the direction of the respective carrier.
There are 4 potential scenarios possible in respect of the next-hop choice:

1. A single router on the link does not create a choice for the host in principle. If the site is complex (multi-hop) then the router itself may need source routing to choose the next hop properly, it is considered resolved and properly discussed in [MHMP Enterprise].

2. A Host on a multi-homing link would be better compliant with [Default Address] section 5 (source address selection) rule 5 (for different interfaces on the host) and rule 5.5 (for different next-hops on the same interface). It would help to properly choose a source address compliant with the next hop chosen first.

3. [MHMP Enterprise] proposes a substitution for rule 5.5 absence on hosts by [Conditional PIO] that should not leave a choice to host for what source address to choose.

4. If the source address would be chosen wrongly (because of no support for rule 5.5 of [Default Address] and no support for [Conditional PIO]) then it is still possible to reroute the packet later by source routing proposed in [MHMP Enterprise]. Albeit, the performance would be affected by pushing traffic through redundant routing hop.

The reversal of choice to source address first would permit the improvement of the functionality (see next section) and simplify the "equal prefixes" case because it is much easier to choose next-hop after source address by simply excluding default routers that do not advertise particular PIOs.

5. Solution for the case "non-equal prefixes"

This case is more complicated. It is not fully resolved yet in the standardization. It is the primary motivation for the development of this document.

It would be too late to try to solve this problem on a router, because the wrong source address may have already been chosen by the host - it would not be possible to contact the appropriate resource in the "walled garden" or filter for any other purpose. Additionally, the wrong choice of source address would not permit traffic engineering and host reaction to network quality of services.

[Provisioning domains] is the capability to virtualize router on the link, i.e. present many virtual routers (one per domain) with different parameters. It is possible to do it explicitly (different options specifically developed for virtualization) or implicitly (emulate many LLAs on the router). It does not help to resolve the MHMP problem because different virtual routers are equal to physical routers. It is still a problem if the host would choose initially the virtual router looking to the next hop because then the host would be restricted for prefixes advertised only from this router which may prevent the host to reach the destination or greatly affect the quality of communication.

Currently, there are two standardized methods to resolve the case of "non-equal prefixes":

1. The same policies could be formatted differently and fed to the host by two mechanisms at the same time: 1) "Routing Information Options" of [Route Preferences] and 2) [Policy by DHCP] to modify policies in [Default Address] selection algorithm. Then the

current priority of mechanisms could be preserved the same: initially [ND] or routing would choose the next hop, then [Default Address] would choose a proper source address. It is the method that is assumed in [MHMP]. This method is complicated and costly, and the probability of acceptance is very low. Moreover, [Policy by DHCP] was not adopted by the market - it is not available on the major operating systems and home gateways.

2. Application developers may use bind() to choose a source address based on many different parameters (including anything from the list below). It would effectively revert the default logic of [Default Address]. Unfortunately, it is difficult to expect it from the client side, which would probably call getaddrinfo() which has a good probability to choose the wrong source address.

Alternatively, [Default Address] section 7 discusses the potential capability to reverse the decision's order: source address may be chosen first, only then to choose next-hop (default router). Then many additional methods are possible for how to choose a source address first:

3. Policies could be supplied by [Policy by DHCP] only to the [Default Address] selection algorithm. This method has a low probability of implementation because of not wide support of DHCPv6 in the industry. This method may be more accepted in the future.
4. It is possible to check the longest match between the source and the destination address to choose the potentially closest address. This method looks most promising, it is partially discussed in [Default Address] section 7.
5. The host could use DNS requests with different source addresses to understand what is visible for a particular source address.
6. URL for configuration information could be supplied in RA - see [Provisioning domains].
7. The host may have local performance management capabilities (packet loss, delay, jitter, etc) to choose the best source for the application.

It is possible to have other methods for how the host could decide locally on the best source address as its first decision. This document is readily extensible in this direction.

The source address selected from the proper carrier is the complete information needed for the host to choose the next hop, but it needs improvement of [ND] and [Default Address].
[ND] default round-robin distribution between available routers should be extended for the host to prioritize default routers that

have announced prefixes used for the source address of the considered packet.
Section 7 of [Default Address] should be extended and recommended for hosts to support MHMP.

This document's proposals are inspired by [Multi-Homing] section 3.2. The difference is that the same rules are formulated not as general advice, but as a particular extension to [ND] and [Default Address] - see section 7. of this document.

6. Resolution for a not-reachable provider

Let's assume the fault situation when one provider is not reachable in the [MHMP] environment. A prefix may be very dynamic for a few reasons. It could be received from some protocols (DHCP-PD, HNCP). The prefix could become invalid (at least for the global Internet connectivity) as a result of the abrupt link loss in the upstream direction to the carrier that distributed this prefix.

Additionally, consider the more complicated case when some hosts on the upstream routed path to this provider may still be reachable using a particular prefix but Internet connectivity is broken later.

Let's consider the problem. Because Internet connectivity is lost for this prefix, it should be announced to hosts with zero Preferred Lifetime. [Route Preferences] gives the possibility to inform hosts that a particular prefix (RIO) is still available on-site but it would be an automation challenge to dynamically calculate and announce the prefix. Additionally, [Route Preferences] should be supported by hosts.
In general, it is not a good idea to involve [ND] in routing. Hence, it is better to support on-site connectivity by ULA which may not be invalidated. There are many reasons to promote [ULA] for internal site connectivity: (1) hosts may not have GUA address at all without initial connection to the provider, (2) PA addresses would be invalidated within 30 days of disconnect anyway, (3) it is not a good idea to use addresses from PA pool that is disconnected from global Internet - hosts may have a better option to get global reachability. ULA has better security (open transport ports that are not accessible from the Internet) which is an additional bonus.
It is effectively the request to join current [CPE Requirements] and [HomeNet Architecture] requirements in sections 2.2, 2.4, 3.4.2 that the subscriber's network should have local ULA addresses.

Prefix deprecation should be done by RA with zero Lifetime for this prefix. It will put the prefix on hosts to the deprecated status

that according to many standards ([ND], [SLAAC], and [Default Address]) would prioritize other addresses. Global communication would be disrupted for this prefix anyway. Local communication for deprecated addresses would continue till normal resolution because the default Valid Lifetime is 30 days. Moreover, if it would happen that this delegated prefix was the only one in the local network (no [ULA] for the same reason), then new sessions would be opened on the deprecated prefix because it is the only address available.
If connectivity would be re-established and the same prefix would be delegated to the link - it would be announced again with the proper preferred lifetime. If a different prefix could be delegated by the PA provider, then the old prefix would stay in deprecated status.
It is an advantage for the host that would know about the global reachability of this prefix (by deprecated status) because the host may use other means for communication at that time.

Such dynamic treatment of prefixes may have the danger of [ND] messages flooding if the link on the path to the PA provider would be oscillating.
[HNCP] section 1.1 states: "it is desirable for ISPs to provide large enough valid and preferred lifetimes to avoid unnecessary HNCP state churn in homes".
It makes sense to introduce dampening for the rate of prefix announcements.

Such conceptual change in the treatment of prefixes would not affect current enterprise installations where prefixes are static.

It is important to mention again that it is the responsibility of the respective protocol (that has delivered the prefix to the considered router) to inform the router that the prefix is not routed anymore to the respective carrier. It is easy to do it in the simplified topology when the only router could correlate uplink status with the DHCP-PD prefix delegated early. Some additional protocols like [HNCP] are needed for a more complex topology.

There is nothing in [ND] or [SLAAC] that prevents us from treating prefixes as something more dynamic than "renumbering" to reflect the dynamic path status to the PA provider. Section 7.3. proposes extensions to [CPE Requirements] and [SLAAC] that follow the logic of this section.

7. Extensions of the existing standards

The solution is about several standard extensions that are needed to fulfill the solutions discussed above. They are split into separate sections for better understanding.

7.1. Preference to choose source address before the next hop.

* Section 7 (Interactions with routing) of [Default Address] has at the beginning:

"This specification of source address selection assumes that routing (more precisely, selecting an outgoing interface on a node with multiple interfaces) is done before source address selection. However, implementations MAY use source address considerations as a tiebreaker when choosing among otherwise equivalent routes."

Replace the above text with the text:

"This specification of source address selection did assume that routing (more precisely, selecting an outgoing interface on a node with multiple interfaces) is done after source address selection. MHMP support strongly demands choosing the source address first. Hence, an implementation SHOULD change the preference to source address choice first. There are a few methods below for how to choose a source address for any particular destination. The list is not exhaustive - it should be augmented later. The implementation MAY develop their method for choosing source address first."

The next 2 paragraphs of the original RFC 6724 should be preserved. The one is about choosing the source address that has the longest much with the destination address. Another one is equivalent to the methods proposed in [Conditional PIO].

Add 4 new methods for source address choice at the end of the section:

"The [Default Address] policy table may be updated by [Policy by DHCP] to guide source address selection.

The implementation may generate DNS requests from an address of every IPv6 PIO available to make sure that a particular source address has the reachability to the resource (split DNS may be implemented for "walled garden").

URL for configuration information could be supplied in RA - see [Provisioning domains].

The host may have local performance management capabilities (packet loss, delay, jitter, etc) to choose the best source for the application."

7.2. Default router choice by host

* Section 6.3.6 (Default Router Selection) of [ND], add an initial policy to default router selection:

0) For the cases when a particular implementation of ND does know the source address at the time of default router selection (it means that the source address was chosen first), then default routers that advertise the prefix for the respective source address SHOULD be preferred over routers that do not advertise the respective prefix.

7.3. Prefixes become dynamic

* This document joins the request to [CPE Requirements] that has been proposed in section 11 (General Requirements for HNCP Nodes) of [HNCP]:

The requirement L-13 to deprecate prefixes is applied to all delegated prefixes in the network from which assignments have been made on the respective interface. Furthermore, the Prefix Information Options indicating deprecation MUST be included in Router Advertisements for the remainder of the prefixes' respective valid lifetime, but MAY be omitted after at least 2 hours have passed.

* Add section 4.2 into [SLAAC]:

4.2 Dynamic Link Renumbering

Prefix delegation (primarily by DHCP-PD) is adopted by the industry as the primary mechanism of PA address delegation in fixed and mobile broadband environments, including cases of small businesses and branches of big enterprises.
The delegated prefix is tied to a dynamic link that has a considerable probability to be disconnected, especially in a mobile environment. The delegated prefix is losing value if the remote site

is disconnected from the prefix provider - this fact should be propagated to all nodes on the disconnected site, including hosts. Information Options indicating deprecation (multicast RA with zero Preferred Lifetime) MUST be sent at least one time. It SHOULD be included in Router Advertisements for the remainder of the prefixes' respective valid lifetime but MAY be omitted after 2 hours of deprecation announcements.

There is a high probability that connectivity to the provider would be restored very soon then the prefix could be announced again to all nodes on the site.

There is the probability that in a small period, the same problem would disconnect the site again (especially for mobile uplink). Such oscillation between available and not available providers could happen frequently which would flood the remote site with [ND] updates.
A dampening mechanism MAY be implemented to suppress oscillation: if the time between a particular prefix announcement and previous deprecation was less than DampeningCheck then delay the next prefix announcement for DampeningDelay and check the need for the prefix announcement after DampeningDelay seconds.
It is recommended for protocol designers implement a dampening mechanism for protocols (like [HNCP]) that would be used to distribute prefix delegation inside the site to relieve the majority of site routers and the protocol itself from the processing of oscillating messages.

* Section 5.1 (Node Configuration Variables) of [SLAAC], add timers:

DampeningCheck - the time between prefix announcement and previous deprecation is checked against this value to decide about the dampening need. The timer should use a 16-bit unsigned integer measured in seconds. The default value is 10 seconds.

DampeningDelay - the delay (penalty) for the next attempt to announce the same prefix again. The timer should use a 16-bit unsigned integer measured in seconds. The default value is 10 seconds.

These timers should be configurable like all other timers in [SLAAC] section 5.1.

7.4. Default router announcement rules

* This document joins [HNCP] section 11 (General Requirements for HNCP Nodes) request to [CPE Requirements]:

The generic requirements G-4 and G-5 are relaxed such that any known default router on any interface is sufficient for a router to announce itself as the default router; similarly, only the loss of all such default routers results in self-invalidation.

7.5. Clean prefixes for the changed default router list

* Section 6.3.5 (Timing out Prefixes and Default Routers) of [ND] has:

"Whenever the Lifetime of an entry in the Default Router List expires, that entry is discarded. When removing a router from the Default Router list, the node MUST update the Destination Cache in such a way that all entries using the router perform next-hop determination again rather than continue sending traffic to the (deleted) router."

Add at the end:

"All prefixes announced by the deprecated default router SHOULD be checked on the announcement from other default routers. If any prefix is not anymore announced from any router - it SHOULD be deprecated."

8. Interoperability analysis

This document mostly intersects with Homenet working group documents [HomeNet Architecture], [HNCP], and [MHMP]. This document simplifies the discussion in the [MHMP] solution of updating two tables on the host (routing and default address selection policy) by reversing the choice for the source address first.

[CPE Requirements] have the assumption of managing simplified topologies by manipulating routing information injection into [ND]. It has been shown in [MHMP] and in this document that it is better to signal reachability information to the host by the deprecation of delegated prefixes. This document joins [MHMP] request to change the approach.

This document does not contradict in any way to [Conditional PIO] or [MHMP Enterprise] that explain in detail the "equal prefixes" case but expend MHMP solution to the "non-equal prefixes" case.

[Happy Eyeballs] are sorting destination addresses. The proposals of this document are coming into the discussion after the destination addresses are chosen. Hence, the [Happy Eyeballs] operation is not impacted.

[Route Preferences] have been avoided as the mechanism for environments with PA address space because it is better to select the source address first for the more general case.
[Route Preferences] could still be applicable for PI (Provider-Independent) address environments where only next-hops need to be chosen properly.

9. Security Considerations

This document does not introduce new vulnerabilities.

10. IANA Considerations

This document has no request to IANA.

11. References

11.1. 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>.

[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>.

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

[SLAAC] 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>.

[Multi-Homing] F. Baker, 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>.

[Default Address] D. Thaler, R. Draves, A. Matsumoto, 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>.

[Policy by DHCP] A. Matsumoto, T. Fujisaki, T. Chown, "Distributing Address Selection Policy Using DHCPv6", RFC 7078 DOI 10.17487/RFC7078, January 2014, <https://www.rfc-editor.org/info/rfc7078>.

[CPE Requirements] 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>.

[MHMP Enterprise] F. Baker, C. Bowers, J. Linkova, "Enterprise Multihoming Using Provider-Assigned IPv6 Addresses without Network Prefix Translation: Requirements and Solutions", RFC 8678 DOI 10.17487/RFC8678, December 2019, <https://www.rfc-editor.org/info/rfc8678>.

[Provisioning domains] P. Pfister, E. Vyncke, T. Pauly, D. Schinazi, W. Shao, " Discovering Provisioning Domain Names and Data", RFC 8801 DOI 10.17487/RFC8801, July 2020, <https://www.rfc-editor.org/info/rfc8801>.

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

11.2. Informative References

[IPv6] S. Deering, R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

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

[Conditional PIO] J. Linkova, M. Stucchi, "Using Conditional Router Advertisements for Enterprise Multihoming", RFC 8475 DOI 10.17487/RFC8475, October 2018, <https://www.rfc-editor.org/info/rfc8475>.

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

[HomeNet Architecture] T. Chown, J. Arkko, A. Brandt, O. Troan, J. Weil, "IPv6 Home Networking Architecture Principles", RFC 7368, DOI 10.17487/RFC7368, October 2014, <https://www.rfc-editor.org/info/rfc7368>.

[Happy Eyeballs] D. Schinazi, T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, April 2016, <https://www.rfc-editor.org/info/rfc8305>.

[HNCP] M. Stenberg, S. Barth, P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <https://www.rfc-editor.org/info/rfc7788>.

12. Acknowledgments

Thanks to the 6man working group for problem discussion.

Authors' Addresses

Eduard Vasilenko
Huawei Technologies
17/4 Krylatskaya st, Moscow, Russia 121614
Email: vasilenko.eduard@huawei.com

Paolo Volpato
Huawei Technologies
Via Lorenteggio 240, 20147 Milan, Italy
Email: paolo.volpato@huawei.com