

Exceptions to the 64-bit Boundary in IPv6 Addressing
draft-farmer-6man-exceptions-64-00

Abstract

This document clarifies exceptions to the 64-bit boundary in IPv6 addressing. The exceptions include, unicast IPv6 addresses with the first three bits 000, manually configured addresses, DHCPv6 assigned addresses, IPv6 on-link determination, and the possibility of an exception specified in separate IPv6 link-type specific documents. Further, operational guidance is provided and [Appendix A](#). discusses the valid options for configuring IPv6 subnets

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 January 24, 2019.

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	2
1.1.	Requirements Language	4
2.	Exceptions to the 64-bit Boundary	5
2.1.	Unicast Addresses with the First Three Bits 000	5
2.2.	Manually Configured Addresses	5
2.3.	DHCPv6 Assigned Addresses	6
2.4.	IPv6 On-link Determination	6
2.5.	IPv6 Link-type Specific Documents	6
3.	Operational Guidance	7
4.	Acknowledgments	8
5.	IANA Considerations	8
6.	Security Considerations	8
7.	References	8
7.1.	Normative References	8
7.2.	Informative References	9
Appendix A.	Options for Configuring IPv6 Subnets	10
	Author's Address	13

[1.](#) Introduction

The 64-bit boundary in IPv6 addressing provides the basis for unicast addresses to be autonomously generated using stateless address auto-configuration (SLAAC) [RFC 4862](#) [[RFC4862](#)]. SLAAC allows hosts to connect to link networks without any pre-configuration, which is especially useful for general-purpose hosts and mobile devices. In this circumstance, unicast addresses have an internal structure composed of 64-bit interface identifiers (IIDs) and therefore 64-bit subnet prefixes, as defined in the IPv6 Addressing Architecture [[RFC4291bis](#)]. For additional discussion of the 64-bit boundary in IPv6 addressing see [RFC 7421](#) [[RFC7421](#)].

However, in other circumstances, such as with manually configured addresses or DHCPv6 [[RFC3315](#)] assigned addresses, unicast addresses are considered to have no internal structure and are assigned to interfaces on hosts as opaque 128-bit quantities without any knowledge of the subnets present on the link network. The idea that unicast addresses may have no internal structure is also defined in IPv6 Addressing Architecture [[RFC4291bis](#)], "a node may consider that unicast addresses (including its own) have no internal structure."

Further, unlike IPv4 where there is a single subnet mask parameter with the two aspects of a subnet, address assignment and on-link determination, tightly coupled together. In IPv6, these two aspects

Farmer

Expires January 24, 2019

[Page 2]

are split into two logically separate parameters serving the two aspects independently. The subnet assignment prefix is used by SLAAC to perform autonomous address assignment. Separately, the on-link prefix is used to determine if an address can be delivered using a directly connected link network. IPv6 Neighbor Discovery (ND) [[RFC4861](#)], the IPv6 subnet model [[RFC5942](#)], SLAAC [[RFC4862](#)] describe and specify the use of these parameters in detail.

Briefly, unicast addresses assigned to interfaces on hosts are not considered on-link unless covered by an on-link prefix advertised through ND Router Advertisement (RA) messages containing Prefix Information Options (PIOs) with the on-link (L) flag set or by manual configuration. Whereas autonomous address assignment uses subnet assignment prefixes that are also advertised through the same ND RA messages and PIOs but with the autonomous (A) flag set instead. While they act independently, most frequently subnets are defined by identical subnet assignment prefixes and on-link prefixes, see [Appendix A](#). for a further decision of this and the other valid options for configuring IPv6 subnets. However, unlike subnet assignment prefixes, which are effectively required to be 64-bits in length, on-link prefixes may have any length between 0 and 128 bits, inclusive. Nevertheless, for consistency with the 64-bit boundary, 64-bit on-link prefix lengths are recommended in most circumstances.

Reinforcing the ideas that on-link prefixes are logically separate and may have any length. On-link prefixes are part of the next-hop determination process in IPv6 ND, which is intrinsically part of routing and forwarding within IPv6, and [BCP 198](#) [[RFC7608](#)] says, "forwarding processes MUST be designed to process prefixes of any length up to /128, by increments of 1."

Finally, SLAAC is currently designed to utilize a single IID length to validate the length of the subnet assignment prefixes provided to it. However, SLAAC itself does not define the IID length or assume it is 64-bits in length. It utilizes the IID length defined in separate link-type specific documents that are intended to be consistent with the standard 64-bit IID length defined in the IPv6 Addressing Architecture [[RFC4291bis](#)]. While this is a possible exception to the 64-bit boundary, currently there are no IPv6 link-type specific documents that specify an IID length other than 64-bits. Effectively requiring 64-bit IIDs, and therefore 64-bit subnet assignment prefixes when use with autonomous address assignment, as performed by SLAAC.

In summary, the essential theory of this document is that the two parameters that define IPv6 subnets, the subnet assignment prefix and the on-link prefix, interact with the 64-bit boundary in subtle but complex ways. Subnet assignment prefixes are the primary parameter

Farmer

Expires January 24, 2019

[Page 3]

used to configure subnets, and when used they are effectively required to be 64-bit in length. However, this does not indicate on-link prefixes are also required to be 64-bits in length. Even when SLAAC is used, and subnets are required to be 64-bits in length, on-link prefixes shorter than 64-bits still seem to be valid. Further, when subnet assignment prefixes are not used to configure subnets, autonomous address assignment is not performed, and either manually configured addresses or DHCPv6 assigned addresses must be used. In this circumstance, subnets are configured solely using on-link prefixes and therefore may have any length between 0 and 128 bits, inclusive. Nevertheless, for consistency with the 64-bit boundary, 64-bit on-link prefix lengths are recommended in most circumstances. Therefore, when subnets are solely configured using on-link prefixes, subnets are only recommended to be 64-bit in length and are not required to be such.

Some have stated, "IPv6 subnets are required to be 64-bits in length." Whereas others counter, "IPv6 subnets are only recommended to be 64-bits in length." However, because of the subtle but complex interaction described above, both of these statements are not entirely correct based on the details of how individual subnets are configured. A more accurate statement is, "when configured using subnet assignment prefixes, IPv6 subnets are required to be 64-bits in length. Otherwise, when configured solely using on-link prefixes, IPv6 subnets are only recommended to be 64-bits in length." Also, it could be said, "standard IPv6 subnets are 64-bits in length," given the 64-bit length is both required or recommended based on the details of how individual subnets are configured. These last two statements seem to more accurately reflect how the protocols that define and implement IPv6 subnets operate. It is hoped that clarifying the following exceptions to the 64-bit boundary and providing clear operational guidance will provide a better understanding of and more clarity to the subtle but complex interaction between the 64-bit boundary in IPv6 addressing and how IPv6 subnets are defined and implemented.

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

Farmer

Expires January 24, 2019

[Page 4]

2. Exceptions to the 64-bit Boundary

2.1. Unicast Addresses with the First Three Bits 000

These are all currently special-purpose IPv6 addresses or are otherwise reserved. Also, they generally not assigned to interfaces on hosts, especially not to general-purpose hosts. Examples of these addresses are the unspecified address, the loopback address, and the IPv4-Mapped IPv6 Address from [RFC4291bis] sections [2.4.2](#), [2.4.3](#), and [2.4.5.2](#) respectively.

Most of these addresses have no internal structure and are considered opaque 128-bit quantities. However, some of these addresses could be presumed to have structure, such as the IPv4-mapped IPv6 address. This structure comes from embedding an IPv4 address within an IPv6 address, but this structure is unrelated to and different from the internal structure, composed of standard IIDs and subnet prefixes, which makes up the 64-bit boundary.

Historically, reservations were also made in this range for the mapping of OSI NSAP and IPX address into IPv6 addresses. They had structures similar to the IPv4-mapped IPv6 address discussed above. However, they have since been deprecated.

Note: ever since [RFC 2373](#) [RFC2373] addresses with the first three bits 000 have been an exception to the 64-bit boundary, and addresses with the first three bits 001 through 111, except for multicast addresses, have been expected to be consistent with the standard 64-bit IID length.

2.2. Manually Configured Addresses

IPv6 addresses manually configured on a node's interface, sometimes known as statically configured, are an exception to the 64-bit boundary as they have no internal structure, are considered opaque 128-bit quantities, and are assigned to node interfaces without any knowledge of the subnets present on the link network.

Manually configured addresses MAY also include an associated an on-link prefix length. This on-link prefix length (n) MAY have any value between 0 and 128 bits, inclusive. If an on-link prefix length is included, the most significant, or leftmost, n-bits of the manually configured address are considered the on-link prefix. Alternatively, if an on-link prefix length is not included, the manually configured address MUST NOT automatically be considered on-link. Nevertheless, for consistency with the 64-bit boundary, 64-bit on-link prefix lengths are recommended in most circumstances. See

Farmer

Expires January 24, 2019

[Page 5]

[section 3](#) for detailed operational guidance regarding on-link prefix lengths.

[2.3.](#) DHCPv6 Assigned Addresses

IPv6 addresses assigned to a host's interface via DHCPv6 [[RFC3315](#)] (Identity Association for Non-temporary Addresses (IA_NA) or Identity Association for Temporary Addresses (IA_TA)) are an exception to the 64-bit boundary as they have no internal structure, are considered opaque 128-bit quantities, and are assigned to host interfaces without any knowledge of the subnets present on the link network. Further, DHCPv6 assigned addresses MUST NOT automatically be considered on-link.

[2.4.](#) IPv6 On-link Determination

IPv6 on-link determination is an exception to the 64-bit boundary, in that IPv6 ND [[RFC4861](#)] does not require on-link prefixes to be 64-bits in length. To the contrary, on-link prefixes MAY have any length between 0 and 128 bits, inclusive. See [section 3](#) for detailed operational guidance regarding the use of on-link prefix lengths.

[2.5.](#) IPv6 Link-type Specific Documents

Separate IPv6 link-type specific documents, sometimes known as "IPv6-over-FOO" documents, specify the IID length utilized by SLAAC to validate the length of subnet assignment prefixes provided. The IID length defined should be consistent with the standard 64-bit IID length specified in the IPv6 addressing architecture [[RFC4291bis](#)]. However, these documents MAY create an exception to the standard 64-bit IID length scoped to a specific link-type technology when justified. Although currently, there are no IPv6 link-type specific documents that specify an IID length other than 64-bits.

When an exception to the standard 64-bit IID is specified in a link-type specific document, valid justification needs to be documented in some detail.

Further, SLAAC is currently designed to validate against only a single IID length per link-type technology. As a result, a link-type technology that specifies a non-standard IID length cannot be directly bridged with another link-type technology that specifies the standard 64-bit IID length without creating confusion about the IID length that is to be used for validation. Therefore, if this type of direct bridging is allowed, then a mechanism to ensure there is no confusion about which IID length SLAAC is to validate against needs to be provided.

Farmer

Expires January 24, 2019

[Page 6]

3. Operational Guidance

At a high-level, this document recommends the following principles for the configuration of IPv6 subnets. The configuration of subnet assignment prefixes is recommended, allowing hosts to use autonomous address assignment. With this configuration, subnet assignment prefixes are required to be 64-bits in length, requiring 64-bit subnets in this circumstance. Further, identical on-link prefixes are recommended, but on-link prefixes are required to be 64-bits or shorter. Otherwise, if subnet assignment prefixes are not configured, then hosts will have to use manually configured addresses or DHCPv6 assigned addresses and subnets are configured solely by on-link prefixes that are recommended to be 64-bits in length, only recommending 64-bit subnets in this circumstance. There are two exceptions to these principles, inter-router point-to-point links with 127-bit prefixes [[RFC6164](#)] and the possible future specification of link-type specific documents based on an IID length that is not 64-bits.

More specifically;

Network operators SHOULD configure routers to advertise to each link network at least one subnet assignment prefix (a PIO with the A flag set). If a subnet assignment prefix is advertised, it MUST be 64-bits in length and an identical on-link prefix (a PIO with the L flag set) SHOULD also be advertised. If an on-link prefix is advertised and is covered by a subnet assignment prefix, the on-link prefix MUST NOT be longer than 64-bits in length. If the specification for the particular link-type is based on an IID length that is not 64-bits, then a length consistent with the specification for the particular link-type MUST be used in the previous requirements instead of 64-bits.

Otherwise, if a subnet assignment prefix is not advertised, network operators SHOULD configure routers to advertise to each link network at least one on-link prefix (a PIO with the L flag set) that is 64-bits in length or provide the same manually configured on-link prefix to each host on the link network that is 64-bits in length. If the specification for the particular link-type is based on an IID length that is not 64-bits, then a length consistent with the specification for the particular link-type SHOULD be used in the previous recommendations instead of 64-bits.

Alternatively, network operators MAY configure point-to-point router links with 127-bit on-link prefixes and no subnet assignment prefix, see [RFC 6164](#) [[RFC6164](#)].

Farmer

Expires January 24, 2019

[Page 7]

[Appendix A](#). discusses in further detail the valid options for configuring IPv6 subnets

4. Acknowledgments

This document was inspired by a series of discussions on the 6MAN and the V6OPS working group mailing lists over a period of approximately two years, including discussions around the following drafts; [draft-jinmei-6man-prefix-clarify](#) [[I-D.jinmei-6man-prefix-clarify](#)], [draft-bourbaki-6man-classless-ipv6](#) [[I-D.bourbaki-6man-classless-ipv6](#)], [draft-jaeggli-v6ops-indefensible-nd](#) [[I-D.jaeggli-v6ops-indefensible-nd](#)]. All basically revolving around the discussion of RFC 4291bis [[RFC4291bis](#)] and its advancement to Internet Standard.

This document was produced using the xml2rfc tool [[RFC2629](#)].

5. IANA Considerations

This memo includes no request to IANA.

6. Security Considerations

This document clarifies exceptions to the 64-bit boundary in IPv6 addressing. These clarifications are not security related and therefore are not expected to introduce any new security considerations.

However, the use of longer on-link prefixes effectively allows the uses of smaller subnets, making it more feasible to perform IPv6 address scans as discussed in [RFC 7707](#) [[RFC7707](#)] and [RFC 7721](#) [[RFC7721](#)]. On the other hand, the use of smaller subnets can be effective mitigation for neighbor cache exhaustion issues as discussed and [RFC 6164](#) [[RFC6164](#)] and [RFC 6583](#) [[RFC6583](#)]. The relative weights applied in this trade-off will vary from situation to situation.

7. References

7.1. 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>>.

Farmer

Expires January 24, 2019

[Page 8]

- [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, <<https://www.rfc-editor.org/info/rfc3315>>.
- [RFC4291bis] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [draft-ietf-6man-rfc4291bis-09](#) (work in progress), July 2017, <<https://tools.ietf.org/id/draft-ietf-6man-rfc4291bis>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [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, <<https://www.rfc-editor.org/info/rfc5942>>.
- [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>>.
- [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>>.
- [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>>.

7.2. Informative References

- [I-D.bourbaki-6man-classless-ipv6] Bourbaki, N., "IPv6 is Classless", [draft-bourbaki-6man-classless-ipv6-03](#) (work in progress), March 2018.

Farmer

Expires January 24, 2019

[Page 9]

[I-D.jaeggli-v6ops-indefensible-nd]

Jaeggli, J., "Indefensible Neighbor Discovery", [draft-jaeggli-v6ops-indefensible-nd-01](#) (work in progress), July 2018.

[I-D.jinmei-6man-prefix-clarify]

Jinmei, T., "Clarifications on On-link and Subnet IPv6 Prefixes", [draft-jinmei-6man-prefix-clarify-00](#) (work in progress), March 2017.

[RFC2373] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 2373](#), DOI 10.17487/RFC2373, July 1998, <<https://www.rfc-editor.org/info/rfc2373>>.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", [RFC 2629](#), DOI 10.17487/RFC2629, June 1999, <<https://www.rfc-editor.org/info/rfc2629>>.

[RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", [RFC 6583](#), DOI 10.17487/RFC6583, March 2012, <<https://www.rfc-editor.org/info/rfc6583>>.

[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, <<https://www.rfc-editor.org/info/rfc7421>>.

[RFC7707] Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", [RFC 7707](#), DOI 10.17487/RFC7707, March 2016, <<https://www.rfc-editor.org/info/rfc7707>>.

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

[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>>.

[Appendix A](#). Options for Configuring IPv6 Subnets

As discussed in the Introduction, IPv6 subnets are defined by two separate parameters, acting independently, the subnet assignment prefix and the on-link prefix. It is possible to configure these parameters with several different relationships to each other. These

Farmer

Expires January 24, 2019

[Page 10]

parameters are primarily advertised in ND RA messages by PIOs, with the A and L flags designating the purpose of the PIO. However, on-link prefixes may also be manually configured.

SLAAC [\[RFC4862\]](#), [section 5.5.3](#) bullet d, validates subnet assignment prefixes against the IID length specified in separate link-type specific documents that are intended to be consistent with the standard 64-bit IID length. Currently, there are no link-type specific documents that specify a non-standard IID length. Therefore subnet assignment prefixes are effectively required to be 64-bit in length. Further, to simplify the following discussion the possibility that a link-type specific document could specify a non-standard IID length is ignored.

Whereas on-link prefixes have no such validation specified in IPv6 ND [\[RFC4861\]](#), this is also confirmed in SLAAC, [section 5.5.3](#) bullet d. Therefore on-link prefixes are not required to be 64-bits in length; they may have any length between 0 and 128 bits, inclusive. Nevertheless, for consistency with the 64-bit boundary, 64-bit on-link prefixes lengths are recommended, except for inter-router point-to-point links with 127-bit prefixes.

The following are the valid options for configuring the two parameters that define an IPv6 subnet;

1. Subnet assignment prefixes and identical on-link prefixes, or;
2. Subnet assignment prefixes and shorter covering on-link prefixes, or;
3. Only subnet assignment prefixes and no on-link prefixes or;
4. Only on-link prefixes and no subnet assignment prefixes

Options 1 through 3, all define subnet assignment prefixes, designating the use of autonomous address assignment, performed by SLAAC, and effectively requiring subnets that are 64-bits in length.

Option 1 is both the most frequently used and the only recommended option, except for inter-router point-to-point links with 127-bit prefixes, it has identical subnet assignment prefixes and on-link prefixes of 64-bits in length. The 64-bit subnets used for autonomous address assignment are considered to be on-link. This option is particularly recommended for networks that are made available to the general public or networks that intend to connect general-purpose hosts or mobile devices.

Farmer

Expires January 24, 2019

[Page 11]

Option 2 is not recommended, but is still valid; it has on-link prefixes shorter than 64-bits, between 0 and 63 bits, inclusive, but covering the subnet assignment prefixes included. The 64-bit subnets used for autonomous address assignment are considered on-link, along with other numerically adjacent subnets. However, these other numerically adjacent subnets are not used for autonomous address assignment unless additional separate 64-bit subnet assignment prefixes are also included.

Option 3 is not recommended, but is still valid; it has subnet assignment prefixes but no on-link prefixes. Therefore the 64-bit subnets used for autonomous address assignment are not considered on-link, and all traffic for the subnets, including host-to-host traffic, must be sent to a default router. See [RFC 8273](#) [[RFC8273](#)] for an example of this option.

Option 4 is not recommended, but is still valid; it has on-link prefixes but no subnet assignment prefixes, and therefore manually configured addresses or DHCPv6 assigned addresses must be used. When DHCPv6 is used a DHCPv6 server, or DHCPv6 relay will also be needed on the link network. The on-link prefixes may have any length between 0 and 128 bits, inclusive. However, 64-bit on-link prefixes are recommended, except for inter-router point-to-point links with 127-bit prefixes. This option effectively results in subnets that are defined only by the on-link prefixes, and therefore the subnets may have any lengths, even though 64-bits is recommended.

Furthermore, Option 4 essentially allows for the use of subnets longer than 64-bits. While this violates the spirit of the 64-bit boundary, technically it is not a violation of the 64-bit boundary; manually configured addresses, DHCPv6 assigned addresses, and on-link determination are all exceptions to the 64-bit boundary defined in this document. Nevertheless, for consistency with the 64-bit boundary, 64-bit on-link prefix lengths are recommended, effectively recommending 64-bit subnets, except for inter-router point-to-point links with 127-bit prefixes.

There can be operationally valid reasons for configuring subnets longer than 64-bits, and when a subnet is solely configured by an on-link prefix, longer subnets are not prohibited. [RFC 6164](#) [[RFC6164](#)] explicitly allows 127-bit prefixes for inter-router point-to-point links. Hence the explicit exceptions included for it. Additionally, [RFC 6583](#) [[RFC6583](#)] discusses "sizing subnets to reflect the number of addresses actually in use" as an operational mitigation for neighbor cache exhaustion issues. [RFC 7421 section 3](#) [[RFC7421](#)] discusses these issues in more detail. Nevertheless, address conservation by itself is never considered a valid reason for configuring subnets

Farmer

Expires January 24, 2019

[Page 12]

longer than 64-bits. Accordingly, if a site needs additional subnets, additional 64-bit subnets are expected to be provided.

Note: some hosts do not provide a mechanism for manually configuring an address on an interface, and other hosts do not implement DHCPv6. Hosts that implement neither and only implement SLAAC do exist and do not function on networks configured based on Option 4.

It possible to simultaneously configure multiple different subnets, associated with a single link network, each based on the same or different options described above. For example, there could be two different subnets based on Option 1 and one based on Option 4, all associated the same link network.

Logically there is another option that could define a subnet, "subnet assignment prefixes and longer covered on-link prefixes," but it does not result in an operationally valid subnet. While SLAAC and ND accept this configuration, it is particularly problematic and is considered an invalid configuration by the detailed operational guidance provided above. It would have on-link prefixes longer than 64-bits, between 65 and 128 bits, inclusive, that would be covered by an included subnet assignment prefix. Its use would result in the 64-bit subnet used for autonomous address assignment being inconsistently considered on-link for some address and not on-link for other addresses within the same subnet. This inconsistency creates a performance differential between addresses within the same subnet, which is inefficient and difficult to troubleshoot.

Author's Address

David Farmer
University of Minnesota
2218 University Ave SE
Minneapolis, MN 55414
US

Phone: +16126260815
Email: farmer@umn.edu

Farmer

Expires January 24, 2019

[Page 13]