6man Working Group
Internet-Draft
Intended status: Informational
Expires: September 14, 2017

T. Jinmei
Infoblox

March 13, 2017

# Clarifications on On-link and Subnet IPv6 Prefixes
# draft-jinmei-6man-prefix-clarify-00

## Abstract

The IPv6 Neighbor Discovery and Stateless Address Autoconfiguration protocols intentionally separate the handling of prefixes for their purposes: these prefixes can be different for the same link even if it may be uncommon in practice; validation for these purposes is expected to be performed separately and independently. Despite the revised text of the latest versions of these protocol specifications, it appears that the idea of this separation can still be easily misunderstood. This document clarifies the idea even more explicitly in order to set the common understanding of the intent of the current specifications.

## 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

Table of Contents



## 1. Introduction

[RFC4861] defines the Prefix Information option (PIO) of the Neighbor Discovery Router Advertisement (RA) message. It is used by receiving hosts for on-link determination, that is, determining which prefixes are to be considered on-link in the link through which the RA is delivered. [RFC4862] also uses the PIO for the purpose of stateless address autoconfiguration (SLAAC). While a single PIO is often used for both purposes, the corresponding specifications intend to handle them separately and independently. In particular, the prefix in a PIO is not expected to be considered to be invalid for on-link determination simply because the prefix fails to validate in SLAAC, or vice versa.

The idea of separating these two purposes had often been misunderstood in prior specifications of [RFC4861] and [RFC4862], so these successor RFCs tried to clarify the intent with additional text. According to a recent discussion at the 6man working group in the context of advancing [RFC4291], however, it turned out that it may still not be uncommon that the separation is misunderstood. There are even reportedly implementations that invalidate a prefix for on-link determination when it is invalid for SLAAC. Although this misunderstanding would normally not cause troubles in practical deployments, it can be a source of operational surprise or interoperability problems once an administrator wants to exploit the idea of this separation more "creatively".

This document clarifies the idea of separating on-link prefixes and SLAAC prefixes according to the original intent of the RFCs with some background history and implementation and testing status so that everyone can have the common understanding of the intent of the current protocol specifications. Even if the separation is not explicitly enforced with normative keywords, the circumstantial evidences discussed in this document hopefully help set the consistent expectation.

It would be worth noting that this document simply and redundantly repeats a point already described in [RFC5942] in a sense: "an IPv6 address isn't automatically associated with an IPv6 on-link prefix." However, this document specifically focuses on the fact that the prefix in a PIO has two independent meaning, while [RFC5942] seems to be more interested in cases like how a host is expected to consider a particular prefix as on-link, especially when the host configures the address through some other way than SLAAC. In particular, it would not be straightforward just from [RFC5942] that a host must not ignore a prefix for on-link determination due to prefix validation for SLAAC. This document specifically addresses such issues.

This document does not update any existing standards. It merely clarifies the intent of some standard documents. This document also does not intend to state the current intent of the standards is good or bad. Instead, it tries to set the accurate understanding of the current specifications in case the working group sees the need for revisiting them.

## 2. Requirements Language and Terminology

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.

For the purpose of clarification, this document introduces the following terminology to clearly distinguish two different types of prefixes that are the main subject of this document:

On-link prefix A prefix advertised in an RA PIO with the L flag being set.

SLAAC subnet prefix A prefix advertised in an RA PIO with the A flag being set.

Note that a single RA PIO can specify an on-link prefix and an SLAAC subnet prefix at the same time, i.e., when both the L and A flags are set. In fact, this is a typical configuration of prefixes advertised via RA messages.

It may also be worth noting that an SLAAC subnet prefix is not directly related to the "subnet prefix" as defined in Section 2.5 of [RFC4291], even if they are actually the same prefix in the end: both define some leading bits of IPv6 addresses followed by a specific length of interface identifier. From the perspective of [RFC4862], however, an SLAAC subnet prefix is just a prefix advertised via an RA PIO with the A flag being set. [RFC4862] does not directly refer to [RFC4291] to introduce the concept of the SLAAC subnet prefix. It would therefore make sense to explicitly call it an "SLAAC" subnet prefix.

For brevity, this document focuses on global unicast prefixes. Discussions on any other types of prefixes, including unicast link-local prefixes, are out of scope of this document.

## 3. Clarifications

[RFC4861] and [RFC4862] intentionally separate the concept and handling of on-link prefixes and SLAAC subnet prefixes, even when they are specified in the same single PIO.

[RFC4862] requires that the sum of the length of SLAAC subnet prefix and the length of the interface identifier be 128 (Section 5.5.3, bullet d), whatever the latter length is according to the link-type specification. This means that, for example, if the length of interface identifiers for the link is specified as 63, 64, or 65, then the SLAAC subnet prefix length must be 65, 64, or 63, respectively; otherwise the PIO is ignored for the purpose of SLAAC.

Note that [RFC4862] itself does not specify a particular prefix length of an SLAAC subnet prefix (which is determined as "128 - length-of-interface-identifier"). [RFC4862] only states in its Section 2 that the length of interface identifiers is defined in a specific link-type document and that it assumes the link-type specification is consistent with the IPv6 addressing architecture. If a link-type specification defines the interface identifier length to be 64 bits, the SLAAC subnet prefix must be 128 - 64 = 64 bits. But this is just logically derived from the link-type specification and the constraint of Section 5.5.3, bullet d; [RFC4862] itself does not directly specify the length of an SLAAC subnet prefix or an interface identifier.

[RFC4861] does not impose any restriction on the length of on-link prefixes. It can be any value between 0 and 128, inclusive. One of the original co-authors of the Neighbor Discovery specification confirmed that an on-link prefix can be different from an SLAAC subnet prefix [rfc2461bis], when the working group discussed a revision of the spec, which subsequently resulted in [RFC4861].

[RFC4861] added the following text to Section 6.3.4 in an attempt of clarifying the separation:

> [...] Similarly, [ADDRCONF] may impose certain restrictions on the prefix length for address configuration purposes. Therefore, the prefix might be rejected by [ADDRCONF] implementation in the host. However, the prefix length is still valid for on-link determination when combined with other flags in the prefix option.

where [ADDRCONF] means [RFC4862]. Along with the specific discussion at that time [rfc2461bis], this added text intends to note that a compliant implementation of [RFC4861] must not ignore an on-link prefix simply because it has an invalid length as an SLAAC subnet prefix per [RFC4862].

Similarly, [RFC4862] introduced the following paragraph in Section 5.5.3 by revising the prior version of it with additional sentences:

> It is the responsibility of the system administrator to ensure that the lengths of prefixes contained in Router Advertisements are consistent with the length of interface identifiers for that link type. It should be noted, however, that this does not mean the advertised prefix length is meaningless. In fact, the advertised length has non-trivial meaning for on-link determination in [RFC4861] where the sum of the prefix length and the interface identifier length may not be equal to 128. Thus, it should be safe to validate the advertised prefix length here, in order to detect and avoid a configuration error specifying an invalid prefix length in the context of address autoconfiguration.

The primary purpose of the added text was to clarify why the SLAAC subnet prefix is needed to be provided with its length and why it has to be validated at all even if it can be derived from the length of the interface identifier. But it also gives another circumstantial evidence that on-link and SLAAC subnet prefixes are independent. In fact, if an invalid length for an SLAAC subnet prefix made it also invalid as an on-link prefix, it would not have to note there is a case where the sum is not equal to 128.

In practice, the same prefix is often used for both on-link and SLAAC subnet purposes. In that case, RAs in that link would be configured to include a single PIO for that prefix with both the L and A flags set.

An on-link prefix could be longer than an SLAAC subnet prefix. For example, when the latter is 2001:db8:1:2::/64, the former could be 2001:db8:1:2:3:4::/80. Perhaps the administrator of the link wanted to force most of on-link communication to 2001:db8:1:2::/64 to be first sent to a default router, while allowing a particular set of hosts (those using the address that match the on-link prefix) to communicate directly. In this case, RAs in that link would have two PIOs, one for the /64 SLLAC subnet prefix with the A flag set but L flag cleared, and the /80 on-link prefix with the L flag set but A flag cleared.

An on-link prefix could also be shorter than an SLAAC subnet prefix. For example, when the latter is 2001:db8:1:2::/64, the former could be 2001:db8:1::/48. Perhaps the administrator of the link configured hosts from multiple /64 SLAAC subnet prefixes that match the /48 on-link prefix (2001:db8:1:1::/64, 2001:db8:1:2::/64, etc), presumably using unicast RAs, but wanted these hosts to communicate directly rather than via a default router. In this case, each of RAs in that link would have two PIOs, one for (one of) the /64 SLLAC subnet prefix with the A flag set but L flag cleared, and the /48 on-link prefix with the L flag set but A flag cleared.

Admittedly, the last two cases are imaginary examples. The author of this document does not know of actual deployments of these configurations. But the point is that the current protocol specifications intentionally allow such configurations.

Another, even more peculiar, example would be that a PIO with both the L and A flags set but its length is different from the length of SLAAC subnet prefix derived from the length of interface identifiers of the link defined in the link-specific documentation. In this case, a receiving host is expected to ignore the prefix as an SLAAC subnet prefix, as clearly described in Section 5.5.3 of [RFC4862]. In that sense this is most likely some kind of operational error: the administrator should either correct the prefix length or clear the A flag. However, the host is still expected to treat it as a valid on-link prefix. In a prior working group discussion it was explicitly confirmed that this is the intent of the protocol specification [rfc2461bis].

Finally, it may also be noteworthy that the fact that an on-link prefix can be of any length between 0 and 128 has nothing to do with how strictly or loosely the subnet prefix length is fixed in the IPv6

addressing architecture [RFC4291]. For example, the fact that the length of an on-link prefix can be 48, 80 or other bits is not contradictory to the fact that [RFC4291] states the interface identifiers (and therefore subnet) of a particular set of addresses are 64 bits in length. Only SLAAC subnet prefixes have an indirect implication with the addressing architecture as noted in Section 2.

## 4. Implementation Status

The separation of on-link prefixes and SLAAC subnet prefixes seems to be widely interpreted and implemented as clarified in this document. The author knows at least three (if not all) broadly used BSD variants, Free/Net/OpenBSDs, have supported the described behavior and they still do. Test 2.2.9 of the widely adopted "IPv6 Ready" test [v6ready] confirms that a compliant host does not ignore an on-link prefix simply because its length is different from the SLAAC subnet prefix length for the link. From these facts it should be safe to say the intent of the specification is generally well understood.

However, a recent discussion reveals that there is a different interpretation of [RFC4862], and there are reportedly implementations of that interpretation. This interpretation applies Section 5.5.3 of [RFC4862] to validation of on-link prefixes simply because the RFC does not explicitly say the restrictions are only applicable to SLAAC subnet prefixes. In particular, those implementations ignore a PIO as both on-link and SLAAC subnet prefix, and do so regardless of the A or L flag values, according to a working discussion [lwip].

As clarified in Section 3, this behavior is clearly against the intent of the specification, if not explicitly violating a statement marked with an [RFC2119] keyword. Besides, it does not make much sense to use the fact that [RFC4862] does not explicitly note Section 5.5.3 is only applicable to SLAAC purposes as the rationale for this behavior. This section also states: "If the Autonomous flag is not set, silently ignore the Prefix Information option." If we allowed the rationale, this restriction would mean any PIO with the L flag set must be ignored unless the A flag is also set. Obviously this does not make sense, which should also mean this section should read in the context of SLAAC only.

In practice, though, the "non-compliant" behavior may not be that harmful in practice. As noted in Section 3, it would be quite unlikely to see an on-link prefix whose length is not equal to the length of SLAAC subnet prefix (or 128 - interface-identifier-length if SLAAC is not used). Unless and until the administrator really wants to employ the unusual length of on-link prefixes, the difference of the behavior will not cause a problem. However, it

should be noted that once an administrator wants to use such an operation, it will be the implementation that is to be blamed, not the administrator.

## 5. Conclusions

[RFC4861] and [RFC4862] separate the concept and validation of on-link and SLAAC subnet prefixes. While these prefixes are often the same in a link, they can also be different. The RFCs do not intend to allow rejecting an on-link prefix due to a restriction on SLAAC subnet prefixes, or vice versa. Whether liked or hated, this is the fact, and the intent has been explicitly discussed and confirmed, and is widely implemented today. There are still reportedly some implementations that misunderstand the intent. While technically they would be considered non-compliant to the current specification, the effect of the violation would be limited in practice.

This document does not take a position about whether the intended concept and behavior described in [RFC4861] and [RFC4862] as clarified in this document is good or bad. It merely clarifies the original intent that seems to be sometimes misunderstood. The working group is free to revisit the original intent if it reaches a consensus as such, but it should be based on the correct understanding of the current specification; otherwise they would not even be able to agree that it is an update to the current specification or a clarification on things implemented incorrectly. This document hopefully helps set such a base understanding for any future discussions.

## 6. Security Considerations

This document only clarifies the intent of some existing standard documents, and does not make updates to them. The clarifications are not security issues. It is therefore not expected to introduce any new security considerations.

## 7. Acknowledgements

Kevin Lahey reviewed the initial draft of this document and provided feedback.

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

## 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>.

[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>.

## 8.2. Informative References

[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>.

[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>.

[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>.

[rfc2461bis] Nordmark, E., "Reception of prefix option with prefix length > 64", February 2004, <https://www.ietf.org/mail-archive/web/ipv6/current/msg01801.html>.

[v6ready] IPv6 Forum, "IPv6 Ready Phase-1/Phase 2 Test Specification Core Protocols", April 2010, <https://www.ipv6ready.org/docs/Core_Conformance_Latest.pdf>.

[lwip] Woodyatt, J., "A proposal for draft-ietf-6man-rfc4291bis-07", March 2017, <https://www.ietf.org/mail-archive/web/ipv6/current/msg26519.html>.

## Author's Address

Tatuya Jinmei
Infoblox