

6MAN Working Group
Internet-Draft
Updates: [RFC4291](#), [RFC4007](#) (if approved)
Intended status: Standards Track
Expires: October 26, 2019

A. Petrescu
CEA, LIST
L. Velvindron
Cyberstorm.mu
N. Kottapalli
Organisation
April 24, 2019

The length of the prefix of an IPv6 link-local address ranges from 10 to 127
[draft-petrescu-6man-ll-prefix-len-12](#)

Abstract

A rejected Errata to [RFC4291](#) "IPv6 Addr Archi" on the topic of link-local addresses 'needs' a draft. This is an answer to that need.

The length of the prefix of an IPv6 link-local address is variable. The minimal value is 10 decimal. The maximum value is 127 decimal.

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 October 26, 2019.

Copyright Notice

Copyright (c) 2019 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.	Definitions and Statements	2
2.	Terminology	3
3.	Context	3
4.	Example of use of LL Prefix Length 32	6
5.	Security Considerations	6
6.	IANA Considerations	7
7.	Contributors	7
8.	Acknowledgements	7
9.	Normative References	7
Appendix A.	ChangeLog	7
	Authors' Addresses	8

[1.](#) Definitions and Statements

The prefix of an IP address is formed by the *n* leftmost bits of the address. (in a left-to-right writing system).

The prefix of an IP address is used for goals such as: identify the type of an IPv6 address (link-local, global, others), identify the belonging of an IP address to a particular subnetwork, assist the forwarding (or not forwarding) decisions, and others.

The minimal length of the prefix of an IPv6 link-local address (the value of *n*) is equal to 10 decimal. The maximum is 127.

The prefix of an IPv6 link-local address is represented textually as "fe80::/n", where *n* MAY be any value between 10 and 127.

Regardless of the prefix length, the leftmost 10 bits of an IPv6 link-local address MUST be set to binary 1111111010 (hexadecimal fe80).

The illustration of an IPv6 link-local address is:

A memo describes the use of IPv6 link-local addresses in applications. The filename of the Internet Draft is [draft-smith-ipv6-link-locals-apps-00](#).

The RFC "IPv6 Address Archi" illustrates the format of the link-local addresses. From the illustration it MAY be understood that the length of the link-local prefix is 10 bits of value 111111010 and 54 0 bits.

IANA lists the "IPv6 prefix", and "Address Block", to be "fe80::/10" on its website. It is possible that in the future the IETF could decide to use the bits 11-53.

The [RFC 2464](#) "IPv6-over-Ethernet" states that the prefix for link-local addresses is "fe80::/64".

[RFC 6874](#), "Representing IPv6 Zone Identifiers in Address Literals and Uniform Resource Identifiers" specifies the link-local addresses to be under prefix "fe80::/10".

[RFC 8415](#) "DHCPv6" considers that link-local addresses are designated by the prefix fe80::/10.

[RFC 4007](#) "IPv6 Scoped Address Architecture" discusses Zone ID. A Zone ID may be used - internally - in the 54bits of a link-local address, even though these 54bits appear to be reset. The document mentions at a point that fe80::1 could be used in two separate physical links (not virtual, like the loopback).

[RFC4291](#) requires that an IPv6 link-local address be assigned on each interface. Yet, it does not require the link-local prefix to be associated to an interface.

[RFC4861](#) requires that the link local prefix be present in the Prefix List associated with an interface, although it does not specify the length of the link local prefix.

[RFC4862](#) "SLAAC" defines how GUAs and LLs self-configure. Whereas the GUA gets its prefix length from the RA (not from an RFC), the LL gets it from [RFC4291](#) (not from RA). They are independent choices based on distinct sources.

Several knowledgeable interpretations state that, generally speaking, the prefix length of link-local addresses is 10, but it is 64 in the particular case of Stateless Address-Autoconfiguration (SLAAC). In this latter case, the prefix is named a "subnet prefix", or "prefix on a link", and it is "fe80::/64".

The term "link-local prefix" is sometimes used to mean the prefix for on-link determination, and is sometimes used to mean the reserved address space for link-local addresses (including all current and future use). The latter is `fe80::/10`. Of which the address architecture spec only gives the addresses that match `fe80::/64` the standard format (by specifying intermediate 54 bits are all 0). As a result the former is (currently) only `fe80::/64`.

For people in the RIR world it's a common thing: you get a prefix from the RIR and then make assignments from it for specific purposes. You can route the aggregate allocation, but you're not allowed to use the unassigned parts (until you make an assignment). In this case `fe80::/10` is the allocation and `fe80::/64` is the assignment.

Independent testing shows that `'ifconfig add fe80:1::1'` works on linux but fails on openbsd. The same command works on a Cisco router.

Interpretations of the situation of linux working ok with `fe80:1::1` call it a violation of standards.

The address `fe80::1/128` is present on the loopback interface of BSD.

Implementations of an IPv6 stack in a particular operating system (linux) allow for the manual configuration of both prefix lengths 64 and 10 for link-local addresses.

In another operating system the prefix length for link-local addresses can not be explicitly specified by the end user, but may be indirectly derived from two distinct textual formats by using an unspecified rule.

In yet another operating system (FreeBSD) an end user can not use a link-local address whose value is `fe80:1::1`; because in that OS the hosts drop incoming packets whose source or destination address matches `fe80::/10` and contains a non-0 value in bits 15-31 (like `fe80:1::1` does). The URL of the C code in OpenBSD that leads to make that packet drop is

https://github.com/openbsd/src/blob/master/sys/netinet6/ip6_input.c

In a particular operating system (openbsd), it is possible to run SLAAC with Interface Identifiers of length different than 64, e.g. 100; this implements [RFC7217](#). In that same operating system it is not possible to use an Interface Identifier of length 100. At the same time, in another operating system (linux) it is possible to use Interface Identifiers of length 100, yet SLAAC does not work with IID that is not 64. In an ideal linux-bsd operating system any length of IID would be possible.

The loopback interface is required to have a link-local address too ([RFC4291](#)), although some OSs don't (linux). The [RFC4007](#) clarifies that, somehow.

Misconfigurations and lack of interoperability MAY arise between computers that use mixed prefix lengths for link-local addresses.

Historical note: earlier, the link-local prefix `fe80::/10` and site-local prefix `fec0::/10` were grouped into a common `fe80::/9`. If bits 10-64 were 0 then the prefix was a link-local, otherwise a site-local. The site-local addresses were later deprecated by [RFC 3879](#).

4. Example of use of LL Prefix Length 32

This figure shows two routers each with two interfaces; one such interface is connected to the other router; there are two interfaces that point elsewhere.

```

i1 ----- i2      i3-----i4
--|Router1|-----|Router2|---
-----

```

i2 address is `fe80:12::1:1/32` ('12' means subnet between R1 and R2, '1' is R1, 2nd '1' is 'front' interface)
i3 address is `fe80:12::2:2/32`

Figure 2: Figure

One router's interface (connected to the other router) uses address `fe80:12::1:1/32` and the other router's corresponding interface uses address `fe80:12::2:2/32`.

5. Security Considerations

The clarification of the definition of the prefix length of the IPv6 link-local prefix at IANA is: call it 'leading bits' and not 'prefix', or state that the IPv6 prefix length of link-local addresses is 10 decimal. This clarification has beneficial impact in the algorithm implementation for calculation of the opaque and stable Interface Identifiers for IPv6 link-local addresses. It also positively impacts some implementations of IPv6 forwarding.

6. IANA Considerations

IANA is requested to change the name of the column head in the table that depicts the "Internet Protocol Version 6 Address Space". The name should be "The n leading bits of an address" instead of "IPv6 Prefix".

The desired effect of this change is that the IPv6 link-local prefix be "fe80::/n" and that the 10 leading bits of this prefix be 1111111010. A second effect would be that the textual representation "fe80::/10" as an IPv6 link-local prefix would disappear from that IANA page.

7. Contributors

Listed from 6man WG discussion.

8. Acknowledgements

The following persons are acknowledged for the discussion that is reflected in this draft. Not all points are reflected. Some points are copied almost entirely.

Ole Troan, Scott Timothy Morizot, Brian Carpenter, Fred Baker, Mark Smith, Peter Occil, Philip Homburg, Albert Manfredi,  –3/4 ’BAE (TATUYA Jinmei), Fernando Gont, Christian Huitema, Simon Hobson, Matthew Petach, Yucel Guven, Sander Steffann, Dennis Ferguson, Musa Stephen Honlue, Fred Templin, Gyan Mishra, Yu Tianpeng.

Peter Paluch submitted the Errata suggestion to [RFC 4291](#) about link-local addresses, and Brian Haberman rejected it, by requiring a draft. Igor Lubashev pointed to that Errata.

9. 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>>.

[Appendix A](#). ChangeLog

The changes are listed in reverse chronological order, most recent changes appearing at the top of the list.

-12: the '64' in GUA vs '64' in LL issued by distinct sources: RA vs [RFC4291](#) respectively; the address fe80::1/128 is present on the

loopback interface of BSD; detailed, again, the distinction for 'on-link' determination; detailed, again, the distinction between 'assignment' and 'allocation'; added the fact that Cisco supports manual assignment of fe80:1::1.

-11: trying the attribute updates=[RFC4291](#),RFC4007 in the rfc tag.

-10: syntax error corrected; more explanation about how FreeBSD C code blocks fe80:1::1; clarification in IANA section, but doubtful.

-09: added a reference to [RFC 4007](#) about Zone ID in LL; added a reference to [draft-bourbaki](#) about IPv6 being classless; added the result of independent testing showing ifconfig add fe80:1::1 works on linux but fails on BSD; added URL to C code in BSD flavor that may be in charge of dropping packets whose src/dst is an LL like fe80:1::1; added two co-authors.

-08: added explanation of which RFC requires the LL address to be present, and which requires the LL prefix to be present; named the OSs, instead of staying generic; explained that the lack of requirement of ll address on lo in [RFC4291](#) is covered by another [RFC4007](#); explained that openbsd allows variable len IID for GUAs but not for LLs, yet linux allows the reverse, and concluded on an obvious ideal.

-07: added the fact that DHCPv6 spec considers the link-local addresses to be fe80::/10; added a valuable explanation of ll behaviour of a particularly important OS.

-04: added an example advantage of using prefix length 32.

-03:

-02: corrected a typo in "fe80::/1" and added a 7-bit encoding for one persons name (in addition to the japanese-shift-jis encoding which is not understood by xml2rfc.)

Authors' Addresses

Alexandre Petrescu
CEA, LIST
CEA Saclay
Gif-sur-Yvette , Ile-de-France 91190
France

Phone: +33169089223
Email: Alexandre.Petrescu@cea.fr

Loganaden Velvindron
Cyberstorm.mu
street
city , region code
country

Phone: +phonenumber
Email: loganaden@gmail.com

Naveen Kottapalli
Organisation
street
City , Region code
Country

Phone: +phonenumber
Email: naveen.sarma@gmail.com