

v6ops Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: November 27, 2015

M. Boucadair
France Telecom
A. Petrescu
CEA, LIST
F. Baker
Cisco Systems
May 26, 2015

IPv6 Prefix Length Recommendation for Forwarding
draft-ietf-v6ops-cidr-prefix-03

Abstract

IPv6 prefix length, as in IPv4, is a parameter conveyed and used in IPv6 routing and forwarding processes in accordance with the Classless Inter-domain Routing (CIDR) architecture. The length of an IPv6 prefix may be any number from zero to 128, although subnets using stateless address autoconfiguration (SLAAC) for address allocation conventionally use a /64 prefix. Hardware and software implementations of routing and forwarding should therefore impose no rules on prefix length, but implement longest-match-first on prefixes of any valid length.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 November 27, 2015.

Internet-Draft

May 2015

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Recommendation	3
3.	IANA Considerations	3
4.	Security Considerations	3
5.	Acknowledgements	4
6.	References	4
6.1.	Normative References	4
6.2.	Informative References	4
	Authors' Addresses	5

[1.](#) Introduction

Discussions on the 64-bit boundary in IPv6 addressing ([\[RFC7421\]](#)) revealed a need for a clear recommendation on which bits must be used by forwarding decision-making processes. However, such a recommendation was out of scope for that document.

Although [Section 2.5 of \[RFC4291\]](#) states "IPv6 unicast addresses are aggregatable with prefixes of arbitrary bit-length, similar to IPv4 addresses under Classless Inter-Domain Routing" (CIDR, [\[RFC4632\]](#)), there is still a misinterpretation that IPv6 prefixes can be either /127 ([\[RFC6164\]](#)) or any length up to /64. This misinterpretation is mainly induced by the 64-bit boundary in IPv6 addressing.

As discussed in [\[RFC7421\]](#), "the notion of a /64 boundary in the address was introduced after the initial design of IPv6, following a

period when it was expected to be at /80". This evolution of the IPv6 Addressing architecture, resulting in [\[RFC4291\]](#), and followed with the addition of /127 prefixes for point-to-point links, clearly demonstrates the intent for future IPv6 developments to have the flexibility to change this part of the architecture when justified.

It is fundamental not to link routing and forwarding to the IPv6 prefix/address semantics [\[RFC4291\]](#). This document includes a recommendation for that aim.

Forwarding decisions rely on the longest-match-first algorithm, which stipulates that, given a choice between two prefixes in the Forwarding Information Base (FIB) of different length that match the destination address in each bit up to their respective lengths, the longer prefix is used. This document's recommendation ([Section 2](#)) is that IPv6 forwarding must follow the longest-match-first rule, regardless of prefix length, unless some overriding policy is configured.

This recommendation does not conflict with the 64-bit boundary for some IPv6 stateless address autoconfiguration (SLAAC, [\[RFC4862\]](#)) based schemes such as [\[RFC2464\]](#). Indeed, [\[RFC7421\]](#) clarifies this is only a parameter in the SLAAC process and other longer prefix lengths are in operational use (e.g., either manually configured or based upon DHCPv6 [\[RFC3315\]](#)).

A historical background of CIDR is documented in [\[RFC1380\]](#) and [Section 2 of \[RFC4632\]](#).

2. Recommendation

IPv6 implementations MUST conform to the rules specified in [Section 5.1 of \[RFC4632\]](#).

Decision-making processes for forwarding MUST NOT restrict the length of IPv6 prefixes by design. In particular, forwarding processes MUST be designed to process prefixes of any length up to /128, by increments of 1.

Policies can be enforced to restrict the length of IP prefixes advertised within a given domain or in a given interconnection link. These policies are deployment-specific and/or driven by

administrative (interconnection) considerations.

3. IANA Considerations

This document does not require any action from IANA.

4. Security Considerations

This document does not introduce security issues in addition to what is discussed in [[RFC4291](#)].

IPv6 security issues, including operational ones, are discussed in [[RFC4942](#)] and [[I-D.ietf-opsec-v6](#)].

5. Acknowledgements

Thanks to Eric Vyncke, Christian Jacquenet, Brian Carpenter, Fernando Gont, Tatuya Jinmei, Lorenzo Colitti, Ross Chandler, David Farmer, David Black, and Barry Leiba for their contributions and comments.

Special thanks to Randy Bush for his support.

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.

[RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", [BCP 122](#), [RFC 4632](#), August 2006.

6.2. Informative References

[[I-D.ietf-opsec-v6](#)]
Chittimaneni, K., Kaeo, M., and E. Vyncke, "Operational

Security Considerations for IPv6 Networks", [draft-ietf-opsec-v6-06](#) (work in progress), March 2015.

- [RFC1380] Gross, P. and P. Almqvist, "IESG Deliberations on Routing and Addressing", [RFC 1380](#), November 1992.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", [RFC 2464](#), December 1998.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), September 2007.
- [RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/ Co-existence Security Considerations", [RFC 4942](#), September 2007.

Boucadair, et al.

Expires November 27, 2015

[Page 4]

Internet-Draft

May 2015

- [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](#), April 2011.
- [RFC7421] Carpenter, B., Chown, T., Gont, F., Jiang, S., Petrescu, A., and A. Yourtchenko, "Analysis of the 64-bit Boundary in IPv6 Addressing", [RFC 7421](#), January 2015.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

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

France

Phone: +33169089223

Email: alexandre.petrescu@cea.fr

Fred Baker

Cisco Systems

Santa Barbara, California 93117

USA

Email: fred@cisco.com