Internet Engineering Task Force (IETF)
Request for Comments: 7608
BCP: 198
Category: Best Current Practice
ISSN: 2070-1721

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

July 2015

IPv6 Prefix Length Recommendation for Forwarding

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.

Status of This Memo

This memo documents an Internet Best Current Practice.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on BCPs is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7608.

Table of Contents

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 in order to support that goal.

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 schemes that based on IPv6 stateless address autoconfiguration (SLAAC) [RFC4862], 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].

1.1. 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].

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. 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 [OPSEC-v6].

4. References

4.1. 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>.

[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>.

[RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, DOI 10.17487/RFC4632, August 2006, <http://www.rfc-editor.org/info/rfc4632>.

4.2. Informative References

[OPSEC-v6] Chittimaneni, K., Kaeo, M., and E. Vyncke, "Operational Security Considerations for IPv6 Networks", Work in Progress, draft-ietf-opsec-v6-06, March 2015.

[RFC1380] Gross, P. and P. Almquist, "IESG Deliberations on Routing and Addressing", RFC 1380, DOI 10.17487/RFC1380, November 1992, <http://www.rfc-editor.org/info/rfc1380>.

[RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <http://www.rfc-editor.org/info/rfc2464>.

[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, <http://www.rfc-editor.org/info/rfc3315>.

[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>.

[RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/ Co-existence Security Considerations", RFC 4942, DOI 10.17487/RFC4942, September 2007, <http://www.rfc-editor.org/info/rfc4942>.

[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, <http://www.rfc-editor.org/info/rfc6164>.

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

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.

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
United States

Email: fred@cisco.com