IPv6 is Classless

Abstract

Over the history of IPv6, various classful address models have been proposed, none of which has withstood the test of time. The last remnant of IPv6 classful addressing is a rigid network interface identifier boundary at /64. This document removes the fixed position of that boundary for interface addressing.

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 14 April 2023.

Table of Contents

1. Introduction

Over the history of the IPv6 protocol, several classful addressing models have been proposed. The most notable example recommended Top-Level Aggregation (TLA) and Next-Level Aggregation (NLA) Identifiers [RFC2450], but was obsoleted by [RFC3587], leaving a single remnant of classful addressing in IPv6: a rigid network interface identifier boundary at /64. This document removes the fixed position of that boundary for interface addressing.

Recent proposed changes to the IP Version 6 Addressing Architecture specification [RFC4291] have caused controversy. While link prefixes of varied lengths, e.g. /127, /126, /124, /120, ... /64 have been successfully deployed for many years, glaring mismatches between a formal specification and long-standing field deployment practices are never wise, not least because of the strong risk of mis-implementation, which can easily result in serious operational problems.

This document also stresses that IPv6 routing subnets may be of any length up to 128, see [RFC7608].

2. Suggested Reading

It is assumed that the reader understands the history of classful addressing in IPv4 and why it was abolished [RFC4632]. Of course, the acute need to conserve address space that forced the adoption of classless addressing for IPv4 does not apply to IPv6, but the arguments for operational flexibility in address assignment remain compelling.

It is also assumed that the reader understands IPv6 [RFC2460], the
IP Version 6 Addressing Architecture [RFC4291], the proposed changes
to RFC4291 [I-D.ietf-6man-rfc4291bis] and RFC2464 [I-D.hinden-6man-
rfc2464bis], [RFC7608] an IPv6 Prefix Length Recommendation for
Forwarding, and the IETF recommendation for the generation of stable
Interface Identifiers [RFC8064].

[I-D.jinmei-6man-prefix-clarify] is also worth reading to clarify
uses of varying prefix lengths on a single link.

3. Problems Reinforced by Classful Addressing

For host computers on local area networks, generation of interface
identifiers is no longer necessarily bound to layer 2 addresses
(MACs) [RFC7217] [RFC8064]. Therefore their length, previously fixed
at 64 bits [RFC7136], is in fact a variably-sized parameter as
explicitly acknowledged in Section 5.5.3(d) of [RFC4862] which
states:

> Note that a future revision of the address architecture [RFC4291]
> and a future link-type-specific document, which will still be
> consistent with each other, could potentially allow for an
> interface identifier of length other than the value defined in
> the current documents. Thus, an implementation should not assume
> a particular constant. Rather, it should expect any lengths of
> interface identifiers

As IPv6 use has evolved and grown, it has become evident that it
faces several scaling and coordination problems. These problems are
analogous to allocation and coordination problems that motivated
IPv4 CIDR allocation and later abundant IPv4 PAT, they include:

> Address allocation models for specific counts of fixed length
> subnets to downstream networks or devices from /48 down to /64
> are based on design assumptions of how subnets are or should be
> allocated and populated within IPv4 networks.

> Hierarchical allocation of fixed-length subnets requires
> coordination between lower / intermediate / upper network
> elements. It has implicit assumption that policies and size
> allocation allowed at the top of the hierarchy will accommodate
> present and future use cases with fixed length subnet allocation.

> Coordination with upstream networks across administrative domains
> for the allocation of fixed length subnets reveals topology and
> intent that may be private in scope, allowing the upstream
> networks to restrict the topology that may be built. Policies for
> hierarchical allocation are applied top-down and amount to
> permission to build a particular topology (for example mobile

device tethering, virtual machine instantiation, containers and
so on).

In the case where a device is given a /64 (e.g. mobile phone
running SLAAC only, not DHCP), there is no protocol allowing them
to provide downstream routed layer 3 subnets, because all they
have is a /64. This applies more to nodes which do not have
DHCPv6.

4. Identifier and Subnet Length Statements

IPv6 unicast interfaces may use any subnet length up to 128 except
for situations where an Internet Standard document may impose a
particular length, for example Stateless Address Autoconfiguration
(SLAAC) [RFC4862], or Using 127-Bit IPv6 Prefixes on Inter-Router
Links [RFC6164].

Additionally, this document clarifies that a node or router MUST
support routing of any valid network prefix length, even if SLAAC or
other standards are in use, because routing could choose to
differentiate at a different granularity than is used by any such
automated link local address configuration tools.

5. Recommendations

For historical reasons, when a prefix is needed on a link, barring
other considerations, a /64 is recommended [RFC7136].

The length of the Interface Identifier in Stateless Address
Autoconfiguration [RFC4862] is a parameter; its length SHOULD be
sufficient for effective randomization for privacy reasons. For
example, 48 bits might be sufficient. But operationally we
recommend, barring strong considerations to the contrary, using 64-
bits for SLAAC in order not to discover bugs where 64 was hard-
coded, and to favor portability of devices and operating systems.

Note that OpenBSD ships with SLAAC for lengths longer than /64.

Nonetheless, there is no reason in theory why an IPv6 node should
not operate with different interface identifier lengths on different
physical interfaces. Thus, a correct implementation of SLAAC must in
fact allow for any prefix length, with the value being a parameter
per interface. For instance, the Interface Identifier length in the
recommended (see [RFC8064]) algorithm for selecting stable interface
identifiers [RFC7217] is a parameter, rather than a hard-coded
value.

6. Security Considerations

Assuming that nodes employ unpredictable interface identifiers [RFC7721], the subnet size may have an impact on some security and privacy properties of a network. Namely, the smaller the subnet size, the more feasible it becomes to perform IPv6 address scans [RFC7707] [RFC7721]. For some specific subnets, such as point to point links, this may be less of an issue.

On the other hand, we assume that a number of IPv6 implementations fail to enforce limits on the size of some of the data structures they employ for communicating with neighboring nodes, such as the Neighbor Cache. In such cases, the use of smaller subnets forces an operational limit on such data structures, thus helping mitigate some pathological behaviors (such as Neighbor Cache Exhaustion attacks).

7. IANA Considerations

This document has no IANA Considerations.

8. Authors

The authors of this document are as follows:

Randy Bush <randy@psg.com>, Internet Initiative Japan

Brian Carpenter <brian.e.carpenter@gmail.com>, University of Auckland

Fernando Gont <fgont@si6networks.com>, SI6 Networks / UTN-FRH

Nick Hilliard <nick@netability.ie>, INEX

Joel Jaeggli <joelja@bogus.com>, Fastly

Geoff Huston <gih@apnic.net>, APNIC

Chris Morrow <morrowc@ops-netman.net>, Google, Inc.

Job Snijders <job@net.net>, NTT Communications

9. References

9.1. Normative References

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6
           (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460,
           December 1998, <https://www.rfc-editor.org/info/rfc2460>.

[RFC4291]
Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <https://www.rfc-editor.org/info/rfc7217>.

[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>.

[RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", RFC 8064, DOI 10.17487/RFC8064, February 2017, <https://www.rfc-editor.org/info/rfc8064>.

9.2. Informative References

[I-D.hinden-6man-rfc2464bis] Crawford, M. and R. M. Hinden, "Transmission of IPv6 Packets over Ethernet Networks", Work in Progress, Internet-Draft, draft-hinden-6man-rfc2464bis-02, 13 March 2017, <https://www.ietf.org/archive/id/draft-hinden-6man-rfc2464bis-02.txt>.

[I-D.ietf-6man-rfc4291bis] Hinden, R. M. and S. E. Deering, "IP Version 6 Addressing Architecture", Work in Progress, Internet-Draft, draft-ietf-6man-rfc4291bis-09, 3 July 2017, <https://www.ietf.org/archive/id/draft-ietf-6man-rfc4291bis-09.txt>.

[I-D.jinmei-6man-prefix-clarify]
Jinmei, T., "Clarifications on On-link and Subnet IPv6 Prefixes", Work in Progress, Internet-Draft, draft-jinmei-6man-prefix-clarify-00, 13 March 2017, <https://www.ietf.org/archive/id/draft-jinmei-6man-prefix-clarify-00.txt>.

[RFC2450] Hinden, R., "Proposed TLA and NLA Assignment Rule", RFC 2450, DOI 10.17487/RFC2450, December 1998, <https://www.rfc-editor.org/info/rfc2450>.

[RFC3587] Hinden, R., Deering, S., and E. Nordmark, "IPv6 Global Unicast Address Format", RFC 3587, DOI 10.17487/RFC3587, August 2003, <https://www.rfc-editor.org/info/rfc3587>.

[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, <https://www.rfc-editor.org/info/rfc4632>.

[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>.

[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>.

[RFC7136] Carpenter, B. and S. Jiang, "Significance of IPv6
Interface Identifiers", RFC 7136, DOI 10.17487/RFC7136,
February 2014, <https://www.rfc-editor.org/info/rfc7136>.

[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>.

Author's Address

Nicolas Bourbaki
The Intertubes
42 Rue du Jour
::1 Sophia-Antipolis
France

Email: bourbaki@bogus.com