

Internet Engineering Task Force (IETF)
Request for Comments: 8215
Category: Standards Track
ISSN: 2070-1721

T. Anderson
Redpill Linpro
August 2017

Local-Use IPv4/IPv6 Translation Prefix

Abstract

This document reserves the IPv6 prefix 64:ff9b:1::/48 for local use within domains that enable IPv4/IPv6 translation mechanisms.

Status of This Memo

This is an Internet Standards Track document.

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 Internet Standards is available in Section 2 of RFC 7841.

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/rfc8215>.

Copyright Notice

Copyright (c) 2017 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. Terminology	2
3. Problem Statement	2
4. Why 64:ff9b:1::/48?	3
4.1. Prefix Length	3
4.2. Prefix Value	4
5. Deployment Considerations	4
6. Checksum Neutrality	5
7. IANA Considerations	6
8. Security Considerations	6
9. References	6
9.1. Normative References	6
9.2. Informative References	7
Acknowledgements	7
Author's Address	7

1. Introduction

This document reserves 64:ff9b:1::/48 for local use within domains that enable IPv4/IPv6 translation mechanisms. This facilitates the coexistence of multiple IPv4/IPv6 translation mechanisms in the same network without requiring the use of a Network-Specific Prefix assigned from the operator's allocated global unicast address space.

2. Terminology

This document uses the following terms:

Network-Specific Prefix (NSP)

A globally unique prefix assigned by a network operator for use with an IPv4/IPv6 translation mechanism [RFC6052].

Well-Known Prefix (WKP)

The prefix 64:ff9b::/96, which is reserved for use with the [RFC6052] IPv4/IPv6 address translation algorithms.

3. Problem Statement

Since the WKP 64:ff9b::/96 was reserved by [RFC6052], several new IPv4/IPv6 translation mechanisms have been defined by the IETF, such as those defined in [RFC6146] and [RFC7915]. These mechanisms target various different use cases. An operator might therefore wish to make use of several of them simultaneously.

The WKP is reserved specifically for use with the algorithms specified in [RFC6052]. More recent RFCs describe IPv4/IPv6

translation mechanisms that use different algorithms. An operator deploying such mechanisms cannot make use of the WKP in a legitimate fashion.

Also, because the WKP is a /96, an operator preferring to use the WKP over an NSP can do so for only one of their IPv4/IPv6 translation mechanisms. All others must necessarily use an NSP.

Section 3.1 of [RFC6052] imposes certain restrictions on the use of the WKP, such as forbidding its use in combination with private IPv4 addresses [RFC1918]. These restrictions might conflict with the operator's desired use of an IPv4/IPv6 translation mechanism.

In summary, there is a need for a local-use prefix that facilitates the coexistence of multiple IPv4/IPv6 translation mechanisms in a single network domain, as well as the deployment of translation mechanisms that do not use the [RFC6052] algorithms or adhere to its usage restrictions.

4. Why 64:ff9b:1::/48?

4.1. Prefix Length

One of the primary goals of this document is to facilitate multiple simultaneous deployments of IPv4/IPv6 translation mechanisms in a single network. The first criterion is therefore that the prefix length chosen must be shorter than the prefix length used by any individual translation mechanism.

The second criterion is that the prefix length chosen is a multiple of 16. This ensures the prefix ends on a colon boundary when representing it in text, easing operator interaction with it.

The [RFC6052] algorithms specifies IPv4/IPv6 translation prefixes as short as /32. In order to facilitate multiple instances of translation mechanisms using /32s, while at the same time aligning on a 16-bit boundary, it would be necessary to reserve a /16. Doing so, however, was considered as too wasteful by the IPv6 Operations Working Group.

The shortest translation prefix that was reported to the IPv6 Operations Working Group as being deployed in a live network was /64. The longest 16-bit-aligned prefix length that can accommodate multiple instances of /64 is /48. The prefix length of /48 was therefore chosen, as it satisfies both the criteria above, while at the same time avoids wasting too much of the IPv6 address space.

4.2. Prefix Value

It is desirable to minimise the amount of additional "pollution" in the unallocated IPv6 address space caused by the reservation made by this document. Ensuring the reserved prefix is adjacent to the 64:ff9b::/96 WKP already reserved by [RFC6052] accomplishes this.

Given the previous decision to use a prefix length of /48, this leaves two options: 64:ff9a:ffff::/48 and 64:ff9b:1::/48.

64:ff9a:ffff::/48 has the benefit that it is completely adjacent to the [RFC6052] WKP. That is, 64:ff9a:ffff::/48 and 64:ff9b::/96 combine to form an uninterrupted range of IPv6 addresses starting with 64:ff9a:ffff:: and ending with 64:ff9b:ffff:ffff.

64:ff9b:1::/48 is, on the other hand, not completely adjacent to 64:ff9b::/96. The range starting with 64:ff9b:1:0:0 and ending with 64:ff9b:0:ffff:ffff:ffff:ffff:ffff would remain unallocated.

This particular drawback is, however, balanced by the fact that the smallest possible aggregate prefix that covers both the [RFC6052] WKP and 64:ff9a:ffff::/48 is much larger than the smallest possible aggregate prefix that covers both the [RFC6052] WKP and 64:ff9b:1::/48. These aggregate prefixes are 64:ff9a::/31 and 64:ff9b::/47, respectively. IPv6 address space is allocated using prefixes rather than address ranges, so it could be argued that 64:ff9b:1::/48 is the option that would cause special-use prefixes reserved for IPv4/IPv6 translation to "pollute" the minimum possible amount of unallocated IPv6 address space.

Finally, 64:ff9b:1::/48 also has the advantage that its textual representation is shorter than 64:ff9a:ffff::/48. While this might seem insignificant, the preference human network operators have for addresses that are simple to type should not be underestimated.

After weighing the above pros and cons, 64:ff9b:1::/48 was chosen.

5. Deployment Considerations

64:ff9b:1::/48 is intended as a technology-agnostic and generic reservation. A network operator may freely use it in combination with any kind of IPv4/IPv6 translation mechanism deployed within their network.

By default, IPv6 nodes and applications must not treat IPv6 addresses within 64:ff9b:1::/48 differently from other globally scoped IPv6 addresses. In particular, they must not make any assumptions regarding the syntax or properties of those addresses (e.g., the

existence and location of embedded IPv4 addresses) or the type of associated translation mechanism (e.g., whether it is stateful or stateless).

64:ff9b:1::/48 or any more-specific prefix may only be used in inter-domain routing if done in accordance with the rules described in Section 3.2 of [RFC6052].

Note that 64:ff9b:1::/48 (or any more-specific prefix) is distinct from the WKP 64:ff9b::/96. Therefore, the restrictions on the use of the WKP described in Section 3.1 of [RFC6052] do not apply to the use of 64:ff9b:1::/48.

Operators tempted to use the covering aggregate prefix 64:ff9b::/47 to refer to all special-use prefixes currently reserved for IPv4/IPv6 translation should be warned that this aggregate includes a range of unallocated addresses (see Section 4.2) that the IETF could potentially reserve in the future for entirely different purposes.

6. Checksum Neutrality

Use of 64:ff9b:1::/48 does not in itself guarantee checksum neutrality, as many of the IPv4/IPv6 translation algorithms it can be used with are fundamentally incompatible with checksum-neutral address translations.

Section 4.1 of [RFC6052] contains further discussion about IPv4/IPv6 translation and checksum neutrality.

The Stateless IP/ICMP Translation algorithm [RFC7915] is one well-known algorithm that can operate in a checksum-neutral manner, when using the [RFC6052] algorithms for all of its address translations. However, in order to attain checksum neutrality, it is imperative that the translation prefix be chosen carefully. Specifically, in order for a 96-bit [RFC6052] prefix to be checksum neutral, all the six 16-bit words in the prefix must add up to a multiple of 0xffff.

The following non-exhaustive list contains examples of translation prefixes that are checksum neutral when used with the [RFC7915] and [RFC6052] algorithms:

- o 64:ff9b:1:fffe::/96
- o 64:ff9b:1:fffd:1::/96
- o 64:ff9b:1:fffc:2::/96
- o 64:ff9b:1:abcd:0:5431::/96

7. IANA Considerations

The IANA has added the following entry to the "IANA IPv6 Special-Purpose Address Registry":

Attribute	Value
Address Block	64:ff9b:1::/48
Name	IPv4-IPv6 Translat.
RFC	RFC 8215
Allocation Date	2017-06
Termination Date	N/A
Source	True
Destination	True
Forwardable	True
Globally Reachable	False
Reserved-by-Protocol	False

The IANA has also added the following footnote to the 0000::/8 entry of the "Internet Protocol Version 6 Address Space" registry:

64:ff9b:1::/48 reserved for Local-Use IPv4/IPv6 Translation [RFC8215].

8. Security Considerations

The reservation of 64:ff9b:1::/48 is not known to cause any new security considerations beyond those documented in Section 5 of [RFC6052].

9. References

9.1. Normative References

- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<https://www.rfc-editor.org/info/rfc6052>>.

9.2. Informative References

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <<https://www.rfc-editor.org/info/rfc6146>>.
- [RFC7915] Bao, C., Li, X., Baker, F., Anderson, T., and F. Gont, "IP/ICMP Translation Algorithm", RFC 7915, DOI 10.17487/RFC7915, June 2016, <<https://www.rfc-editor.org/info/rfc7915>>.

Acknowledgements

The author would like to thank Fred Baker, Mohamed Boucadair, Brian E. Carpenter, Pier Carlo Chiodi, Joe Clarke, David Farmer, Suresh Krishnan, Warren Kumari, Holger Metschulat, Federico Santandrea, and David Schinazi for contributing to the creation of this document.

Author's Address

Tore Anderson
Redpill Linpro
Vitaminveien 1A
0485 Oslo
Norway

Phone: +47 959 31 212
Email: tore@redpill-linpro.com
URI: <http://www.redpill-linpro.com>

