

Internet Engineering Task Force (IETF)
Request for Comments: 7526
BCP: 196
Obsoletes: 3068, 6732
Category: Best Current Practice
ISSN: 2070-1721

O. Troan
Cisco
B. Carpenter, Ed.
Univ. of Auckland
May 2015

Deprecating the Anycast Prefix for 6to4 Relay Routers

Abstract

Experience with the 6to4 transition mechanism defined in RFC 3056 ("Connection of IPv6 Domains via IPv4 Clouds") has shown that the mechanism is unsuitable for widespread deployment and use in the Internet when used in its anycast mode. Therefore, this document requests that RFC 3068 ("An Anycast Prefix for 6to4 Relay Routers") and RFC 6732 ("6to4 Provider Managed Tunnels") be made obsolete and moved to Historic status. It recommends that future products should not support 6to4 anycast and that existing deployments should be reviewed. This complements the guidelines in RFC 6343.

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/rfc7526>.

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	3
1.1. Related Work	3
2. Conventions	4
3. 6to4 Operational Problems	4
4. Deprecation	5
5. Implementation Recommendations	5
6. Operational Recommendations	6
7. IANA Considerations	7
8. Security Considerations	7
9. References	7
9.1. Normative References	7
9.2. Informative References	8
Acknowledgements	9
Authors' Addresses	10

1. Introduction

The original form of the 6to4 transition mechanism [RFC3056] relies on unicast addressing. However, its extension specified in "An Anycast Prefix for 6to4 Relay Routers" [RFC3068] has been shown to have severe practical problems when used in the Internet. This document requests that RFCs 3068 and 6732 be moved to Historic status, as defined in Section 4.2.4 of [RFC2026]. It complements the deployment guidelines in [RFC6343].

6to4 was designed to help transition the Internet from IPv4 to IPv6. It has been a good mechanism for experimenting with IPv6, but because of the high failure rates seen with anycast 6to4 [HUSTON], end users may end up disabling IPv6 on hosts; this has resulted in some content providers being reluctant to make content available over IPv6 in the past.

[RFC6343] analyzes the known operational issues in detail and describes a set of suggestions to improve 6to4 reliability, given the widespread presence of hosts and customer premises equipment that support it. The advice to disable 6to4 by default has been widely adopted in recent operating systems, and the failure modes have been widely hidden from users by many browsers adopting the "Happy Eyeballs" approach [RFC6555].

Nevertheless, a measurable amount of 6to4 traffic is still observed by IPv6 content providers. The remaining successful users of anycast 6to4 are likely to be on hosts using the obsolete policy table [RFC3484] (which prefers 6to4 above IPv4) and running without Happy Eyeballs. Furthermore, they must have a route to an operational anycast relay and they must be accessing an IPv6 host that has a route to an operational return relay.

However, experience shows that operational failures caused by anycast 6to4 have continued despite the advice in RFC 6343 being available.

1.1. Related Work

"IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification" [RFC5969] explicitly builds on the 6to4 mechanism, using a service provider prefix instead of 2002::/16. However, the deployment model is based on service provider support such that 6rd avoids the problems observed with anycast 6to4.

The framework for "6to4 Provider Managed Tunnels" [RFC6732] is intended to help a service provider manage 6to4 anycast tunnels. This framework only exists because of the problems observed with anycast 6to4.

2. Conventions

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

In this document, the word "deprecate" and its derivatives are used only in their generic sense of "criticize or express disapproval" and do not have any specific normative meaning. A deprecated function might exist in the Internet for many years to allow backwards compatibility.

3. 6to4 Operational Problems

6to4 is a mechanism designed to allow isolated IPv6 islands to reach each other using IPv6-over-IPv4 automatic tunneling. To reach the native IPv6 Internet, the mechanism uses relay routers in both the forward and reverse direction. The mechanism is supported in many IPv6 implementations. With the increased deployment of IPv6, the mechanism has been shown to have a number of shortcomings.

In the forward direction, a 6to4 node will send IPv4-encapsulated IPv6 traffic to a 6to4 relay that is connected to both the 6to4 cloud and native IPv6. In the reverse direction, a 2002::/16 route is injected into the native IPv6 routing domain to attract traffic from native IPv6 nodes to a 6to4 relay router. It is expected that traffic will use different relays in the forward and reverse direction.

One model of 6to4 deployment, described in Section 5.2 of RFC 3056, suggests that a 6to4 router should have a set of managed connections (via BGP connections) to a set of 6to4 relay routers. While this makes the forward path more controlled, it does not guarantee a functional reverse path. In any case, this model has the same operational burden as manually configured tunnels and has seen no deployment in the public Internet.

RFC 3068 adds an extension that allows the use of a well-known IPv4 anycast address to reach the nearest 6to4 relay in the forward direction. However, this anycast mechanism has a number of operational issues and problems, which are described in detail in Section 3 of [RFC6343]. This document is intended to deprecate the anycast mechanism.

Peer-to-peer usage of the 6to4 mechanism exists in the Internet, likely unknown to many operators. This usage is harmless to third parties and is not dependent on the anycast 6to4 mechanism that this document deprecates.

4. Deprecation

This document formally deprecates the anycast 6to4 transition mechanism defined in [RFC3068] and the associated anycast IPv4 address 192.88.99.1. It is no longer considered to be a useful service of last resort.

The prefix 192.88.99.0/24 MUST NOT be reassigned for other use except by a future IETF Standards Action.

The basic unicast 6to4 mechanism defined in [RFC3056] and the associated 6to4 IPv6 prefix 2002::/16 are not deprecated. The default address selection rules specified in [RFC6724] are not modified.

In the absence of 6to4 anycast, "6to4 Provider Managed Tunnels" [RFC6732] will no longer be necessary, so they are also deprecated by this document.

Incidental references to 6to4 should be reviewed and possibly removed from other IETF documents if and when they are updated. These documents include RFC 3162, RFC 3178, RFC 3790, RFC 4191, RFC 4213, RFC 4389, RFC 4779, RFC 4852, RFC 4891, RFC 4903, RFC 5157, RFC 5245, RFC 5375, RFC 5971, RFC 6071, and RFC 6890.

5. Implementation Recommendations

It is NOT RECOMMENDED to include the anycast 6to4 transition mechanism in new implementations. If included in any implementations, the anycast 6to4 mechanism MUST be disabled by default.

In host implementations, unicast 6to4 MUST also be disabled by default. All hosts using 6to4 MUST support the IPv6-address-selection policy described in [RFC6724].

In router implementations, 6to4 MUST be disabled by default. In particular, enabling IPv6 forwarding on a device MUST NOT automatically enable 6to4.

6. Operational Recommendations

This document does not imply a recommendation for the generalized filtering of traffic or routes for 6to4 or even anycast 6to4. It simply recommends against further deployment of the anycast 6to4 mechanism, calls for current 6to4 deployments to evaluate the efficacy of continued use of the anycast 6to4 mechanism, and makes recommendations intended to prevent any use of 6to4 from hampering broader deployment and use of native IPv6 on the Internet as a whole.

Networks SHOULD NOT filter out packets whose source address is 192.88.99.1, because this is normal 6to4 traffic from a 6to4 return relay somewhere in the Internet. This includes ensuring that traffic from a local 6to4 return relay with a source address of 192.88.99.1 is allowed through anti-spoofing filters (such as those described in [RFC2827] and [RFC3704]) or through Unicast Reverse Path Forwarding (uRPF) checks [RFC5635].

The guidelines in Section 4 of [RFC6343] remain valid for those who choose to continue operating anycast 6to4 despite its deprecation.

Current operators of an anycast 6to4 relay with the IPv4 address 192.88.99.1 SHOULD review the information in [RFC6343] and the present document, and then consider carefully whether the anycast relay can be discontinued as traffic diminishes. Internet service providers that do not operate an anycast relay but do provide their customers with a route to 192.88.99.1 SHOULD verify that it does in fact lead to an operational anycast relay, as discussed in Section 4.2.1 of [RFC6343]. Furthermore, Internet service providers and other network providers MUST NOT originate a route to 192.88.99.1, unless they actively operate and monitor an anycast 6to4 relay service as detailed in Section 4.2.1 of [RFC6343].

Operators of a 6to4 return relay responding to the IPv6 prefix 2002::/16 SHOULD review the information in [RFC6343] and the present document, and then consider carefully whether the return relay can be discontinued as traffic diminishes. To avoid confusion, note that nothing in the design of 6to4 assumes or requires that return packets are handled by the same relay as outbound packets. As discussed in Section 4.5 of RFC 6343, content providers might choose to continue operating a return relay for the benefit of their own residual 6to4 clients. Internet service providers SHOULD announce the IPv6 prefix 2002::/16 to their own customers if and only if it leads to a correctly operating return relay as described in RFC 6343. IPv6-only service providers, including those operating a NAT64 service [RFC6146], are advised that their own customers need a route to such a relay in case a residual 6to4 user served by a different service provider attempts to communicate with them.

Operators of "6to4 Provider Managed Tunnels" [RFC6732] SHOULD carefully consider when this service can be discontinued as traffic diminishes.

7. IANA Considerations

The document creating the "IANA IPv4 Special-Purpose Address Registry" [RFC6890] included the 6to4 Relay Anycast prefix (192.88.99.0/24) as Table 10. Per this document, IANA has marked the 192.88.99.0/24 prefix (originally defined by [RFC3068]) as "Deprecated (6to4 Relay Anycast)" and added a reference to this RFC. The Boolean values for the address block 192.88.99.0/24 have been removed. Redlegation of this prefix for any use requires justification via an IETF Standards Action [RFC5226].

8. Security Considerations

There are no new security considerations pertaining to this document. General security issues with tunnels are listed in [RFC6169] and more specifically to 6to4 in [RFC3964] and [RFC6324].

9. References

9.1. Normative References

- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, DOI 10.17487/RFC2026, October 1996, <<http://www.rfc-editor.org/info/rfc2026>>.
- [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>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, DOI 10.17487/RFC3056, February 2001, <<http://www.rfc-editor.org/info/rfc3056>>.
- [RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, DOI 10.17487/RFC3068, June 2001, <<http://www.rfc-editor.org/info/rfc3068>>.

- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<http://www.rfc-editor.org/info/rfc3704>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [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, <<http://www.rfc-editor.org/info/rfc6146>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<http://www.rfc-editor.org/info/rfc6724>>.
- [RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <<http://www.rfc-editor.org/info/rfc6890>>.

9.2. Informative References

- [HUSTON] Huston, G., "Flailing IPv6", The ISP Column, December 2010, <<http://www.potaroo.net/ispcol/2010-12/6to4fail.html>>.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, DOI 10.17487/RFC3484, February 2003, <<http://www.rfc-editor.org/info/rfc3484>>.
- [RFC3964] Savola, P. and C. Patel, "Security Considerations for 6to4", RFC 3964, DOI 10.17487/RFC3964, December 2004, <<http://www.rfc-editor.org/info/rfc3964>>.
- [RFC5635] Kumari, W. and D. McPherson, "Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF)", RFC 5635, DOI 10.17487/RFC5635, August 2009, <<http://www.rfc-editor.org/info/rfc5635>>.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, DOI 10.17487/RFC5969, August 2010, <<http://www.rfc-editor.org/info/rfc5969>>.

- [RFC6169] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns with IP Tunneling", RFC 6169, DOI 10.17487/RFC6169, April 2011, <<http://www.rfc-editor.org/info/rfc6169>>.
- [RFC6324] Nakibly, G. and F. Templin, "Routing Loop Attack Using IPv6 Automatic Tunnels: Problem Statement and Proposed Mitigations", RFC 6324, DOI 10.17487/RFC6324, August 2011, <<http://www.rfc-editor.org/info/rfc6324>>.
- [RFC6343] Carpenter, B., "Advisory Guidelines for 6to4 Deployment", RFC 6343, DOI 10.17487/RFC6343, August 2011, <<http://www.rfc-editor.org/info/rfc6343>>.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, DOI 10.17487/RFC6555, April 2012, <<http://www.rfc-editor.org/info/rfc6555>>.
- [RFC6732] Kuarsingh, V., Ed., Lee, Y., and O. Vautrin, "6to4 Provider Managed Tunnels", RFC 6732, DOI 10.17487/RFC6732, September 2012, <<http://www.rfc-editor.org/info/rfc6732>>.

Acknowledgements

The authors would like to acknowledge Tore Anderson, Mark Andrews, Dmitry Anipko, Jack Bates, Cameron Byrne, Ben Campbell, Lorenzo Colitti, Gert Doering, Nick Hilliard, Philip Homburg, Ray Hunter, Joel Jaeggli, Victor Kuarsingh, Kurt Erik Lindqvist, Jason Livingood, Jeroen Massar, Keith Moore, Tom Petch, Daniel Roesen, Mark Townsley, and James Woodyatt for their contributions and discussions on this topic.

Special thanks go to Fred Baker, David Farmer, Wes George, and Geoff Huston for their significant contributions.

Many thanks to Gunter Van de Velde for documenting the harm caused by non-managed tunnels and stimulating the creation of this document.

Authors' Addresses

Ole Troan
Cisco
Oslo
Norway

EMail: ot@cisco.com

Brian Carpenter (editor)
Department of Computer Science
University of Auckland
PB 92019
Auckland 1142
New Zealand

EMail: brian.e.carpenter@gmail.com

