

Internet Engineering Task Force (IETF)
Request for Comments: 6791
Updates: 6145
Category: Standards Track
ISSN: 2070-1721

X. Li
C. Bao
CERNET Center/Tsinghua University
D. Wing
R. Vaithianathan
Cisco
G. Huston
APNIC
November 2012

Stateless Source Address Mapping for ICMPv6 Packets

Abstract

A stateless IPv4/IPv6 translator may receive ICMPv6 packets containing non-IPv4-translatable addresses as the source. These packets should be passed across the translator as ICMP packets directed to the IPv4 destination. This document presents recommendations for source address translation in ICMPv6 headers to handle such cases.

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 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/rfc6791>.

Copyright Notice

Copyright (c) 2012 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. Notational Conventions	3
3. Problem Statement and Considerations	3
3.1. Considerations	3
3.2. Recommendations	3
4. ICMP Extension	4
5. Stateless Address Mapping Algorithm	4
6. Security Considerations	4
7. Acknowledgments	4
8. References	5
8.1. Normative References	5
8.2. Informative References	5

1. Introduction

Section 5.3 of "IP/ICMP Translation Algorithm" [RFC6145] states that "the IPv6 addresses in the IPv6 header may not be IPv4-translatable addresses and there will be no corresponding IPv4 addresses representing this IPv6 address. In this case, the translator can do stateful translation. A mechanism by which the translator can instead do stateless translation of this address is left for future work." This document, "Stateless Source Address Mapping for ICMPv6 Packets", provides recommendations for this case.

For the purposes of this document, the term "IPv4-translatable IPv6 address" is as defined in Section 2.2 of [RFC6052].

2. Notational Conventions

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

3. Problem Statement and Considerations

When a stateless IPv4/IPv6 translator receives an ICMPv6 message [RFC4443] (for example, "Packet Too Big") sourced from a non-IPv4-translatable IPv6 address and bound for an IPv4-translatable IPv6 address, the translator needs to pick a source address with which to generate an ICMP message. For the reasons discussed below, this choice is problematic.

3.1. Considerations

The source address used SHOULD NOT cause the ICMP packet to be discarded. It SHOULD NOT be drawn from [RFC1918] or [RFC6598] address space, because that address space is likely to be subject to unicast Reverse Path Forwarding (uRPF) [RFC3704] filtering.

IPv4/IPv6 translation is intended for use in contexts where IPv4 addresses may not be readily available. Therefore, it is not considered appropriate to assign IPv4-translatable IPv6 addresses for all internal points in the IPv6 network that may originate ICMPv6 messages.

Another consideration for source selection is that it should be possible for the IPv4 recipients of the ICMP message to be able to distinguish between different IPv6 network origination of ICMPv6 messages (for example, to support a traceroute diagnostic utility that provides some limited network-level visibility across the IPv4/IPv6 translator). This consideration implies that an IPv4/IPv6 translator needs to have a pool of IPv4 addresses for mapping the source address of ICMPv6 packets generated from different origins, or to include the IPv6 source address information for mapping the source address by others means. Currently, the TRACEROUTE and MTR [MTR] are the only consumers of translated ICMPv6 messages that care about the ICMPv6 source address.

3.2. Recommendations

The recommended approach to source selection is to use a single (or small pool of) public IPv4 address as the source address of the translated ICMP message and leverage the ICMP extension [RFC5837] to include the IPv6 address as an Interface IP Address Sub-Object.

4. ICMP Extension

In the case of either a single public IPv4 address (the IPv4 interface address or loopback address of the translator) or a pool of public IPv4 addresses, the translator SHOULD implement the ICMP extension defined by [RFC5837]. The ICMP message SHOULD include the Interface IP Address Sub-Object and specify the source IPv6 addresses of the original ICMPv6. When an enhanced traceroute application is used, it can derive the real IPv6 source addresses that generated the ICMPv6 messages. Therefore, it would be able improve on visibility towards the origin rather than simply blackholing at or beyond the translator. In the future, a new ICMP extension whose presence indicates that the packet has been translated and that the source address belongs to the translator, not the originating node, can also be considered.

5. Stateless Address Mapping Algorithm

If a pool of public IPv4 addresses is configured on the translator, it is RECOMMENDED to randomly select the IPv4 source address from the pool. Random selection reduces the probability that two ICMP messages elicited by the same TRACEROUTE might specify the same source address and, therefore, erroneously present the appearance of a routing loop.

[RFC5837] extensions and an enhanced traceroute application, if used, will reveal the IPv6 source addresses that generated the original ICMPv6 messages.

6. Security Considerations

This document recommends the generation of IPv4 ICMP messages from IPv6 ICMP messages. These messages would otherwise have been discarded. New considerations are not expected to result from this change. As with a number of ICMP messages, a spoofed source address may result in replies arriving at hosts that did not expect them using the facility of the translator.

7. Acknowledgments

The authors would like to acknowledge the following contributors of this document: Kevin Yin, Chris Metz, Neeraj Gupta, and Joel Jaeggli. The authors would also like to thank Ronald Bonica, Ray Hunter, George Wes, Yu Guanghui, Sowmini Varadhan, David Farmer, Fred Baker, Leo Vegoda, Joel Jaeggli, Henrik Levkowetz, Randy Bush, and Warren Kumari for their comments and suggestions.

8. References

8.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [RFC5837] Atlas, A., Ed., Bonica, R., Ed., Pignataro, C., Ed., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, April 2010.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, October 2010.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011.
- [RFC6598] Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space", BCP 153, RFC 6598, April 2012.

8.2. Informative References

- [MTR] "BitWizard B.V. - The Linux Experts",
<<http://www.bitwizard.nl/mtr/>>.

Authors' Addresses

Xing Li
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
China

Phone: +86 10-62785983
EMail: xing@cernet.edu.cn

Congxiao Bao
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
China

Phone: +86 10-62785983
EMail: congxiao@cernet.edu.cn

Dan Wing
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
United States

EMail: dwing@cisco.com

Ramji Vaithianathan
Cisco Systems, Inc.
A 5-2, BGL 12-4, SEZ Unit,
Cessna Business Park, Varthur Hobli
Sarjapur Outer Ring Road
Bangalore Karnataka 560 103
India

Phone: +91 80 4426 0895
EMail: rvaithia@cisco.com

Geoff Huston
APNIC

EMail: gih@apnic.net

