Network Working Group
Internet-Draft
Intended status: BCP
Expires: January 12, 2013

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

July 11, 2012

Stateless Source Address Mapping for ICMPv6 Packets
draft-ietf-v6ops-ivi-icmp-address-03

Abstract

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

Status of this Memo

This Internet-Draft is submitted to IETF 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 http://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 January 12, 2013.

Table of Contents

1. Introduction

The IP/ICMP translation document of IPv4/IPv6 translation [RFC6145] states that "the IPv6 addresses in the ICMPv6 header may not be IPv4-translatable addresses and there will be no corresponding IPv4 addresses represented of this IPv6 address. In this case, the translator can do stateful translation. A mechanism by which the translator can instead do stateless translation is left for future work." This document presents recommendations for this case.

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 an non-IPv4-translatable IPv6 address, directed to an IPv4-translatable IPv6 address, it needs to generate an ICMP message. For the reasons discussed below, choosing the source IPv4 address of this ICMP message is problematic.

The address used should not cause the ICMP packet to be a candidate for discarding, particularly in the contest of uRPF filters [RFC3704]. This consideration precludes the use of private IPv4 address space [RFC1918] in this context.

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

It is also an objective that it is possible for the IPv4 recipient of the ICMP message be able to distinguish between different IPv6 ICMPv6 originations (for example, to support a traceroute diagnostic utility that provides some limited network level visibility across the IPv4/IPv6 translator). This 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 originations, or to include the IPv6 source address information for mapping the source address by others means.

These considerations leads to the recommendation of using the a single (or small pool) of public IPv4 address as the source address of translated ICMP and using ICMP extension [RFC5837] to include IPv6 address as Interface IP Address Sub-Object.

4. ICMP Extension

No matter a single public IPv4 address (in the IPv4 interface or a loopback address of the translator) or a pool of public IPv4 addresses are used, the translator SHOULD implement ICMP extension defined by [RFC5837]. The resulting ICMP extension SHOULD include the Interface IP Address Sub-Object that specify the source IPv6 addresses in the original ICMPv6. When an enhanced traceroute application is used, it can get the real IPv6 source addresses which generate the ICMPv6 messages. Therefore, it is able to traceback to their origins and take filtering/rate-limiting actions if necessary.

5. Stateless Address Mapping Algorithm

If a pool of public IPv4 addresses is configured in the translator, it is RECOMMENDED to randomly pick up the IPv4 public address in the pool. This can somehow avoid the appearance of a routing loop to tools such as traceroute. An enhanced traceroute application is still RECOMMENDED in order to obtain the real IPv6 source addresses which generate the ICMPv6 messages.

6. Security Considerations

This document does not introduce any new security considerations.

7. IANA Considerations

None.

8. Acknowledgments

The authors would like to acknowledge the following contributors of this document: Kevin Yin, Chris Metz and Neeraj Gupta. 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, Henrik Levkowetz, Randy Bush and Warren Kumari for their comments and suggestions.

9. Normative References

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., 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, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

[RFC5837] Atlas, A., Bonica, R., Pignataro, C., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, April 2010.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011.

Authors' Addresses

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

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

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

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

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

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