Network Working Group
Internet-Draft
Intended status: BCP
Expires: August 28, 2012

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

February 25, 2012

Stateless Source Address Mapping for ICMPv6 Packets
draft-ietf-v6ops-ivi-icmp-address-01

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 the IPv4-translatable destination. This document discusses the considerations and presents a stateless address mapping algorithm 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 August 28, 2012.

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 defines such a stateless translation mechanism.

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 (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 a IPv4/IPv6 translator needs to have a pool of IPv4 addresses to be used for mapping the source address of ICMPv6 packets generated from different originations.

These addresses are for use in the source address of ICMP packets, and therefore are not intended to be used as a destination address for any packet. It is therefore possible to use a common address

pool for the IPv4/IPv6 translation protocol, and, considering an objective of constraining the use of these IPv4 addresses in this application, it is feasible to use a common address pool for mapping the source addresses of non-translatable ICMPv6 packets as a part of the protocol specification.

These considerations leads to the recommendation of drawing an IPv4 /24 prefix from the IANA Special Purpose Address Registry as a "Well-Known Prefix" for use by IPv4/IPv6 translators for the purpose of mapping otherwise untranslatable IPv6 source addresses of ICMPv6 messages to IPv4 ICMP messages.

The ICMP extension defined by [RFC5837] provides a mechanism to process the ICMPv4 messages that contain IP Address Sub-Objects that specify IPv6 addresses. However, an enhanced traceroute application must be used, which has not yet been widely available. In this document, a combined approach is proposed, i.e. non IPv4-translatable address is mapped to IANA-assigned /24, and the resulting ICMP is extended according to [RFC5837]. Therefore, ordinary ICMP processing tools (traceroute) can be utilized in normal cases and when DDoS happens, enhanced ICMP process tools can be utilized to identify the real source.

4. Routing Considerations

Addresses from the assigned address prefix are intended to be used as source addresses and not as destination addresses in the context of the public network. As packets passing through the public network need to pass through conventional packet filters, including uRPF filters [RFC3704], this implies that the assigned address may be used in routing advertisements. Such routing advertisements are non-exclusive and should be accepted from any originating AS in an anycast fashion.

5. Stateless Address Mapping Algorithm

When an IPv4 /24 prefix is allocated to represent the source address of ICMP, the translator MUST copy the "Hop Count" in the IPv6 header of the ICMPv6 to the Last Octet. When routers emit ICMPv6 packets with the same hop count, as the ICMPv6 packet is routed through the network its hop count is decreased. However, if the routers emit ICMPv6 packets with different hop counts, it may give the appearance of a routing loop to tools such as traceroute. That minor side-effect in that particular case cannot be avoided while still being stateless.

6. ICMP Extension

When translator is configured to use the IANA-assigned /24 to map non IPv4-translatable address, the translator MUST implement ICMP extension defined by [RFC5837]. The resulting ICMP extension MUST include the IP address Sub-Objects that specify the source IPv6 addresses in the original ICMPv6. Therefore, an enhanced traceroute application can get the real IPv6 source addresses which generate the ICMPv6 messages, be able to traceback to their origins and take filtering/rate-limiting actions if necessary.

7. Security Considerations

The use of an address for source addresses in ICMP packets is considered "safe" in so far as ICMP packets are not intended to generate responses directed to the source address.

However it is possible to use this address as a means of gaining anonymity when launching a denial of service attacks by using this address as the source address for other forms of malicious traffic. Packet firewall filters should be configured to treat addresses in the IANA-assigned /24 network as martian addresses by discarding all non-ICMP packets that use the IANA-assigned /24 network as a source address, and all packets that use the IANA-assigned /24 network as a destination address.

7.1. Filtering Recommendations

- SHOULD allow ICMP type 3 - Destination Unreachable (inc PTB).
- SHOULD allow ICMP type 11 - Time Exceeded.
- MAY allow ICMP type 12 - Parameter Problem.
- SHOULD NOT allow any of the various ICMP request messages.

7.2. Rate-limiting Recommendations

The rate limiting of traffic from the prefix SHOULD also be enabled as additional countermeasure against abuse of this prefix. The methods presented in [RFC4443] [RFC5597] [RFC6192] [RFC6398] [RFC6450] can be used.

7.3. RFC5837 Recommendations

Advanced filtering and rate-limiting techniques which can process the ICMP extension defined in [RFC5837] MAY also be used to control the

source of the ICMP.

8. IANA Considerations

IANA is requested to make a permanent assignment of a /24 from the IPv4 Special Purpose Address Registry [RFC5736]. The assigned address is to be used in the context of generating an IPv4 source address for mapped ICMPv6 packets being passed through a stateless IPv4/IPv6 translator. The assignment is under the category of a specialized use of a designated address block in an anycast context associated with an Internet Standards Track protocol.

The IANA IPv4 Special Purpose Address Registry records are:

- Prefix 192.70.192.0/24
- Description: To be used in the context of generating an IPv4 source address for mapped ICMPv6 packets being passed through a stateless IPv4/IPv6 translator.
- Begin: 2011-06-01
- End: Never
- Purpose: Stateless ICMPv6/ICMP translation
- Contact: See RFC
- Scope: Addresses from the assigned address prefix are intended to be used as source addresses and not as destination addresses in the context of the public network.
- RFC: This draft.

9. Acknowledgments

The authors would like to acknowledge the following contributors of this document: Kevin Yin, Chris Metz and Neeraj Gupta.

10. References

10.1. 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.

[RFC5597] Denis-Courmont, R., "Network Address Translation (NAT) Behavioral Requirements for the Datagram Congestion Control Protocol", BCP 150, RFC 5597, September 2009.

[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.

[RFC6398] Le Faucheur, F., "IP Router Alert Considerations and Usage", BCP 168, RFC 6398, October 2011.

[RFC6450] Venaas, S., "Multicast Ping Protocol", RFC 6450, December 2011.

10.2. Informative References

[RFC5736] Huston, G., Cotton, M., and L. Vegoda, "IANA IPv4 Special Purpose Address Registry", RFC 5736, January 2010.

[RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, March 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