v6ops WG
Internet-Draft
Intended status: Standards Track
Expires: August 19, 2010

S. Gundavelli
M. Townsley
O. Troan
W. Dec
Cisco

February 15, 2010

Unicast Transmission of IPv6 Multicast Messages on Link-layer
draft-gundavelli-v6ops-l2-unicast-00.txt

Abstract

When transmitting an IPv6 packet to a multicast group address, the destination address in the link-layer header is typically set to the corresponding mapped address of that multicast group address. However, it is not mandatory that the destination address in the link-layer header is always a mapped multicast equivalent of its IP destination address. There are various deployment scenarios where there a need to transmit an IPv6 multicast message as an unicast message on the link-layer. Unfortunately, the IPv6 specifications do not clearly state this. This document explicitly clarifies this aspect and makes such packet construct and transmission legal and valid.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on August 19, 2010.

Table of Contents

1. Introduction

When transmitting an IPv6 packet to a multicast group address, the destination address in the link-layer header is typically set to the corresponding mapped address of that multicast group address. However, it is not mandatory that the destination address in the link-layer header is always a mapped multicast equivalent of its IP destination address. There are various deployment scenarios where there a need to transmit an IPv6 multicast message as an unicast message on the link-layer. Such a message construct is valid and is used in various protocols, such as in ISATAP [RFC5214] for sending a unicast Router Advertisement message on ISTAP interfaces. Evidentially, some of the IPv6 specifications, such as [RFC4861], or [RFC2464] do not make any assumption on such tight relationship and it does require the receiving IPv6 node to explicitly apply any such checks across protocol layers. However, it is ambiguous from the protocol specification perspective, on the legality of such transmission and any discussions on this subject always resulted in differing opinions. Therefore, it is the intent of this document to make the specification clear on this aspect.

There are many deployment scenarios where there is a need to transmit IPv6 multicast Neighbor Discovery packets as unicast packets on the link-layer. For example, an 802.11 wireless access point may be hosting multiple IPv6 subnets and it would need the ability to selectively advertise hosted IPv6 prefixes on a per-node basis. Such segregation can only be possible by ensuring the Router Advertisements received by any IPv6 node includes only those prefixes that are associated with their respective layer-3 subnet. This essentially requires the ability to transmit multicast messages as unicast messages on the link-layer.

The function of the link-layer is purely for transmitting the frame to a peer or to a set of peers on a given media. It is inconsequential for the network layer protocols to go across the layers and check the semantics of message delivery in the link-layer header. Any such check is a violation of the principles of protocol layering and does not serve any purpose. Unfortunately, [RFC4861] or [RFC2464] does not explicitly state this. However, we have verified this on many open source and commercial IPv6 implementations on the behavior of the existing IPv6 stacks, firewalls and we could not find any implementation that drops IPv6 packets sent to a multicast IPv6 destination address, but with a unicast destination address in the link-layer header. Case and Point:

- Cisco IOS Operating System

- Linux Operating System with 2.6 Kernel
- BSD Variants based on IPv6 KAME implementation
- Microsoft Windows Vista

As a result of this analysis, it appears to be quite safe to explicitly state that such message construct is valid, so future implementations do not drop packets based on these checks. This document updates [RFC4861] for allowing this change.

2. Requirements Language

In this document, the key words "MAY", "MUST, "MUST NOT", "OPTIONAL", "RECOMMENDED", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [RFC2119].

3. Sending and Receiving IPv6 Multicast Packets

The following considerations MUST be applied by all IPv6 nodes when sending and receiving IPv6 multicast messages.

- An IPv6 node MAY choose to unicast an IPv6 multicast message on the link-layer. In this case, the destination address in the IPv6 header will be a multicast group address, but the destination address in the link-layer header will be an unicast address. It is up to to the system architecture as when to transmit a IPv6 multicast message as an link-layer unicast message, as long as there is no real impact to the multicast communication.

- An IPv6 receiver node SHOULD NOT drop a received IPv6 multicast message containing a multicast destination address in the IPv6 header, but with a unicast destination address in the link-layer header, withstanding all other validity considerations as specified in the relevant IPv6 standards specifications.

4. IANA Considerations

This specification does not require any IANA actions.

5. Security Considerations

This document is about a clarification to the construction and processing rules of IPv6 multicast messages. This clarification makes it valid for an IPv6 receiver node to consider a received IPv6 multicast message with a multicast destination address in the IPv6 header, but containing an unicast destination address in the link-layer header, to be valid withstanding all other validity considerations specified in the IPv6 standards specifications. This change follows the principles of protocol layer design more tightly and does not introduce any security vulnerabilities.

Network firewalls and Deep Packet inspection tools that perform validity checks on link-layer and IP layer headers may have to modified to allow such packet transmission. However, the authors of this document could not find a single such implementation that rejects packets based on this check.

6. Acknowledgements

The authors would like to acknowledge Hemant Singh, Eric Levy, Pascal Thubert and Eric Voit for all the discussions on this topic.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

7.2. Informative References

[RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.

[RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.

Authors' Addresses

Sri Gundavelli
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

Email: sgundave@cisco.com

Mark Townsley
Cisco

Email: townsley@cisco.com

Ole Troan
Cisco
Skoyen Atrium, Drammensveien 145A
Oslo,   N-0277
Norway

Email: otroan@cisco.com

Wojciech Dec
Cisco
Haarlerbergweg 13-19
Amsterdam, Noord-Holland 1101 CH
Netherlands

Email: wdec@cisco.com