

Workgroup: Network Working Group
Internet-Draft:
draft-karstens-pim-ipv6-zeroconf-assignment-00
Published: 23 October 2022
Intended Status: Informational
Expires: 26 April 2023

A	N. Karstens	D. Farinacci	M. McBride
	uGarmin International	lispers.net	Futurewei
	t		
	h		
	o		
	r		
	s		
	:		

Zero-Configuration Assignment of IPv6 Multicast Addresses

Abstract

Marine networks contain a combination of sensors, controls, and displays. The latest marine industry standards require IPv6. The most optimal way to distribute sensor data to all displays on the network is multicast. However, use of traditional switches can be problematic (overwhelm links) when both high-bandwidth and low-bandwidth devices are installed. To solve this problem, the network requires switches with multicast snooping. However, source-specific multicast (SSM) is not supported on marine switches so the destination address is the only way to differentiate multicast streams. This limitation creates several challenges including with the pre-allocation of addresses. The solution, described in this draft, provides a decentralized, zero-configuration method for dynamically assigning multicast addresses through defining an extension to the multicast portion of the IPv6 addressing architecture along with a new IANA registry.

Status of This Memo

This Internet-Draft is submitted 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 <https://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 26 April 2023.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
 - [1.1. Requirements Language](#)
- [2. Technical Background](#)
- [3. Design Goals](#)
- [4. Method](#)
- [5. IANA Considerations](#)
- [6. Security Considerations](#)
- [7. Acknowledgement](#)
- [8. References](#)
 - [8.1. Normative References](#)
 - [8.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

Marine networks contain a combination of sensors, controls, and displays. Installations vary widely depending on the design and intended purpose of the boat and the amount of redundancy required. Sensors on these networks can be a mix of low-cost, low-bandwidth devices, like temperature or fluid sensors, and high-bandwidth devices, like radar, sonar, and video cameras. In most cases these networks use a single subnet. The latest marine industry standards require IPv6.

The most optimal way to distribute sensor data to all displays on the network is multicast. However, use of traditional switches can be problematic when both high-bandwidth and low-bandwidth devices are installed. Low-bandwidth devices are commonly designed with a low-speed link to reduce cost, and the multicast stream from the high-bandwidth device can overwhelm this link. To solve this problem, the network requires switches with multicast snooping [[RFC4541](#)], which directs multicast streams only to the ports leading to devices that request the data.

Switch parts at the required price point do not support source-specific multicast, so the destination address is the only way to differentiate multicast streams. This presents several challenges. First, defining an industry standard set of pre-allocated addresses is not practical due to the wide variety of network designs. Users in the marine industry would not find static assignment to be acceptable. MADCAP [[RFC2730](#)] could be used to dynamically assign addresses, but its reliance on a dedicated server results in a single point of failure for the system, which is not acceptable in the marine environment.

The solution, proposed in this draft, is a decentralized, zero-configuration method for dynamically assigning multicast addresses. This document defines an extension to the multicast portion of the IPv6 addressing architecture [[RFC4291](#)]. The current architecture does not account for potential address collisions when IPv6 multicast packets are transmitted on the data link layer. This extension defines a collision detection mechanism that utilizes Multicast DNS [[RFC6762](#)] to distribute a database of dynamically assigned multicast Ethernet addresses.

It also proposes a new IANA registry based on amendments to [Section 4.3](#) of [[RFC3307](#)]. This will allow for different methods of dynamically allocating IPv6 multicast addresses to coexist on the same network.

1.1. Requirements Language

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 BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2. Technical Background

Link-scoped IPv6 multicast addresses [[RFC4489](#)] are an effective way to dynamically allocate multicast addresses on the local link. Because this method utilizes SLAAC it is also a zero-configuration technology.

However, according to [[RFC4541](#)], [Section 4](#), most switch vendors forward multicast traffic based only on the MAC address (see the results for Q2 and Q3). There is a problem when transmitting link-scoped IPv6 multicast addresses on Ethernet. According to [[RFC2464](#)], [Section 7](#), the destination multicast Ethernet address is generated by combining the hexadecimal value 3333 with the last four octets of the destination multicast IPv6 address. These last four octets correspond with the group ID in the link-scoped IPv6 multicast address, meaning that any two applications that happen to choose the same group ID will transmit using the same destination multicast Ethernet address. This prevents multicast snooping switches from directing traffic only to devices interested in the data, and may result in a low-bandwidth link being saturated by a high-bandwidth stream.

3. Design Goals

The primary goal is to define a zero-configuration method for dynamically assigning IPv6 multicast addresses and preventing collisions at the Ethernet layer. This method must allow for multiple streams to be transmitted from the same host by different applications that are not cooperating.

A secondary goal is to allow several methods for dynamically assigning IPv6 multicast addresses to coexist on the same network without user configuration.

Advertising the data contained in each multicast stream is outside the scope of this document.

4. Method

When an application is preparing to transmit a multicast stream it generates a link-scoped IPv6 multicast address. The IID is set to the intended source address for the multicast stream. The group ID is a random value in the range reserved for mDNS-based dynamic IPv6 multicast address allocation algorithms (see below). The application then calculates the multicast Ethernet address that will be used to transmit the data [RFC2464], Section 7 and generates a string akin to a reverse mapping domain using a new "eth-addr.arpa" special-use domain.

For example, given a source address of FE80::A12:34FF:FE56:7890, the IPv6 multicast address may be FF32:00FF:A12:34FF:FE56:7890:CFED:2468, the multicast Ethernet address 33:33:CF:ED:24:68, and the string "8.6.4.2.d.e.f.c.3.3.3.3.eth-addr.arpa".

The application then uses the mDNS probing algorithm described in [RFC6762], Section 8.1 to continuously query for a PTR record with the generated string for the name. If the probing algorithm completes without any conflict, then the application begins advertising its own PTR record using that name. The PTRDNAME field is the concatenation of the device's host name, the colon character (:), and the source port of the multicast stream. Integrating the source port in this manner allows for multiple applications to be on the same host. It may then begin transmitting multicast data using that address.

The application should retain the group ID value in long-term storage and use it the next time the multicast stream is transmitted.

If at any point the query returns a result from a different host, then the application stops transmitting that multicast stream and start the process over using a different group ID.

The host should monitor the bus for traffic that uses the same destination multicast Ethernet address, but a different destination multicast IPv6 address. If this is detected then the application acts as if the collision had been detected from the mDNS query.

5. IANA Considerations

The special-use domain "eth-addr.arpa" should be registered in the .arpa registry (https://www.iana.org/domains/arpa) and the "Special-Use Domain Names" registry (https://www.iana.org/assignments/special-use-domain-names).

IANA should create a new registry of ranges for dynamic multicast group IDs that is based on the description in [RFC3307], Section 4.3. The registry should contain the following entries:

0x80000000-0xBFFFFFFF	MADCAP [RFC2730]
0xC0000000-0xCFFFFFFF	

	mDNS-based zero-configuration algorithm described above
0xD0000000-0xFEFFFFFF	Reserved for future zero-configuration algorithms
0xFF000000-0xFFFFFFFF	Solicited-node multicast addresses [RFC4291], Section 2.7.1

Table 1

6. Security Considerations

This algorithm only works in environments where all hosts are cooperating. Malicious hosts could deny service by either repeatedly responding to queries for a given address or by flooding the network with traffic.

7. Acknowledgement

Special thanks to the National Marine Electronics Association for their contributions in developing marine industry standards and their support for this research.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/info/rfc2464>>.
- [RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, DOI 10.17487/RFC3307, August 2002, <<https://www.rfc-editor.org/info/rfc3307>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4489] J-Park, S., M-Shin, K., and J. H-Kim, "A Method for Generating Link-Scoped IPv6 Multicast Addresses", RFC 4489, DOI 10.17487/RFC4489, April 2006, <<https://www.rfc-editor.org/info/rfc4489>>.
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, DOI 10.17487/RFC4541, May 2006, <<https://www.rfc-editor.org/info/rfc4541>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.

[RFC8174]

Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

8.2. Informative References

[RFC2730] Hanna, S., Patel, B., and M. Shah, "Multicast Address Dynamic Client Allocation Protocol (MADCAP)", RFC 2730, DOI 10.17487/RFC2730, December 1999, <<https://www.rfc-editor.org/info/rfc2730>>.

Authors' Addresses

Nate Karstens
Garmin International

Email: nate.karstens@gmail.com

Dino Farinacci
lispers.net

Email: farinacci@gmail.com

Mike McBride
Futurewei

Email: michael.mcbride@futurewei.com