

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 21 August 2026

N. Karstens
Garmin
D. Farinacci
lispers.net
M. McBride
Futurewei
17 February 2026

Zeroconf Multicast Address Allocation Problem Statement and Requirements
draft-ietf-pim-zeroconf-mcast-addr-alloc-ps-13

Abstract

This document surveys current problems with existing protocols for automatically assigning multicast IP addresses in zero-configuration ("zeroconf") networking environments. It addresses key challenges, such as link-layer address collisions, hardware limitations, multicast snooping inefficiencies, and the need to avoid manual configuration. Based on these challenges, it derives requirements for a lightweight, decentralized solution for dynamically allocating unique multicast group addresses without central coordination.

The document presents explicit requirements covering discovery, allocation, conflict detection and resolution, and lease management. It also evaluates considerations specific to IPv6 and IPv4 multicast address ranges, and identifies approaches that are unsuited for zeroconf deployment. This foundation serves as a reference for developing future solutions for multicast address allocation that operate autonomously within local networks.

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 21 August 2026.

Copyright Notice

Copyright (c) 2026 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	2
1.1. Requirements Language	3
2. Link-Layer Address Collisions	4
3. Solution Requirements	5
4. IPv6 Considerations	6
5. IPv4 Considerations	7
6. Security Considerations	8
7. IANA Considerations	8
8. Acknowledgement	8
9. References	8
9.1. Normative References	8
9.2. Informative References	9
Appendix A. Excluded Solutions	10
Authors' Addresses	10

1. Introduction

Multicast communication is commonly used in networks that need to distribute data from one sender to multiple receivers efficiently. In some environments, such as small or isolated networks, multicast must operate without centralized servers or manual configuration. These are referred to as zero-configuration (zeroconf) multicast networks.

One example of such an environment is marine networks, which typically include a mix of sensors, controls, and displays. These networks vary in complexity depending on the size and design of the vessel. Devices may range from low-cost temperature or fluid sensors to high-bandwidth sources such as radar, sonar, or video feeds. Many marine networks are built on a single subnet and rely on Layer 2 Ethernet switches to connect devices.

In these networks, multicast is the most efficient method for distributing sensor data to multiple displays. However, challenges arise when high-bandwidth multicast streams overload links to low-bandwidth devices. Cost-effective switches often do not support source-specific multicast (SSM) because their address table only maps destination MAC addresses. Instead, each multicast stream is assigned a unique destination multicast IP address, and IGMP/MLD snooping [RFC4541] is used to control multicast delivery. This method introduces limitations, especially in environments where switch hardware lacks advanced multicast filtering capabilities.

While marine networks illustrate these issues well, the challenges they face are not unique. Many other zeroconf multicast environments, such as industrial automation, small-scale AV systems, or ad hoc sensor networks, share similar constraints.

[RFC2730] (MADCAP) describes a method for multicast IP address allocation, but its server-based model does not suit a zeroconf environment. [RFC3306] and [RFC4489] both discuss similar approaches to host-based multicast IPv6 address allocation, but neither adequately prevents link-layer address collisions. Although [RFC3307] establishes a framework to avoid multicast address collisions at both IPv6 and link layers, additional refinement is needed to put this into practice (see Section 4).

This document outlines the problem space for zeroconf multicast address allocation, describes the key limitations of current protocols, details sources of link-layer address collisions, and defines a set of requirements for decentralized, zero-configuration multicast address allocation solutions. Multicast applications deployed in a zeroconf environment can use these solutions to dynamically allocate multicast group addresses. The manner in which a multicast group address is used after allocation is outside the scope of this document.

The organization of IPv6 multicast and immense volume of possible multicast addresses makes IPv6 well-suited for zeroconf multicast networks. However, IPv4 is still considered by this document for networks where transitioning to IPv6 remains impractical.

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.

The requirements language is used in Section 3 and applies to implementations conformant to the listed requirements.

2. Link-Layer Address Collisions

Link-layer address collisions are a key concern in multicast networks, particularly when devices rely on zero-configuration operation. Collisions occur when two or more multicast groups are assigned the same link-layer (MAC) address, leading to performance or forwarding issues. This section outlines three scenarios where such collisions can cause problems.

First, many host network interfaces allow filtering of multicast traffic directly in hardware. When an application joins a multicast group, the host network stack typically programs the hardware to accept only traffic for that group. However, if two groups share the same link-layer address, the network interface cannot distinguish them. The network stack is then forced to process unwanted traffic in software, reducing performance and increasing CPU usage.

Second, link-layer address collisions reduce the benefit of using multicast snooping switches on a network. Section 4 of [RFC4541] indicates some switches forward multicast traffic based solely on the link-layer address, without considering the network-layer group (see the results for Q2 and Q3). Although the survey is outdated, switches with these limitations still exist in some of the target deployments. In such cases, if two multicast streams share the same MAC address, traffic may be sent to devices that did not request it. This is especially problematic when low-bandwidth links are overwhelmed by high-bandwidth streams. Additional concerns related to the overlap of IPv6 and link-layer addresses are discussed in [RFC4541], Section 3.

Third, the internal design of some switches can also contribute to collisions. For example, certain switch implementations use a hash table with fixed buckets to store forwarding entries based on MAC addresses (see [US6690667B1]). If multiple addresses hash to the same location and the bucket fills up, additional entries may be dropped or rejected, resulting in forwarding failures.

These examples highlight why a collision-resistant multicast address allocation mechanism is essential in zeroconf environments. The success of this approach depends on the capabilities of the hardware (such as the number of addresses supported by the filters in host network interfaces and the size of the tables maintained in switches/routers) and on the number of multicast groups that are subscribed to by a particular receiver. When the number of entries in the filter tables exceeds availability, a typical fallback mechanism is to

bypass the filter entirely and flood traffic to the receiver. This subsequently incurs a per-packet cost for any software-based filtering that is needed.

3. Solution Requirements

A solution intended for decentralized, zero-configuration multicast IP address assignment is expected to operate in dynamic, infrastructure-free environments. To be effective in such contexts, the solution needs to exhibit the following characteristics:

1. [REQ-1] Unique Address Assignment: Use of the solution SHALL result in a unique address being assigned to the multicast group at both the network and link layers.
2. [REQ-2] Resilience to Single Points of Failure: The solution SHALL be designed to minimize introduction of a single point of failure, ensuring that operation continues even if individual devices or links become unavailable.
3. [REQ-3] Zero User Configuration: It SHALL operate without requiring user or administrator configuration, allowing seamless deployment in unmanaged networks.
4. [REQ-4] Coexistence with Multicast Address Allocation Solutions: The design SHALL allow coexistence with other multicast IP address allocation solutions, including both manual assignment and existing dynamic protocols.
5. [REQ-5] Single-Subnet Operation: It SHALL support operation within a single IP subnet, which is typical in link-local or isolated network environments.
6. [REQ-6] No External Connectivity: The solution SHALL NOT require Internet access or connectivity to external infrastructure.
7. [REQ-7] Supports Multiple Host Applications: It SHALL support multiple applications on the same host, each independently allocating multicast addresses and transmitting to those addresses.
8. [REQ-8] Collision Detection and Resolution: The solution SHALL include mechanisms to detect and resolve multicast address collisions at both the network and link layers.

Note: In rare cases, collisions may arise after a temporary network partition, when different parts of the network allocate the same multicast address independently. Upon reconnection, such collisions SHALL be detectable and resolved gracefully by ensuring that conflicting streams are migrated to unique addresses.

In addition to the above, the following characteristics are considered desirable, but are left as recommendations to allow for flexibility in solution design:

1. [CONS-1] Multi-Subnet Support: Operation across multiple subnets is beneficial in more complex or routed environments and SHOULD be supported.
2. [CONS-2] Standards Compatibility: The solution SHOULD aim to minimize the need for changes to existing protocols or standards that affect backwards compatibility or deployment in existing networks.
3. [CONS-3] Cross-Platform Availability: It SHOULD use capabilities that are widely available across host platforms and operating systems.
4. [CONS-4] Minimal Dependency on Manufacturing Data: It SHOULD avoid reliance on pre-loaded configuration or device-specific manufacturing data.
5. [CONS-5] Low Overhead: The solution SHOULD minimize the volume and frequency of network traffic generated during normal operation.
6. [CONS-6] Advertisement: The solution SHOULD describe a mechanism for advertising and discovering multicast addresses allocated for an application.
7. [CONS-7] Network Topology: To allow compatibility with a variety of networks, the solution SHOULD work independent of the dynamics of the underlying topology and adjacencies.

4. IPv6 Considerations

The rules for IPv6 multicast addresses, described in [RFC3307], are comprehensive and well-organized. These rules must be followed by any solution allocating IPv6 multicast addresses, including solutions that meet the requirements defined in this document. However, some aspects of its current organization need to be improved to ensure that a zeroconf multicast address assignment solution can coexist with other IPv6 multicast address assignment protocols.

For example, Section 2 of [RFC3307] explains that the last 32 bits of an IPv6 multicast address, called the group ID, are mapped directly to the Ethernet MAC address. Different parts of the group ID range are assigned based on how the address is allocated. Section 4.3 of [RFC3307] describes two ways to assign group IDs dynamically: one where a server assigns addresses, and one where hosts assign addresses themselves. However, both methods use the same group ID range, which creates a risk of address collisions if both are used at the same time.

An additional concern is that the group IDs used for this dynamic range overlap with the range used for Solicited-Node multicast addresses, a special type of multicast used by IPv6 for neighbor discovery (see Section 2.7.1 of [RFC4291]). This overlap increases the risk of unintentional conflicts with link-layer addresses.

Note that [RFC3307] focuses on 48-bit addresses on Ethernet, but similar issues would be seen on any medium that generates link-layer multicast addresses by truncating an IPv6 multicast address.

5. IPv4 Considerations

In IPv4, multicast addresses can sometimes cause conflicts at the data link layer. For example, as explained with Ethernet in Section 6.4 of [RFC1112], this happens because only the lower 23 bits of an IPv4 multicast address are used to generate the Ethernet multicast address. Since an IPv4 multicast address is 32 bits and starts with a fixed 4-bit prefix (leaving 28 bits), this means up to $2^{(28-23)} = 32$ different multicast IP addresses can map to the same Ethernet address. As a result, devices may receive multicast traffic they didn't ask for.

The address allocation guidelines in [RFC5771] did not account for this type of collision when they were created. Because of this limitation, the recommended approach for new designs that need dynamic multicast IP address assignment is to use IPv6 instead of IPv4.

However, if using IPv4 is necessary, then multicast addresses SHOULD be chosen carefully from within the Administratively Scoped Block (239.0.0.0/8). Additionally, solutions for zeroconf multicast address allocation SHOULD try to avoid using addresses that may already be in use by other applications on the same network, to minimize the risk of conflicts.

Zeroconf coexistence with other IPv4 multicast address allocation solutions may not be possible, in which case it may be necessary to require manual configuration or to limit the solutions that are deployed.

6. Security Considerations

Zeroconf multicast address allocation mechanisms are vulnerable to accidental or malicious address collisions, which may lead to denial of service or misdirection of traffic. Solutions derived from these requirements SHALL include measures for collision detection and conflict resolution [REQ-7], and SHOULD include measure to prevent unauthorized address use. Specific security mechanisms are outside the scope of this document.

7. IANA Considerations

This document has no IANA actions.

8. Acknowledgement

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

Thanks also to the members of the PIM working group for their early brainstorming sessions and review of this draft, and to Gunter van de Velde for his review and suggestions.

9. References

9.1. Normative References

- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <<https://www.rfc-editor.org/info/rfc1112>>.
- [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>>.
- [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>>.
- [RFC5771] Cotton, M., Vegoda, L., and D. Meyer, "IANA Guidelines for IPv4 Multicast Address Assignments", BCP 51, RFC 5771, DOI 10.17487/RFC5771, March 2010, <<https://www.rfc-editor.org/info/rfc5771>>.
- [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>>.

9.2. 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>>.
- [RFC3306] Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6 Multicast Addresses", RFC 3306, DOI 10.17487/RFC3306, August 2002, <<https://www.rfc-editor.org/info/rfc3306>>.
- [RFC4489] Park, J., Shin, M., and 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>>.
- [US6690667B1] Warren, D., "United States Patent 6690667B1: Switch with adaptive address lookup hashing scheme", 10 February 2004.

Appendix A. Excluded Solutions

The way multicast IP addresses are mapped to link-layer multicast addresses is already defined in existing standards, such as [RFC1112] for IPv4 over Ethernet and [RFC2464] for IPv6 over Ethernet. These standards specify a fixed prefix used in creating the Ethernet multicast address. Changing this prefix would open the door to new solutions, but those are not being considered here for practical reasons.

One idea is to reduce the size of the fixed prefix, which would leave more bits available for the group ID. This would make address collisions less likely. Another idea is to create a new protocol that dynamically maps multicast IP addresses to link-layer addresses, similar to how DHCP assigns IP addresses. This protocol could work locally on a subnet, and routers could adjust the mapping for incoming multicast traffic at the network edge.

However, these ideas would require significant changes to how network devices handle multicast traffic. Since existing hardware and operating systems are built around the current standards, it's unlikely that such changes would be widely supported anytime soon.

Another potential solution for IPv4 was to assign 32 separate, non-overlapping address ranges to avoid collisions altogether (e.g., assign 224.0.0.254, 224.128.0.254, 225.0.0.254, etc.). But this was rejected because [RFC5771] discourages new allocations, given how limited the IPv4 multicast address space already is.

Authors' Addresses

Nate Karstens
Garmin International, Inc.
1200 E. 151st St.
Olathe, KS 66062-3426
United States of America
Email: nate.karstens@gmail.com

Dino Farinacci
lispers.net
San Jose, CA
United States of America
Email: farinacci@gmail.com

Mike McBride
Futurewei
United States of America
Email: michael.mcbride@futurewei.com