Network Working Group
Internet-Draft
Intended status: Informational
Expires: 3 April 2026

N. Karstens
Garmin
D. Farinacci
lispers.net
M. McBride
Futurewei

30 September 2025

Zeroconf Multicast Address Allocation Problem Statement and Requirements
draft-ietf-pim-zeroconf-mcast-addr-alloc-ps-07

Abstract

This document defines the problem space and associated requirements for automatically assigning multicast addresses in zero-configuration ("zeroconf") networking environments. It addresses key challenges, such as 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 protocol capable of 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 multicast address allocation protocols that operate autonomously within local networks.

Table of Contents

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. Most 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), so IGMP 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. This document outlines the problem space for zeroconf multicast address allocation, describes the key limitations of current solutions such as MADCAP [RFC2730], and defines a set of requirements for a decentralized, zero-configuration multicast address allocation protocol.

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. Address Collisions

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

Second, networks that use multicast snooping switches are particularly vulnerable. As described in [RFC4541], Section 4, many switches forward multicast traffic based solely on the link-layer address, without considering the network-layer group (see the results for Q2 and Q3). 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 [US6690667B1] use hash tables to store forwarding entries based on MAC addresses. If multiple addresses hash to the same location and the table 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.

3. Protocol Requirements

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

1. Resilience to Failure: The protocol should function without reliance on a single point of failure, ensuring that operation continues even if individual devices or links become unavailable.

2. Zero User Configuration: It should operate without requiring user or administrator configuration, allowing seamless deployment in unmanaged networks.

3. Protocol Coexistence: The design should allow coexistence with other multicast address allocation methods, including both manual assignment and existing dynamic protocols.

4. Single-Subnet Operation: It should support effective operation within a single IP subnet, which is typical in link-local or isolated network environments.

5. No External Connectivity: The protocol should not require Internet access or connectivity to external infrastructure.

6. Host-Level Multiplexing: It should support multiple applications on the same host, each independently allocating and using multicast addresses.

7. Collision Detection and Resolution: The protocol should include mechanisms to detect and resolve multicast address collisions, including those that may occur due to network partitions and subsequent re-merging of segments.

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 should be detectable and resolved gracefully.

In addition to the above, the following characteristics are considered desirable:

1. Multi-Subnet Support: Support for operation across multiple subnets is beneficial in more complex or routed environments.

2. Standards Compatibility: The protocol should aim to minimize the need for changes to existing protocols or standards.

3. Cross-Platform Availability: It should use capabilities that are widely available across platforms and operating systems.

4. Support for Unprivileged Applications: The protocol should function without requiring elevated privileges, enabling broader applicability in user-space applications.

5. Minimal Dependency on Manufacturing Data: It should avoid reliance on pre-loaded configuration or device-specific manufacturing data.

6. Low Overhead: The protocol should minimize the volume and frequency of network traffic generated during normal operation.

4. IPv6 Considerations

The rules for IPv6 multicast addresses, described in [RFC3307], are comprehensive and well-organized. However, some aspects of its current organization need to be improved to ensure that a zeroconf multicast address assignment protocol can coexist with other IPv6 multicast protocols.

For example, section 2 of this RFC 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 the same RFC 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 this dynamic range overlaps 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.

A solution to these issues is presented in [I-D.ietf-pim-updt-ipv6-dyn-mcast-addr-grp-id].

5. IPv4 Considerations

In IPv4, multicast addresses can sometimes cause conflicts at the Ethernet (link-layer) level. As explained 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, this means up to 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 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, the protocol 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.

6. Excluded Solutions

The way multicast IP addresses are mapped to Ethernet (link-layer) multicast addresses is already defined in existing standards: [RFC1112] for IPv4 and [RFC2464] for IPv6. 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 Ethernet 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. But this was rejected because [RFC5771] discourages new allocations, given how limited the IPv4 multicast address space already is.

7. Security Considerations

Security considerations will be discussed by any proposed zero-configuration multicast address allocation algorithm.

8. IANA Considerations

This document has no IANA actions.

9. 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.

10. References

10.1. 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>.

[RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, DOI 10.17487/RFC3307, August 2002, <https://www.rfc-editor.org/info/rfc3307>.

[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>.

10.2. Informative References

[I-D.ietf-pim-updt-ipv6-dyn-mcast-addr-grp-id] Karstens, N., Farinacci, D., and M. McBride, "Updates to Dynamic IPv6 Multicast Address Group IDs", Work in Progress, Internet-Draft, draft-ietf-pim-updt-ipv6-dyn-mcast-addr-grp-id-07, 25 August 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-pim-updt-ipv6-dyn-mcast-addr-grp-id-07>.

[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>.

[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>.

[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>.

[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>.

[US6690667B1]
Warren, D., "United States Patent 6690667B1: Switch with adaptive address lookup hashing scheme", 10 February 2004.

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