

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 4 October 2026

N. Karstens
Garmin
D. Farinacci
lispers.net
M. McBride
Futurewei
2 April 2026

Zero-Configuration Assignment of IPv6 Multicast Addresses Using mDNS
draft-ietf-pim-ipv6-zeroconf-assignment-10

Abstract

This document describes a zero-configuration protocol for dynamically assigning IPv6 multicast addresses that are unique at the link-layer. Applications randomly assign multicast group IDs from a specified range and prevent collisions by using Multicast DNS (mDNS) to publish resource records under a new "eth-addr.arpa" domain. This protocol satisfies all of the criteria listed in draft-ietf-pim-zeroconf-mcast-addr-alloc-ps.

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 4 October 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. Procedure	3
2.1. Veto Records	5
3. Use on Networks with Multiple Subnets	5
4. Evaluation of Solution	5
4.1. Collision Detection After Partition Repair	7
5. IANA Considerations	7
5.1. Domain Name Reservation Considerations	8
6. Security Considerations	8
7. Acknowledgement	9
8. References	9
8.1. Normative References	9
8.2. Informative References	9
Authors' Addresses	10

1. Introduction

[I-D.ietf-pim-zeroconf-mcast-addr-alloc-ps] includes a problem statement and requirements for a zero-configuration method for dynamically assigning multicast addresses. This document describes a process that fulfills these requirements by having applications randomly assign IPv6 multicast group IDs from a specified range and using mDNS [RFC6762] to prevent collisions in both IPv6 and link-layer addresses.

Note that DNS-based Service Discovery (DNS-SD) [RFC6763] uses several different DNS resource record types, published using either Unicast or Multicast DNS, to facilitate service discovery. This document uses a single DNS resource record type (PTR), published using Multicast DNS (mDNS), to coordinate IPv6 multicast address assignment in a zero-configuration environment. The DNS resource records in this protocol may be published alongside records for other domain name services, such as DNS-SD, or they may be published alone. mDNS is used rather than a new protocol with the expectation that functionality for address assignment can be achieved using existing mDNS implementations.

This protocol is well-suited for networks that rely on IPv6 multicast and already deploy mDNS.

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. Procedure

When an application is preparing to transmit a multicast stream, it SHALL first generate a random group ID in the range 0x90000000-0x9FFFFFFF. IANA is REQUESTED to assign this range from the "Dynamic Multicast Group IDs" registry (see Section 5). The application SHALL then generate a link-scoped IPv6 multicast address by combining the group ID with the Interface Identifier (IID) of the intended source address for the multicast stream, according to the format given in [RFC4489], Section 3. The application SHALL then calculate the corresponding multicast Ethernet address that will be used to transmit the data (see [RFC2464], Section 7).

Once the multicast Ethernet address is determined, the application SHALL then use this to construct a domain name by encoding nibbles from the multicast Ethernet address as a sequence of lower-case hexadecimal digits, separated by dots, in reverse order, and ending with the suffix ".eth-addr.arpa" (a new domain in the .arpa registry).

For example, given a source address of fe80::a12:34ff:fe56:7890, the IPv6 multicast address may be ff32:ff:a12:34ff:fe56:7890:9abc:def0, the group ID 9abc:def0, the multicast Ethernet address 33:33:9A:BC:DE:F0, and the resulting string is "0.f.e.d.c.b.a.9.3.3.3.eth-addr.arpa".

Note that this is similar to how DNS resource records used for reverse mapping IPv6 addresses are created (see [RFC3596], Section 2.5).

Once the domain name is created, the application SHALL then use the mDNS probing algorithm described in [RFC6762], Section 8.1 to continuously query for a PTR record with this domain name. If the probing algorithm completes without any conflict, then the application SHALL begin advertising its own unique PTR record using that name. The PTRDNAME field of this record SHALL consist of a unique application identifier, in the form of one or more DNS labels, followed by the device's host name (for example, "application.example.local."). Integrating a unique identifier in this manner allows for multiple applications to be on the same host. Note that A/AAAA records may also be published for this host name ("example.local."), though this is not a requirement for this design.

Because this protocol is focused specifically on allocating IPv6 multicast addresses, resource records MUST be published using the IPv6 multicast address for mDNS. In order to be compatible with existing mDNS implementations, resource records MAY also be published using the IPv4 multicast address for mDNS.

mDNS implementations used with this protocol MUST facilitate interaction between multiple applications using the same mDNS implementation, and with other mDNS implementations on the same host. This ensures that multiple applications that happen to generate the same group ID will detect this when probing and use a different group ID.

Once the PTR record is advertised, the host MAY then begin transmitting multicast data using the generated address.

The application SHALL retain the group ID value in persistent storage and use it the next time the multicast stream is transmitted. This allows the network to quickly settle on a configuration that will never have another collision as long as the network is unchanged.

If a conflict is detected at any point, then the application SHALL stop transmitting that multicast stream and start the process over using a different group ID. As before, this new group ID is also retained in persistent storage, overwriting any group ID previously saved for this multicast stream.

The host network stack may optionally monitor the network for traffic that uses the same destination multicast Ethernet address, but a different destination multicast IPv6 address. If this is detected, then the application SHALL respond the same as with a collision.

While intended primarily for allocating IPv6 multicast addresses on the same subnet (link-local scope), the same technique could also apply to a larger network as long as mDNS traffic is routed between subnets (for any scope excluding global scope).

2.1. Veto Records

[I-D.ietf-pim-zeroconf-mcast-addr-alloc-ps], Section 2 describes collisions occurring in the network infrastructure (for example, due to switch design). When an infrastructure component detects a collision it cannot resolve, it triggers a conflict with the application by publishing a veto record. A veto record is a unique PTR record using the string generated for the address as its name and the PTRDNAME field set to the string "veto", formatted as a DNS label. The veto record is published without probing.

Applications SHALL respond to the conflict the same as to a collision, eventually resulting in a new group ID being acquired. As before, the application retains its new group ID in persistent storage, ensuring the same conflict is not repeated in the future.

3. Use on Networks with Multiple Subnets

The protocol can be extended across multiple subnets if PTR records are distributed between subnets (for example, by using an mDNS reflector or the Discovery Proxy described in [RFC8766]). Note that the stream MUST use a unicast-prefix based IPv6 multicast address ([RFC3306]), as the link-scoped IPv6 multicast address used above is scoped to the local link.

The protocol's reliance on cooperating hosts (see Section 6) makes it unsuited for use on the global Internet. [RFC8815] recommends Source-Specific Multicast in this environment.

4. Evaluation of Solution

[I-D.ietf-pim-zeroconf-mcast-addr-alloc-ps] contains a list of criteria to evaluate potential solutions. The following is an analysis of how this protocol satisfies the requirements listed in that document:

1. [REQ-1] Unique Address Assignment: Use of the protocol results 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 protocol uses mDNS, which uses a peer-to-peer communication model and so has no single points of failure, as long as network connectivity is not interrupted.
3. [REQ-3] Zero User Configuration: The protocol operates without requiring user or administrator configuration.
4. [REQ-4] Coexistence with Multicast Address Allocation Solutions: This document assigns a range of group IDs from the "Dynamic Multicast Group IDs" registry for use with the protocol, which allows the protocol to coexist with other multicast IP address allocation solutions (see [I-D.ietf-pim-updt-ipv6-dyn-mcast-addr-grp-id]).
5. [REQ-5] Single-Subnet Operation: The protocol supports operation within a single IPv6 subnet.
6. [REQ-6] No External Connectivity: The protocol does not require Internet access or connectivity to external infrastructure.
7. [REQ-7] Supports Multiple Host Applications: Multiple applications on the same host are supported by 1) incorporating a unique application identifier into the PTR record's PTRDNAME field and 2) by requiring the mDNS implementation(s) on a host to facilitate interaction between those applications.
8. [REQ-8] Collision Detection and Resolution: The protocol includes a mechanism to detect and resolve multicast address collisions at both the network and link layers. See Section 4.1 for additional discussion about detecting and resolving collisions after a network partition is repaired.

The following is an analysis of how this protocol satisfies the desirable characteristics listed in [I-D.ietf-pim-zeroconf-mcast-addr-alloc-ps]:

1. [CONS-1] Multi-Subnet Support: The protocol can operate across multiple subnets (see Section 3).
2. [CONS-2] Standards Compatibility: The protocol uses existing protocols without requiring any changes.
3. [CONS-3] Cross-Platform Availability: The protocol uses mDNS, which is implemented on many host platforms and operating systems.

4. [CONS-4] Minimal Dependency on Manufacturing Data: The protocol does not rely on pre-loaded configuration or device-specific manufacturing data.
5. [CONS-5] Low Overhead: The protocol uses mDNS, which is designed to minimize the volume and frequency of network traffic generated during normal operation.
6. [CONS-6] Advertisement: This document does not prescribe any mechanism for advertising the IPv6 address used for the multicast stream. However, because the protocol already uses mDNS, a natural option for doing so would be to use DNS-SD ([RFC6763]) to advertise a "_udp" service and publish the address used for the multicast stream in a TXT record.
7. [CONS-7] Network Topology: Veto records ensure that the protocol works regardless of the underlying topology and adjacencies.

4.1. Collision Detection After Partition Repair

Because mDNS is designed to be a low-bandwidth protocol, it can take a significant amount of time to detect a resource record collision after a network partition is repaired. This is not a concern on networks where all multicast streams are established before any likely partition event because all group IDs will have been selected and stored for future use.

It is a greater concern on networks where multicast streams may be established at any time. Deployments on these networks may consider engaging supplemental detection and resolution mechanisms. Specifics of this functionality are left as a future enhancement.

5. IANA Considerations

IANA should allocate a block of group IDs from the "Dynamic Multicast Group IDs" registry in the "IPv6 Multicast Address Space" registry group that was created by [I-D.ietf-pim-updt-ipv6-dyn-mcast-addr-grp-id]. The range of this block should be 0x90000000-0x9FFFFFFF and the description should be the title of this document.

The domain "eth-addr.arpa" should be registered in the .arpa registry (<https://www.iana.org/domains/arpa>). The usage should indicate that it is for mapping Ethernet addresses to local domain names and it should reference this document.

The special-use domain "9.3.3.3.3.eth-addr.arpa." should be registered in the "Special-Use Domain Names" registry (<https://www.iana.org/assignments/special-use-domain-names>). This domain should not be delegated.

5.1. Domain Name Reservation Considerations

[RFC6761], Section 5 includes a list of questions that must be answered when reserving a new Special-Use Domain Name. The answers to these questions for the "9.3.3.3.3.eth-addr.arpa" Special-Use Domain, and any names falling within this domain, are as follows:

1. Users are free to use these names as they would for any other reverse-mapping domain. Because this mapping is closely-related to the local network, users SHOULD be aware that these names are likely to yield different results on different networks.
2. Application software SHOULD recognize these names as a reverse-mapping domain and MAY associate them with the protocol described in this document.
3. Name resolution APIs and libraries SHOULD recognize these names as being used exclusively with mDNS, and so SHOULD NOT send queries for these names to their configured unicast DNS server(s).
4. Caching DNS servers SHOULD recognize these names as being used exclusively with mDNS, SHOULD NOT attempt to look up resource records associated with these names, and SHOULD respond to any query with NXDOMAIN.
5. Authoritative DNS servers SHOULD recognize these names as being used exclusively with mDNS and SHOULD respond to any query with NXDOMAIN.
6. DNS server operators MUST NOT configure an authoritative DNS server to answer queries for these names.
7. DNS Registries/Registrars MUST NOT register 9.3.3.3.3.eth-addr.arpa names.

6. Security Considerations

As with mDNS itself, this protocol only works in environments where all hosts are cooperating. Malicious hosts could deny service by repeatedly triggering address conflicts (for example, by publishing veto records).

7. Acknowledgement

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

Thanks also to the members of the PIM working group for their early brainstorming sessions and review of this draft, and to Esko Dijk for his review of the draft.

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>>.
- [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>>.
- [RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <<https://www.rfc-editor.org/info/rfc6761>>.
- [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

- [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-13, 10 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-pim-updt-ipv6-dyn-mcast-addr-grp-id-13>>.
- [I-D.ietf-pim-zeroconf-mcast-addr-alloc-ps]
Karstens, N., Farinacci, D., and M. McBride, "Zeroconf Multicast Address Allocation Problem Statement and Requirements", Work in Progress, Internet-Draft, draft-ietf-pim-zeroconf-mcast-addr-alloc-ps-13, 17 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-pim-zeroconf-mcast-addr-alloc-ps-13>>.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", STD 88, RFC 3596, DOI 10.17487/RFC3596, October 2003, <<https://www.rfc-editor.org/info/rfc3596>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.
- [RFC8766] Cheshire, S., "Discovery Proxy for Multicast DNS-Based Service Discovery", RFC 8766, DOI 10.17487/RFC8766, June 2020, <<https://www.rfc-editor.org/info/rfc8766>>.
- [RFC8815] Abrahamsson, M., Chown, T., Giuliano, L., and T. Eckert, "Deprecating Any-Source Multicast (ASM) for Interdomain Multicast", BCP 229, RFC 8815, DOI 10.17487/RFC8815, August 2020, <<https://www.rfc-editor.org/info/rfc8815>>.

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