

IPv6 Maintenance  
Internet-Draft  
Updates: 2464, 6085, 7371 (if approved)  
Intended status: Standards Track  
Expires: 29 July 2026

D. Lamparter  
NetDEF, Inc.  
25 January 2026

Sub-Link Scoped IPv6 Multicast Addressing  
draft-equinox-6man-sub-link-scope-multicast-01

## Abstract

The IPv6 addressing architecture for multicast has the scope of a multicast group embedded in its address, with the smallest non-reserved scopes being interface-local and link-local, numbered 1 and 2. This document suggests the introduction of a scope inbetween these two, for use with lower-layer transport multicast that reaches parts of a link. Since there is no room to insert a scope value for this, a separate address block is used. A mapping for Ethernet as lower-layer transport is provided.

## 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 29 July 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
2. Requirements Language . . . . .	3
3. Behavior of Sub-Link scoped multicast groups / addresses . . . . .	3
4. Sub-Link scoped multicast address structure . . . . .	4
5. MLD applicability . . . . .	5
6. Operating system API considerations . . . . .	5
7. Bindings to Ethernet . . . . .	6
7.1. Ethernet group usage guidance . . . . .	7
7.2. Default groups . . . . .	8
8. Group identifiers . . . . .	8
8.1. All Nodes and All Routers addresses . . . . .	9
9. Security Considerations . . . . .	9
10. Privacy Considerations . . . . .	9
11. IANA Considerations . . . . .	9
11.1. Variable Scope Multicast Addresses . . . . .	9
11.2. Link Types for Sub-Link Scope Multicast Addresses . . . . .	10
11.3. Sub-Link Scope Multicast Addresses . . . . .	10
11.3.1. Initial contents . . . . .	11
12. References . . . . .	11
12.1. Normative References . . . . .	11
12.2. Informative References . . . . .	12
Acknowledgements . . . . .	13
Implementation note . . . . .	13
Revision history (TO BE REMOVED) . . . . .	13
Author's Address . . . . .	13

## 1. Introduction

// This draft lives at <https://github.com/eqvinox/6man-sub-link-scope-multicast>

A major application of IPv6 multicast is in discovery protocols to find other systems participating in the same protocol on the same link. These applications commonly use an IPv6 multicast address in the ff02::/16 range, i.e. scoped to the link.

In some cases however, it is useful to further limit the scope of discovery for an application. In particular, a device's immediate attachment segment to a layer 2 domain (i.e. switch) is useful for hybrid layer 2/3 setups (e.g. EVPN [EVPN]), as well as for situations where the first layer 2 hop might be trusted but other participants in the broadcast domain are not.

Ongoing work in this area has resorted to using LLDP [LLDP] as a transport and encapsulating their data, e.g. [BGP-LLDP] and [HOSTRT-LLDP]. However, LLDP was designed as a layer 2 discovery protocol, and its use in such applications has drawbacks like limiting choice on the actual scope getting used, interacting nontrivially with STP, complicating security considerations, and first and foremost creating dependencies between components that are normally independent of each other.

The desirable functionality in these cases is not necessarily LLDP itself, but rather the limited scope of propagation for the discovery protocol. This document exposes these scopes for use in IPv6 multicast.

## 2. 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.

## 3. Behavior of Sub-Link scoped multicast groups / addresses

The representation of addresses in this document introduces two fields, "linktype" and "linkspecific". The former is used to index into new IANA-managed allocations of values identifying lower-layer link technologies. Given that value, the latter is then used to identify a particular lower-link multicast behavior.

Packets with a multicast destination address of this structure behave as follows:

1. Packets indicating an unknown linktype, and packets indicating a linktype that does not match the link in use MUST be discarded, and MUST NOT be sent.

2. If for a received packet, the scope used to send it can be identified, e.g. by a lower-layer destination address or other field indicating scope of distribution, packets where this does not match the "linkspecific" value MUST be discarded and MUST NOT be sent. This also includes discarding packets that used unicast on the lower layer.
  3. "Optimizing" multicast addresses mapping to lower-layer unicast (e.g. [MCinUC]) MUST NOT be applied.
  4. Packets with "linkspecific" values not specified by a link layer's bindings MUST be discarded and MUST NOT be sent.
  5. All mechanisms governing multicast packets with link-local scope apply. In particular, they MUST NOT be forwarded onto another link by a multicast router.
  6. Packets discarded by the above requirements SHOULD be counted and/or logged, but if logging is implemented it MUST be limited as to prevent denial of service (CPU and log disk space particularly) attacks. Not reporting these mismatches deprives the operator of an indicator that some security breach might be attempted and should only be considered if the node has no good way to report them.
4. Sub-Link scoped multicast address structure

[MCASTARCH] defines the structure of IPv6 multicast addresses as follows:

8	4	4	4	4	8	64	32
11111111	ff1	scop	ff2	rsvd	plen	network prefix	group ID

With ff1 as follows:

X	Y	P	T
---	---	---	---

The combination of P = 1, T = 0 was explicitly forbidden. This document redefines that combination to indicate a sub-link scoped multicast address.

For an IPv6 multicast address with this combination of bits, the scope value MUST be set to 2. The following further structure is defined:

8	4	4	4	4	72 bits	32 bits
11111111	XY10	0010	ff2	type	linkspecific	group ID

The (non-constant) fields are as follows:

X, Y and ff2 As in [MCASTARCH]

type ("linktype") Values from a newly established registry that identifies the lower-layer link technology in use, and therefore the meaning of the next field.

linkspecific Lower-layer link specific information identifying which lower-layer multicast group/mechanism is to be used for this group.

group ID Identifies the multicast group as with other types of multicast group addresses.

No meaning is defined for the X and Y bits, as well as the flags in the ff2 field. They MUST be sent as zero and packets received with nonzero values MUST be discarded until some other document assigns some meaning to them. (The use of an embedded RP address is nonsensical for a multicast group that is never forwarded, as such the interpretation of Y to signal an embedded RP address is not applicable here.)

There is no conflict with the "plen"/"network prefix" fields used in other multicast addresses since the scope defined here is less than even a link. Deriving unique addresses on a larger scale is thus unnecessary.

The use of P = 1, T = 0 with other scope values other than 2 is not specified by this document, currently reserved, and available for future use elsewhere.

## 5. MLD applicability

to be decided

## 6. Operating system API considerations

While multicast group addresses as outlined in this document fit the existing multicast socket interface outlined in [BASICSOCKET] and [SSMSOCKET], the following considerations apply:

1. since packets with mismatching lower layer vs. IPv6 indicators are already required to be discarded, applications can expect a packet received with a sub-link-scope multicast addresses have in fact been limited to the indicated scope of forwarding.
2. requests by an application to join a group with unsupported or invalid "linktype" or "linkspecific" value SHOULD be rejected with an error. Not rejecting such values runs the risk of complicating future introductions of new identifiers, and introduces a risk of generating invalid packets by confusing link-layer technologies on send.
3. for transmission, the sub-link behavior is not separately requested, the destination address is the indicator and any packets to an address described in this document MUST use the appropriate lower-layer delivery without further action by the application.
4. the API MUST provide a way to determine whether the behavior described in this document is in fact supported. If the system was previously rejecting multicast addresses in the ff22::/16 range, the fact that they are now accepted is sufficient. However, most known systems are in fact accepting applications joining or sending to these groups. In that case, an explicit query for support for this mechanism MUST be provided.

## 7. Bindings to Ethernet

Ethernet is by far the most commonly used lower-layer link technology carrying IPv6 at this point. For Ethernet's less than whole link multicast addresses, [IPV6oETH] is updated for the following more specific address format and mapping:

8	4	4	4	4	24 bits	48 bits	32 bits
11111111	淡淡	10	0010	淡淡淡淡	1110	MAC_suffix	reserved   group ID

maps to: 01-80-C2-MAC\_suffix

(The "淡" symbol is used to indicate reserved flag fields that are currently required to be zero.)

The (non-constant) fields are as follows:

linktype Set to Eh, indicating the lower-layer is Ethernet and its multicast capabilities are being expressed.

**MAC\_suffix** The second half of an Ethernet address with special behavior defined in 802.1Q. The first half of such addresses is 01:80:C2. Concatenating the two halves forms a full Ethernet address.

**reserved** No meaning is currently assigned to these bits. They MUST be sent as zero and packets with nonzero values in this field MUST be discarded.

**group ID** As elsewhere.

While 802.1Q does not currently define any specific behavior outside of the range 01:80:C2:00:00:00 to 01:80:C2:00:00:3F, the entire block is made representable in the interest of future proofing.

This section is applicable to all link layers using MAC-48 addresses and the forwarding behavior described in 802.1Q. This notably also includes 802.11, despite the additional multicast considerations for wireless networks.

Since the destination address in a received Ethernet frame indicates which multicast scope it was distributed to, it MUST be verified to match the MAC suffix in the IPv6 address as noted in Section 3. [MCinUC] MUST NOT be applied.

When joining any of these groups on a layer 2 forwarding device's IPv6 stack, the join MUST NOT affect forwarding behavior for Ethernet frames addressed to these Ethernet multicast addresses, including IPv6 packets. Whether or not these groups are forwarded or not is solely defined by IEEE 802.1Q.

### 7.1. Ethernet group usage guidance

Since the 802.1Q definitions have mostly been made with an intent of a specific use case, they do not directly express a forwarding scope, rather a forwarding scope is derived from their use. However, some of the addresses have shifted into functioning as scope indicator. The use of such addresses is RECOMMENDED and at the time of creation of this document they were, in ascending distribution size:

01:80:C2:00:00:0E - ff22:e00:e::/96 never forwarded, not even by Two-Port MAC Relays (TPMRs; intelligent media converters.) Originally intended for LLDP. Note that this group may bypass the port-blocked state in STP since the forwarding loops avoided by STP are not a concern when forwarding is as restricted as with this group, and LLDP was intended to work on ports that are blocked in STP.

01:80:C2:00:00:03 - ff22:e00:3::/96 Forwarded only by TPMRs, terminates at nearest multi-port relay of any kind. Primarily used for 802.1X port authentication (and therefore also MACsec key agreement). Since TPMRs are more or less

01:80:C2:00:00:00 - ff22:e00:0::/96 Forwarded until closest "Customer Bridge", but in practice refers to STP-enabled switches, i.e. is dropped by any switch running STP but forwarded if STP is disabled.

The above list is a recommendation only, and ultimately it is up to the architects developing a particular use of sub-link scoped multicast to choose an appropriate group given the constraints and behavior in the IEEE 802 standards.

## 7.2. Default groups

The following groups SHOULD be automatically joined by all IPv6 nodes implementing this specification, since they are useful to operators for OAM purposes:

- \* ff22:e00:3::1 (All Nodes on Ethernet segment to next multiport switch)
- \* ff22:e00:e::1 (All Nodes on immediate Ethernet segment)

Additionally, the following group SHOULD be joined by all IPv6 routers implementing this specification:

- \* ff22:e00:3::2 (All Routers on Ethernet segment to next multiport switch)

(A TPMR also being a router is considered self-contradictory, it would no longer fulfill the IEEE 802.1Q criteria for a TPMR. The ff22:e00:e::2 group is therefore omitted.)

Some implementations may choose to not join these groups out of an abundance of caution for security. It should however be noted that the limited scope of these groups inherently makes them quite difficult to abuse.

## 8. Group identifiers

Group identifiers for multicast addresses in this range function the same way as identifiers for other scopes. In particular:



1. Variable Scope Multicast Addresses allocated in [IANA-IPv6MC] are applicable to this range. The leading "FF0X:" prefix on addresses in that range is extended to also include "FF22:".
2. As with other scopes, a scope specific registry will track group identifiers used specifically with this scope.

#### 8.1. All Nodes and All Routers addresses

The ::1 and ::2 group identifiers are assigned to the same function they have in the link-local scope. However, since there is a multitude of group addresses with these group identifiers, but different link layer specific values in the upper part of the address, recommendations to automatically join some of these groups are only made in the definitions of link layer specific mappings.

#### 9. Security Considerations

Exposing the limited scope multicast functionality from lower layers can be used to improve security properties of discovery protocols since there is then a guarantee that the packet originated from a limited set of devices.

However, to achieve this gain in security, it is paramount that operating systems in fact implement the discard requirements described in section Section 3 MUST absolutely be enforced by all implementations.

#### 10. Privacy Considerations

\_TBD\_

#### 11. IANA Considerations

This document requests the update of one IANA registry and creation of two new IANA registries, all in the IPv6 Multicast Address Space Registry group.

##### 11.1. Variable Scope Multicast Addresses

The description and applicability of the "Variable Scope Multicast Addresses" registry is modified to add to the note:

| The addresses assigned here are also applicable to the Sub-Link  
| scope, with the FF0X: prefix being replaced with "FF22:" in that  
| case.

## 11.2. Link Types for Sub-Link Scope Multicast Addresses

New registry:

Name Link Types for IPv6 Sub-Link-Scoped Multicast Addresses

Registry group IPv6 Multicast Address Space Registry

Registration procedure as indicated in initial table contents

Field names:

- \* "linktype" (hex)
- \* Lower-Layer protocol
- \* Reference/Procedure

The initial contents of the registry are:

"linktype" (hex)	Lower-Layer protocol	Reference/Procedure
0 寡ヲ D	reserved	IETF Review
E	Ethernet (and compatible)	[this document]
F	Experimental	Experimental

Table 1

## 11.3. Sub-Link Scope Multicast Addresses

IANA is requested to create a new registry called "Sub-Link Scope Multicast Addresses". The fields are as follows:

Registration Procedure Expert Review

Reference [this document]

Note These permanently assigned multicast addresses are valid over a specified scope value. Bits 16 to 95 (0-based counting) of addresses in this scope are populated with a link layer identifier and link layer specific scope information as documented in the reference document.

Fields Copied from "Link-Local Scope Multicast Addresses" registry.

### 11.3.1. Initial contents

As is done with the other multicast address space registries, entries on the "Variable Scope Multicast Addresses" registry have a corresponding "variable scope allocation" item in this registry. Therefore, this registry receives a copy of all of these items, e.g. by duplicating them from one of the existing registries. This is a large number of entries not listed here for brevity and avoiding it becoming outdated. These items have only the Address(es) column and the Description column filled in, the latter with "variable scope allocation".

Any change to the "Variable Scope Multicast Addresses" registry is requested to be accompanied by a corresponding update in this registry.

The following entries are added on top of the variable scope allocation entries:

Address(es)	Description	Reference/Procedure
FF22:::0:1	All Nodes in this Scope	[this document]
FF22:::0:2	All Routers in this Scope	[this document]

Table 2

For the above entries, the "Change Controller" is IETF, the "Last Reviewed" field is left empty, and "Date Registered" is filled appropriately for this document.

## 12. References

### 12.1. 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>>.
- [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>>.

[IPV6oETH] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/rfc/rfc2464>>.

[MCinUC] Gundavelli, S., Townsley, M., Troan, O., and W. Dec, "Address Mapping of IPv6 Multicast Packets on Ethernet", RFC 6085, DOI 10.17487/RFC6085, January 2011, <<https://www.rfc-editor.org/rfc/rfc6085>>.

[MCASTARCH] Boucadair, M. and S. Venaas, "Updates to the IPv6 Multicast Addressing Architecture", RFC 7371, DOI 10.17487/RFC7371, September 2014, <<https://www.rfc-editor.org/rfc/rfc7371>>.

[IANA-IPv6MC] IANA, "IPv6 Multicast Address Space Registry", <<https://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xhtml>>.

## 12.2. Informative References

[BASICSOCKET] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", RFC 3493, DOI 10.17487/RFC3493, February 2003, <<https://www.rfc-editor.org/rfc/rfc3493>>.

[SSMSOCKET] Thaler, D., Fenner, B., and B. Quinn, "Socket Interface Extensions for Multicast Source Filters", RFC 3678, DOI 10.17487/RFC3678, January 2004, <<https://www.rfc-editor.org/rfc/rfc3678>>.

[EVPN] Sajassi, A., Aggarwal, R., Uttaro, J., Bitar, N., Henderickx, W., and A. Isaac, "Requirements for Ethernet VPN (EVPN)", RFC 7209, DOI 10.17487/RFC7209, May 2014, <<https://www.rfc-editor.org/rfc/rfc7209>>.

[LLDP] IEEE 802, "IEEE 802.1AB-2016 婰IEEE Standard for Local and metropolitan area networks — Station and Media Access Control Connectivity Discovery", 2016.

[BGP-LLDP] Lindem, A., Patel, K., Zandi, S., Haas, J., and X. Xu, "BGP Logical Link Discovery Protocol (LLDP) Peer Discovery", Work in Progress, Internet-Draft, draft-acee-idr-lldp-peer-discovery-21, 22 December 2025, <<https://datatracker.ietf.org/doc/html/draft-acee-idr-lldp-peer-discovery-21>>.

[HOSTRT-LLDP] Filsfils, C., Camarillo, P., and D. Bernier, "Lightweight Host Routing using LLDP", Work in Progress, Internet-Draft, draft-filsfils-rtgwg-lightweight-host-routing-00, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-filsfils-rtgwg-lightweight-host-routing-00>>.

## Acknowledgements

The author(s) would like to thank the following people for their input, review, comments and discussion: Erik Auerswald, Erik Kline, and Michael Richardson.

## Implementation note

As there will be some delay until operating systems provide this functionality, it is worth pointing out that it can be emulated by using whatever lower-layer socket interface is also used by LLDP. In most cases this is an interface to Ethernet frames. The application needs then handle the extra headers, i.e. adding and removing Ethernet, IPv6, and likely UDP headers. This is likely quite doable since the values in these headers will be almost entirely static.

It also bears noting that while the send direction could also be handled on a "regular" IPv6 socket with e.g. a socket option to change the output lower-layer address, the receive direction is rather tricky to handle if not done at an early, low level, in particular if the socket option were to diverge across multiple applications. Using a dedicated address range avoids these complications.

## Revision history (TO BE REMOVED)

- \* -01: fix copypaste snafu in router group addresses; improve IANA considerations, note sockopt approach, explain SHOULDs
- \* -00: initial revision.

## Author's Address

David 'equinox' Lamparter  
NetDEF, Inc.  
San Jose,  
United States of America  
Email: equinox@diac24.net, equinox@opensourcerouting.org