

IDMR Working Group
Internet Draft
[draft-ietf-idmr-igmp-mrdisc-08.txt](#)
January 2002
Expires July 2002

S. Biswas
Nortel Networks
B. Cain
B. Haberman

IGMP Multicast Router Discovery

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts. 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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

The concept of IGMP snooping requires the ability to identify the location of multicast routers. Since IGMP snooping is not standardized, there are many mechanisms in use to identify the multicast routers. However, this scenario can lead to interoperability issues between multicast routers and layer-2 switches from different vendors.

This document introduces a general mechanism that allows for the discovery of multicast routers. By introducing these new messages, IGMP snooping devices have a uniform means of identifying multicast routers without dependency on particular routing protocols. In addition, other devices that may need to discover multicast routers can use these messages.

Table of Contents

1. Introduction.....	2
--------------------------------------	-------------------

2. Protocol Overview.....	3
3. Multicast Router Advertisement.....	4
3.1 Overview.....	4

3.2 IP Header Fields.....	4
3.3 Multicast Router Advertisement Message Format.....	4
3.4 Sending Multicast Router Advertisements.....	6
3.5 Receiving Multicast Router Advertisements.....	6
3.6 Multicast Router Advertisement Configuration Variables.....	6
4. Multicast Router Solicitation.....	7
4.1 Overview.....	8
4.2 IP Header Fields.....	8
4.3 Multicast Router Solicitation Message Format.....	8
4.4 Sending Multicast Router Solicitations.....	9
4.5 Receiving Multicast Router Solicitations.....	9
4.6 Multicast Router Solicitation Configuration Variables.....	9
5. Multicast Router Termination.....	9
5.1 Overview.....	9
5.2 IP Header Fields.....	10
5.3 Multicast Router Termination Message Format.....	10
5.4 Sending Multicast Router Termination Messages.....	10
5.5 Receiving Multicast Router Termination Messages.....	11
6. Multicast Router Discovery Protocol Constants.....	11
7. Mandatory Advertisement Options.....	11
7.1 Overview of Options.....	11
7.2 Query Interval Advertisement Option.....	11
7.3 Robustness Variable Advertisement Option.....	12
8. IPv6 Support.....	12
8.1 Router Advertisement Message.....	12
8.2 Router Solicitations.....	13
9. Security Considerations.....	13
10. IANA Considerations.....	14
11. Acknowledgements.....	14
12. References.....	14
13. Authors' Addresses.....	14
14. Full Copyright Statement.....	15

[1. Introduction](#)

Multicast router discovery messages are useful for discovering multicast capable routers. This capability is useful in a layer-2 bridging domain with "IGMP snooping" type of schemes. By listening to multicast router discovery messages, layer-2 devices can determine where to send multicast source data and IGMP Host Membership Reports [[RFC1112](#)] [[RFC2236](#)]. Multicast source data and IGMP Host Membership

Reports must be received by all multicast routers on a segment. Using IGMP Host Membership Queries to discover multicast routers is not useful because of query suppression in IGMP.

The use of the multicast router advertisement is not precluded from being used for other purposes. Extensible options have been included in the advertisement message for future enhancements.

The following are justifications for inventing another router discovery protocol:

Biswas, Cain, Haberman

2

Internet Draft

IGMP Multicast Router Discovery

January 2002

- o Using ICMP router discovery is not an appropriate solution for multicast router discovery because: 1.) It may confuse hosts listening to ICMP router advertisements; unicast and multicast topologies may not be congruent. 2.) There is no way to tell from an ICMP router advertisement if a router is running a multicast routing protocol.
- o By making multicast router discovery messages extensible, future enhancements can be made.
- o By inventing a generic IP layer message, multiple types of messages per link layer are not needed (i.e. including this functionality as part of IP is better than inventing N discovery protocols for N layer-2 technologies).

Although multicast router discovery messages could be sent as ICMP messages, IGMP was chosen because IGMP snooping switches already snoop IGMP messages and because the intended first use of these protocol messages is multicast specific.

[2. Protocol Overview](#)

IGMP Multicast Router Discovery consists of three messages for discovering multicast routers. The Multicast Router Advertisement is sent by routers to advertise IP multicast forwarding enabled on an interface. The Multicast Router Solicitation is used by routers to solicit Multicast Router Advertisements. The Multicast Router Termination message is sent when a router terminates its multicast routing functions.

Multicast routers send Multicast Router Advertisements (hereafter called advertisements) periodically on all interfaces on which multicast forwarding is enabled.

Multicast Router Advertisements are also sent in response to Multicast Router Solicitations (hereafter called solicitations). These are sent to solicit a response of Multicast Router Advertisements from all multicast routers on a subnet. Solicitations are sent to the ALL-Systems (224.0.0.1) multicast group.

Multicast Router Solicitations are sent whenever a device wishes to discover multicast routers on a directly attached subnet.

Multicast Router Termination messages are sent when a router terminates its multicast routing functions.

All IGMP Multicast Router Discovery messages are sent with an IP TTL of 1 and contain the IP Router Alert Option [[RFC2113](#)] in their IP header. All IGMP Multicast Router Discovery messages are sent with to the All-Systems multicast group (224.0.0.1).

Biswas, Cain, Haberman

3

Internet Draft

IGMP Multicast Router Discovery

January 2002

Other non-IP forwarding devices (e.g. layer-2 switches) may send Multicast Router Solicitations to solicit Multicast Router Advertisements.

[3.](#) Multicast Router Advertisement

3.1 Overview

Multicast Router Advertisements are sent periodically on all router interfaces on which multicast forwarding is enabled. They are also sent in response to Multicast Router Solicitations.

Router advertisements are sent upon expiration of a periodic timer, when a router starts up, and when a router interface (that has IP multicast forwarding enabled) initializes/restarts. Advertisements are sent as IGMP messages to the All-Systems multicast address (224.0.0.1) and should be rate-limited.

Router advertisements may contain any number of options. Two options are defined in this document and MUST be supported by any implementation of IGMP multicast router discovery. These options are described in [Section 5](#). Additional options may be defined as needed by future work.

3.2 IP Header Fields

3.2.1 Source Address

An IP address belonging to the interface from which this message is sent. If multiple source addresses are configured on an interface, then the one chosen is implementation dependent.

3.2.2 Destination Address

Router Advertisements are sent to the All-Systems multicast address (224.0.0.1).

3.2.3 Time-to-Live

The Time-to-Live field MUST be 1.

3.2.4 Protocol

The protocol field is set to IGMP (2).

3.3 Multicast Router Advertisement Message Format

0 1 2 3

Biswas, Cain, Haberman

4

Internet Draft

IGMP Multicast Router Discovery

January 2002

```

 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type      | Ad. Interval |      Checksum      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Unused      | Number of Options (N) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Option [1] (TLV encoded)      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Option [...] (TLV encoded)      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Option [N] (TLV encoded)      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

3.3.1 Type Field

The type field is set to XX (to be assigned by IANA).

3.3.2 Advertisement Interval

This specifies the periodic time interval at which Multicast Router Advertisements are sent in units of seconds. This value is set to the configured MaxAdvertisementInterval variable.

3.3.3 Checksum

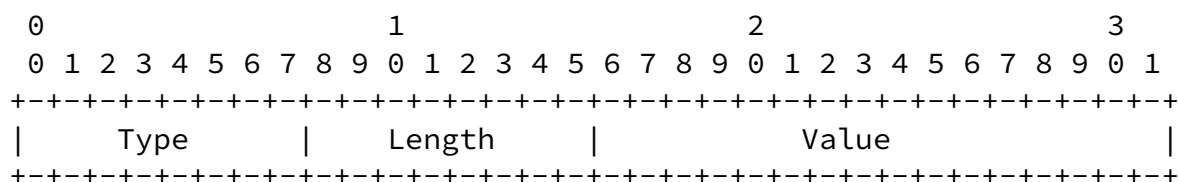
The checksum is the 16-bit one's complement of the one's complement sum of the IGMP message, starting with the IGMP type. For computing the checksum, the Checksum field is set to 0.

3.3.4 Number of Options (N)

The number of options contained in the router advertisement. If no options are sent this field MUST be set to 0.

3.3.5 Option[1..N]

Options are encoded as TLV in the following manner:



If the Number of Options field is not zero, a receiver MUST examine all options. No strict ordering of options is enforced.

Type: Set to option type being advertised

Length: Length in bytes of Value field

Value: Option dependent value

3.4 Sending Multicast Router Advertisements

Router Advertisements are sent when the following events occur:

- o When the periodic advertisement interval timer expires.
Note that it is not strictly periodic because the advertisement interval is a random number between MaxAdvertisementInterval and MinAdvertisementInterval.
- o After waiting for a random delay less than

MaxInitialAdvertisementInterval when an interface first comes up, is (re)initialized, or IGMP Multicast Router Discovery is enabled. A router may send up to a maximum of MaxInitialAdvertisements advertisements, waiting for a random delay less than MaxInitialAdvertisementInterval between each successive advertisement. Multiple messages are sent for robustness in the face of packet loss on the network.

This is to prevent an implosion of router advertisements. An example of this occurring would be when many routers are powered on at the same time. When a solicitation is received, a router advertisement is sent in response with a random delay less than MAX_RESPONSE_DELAY. If a solicitation is received while an advertisement is pending (because of a recent solicitation), that solicitation will be ignored.

Whenever an advertisement is sent, the periodic advertisement interval timer must be reset.

3.5 Receiving Multicast Router Advertisements

Upon receiving a router advertisement, devices will validate the message by the following criteria:

1. Verifying the IGMP checksum
2. IP Destination Address = All-Systems multicast address

A router advertisement not meeting the validity requirements should be silently discarded or logged in a rate-limited manner. Devices MUST process all options, discarding options that are not recognized.

If a router advertisement is not received for a particular neighbor within NeighborDeadInterval time interval, then the neighbor is considered to be unreachable.

3.6 Multicast Router Advertisement Configuration Variables

A router that implements multicast router discovery MUST allow for the following variables to be configured by system management; default values are specified so as to make it unnecessary to configure any of these variables in many cases.

For each interface the following configurable variables are defined:

3.6.1 MaxAdvertisementInterval

The maximum time allowed between sending router advertisements from the interface, in seconds. Must be no less than 2 seconds and no greater than 180 seconds.

Default: 20 seconds.

3.6.2 MinAdvertisementInterval

The minimum time allowed between sending unsolicited router advertisements from the interface, in seconds. Must be no less than 3 seconds and no greater than MaxAdvertisementInterval.

Default: $0.75 * \text{MaxAdvertisementInterval}$

3.6.3 MaxInitialAdvertisementInterval

The first router advertisement out of an interface is sent after waiting for a random interval less than this variable. This will prevent a flood of router advertisements when many routers start up at the same time.

Default: 2 seconds

3.6.4 MaxInitialAdvertisements

The maximum number of router advertisements that will be sent on a subnet after a router boots.

Default: 3

3.6.5 NeighborDeadInterval

The maximum time allowed before a neighbor can be declared "dead". This variable is defined in seconds. In order for all devices to have a consistent state, it is necessary for the MaxAdvertisementInterval to be configured the same on all routers per subnet.

Default: $3 * \text{MaxAdvertisementInterval}$

[4. Multicast Router Solicitation](#)

4.1 Overview

Multicast Router Solicitations are used to solicit Multicast Router Advertisements. These messages are used when a device wishes to discover multicast routers. Upon receiving a solicitation on an interface with IP multicast forwarding enabled and multicast router discovery enabled, a router will respond with an advertisement.

Router solicitations may be sent when a device starts up, when an interface (re)initializes, or when IGMP Multicast Router Discovery is enabled. Solicitations are sent as IGMP messages to the All-Systems multicast address (224.0.0.1) and must be rate-limited.

4.2 IP Header Fields

4.2.1 Source Address

An IP address belonging to the interface from which this message is sent. If multiple source addresses are configured on an interface, then the one chosen is implementation dependent.

If the solicitation is being sent from a device that does not have an IP address (i.e. non-managed layer-2 switch), then the source address should be set to all zeros.

4.2.2 Destination Address

Solicitation messages are sent to the All-Systems multicast address (224.0.0.1).


4.2.3 Time-to-Live

The time-to-live field MUST be 1.

4.2.4 Protocol

The protocol field is set to IGMP (2).

4.3 Multicast Router Solicitation Message Format



4.3.1 Type Field

The type field is set to YY (to be assigned by IANA).

4.3.2 Reserved Field

Biswas, Cain, Haberman

8

Internet Draft

IGMP Multicast Router Discovery

January 2002

Sent as 0; ignored on reception.

4.3.3 Checksum

The checksum is the 16-bit one's complement of the one's complement sum of the IGMP message, starting with the IGMP type. For computing the checksum, the Checksum field is set to 0.

4.4 Sending Multicast Router Solicitations

Router solicitations are sent when the following events occur:

1. After waiting for a random delay less than SOLICITATION_INTERVAL when an interface first comes up, is (re)initialized, or IGMP Multicast Router Discovery is enabled. A device may send up to a maximum of MAX_SOLICITATIONS, waiting for a random delay less than SOLICITATION_INTERVAL between each successive solicitation.
2. Optionally, for an implementation specific event. Solicitations MUST be rate-limited; no more than MAX_SOLICITATIONS MUST be sent in SOLICITATION_INTERVAL seconds.

4.5 Receiving Multicast Router Solicitations

Upon receiving a router solicitation, routers will validate the message by:

1. Verifying the IGMP checksum
2. IP Destination Address = All-Systems multicast address

A router solicitation not meeting the validity requirements should be silently discarded or logged in a rate-limited manner.

Solicitation message IP source addresses MUST NOT be used as part of the validity check.

4.6 Multicast Router Solicitation Configuration Variables

There are no configurable variables with respect to router

5.3.2 Reserved Field

Sent as 0; ignored on reception.

5.3.3 Checksum

The checksum is the 16-bit one's complement of the one's complement sum of the IGMP message, starting with the IGMP type. For computing the checksum, the Checksum field is set to 0.

5.4 Sending Multicast Router Termination Messages

Multicast Router Termination messages are sent for three reasons:

1. Multicast forwarding is disabled on the interface
2. The interface is administratively disabled

Biswas, Cain, Haberman

10

Internet Draft IGMP Multicast Router Discovery January 2002

3. The router is gracefully shutdown

5.5 Receiving Multicast Router Termination Messages

Upon receiving a termination message, routers will validate the message by the following criteria:

1. Verifying the IGMP checksum
2. IP Destination Address = All-Systems multicast address

A termination message not meeting the validity requirements should be silently discarded or logged in a rate-limited manner.

[6. Multicast Router Discovery Protocol Constants](#)

o MAX_RESPONSE_DELAY	2 seconds
o MAX_SOLICITATION_DELAY	1 second
o SOLICITATION_INTERVAL	3 seconds
o MAX_SOLICITATIONS	3 transmissions

[7. Mandatory Advertisement Options](#)

7.1 Overview of Options

The following options MUST be supported by an implementation of IGMP Multicast Router Discovery: Query Interval Advertisement Option and Robustness Variable Advertisement Option. These options advertise specific IGMP variables and are sent in an advertisement depending on the version of IGMP enabled on an interface. Although no requirements exist for multicast routers at this time, it is assumed that all multicast routers support the IGMP protocol.

7.2 Query Interval Advertisement Option

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type=1           | Length=2       |   IGMP Query Interval       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

If a multicast router has any version of IGMP [[RFC1112](#)] enabled on an interface on which IGMP Multicast Router Discovery is also enabled, it MUST send all advertisements with the Query Interval Advertisement Option. This option contains the IGMP "Query Interval" configured on the interface on which advertisements are sent.

This option is sent regardless of whether the router is currently the IGMP querier for the subnet. This option is sent regardless of what version of IGMP the router is running.

IGMP Query Interval field is equal (in seconds) to the configured IGMP "query interval" on the interface from which the advertisement was sent.

7.3 Robustness Variable Advertisement Option

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type=2           | Length=2       |   Robustness Variable       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

If a multicast router has IGMPv2 [[IGMPv2](#)] or IGMPv3 [[IGMPv3](#)] enabled on an interface on which IGMP Multicast Router Discovery is also enabled, it MUST send all advertisements with the Robustness Variable Advertisement Option. This option contains the IGMP "Robustness

Variable" configured on the interface on which advertisements are sent.

This option is sent regardless of whether the router is currently the IGMP querier for the subnet. This option may be omitted if IGMPv1 is enabled on the interface.

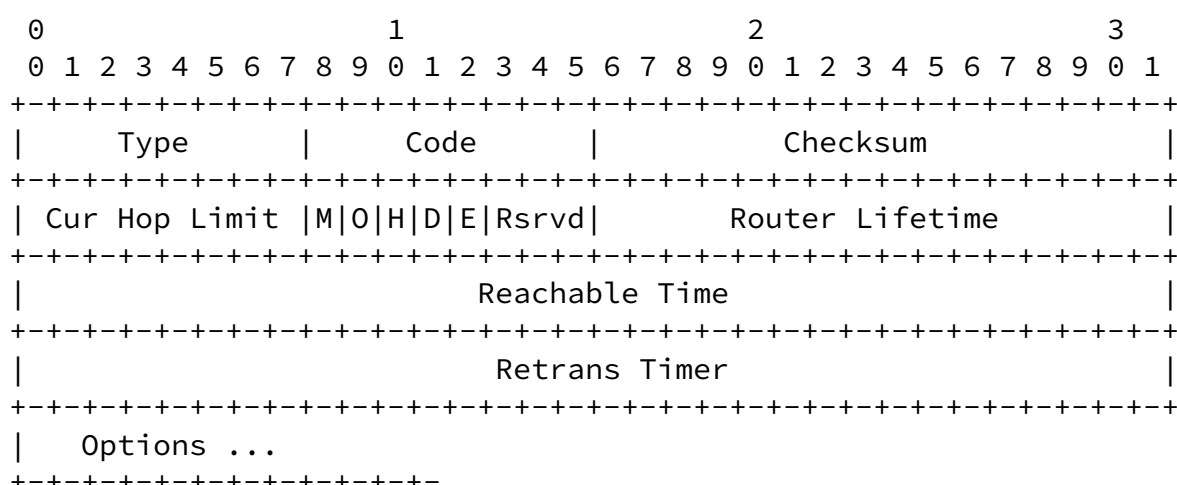
Robustness Variable is an integer that MUST not be zero [IGMPv2] and is equal to the IGMPv2 robustness variable.

8. IPv6 Support

The Multicast Router Discovery function for IPv6 is accomplished using the Neighbor Discovery Protocol for IPv6 [[RFC2461](#)] (hereafter called NDP). Specifically, the Router Advertisement message contains new fields to support the discovery of multicast routers. For this reason, the timing mechanisms defined for NDP will be used instead of those defined in this document for IPv4 support.

8.1 Router Advertisement Message

The Router Advertisement message contains two new fields to support the multicast router discovery mechanism. The modified message format is:



The two new fields are the 'D' and 'E' bits. All other fields retain their definitions and functions as described in [Section 4.2](#) of the NDP specification [[RFC2461](#)].

8.1.1 Discovery (D) bit

The 'D' bit is used by a router to indicate support for the Multicast Router Discovery protocol. A value of '1' indicates that the router supports the discovery protocol. A value of '0' indicates no support. This allows for backwards compatibility of the Router Advertisement message.

8.1.2 Enabled (E) bit

The 'E' bit indicates whether multicast routing is enabled on the router's interface. A value of '1' indicates that multicast forwarding is enabled on the router's interface. A value of '0' indicates that multicast forwarding is disabled.

When the state of multicast forwarding changes on an interface, a router must stop its Router Advertisement timer, transmit a Router Advertisement with the 'E' bit set to the value associated with the new multicast forwarding state, and restart its Router Advertisement timer.

8.2 Router Solicitations

An NDP Router Solicitation message can be sent to solicit a Router Advertisement message in order to determine the multicast forwarding state of a router. The periodic transmission of solicitation messages is outlined in [RFC 2461](#).

[9](#). Security Considerations

The Multicast Router Advertisement message may allow rogue machines to masquerade as multicast routers. This could allow those machines to eavesdrop on multicast data transmission or create a denial of

service attack on multicast flows. However, these new messages are extensible and that allows for the introduction of security associations into the protocol. These security extensions could be used to authenticate or encrypt the messages.

[10](#). IANA Considerations

This document introduces three new IGMP messages. Each of these messages requires a new IGMP 'Type' value.

11. Acknowledgements

ICMP Router Discovery [[RFC1256](#)] was used as a general model for IGMP Multicast Router Discovery.

12. References

- [RFC1256] Deering, S., "ICMP Router Discovery Messages", [RFC 1256](#), September 1991.
- [RFC1112] Deering, S., "Host Extensions for IP Multicasting", [RFC 1112](#), August 1989.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", [RFC 2236](#), November 1997.
- [IGMPv3] Cain, B., et al, "Internet Group Management Protocol, Version 3", work in progress, January 2002.
- [RFC2113] Katz, D., "IP Router Alert Option," [RFC 2113](#), April 1996.
- [RFC2461] Narten, T., Nordmark, E., and Simpson, W., "Neighbor Discovery IP Version 6 (IPv6)", [RFC 2461](#), December 1998.

13. Authors' Addresses

Shantam Biswas
Nortel Networks
600 Technology Park Drive
Billerica, MA 01821
Email: sbiswas@baynetworks.com
Phone: 1-978-916-8048

Brad Cain
Email: bcain@mediaone.net

Brian Haberman
Email: haberman@innovationslab.net
Phone: 1-919-949-4828

14. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

