Network Working Group
Request for Comments: 4286
Category: Standards Track

B. Haberman
JHU APL
J. Martin
Netzwert AG

December 2005

# Multicast Router Discovery

## Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

## Copyright Notice



## Abstract

The concept of Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) snooping requires the ability to identify the location of multicast routers. Since snooping is not standardized, there are many mechanisms in use to identify the multicast routers. However, this can lead to interoperability issues between multicast routers and snooping switches from different vendors.

This document introduces a general mechanism that allows for the discovery of multicast routers. This new mechanism, Multicast Router Discovery (MRD), introduces a standardized means of identifying multicast routers without a dependency on particular multicast routing protocols.

Table of Contents

## 1. Introduction

Multicast Router Discovery (MRD) messages are useful for determining which nodes attached to a switch have multicast routing enabled. This capability is useful in a layer-2 bridging domain with snooping switches. By utilizing MRD messages, layer-2 switches can determine where to send multicast source data and group membership messages [1] [2]. Multicast source data and group membership reports must be received by all multicast routers on a segment. Using the group membership protocol Query messages to discover multicast routers is insufficient due to query suppression.

Although MRD messages could be sent as ICMP messages, the group management protocols were chosen since this functionality is multicast specific. The addition of this functionality to the group membership protocol also allows operators to have congruence between MRD problems and data forwarding issues.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [3].

## 2. Protocol Overview

Multicast Router Discovery consists of three messages for discovering multicast routers. The Multicast Router Advertisement is sent by routers to advertise that IP multicast forwarding is enabled. Devices may send Multicast Router Solicitation messages in order to solicit Advertisement messages from multicast routers. The Multicast Router Termination messages are sent when a router stops IP multicast routing functions on an interface.

Multicast routers send unsolicited Advertisements periodically on all interfaces on which multicast forwarding is enabled. Advertisement messages are also sent in response to Solicitations. In addition to advertising the location of multicast routers, Advertisements also

convey useful information concerning group management protocol variables. This information can be used for consistency checking on the subnet.

A device sends Solicitation messages whenever it wishes to discover multicast routers on a directly attached link.

A router sends Termination messages when it terminates multicast routing functionality on an interface.

All MRD messages are sent with an IPv4 Time to Live (TTL) or IPv6 Hop Limit of 1 and contain the Router Alert Option [4] [5]. All MRD messages SHOULD be rate-limited as per the MaxMessageRate variable.

Advertisement and Termination messages are sent to the All-Snoopers multicast address.

Solicitation messages are sent to the All-Routers multicast address.

Any data beyond the fixed message format MUST be ignored.

## 3. Multicast Router Advertisement

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

Advertisements are sent

1. Upon the expiration of a periodic (modulo randomization) timer

2. As part of a router’s start-up procedure

3. During the restart of a multicast forwarding interface

4. On receipt of a Solicitation message

All Advertisements are sent as Internet Group Management Protocol (for IPv4) or Multicast Listener Discovery (for IPv6) messages to the All-Snoopers multicast address. These messages SHOULD be rate-limited as per the MaxMessageRate variable.

### 3.1. Advertisement Configuration Variables

An MRD implementation MUST support the following variables being configured by system management. Default values are specified to make it unnecessary to configure any of these variables in many cases.

#### 3.1.1. AdvertisementInterval

This variable is the base interval (in integer seconds) between the transmissions of unsolicited Advertisements on an interface. This value MUST be no less than 4 seconds and no greater than 180 seconds.

Default: 20 seconds

#### 3.1.2. AdvertisementJitter

This is the maximum time (in seconds) by which the AdvertisementInterval is perturbed for each unsolicited Advertisement. Note that the purpose of this jitter is to avoid synchronization of multiple routers on a network, hence choosing a value of zero is discouraged. This value MUST be an integer no less than 0 seconds and no greater than AdvertisementInterval.

The AdvertisementJitter MUST be 0.025*AdvertisementInterval

#### 3.1.3. MaxInitialAdvertisementInterval

The first unsolicited Advertisement transmitted on an interface is sent after waiting a random interval (in seconds) less than this variable. This prevents a flood of Advertisements when multiple routers start up at the same time.

Default: 2 seconds

#### 3.1.4. MaxInitialAdvertisements

This variable is the maximum number of unsolicited Advertisements that will be transmitted by the advertising interface when MRD starts up.

Default: 3

#### 3.1.5. NeighborDeadInterval

The NeighborDeadInterval variable is the maximum time (in seconds) allowed to elapse (after receipt of the last valid Advertisement) before a neighboring router is declared unreachable. This variable is maintained per neighbor. An MRD receiver should set the NeighborDeadInterval to 3 times the sum of Advertisement Interval Field received plus the AdvertisementJitter calculated from the received Advertisement Interval Field. This ensures consistent behavior between multiple devices on a network.

Default : 3 * (Advertisement Interval Field + calculated AdvertisementJitter)

### 3.1.6. MaxMessageRate

The MaxMessageRate variable is the maximum aggregate number of messages an MRD implementation SHOULD send (per second) per interface or per management or logging destination.

Default: 10

## 3.2. Advertisement Packet Format

The Advertisement message has the following format:

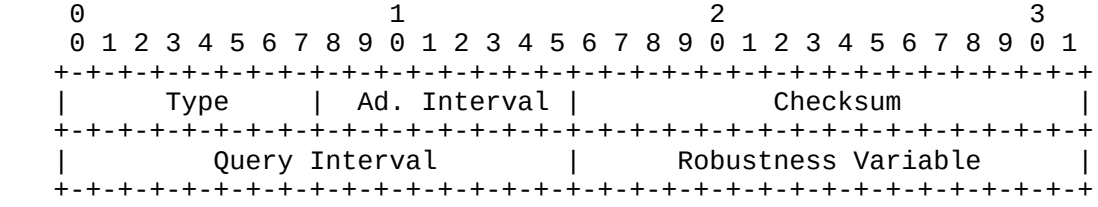


### 3.2.1. Type Field

The Type field identifies the message as an Advertisement. It is set to 0x30 for IPv4 and 151 for IPv6.

### 3.2.2. Advertisement Interval Field

This field specifies the periodic time interval at which unsolicited Advertisement messages are transmitted in units of seconds. This value is set to the configured AdvertisementInterval.

### 3.2.3. Checksum Field

The checksum field is set as follows:

1. For IPv4 it is the 16-bit one’s complement of the one’s complement sum of the IGMP message, starting with the Type field. For computing the checksum, the checksum field is set to 0.

2. For IPv6 it is ICMPv6 checksum as specified in [6].

#### 3.2.4. Query Interval Field

The Query Interval field is set to the Query Interval value (in seconds) in use by IGMP or MLD on the interface. If IGMP or MLD is not enabled on the advertising interface, this field MUST be set to 0. Note that this is the Querier's Query Interval (QQI), not the Querier's Query Interval Code (QQIC) as specified in the IGMP/MLD specifications.

#### 3.2.5. Robustness Variable Field

This field is set to the Robustness Variable in use by IGMPv2 [2], IGMPv3 [7], or MLD [8] [9] on the advertising interface. If IGMPv1 is in use or no group management protocol is enabled on the interface, this field MUST be set to 0.

### 3.3. IP Header Fields

#### 3.3.1. Source Address

The IP source address is set to an IP address configured on the advertising interface. For IPv6, a link-local address MUST be used.

#### 3.3.2. Destination Address

The IP destination address is set to the All-Snoopers multicast address.

#### 3.3.3. Time-to-Live / Hop Limit

The IPv4 TTL and IPv6 Hop Limit are set to 1.

#### 3.3.4. IPv4 Protocol

The IPv4 Protocol field is set to IGMP (2).

#### 3.3.5. IPv6 Next Header

The ICMPv6 header is identified by a Next Header value of 58 in the immediately preceding header [6].

### 3.4. Sending Multicast Router Advertisements

Advertisement messages are sent when the following events occur:

1. The expiration of the periodic advertisement interval timer. Note that this timer is not strictly periodic since the base AdvertisementInterval is varied at each interval by a random value no more than plus or minus AdvertisementJitter seconds.

2. After a random delay less than MaxInitialAdvertisementInterval when an interface is first enabled, is (re-)initialized, or MRD is enabled. A router may send up to a maximum of MaxInitialAdvertisements Advertisements, waiting for a random delay less than MaxInitialAdvertisementInterval between each successive message. Multiple Advertisements are sent for robustness in the face of packet loss on the network.

This is to prevent an implosion of Advertisements. An example of this occurring would be when many routers are powered on at the same time. When a Solicitation is received, an Advertisement is sent in response with a random delay less than MAX_RESPONSE_DELAY. If a Solicitation is received while an Advertisement is pending, that Solicitation MUST be ignored.

Changes in the Query Interval or Robustness Variable MUST NOT trigger a new Advertisement; however, the new values MUST be used in all future Advertisement messages.

When an Advertisement is sent, the periodic advertisement interval timer MUST be reset.

### 3.5. Receiving Multicast Router Advertisements

Upon receiving an Advertisement message, devices validate the message with the following criteria:

1. The checksum is correct

2. The IP destination address is equal to the All-Snoopers multicast address

3. For IPv6, the IP source address is a link-local address

An Advertisement not meeting the validity requirements MUST be silently discarded and may be logged in a rate-limited manner as per the MaxMessageRate variable.

If an Advertisement is not received for a particular neighbor within a NeighborDeadInterval time interval, then the neighbor is considered unreachable.

## 4. Multicast Router Solicitation

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

Solicitations may be sent when these occur:

1. An interface is (re-)initialized

2. MRD is enabled

Solicitations are sent to the All-Routers multicast address and SHOULD be rate-limited, as per the MaxMessageRate variable.

### 4.1. Solicitation Packet Format

The Solicitation message has the following format:

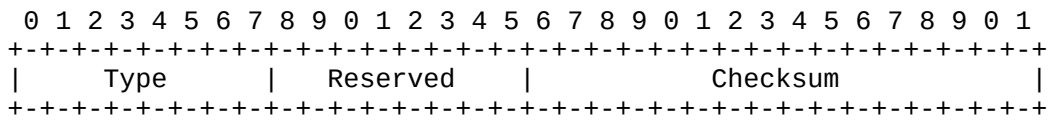


#### 4.1.1. Type Field

The Type field identifies the message as a Solicitation. It is set to 0x31 for IPv4 and 152 for IPv6.

#### 4.1.2. Reserved Field

The Reserved field is set to 0 on transmission and ignored on reception.

#### 4.1.3. Checksum Field

The checksum field is set as follows:

- For IPv4 it is the 16-bit one’s complement of the one’s complement sum of the IGMP message, starting with the Type field. For computing the checksum, the checksum field is set to 0.

- For IPv6 it is ICMPv6 checksum as specified in [6].

## 4.2. IP Header Fields

### 4.2.1. Source Address

The IP source address is set to an IP address configured on the soliciting interface. For IPv6, a link-local address MUST be used.

### 4.2.2. Destination Address

The IP destination address is set to the All-Routers multicast address.

### 4.2.3. Time-to-Live / Hop Limit

The IPv4 TTL and IPv6 Hop Limit are set to 1.

### 4.2.4. IPv4 Protocol

The IPv4 Protocol field is set to IGMP (2).

### 4.2.5. IPv6 Next Header

The ICMPv6 header is identified by a Next Header value of 58 in the immediately preceding header [6].

## 4.3. Sending Multicast Router Solicitations

Solicitation messages are sent when the following events occur:

- After waiting for a random delay less than MAX_SOLICITATION_DELAY when an interface first becomes operational, is (re-)initialized, or MRD is enabled. A device may send up to a maximum of MAX_SOLICITATIONS, waiting for a random delay less than MAX_SOLICITATION_DELAY between each solicitation.

- Optionally, for an implementation specific event.

Solicitations MUST be rate-limited as per the MaxMessageRate variable; the implementation MUST send no more than MAX_SOLICITATIONS in MAX_SOLICITATION_DELAY seconds.

## 4.4. Receiving Multicast Router Solicitations

A Solicitation message MUST be validated before a response is sent. A router MUST verify the following:

- The checksum is correct.
- The IP destination address is the All-Routers multicast address.
- For IPv6, the IP source address MUST be a link-local address.

Solicitations not meeting the validity requirements SHOULD be silently discarded and may be logged in a rate-limited manner as per the MaxMessageRate variable.

## 5. Multicast Router Termination

The Multicast Router Termination message is used to expedite the notification of a change in the status of a router's multicast forwarding functions. Multicast routers send Terminations when multicast forwarding is disabled on the advertising interface.

### 5.1. Termination Packet Format

The Termination message has the following format:

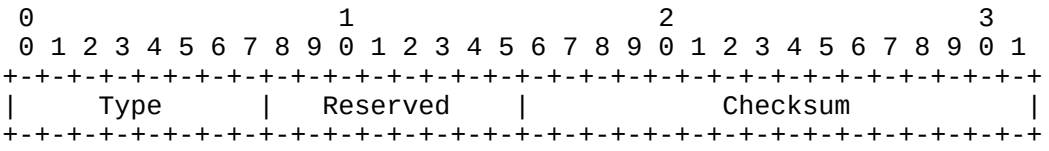


#### 5.1.1. Type Field

The Type field identifies the message as a Termination. It is set to 0x32 for IPv4 and 153 for IPv6.

#### 5.1.2. Reserved Field

The Reserved field is set to 0 on transmission and ignored on reception.

#### 5.1.3. Checksum Field

The checksum field is set as follows:

- For IPv4 it is the 16-bit one's complement of the one's complement sum of the IGMP message, starting with the Type field. For computing the checksum, the checksum field is set to 0.
- For IPv6 it is ICMPv6 checksum as specified in [6].

### 5.2. IP Header Fields

#### 5.2.1. Source Address

The IP source address is set to an IP address configured on the advertising interface. For IPv6, a link-local address MUST be used.

#### 5.2.2. Destination Address

The IP destination address is set to the All-Snoopers multicast address.

#### 5.2.3. Time-to-Live / Hop Limit

The IPv4 TTL and IPv6 Hop Limit are set to 1.

#### 5.2.4. IPv4 Protocol

The IPv4 Protocol field is set to IGMP (2).

#### 5.2.5. IPv6 Next Header

The ICMPv6 header is identified by a Next Header value of 58 in the immediately preceding header [6].

### 5.3. Sending Multicast Router Terminations

Termination messages are sent by multicast routers when

- Multicast forwarding is disabled on an interface
- An interface is administratively disabled
- The router is gracefully shut down
- MRD is disabled

The sending of Termination messages SHOULD be rate-limited as per the MaxMessageRate variable.

### 5.4. Receiving Multicast Router Terminations

Upon receiving a Termination message, devices validate the message. The validation criteria are the following:

- Checksum MUST be correct.

- IP destination address MUST equal the All-Snoopers multicast address.
- For IPv6, the IP source address MUST be a link-local address.

Termination messages not meeting the validity requirements MUST be silently discarded and may be logged in a rate-limited manner as per the MaxMessageRate variable.

If the message passes these validation steps, a Solicitation is sent. If an Advertisement is not received within NeighborDeadInterval, the sending router is removed from the list of active multicast routers.

## 6. Protocol Constants

The following list identifies constants used in the MRD protocol. These constants are used in the calculation of parameters.

- MAX_RESPONSE_DELAY 2 seconds
- MAX_SOLICITATION_DELAY 1 second
- MAX_SOLICITATIONS 3 transmissions

## 7. Security Considerations

As MRD is a link-local protocol, there is no circumstance in which it would be correct for an MRD receiver to receive MRD traffic from an off-network source. For IPv6, MRD messages MUST have a valid link-local source address. Any messages received without a valid link-local source address MUST be discarded. Similarly, for IPv4, the MRD receiver MUST determine if the source address is local to the receiving interface, and MUST discard any messages that have a non-local source. Determining what networks are local may be accomplished through configuration information or operational capabilities.

Rogue nodes may attempt to attack a network running MRD by sending spoofed Advertisement, Solicitation, or Termination messages. Each type of spoofed message can be dealt with using existing technology.

A rogue node may attempt to interrupt multicast service by sending spoofed Termination messages. As described in Section 5.4, all Termination messages are validated by sending a Solicitation message. By sending a Solicitation, the node will force the transmission of an Advertisement by an active router.

Spoofed Solicitation messages do not cause any operational harm. They may be used as a flooding mechanism to attack a multicast router. This attack can be mitigated through the rate-limiting recommendation for all MRD messages.

The Multicast Router Advertisement message may allow rogue machines to masquerade as multicast routers. This could allow those machines to eavesdrop on multicast data transmissions. Additionally, it could constitute a denial of service attack to other hosts in the same snooping domain or sharing the same device port in the presence of high-rate multicast flows.

The technology available in SEND [10] can be utilized to address spoofed Advertisement messages in IPv6 networks. IPv6 Multicast routers in an MRD-enabled network can use SEND-based link-local addresses as the IPv6 source address for MRD messages. When a switch receives an initial Advertisement, it can use the information in the SEND-based address to challenge the router to authenticate itself. It should be noted that this approach only applies to IPv6 networks.

Another solution that supports both IPv4 and IPv6 is to use IPsec in Encapsulating Security Payload (ESP) mode [11] to protect against attacks by ensuring that messages came from a system with the proper key. When using IPsec, the messages sent to the All-Snoopers address should be authenticated using ESP. Should encryption not be desired, ESP with a null encryption algorithm and a symmetric authentication algorithm, such as HMAC-SHA-1, is viable. For keying, a symmetric signature algorithm with a single manually configured key is used for routers sending Advertisements. This allows validation that the MRD message was sent by a system with the key. It should be noted that this does not prevent a system with the key from forging a message and it requires the disabling of IPsec's Replay Protection. It is the responsibility of the network administrator to ensure that the same key is present on all possible MRD participants.

## 8. IANA Considerations

This document introduces three new IGMP messages. Each of these messages requires a new IGMP Type value. The IANA has assigned three new IGMP Type values to the Multicast Router Discovery Protocol:

| IGMP Type | Section | Message Name |
|---|---|---|
| 0x30 | Section 3.2.1 | Multicast Router Advertisement |
| 0x31 | Section 4.1.1 | Multicast Router Solicitation |
| 0x32 | Section 5.1.1 | Multicast Router Termination |

This document also introduces three new MLD messages. Each of these messages requires a new ICMPv6 Type value. The IANA has assigned three new ICMPv6 Type values from the Informational range:

| ICMPv6 Type | Section | Message Name |
|---|---|---|
| 151 | Section 3.2.1 | Multicast Router Advertisement |
| 152 | Section 4.1.1 | Multicast Router Solicitation |
| 153 | Section 5.1.1 | Multicast Router Termination |

This document also requires the assignment of an All-Snoopers multicast address for IPv4. This multicast address is in the 224.0.0/24 range since it is used for link-local, control messages. The IPv4 multicast address for All-Snoopers is 224.0.0.106.

A corresponding IPv6 multicast address has also been assigned. Following the guidelines in [12], the IPv6 multicast address is a link-local in scope and has a group-ID value equal to the low-order 8 bits of the requested IPv4 multicast address. The IPv6 multicast address is FF02:0:0:0:0:0:0:6A.

## 9. Acknowledgements

Brad Cain and Shantam Biswis are the authors of the original Multicast Router Discovery proposal.

ICMP Router Discovery [13] was used as a general model for Multicast Router Discovery.

Morten Christensen, Pekka Savola, Hugh Holbrook, and Isidor Kouvelas provided helpful feedback on various versions of this document.

## 10. References

### 10.1. Normative References

[1] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.

[2] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[4] Katz, D., "IP Router Alert Option", RFC 2113, February 1997.

[5] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", RFC 2711, October 1999.

[6] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 2463, December 1998.

[7] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[8] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.

[9] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[10] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[11] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.

[12] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, August 2002.

## 10.2. Informative Reference

[13] Deering, S., "ICMP Router Discovery Messages", RFC 1256, September 1991.

## Authors' Addresses

Brian Haberman
Johns Hopkins University Applied Physics Lab
11100 Johns Hopkins Road
Laurel, MD 20723-6099
US

Phone: +1 443 778 1319
EMail: brian@innovationslab.net

Jim Martin
Netzwert AG
An den Treptowers 1
D-12435 Berlin
Germany

Phone: +49.30/5 900 80-1180
EMail: jim@netzwert.ag

## Full Copyright Statement



## Intellectual Property



## Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.