

INTERNET-DRAFT

Rolland Vida, LIP6
Luis Costa, LIP6
Remi Zara, LIP6
Serge Fdida, LIP6
Steve Deering, Cisco Systems
Bill Fenner, AT&T Labs - Research
Isidor Kouvelas, Cisco Systems
Brian Haberman, Nortel Networks

Expires August 2001

February 2001

Multicast Listener Discovery Version 2 (MLDv2) for IPv6
<[draft-vida-mld-v2-00.txt](#)>

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

This document specifies Version 2 of the Multicast Listener Discovery protocol, MLDv2. MLD is the protocol used by an IPv6 router to discover the presence of multicast listeners (that is, nodes wishing to receive multicast packets) on its directly attached links, and to discover specifically which multicast addresses are of interest to those neighboring nodes.

MLDv2 is derived from version 3 of IPv4's Internet Group Management Protocol, IGMPv3. Compared to the previous version, MLDv2 adds support for "source filtering", that is, the ability for a node to report interest in listening to packets *only* from specific source addresses, or from *all but* specific source addresses, sent to a particular multicast address. That information may be used by multicast routing protocols to avoid delivering multicast packets

INTERNET-DRAFT

MLDv2

February 2001

from specific sources to links where there are no interested listeners. When compared to IGMPv3, one important difference to note is that MLDv2 uses ICMPv6 (IP Protocol 58) message types, rather than IGMP (IP Protocol 2) message types.

The capitalized 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 [RFC 2119]. Due to the lack of italics, emphasis is indicated herein by bracketing a word or phrase in "*" characters.

Table of Contents

1.	Introduction.	4
2.	The API for requesting IP multicast reception	4
3.	Multicast listening state maintained by nodes	6
3.1.	Socket State.	6
3.2.	Interface State	6
4.	Message formats	8
4.1.	Multicast Listener Query Message.	9
4.1.1.	Code	10
4.1.2.	Checksum	11
4.1.3.	Maximum Response Code.	11
4.1.4.	Reserved	11
4.1.5.	Multicast Address.	11
4.1.6.	Resv (Reserved).	11
4.1.7.	S Flag (Suppress Router-Side Processing)	12
4.1.8.	QRV (Querier's Robustness Variable).	12
4.1.9.	QQIC (Querier's Query Interval Code)	12
4.1.10.	Number of Sources (N).	12
4.1.11.	Source Address [i]	13
4.1.12.	Additional Data.	13
4.1.13.	Query Variants	13
4.1.14.	Destination Addresses for Queries.	14
4.2.	Version 2 Multicast Listener Report Message	14
4.2.1.	Reserved	15
4.2.2.	Checksum	16
4.2.3.	Number of Mcast Address Records (M).	16
4.2.4.	Multicast Address Record	16
4.2.5.	Record Type.	16
4.2.6.	Aux Data Len	16
4.2.7.	Number of Sources (N).	16
4.2.8.	Multicast Address.	16
4.2.9.	Source Address [i]	16
4.2.10.	Auxiliary Data	17
4.2.11.	Additional Data.	17
4.2.12.	Multicast Address Record Types	17

4.2.13.	Destination Addresses for Reports.	19
4.2.14.	Notation for Multicast Address Records	19
4.2.15.	Multicast Listener Report Size	19

5.	Description of the protocol for multicast address listeners	20
5.1.	Action on Change of Interface State	21
5.2.	Action on Reception of a Query.	23
6.	Description of the protocol for multicast routers	25
6.1.	Conditions for MLD Queries.	26
6.2.	MLD State Maintained by Multicast Routers	27
6.2.1.	Definition of Router Filter Mode.	27
6.2.2.	Definition of Multicast Address Timers.	28
6.2.3.	Definition of Source Timers	29
6.3.	MLDv2 Source Specific Forwarding Rules.	30
6.4.	Action on Reception of Reports.	31
6.4.1.	Reception of Current State Records.	31
6.4.2.	Reception of Filter Mode Change and Source List Change Records.	32
6.5.	Switching Router Filter Modes	34
6.6.	Action on Reception of Queries.	34
6.6.1.	Timer Updates	34
6.6.2.	Querier Election.	35
6.6.3.	Building and Sending Specific Queries	35
6.6.3.1.	Building and Sending Multicast Address Specific Queries.	35
6.6.3.2.	Building and Sending Multicast Address and Source Specific Queries	35
7.	Interoperation with older versions of MLD	36
7.1.	Query Version Distinctions.	36
7.2.	Multicast Address Listener Behavior	36
7.2.1.	In the Presence of Older Version Queriers	36
7.2.2.	In the Presence of Older Version Multicast Address Listeners	37
7.3.	Multicast Router Behavior	37
7.3.1.	In the Presence of Older Version Queriers	37
7.3.2.	In the Presence of Older Version Multicast Address Listeners	38
8.	List of timers, counters and their default values	39
8.1.	Robustness Variable	39
8.2.	Query Interval.	39
8.3.	Query Response Interval	39
8.4.	Multicast Address Listener Interval	39
8.5.	Other Querier Present Interval.	40
8.6.	Startup Query Interval.	40
8.7.	Startup Query Count	40
8.8.	Last Listener Query Interval.	40
8.9.	Last Listener Query Count	40
8.10.	Unsolicited Report Interval	41
8.11.	Older Version Querier Present Timeout	41
8.12.	Older Host Present Interval	41

8.13.	Configuring timers.	41
8.13.1.	Robustness Variable.	41
8.13.2.	Query Interval	42
8.13.3.	Maximum Response Delay	42
9.	Security considerations.	42
9.1.	Query Message	43

Vida et al.

[Page 3]

INTERNET-DRAFT

MLDv2

February 2001

9.2.	Current State Report Message.	44
9.3.	State Change Report Message	44
10.	IANA considerations	44
11.	Acknowledgements.	45
12.	References.	45
Appendix A.	Design rationale	46
A.1	The Need for State Change Messages.	46
A.2	Host Suppression.	46
A.3	Switching router filter modes from EXCLUDE to INCLUDE	47
Author's addresses.	47

1. INTRODUCTION

The Multicast Listener Discovery Protocol (MLD) is used by IPv6 routers to discover the presence of multicast listeners (that is, nodes wishing to receive multicast packets) on their directly attached links, and to discover specifically which multicast addresses are of interest to those neighboring nodes. Note that a multicast router may itself be a listener of one or more multicast addresses, in which case it performs both the "multicast router part" of the protocol (to collect the multicast listener information needed by its multicast routing protocol) and the "multicast address listener part" of the protocol (to inform itself and other neighboring multicast routers of its listening state).

This document specifies Version 2 of MLD. The previous version of MLD [[MLDv1](#)] became an Internet Standard and is specified in [RFC 2710](#). In this document we will refer to it as MLDv1.

Version 2 of the MLD protocol, when compared to the previous version, adds support for "source filtering", that is, the ability for a node to report interest in listening to packets *only* from specific source addresses, or from *all but* specific source addresses, sent to a particular multicast address. Version 2 is designed to be interoperable with the previous version.

2. THE API FOR REQUESTING IP MULTICAST RECEPTION

Within an IP system, there is (at least conceptually) an Application Programming Interface or API used by upper-layer protocols or application programs to ask the IP layer to enable and disable

reception of packets sent to specific IP multicast addresses. In order to take full advantage of the capabilities of MLDv2, a node's IP API must support the following operation (or any logical equivalent; for example, see [[FILTER-API](#)]):

```
IPv6MulticastListen ( socket, interface, IPv6 multicast address,  
                      filter mode, source list )
```

where:

Vida et al.

[Page 4]

INTERNET-DRAFT

MLDv2

February 2001

"socket" is an implementation-specific parameter used to distinguish among different requesting entities (e.g., programs or processes) within the node; the socket parameter of BSD Unix system calls is a specific example.

"interface" is a local identifier of the network interface on which reception of the specified multicast address is to be enabled or disabled. Interfaces may be physical (e.g., an Ethernet interface) or virtual (e.g., the endpoint of a Frame Relay virtual circuit or the endpoint of an IP-in-IP "tunnel"). An implementation may allow a special "unspecified" value to be passed as the interface parameter, in which case the request would apply to the "primary" or "default" interface of the node (perhaps established by system configuration). If reception of the same multicast address is desired on more than one interface, IPv6MulticastListen is invoked separately for each desired interface.

"IPv6 multicast address" is the multicast address to which the request pertains. If reception of more than one multicast address on a given interface is desired, IPv6MulticastListen is invoked separately for each desired address.

"filter mode" may be either INCLUDE or EXCLUDE. In INCLUDE mode, reception of packets sent to the specified multicast address is requested *only* from those source addresses listed in the source list parameter. In EXCLUDE mode, reception of packets sent to the given multicast address is requested from all source addresses *except* those listed in the source list parameter.

"source list" is an unordered list of zero or more unicast addresses from which multicast reception is desired or not desired, depending on the filter mode. An implementation MAY impose a limit on the size of source lists, but that limit MUST NOT be less than 64 addresses per list. When an operation causes the source list size limit to be exceeded, the API MUST return an error.

For a given combination of socket, interface, and IPv6 multicast address, only a single filter mode and source list can be in effect at any one time. However, either the filter mode or the source list, or both, may be changed by subsequent IPv6MulticastListen requests

that specify the same socket, interface, and IPv6 multicast address.

The previous version of MLD did not support source filters and had a simpler API consisting of Start Listening and Stop Listening operations to enable and disable listening to a given multicast address (from *all* sources) on a given interface. Those Start and Stop Listening operations are supported by the new API as follows:

The Start Listening operation is equivalent to:

```
IPv6MulticastListen ( socket, interface, IPv6 multicast address,  
                      EXCLUDE, {} )
```

Vida et al.

[Page 5]

INTERNET-DRAFT

MLDv2

February 2001

and the Stop Listening operation is equivalent to:

```
IPv6MulticastListen ( socket, interface, IPv6 multicast address,  
                      INCLUDE, {} )
```

where {} is an empty source list.

It is recommended that implementations continue to support the old API, (perhaps as calls on the new API) for compatibility with pre-existing IPv6 multicast applications.

3. MULTICAST LISTENING STATE MAINTAINED BY NODES

3.1. Socket State

For each socket on which IPv6MulticastListen has been invoked, the node records the desired multicast listening state for that socket. That state conceptually consists of a set of records of the form:

```
(interface, IPv6 multicast address, filter mode, source list)
```

The socket state evolves in response to each invocation of IPv6MulticastListen on the socket, as follows:

- o If the requested filter mode is INCLUDE *and* the requested source list is empty, then the entry corresponding to the requested interface and multicast address is deleted if present. If no such entry is present, the request is ignored.
- o If the requested filter mode is EXCLUDE *or* the requested source list is non-empty, then the entry corresponding to the requested interface and multicast address, if present, is changed to contain the requested filter mode and source list. If no such entry is present, a new entry is created, using the parameters specified in the request.

3.2. Interface State

In addition to the per-socket multicast listening state, a node must also maintain or compute multicast listening state for each of its interfaces. That state conceptually consists of a set of records of the form:

(IPv6 multicast address, filter mode, source list)

This per-interface state is derived from the per-socket state, but may differ from the per-socket state when different sockets have differing filter modes and/or source lists for the same multicast

Vida et al.

[Page 6]

INTERNET-DRAFT

MLDv2

February 2001

address and interface. For example, suppose one application or process invokes the following operation on socket s1:

IPv6MulticastListen (s1, i, m, INCLUDE, {a, b, c})

requesting reception on interface i of packets sent to multicast address m, *only* if they come from source a, b, or c. Suppose another application or process invokes the following operation on socket s2:

IPv6MulticastListen (s2, i, m, INCLUDE, {b, c, d})

requesting reception on the same interface i of packets sent to the same multicast address m, *only* if they come from sources b, c, or d. In order to satisfy the reception requirements of both sockets, it is necessary for interface i to receive packets sent to m from any one of the sources a, b, c, or d. Thus, in this example, the listening state of interface i for multicast address m has filter mode INCLUDE and source list {a, b, c, d}.

After a multicast packet has been accepted from an interface by the IP layer, its subsequent delivery to the application or process listening on a particular socket depends on the multicast listening state of that socket [and possibly also on other conditions, such as what transport-layer port the socket is bound to]. So, in the above example, if a packet arrives on interface i, destined to multicast address m, with source address a, it may be delivered on socket s1 but not on socket s2. Note that MLDv2 messages are not subject to source filtering and must always be processed by hosts and routers.

Filtering of packets based upon a socket's multicast reception state is a new feature of this API. The previous API described no filtering based upon multicast listening state; rather, a Start Listening operation on a socket simply caused the node to start to listen to a multicast address on the given interface, and packets

sent to that multicast address could be delivered to all sockets whether they had started to listen or not.

The general rules for deriving the per-interface state from the per-socket state are as follows: for each distinct (interface, IPv6 multicast address) pair that appears in any socket state, a per-interface record is created for that multicast address on that interface. Considering all socket records containing the same (interface, IPv6 multicast address) pair,

- o if **any** such record has a filter mode of EXCLUDE, then the filter mode of the interface record is EXCLUDE, and the source list of the interface record is the intersection of the source lists of all socket records in EXCLUDE mode, minus those source addresses that appear in any socket record in INCLUDE mode. For example, if the socket records for multicast address m on interface i are:

Vida et al.

[Page 7]

INTERNET-DRAFT

MLDv2

February 2001

```
from socket s1: ( i, m, EXCLUDE, {a, b, c, d} )
from socket s2: ( i, m, EXCLUDE, {b, c, d, e} )
from socket s3: ( i, m, INCLUDE, {d, e, f} )
```

then the corresponding interface record on interface i is:

```
( m, EXCLUDE, {b, c} )
```

- o if **all** such records have a filter mode of INCLUDE, then the filter mode of the interface record is INCLUDE, and the source list of the interface record is the union of the source lists of all the socket records. For example, if the socket records for multicast address m on interface i are:

```
from socket s1: ( i, m, INCLUDE, {a, b, c} )
from socket s2: ( i, m, INCLUDE, {b, c, d} )
from socket s3: ( i, m, INCLUDE, {e, f} )
```

then the corresponding interface record on interface i is:

```
( m, INCLUDE, {a, b, c, d, e, f} )
```

If system resource limits are reached when an interface state source list is calculated, an error MUST be returned to the application which requested the operation.

The above rules for deriving the interface state are (re-)evaluated whenever an IPv6MulticastListen invocation modifies the socket state by adding, deleting, or modifying a per-socket state record. Note that a change of socket state does not necessarily result in a change of interface state.

4. MESSAGE FORMATS

MLDv2 is a sub-protocol of ICMPv6, that is, MLDv2 message types are a subset of the set of ICMPv6 messages, and MLDv2 messages are identified in IPv6 packets by a preceding Next Header value of 58. All MLDv2 messages described in this document are sent with a link-local IPv6 Source Address, an IPv6 Hop Limit of 1, and an IPv6 Router Alert option [[IPv6-ALERT](#)] in a Hop-by-Hop Options header. (The Router Alert option is necessary to cause routers to examine MLDv2 messages sent to IPv6 multicast addresses in which the routers themselves have no interest.)

There are two MLD message types of concern to the MLDv2 protocol described in this document:

- o Multicast Listener Query (Type = decimal 130)
- o Version 2 Multicast Listener Report (Type = decimal 136)

Vida et al.

[Page 8]

INTERNET-DRAFT

MLDv2

February 2001

An implementation of MLDv2 must also support the following two message types, for interoperation with the previous version of MLD (see [section 7](#)):

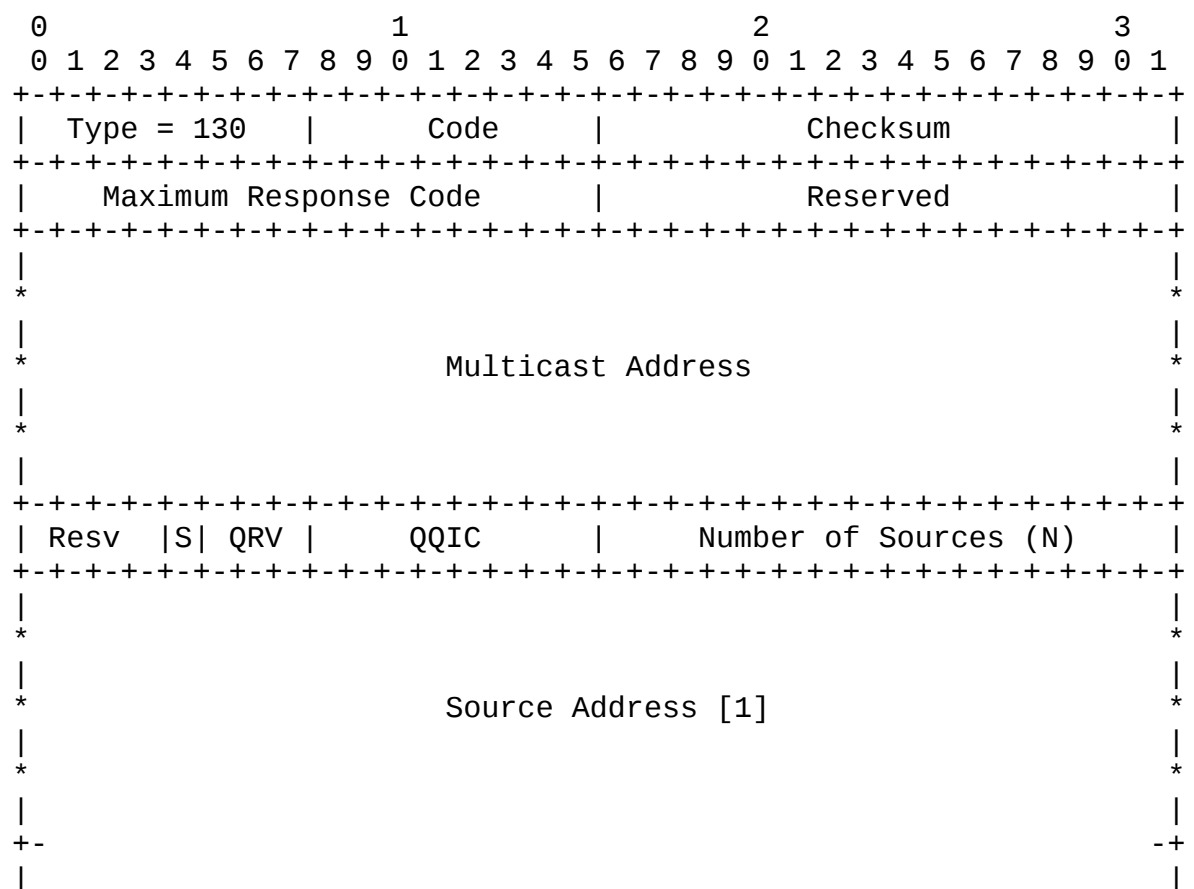
- o Version 1 Multicast Listener Report (Type = decimal 131) [[MLDv1](#)]
- o Version 1 Multicast Listener Done (Type = decimal 132) [[MLDv1](#)]

Unrecognized message types MUST be silently ignored. Other message types may be used by newer versions or extensions of MLD, by multicast routing protocols, or for other uses.

In this document, unless otherwise qualified, the capitalized words "Query" and "Report" refer to MLD Multicast Listener Queries and MLD Version 2 Multicast Listener Reports, respectively.

4.1. Multicast Listener Query Message

Multicast Listener Queries are sent by multicast routers to query the multicast listening state of neighboring interfaces. Queries have the following format:




```

|1| exp |          mant          |
+--+--+--+--+--+--+--+--+--+--+

```

Maximum Response Delay = (mant | 0x1000) << (exp+3)

Small values of Maximum Response Delay allow MLDv2 routers to tune the "leave latency" (the time between the moment the last node on a link ceases listening to a specific multicast address and the moment the routing protocol is notified that there are no more listeners for that address). Larger values, especially in the exponential range, allow tuning of the burstiness of MLD traffic on a link.

[4.1.4.](#) Reserved

Initialized to zero by the sender; ignored by receivers.

[4.1.5.](#) Multicast Address

The Multicast Address field is set to zero when sending a General Query, and set to the multicast address being queried when sending a Multicast Address Specific Query or Multicast Address and Source Specific Query (see [section 4.1.10](#), below).

[4.1.6.](#) Resv (Reserved)

Initialized to zero by the sender; ignored by receivers.

Vida et al.

[Page 11]

[4.1.7.](#) S Flag (Suppress Router-Side Processing)

When set to one, the S Flag indicates to any receiving multicast routers that they have to suppress the normal timer updates they perform upon hearing a Query. It does not, however, suppress the querier election or the normal "host-side" processing of a Query that a router may be required to perform as a consequence of itself being a multicast listener.

[4.1.8.](#) QRV (Querier's Robustness Variable)

If non-zero, the QRV field contains the [Robustness Variable] value used by the Querier, i.e., the sender of the Query. If the Querier's [Robustness Variable] exceeds 7, the maximum value of the QRV field, the QRV is set to zero. Routers adopt the QRV value from the most recently received Query as their own [Robustness Variable] value, unless that most recently received QRV was zero, in which case the receivers use the default [Robustness Variable] value specified in [section 8.1](#) or a statically configured value.

[4.1.9.](#) QQIC (Querier's Query Interval Code)

The Querier's Query Interval Code field specifies the [Query Interval] used by the querier. The actual interval, called the Querier's Query Interval (QQI), is represented in units of seconds and is derived from the Querier's Query Interval Code as follows:

If $QQIC < 128$, $QQI = QQIC$

If $QQIC \geq 128$, QQIC represents a floating-point value as follows:

```

  0 1 2 3 4 5 6 7
+-+--+--+--+--+--+
|1| exp | mant  |
+-+--+--+--+--+--+
```

$QQI = (mant \mid 0x10) \ll (exp + 3)$

Multicast routers that are not the current querier adopt the QQI value from the most recently received Query as their own [Query Interval] value, unless that most recently received QQI was zero, in which case the receiving routers use the default [Query Interval] value specified in [section 8.2](#).

[4.1.10.](#) Number of Sources (N)

The Number of Sources (N) field specifies how many source addresses are present in the Query. This number is zero in a General Query or a Multicast Address Specific Query, and non-zero in a Multicast

Vida et al.

[Page 12]

INTERNET-DRAFT

MLDv2

February 2001

Address and Source Specific Query. This number is limited by the MTU of the link over which the Query is transmitted. For example, on an Ethernet with an MTU of 1500 octets, the IPv6 header (40 octets) including the Router Alert option in the Hop-By-Hop Extension Header (8 octets) consumes 48 octets, and the MLD fields up to including the Number of Sources (N) field consume 28 octets, leaving 1424 octets for source addresses, which limits the number of source addresses to 89 (1424/16).

[4.1.11.](#) Source Address [i]

The Source Address [i] fields are a vector of n unicast addresses, where n is the value in the Number of Sources (N) field.

[4.1.12.](#) Additional Data

If the Payload Length field in the IPv6 header of a received Query indicates that there are additional octets of data present, beyond the fields described here, MLDv2 implementations MUST include those octets in the computation to verify the received MLD Checksum, but MUST otherwise ignore those additional octets. When sending a Query, an MLDv2 implementation MUST NOT include additional octets beyond the fields described here.

4.1.13. Query Variants

There are three variants of the Query message:

- o A "General Query" is sent by a multicast router to learn which multicast addresses have listeners on an attached link. In a General Query, both the Multicast Address field and the Number of Sources (N) field are zero.
- o A "Multicast Address Specific Query" is sent by a multicast router to learn if a particular multicast address has any listeners on an attached link. In a Multicast Address Specific Query, the Multicast Address field contains the multicast address of interest, and the Number of Sources (N) field contains zero.
- o A "Multicast Address and Source Specific Query" is sent by a multicast router to learn if any of the sources from the specified list for the particular multicast address has any listeners on an attached link or not. In a Multicast Address and Source Specific Query, the Multicast Address field contains the multicast address of interest, and the Source Address [i] field(s) contain the source address(es) of interest.

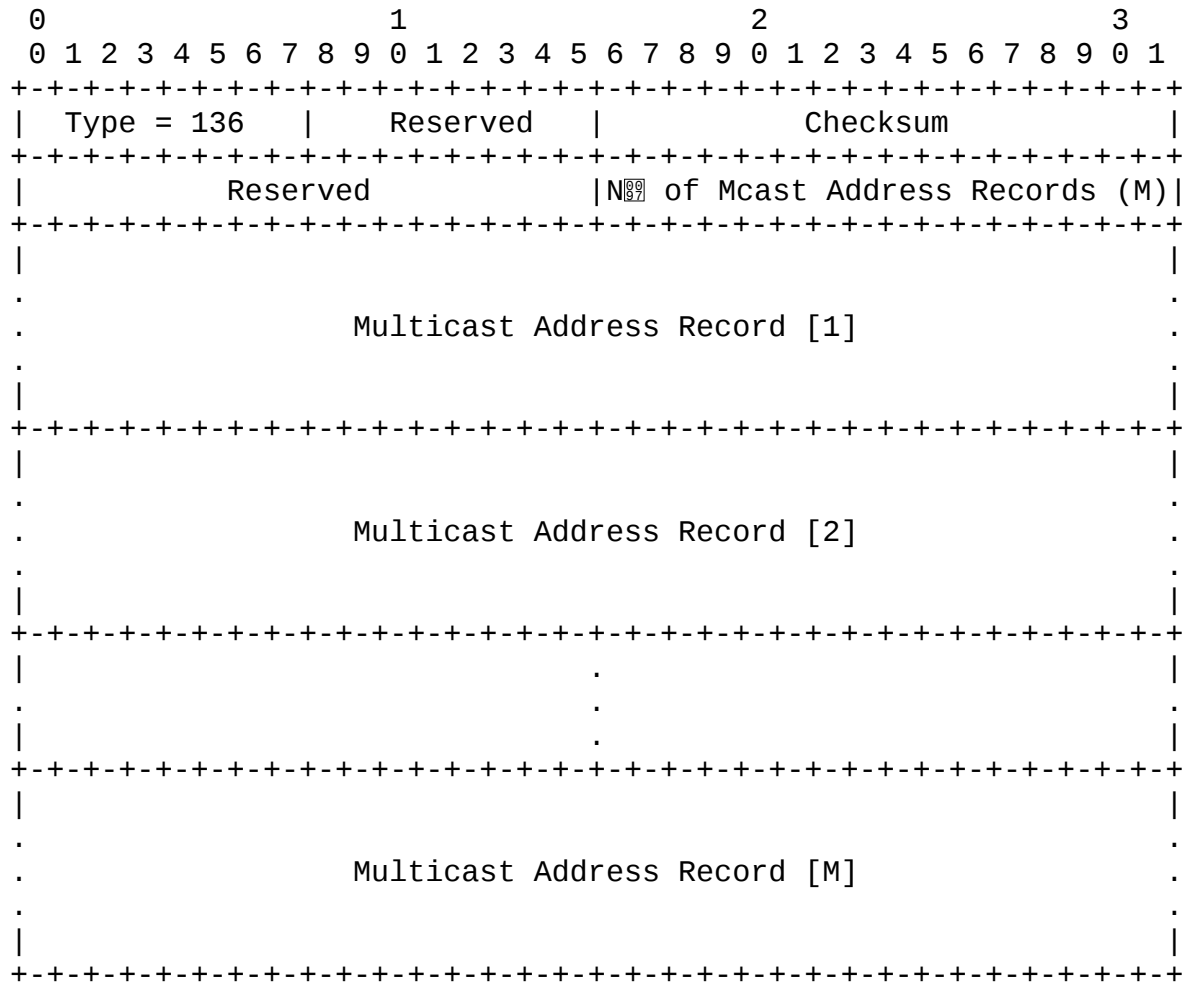
4.1.14. Destination Addresses for Queries

In MLDv2, General Queries are sent to the link-scope all-nodes multicast address (FF02::1). Multicast Address Specific and Multicast Address and Source Specific Queries are sent with an IP destination address equal to the multicast address of interest. *However*, a node MUST accept and process any Query whose IP Destination Address field contains *any* of the addresses (unicast or multicast) assigned to the interface on which the Query arrives.

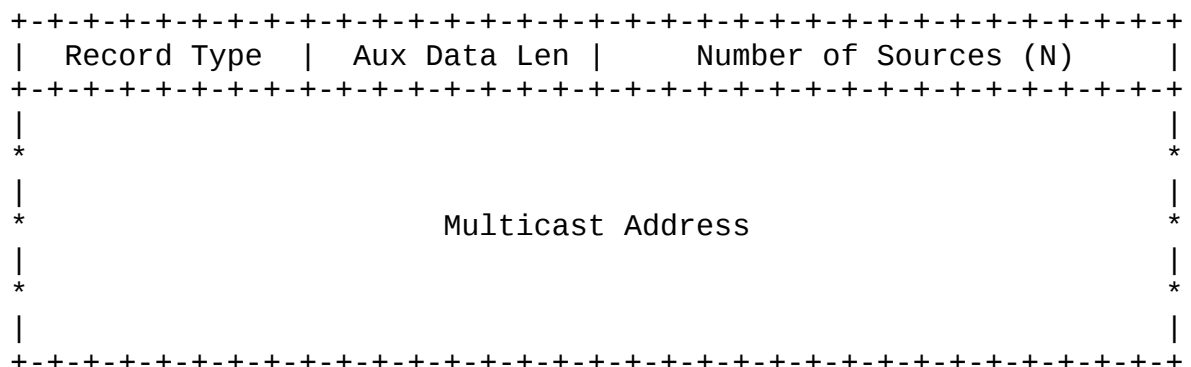
4.2. Version 2 Multicast Listener Report Message

Version 2 Multicast Listener Reports are sent by IP nodes to report

(to neighboring routers) the current multicast listening state, or changes in the multicast listening state, of their interfaces. Reports have the following format:



where each Multicast Address Record has the following internal format:



[4.2.3.](#) N^{of} of Mcast Address Records (M)

The N^{of} of Mcast Address Records (M) field specifies how many Multicast Address Records are present in this Report.

[4.2.4.](#) Multicast Address Record

Each Multicast Address Record is a block of fields containing information on the sender listening to a single multicast address on the interface from which the Report is sent.

[4.2.5.](#) Record Type

See [section 4.2.12](#), below.

[4.2.6.](#) Aux Data Len

The Aux Data Len field contains the length of the Auxiliary Data Field in this Multicast Address Record, in units of 32-bit words. It may contain zero, to indicate the absence of any auxiliary data.

[4.2.7.](#) Number of Sources (N)

The Number of Sources (N) field specifies how many source addresses are present in this Multicast Address Record.

[4.2.8.](#) Multicast Address

The Multicast Address field contains the multicast address to which this Multicast Address Record pertains.

[4.2.9.](#) Source Address [i]

The Source Address [i] fields are a vector of n unicast addresses, where n is the value in this record's Number of Sources (N) field.

[4.2.10.](#) Auxiliary Data

The Auxiliary Data field, if present, contains additional information pertaining to this Multicast Address Record. The protocol specified in this document, MLDv2, does not define any auxiliary data.

Therefore, implementations of MLDv2 MUST NOT include any auxiliary data (i.e., MUST set the Aux Data Len field to zero) in any transmitted Multicast Address Record, and MUST ignore any auxiliary data present in any received Multicast Address Record. The semantics and internal encoding of the Auxiliary Data field are to be defined by any future version or extension of MLD that uses this field.

4.2.11. Additional Data

If the Payload Length field in the IPv6 header of a received Report indicates that there are additional octets of data present, beyond the last Multicast Address Record, MLDv2 implementations MUST include those octets in the computation to verify the received MLD Checksum, but MUST otherwise ignore those additional octets. When sending a Report, an MLDv2 implementation MUST NOT include additional octets beyond the last Multicast Address Record.

4.2.12. Multicast Address Record Types

There are a number of different types of Multicast Address Records that may be included in a Report message:

- o A "Current State Record" is sent by a node in response to a Query received on an interface. It reports the current listening state of that interface, with respect to a single multicast address. The Record Type of a Current State Record may be one of the following two values:

Value	Name and Meaning
-------	------------------

-----	-----
-------	-------

- | | |
|---|--|
| 1 | MODE_IS_INCLUDE - indicates that the interface has a filter mode of INCLUDE for the specified multicast address. The Source Address [i] fields in this Multicast Address Record contain the interface's source list for the specified multicast address, if it is non-empty. |
| 2 | MODE_IS_EXCLUDE - indicates that the interface has a filter mode of EXCLUDE for the specified multicast address. The Source Address [i] fields in this Multicast Address Record contain the interface's source list for the specified multicast address, if it is non-empty. |

- o A "Filter Mode Change Record" is sent by a node whenever a local invocation of IPv6MulticastListen causes a change of the filter mode (i.e., a change from INCLUDE to EXCLUDE, or from EXCLUDE to

INCLUDE) of the interface-level state entry for a particular multicast address. The Record is included in a Report sent from the

interface on which the change occurred. The Record Type of a Filter Mode Change Record may be one of the following two values:

- 3 `CHANGE_TO_INCLUDE_MODE` - indicates that the interface has changed to `INCLUDE` filter mode for the specified multicast address. The Source Address [i] fields in this Multicast Address Record contain the interface's new source list for the specified multicast address, if it is non-empty.
 - 4 `CHANGE_TO_EXCLUDE_MODE` - indicates that the interface has changed to `EXCLUDE` filter mode for the specified multicast address. The Source Address [i] fields in this Multicast Address Record contain the interface's new source list for the specified multicast address, if it is non-empty.
- o A "Source List Change Record" is sent by a node whenever a local invocation of `IPv6MulticastListen` causes a change of source list that is *not* coincident with a change of filter mode, of the interface level state entry for a particular multicast address. The Record is included in a Report sent from the interface on which the change occurred. The Record Type of a Source List Change Record may be one of the following two values:
- 5 `ALLOW_NEW_SOURCES` - indicates that the Source Address [i] fields in this Multicast Address Record contain a list of the additional sources that the node wishes to listen to, for packets sent to the specified multicast address. If the change was to an `INCLUDE` source list, these are the addresses that were added to the list; if the change was to an `EXCLUDE` source list, these are the addresses that were deleted from the list.
 - 6 `BLOCK_OLD_SOURCES` - indicates that the Source Address [i] fields in this Multicast Address Record contain a list of the sources that the node no longer wishes to listen to, for packets sent to the specified multicast address. If the change was to an `INCLUDE` source list, these are the addresses that were deleted from the list; if the change was to an `EXCLUDE` source list, these are the addresses that were added to the list.

If a change of source list results in both allowing new sources and blocking old sources, then two Multicast Address Records are sent for the same multicast address, one of type `ALLOW_NEW_SOURCES` and one of type `BLOCK_OLD_SOURCES`.

We use the term "State Change Record" to refer to either a Filter Mode Change Record or a Source List Change Record.

Unrecognized Record Type values MUST be silently ignored.

4.2.13. Destination Addresses for Reports

Version 2 Multicast Listener Reports are sent with an IP destination address to which all MLDv2-capable multicast routers listen (see [section 10](#) for IANA considerations related to this special destination address). A node that is operating in version 1 compatibility mode sends version 1 Reports to the multicast address specified in the Multicast Address field of the Report. In addition, a node MUST accept and process any version 1 Report whose IP Destination Address field contains *any* of the IPv6 addresses (unicast or multicast) assigned to the interface on which the Report arrives.

4.2.14. Notation for Multicast Address Records

In the rest of this document, we use the following notation to describe the contents of a Multicast Address Record pertaining to a particular multicast address:

```
IS_IN ( x ) - Type MODE_IS_INCLUDE, source addresses x
IS_EX ( x ) - Type MODE_IS_EXCLUDE, source addresses x
TO_IN ( x ) - Type CHANGE_TO_INCLUDE_MODE, source addresses x
TO_EX ( x ) - Type CHANGE_TO_EXCLUDE_MODE, source addresses x
ALLOW ( x ) - Type ALLOW_NEW_SOURCES, source addresses x
BLOCK ( x ) - Type BLOCK_OLD_SOURCES, source addresses x
```

where x is either:

- o a capital letter (e.g., "A") to represent the set of source addresses,

or

- o a set expression (e.g., "A+B"), where "A+B" means the union of sets A and B, "A*B" means the intersection of sets A and B, and "A-B" means the removal of all elements of set B from set A.

4.2.15. Multicast Listener Report Size

If the set of Multicast Address Records required in a Report does not fit within the size limit of a single Report message (as determined by the MTU of the link on which it will be sent), the Multicast Address Records are sent in as many Report messages as needed to report the entire set.

If a single Multicast Address Record contains so many source addresses that it does not fit within the size limit of a single Report message, if its Type is not MODE_IS_EXCLUDE or CHANGE_TO_EXCLUDE_MODE, it is split into multiple Multicast Address Records, each containing a different subset of the source addresses and each sent in a separate Report message. If its Type is

MODE_IS_EXCLUDE or CHANGE_TO_EXCLUDE_MODE, a single Multicast Address Record is sent, containing as many source addresses as can fit, and the remaining source addresses are not reported; though the choice of which sources to report is arbitrary, it is preferable to report the same set of sources in each subsequent report, rather than reporting different sources each time.

5. DESCRIPTION OF THE PROTOCOL FOR MULTICAST ADDRESS LISTENERS

MLD is an asymmetric protocol, specifying separate behaviors for multicast address listeners -- that is, hosts or routers that listen to multicast packets -- and multicast routers. This section describes the part of MLDv2 that applies to all multicast address listeners. (Note that a multicast router that is also a multicast address listener performs both parts of MLDv2, receiving and responding to its own MLD message transmissions as well as to those of its neighbors.) The multicast router part of MLDv2 is described in [section 6](#).

A node performs the protocol described in this section over all interfaces on which multicast reception is supported, even if more than one of those interfaces is connected to the same link.

For interoperability with multicast routers running the older version of MLD, nodes maintain a MulticastRouterVersion variable for each interface on which multicast reception is supported. This section describes the behavior of multicast address listener nodes on interfaces for which MulticastRouterVersion = 2. The algorithm for determining MulticastRouterVersion, and the behavior for version 1, are described in [section 7](#).

The link-scope all-nodes multicast address, (FF02::1), is handled as a special case. On all nodes -- that is all hosts and routers, including multicast routers -- listening to packets destined to the all-nodes multicast address, from all sources, is permanently enabled on all interfaces on which multicast listening is supported. No MLD messages are ever sent regarding the all-nodes multicast address.

There are two types of events that trigger MLDv2 protocol actions on an interface:

- o a change of the interface listening state, caused by a local invocation of IPv6MulticastListen.
- o reception of a Query.

(Received MLD messages of types other than Query are silently ignored, except as required for interoperation with the earlier version of MLD.)

The following subsections describe the actions to be taken for each

INTERNET-DRAFT

MLDv2

February 2001

case. Timer and counter names appear in square brackets. Default values for those timers and counters are specified in [section 8](#).

[5.1](#). Action on Change of Interface State

An invocation of IPv6MulticastListen may cause the multicast listening state of an interface to change, according to the rules in [section 3.2](#). Each such change affects the per-interface entry for a single multicast address.

A change of interface state causes the node to immediately transmit a State Change Report from that interface. The type and contents of the Multicast Address Record(s) in that Report are determined by comparing the filter mode and source list for the affected multicast address before and after the change, according to the table below. If no interface state existed for that multicast address before the change (i.e., the change consisted of creating a new per-interface record), or if no state exists after the change (i.e., the change consisted of deleting a per-interface record), then the "non-existent" state is considered to have a filter mode of INCLUDE and an empty source list.

Old State -----	New State -----	State Change Record Sent -----
INCLUDE (A)	INCLUDE (B)	ALLOW (B-A), BLOCK (A-B)
EXCLUDE (A)	EXCLUDE (B)	ALLOW (A-B), BLOCK (B-A)
INCLUDE (A)	EXCLUDE (B)	TO_EX (B)
EXCLUDE (A)	INCLUDE (B)	TO_IN (B)

If the computed source list for either an ALLOW or a BLOCK State Change Record is empty, that record is omitted from the Report message.

To cover the possibility of the State Change Report being missed by one or more multicast routers, it is retransmitted [Robustness Variable] - 1 more times, at intervals chosen at random from the range (0, [Unsolicited Report Interval]).

If more changes to the same interface state entry occur before all the retransmissions of the State Change Report for the first change have been completed, each such additional change triggers the immediate transmission of a new State Change Report.

The contents of the new transmitted report are calculated as follows:

- o As was done with the first report, the interface state for the affected multicast address before and after the latest change is compared.

- o The report records expressing the difference are built according to the table above. However these records are not transmitted in a message but instead merged with the contents of the pending report, to create the new State Change Report.

The rules for merging the difference report resulting from the state change and the pending report are described below.

The transmission of the merged State Change Report terminates retransmissions of the earlier State Change Reports for the same multicast address, and becomes the first of [Robustness Variable] transmissions of the new State Change Reports.

Each time a source is included in the difference report calculated above, retransmission state for that source needs to be maintained until [Robustness Variable] State Change Reports have been sent by the node. This is done in order to ensure that a series of successive state changes do not break the protocol robustness.

If the interface listening state change that triggers the new report is a filter mode change, then the next [Robustness Variable] State Change Reports will include a Filter Mode Change Record. This applies even if any number of source list changes occur in that period. The node has to maintain retransmission state for the multicast address until the [Robustness Variable] State Change Reports have been sent. When [Robustness Variable] State Change Reports with Filter Mode Change Records have been transmitted after the last filter mode change, and if source list changes to the interface listening have scheduled additional reports, then the next State Change Report will include Source List Change Records.

Each time a State Change Report is transmitted, the contents are determined as follows. If the report should contain a Filter Mode Change Record, then if the current filter mode of the interface is INCLUDE, a TO_IN record is included in the report, otherwise a TO_EX record is included. If instead the report should contain Source List Change Records, an ALLOW and a BLOCK record are included. The contents of these records are built according to the table below.

Record	Sources included
-----	-----
TO_IN	All in the current interface state that must be forwarded
TO_EX	All in the current interface state that must be blocked
ALLOW	All with retransmission state that must be forwarded
BLOCK	All with retransmission state that must be blocked

If the computed source list for either an ALLOW or a BLOCK record is empty, that record is omitted from the State Change Report.

Note: When the first State Change Report is sent, the non-existent pending report to merge with can be treated as a Source Change Report with empty ALLOW and BLOCK records (no sources have

Vida et al.

[Page 22]

INTERNET-DRAFT

MLDv2

February 2001

retransmission state).

5.2. Action on Reception of a Query

When a node receives a Query, it does not respond immediately. Instead, it delays its response by a random amount of time, bounded by the Maximum Response Delay value derived from the Maximum Response Code in the received Query message. A node may receive a variety of Queries on different interfaces and of different kinds (e.g., General Queries, Multicast Address Specific Queries, and Multicast Address and Source Specific Queries), each of which may require its own delayed response.

Before scheduling a response to a Query, the node must first consider previously scheduled pending responses and in many cases schedule a combined response. Therefore, the node must be able to maintain the following state:

- o A timer per interface for scheduling responses to General Queries.
- o A per-multicast-address-and-interface timer for scheduling responses to Multicast Address Specific and Multicast Address and Source Specific Queries.
- o A per-multicast-address-and-interface list of sources to be reported in the response to a Multicast Address and Source Specific Query.

When a new Query with the Router Alert option arrives on an interface, provided the node has state to report, a delay for a response is randomly selected in the range (0, [Maximum Response Delay]) where Maximum Response Delay is derived from Maximum Response Code in the received Query message. The following rules are then used to determine if a Report needs to be scheduled or not and the type of Report to schedule:

1. If there is a pending response to a previous General Query scheduled sooner than the selected delay, no additional response needs to be scheduled.
2. If the received Query is a General Query, the interface timer is used to schedule a response to the General Query after the

selected delay. Any previously pending response to a General Query is canceled.

3. If the received Query is a Multicast Address Specific Query or a Multicast Address and Source Specific Query and there is no pending response to a previous Query for this multicast address, then the multicast address timer is used to schedule a report. If the received Query is a Multicast Address and Source Specific Query, the list of queried sources is recorded to be used when generating a response.

Vida et al.

[Page 23]

INTERNET-DRAFT

MLDv2

February 2001

4. If there is already a pending response to a previous Query scheduled for this multicast address, and either the new Query is a Multicast Address Specific Query or the recorded source list associated with the multicast address is empty, then the multicast address source list is cleared and a single response is scheduled using the multicast address timer. The new response is scheduled to be sent at the earliest of the remaining time for the pending report and the selected delay.
5. If the received Query is a Multicast Address and Source Specific Query and there is a pending response for this multicast address with a non-empty source list, then the multicast address source list is augmented to contain the list of sources in the new Query and a single response is scheduled using the multicast address timer. The new response is scheduled to be sent at the earliest of the remaining time for the pending report and the selected delay.

When the timer in a pending response record expires, the node transmits, on the associated interface, one or more Report messages carrying one or more Current State Records (see [section 4.2.12](#)), as follows:

1. If the expired timer is the interface timer (i.e., it is a pending response to a General Query), then one Current State Record is sent for each multicast address for which the specified interface has listening state, as described in [section 3.2](#). The Current State Record carries the multicast address and its associated filter mode (MODE_IS_INCLUDE or MODE_IS_EXCLUDE) and source list. Multiple Current State Records are packed into individual Report messages, to the extent possible.
2. If the expired timer is a multicast address timer and the list of recorded sources for that multicast address is empty (i.e., it is a pending response to a Multicast Address Specific Query), then if and only if the interface has listening state for that multicast address, a single Current State Record is sent for that address. The Current State Record carries the multicast address and its associated filter mode (MODE_IS_INCLUDE or MODE_IS_EXCLUDE) and source list, if any.

3. If the expired timer is a multicast address timer and the list of recorded sources for that multicast address is non-empty (i.e., it is a pending response to a Multicast Address and Source Specific Query), then if and only if the interface has listening state for that multicast address, the contents of the responding Current State Record is determined from the interface state and the pending response record, as specified in the following table:

interface state	set of sources in the pending response record	Current State Record
-----	-----	-----
INCLUDE (A)	B	IS_IN (A*B)
EXCLUDE (A)	B	IS_IN (B-A)

If the resulting Current State Record has an empty set of source addresses, then no response is sent.

Finally, after any required Report messages have been generated, the source lists associated with any reported multicast addresses are cleared.

6. DESCRIPTION OF THE PROTOCOL FOR MULTICAST ROUTERS

The purpose of MLD is to enable each multicast router to learn, for each of its directly attached links, which multicast addresses have listeners on that link. MLD version 2 adds the capability for a multicast router to also learn which *sources* have listeners among the neighboring nodes, for packets sent to any particular multicast address. The information gathered by MLD is provided to whichever multicast routing protocol is being used by the router, in order to ensure that multicast packets are delivered to all links where there are interested listeners.

This section describes the part of MLDv2 that is performed by multicast routers. Multicast routers may themselves become multicast address listeners, and therefore also perform the multicast listener part of MLDv2, described in [section 5](#).

A multicast router performs the protocol described in this section over each of its directly attached links. If a multicast router has more than one interface to the same link, it only needs to operate this protocol over one of those interfaces.

For each interface over which the router is operating the MLD protocol, the router must configure that interface to listen to all link-layer multicast addresses that can be generated by IPv6 multicasts. For example, an Ethernet-attached router must set its Ethernet address reception filter to accept all Ethernet multicast addresses that start with the hexadecimal value 3333 [[IPv6-ETHER](#)]; in the case of an Ethernet interface that does not support the filtering of such a range of multicast address, it must be configured to accept ALL Ethernet multicast addresses, in order to meet the requirements of MLD.

On each interface over which this protocol is being run, the router MUST enable reception of link-scope all MLDv2-capable routers multicast address from all sources (and MUST perform the multicast address listener part of MLDv2 for that address on that interface

Vida et al.

[Page 25]

INTERNET-DRAFT

MLDv2

February 2001

Multicast routers need to know only that **at least one** node on an attached link is listening to packets for a particular multicast address from a particular source; a multicast router is not required to keep track of the interests of each individual neighboring node.

MLDv2 is backward compatible with the previous version of the MLD protocol. In order to remain backward compatible with older MLD nodes, MLDv2 multicast routers MUST also implement the previous version of the protocol (see [section 7](#)).

[6.1](#). Conditions for MLD Queries

Multicast routers send General Queries periodically to request Multicast Address Listener information from an attached link. These queries are used to build and refresh the Multicast Address Listener state of routers on attached links.

Nodes respond to these queries by reporting their Multicast Address Listening state (and set of sources they listen to) with Current State Multicast Address Records in MLDv2 Multicast Listener Reports.

As a listener of a multicast address, a node may express interest in listening or not listening to traffic from particular sources. As the desired listening state of a node changes, it reports these changes using Filter Mode Change Records or Source List Change Records. These records indicate an explicit state change in a multicast address at a node in either the Multicast Address Record's source list or its filter mode. When Multicast Address Listening is terminated at a node or traffic from a particular source is no longer desired, a multicast router must query for other listeners of the multicast address or of the source before deleting the multicast address (or source) from its Multicast Address Listener state and

pruning its traffic.

To enable all nodes on a link to respond to changes in multicast address listening, multicast routers send specific queries. A Multicast Address Specific Query is sent to verify that there are no nodes that listen to the specified multicast address or to "rebuild" the listening state for a particular multicast address. Multicast Address Specific Queries are sent when a router receives a State Change Record indicating that a node ceases to listen to a multicast address. They are also sent in order to enable a fast transition of a router from EXCLUDE to INCLUDE mode, in case a received State Change Record motivates this action.

A Multicast Address and Source Specific Query is used to verify that there are no nodes on a link which listen to traffic from a specific set of sources. Multicast Address and Source Specific Queries list sources for a particular multicast address which have been requested

Vida et al.

[Page 26]

INTERNET-DRAFT

MLDv2

February 2001

to no longer be forwarded. This query is sent by a multicast router to learn if any node listens to packets sent to the specified multicast address from the specified source addresses. Multicast Address and Source Specific Queries are only sent in response to State Change Records and never in response to Current State Records. [Section 4.1.13](#) describes each query in more detail.

[6.2.](#) MLD State Maintained by Multicast Routers

Multicast routers implementing MLDv2 keep state per multicast address per attached link. This multicast address state consists of a filter mode, a list of sources, and various timers. For each attached link running MLD, a multicast router records the listening state for that link. That state conceptually consists of a set of records of the form:

(IPv6 multicast address, multicast address timer,
filter mode, (source records))

Each source record is of the form:

(IPv6 source address, source timer)

If all sources for a multicast address are listened to, an empty source record list is kept with filter mode set to EXCLUDE. This means that nodes on this link want all sources for this multicast address to be forwarded. This is the MLDv2 equivalent of an MLDv1 listening state.

[6.2.1.](#) Definition of Router Filter Mode

To reduce internal state, MLDv2 routers keep a filter mode per multicast address per attached link. This filter mode is used to condense the total listening state of a multicast address to a minimum set such that all nodes' listening states are covered. This filter mode may change in response to the reception of particular types of Multicast Address Records or when certain timer conditions occur. In the following sections, we use the term "router filter mode" to refer to the filter mode of a particular multicast address within a router. [Section 6.4](#) describes the changes of a router filter mode per Multicast Address Record received.

Conceptually, when a Multicast Address Record is received, the router filter mode for that multicast address is updated to cover all the requested sources using the least amount of state. As a rule, once a Multicast Address Record with a filter mode of EXCLUDE is received, the router filter mode for that multicast address will be EXCLUDE.

When a router filter mode for a multicast address is EXCLUDE, the source record list contains two types of sources. The first type

Vida et al.

[Page 27]

INTERNET-DRAFT

MLDv2

February 2001

is the set which represents conflicts in the desired reception state; this set must be forwarded by some router on the network. It serves essentially to rebuild the set of sources not forwarded by the router when certain timer conditions occur. It will be also used as the new source list when the router switches back to INCLUDE mode, after its multicast address timer has expired.

The second type of sources is the set of sources which all nodes have requested not to be forwarded. The rules for updating the two sets of the source record list when the router is in EXCLUDE mode will be described in [section 6.4](#). [Appendix A](#) describes the reasons for keeping the first set of sources when in EXCLUDE mode.

When a router filter mode for a multicast address is INCLUDE, the source record list is the list of sources that have listeners for the specific multicast address. Each source in a source record must be forwarded by some router on the network.

Because a reported Multicast Address Record with a filter mode of EXCLUDE will cause a router to transition its filter mode for that multicast address to EXCLUDE, a mechanism for transitioning a router's filter mode back to INCLUDE must exist. If all nodes with a multicast address record having filter mode set to EXCLUDE cease reporting, it is desirable for the router filter mode for that multicast address to transition back to INCLUDE mode. This transition occurs when the Multicast Address Timer expires and is explained in detail in [section 6.5](#).

6.2.2. Definition of Multicast Address Timers

The Multicast Address Timer is only used when a multicast address is in EXCLUDE mode and it represents the time for the *filter mode* of the multicast address to expire and switch to INCLUDE mode. We define a multicast address timer as a decrementing timer with a lower bound of zero kept per multicast address per attached link. Multicast address timers are updated according to the types of Multicast Address Records received.

A Multicast Address Timer expiring when a router filter mode for the multicast address is EXCLUDE means there are no more listeners on the attached link in EXCLUDE mode. At this point, a router will transition to INCLUDE filter mode. [Section 6.5](#) describes the actions taken when a Multicast Address Timer expires while in EXCLUDE mode.

The following table summarizes the role of the Multicast Address Timer. [Section 6.4](#) describes the details of setting the Multicast Address Timer per type of Multicast Address Record received.

Multicast Address Filter Mode -----	Multicast Address Timer Value -----	Actions/Comments -----
INCLUDE	Timer >= 0	All listeners in INCLUDE mode.
EXCLUDE	Timer > 0	At least one listener in EXCLUDE mode.
EXCLUDE	Timer == 0	No more listeners in EXCLUDE mode to the multicast address. If all source timers have expired then delete Multicast Address Record. If there are still source record timers running, switch to INCLUDE filter mode using those source records with running timers as the INCLUDE source record state.

6.2.3. Definition of Source Timers

A source timer is kept per source record and is a decrementing timer with a lower bound of zero. Source timers are updated according to the type and filter mode of the Multicast Address Record received. [Section 6.4](#) describes the setting of source timers per type of Multicast Address Records received.

When the router filter mode for a multicast address is INCLUDE, a source record with a running timer means that there are currently one or more nodes (in INCLUDE filter mode) which listen to that source. If a source timer expires, the router concludes that this particular source has no longer listeners on the attached link, and deletes the associated source record.

Source timers are treated differently when a router filter mode for a multicast address is EXCLUDE. If a source record has a running timer it means that at least one system desires the source. It should therefore be forwarded by some router on the network. [Appendix A](#) describes the reasons for keeping state for sources that have been requested to be forwarded while in EXCLUDE state.

If a source timer expires with a router filter mode for the multicast address of EXCLUDE, the router informs the routing protocol that there is no longer a listener on the link interested in traffic from

Vida et al.

[Page 29]

INTERNET-DRAFT

MLDv2

February 2001

this source.

When a router filter mode for a multicast address is EXCLUDE, source records are only deleted when the Multicast Address Timer expires or when newly received Multicast Address Records modify the source record list of the router. [Section 6.3](#) describes the actions that should be taken depending on the value of the source timer.

6.3. MLDv2 Source Specific Forwarding Rules

When a multicast router receives a datagram from a source destined to a particular multicast address, a decision has to be made whether to forward the datagram on an attached link or not. The multicast routing protocol in use is in charge of this decision, and should use the MLDv2 information to ensure that all sources/multicast addresses that have listeners on a link are forwarded to that link. MLDv2 information does not override multicast routing information; for example, if the MLDv2 filter mode for a multicast address is EXCLUDE, a router may still forward packets for excluded sources to a transit link.

To summarize, the following table describes the forwarding

suggestions made by MLDv2 to the routing protocol for traffic originating from a source destined to a multicast address. It also summarizes the actions taken upon the expiration of a source timer based on the router filter mode of the multicast address.

Multicast Address Filter Mode -----	Source Timer Value -----	Action -----
INCLUDE	TIMER > 0	Suggest to forward traffic from source
INCLUDE	TIMER == 0	Suggest to stop forwarding traffic from source and remove source record. If there are no more source records, delete multicast address record
INCLUDE	No source element	Suggest to not forward traffic from source
EXCLUDE	TIMER > 0	Suggest to forward traffic from source
EXCLUDE	TIMER == 0	Suggest to not forward traffic from source (DO NOT remove record)

EXCLUDE	No Source Element	Suggest to forward traffic from source
---------	-------------------	--

[6.4.](#) Action on Reception of Reports

[6.4.1.](#) Reception of Current State Records

When receiving Current State Records, a router updates both its multicast address and source timers. In some circumstances, the reception of a type of multicast address record will cause the router filter mode for that multicast address to change. The table below describes the actions, with respect to state and timers, that occur to a router's state upon reception of Current State Records.

The following notation is used to describe the updating of source timers. The notation (A, B) will be used to represent the total number of sources for a particular multicast address, where

A = set of source records whose source timers > 0
 (Sources that at least one node has requested to be forwarded)
 B = set of source records whose source timers = 0
 (Sources that MLD will suggest to the routing protocol not to forward)

Note that there will be two sets only when a router's filter mode for a multicast address is EXCLUDE. When the filter mode is INCLUDE, a single set is used to describe the set of sources requested to be forwarded (e.g. simply (A)).

In the following tables, abbreviations are used for several variables (all of which are described in detail in [section 8](#)). The variable MALI is an abbreviation for the Multicast Address Listening Interval which is the time in which multicast address listening will time out. The variable LLQI is an abbreviation for the Last Listener Query Interval (default 1s) which is the Maximum Response Delay used to derive the Maximum Response Code in Multicast Address or Multicast Address and Source Specific Queries. (Note that for values of LLQI greater than 32768 milliseconds, a limited set of values can be represented, corresponding to sequential values of Maximum Response Code. When converting a configured time to a Maximum Response Code value, it is recommended to use the exact value if possible, or the next higher value if the requested value is not exactly representable.)

Within the "Actions" section of the router state tables, we use the notation '(A)=J', which means that the set A of source records should have their source timers set to value J. 'Delete (A)' means that the set A of source records should be deleted. 'Multicast Address Timer = J' means that the Multicast Address Timer for the multicast address should be set to value J.

Vida et al.

[Page 31]

INTERNET-DRAFT

MLDv2

February 2001

Router State -----	Report Received -----	New Router State -----	Actions -----
INCLUDE (A)	IS_IN (B)	INCLUDE (A+B)	(B)=MALI
INCLUDE (A)	IS_EX (B)	EXCLUDE (A*B, B-A)	(B-A)=0 Delete (A-B) Multicast Address Timer = MALI
EXCLUDE (X,Y)	IS_IN (A)	EXCLUDE (X+A, Y-A)	(A)=MALI
EXCLUDE (X,Y)	IS_EX (A)	EXCLUDE (A-Y, Y*A)	(A-X-Y)=MALI Delete (X-A) Delete (Y-A) Multicast Address Timer = MALI

6.4.2. Reception of Filter Mode Change and Source List Change Records

When a change in the global state of a multicast address occurs in a node, the node sends either a Source List Change Record or a Filter Mode Change Record for that multicast address. As with Current State Records, routers must act upon these records and possibly change their own state to reflect the new listening state of the link.

Routers must query sources or multicast addresses that are requested to be no longer forwarded. When a router queries or receives a query for a specific set of sources, it lowers its source timers for those sources to a small interval of Last Listener Query Interval milliseconds. If multicast address records are received in response to the queries which express interest in listening the queried sources, the corresponding timers are updated.

Multicast Address Specific queries can also be used in order to enable a fast transition of a router from EXCLUDE to INCLUDE mode, in case a received Multicast Address Record motivates this action. The Multicast Address Timer for that multicast address will be lowered to a small interval of Last Listener Query Interval milliseconds. If any multicast address records expressing EXCLUDE mode interest in the multicast address are received within this interval, the multicast address timer is updated and the suggestion to the routing protocol to forward the multicast address stands without any interruption. If not, the router will switch to INCLUDE filter mode for that multicast address.

During a query period (i.e. Last Listener Query Interval milliseconds) the MLD component in the router continues to suggest to the routing protocol that it forwards traffic from the multicast addresses or sources that it is querying. It is not until after Last Listener Query Interval milliseconds without receiving a record

Vida et al.

[Page 32]

expressing interest in the queried multicast address or sources that the router may prune the multicast address or sources from the link.

The following table describes the changes in multicast address state and the action(s) taken when receiving either Filter Mode Change or Source List Change Records. This table also describes the queries which are sent by the router in Querier state when a particular report is received.

We use the following notation for describing the queries which are sent. We use the notation 'Q(MA)' to describe a Multicast Address Specific Query to the MA multicast address. We use the notation 'Q(MA,A)' to describe a Multicast Address and Source Specific Query to the MA multicast address with source list A. If source list A is null as a result of the action (e.g. A*B) then no query is sent as a

result of the operation.

In order to maintain protocol robustness, queries defined in the Actions column of the table below need to be transmitted [Last Listener Query Count] times within the [Last Listener Query Interval] period. If while scheduling new queries, there are already pending queries to be retransmitted for the same multicast address, the new and pending queries have to be merged. In addition, received host reports for a multicast address with pending queries may affect the contents of those queries. [Section 6.7.](#) describes the process of building and maintaining the state of pending queries.

Router State	Report Received	New Router State	Actions
-----	-----	-----	-----
INCLUDE (A)	ALLOW (B)	INCLUDE (A+B)	(B)=MALI
INCLUDE (A)	BLOCK (B)	INCLUDE (A)	Send Q(MA,A*B)
INCLUDE (A)	TO_EX (B)	EXCLUDE (A*B,B-A)	(B-A)=0 Delete (A-B) Send Q(MA,A*B) Multicast Address Timer = MALI
INCLUDE (A)	TO_IN (B)	INCLUDE (A+B)	(B)=MALI Send Q(MA,A-B)
EXCLUDE (X,Y)	ALLOW (A)	EXCLUDE (X+A,Y-A)	(A)=MALI
EXCLUDE (X,Y)	BLOCK (A)	EXCLUDE (X+(A-Y),Y)	(A-X-Y)=LLQI Send Q(MA,A-Y)
EXCLUDE (X,Y)	TO_EX (A)	EXCLUDE (A-Y,Y*A)	(A-X-Y)=LLQI Delete (X-A) Delete (Y-A) Send Q(MA,A-Y) Multicast Address Timer = MALI

Vida et al.

[Page 33]

INTERNET-DRAFT

MLDv2

February 2001

EXCLUDE (X,Y)	TO_IN (A)	EXCLUDE (X+A,Y-A)	(A)=MALI Send Q(MA,X-A) Send Q(MA)
---------------	-----------	-------------------	--

[6.5.](#) Switching Router Filter Modes

The multicast address timer is used as a mechanism for transitioning the router filter mode from EXCLUDE to INCLUDE.

When a multicast address timer expires with a router filter mode of

EXCLUDE, a router assumes that there are no nodes with a *filter mode* of EXCLUDE present on the attached link. When a router's filter mode for a multicast address is EXCLUDE and the multicast address timer expires, the router filter mode for the multicast address transitions to INCLUDE.

A router uses source records with running source timers as its state for the switch to a filter mode of INCLUDE. If there are any source records with source timers greater than zero (i.e. requested to be forwarded), a router switches to filter mode of INCLUDE using those source records. Source records whose timers are zero (from the previous EXCLUDE mode) are deleted.

For example, if a router's state for a multicast address is EXCLUDE(X,Y) and the multicast address timer expires for that multicast address, the router switches to filter mode of INCLUDE with state INCLUDE(X).

[6.6.](#) Action on Reception of Queries

[6.6.1.](#) Timer Updates

When a router sends or receives a query with a clear Suppress Router-Side Processing flag, it must update its timers to reflect the correct timeout values for the multicast address or sources being queried. The following table describes the timer actions when sending or receiving a Multicast Address Specific or Multicast Address and Source Specific Query with the Suppress Router-Side Processing flag not set.

Query	Action
-----	-----
Q(MA,A)	Source Timers for sources in A are lowered to LLQI
Q(MA)	Multicast Address Timer is lowered to LLQI

When a router sends or receives a query with the Suppress Router-Side Processing flag set, it will not update its timers.

[6.6.2.](#) Querier Election

MLDv2 elects a single router to be in Querier state per subnet using the same querier election mechanism as MLDv1, namely by IPv6 address. When a router receives a query with a lower IPv6 address, it sets the Other Querier Present timer to Other Querier Present Timeout and ceases to send queries on the link if it was the previously elected querier. After its Other Querier Present timer expires, it should

begin sending General Queries.

If a router receives an older version query, it MUST use the older version of MLD on the link. For a detailed description of compatibility issues between MLD versions see [section 7](#).

[6.6.3](#) Building and Sending Specific Queries

[6.6.3.1](#). Building and Sending Multicast Address Specific Queries

When a table action "Send Q(MA)" is encountered, then the multicast address timer must be lowered to LLQI. The router must then immediately send a Multicast Address Specific query as well as schedule [Last Listener Query Count - 1] query retransmissions to be sent within [Last Listener Query Interval].

When transmitting a Multicast Address Specific Query, if the multicast address timer is larger than LLQI, the "Suppress Router-Side Processing" bit is set in the query message.

[6.6.3.2](#). Building and Sending Multicast Address and Source Specific Queries

When a table action "Send Q(MA,X)" is encountered by a querier in the table in [section 6.4.2](#), the following actions must be performed for each of the sources in X that send to multicast address MA, with source timer larger than LLQI:

- o Lower source timer to LLQI.
- o Set number of retransmissions for each source to [Last Listener Query Count].

The router must then immediately send a Multicast Address and Source Specific Query as well as schedule [Last Listener Query Count -1] query retransmissions to be sent within [Last Listener Query Interval]. The contents of these queries are calculated as follows.

When building a Multicast Address and Source Specific Query for a multicast address MA, two separate query messages are sent for the multicast address. The first one has the "Suppress Router-Side

Processing" bit set and contains all the sources with retransmission state and timers greater than LLQI. The second has the "Suppress Router-Side Processing" bit clear and contains all the sources with retransmission state and timers lower or equal to LLQI. If either of the two calculated messages does not contain any sources, then its

transmission is suppressed.

Note: If a Multicast Address Specific query is scheduled to be transmitted at the same time as a Multicast Address and Source specific query for the same multicast address, then transmission of the Multicast Address and Source specific message with the "Suppress Router-Side Processing" bit set may be suppressed.

7. INTEROPERATION WITH OLDER VERSIONS OF MLD

MLD version 2 hosts and routers interoperate with hosts and routers that have not yet been upgraded to MLDv2. This compatibility is maintained by hosts and routers taking appropriate actions depending on the versions of MLD operating on hosts and routers within a network.

7.1. Query Version Distinctions

The MLD version of a Multicast Listener Query message is determined as follows:

MLDv1 Query: length = 24 octets

MLDv2 Query: length \geq 28 octets

Query messages that do not match any of the above conditions (e.g., a Query of length 26 octets) MUST be silently ignored.

7.2. Multicast Address Listener Behavior

7.2.1. In the Presence of Older Version Queriers

In order to be compatible with older version routers, MLDv2 hosts MUST operate in version 1 compatibility mode. MLDv2 hosts MUST keep state per local interface regarding the compatibility mode of each attached link. A host's compatibility mode is determined from the Host Compatibility Mode variable which can be in one of the two states: MLDv1 or MLDv2. This variable is kept per interface and is dependent on the version of General Queries heard on that interface as well as on the Older Version Querier Present timer for the interface. This timer is set to Older Version Querier Present Timeout seconds whenever an MLDv1 Multicast Address Listener Query is received.

Vida et al.

[Page 36]

The Host Compatibility Mode of an interface changes whenever an older version query (than the current compatibility mode) is heard or when

certain timer conditions occur. When Older Version Querier Present timer expires, a host switches to Host Compatibility mode of MLDv2.

The Host Compatibility Mode variable is based on whether an older version query was heard in the last Older Version Querier Present Timeout seconds. The Host Compatibility Mode is set depending on the following:

Host Compatibility Mode -----	Timer State -----
MLDv2 (default)	Older Version Querier Present = 0
MLDv1	Older Version Querier Present > 0

If a host receives a query which causes its Older Version Querier Present timer to be updated and correspondingly its compatibility mode, it should switch compatibility modes immediately.

When Host Compatibility Mode is MLDv2, a host acts using the MLDv2 protocol on that interface. When Host Compatibility Mode is MLDv1, a host acts in MLDv1 compatibility mode, using only the MLDv1 protocol, on that interface.

An MLDv1 router will send General Queries with the Maximum Response Code set to the desired Maximum Response Delay, i.e. the full range of this field is linear and the exponential algorithm described in [section 4.1.3](#). is not used.

Whenever a host changes its compatibility mode, it cancels all its pending response and retransmission timers.

[7.2.2](#). In the Presence of Older Version Multicast Address Listeners

An MLDv2 host may be placed on a link where there are hosts that have not yet been upgraded to MLDv2. A host MAY allow its MLDv2 Multicast Listener Report to be suppressed by Version 1 Multicast Listener Report.

[7.3](#). Multicast Router Behavior

[7.3.1](#). In the Presence of Older Version Queriers

MLDv2 routers may be placed on a network where there is at least one router that has not yet been upgraded to MLDv2. The following requirements apply:

- o If an older version of MLD is present on routers, the querier MUST use the lowest version of MLD present on the network. This must be administratively assured; routers that desire to be compatible with MLDv1 MUST have a configuration option to act in MLDv1 compatibility mode. When in MLDv1 mode, routers MUST send periodic Queries truncated at the Multicast Address field (i.e. 24 bytes long), and SHOULD also warn about receiving an MLDv2 query (such warnings must be rate-limited). They also MUST fill in the Maximum Response Delay in the Maximum Response Code field, i.e. the exponential algorithm described in [section 4.1.3](#). is not used.
- o If a router is not explicitly configured to use MLDv1 and hears an MLDv1 Query, it SHOULD log a warning. These warnings MUST be rate-limited.

[7.3.2](#). In the Presence of Older Version Multicast Address Listeners

MLDv2 routers may be placed on a network where there are hosts that have not yet been upgraded to MLDv2. In order to be compatible with older version hosts, MLDv2 routers MUST operate in version 1 compatibility mode. MLDv2 routers keep a compatibility mode per multicast address record. The compatibility mode of a multicast address is determined from the Multicast Address Compatibility Mode variable which can be in one of the two following states: MLDv1 or MLDv2. This variable is kept per multicast address record and is dependent on the version of Multicast Listener Reports heard for that multicast address as well as the Older Version Host Present timer for the multicast address. This timer is set to Older Version Host Present Timeout seconds whenever an MLDv1 Multicast Listener Report is received.

The Multicast Address Compatibility Mode of a multicast address record changes whenever an older version report (than the current compatibility mode) is heard or when certain timer conditions occur. When Older Version Host Present timer expires, a router switches to Multicast Address Compatibility mode of MLDv2.

The Multicast Address Compatibility Mode variable is based on whether an older version report was heard in the last Older Version Host Present Timeout seconds. The Multicast Address Compatibility Mode is set depending on the following:

Multicast Address Compatibility Mode	Timer State
-----	-----
MLDv2 (default)	Older Version Host Present = 0
MLDv1	Older Version Host Present > 0

If a router receives a report which causes its older Host Present timer to be updated and correspondingly its compatibility mode, it SHOULD switch compatibility modes immediately.

When Multicast Address Compatibility Mode is MLDv2, a router acts using the MLDv2 protocol for that multicast address. When Multicast Address Compatibility Mode is MLDv1, a router acts in MLDv1 compatibility mode, using only the MLDv1 protocol for that multicast address.

8. LIST OF TIMERS, COUNTERS, AND THEIR DEFAULT VALUES

Most of these timers are configurable. If non-default settings are used, they MUST be consistent among all nodes on a single link. Note that parentheses are used to group expressions to make the algebra clear.

8.1. Robustness Variable

The Robustness Variable allows tuning for the expected packet loss on a link. If a link is expected to be lossy, the Robustness Variable may be increased. MLD is robust to (Robustness Variable - 1) packet losses. The Robustness Variable MUST NOT be zero, and SHOULD NOT be one. Default: 2.

8.2. Query Interval

The Query Interval is the interval between General Queries sent by the Querier. Default: 125 seconds.

By varying the [Query Interval], an administrator may tune the number of MLD messages on the link; larger values cause MLD Queries to be sent less often.

8.3. Query Response Interval

The Maximum Response Delay used to calculate the Maximum Response Code inserted into the periodic General Queries. Default: 10000 (10 seconds)

By varying the [Query Response Interval], an administrator may tune the burstiness of MLD messages on the link; larger values make the traffic less bursty, as host responses are spread out over a larger interval. The number of seconds represented by the [Query Response Interval] must be less than the [Query Interval].

8.4. Multicast Address Listener Interval

The Multicast Address Listener Interval (MALI) is the amount of time that must pass before a multicast router decides there are no more listeners of a multicast address or a particular source on a link.

INTERNET-DRAFT

MLDv2

February 2001

This value MUST be ((the Robustness Variable) times (the Query Interval)) plus (one Query Response Interval).

8.5. Other Querier Present Timeout

The Other Querier Present Timeout is the length of time that must pass before a multicast router decides that there is no longer another multicast router which should be the querier. This value MUST be ((the Robustness Variable) times (the Query Interval)) plus (one half of one Query Response Interval).

8.6. Startup Query Interval

The Startup Query Interval is the interval between General Queries sent by a Querier on startup. Default: 1/4 the Query Interval.

8.7. Startup Query Count

The Startup Query Count is the number of Queries sent out on startup, separated by the Startup Query Interval. Default: the Robustness Variable.

8.8. Last Listener Query Interval

The Last Listener Query Interval is the Maximum Response Delay used to calculate the Maximum Response Code inserted into Multicast Address Specific Queries sent in response to Version 1 Multicast Listener Done messages. It is also the Maximum Response Delay used in calculating the Maximum Response Code inserted into Multicast Address and Source Specific Query messages. Default: 10000 (1 second).

This value may be tuned to modify the "leave latency" of the link. A reduced value results in reduced time to detect the departure of the last listener for a multicast address or source.

8.9. Last Listener Query Count

The Last Listener Query Count is the number of Multicast Address Specific Queries sent before the router assumes there are no local listeners. The Last Listener Query Count is also the number of Multicast Address and Source Specific Queries sent before the router assumes there are no listeners for a particular source. Default: the Robustness Variable.

INTERNET-DRAFT

MLDv2

February 2001

[8.10.](#) Unsolicited Report Interval

The Unsolicited Report Interval is the time between repetitions of a node's initial report of interest in a multicast address. Default: 1 second.

[8.11.](#) Older Version Querier Present Timeout

The Older Version Querier Present Timeout is the time-out for transitioning a host back to MLDv2 mode once an older version query is heard. When an older version query is received, hosts set their Older Version Querier Present Timer to Older Version Querier Present Timeout.

This value MUST be ((the Robustness Variable) times (the Query Interval in the last Query received)) plus (one Query Response Interval).

[8.12.](#) Older Version Host Present Timeout

The Older Version Host Present Timeout is the time-out for transitioning a multicast address back to MLDv2 mode once an older version report is sent for that multicast address. When an older version report is received, routers set their Older Version Host Present Timer to Older Version Host Present Timeout.

This value MUST be ((the Robustness Variable) times (the Query Interval)) plus (one Query Response Interval).

[8.13.](#) Configuring timers

This section is meant to provide advice to network administrators on how to tune these settings to their network. Ambitious router implementations might tune these settings dynamically based upon changing characteristics of the network.

[8.13.1.](#) Robustness Variable

The Robustness Variable tunes MLD to expected losses on a link. MLDv2 is robust to (Robustness Variable - 1) packet losses, e.g. if the Robustness Variable is set to the default value of 2, MLDv2 is robust to a single packet loss but may operate imperfectly if more

losses occur. On lossy links, the Robustness Variable should be increased to allow for the expected level of packet loss. However, increasing the Robustness Variable increases the leave latency of the link (the time between when the last listener stops listening to a source or multicast address and when the traffic stops flowing).

[8.13.2.](#) Query Interval

The overall level of periodic MLD traffic is inversely proportional to the Query Interval. A longer Query Interval results in a lower overall level of MLD traffic. The Query Interval **MUST** be equal to or longer than the Maximum Response Delay used to calculate the Maximum Response Code inserted in General Query messages.

[8.13.3.](#) Maximum Response Delay

The burstiness of MLD traffic is inversely proportional to the Maximum Response Delay. A longer Maximum Response Delay will spread Report messages over a longer interval. However, a longer Maximum Response Delay in Multicast Address Specific and Multicast Address And Source Specific Queries extends the leave latency (the time between when the last listener stops listening to a source or multicast address and when the traffic stops flowing.) The expected rate of Report messages can be calculated by dividing the expected number of Reporters by the Maximum Response Delay. The Maximum Response Delay may be dynamically calculated per Query by using the expected number of Reporters for that Query as follows:

Query Type -----	Expected number of Reporters -----
General Query	All nodes on link
Multicast Address Specific Query	All nodes on the link that had expressed interest in the multicast address
Multicast Address and Source Specific Query	All nodes on the link that had expressed interest in the source and multicast address

A router is not required to calculate these populations or tune the Maximum Response Delay dynamically; these are simply guidelines.

[9.](#) SECURITY CONSIDERATIONS

IPSEC in Authentication Header mode [[AH](#)] may be used to protect

against remote attacks by ensuring that MLDv2 messages came from a node on the LAN (or, more specifically, a node with the proper key). When using IPSEC, the messages sent to FF02::1 and the all MLDv2 capable routers address should be authenticated using AH. When keying, there are two possibilities:

1. Use a symmetric signature algorithm with a single key for the LAN (or a key for each multicast address). This allows validation that a packet was sent by a node with the key. This has the limitation that any node with the key can forge a message; it is

Vida et al.

[Page 42]

INTERNET-DRAFT

MLDv2

February 2001

not possible to authenticate the individual sender precisely.

2. When appropriate key management standards have been developed, use an asymmetric signature algorithm. All nodes need to know the public key of all routers, and all routers need to know the public key of all nodes. This requires a large amount of key management but has the advantage that senders can be authenticated individually so e.g. a host cannot forge a message that only routers should be allowed to send.

This solution only directly applies to Query and Done messages in MLDv1, since Reports are sent to the multicast address being reported and it is not feasible to agree on a key for host-to-router communication for arbitrary multicast addresses.

We consider the ramifications of a forged message of each type.

9.1. Query Message

A forged Query message from a machine with a lower IPv6 address than the current Querier will cause Querier duties to be assigned to the forger. If the forger then sends no more Query messages, other routers' Other Querier Present timer will time out and one will resume the role of Querier. During this time, if the forger ignores Multicast Listener Done Messages, traffic might flow to multicast addresses with no listeners for up to [Multicast Address Listener Interval].

A DoS attack on a node could be staged through forged Multicast Address and Source Specific Queries. The attacker can find out about the listening state of a specific node with a general query. After that it could send a large number of Multicast Address and Source Specific queries, each with a large source list and/or long Maximum Response Delay. The node will have to store and maintain the sources specified in all of those queries for as long as it takes to send the delayed response. This would consume both memory and CPU cycles in order to augment the recorded sources with the source lists included in the successive queries.

To protect against such a DoS attack, a node stack implementation could restrict the number of Multicast Address and Source Specific Queries per multicast address within this interval, and/or record only a limited number of sources.

Forged Query messages from the local network can be easily traced. There are two measures necessary to defend against externally forged Queries:

- o Routers SHOULD NOT forward Queries. This is easier for a router to accomplish if the Query carries the Router Alert option.

Vida et al.

[Page 43]

INTERNET-DRAFT

MLDv2

February 2001

- o Hosts SHOULD Ignore v2 Queries without the Router Alert option.

9.2. Current State Report messages

A forged Report message may cause multicast routers to think there are listeners of a multicast address on a link when there are not. Forged Report messages from the local network are meaningless, since listening to a multicast address on a host is generally an unprivileged operation, so a local user may trivially gain the same result without forging any messages. Forged Report messages from external sources are more troublesome; there are two defenses against externally forged Reports:

- o Ignore the Report if you cannot identify the source address of the packet as belonging to a link assigned to the interface on which the packet was received. This solution means that Reports sent by mobile hosts without addresses on the local network will be ignored.
- o Ignore Report messages without Router Alert options [[IPv6-ALERT](#)], and require that routers not forward Report messages. (The requirement is not a requirement of generalized filtering in the forwarding path, since the packets already have Router Alert options in them).

A forged Version 1 Report Message may put a router into "older version listener present" state for a particular multicast address, meaning that the router will ignore MLDv2 source specific state messages. This can cause traffic to flow from unwanted sources for up to [Multicast Address Listener Interval]. This can be solved by providing routers with a configuration switch to ignore Version 1 messages completely. This breaks automatic compatibility with Version 1 hosts, so should only be used in situations where source include and exclude is critical.

9.3. State Change Report messages

A forged State Change Report message will cause the Querier to send out Multicast Address Specific or Multicast Address and Source Specific Queries for the multicast address in question. This causes extra processing on each router and on each listener of the multicast address, but can not cause loss of desired traffic. There are two defenses against externally forged State Change Report messages:

- o Ignore the State Change Report message if you cannot identify the source address of the packet as belonging to a link assigned to the interface on which the packet was received. This solution means that State Change Report messages sent by mobile nodes without addresses on the local link will be ignored.

Vida et al.

[Page 44]

INTERNET-DRAFT

MLDv2

February 2001

- o Ignore State Change Report messages without Router Alert options [[IPv6-ALERT](#)] and require that routers not forward State Change Report messages. (The requirement is not a requirement of generalized filtering in the forwarding path, since the packets already have Router Alert options in them).

10. IANA CONSIDERATIONS

A special IP destination address called *all MLDv2-capable routers* should be allocated by IANA. Version 2 Multicast Listener Reports will be sent to this special address.

11. ACKNOWLEDGMENTS

This document is the translation of the [[IGMPv3](#)] Internet Draft for IPv6 semantics. It was elaborated based on the translation of [IGMPv2] into [[MLDv1](#)].

12. REFERENCES

- [ADDR-ARCH] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 2373](#), July 1998.
- [AH] Kent, S. and R. Atkinson, "IP Authentication Header", [RFC 2402](#), November 1998.
- [FILTER-API] Thaler, D., B. Fenner, and B. Quinn, "Socket Interface Extensions for Multicast Source Filters", Work in progress, <[draft-ietf-idmr-msf-api-01.txt](#)>, June 2000.

- [ICMPv6] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 2463](#), December 1998.
- [IGMPv3] Cain, B., Deering, S., Fenner, B., Kouvelas, I., Thyagarajan, A., "Internet Group Management Protocol, Version 3", Work in progress, [<draft-ietf-idmr-igmp-v3-06.txt>](#), January 2001.
- [IPv6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [IPv6-ALERT] Partridge, C., Jackson, A., "IPv6 Router Alert Option", [RFC 2711](#), November 1999.
- [IPv6-ETHER] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", [RFC 2464](#), December, 1998.

Vida et al.

[Page 45]

INTERNET-DRAFT

MLDv2

February 2001

- [IPv6-SOCKET] Gilligan, R., et al., "Basic Socket Interface Extensions for IPv6", [RFC 2553](#), March 1999.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), [BCP14](#), March 1997.
- [MLDv1] Deering, S., Fenner, W., Haberman, B., "Multicast Listener Discovery (MLD) for IPv6", [RFC 2710](#), November 1999.
- [STD-PROC] Bradner, S., "The Internet Standards Process - Revision 3", [BCP 9](#), [RFC 2026](#), November 1996.

APPENDIX A. DESIGN RATIONALE

[A.1](#) The Need for State Change Messages

MLDv2 specifies two types of Multicast Listener Reports: Current State and State Change. This section describes the rationale for the need for both these types of Reports.

Routers need to distinguish Multicast Listener Reports that were sent in response to Queries from those that were sent as a result of a change in interface state. Multicast Listener Reports that are sent in response to Multicast Address Listener Queries are used mainly to refresh the existing state at the router; they typically do not cause transitions in state at the router. Multicast Listener Reports that are sent in response to changes in interface state require the router

to take some action in response to the received report (see [Section 6.4](#)).

The inability to distinguish between the two types of reports would force a router to treat all Multicast Listener Reports as potential changes in state and could result in increased processing at the router as well as an increase in MLD traffic on the link.

[A.2](#) Host Suppression

In MLDv1, a host would cancel sending a pending multicast listener report if a similar report was observed from another listener on the link. In MLDv2, this suppression of multicast listener reports has been removed. The following points explain the reasons behind this decision.

1. Routers may want to track per-host multicast listener status on an interface. This allows routers to implement fast leaves (e.g. for layered multicast congestion control schemes) as well as track listener status for possible accounting purposes.

Vida et al.

[Page 46]

INTERNET-DRAFT

MLDv2

February 2001

2. Multicast Listener Report suppression does not work well on bridged LANs. Many bridges and Layer2/Layer3 switches that implement MLD snooping do not forward MLD messages across LAN segments in order to prevent multicast listener report suppression. Removing multicast listener report suppression eases the job of these MLD snooping devices.
3. By eliminating multicast listener report suppression, hosts have fewer messages to process; this leads to a simpler state machine implementation.
4. In MLDv2, a single multicast listener report now bundles multiple multicast address records to decrease the number of packets sent. In comparison, the previous version of MLD required that each multicast address be reported in a separate message.

[A.3](#) Switching router filter modes from EXCLUDE to INCLUDE

If there exist nodes in both EXCLUDE and INCLUDE modes for a single multicast address on a link, the router must be in EXCLUDE mode as well (see [section 6.2.1](#)). In EXCLUDE mode, a router forwards traffic from all sources unless that source exists in the exclusion source list. If all nodes in EXCLUDE mode cease to exist, it would be desirable for the router to switch back to INCLUDE mode seamlessly without interrupting the flow of traffic to existing listeners.

One of the ways to accomplish this is for routers to keep track of all sources that nodes that are in INCLUDE mode listen to, even though the router itself is in EXCLUDE mode. If the multicast address timer now expires in EXCLUDE mode, it implies that there are no nodes in EXCLUDE mode on the link (otherwise a multicast listener report from that node would have refreshed the multicast address timer). The router can then switch to INCLUDE mode seamlessly with the list of sources currently being forwarded in its source list.

AUTHORS' ADDRESSES

Rolland Vida
LIP6, Universite Pierre et Marie Curie
8, rue du Capitaine Scott
75015 Paris, France
phone: +33-1.44.27.71.26
email: Rolland.Vida@lip6.fr

Luis Henrique Maciel Kosmowski Costa
LIP6, Universite Pierre et Marie Curie
8, rue du Capitaine Scott
75015 Paris, France
phone: +33-1.44.27.87.72
email: Luis.Costa@lip6.fr

Vida et al.

[Page 47]

INTERNET-DRAFT

MLDv2

February 2001

Remi Zara
LIP6, Universite Pierre et Marie Curie
8, rue du Capitaine Scott
75015 Paris, France
phone: +33-1.44.27.71.26
email: Remi.Zara@lip6.fr

Serge Fdida
LIP6, Universite Pierre et Marie Curie
8, rue du Capitaine Scott
75015 Paris, France
phone: +33-1.44.27.30.58
email: Serge.Fdida@lip6.fr

Steve Deering
Cisco Systems, Inc.
170 Tasman Drive
San Jose, CA 95134-1706
phone: +1-408-527-8213
email: deering@cisco.com

Bill Fenner
AT&T Labs - Research
75 Willow Rd.

Menlo Park, CA 94025
phone: +1-650-330-7893
email: fenner@research.att.com

Isidor Kouvelas
Cisco Systems, Inc.
170 Tasman Drive
San Jose, CA 95134-1706
phone: +1-408-525-0727
email: kouvelas@cisco.com

Brian Haberman
Nortel Networks
4309 Emperor Blvd. Suite 200
Durham, NC 27703
phone: +1-919-992-4439
email: haberman@nortelnetworks.com